




Cisco UCS Integrated Infrastructure for Big Data with Cloudera for Enterprise Data Hub

With optional guidelines for Scaling with Cisco Application Centric Infrastructure (ACI)

Last Updated: May 22, 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.



Karthik Kulkarni

Karthik Kulkarni, Technical Marketing Engineer, Data Center Solutions Group (Cisco Systems)

Karthik Kulkarni 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.



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.

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 the contributions of Amrit Kharel, Ashwin Manjunatha, and Sindhu Sudhir in developing this document.



Cisco UCS Integrated Infrastructure for Big Data with Cloudera for Enterprise Data Hub

Audience

This document describes the architecture and deployment procedures of Cloudera on Cisco UCS Integrated Infrastructure for Big Data with Application Centric Infrastructure (ACI). The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering and customers who want to deploy Cloudera on Cisco UCS Integrated Infrastructure for Big Data with Application Centric Infrastructure (ACI).

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. Cloudera is the leading provider of enterprise-grade Hadoop infrastructure software and services, and the leading contributor to the Apache Hadoop project overall. Cloudera provides an enterprise-ready Hadoop-based solution known as Cloudera Enterprise, which includes their market leading open source Hadoop distribution (CDH), their comprehensive management system (Cloudera Manager), and technical support. The combination of Cisco UCS Servers along with Application Centric Infrastructure (ACI) and Cloudera provides industry-leading platform for Hadoop based applications.

This solution is based on Cisco UCS Integrated Infrastructure for Big Data with Application Centric Infrastructure (ACI), with multiple Cisco UCS Fabric Interconnect domains. Each Fabric Interconnect domain consists of 5 racks of servers (total of 80 Cisco UCS C240 M4 servers) along with a pair of Fabric Interconnect. These domains are inter-connected through ACI.

Cisco UCS Integrated Infrastructure for Big Data

The Cisco UCS solution for Cloudera 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 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 3 Cisco UCS VIC 1227

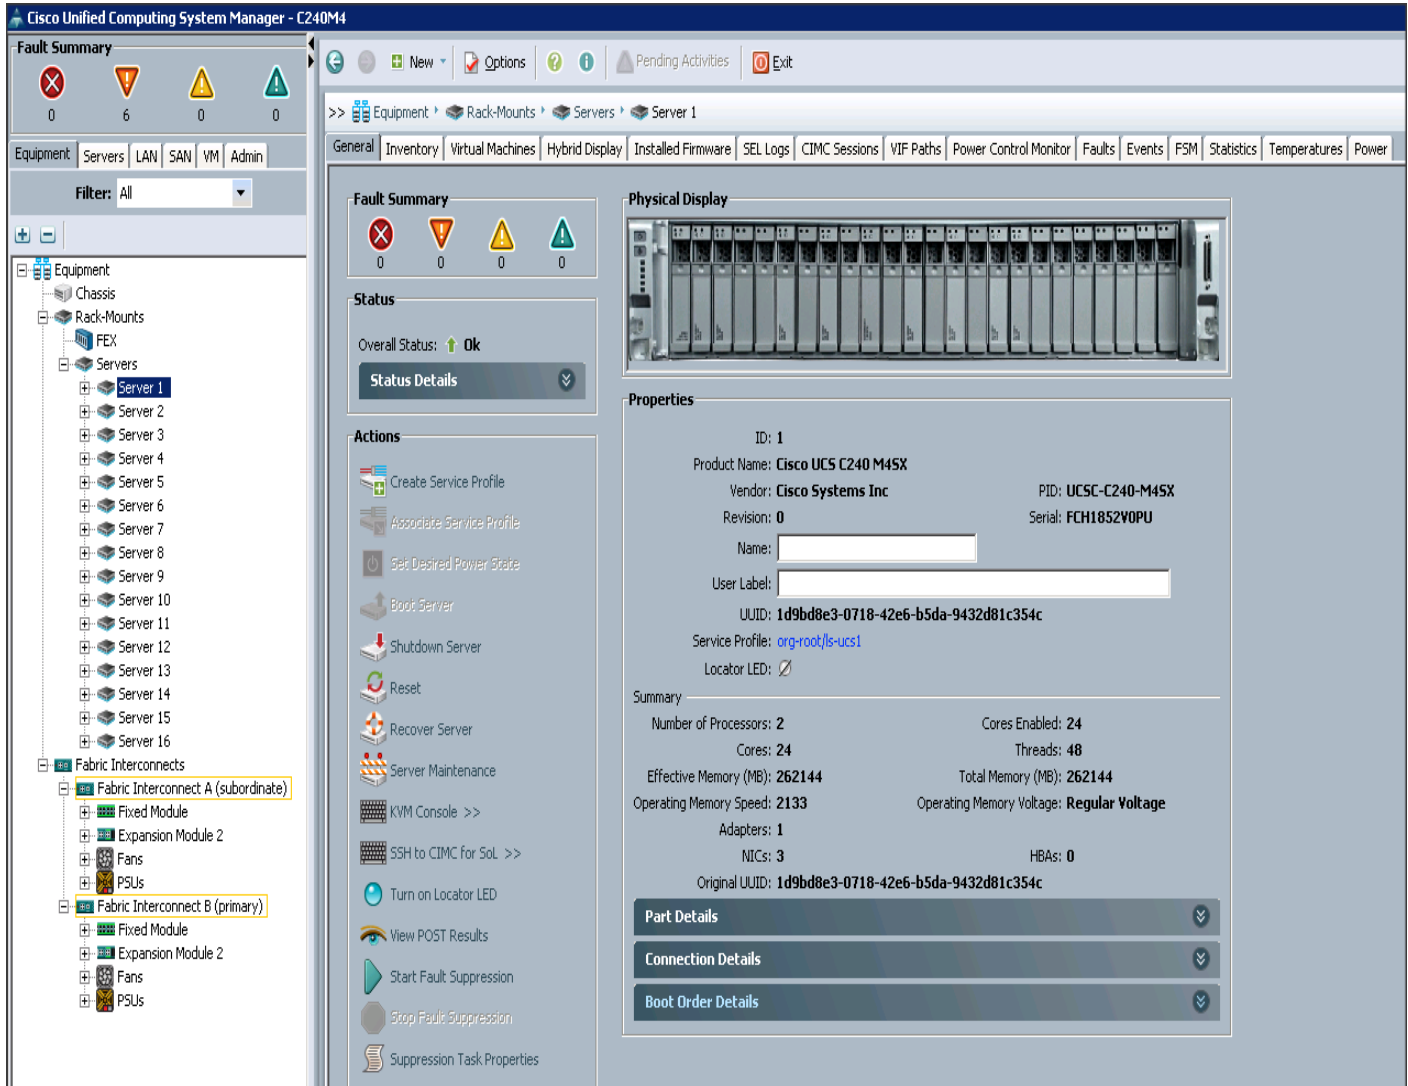


Cisco UCS Manager

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 4 Cisco UCS Manager



Cloudera's Distribution Including Apache Hadoop (CDH) 5.3.2

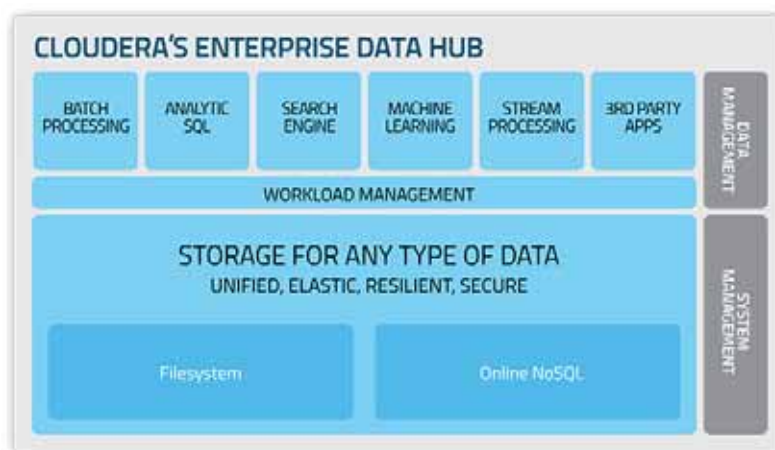
Built on the transformative Apache Hadoop open source software project, Cloudera Enterprise is an hardened distribution of Apache Hadoop and related projects that are designed to meet the demanding needs of enterprise customers. Cloudera is one of the largest contributors to the Hadoop ecosystem, and has created a rich suite of complementary open source projects that are included in Cloudera Enterprise.

All the facets of integration and the entire solution is thoroughly tested and documented. The readily available pre-integrated solution in CDH helps in building your Hadoop deployment model with ease. CDH provides a coherent and practical solution to solve real business problems.

Cloudera Enterprise, with Apache Hadoop at the core, is:

- **Unified** – one integrated system, bringing diverse users and application workloads to one pool of data on common infrastructure; no data movement required
- **Secure** – perimeter security, authentication, granular authorization, and data protection
- **Governed** – enterprise-grade data auditing, data lineage, and data discovery
- **Managed** – native high-availability, fault-tolerance and self-healing storage, automated backup and disaster recovery, and advanced system and data management
- **Open** – Apache-licensed open source to ensure your data and applications remain yours, and an open platform to connect with all of your existing investments in technology and skills

Figure 5 Enterprise Data Hub Architecture by Cloudera



Cloudera provides a scalable, flexible, integrated platform that makes it easy to manage rapidly increasing volumes and varieties of data in your enterprise. Industry-leading Cloudera products and solutions enable you to deploy and manage Apache Hadoop and related projects, manipulate and analyze your data, and keep that data secure and protected.

Cloudera provides the following products and tools:

- **CDH**—The Cloudera distribution of Apache Hadoop and other related open-source projects, including Cloudera Impala and Cloudera Search. CDH also provides security and integration with numerous hardware and software solutions.
 - **Cloudera Impala**—A massively parallel processing SQL engine for interactive analytics and business intelligence. Its highly optimized architecture makes it ideally suited for traditional BI-style queries with joins, aggregations, and subqueries. It can query Hadoop data files from a variety of sources, including those produced by MapReduce jobs or loaded into Hive tables. The YARN and Llama resource management components let Impala coexist on clusters running batch workloads concurrently with Impala SQL queries. You can manage Impala alongside other Hadoop components through the Cloudera Manager user interface, and secure its data through the Sentry authorization framework.
 - **Cloudera Search**—Provides near real-time access to data stored in or ingested into Hadoop and HBase. Search provides near real-time indexing, batch indexing, full-text exploration and navigated drill-down, as well as a simple, full-text interface that requires no SQL or

programming skills. Fully integrated in the data-processing platform, Search uses the flexible, scalable, and robust storage system included with CDH. This eliminates the need to move large data sets across infrastructures to perform business tasks.

- **Cloudera Manager**—A sophisticated application used to deploy, manage, monitor, and diagnose issues with your CDH deployments. Cloudera Manager provides the Admin Console, a web-based user interface that makes administration of your enterprise data simple and straightforward. It also includes the Cloudera Manager API, which you can use to obtain cluster health information and metrics, as well as configure Cloudera Manager.
- **Cloudera Navigator**—An end-to-end data management tool for the CDH platform. Cloudera Navigator enables administrators, data managers, and analysts to explore the large amounts of data in Hadoop. The robust auditing, data management, lineage management, and life cycle management in Cloudera Navigator allow enterprises to adhere to stringent compliance and regulatory requirements.

Cisco Application Centric Infrastructure (ACI) Overview

ACI provides network the ability to deploy and respond to the needs of applications, both in the data center and in the cloud. The network must be able to deliver the right levels of connectivity, security, compliance, firewalls, and load balancing, and it must be able to do this dynamically and on-demand.

This is accomplished through centrally defined policies and application profiles. The profiles are managed by new Application Policy Infrastructure Controller [APIC] and distributed to switches like the Cisco Nexus 9000 Series. Cisco Nexus 9000 Series Switches and the Cisco Application Policy Infrastructure Controller (APIC) are the building blocks for ACI.

ACI is software-defined networking (SDN) plus a whole lot more. Most SDN models stop at the network. ACI extends the promise of SDN—namely agility and automation—to the applications themselves. Through a policy-driven model, the network can cater to the needs of each application, with security, network segmentation, and automation at scale. And it can do so across physical and virtual environments, with a single pane of management.

The ACI fabric supports more than 64,000 dedicated tenant networks. A single fabric can support more than one million IPv4/IPv6 endpoints, more than 64,000 tenants, and more than 200,000 10G ports. The ACI fabric enables any service (physical or virtual) anywhere with no need for additional software or hardware gateways to connect between the physical and virtual services and normalizes encapsulations for Virtual Extensible Local Area Network (VXLAN) / VLAN / Network Virtualization using Generic Routing Encapsulation (NVGRE).

The ACI fabric decouples the endpoint identity and associated policy from the underlying forwarding graph. It provides a distributed Layer 3 gateway that ensures optimal Layer 3 and Layer 2 forwarding. The fabric supports standard bridging and routing semantics without standard location constraints (any IP address anywhere), and removes flooding requirements for the IP control plane Address Resolution Protocol (ARP) / Generic Attribute Registration Protocol (GARP). All traffic within the fabric is encapsulated within VXLAN.

Architectural Benefits of using Fabric Interconnect with Cisco ACI

UCS Servers are connected directly to Fabric Interconnect (FI) which in-turn connects to ACI (N9K switches). This mode allows using the UCS Manager capabilities in FI for provisioning the servers within a domain. This topology can scale up to 5760 servers for a fully populated pair of Nexus 9508s with all the eight linecards; details of which are discussed in “Scaling section”. Benefits of ACI architecture are discussed in the next section.

Centralized Management for the Entire Network

Cisco ACI treats the network as a single entity rather than a collection of switches. It uses a central controller to implicitly automate common practices such as Cisco ACI fabric startup, upgrades, and individual element configuration. The Cisco Application Policy Infrastructure Controller (Cisco APIC) is this unifying point of automation and management for the Application Centric Infrastructure (ACI) fabric. This architectural approach dramatically increases the operational efficiency of networks, by reducing the time and effort needed to make modifications to the network and, also, for root cause analysis and issue resolution.

Performance Oriented Fabric

The Cisco ACI Fabric incorporates numerous capabilities that can help provide performance improvements to applications.

Dynamic Load Balancing (DLB): The ACI fabric provides several load balancing options for balancing the traffic among the available uplinks. Static hash load balancing is the traditional load balancing mechanism used in networks where each flow is allocated to an uplink based on a hash of its 5-tuple. This load balancing gives a distribution of flows across the available links that is roughly even. Usually, with a large number of flows, the even distribution of flows results in an even distribution of bandwidth as well. However, if a few flows are much larger than the rest, static load balancing might give suboptimal results. Dynamic load balancing (DLB) adjusts the traffic allocations according to congestion levels. It measures the congestion across the available paths and places the flows on the least congested paths, which results in an optimal or near optimal placement of the data.

Dynamic Packet Prioritization (DPP), while not a load balancing technology, uses some of the same mechanisms as DLB in the switch. DPP configuration is exclusive of DLB. DPP prioritizes short flows higher than long flows; a short flow is less than approximately 15 packets. Short flows are more sensitive to latency than long ones. DPP can improve overall application performance.

Together these technologies enable performance enhancements to applications, including Big Data workloads. More information on these technologies and results associated with performance analysis can be found in the following paper that recently won the best paper award at ACM SIGCOMM 2014:

Application-Centric Policy Model

The Cisco ACI policy model is designed top down using a promise theory model to control a scalable architecture of defined network and service objects. This model provides robust repeatable controls, multitenancy, and minimal requirements for detailed knowledge by the control system known as the Cisco APIC. The model is designed to scale beyond current needs to the needs of private clouds, public clouds, and software-defined data centers.

The policy enforcement model within the fabric is built from the ground up in an application-centric object model. This provides a logical model for laying out applications, which will then be applied to the fabric by the Cisco APIC. This helps to bridge the gaps in communication between application requirements and the network constructs that enforce them. The Cisco APIC model is designed for rapid provisioning of applications on the network that can be tied to robust policy enforcement while maintaining a workload anywhere approach.

Multi-Tenant and Mixed Workload Support

Cisco ACI is built to incorporate secure multi-tenancy capabilities. The fabric enables customers to host multiple concurrent Big Data workloads on a shared infrastructure. ACI provides the capability to enforce proper isolation and SLA's for workloads of different tenants. These benefits extend beyond multiple Big Data workloads – Cisco ACI allows the same cluster to run a variety of different application workloads, not just Big Data, with the right level of security and SLA for each workload.

Extensibility and Openness

ACI supports an open ecosystem embracing open APIs, open source, and open standards. This provides the broadest choice in data center management and infrastructure. ACI supports embracing open APIs, open source, and open standards. This provides the broadest choice in data center management and infrastructure.

Easy Migration to 40Gbps in the Network

Cisco QSFP BiDi technology removes 40-Gbps cabling cost barriers for migration from 10-Gbps to 40-Gbps connectivity in data center networks. Cisco QSFP BiDi transceivers provide 40-Gbps connectivity with immense savings and simplicity compared to other 40-Gbps QSFP transceivers. The Cisco QSFP BiDi transceiver allows organizations to migrate the existing 10-Gbps cabling infrastructure to 40 Gbps at no cost and to expand the infrastructure with low capital investment. Together with Cisco Nexus 9000 Series Switches, which introduce attractive pricing for networking devices, Cisco QSFP BiDi technology provides a cost-effective solution for migration from 10-Gbps to 40-Gbps infrastructure.

Cisco ACI Building blocks

Cisco ACI consists of:

- The Cisco Nexus 9000 Series Switches.
- A centralized policy management and Cisco Application Policy Infrastructure Controller (APIC).

Cisco Nexus 9000 Series Switches

The 9000 Series Switches offer both modular (9500 switches) and fixed (9300 switches) 1/10/40/100 Gigabit Ethernet switch configurations designed to operate in one of two modes:

- Cisco NX-OS mode for traditional architectures and consistency across the Cisco Nexus portfolio.
- ACI mode to take full advantage of the policy-driven services and infrastructure automation features of ACI.

The ACI-Ready Cisco Nexus 9000 Series provides:

- Accelerated migration to 40G: zero cabling upgrade cost with Cisco QSFP+ BiDi Transceiver Module innovation.
- Switching platform integration: Nexus 9000 Series enables a highly scalable architecture and is software upgradable to ACI.
- Streamline Application Management: Drastically reduce application deployment time and get end to end application visibility.

This architecture consists of Cisco Nexus 9500 series switches acting as the spine and Cisco Nexus 9300 series switches as leaves.

Cisco Nexus 9508 Spine Switch

The Cisco Nexus 9508 Switch offers a comprehensive feature set, high resiliency, and a broad range of 1/10/40 Gigabit Ethernet line cards to meet the most demanding requirements of enterprise, service provider, and cloud data centers. The Cisco Nexus 9508 Switch is an ACI modular spine device enabled by a non-blocking 40 Gigabit Ethernet line card, supervisors, system controllers, and power supplies.

The Cisco Nexus 9500 platform internally uses a Clos fabric design that interconnects the line cards with rear-mounted fabric modules. The Cisco Nexus 9500 platform supports up to six fabric modules, each of which provides up to 10.24-Tbps line-rate packet forwarding capacity. All fabric cards are directly connected to all line cards. With load balancing across fabric cards, the architecture achieves optimal bandwidth distribution within the chassis.

Figure 6 Cisco Nexus 9508 Switch



ACI Spine Line Card for Cisco Nexus 9508

There are multiple spine line cards supported on Nexus 9508. This architecture uses N9K-X9736PQ: 40 Gigabit Ethernet ACI Spine Line Card.

- 36-port 40 Gigabit Ethernet QSFP+ line card
- Non-blocking
- Designed for use in an ACI spine switch role
- Works only in ACI mode
- Cannot mix with non-spine line cards
- Supported in 8-slot chassis

Figure 7 Cisco N9K-X9736PQ Linecard



Cisco Nexus 9396 Leaf Switch

The Cisco Nexus 9396PX Switch delivers comprehensive line-rate layer 2 and layer 3 features in a two-rack-unit (2RU) form factor. It supports line rate 1/10/40 GE with 960 Gbps of switching capacity. It is ideal for top-of-rack and middle-of-row deployments in both traditional and Cisco Application Centric Infrastructure (ACI)-enabled enterprise, service provider, and cloud environments.

Figure 8 Cisco Nexus 9396PX Switch



Cisco Application Policy Infrastructure Controller (APIC)

The Application Centric Infrastructure is a distributed, scalable, multitenant infrastructure with external end-point connectivity controlled and grouped through application-centric policies. The APIC is the unified point of automation, management, monitoring, and programmability for the Cisco Application Centric Infrastructure. The APIC supports the deployment, management, and monitoring of any application anywhere, with a unified operations model for physical and virtual components of the infrastructure. The APIC programmatically automates network provisioning and Control that is based on the application requirements and policies. It is the central control engine for the broader cloud network; it simplifies management and allows flexibility in how application networks are defined and automated. It also provides northbound REST APIs. The APIC is a distributed system that is implemented as a cluster of many controller instances.

Figure 9 APIC Appliance



ACI Topology

ACI topology is spine-leaf architecture. Each leaf is connected to each spine. It uses internal routing protocol; Intermediate System to Intermediate System (IS-IS) to establish IP connectivity throughout the fabric among all the nodes including spine and leaf. To transport tenant traffic across the IP fabric, integrated VxLAN overlay is used. The broadcast ARP traffic coming from the end point or hosts to the leaf are translated to unicast ARP in the fabric.

The forwarding is done as a host based forwarding. In the leaf layer the user information such as username, IP address, locations, policy groups etc., are decoupled from the actual forwarding path and encode them into the fabric VxLAN header and is forwarded to the desired destination.


Each spine has the complete forwarding information about the end hosts that are connected to the fabric and on every leaf have the cached forwarding information. The leaf only needs to know the hosts it needs to talk to. For example if Server Rack-1 has to send some information to Server Rack-2, When packet comes in the ingress leaf (LEAF_1) it will encapsulate the information into the VxLAN header and

forward that information to LEAF_2. If the LEAF_1 does not have information about the LEAF_2, it uses Spine as a proxy and since Spine has all the complete information about the entire end host connected to the fabric, it will resolve the egress leaf and forward the packet to the destination.

To the outside world, routing protocols can be used to learn outside prefixes or static routing can be used instead. The outside learned routes will be populated into the fabric or to the other leafs with Multiprotocol BGP (M-BGP). In M-BGP topology the spine nodes acts as route reflectors.

The Network topology of ACI is as depicted below:

Figure 10 Network topology based on Cisco ACI



The Cisco ACI infrastructure incorporates the following components:

- Two Cisco Nexus 9508 Spine Switch
 - ACI Spine Line Card for Nexus 9508
- Cisco Nexus 9396 Leaf Switch for Data Traffic
- Cisco APIC-L1-Cluster with three APIC-L1 appliances

Solution Overview

This CVD describes architecture and deployment procedures for Cloudera (CDH 5.3.2) on 160 Cisco UCS C240 M4 server based on Cisco UCS Integrated Infrastructure for Big Data with two domains (each Fabric-Interconnect domain has 80 servers under a pair of Fabric Interconnect) interconnected through ACI. The Cisco UCS Integrated Infrastructure with ACI brings together a highly scalable architecture designed to meet a variety of scale-out application demands with seamless data integration and management integration capabilities.



Note

System Architecture and Scaling sections discussed below describe three Cisco UCS Fabric Interconnect domains under a pair for Cisco Nexus 9396.

Further, the CVD describes in detail the process of creating the Application Network Profile in the ACI for Big Data application. Application Network Profiles (Application Network Profile is a collection of EPGs, their connections, and the policies that define those connections described in detail later) are the logical representation of an application (here Big Data) and its interdependencies in the network fabric.

Application Network Profiles are designed to be modeled in a logical way that matches the way that applications are designed and deployed. The configuration and enforcement of policies and connectivity is handled by the system rather than manually by an administrator.

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 *Compute Nodes used for the Big Data Cluster with ACI*

Performance Optimized	Capacity Optimized
16 Cisco UCS C240 M4 Rack Servers (SFF), each with:	16 Cisco UCS C240 M4 Rack Servers (LFF), each with:
2 Intel Xeon processors E5-2680 v3 CPUs	2 Intel Xeon processors E5-2620 v3 CPUs
256 GB of memory	128 GB of memory
Cisco 12-Gbps SAS Modular Raid Controller with 2-GB flash-based write cache (FBWC)	Cisco 12-Gbps SAS Modular Raid Controller with 2-GB FBWC
24 1.2-TB 10K SFF SAS drives (460 TB total)	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	2 120-GB 6-Gbps 2.5-inch Enterprise Value SATA SSDs for Boot
Cisco UCS VIC 1227 (with 2 10 GE SFP+ ports)	Cisco UCS VIC 1227 (with 2 10 GE SFP+ ports)

This CVD uses Performance Optimized configuration.



Note

This CVD describes the install process of CDH 5.3.2 for a 160 node (3 Master nodes in High Availability + 157 Data node) of Performance Optimized Cluster configuration.

The Performance cluster configuration consists of the following:

- Four Cisco UCS 6296UP Fabric Interconnects
- 160 UCS C240 M4 Rack-Mount servers (16 per rack)
- Ten Cisco R42610 standard racks
- Eighteen Vertical Power distribution units (PDUs) (Country Specific)
- Two Cisco Nexus 9508 Spine Switch
 - ACI Spine Line Card for Nexus 9508
- Cisco Nexus 9396 Leaf Switch for Data Traffic
- Cisco APIC-L1-Cluster with three APIC-L1 appliances

Table 2 *Hardware Component Details*

Hardware	Role	Quantity
NK-C9508	Spine	2

Table 2 Hardware Component Details

N9K-X9736PQ	36 ports 40 Gig QSFP+ Line Card for the Spine	2
N9K-C9396PX	Leaf	2
UCS FI 6296UP	Fabric Interconnect	4
APIC-L1	APIC Appliance	3
UCS C240 M4	Rack Server	160
QSFP-H40G	40 Gig connectivity	26
SFP-H10GB	10 Gig Connectivity	320(Servers) +3 (APICs) + 56 (FI Uplink)

**Note**

For more details on Connecting Application Centric Infrastructure (ACI) to Outside Layer 2 and 3 Networks can be found at:

<http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrast-structure/white-paper-c07-732033.html>

Physical Layout for the Solution

Physical Layout for the solution is as shown in the following table. Each rack consists of two vertical PDUs. The solution consists of 5 Cisco R42610 racks. The Nexus 9396 leaf switch and the Fabric Interconnect is distributed across rack1 and rack2, the APIC appliances are distributed across rack2 to rack4. Similarly, nexus 9508 spine switch is mounted in rack2 for easier caballing between the spine and leaf switches. The rest of the spaces in the 5 racks are used for mounting 80 servers. All the Switches and UCS Servers are dual connected to vertical PDUs for redundancy; thereby, ensuring availability during power source failure.

For second pod, only two FI's are required in this domain because the uplink from the FI is connected to the leaf switches in pod1 and rest of the space is used to mount another 80 servers.

Table 3 Rack 1-5

Slot	Rack 1	Rack 2	Rack 3	Rack 4	Rack 5
1	N9K-C9396PX	N9K-C9396PX	APIC-L1	APIC-L1	UCS C240M4
2					
3	FI-A	FI- B	UCS C240M4	UCS C240M4	UCS C240M4
4					
5	UCS C240M4	APIC-L1	UCS C240M4	UCS C240M4	UCS C240M4
6					
7	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
8					
9	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
10					

Table 3 **Rack 1-5**

11	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
12					
13	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
14					
15	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
16					
17	UCS C240M4	N9k-C9508	UCS C240M4	UCS C240M4	UCS C240M4
18					
19	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
20					
21	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
22					
23	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
24					
25	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
26					
27	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	
28					
29	UCS C240M4	N9k-C9508	UCS C240M4	UCS C240M4	UCS C240M4
30					
31	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
32					
33	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
34					
35	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
36					
37	UCS C240M4		UCS C240M4	UCS C240M4	UCS C240M4
38					
39	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	
40					
41	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	
42					

Table 4 Rack 1-5

Slot	Rack 1	Rack 2	Rack 3	Rack 4	Rack 5
1		FI-C	FI-D		
2					
3					
4					
5					
6					
7					
8					
9					
10					
11	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
12					
13	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
14					
15	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
16					
17	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
18					
19	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
20					
21	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
22					
23	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
24					
25	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
26					
27	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
28					
29	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
30					
31	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
32					
33	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
34					

Table 4 *Rack 1-5*

35	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
36					
37	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
38					
39	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
40					
41	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4	UCS C240M4
42					



Note

Number of servers in a Rack can be reduced to 16 or less based on the power requirements.

Software Distributions and Versions

The software distribution's required versions are listed below.

Cloudera Enterprise

The Cloudera software for Cloudera Distribution for Apache Hadoop is version 5.3.2. For more information on CDH, visit: www.cloudera.com

Red Hat Enterprise Linux (RHEL)

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

Software Versions

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

Table 5 *Software Component Details*

Layer	Component	Version or Release
Network	Cisco ACI OS	11.0 (2m)
	APIC OS	1.0 (1e)
	Cisco UCS 6296UP	UCS 2.2(3d)A
	Cisco UCS VIC1227 Firmware	4.0(1d)
	Cisco UCS VIC1227 Driver	2.1.1.66
Compute	Cisco UCS C240-M4	C240M4.2.0.3d
Storage	LSI SAS 3108	24.5.0-0020
Software	Red Hat Enterprise Linux Server	6.5 (x86_64)
	CDH	5.3.2
	UCS Manager	2.2(3d)

**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&relind=AVAILABLE&rellifecycle=&reltype=latest>
- The Latest Supported RAID controller Driver is already included with the RHEL 6.5 operating system.
- C240/C220 M4 Rack Servers are supported from UCS firmware 2.2(3d) onwards.

System Architecture

The ACI fabric consists of three major components: the Application Policy Infrastructure Controller (APIC), spine switches, and leaf switches. These three components handle both the application of network policy and the delivery of packets.

The system architecture consists of 3 domains (3 pair of FIs) connecting to ACI having two Cisco Nexus 9508 switches acting as a Spine and two Cisco Nexus 9396 as the leaf switches and three APIC-L1 as an APIC appliance.

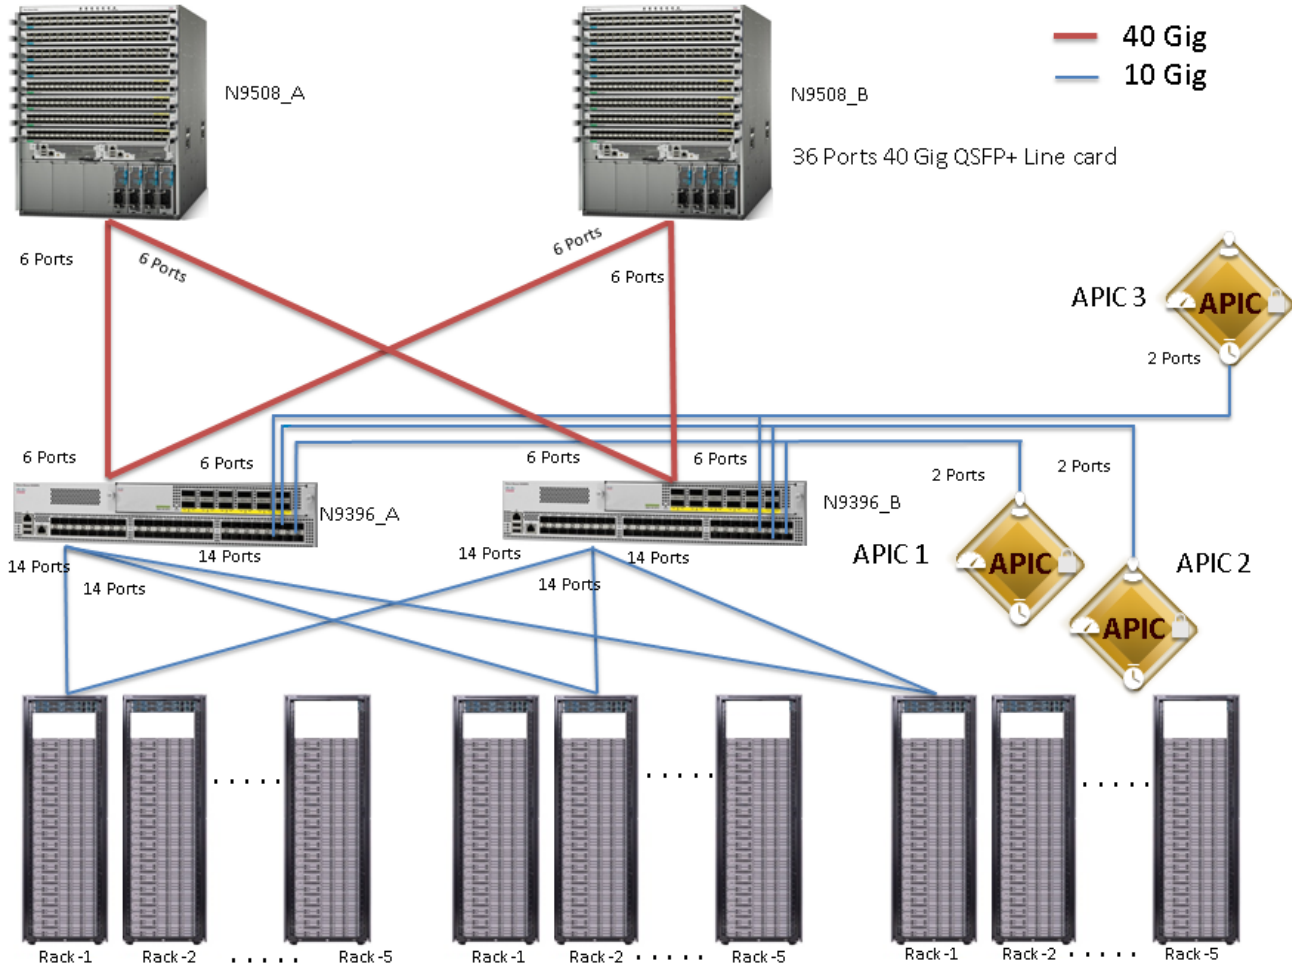
System architecture can be explained as:

- The 80 server are rack mounted and are connected to a pair of Cisco UCS FIs representing a domain through 10GE link (dual 10GE link to a pair of FI).
- 3 such domains are connected to a pair of Nexus 9396 which is the ACI Fabric leaf nodes. Here 10GEx14 links from each FI are connected to Nexus 9396. This is done through a port-channel of 7 ports connected to each of the Nexus 9396.

- Nexus 9396 receives the 14x10GE from each pair of FI as a vPC (Virtual Port-Channel), that is, all the 7 ports set from each of the FIs as an uplink to the leaf. There are 6 vPC for the 3 domains in each of 9396 connecting to the 3 pair of FIs.
- Each leaf is connected to Spines via 12 x 40 Gig connectivity cables.
- The three APIC's are connected to two leaves (Nexus 9396) via 10 Gig SFP cable.


The figure below shows the overall system architecture and physical layout of the solution.

Figure 11 System Architecture




The figure below show the connectivity between the leaf switches and fabric interconnect, where port channeling has been configured on Fabric Interconnect. This port channeling helps to aggregate the bandwidth towards the uplink leaf switches.

Figure 12 Fabric Interconnect Connectivity




The figure below show the connectivity between the leaf switches and fabric interconnect, where vPC has been configured on leaf switches through the APIC. These vPC ports are the same ports that were configured as port-channel in fabric interconnect.

Figure 13 vPC Connectivity



The figure below shows the connectivity between the one C240 M4 servers and two Fabric Interconnects.

Figure 14 Cisco UCS C240 M4 Server Connectivity



Scaling the Architecture

Here the UCS Servers are directly connected to Cisco UCS Fabric Interconnect (FI) which in-turn connects to leaf switches (Nexus 9396). This mode allows using the UCS Manager capabilities in FI for provisioning the servers. Up to 5 Racks (each with 16 servers) are connected to a Pair of FI forming a single domain and three such domains are connected to a pair of Leaf 9396 (every domain or pair of Cisco FI has 14 uplinks to Nexus 9396). This topology has no network over-subscription within a domain (80 servers under a pair of FI). The over-subscription ratio between domains is 5.7:1 and can scale up to 5760 servers for a fully populated pair of Nexus 9508 with all the 8 line cards in use.

Figure 15 Scaling Architecture

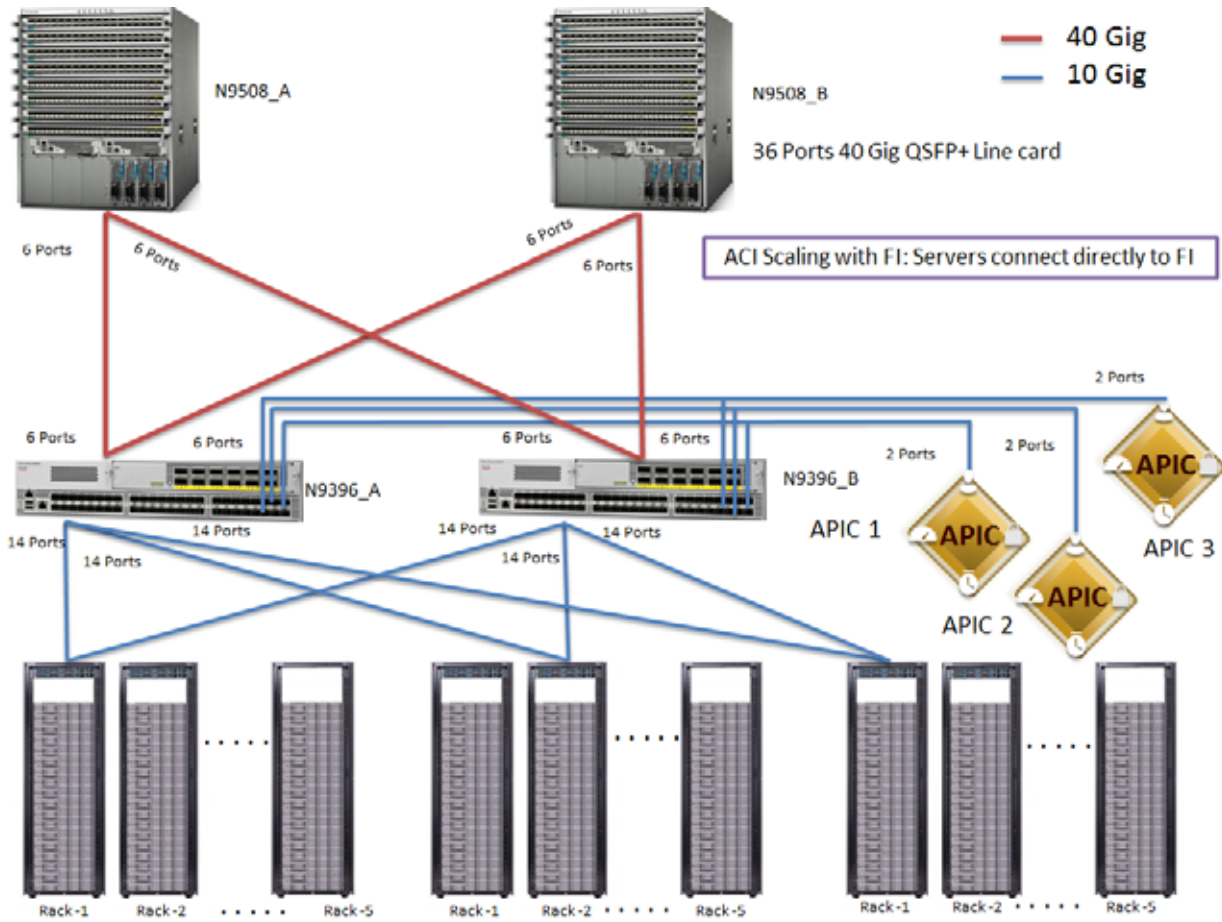


Table 6 *Cisco Nexus 9508 – Cisco Nexus 9396PX Connectivity*

SPINE	Line Card Pair	Ports Used	POD	Servers	LEAF
N9508_A	Line Card 1	1-6			9396_1A
	Line Card 1	7-12	1	240	9396_1B
	Line Card 1	13-18			9396_2A
	Line Card 1	19-24	2	480	9396_2B
	Line Card 1	25-30			9396_3A
	Line Card 1	31-36	3	720	9396_3B


	Line Card 8	1-6			9396_22A
	Line Card 8	7-12	22	5280	9396_22B
	Line Card 8	13-18			9396_23A
	Line Card 8	19-24	23	5520	9396_23B
	Line Card 8	25-30			9396_24A
	Line Card 8	31-36	24	5760	9396_24B
	N9508_B	Line Card 1	1-6		
Line Card 1		7-12	1	240	9396_1B
Line Card 1		13-18			9396_2A
Line Card 1		19-24	2	480	9396_2B
Line Card 1		25-30			9396_3A
Line Card 1		31-36	3	720	9396_3B
...	
Line Card 8		1-6			9396_22A
Line Card 8		7-12	22	5280	9396_22B
Line Card 8		13-18			9396_23A
Line Card 1		19-24	23	5520	9396_23B
Line Card 8		25-30			9396_24A
Line Card 8		31-36	24	5760	9396_24B

Table 7 Cisco Nexus 9396 - Cisco Fabric Interconnect Connectivity

LEAF	Ports Used	FI	Servers
9396_1A	1-14	FI_1A	1-80
9396_1A	15-28	FI_2A	81-160
9396_1A	29-42	FI_3A	161-240
9396_1A	43	APIC	
	44-48	Unused	
9396_1B	1-14	FI_1B	1-80
9396_1B	15-28	FI_2B	81-160
9396_1B	29-42	FI_3B	161-240
9396_1B	43	APIC	
	44-48	Unused	

Scaling the Architecture Further with Additional Spines Switches

The physical network of the Cisco Application Centric Infrastructure is built around leaf-spine architecture. It is possible to scale this infrastructure, immensely, by adding additional Spine switches. The ACI infrastructure supports up to 12 spine switches.

Figure 16 Cisco ACI Fabric with Multiple Spine Switches

With a 12-spine design, each leaf switch can be connected up to 12 spine switches. Allowing for tens of thousands of servers to be part of this infrastructure – being interconnected by a non-blocking fabric.

Network Configuration

The network configuration includes configuring the APIC, leaf, spine switches and Fabric Interconnect and deploying various application profiles and policies. In order to achieve this we first need to register the connected Nexus 9K switches to the APIC so that these switches become the part of the ACI fabric. Once the switch is registered the communication between the spine and leaf are completed.

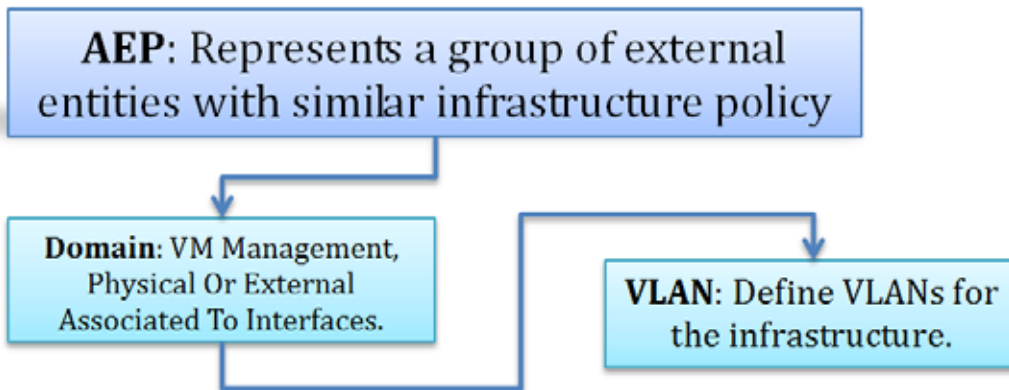
The admin is the only account enabled by default after the APIC is configured and it is always a good practice to create other user accounts with different privilege levels to make the APIC and the network secure. For this purpose we create a local or remote user depending on requirement.

Adding a management access is required in the ACI to let ACI know about any physical or virtual domain that is connected to it. By adding management access, APIC will control the physical interface and assign the policies to this interface. This is achieved by configuring Attachable Access Entity Profile (AEP). AEP requires having the domain and vlan pool that the ACI fabric will be using to communicate with various devices attached to it.



Note For more detail on AEP please refer “Adding Management Access” section.

Figure 17 Attachable Access Entity Profile for Communication with Other Devices



In this CVD, two pair of FIs representing two domains are connected to the pair of leaf switch. The uplink in the FIs is connected to the leaf via the port channeling (created in FI) and vPC is created at the leaf switches. The vPC allows single device to use a PortChannel across two upstream devices, eliminating Spanning Tree Protocol blocked ports which in turns provides a loop-free topology. With the use of vPC provides high availability and link-level resiliency.

Depending on the number of VLANs created in the FI, to trunk these vlans across the ACI fabric an Attachable Entity Profile (AEP) is required. An AEP provisions the VLAN pool (and associated VLANs) on the leaf, these VLAN pools are defined under the domain created within the AEP. A domain could be various external entities such as bare metal servers, hypervisors, VM management, Layer 2 or Layer 3 domains. The VLANs are not actually enabled on the port. No traffic flows unless an EPG is deployed on the port. An EPG acts as a separate entity which is analogous to VLAN. A tenant needs to be created before an EPG is defined.

A tenant contains policies that enable qualified users domain-based access control. Application profile, security policies and network are the elements of Tenants. An EPG for each VLAN is created under the application profile. Since EPG represent VLANs, a switch virtual interface (SVI) is needed to provide the Layer 3 processing for packets from all switch ports associated with the VLAN. A bridge domain

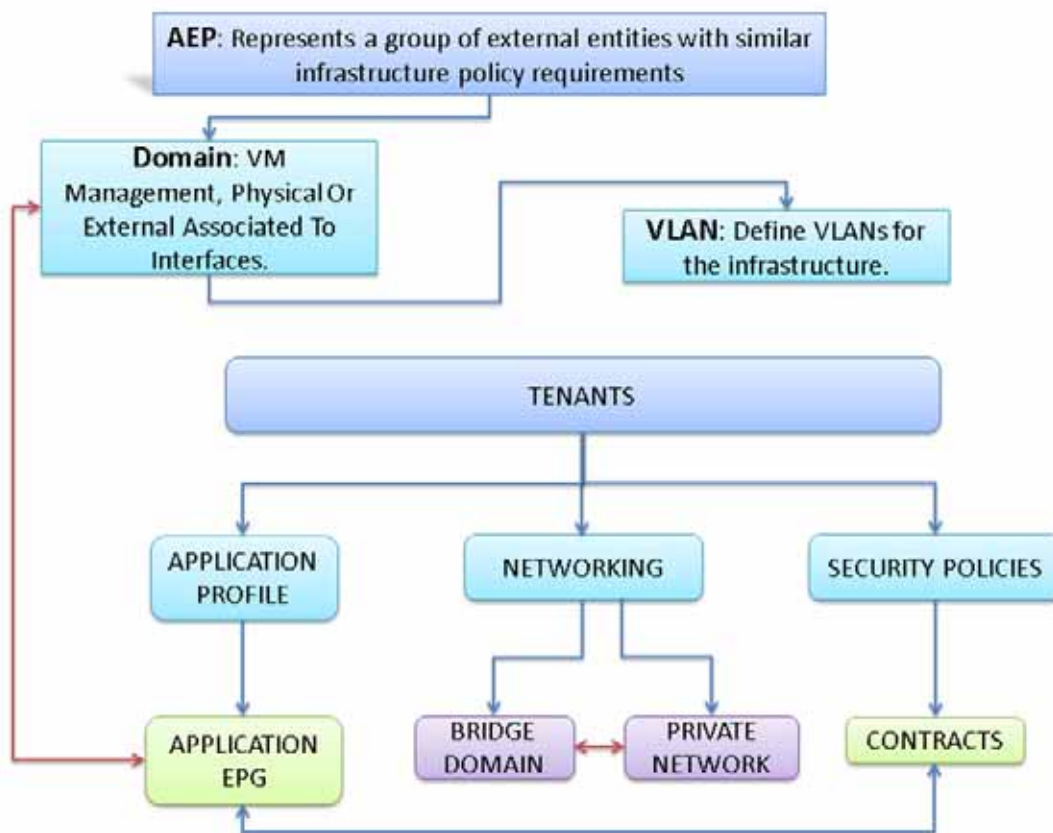
needs to be created which acts as switch virtual interface (SVI) for this purpose. Now, for the inter-Vlan communication, contracts need to be created to achieve communication among each EPG. Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers.



Note For more details on Tenant please refer to the “Adding Tenant” section.

The relationship between the AEP, its elements and tenants is show in the flowchart below.

Figure 18 Flowchart Showing AEP, AEP Elements and Tenants



IP Address Assignment

The IP address schemes of UCS and ACI management are configured as out of band management access through the management switch.

APIC 10.0.130.71/24

DOMAIN - 1	DOMAIN - 2
UCSM 10.0.141.5/24	UCSM 10.0.141.10/24
FI-A 10.0.141.6/24	FI-C 10.0.141.8/24

DOMAIN - 1	DOMAIN - 2
FI-B 10.0.141.7/24	FI-D 10.0.141.9/24
KVM 10.0.141.11/24 – 10.0.141.90/24	KVM 10.0.141.91/24 – 10.0.141.170/24

Table 8 IP Address Assignment for Domain 1 and 2

VLAN	Domain - 1	Domain - 2
VLAN 160	10.0.145.45 - 124/24	10.0.145.125 -204 /24
VLAN 11	10.0.146.45 -124 /24	10.0.146.125 -204/24
VLAN 12	10.0.147.45 -124 /24	10.0.147.125 -204/24

Configuration of APIC

This section describes loading and configuring the APIC.

Once the APIC appliance is booted for the first time, the APIC console presents a series of initial setup options. For many options, you can press Enter to choose the default setting that is displayed in brackets. At any point in the setup dialog, you can restart the dialog from the beginning by pressing Ctrl-C.

Shown below is the initial configuration of the APIC.

```

Enter the fabric name [ACI Fabric1]:
Enter the number of controllers in the fabric (1-9) [3]:3
Enter the controller ID (1-3) [1]:1
Enter the controller name [apic1]:APIC_1
Enter address pool for TEP addresses [10.0.0.0/16]:
Enter the VLAN ID for infra network (1-4094) [4]: 130

Out-of-band management configuration...
Enter the IP address for out-of-band management: 10.0.130.71/24
Enter the IP address of the default gateway [None]: 10.0.130.1
Administrator user configuration...
Enable strong passwords? [Y]
Enter the password for admin:
    
```

Below is the screenshot of the configuration

Figure 19 APIC Initial Configuration

```

10.0.130.68 - XVM Console
File View Macros Tools Power Virtual Media Help

Reenter the password for admin:

Cluster configuration ...
Fabric name: BIG_DATA
Number of controllers: 3
Controller name: APIC
Controller ID: 1
TEP address pool: 10.0.0.0/16
Infra VLAN ID: 130
Multicast address pool: 225.0.0.0/15

Out-of-band management configuration ...
Management IP address: 10.0.130.71/24
Default gateway: 10.0.130.1
Interface speed/duplex mode: auto

admin user configuration ...
Strong Passwords: N
User name: admin
Password: ****


The above configuration will be applied ...

Would you like to edit the configuration? (y/n) [n]:

```

Once the configuration is completed, the APIC will Boot its APIC IOS Image and will ask for the login information. The default username is “admin” and the password is the one that was set during the initial configuration.

Figure 20 *APIC Login Screen*



Switch Discovery with the APIC

The APIC is a central point of automated provisioning and management for all the switches that are part of the ACI fabric. A single data center might include multiple ACI fabrics, each with their own APIC cluster and Cisco Nexus 9000 Series switches that are part of the fabric. To ensure that a switch is managed only by a single APIC cluster, each switch must be registered with that specific APIC cluster that manages the fabric. The APIC discovers new switches that are directly connected to any switch it currently manages. Each APIC instance in the cluster first discovers only the leaf switch to which it is directly connected. After the leaf switch is registered with the APIC, the APIC discovers all spine switches that are directly connected to the leaf switch. As each spine switch is registered, that APIC discovers all the leaf switches that are connected to that spine switch. This cascaded discovery allows the APIC to discover the entire fabric topology in a few simple steps.

Switch Registration with the APIC Cluster

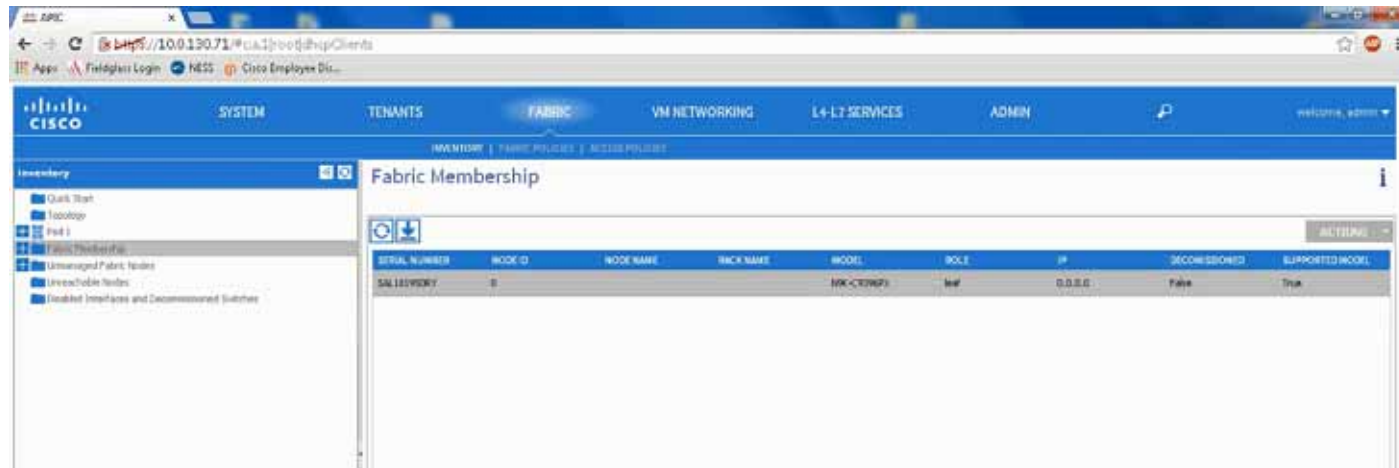
Once the switch is discovered by the APIC cluster it needs to be registered in the APIC to make it as a part of the fabric.

Prerequisite: All switches must be physically connected and booted with the correct ACI Image.

Using a web browser connect to the out-of-band management ip address [10.0.130.71] configured in the initial configuration.

1. On the menu bar, choose **FABRIC > INVENTORY**. In the Navigation pane, choose the appropriate pod.
2. In the Navigation pane, expand the pod, and click **Fabric Membership**. In the Work pane, in the Fabric Membership table, a single leaf switch is displayed with an ID of 0. It is the leaf switch that is connected to APIC.

Figure 21 Switch Discovery

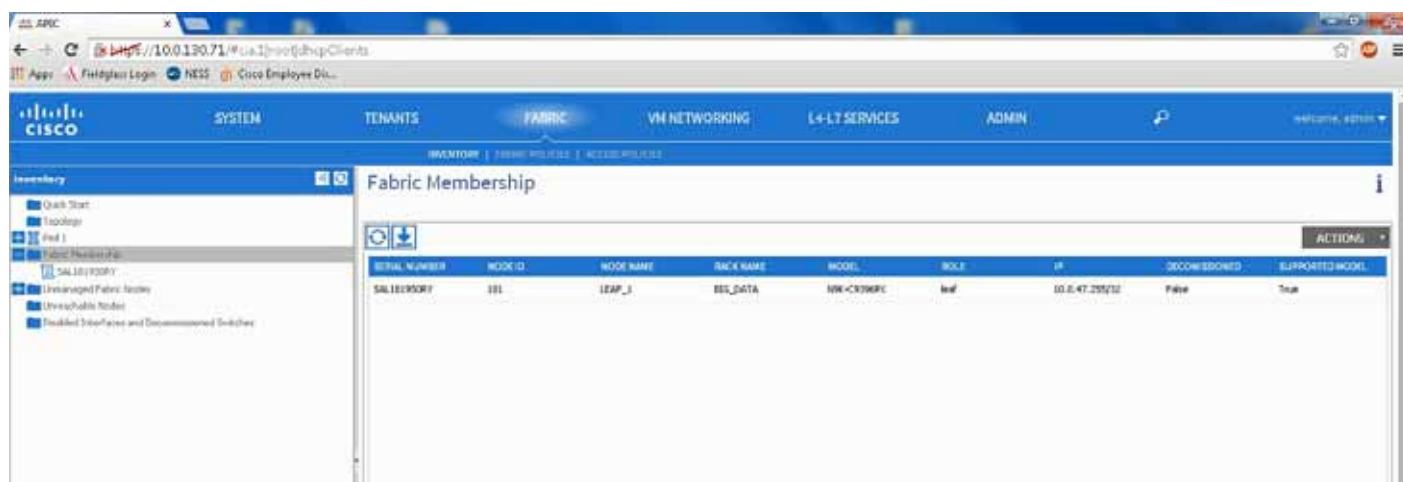


3. To configure the ID, double-click the leaf switch row, and perform the following actions:
 - a. In the ID field, add the appropriate ID (leaf1 is ID 101, leaf2 is ID 102 and leaf3 is ID103). The ID must be a number that is greater than 100 because the first 100 IDs are for APIC appliance nodes.
 - b. In the Switch Name field, add the name of the switch, and click **Update**. After an ID is assigned, it cannot be updated. The switch name can be updated by double-clicking the name and updating the Switch Name field.



Note The Success dialog box is displayed. An IP address gets assigned to the switch, and in the Navigation pane, the switch is displayed under the pod.

Figure 22 Switch Registration



4. Monitor the Work pane until one or more spine switches appear.
5. To configure the ID, double-click the spine switch row and perform the following actions:
 - a. In the ID field, add the appropriate ID (spine1 is ID 201 and spine 2 is ID 202). The ID must be a number that is greater than 100.

- b. In the Switch Name field, add the name of the switch, and click **Update**.
The Success dialog box is displayed. An IP address gets assigned to the switch, and in the Navigation pane, the switch is displayed under the pod. Wait until all remaining switches appear in the Node Configurations table.
- 6. For each switch listed in the Fabric Membership table, perform the following steps:
 - a. Double-click the switch, enter an ID and a Name, and click **Update**.
 - b. Repeat for the next switch in the list.

Validating the Switches

1. On the menu bar, choose **FABRIC > INVENTORY**, and in the Navigation pane, under Pod 1, expand Fabric Membership.
2. The switches in the fabric are displayed with their node IDs. In the Work pane, all the registered switches are displayed with the IP addresses that are assigned to them.


Figure 23 Switch Validation

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMMISSIONED	SUPPORT
FGE18200AW0	201	SPINE_1	BIG_DATA	N9K-C9508	spine	10.0.168.94/32	False	True
FGE18200AWL	202	SPINE_2	BIG_DATA	N9K-C9508	spine	10.0.168.65/32	False	True
SAL1816QWFA	103	LEAF_3	BIG_DATA	N9K-C93128TX	leaf	10.0.168.64/32	False	True
SAL181950M8	102	LEAF_2	BIG_DATA	N9K-C9396PX	leaf	10.0.168.93/32	False	True
SAL181950RY	101	LEAF_1	BIG_DATA	N9K-C9396PX	leaf	10.0.168.95/32	False	True

Validating the Fabric Topology

1. On the menu bar, choose **FABRIC > INVENTORY**.
2. In the Navigation pane, choose the pod that you want to view.
3. In the Work pane, click the **TOPOLOGY** tab.
The displayed diagram shows all attached switches, APIC instances, and links.
4. (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
To return to the topology diagram, in the upper left corner of the Work pane click the **Previous View** icon.
5. (Optional) To refresh the topology diagram, in the upper left corner of the Work pane, click the **Refresh** icon.

Figure 24 Fabric Topology



Creating User Accounts

The admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

1. On the menu bar, choose **ADMIN > AAA**.
2. In the Navigation pane, click **AAA Authentication**.
In the Work pane, the AAA Authentication dialog box is displayed.
3. Verify that in the default Authentication field, the Realm field displays as Local.

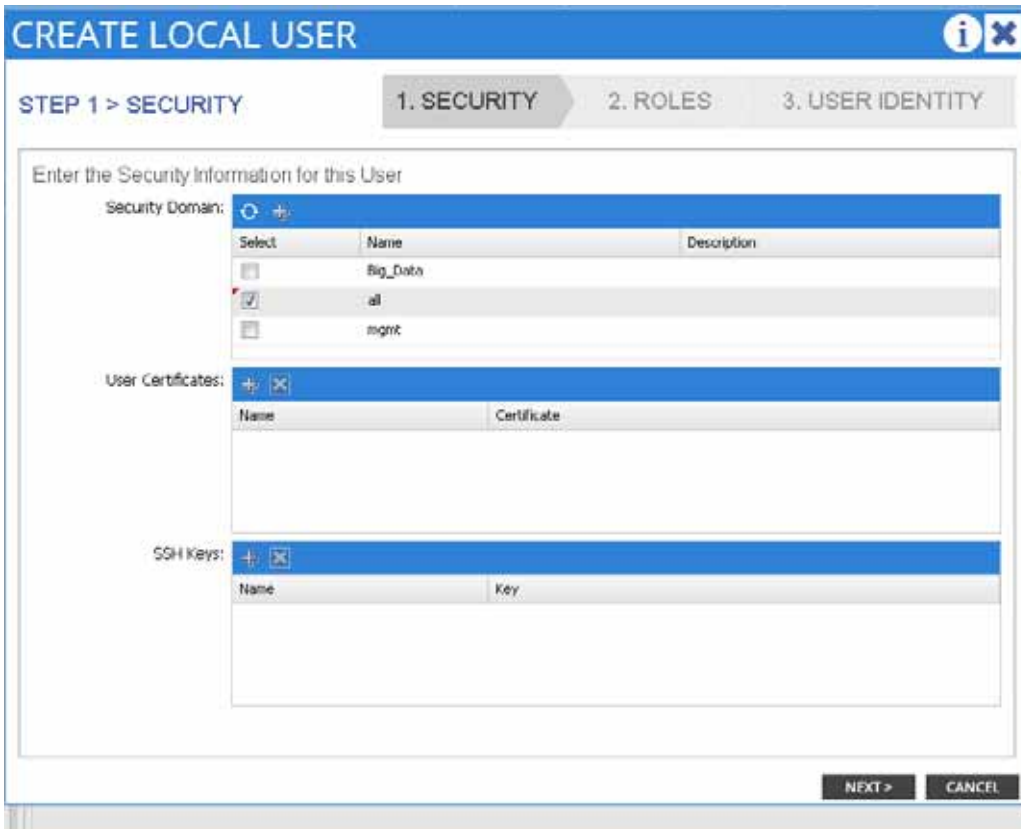
Figure 25 AAA Authentication

AAA Authentication

4. In the Navigation pane, expand **Security Management > Local Users**.
The admin user is present by default.

5. In the Navigation pane, right-click Create Local User.
The Create Local User dialog box opens.
6. Under the Security dialog box, choose the desired security domain for the user, and click **Next**.

Figure 26 Creating Local User



The Roles dialog box opens.

7. In the Roles dialog box, click the radio buttons to choose the roles for your user, and click **Next**.
You can provide read-only or read/write privileges.
8. In the User Identity dialog box, perform the following actions:
 - a. In the Login ID field, add an ID.
 - b. In the Password field, type the password.
 - c. In the Confirm Password field, confirm the password.
 - d. Click **Finish**.
 - e. Type other parameters if desired.

Figure 27 User Identity

CREATE LOCAL USER

STEP 3 > USER IDENTITY

1. SECURITY 2. ROLES 3. USER IDENTITY

Specify the User Identity

Login ID:

Password:

Confirm Password:

First Name:

Last Name:

Phone:

Email:

Description:

Account Status: Inactive Active

Account Expires: Yes No

FINISH **CANCEL**

- In the Navigation pane, click the name of the user that you created. In the Work pane, expand the + sign next to your user in the Security Domains area.

The access privileges for your user are displayed.

Adding Management Access

Attach Entity Profiles (AEP)

The ACI fabric provides multiple attachment points that connect through leaf ports to various external entities such as baremetal servers, hypervisors, Layer 2 switches (for example, the Cisco UCS fabric interconnect), and Layer 3 routers (for example Cisco Nexus 7000 Series switches). These attachment points can be physical ports, port channels, or a virtual port channel (vPC) on the leaf switches.

An attachable entity profile (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies, for example, Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), maximum transmission unit (MTU), and Link Aggregation Control Protocol (LACP).

An AEP is required to deploy any VLAN pools on the leaf switches. It is possible to reuse the encapsulation pools (for example, VLAN) across different leaf switches. An AEP implicitly provides the scope of the VLAN pool (associated to the VMM domain) to the physical infrastructure.



Note

- An AEP provisions the VLAN pool (and associated VLANs) on the leaf. The VLANs are not actually enabled on the port. No traffic flows unless an EPG is deployed on the port.

- Without VLAN pool deployment using an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
- A particular VLAN is provisioned or enabled on the leaf port based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter.
- A leaf switch does not support overlapping VLAN pools. Different overlapping VLAN pools must not be associated with the same AEP.

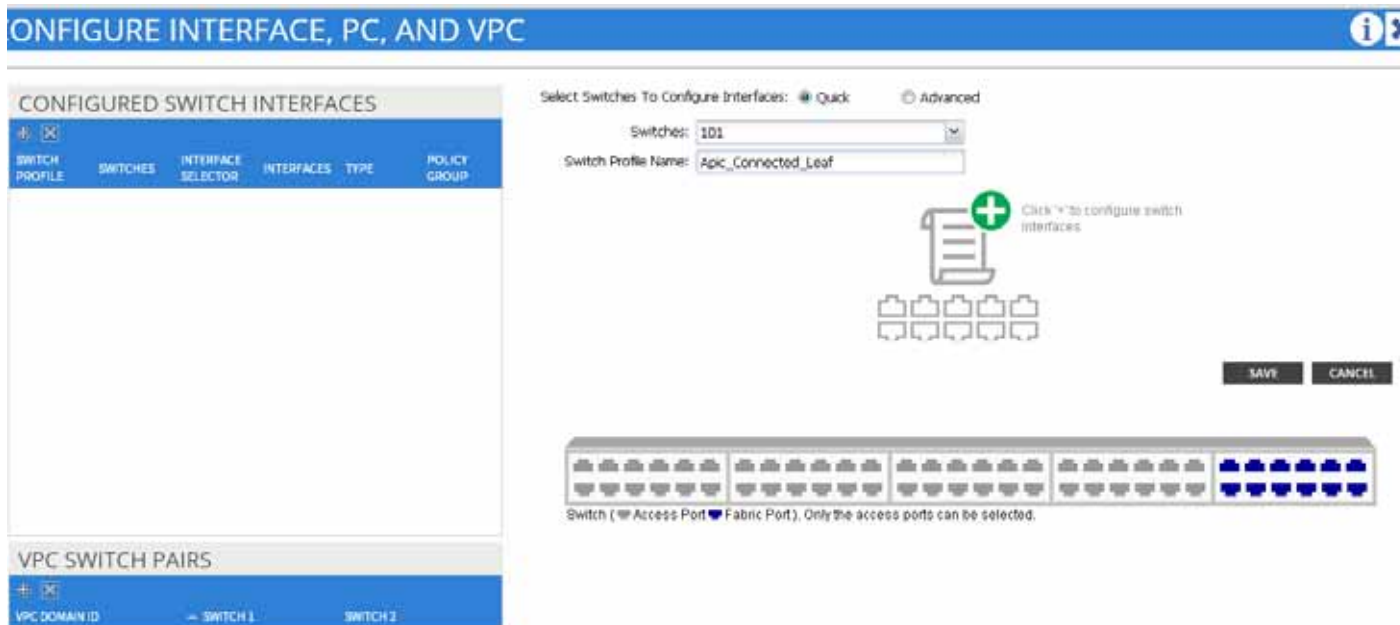
Configuring In-Band Management Access Using the GUI

The In-Band management access is required to establish the communication between the APIC and the ACI fabric.

1. On the menu bar, choose **FABRIC > Access Policies**. In the Work pane, click **Configure an Interface, PC and VPC**.
2. In the Configure Interface, PC, and VPC dialog box, click the large + icon next to the switch diagram to create a new profile and configure VLANs for the APIC.

In the Switches field, from drop-down list, check the check boxes for the switches to which the APICs are connected. (leaf1 and leaf2).

Figure 28 Configuring Interface, PC, and VPC



3. In the Switch Profile Name field, enter a name for the profile (Apic_Connected_Leaf).
4. Click the + icon to configure the ports.
5. Verify that in the Interface Type area, the **Individual** radio button is selected.

Figure 29 Configuring APIC Interface

The screenshot displays the 'CONFIGURE INTERFACE, PC, AND VPC' configuration page. On the left, there is a table for 'CONFIGURED SWITCH INTERFACES' with columns for SWITCH PROFILE, SWITCHES, INTERFACE SELECTOR, INTERFACES, TYPE, and POLICY GROUP. Below this is a section for 'VPC SWITCH PAIRS' with columns for VPC DOMAIN ID, SWITCH 1, and SWITCH 2. The main configuration area includes:

- 'Select Switches To Configure Interfaces' with radio buttons for 'Quick' (selected) and 'Advanced'.
- 'Switches' dropdown set to '101'.
- 'Switch Profile Name' dropdown set to 'Apic_Connected_Leaf'.
- 'Interface Type' radio buttons for 'Individual' (selected), 'PC', and 'VPC'.
- 'Interfaces' text field containing '1/48'.
- 'Interface Selector Name' text field containing 'Apic_Connected_Port'.
- 'Interface Policy Group' dropdown menu showing 'select type to create', 'INBAND', and 'Create Interface Policy Group'.
- A visual representation of a switch with 48 ports, where the first 48 ports are highlighted in blue.
- A note: 'Switch (☑ Access Port ▼ Fabric Port). Only the access ports can be selected.'

6. In the Interfaces field, enter the ports to which APICs are connected (1/48).
7. In the Interface Selector Name field, enter the name of the port profile (Apic_Connected_Port).
8. In the Interface Policy Group field, from drop-down list, choose **Create Interface Policy Group**.
9. In the Create Access Port Policy Group dialog box, perform the following actions:
 - a. In the Name field, enter the name of the policy group (INBAND).
You can leave the default values in the rest of the fields as they are.
 - b. In the Attached Entity Profile field, choose create **Attachable Access Entity Profile**.

Figure 30 Creating Access Port Policy Group

10. In the Create Attachable Access Entity Profile dialog box, perform the following actions:
 - a. In the Name field, enter a name (INBAND).
 - b. Expand Domains to be Associated to Interfaces field. In the Domain Profile field, from the drop-down list, choose Create Physical Domain.

Figure 31 Creating Attachable Access Entity Profile

- c. In the Create Physical Domain dialog box, in the Name field, enter the name (INBAND).
- d. In the VLAN Pool field, from the drop-down list, choose **Create VLAN Pool**.

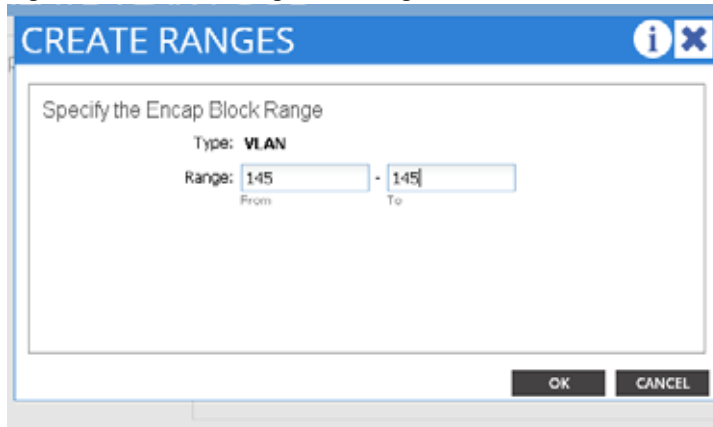
Figure 32 Creating Physical Domain

- e. In the Create VLAN Pool dialog box, in the Name field, enter the pool name (INBAND).
- f. In the Allocation Mode area, click the **Static Allocation** radio button.

Figure 33 Creating Vlan Pool

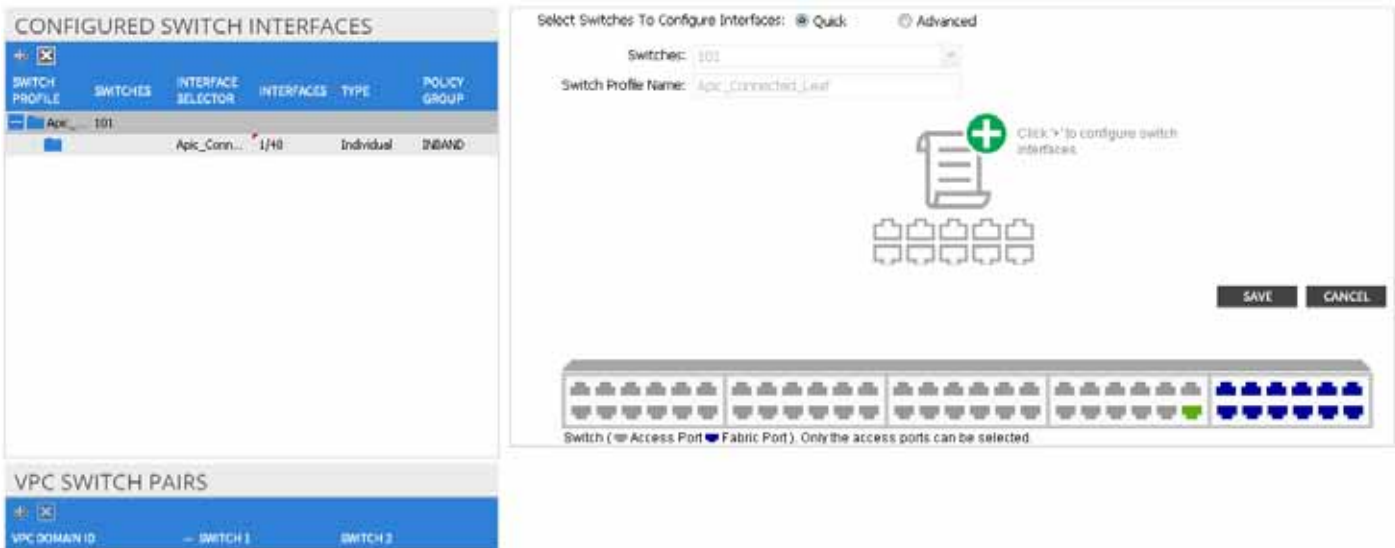
- g. Expand Encap Blocks. In the Create Ranges dialog box, in the Range fields, add a VLAN range (145-145).

Figure 34 Creating VLAN Range



- h. In the Create VLAN Pool dialog box, click **Submit**.
- i. In the Create Physical Domain dialog box, click **Submit**.
- j. In the Create Attachable Access Entity Profile dialog box, click **Update** and then **Submit**.
- k. In the Create Access Port Policy Group dialog box, click **Submit**.
- l. In the Configure Interface, PC, and VPC dialog box, click **Save**.

Figure 35 Saving the Configuration



The VLAN and the ports to which the APIC is connected are now configured.

Configuring VPC Ports for Fabric Interconnect

In order to configure vPC we need to create CDP Policy, LLDP Policy and LACP Policy that can be applied to the vPC ports.

- The APIC does not manage fabric interconnects and the rack servers, so these services must be configured from UCSM

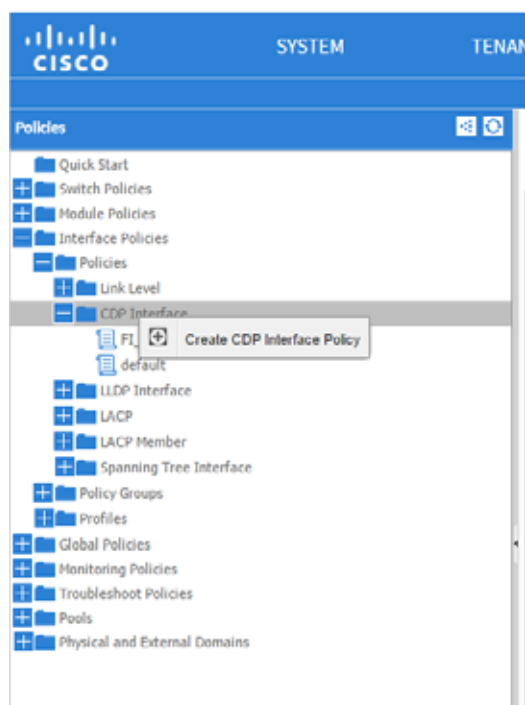
- Create VLAN pools that are associated on the fabric interconnect uplink to the leaf switch on the fabric interconnect.
- Cisco UCS C-series server when used along with ACI, Link Layer Discovery Protocol (LLDP) is not supported and must be disabled.
- Cisco Discovery Protocol (CDP) is disabled by default in the Cisco UCS Manager Fabric interconnects. In the Cisco UCS Manager, you must enable CDP by creating a policy under Network Control Policies > CDP.

The above steps are explained in detail further below

Creating CDP Policy group

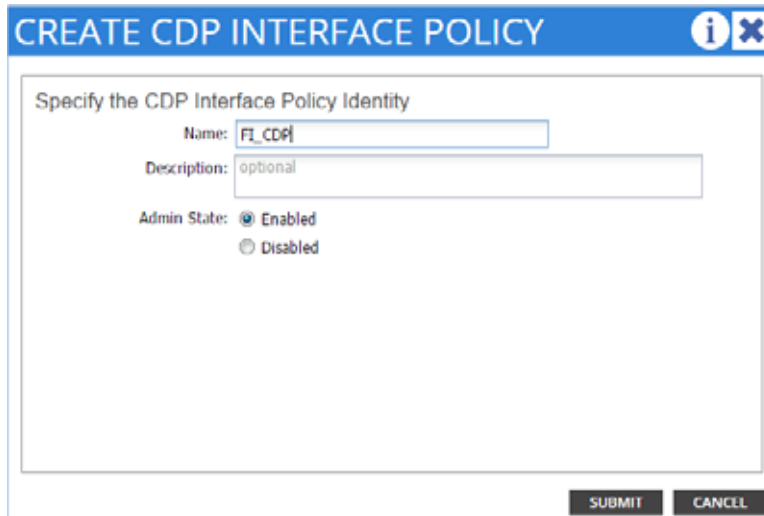
1. On the menu bar, choose **FABRIC > ACCESS POLICIES**.
2. In the Navigation pane, expand the Interface Policies and expand the Policies again.
3. Right-click the CDP Interface and select **Create CDP Interface Policy**.

Figure 36 Create CDP Interface Policy



4. In the Create CDP Interface Policy dialogue box, enter “FI_CDP” as the policy name, set Admin State “Enabled” and click **submit**.
5. This will create the CDP policy group.

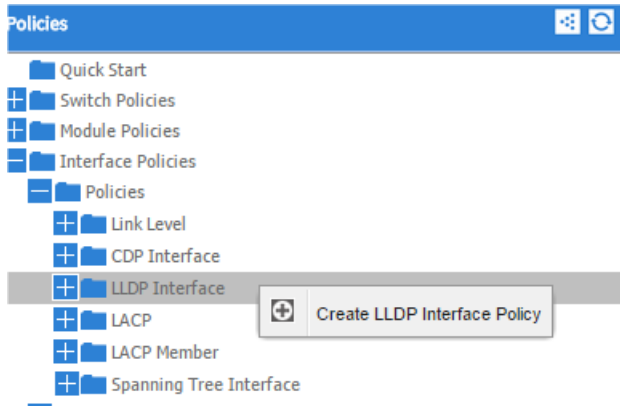
Figure 37 CDP Enabled



Creating LLDP Policy group

1. On the menu bar, choose **FABRIC > ACCESS POLICIES**.
2. In the Navigation pane, expand the Interface Policies and expand the Policies again.
3. Right-click the LLDP Interface and select **Create LLDP Interface Policy**.

Figure 38 Create LLDP Interface Policy



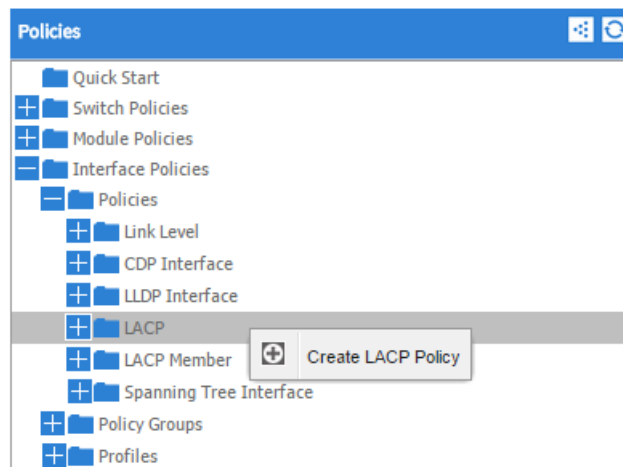
4. In the Create LLDP Interface Policy dialogue box, enter “FI_LLDP” as the policy name, set both the Receive and Transmit State “Disabled” and click **submit**.
5. This will create the LLDP policy group.

Figure 39 LLDP Disabled

Creating LACP Policy

1. On the menu bar, choose **FABRIC > ACCESS POLICIES**.
2. In the Navigation pane, expand the Interface Policies and expand the Policies again.
3. Right-click the LACP and select **Create LACP Policy**.

Figure 40 Create LACP Policy



4. In Create LACP Policy window, enter the name “LACP_Active”. In the mode select the “Active” radio button and click **submit**.

Figure 41 Creating LACP Policy

CREATE LACP POLICY i x

Specify the LACP Policy Identity

Name:

Description:

Mode: Mac Pinning
 Passive
 Off
 Active

Control: Graceful Convergence
 Suspend Individual Port
 Load Defier Member Ports
 Fast Select Hot Standby Ports

Minimum Number of Links: ▲▼

Maximum Number of Links: ▲▼

Creating vPC

1. Expand the Configured Switch Interfaces area to configure the VPCs for the server ports, and perform the following actions:
 - a. In the Switches drop-down list, check the check boxes for the switches that you want to connect to the Fabric Inteconnect. (LEAF_1 & LEAF_2).
 - b. In the Switch Profile Name field, enter a name for the profile (FI_Connected_Leaves).

Figure 42 Configuring vPC Ports

Select Switches To Configure Interfaces: Quick Advanced

Switches: 101-102

Switch Profile Name: FI_Connected_Leaves

Interface Type: Individual PC VPC

Interfaces: 1/1-7
Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Interface Selector Name: VPC_1

VPC Policy Group: select or type to pre-provision

Color: [dropdown]

SAVE CANCEL

Switch (Access Port Fabric Port). Only the access ports can be selected.

- c. Click the + icon to configure the ports.
- d. In the Interface Type area, verify the **VPC** radio button is selected.
- e. In the Interfaces field, enter the ports to which the servers are connected.
- f. In the Interface Selector Name field, enter the name of the port profile (VPC_1).
- g. In the VPC Policy Group field, from the drop-down list, choose **Create VPC Interface Policy Group**.
- h. In Create VPC Interface Policy Group window, enter the name “VPC1”.



Note Create separate VPC interface policy group for each VPC link.

Figure 43 Creating VPC Interface Policy Group

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: VPC_1

Description: optional

Link Level Policy: select or type to pre-provision

CDP Policy: FI_CDP

LLDP Policy: FI_LLDP

STP Interface Policy: select or type to pre-provision

LACP Policy: LACP_Active

Monitoring Policy: select or type to pre-provision

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: select an option

SUBMIT CANCEL

- i. In the CDP Policy field, from the drop-down list, choose “FI_CDP”.
- j. In the LLDP Policy field, from the drop-down list, choose “FI_LLDP”.
- k. In the LACP policy field, from the drop-down list choose “LACP_Active” and click **Submit**.
- l. In the Create VPC Interface Policy Group window click **Submit**.
- m. In the Configure Interface, PC, and VPC dialog box, click **Save** and click **Save** again.
- n. In the Configure Interface, PC, and VPC dialog box, click **Submit**.
- o. Repeat step “C” to “M” to create VPC port for all the Fabric Interconnects connected to the ACI fabric. Once all the FI vPC port is configured, the configured switch interface window should look like fig below.

Configuring vPC Leaf Pairing

1. In the Configure Interface, PC, and VPC dialog box, click on the “+” on VPC DOMAIN ID.

Figure 44 *Creating vPC Domain*

The screenshot shows two configuration panels. The top panel, titled 'CONFIGURED SWITCH INTERFACES', contains a table with the following data:

SWITCH PROFILE	SWITCHES	INTERFACE SELECTOR	INTERFACES	TYPE	POLICY GROUP
JB_C...	102	JB_Connec...	1/48	Individual	Mgmt
FI_C...	101,102	P01	1/1-4	VPC	VPC_1
		P02	1/5-8	VPC	VPC_2
		P03	1/33-36	VPC	VPC_3
		P04	1/37-40	VPC	VPC_4

The bottom panel, titled 'VPC SWITCH PAIRS', has the following fields:

- VPC DOMAIN ID: 145
- SWITCH 1: 101
- SWITCH 2: 102

2. In the VPC Domain ID field, enter “145”.
3. In the “Switch A” drop down box, select node “101”.
4. In the “Switch B” drop down box, select node “102” and click **Save** and **Submit**.

Figure 45 *Creating vPC Peer*

The dialog box contains the following text and fields:

Select two switches to be paired for VPC.
Only switches with interfaces in the same VPC policy group can be paired together.

VPC Domain ID:

Switch 1:

Switch 2:

Interfaces in VPC: switch: 101, interfaces: 1/31
switch: 102, interfaces: 1/31

SAVE CANCEL



Note

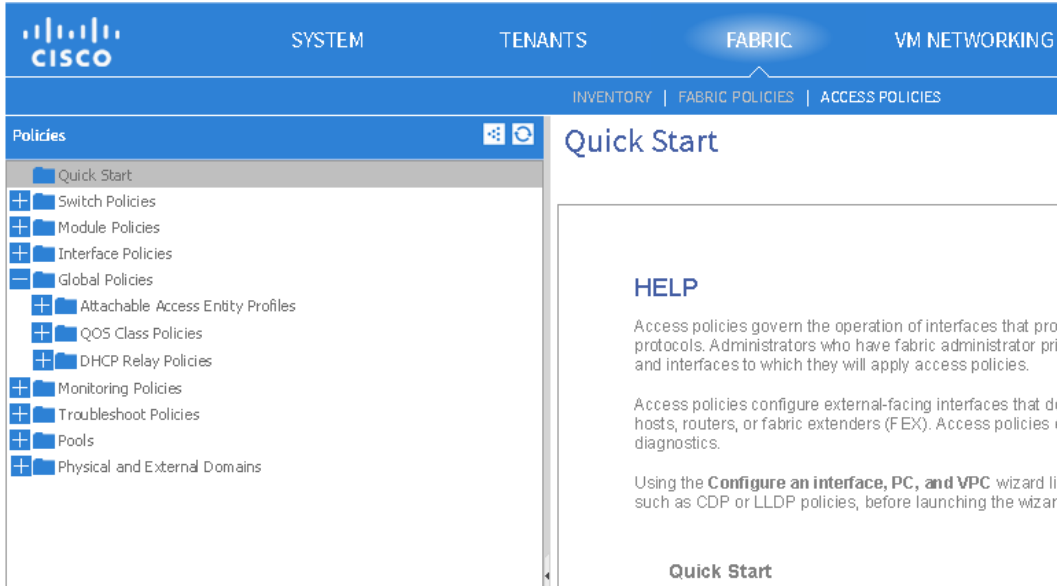
The vPC created here will not come up until the port-channel in Fabric Interconnect uplink ports is created.

Creating Attachable Entity Profile

1. On the menu bar, choose **FABRIC > Access Policies**. In the Work pane, expand Global Policies.

2. Select Attachable Access Entity Profile and right-click on it and select **Create Attachable Access Entity Profile**.

Figure 46 Fabric Window



3. Create Attachable Access Entity Profile window opens, in the name field enter FI_AEP and click + Domains (VMM, Physical Or External) To Be Associated To Interfaces.

Figure 47 Creating Attachable Access Entity Profile for vPC

CREATE ATTACHABLE ACCESS ENTITY PROFILE

STEP 1 > PROFILE

1. PROFILE 2. ASSOCIATION TO INTERFACES

Specify the name, domains and infrastructure encaps

Name: FI_AEP

Description: optional

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation
select or type to pre-provision	

UPDATE CANCEL

< PREVIOUS NEXT > CANCEL

4. In the Domain Profile field, from the drop-down list, choose **Create Physical Domain**.
5. In the Create Physical Domain dialog box, in the Name field, enter the name (UCS_FI).

Figure 48 Creating Physical Domain

CREATE PHYSICAL DOMAIN

Specify the domain name and the VLAN Pool

Name: UCS_FI

VLAN Pool: select an option

Mgmt

Create VLAN Pool

SUBMIT CANCEL

6. In the VLAN Pool field, from the drop-down list, choose **Create VLAN Pool**.
7. In the Create VLAN Pool dialog box, in the Name field, enter the pool name (UCS_FI).

- In the Allocation Mode area, click the **Static Allocation** radio button.

Figure 49 Creating VLAN Pool

The screenshot shows a dialog box titled "CREATE VLAN POOL". It contains the following fields and options:

- Name:** UCS_FI
- Description:** optional
- Allocation Mode:**
 - Dynamic Allocation
 - Static Allocation
- Encap Blocks:** A section with a plus icon and a minus icon. Below it is a table with one row:

VLAN Range

At the bottom right, there are "SUBMIT" and "CANCEL" buttons.

- Expand Encap Blocks. In the Create Ranges dialog box, in the Range fields, add a VLAN range (11-12) and click **OK**.

Figure 50 Assigning VLAN Range

The screenshot shows a dialog box titled "CREATE RANGES". It contains the following fields and options:

- Specify the Encap Block Range**
- Type:** VLAN
- Range:** 11 - 12
 - From: 11
 - To: 12

At the bottom right, there are "OK" and "CANCEL" buttons.

- Repeat step 9 again to create VLAN 160.

VLAN Assignments are as follows:

- Vlan 160 for Management
- Vlan 11 for HDFS
- Vlan 12 for DATA

More detail is provided in the FI configuration section.

Figure 51 VLAN Pool

VLAN Pool - UCS_FI (Static Allocation)

PROPERTIES

Name: **UCS_FI**

Description: optional

Allocation Mode: **Static Allocation**

Encap Blocks: **+** **x**

VLAN RANGE	
[11-12]	
[160]	

Domains: **-** **DN**

uni/phys-UCS_FI	
-----------------	--

SUBMIT **RESET**

11. In the Create VLAN Pool dialog box, click **Submit**.
12. In the Create Physical Domain dialog box, click **Submit**.
13. In the Create Attachable Access Entity Profile dialog box, click **Update**.
14. Create Attachable Access Entity Profile window opens, in the name field enter FI_AEP and click “+” Domains (VMM, Physical Or External) To Be Associated To Interfaces.
15. In the Domain Profile field, from the drop-down list, choose **Create Layer 2 Domain**.
16. In the Create Layer 2 Domain dialog box, in the Name field, enter the name (FI).
17. From the VLAN Pool drop-down list choose “UCS_FI” and click **Submit**.

Figure 52 Creating Layer 2 Domain

VLAN Pool - UCS_FI (Static Allocation)

The screenshot displays the configuration page for a VLAN Pool named 'UCS_FI'. The 'PROPERTIES' section includes the following details:

- Name:** UCS_FI
- Description:** optional
- Allocation Mode:** Static Allocation
- Encap Blocks:**
 - VLAN RANGE
 - [11-12]
 - [160]
- Domains:**
 - DN
 - uni/phys-UCS_FI

At the top right, there are tabs for 'POLICY', 'OPERATIONAL', 'FAULTS', and 'HISTORY', along with an 'ACTIONS' button. At the bottom right, there are 'SUBMIT' and 'RESET' buttons.

18. In Create Attachable Access Entity Profile window click **Next**.
19. In each of the “Interface Policies” that was created for the vPC, select the radio button **All** and click **Finish**.

Figure 53 Associating the Interface to AEP

Creating Tenants, Private Network, and Bridge Domains

Tenants Overview

- A tenant contains policies that enable qualified users domain-based access control. Qualified users can access privileges such as tenant administration and networking administration.
- A user requires read/write privileges for accessing and configuring policies in a domain. A tenant user can have specific privileges into one or more domains.
- In a multi-tenancy environment, a tenant provides group user access privileges so that resources are isolated from one another (such as for endpoint groups and networking). These privileges also enable different users to manage different tenants.

Creating a Tenant, Private Network, and Bridge Domain Using the GUI

Create and specify a network and a bridge domain for the tenant. The defined bridge domain element subnets reference a corresponding Layer 3 context.

1. On the menu bar, choose **TENANTS**, and perform the following actions:
 - a. Click **Add Tenant**.
 - b. The Create Tenant dialog box opens.
 - c. In the Name field, add the tenant name (Big_Data), and click **Next**.
2. Create a security domain so that it allows only users in that security domain to have access.

Click the + sign next to Security Domains to open the Create Security Domain dialog box, and perform the following actions:

- a. In the Name field, specify the security domain name. (Big_Data)
- b. Click **Submit**. In the Create Tenant dialog box, check the check box for the security domain that you created, and click **Next**.

Figure 54 Creating Tenant

CREATE TENANT
i X

STEP 1 > TENANT
1. TENANT
2. NETWORK

Tenant Identity

Specify tenant details

Name:

Description:

Tags: enter tags separated by comma

Monitoring Policy:

Security Domains: +

Select	Name	Description
<input checked="" type="checkbox"/>	Big_Data	
<input type="checkbox"/>	all	
<input type="checkbox"/>	mgmt	

< PREVIOUS
NEXT >
CANCEL

3. In the Network window, perform the following actions:
 - a. Click the + sign to add the network.
 - b. In the Create New Network area, specify the private tenant network name (PVN_1) and click **Next**.

Figure 55 Creating Tenant Network

CREATE TENANT

STEP 2 > NETWORK

1. TENANT 2. NETWORK

TENANT BIG_DATA NETWORK > BRIDGE DOMAIN

CREATE NEW NETWORK

Specify Tenant Network

Name:

Policy Enforcement: Enforced Unenforced

Description:

BGP Timers:

OSPF Timers:

Monitoring Policy:

DNS Labels:

enter names separated by comma

Create A Bridge Domain:

- Specify the bridge domain in the Name field (BD_1), click **OK**. Click **Next**, and perform the following actions:

Figure 56 Creating Bridge Domain

TENANT BIG_DATA
CREATE NEW NETWORK

NETWORK > BRIDGE DOMAIN

Specify Bridge Domain for the Network

Name:

Description:

Forwarding:

IGMP Snoop Policy:

Config BD MAC Address:

Subnets:

Gateway Address	Scope	Subnet Control

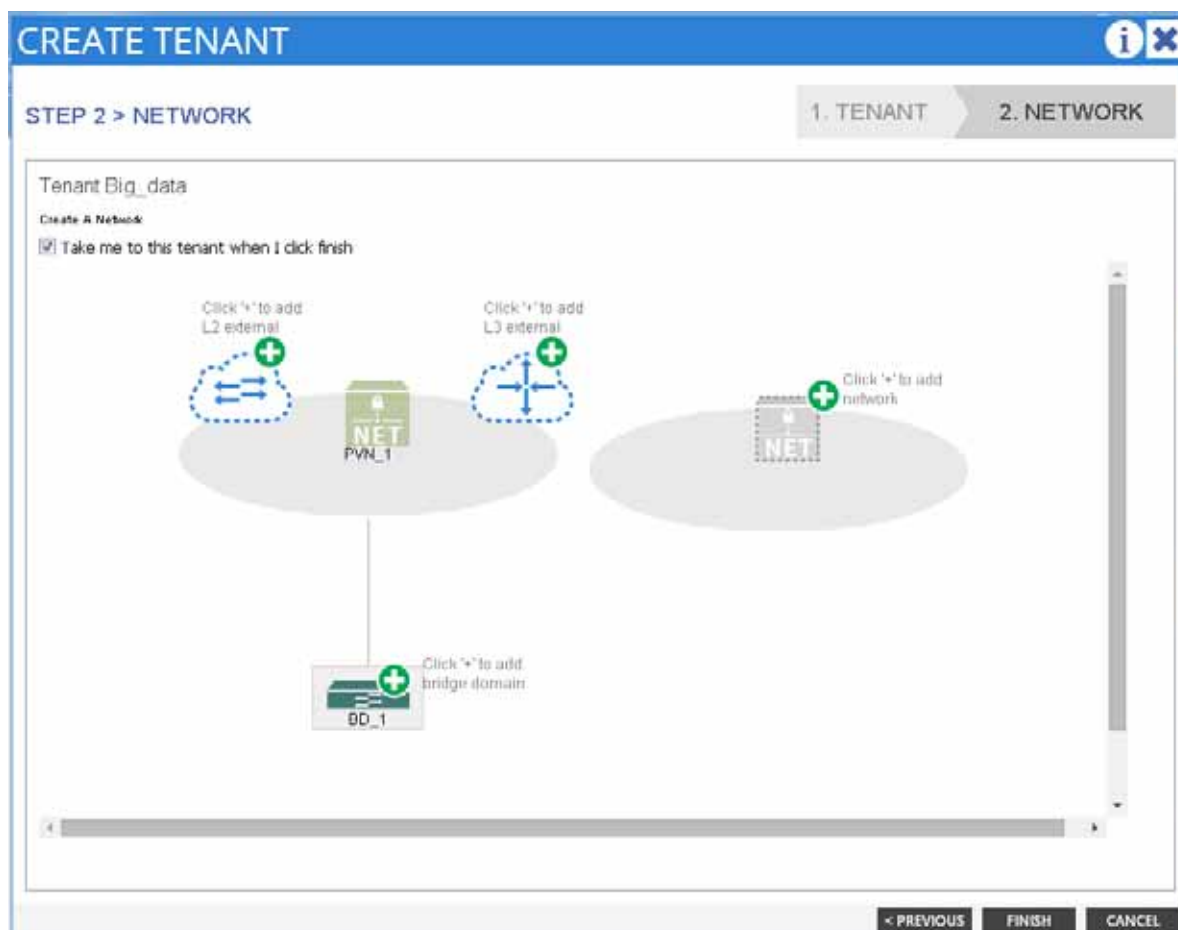
DHCP Labels:

Name	Scope	DHCP Option Policy

< PREVIOUS OK CANCEL

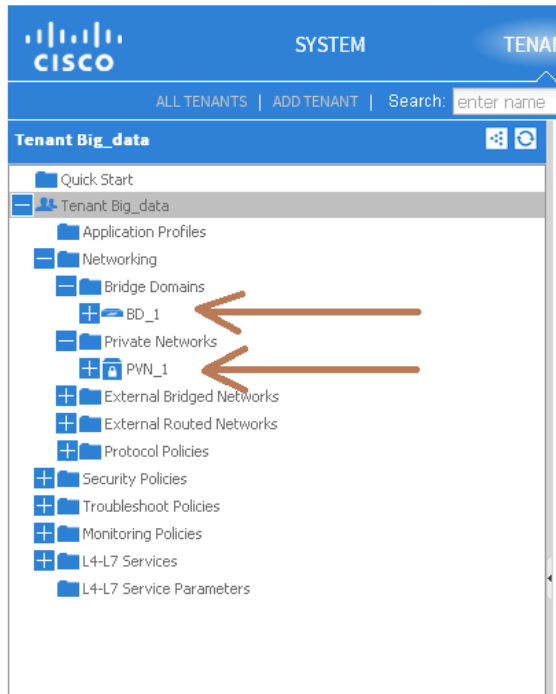
- a. Confirm that the private network (PVN_1) is created and is associated with the bridge domain (BD_1).
- b. In the Application Profile window, click **Finish**.

Figure 57 Confirming the Association



- To validate that the tenant has a private network and bridge domain, in the submenu bar under the Tenants tab, click the new tenant name that was created earlier. In the Navigation pane, expand the tenant name. Under Bridge Domains, the new bridge domain is displayed. Under Private Networks, the new network is displayed.

Figure 58 Validating the Bridge Domain and Private Network



6. Select the bridge domain created earlier (BD_1) and check the L2 Unknown Unicast to Flood, and check the ARP Flooding checkbox and click **Submit**.

Figure 59 ARP Flooding

The screenshot displays the configuration for Bridge Domain - BD_1. The left sidebar shows the navigation tree with 'BD_1' selected under 'Bridge Domains'. The main panel displays the 'PROPERTIES' for 'BD_1'. Key settings include: Name: BD_1, Description: optional, Unknown Unicast Traffic Class ID: 49153, Segment: 15204288, Multicast Address: 225.0.176.16, Network: PVN_1, Custom MAC Address: 00:22:BD:F8:19:FF, L2 Unknown Unicast: Flood (selected), Unknown Multicast Flooding: Flood (selected), ARP Flooding: checked, Unicast Routing: checked, IGMP Snoop Policy: select or type to pre-provi, End Point Retention Policy: select or type to pre-provi, and Associated L3 Outs: L3 OUT. Two red arrows point to the 'Flood' radio button for L2 Unknown Unicast and the 'checked' checkbox for ARP Flooding.

7. Expand the Bridge Domain and BD_1, right-click on the Subnets and select **Create Subnet**.
8. In the Create Subnet dialogue box, enter the gateway IP 10.0.145.1/24 and click **Submit**.

This IP address (10.0.145.1/24) is assigned to the bridge domain that typically is used as the switch virtual interface (SVI) in a traditional switch configuration.

Figure 60 Create Subnet

9. Repeat step 7 again to create two more subnets for other two VLANs.

Figure 61 Created Subnets

Subnets

GATEWAY ADDRESS	SCOPE	SUBNET CONTROL
10.0.145.1/24	Private Subnet	
10.0.146.1/24	Private Subnet	
10.0.147.1/24	Private Subnet	

Creating an Application Profile Using the GUI

1. On the menu bar, choose **TENANTS**. In the Navigation pane, expand the tenant, right-click Application Profiles, and click **Create Application Profile**.
2. In the Create Application Profile dialog box, in the Name field, add the application profile name (CLOUDERA).

Figure 62 Creating Application Profile

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name:

Description:

Tags:
enter tags separated by comma

Monitoring Policy:

EPGs

Name	Description

Contracts

Create EPGs on the left table to add contracts

Creating EPGs Using the GUI

1. Expand **Tenant BIG_DATA > Application Profiles > CLOUDERA**, right-click on the Application EPGs and select **Create Application EPG**. In the Create Application EPG dialog box, perform the following actions:
 - a. In the **Name** field, add the EPG name (Mgmt).
 - b. In the **Bridge Domain** field, choose the bridge domain from the drop-down list (BD_1).
 - c. Expand **Associated Domain Profiles** and from the drop-down list, choose the Domain Profile name (Mgmt).
 - d. From the Deployment Immediacy and Resolution Immediacy drop-down list select Immediate.
 - e. Click **Update**, and click **Finish**.

Figure 63 Creating EPG

CREATE APPLICATION EPG

STEP 1 > IDENTITY 1. IDENTITY

Specify the EPG Identity

Name: Mgmt

Description: optional

Tags:

QoS class: Unspecified

Custom QoS:

Bridge Domain: BD_1

Monitoring Policy:

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
FI_External	Immediate	Immediate

Statically Link with Leaves/Paths:

< PREVIOUS FINISH CANCEL

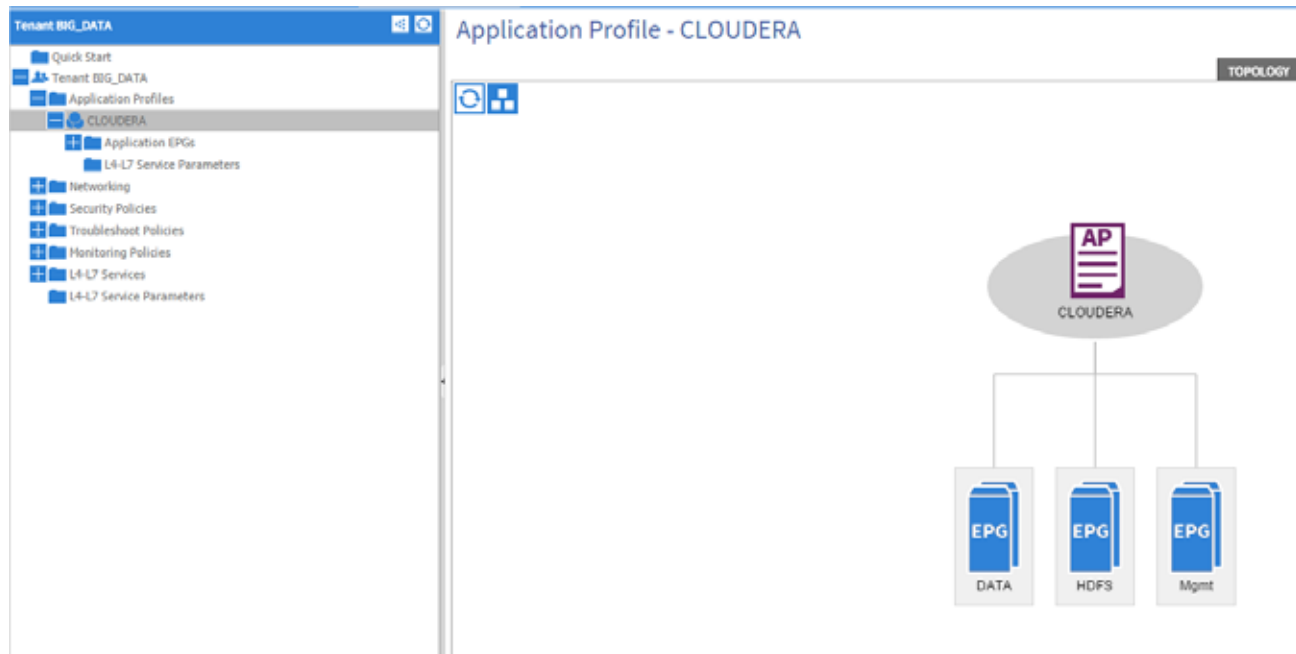
- Repeat step 1 to create two more EPGs named DATA and HDFS.



Note On Cloudera Security: When deploying Cloudera with Security only one VLAN on one vNIC is supported at the UCS Manager and UCS Server level. Under this scenario, there is no need to create EPGs DATA and HDFS at the ACI level. If already created, there is no problem. This could be left as is as from the downstream, only data from mgmt VLAN will be forwarded upstream.

- Once all three EPGs are created, these EPGs are associated with the Application Profile CLOUDERA.

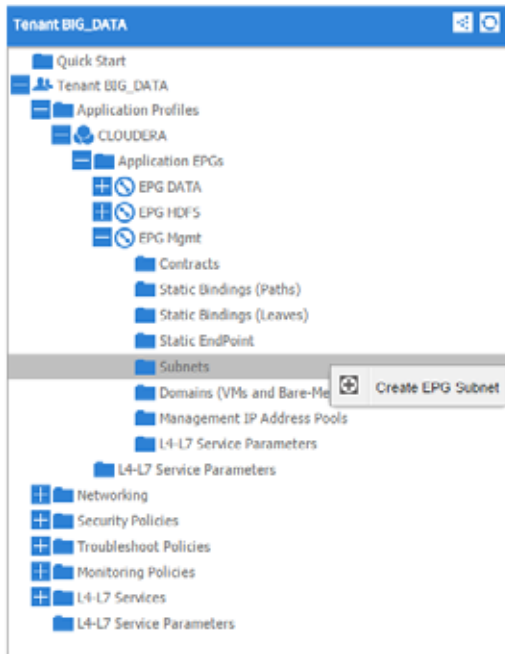
Figure 64 EPG Associated with Application Profile CLOUDERA



Configuring EPGs

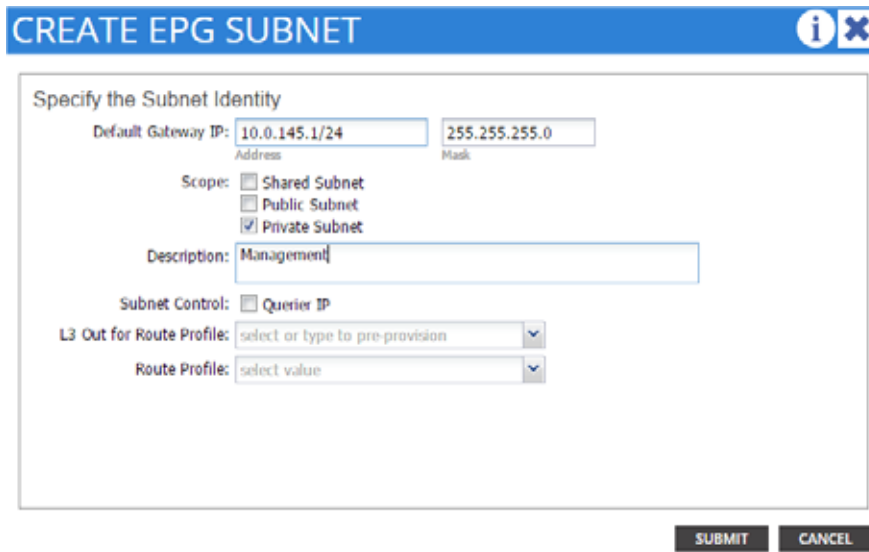
1. Expand **Tenant BIG_DATA > Application Profiles > CLOUDERA > Application EPGs > EPG Mgmt.**
2. Right-click the Subnets and select **Create EPG Subnet.**

Figure 65 Creating EPG Subnet



3. In the Create EPG Subnet dialogue box, enter the Default Gateway IP as 10.0.145.1/24 and click **Submit**.

Figure 66 Defining Subnet Address



4. Similarly configure subnet for other EPGs with appropriate subnet address. For more detail, navigate to the IP Address Assignment section.

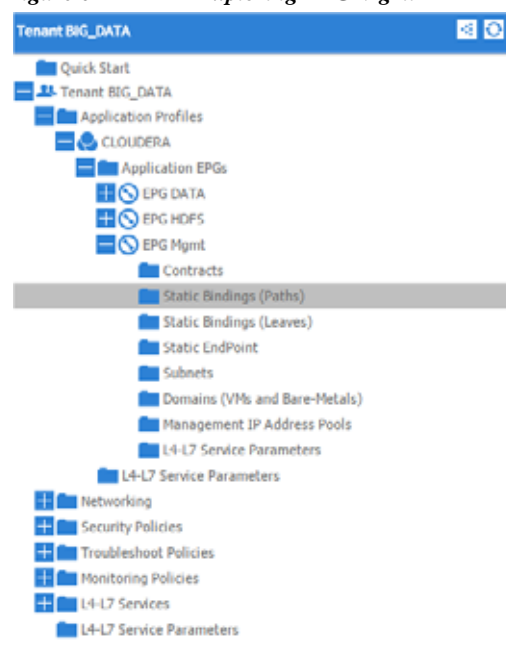
Creating the Static Binding for the Leaves and Paths

The static binding for the leaves are required to associate the physical interfaces with the EPGs.

No traffic flows unless an EPG is deployed on the port. Without VLAN pool deployment using an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned. A particular VLAN is provisioned or enabled on the leaf port based on EPG events by statically binding on a leaf port.

1. On the menu bar, choose TENANTS and the tenant name on which you want to operate. In the Navigation pane, expand the **Tenants > Application Profiles > CLOUDERA > Application EPGs > EPG Mgmt** and select **Static Bindings (Paths)**.

Figure 67 Exploring EPG Mgmt



2. Right-click the Static Bindings (paths) and select **Deploy Static EPG on PC, VPC, or Interface**.
3. In the Path Type: select the **Virtual Port Channel** radio button.
4. From the Path: drop down list select the appropriate nodes and port where the FI's are connected. On Encap field use vlan-160, on Deployment Immediacy select the **Immediate** radio button and on Mode select the Untagged and click **Submit**.

Figure 68 Deploying Static EPG on vPC

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: Port
 Direct Port Channel
 Virtual Port Channel

Path: topology/pod-1/paths-101/pathep-[VPC_1]

Encap: vlan-160
For example, vlan-1

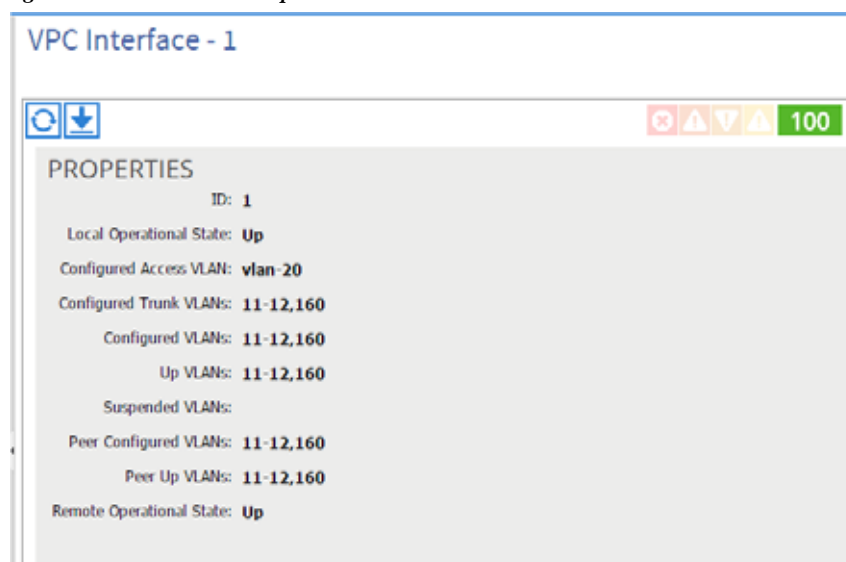
Deployment Immediacy: Immediate
 On Demand

Mode: Tagged
 Untagged
 802.1P Tag

SUBMIT CANCEL

5. Repeat step 2, 3 & 4 for all the vPC ports created.
6. Similarly, statically bind the ports in other EPGs created using the appropriate VLAN numbers (12 for HDFS and 11 for DATA).
7. Once the Static binding for all the EPG is configured properly, verify that the VPC ports created earlier are trunking the appropriate VLANS. This can be verified by the following steps:
 - a. On the menu bar, choose **FABRIC > Access Policies**.
 - a. Expand **Pod 1 > LEAF_1 (Node-101) > Interfaces > VPC Interfaces > 1** (This number might be different in different setups). Select any of the Interfaces to view the properties.

Figure 69 VPC Properties

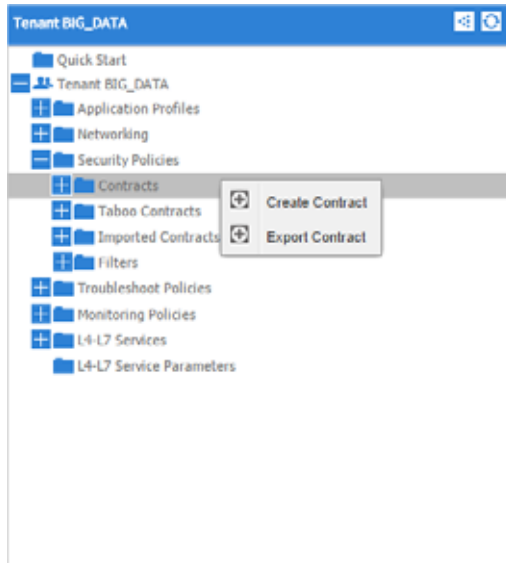


Creating Contracts

Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

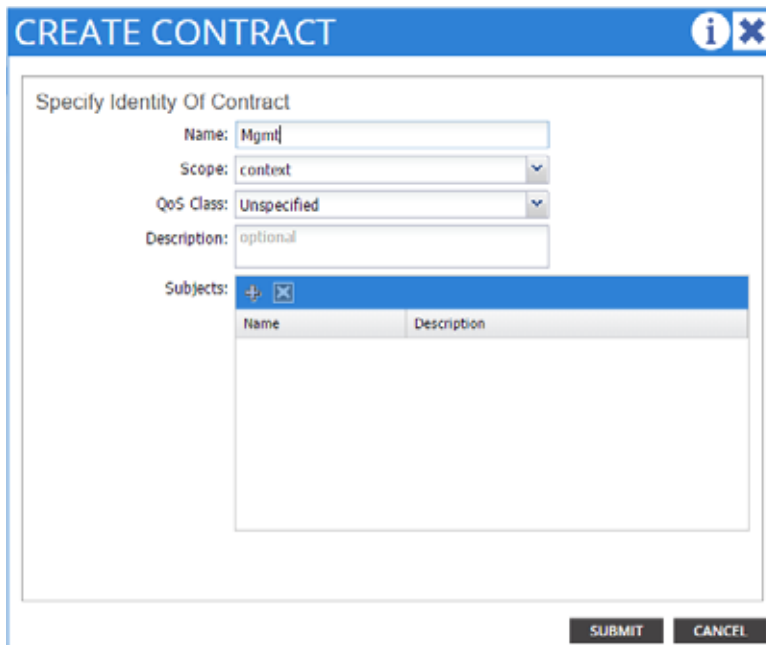
1. On the menu bar, choose **TENANTS** and the tenant name on which you want to operate. In the Navigation pane, expand the **Tenant > Security Policies**.
 - a. Right-click **Contracts > Create Contract**.

Figure 70 Create Contract



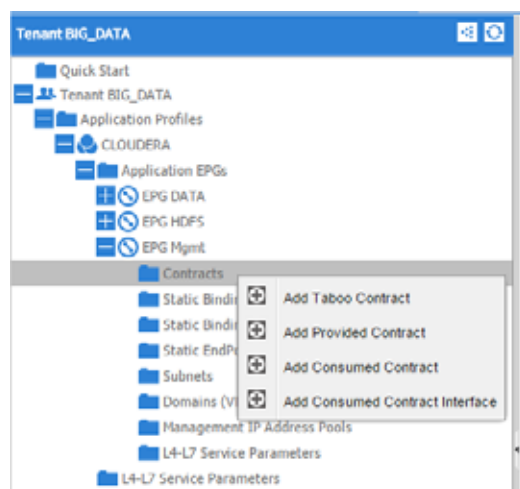
- b. In the Create Contract dialog box, In the Name field, enter the contract name (Mgmt) and click **Submit**.

Figure 71 Enter Contract Details



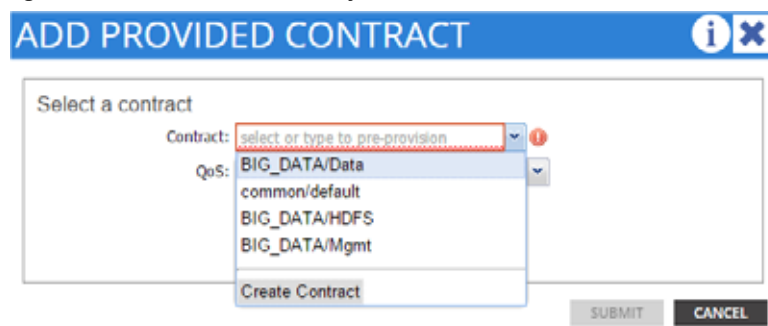
2. Create two more contracts for Data and for HDFS following the same steps in this procedure.
3. On the menu bar, choose TENANTS and the tenant name on which you want to operate. In the Navigation pane, expand the **Tenant > Application Profiles > CLOUDERA > Application EPGs > EPG Mgmt**.
4. Right-click the contract and select **Add Provided Contract**.

Figure 72 *Add Provided Contract*



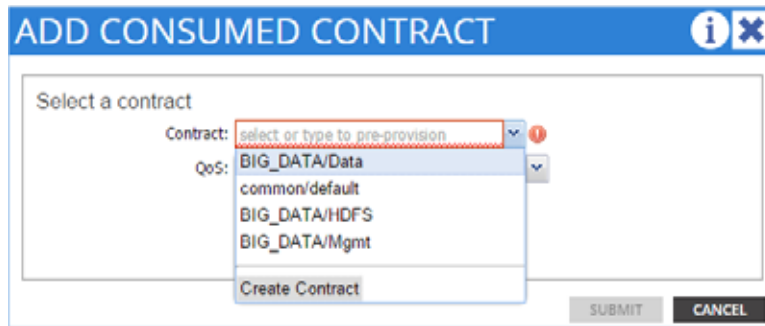
5. In the add provided contract dialogue box, from the contract drop-down list choose BIG_DATA/Mgmt and click **Submit**.

Figure 73 *Select Contract for Provided Contract*



6. Right-click the contract and select **Add Consumed Contract**.
7. In the add consumed contract dialogue box, from the contract drop-down list choose BIG_DATA/DATA and click Submit.
8. Right-click the contract and select **Add Consumed Contract**.
9. In the add consumed contract dialogue box, from the contract drop-down list choose BIG_DATA/HDFS and click **Submit**.

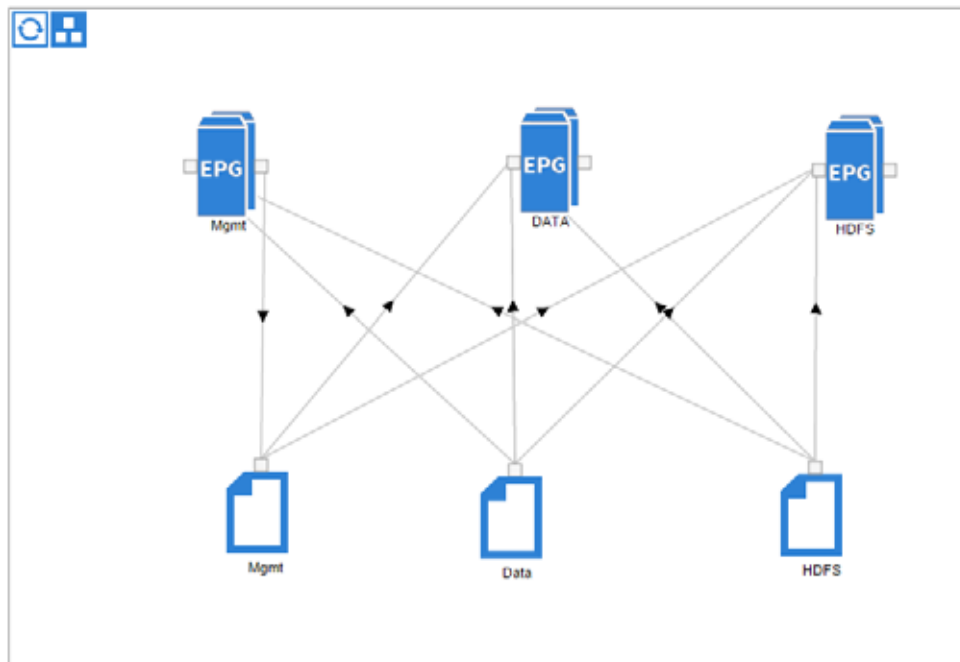
Figure 74 Select Contract for Consumed Contract



10. For EPG DATA, add provided contract BIG_DATA/Data and consumed contract BIG_DATA/Mgmt and BIG_DATA/HDFS.
11. For EPG HDFS, add provided contract BIG_DATA/HDFS and consumed contract BIG_DATA/Mgmt and BIG_DATA/Data.
12. Once all the contract is configured, in the Navigation pane, expand the tenant **BIG_DATA** > **Application Profiles** > **CLOUDERA** and select **Application EPGs**, the window should appear as follows.

Figure 75 Application EPGs After all the Contracts are Configured

Application EPGs



This will complete the Network configuration with three EPGs for each VLANs, a Private Network and a Bridge Domain.

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)

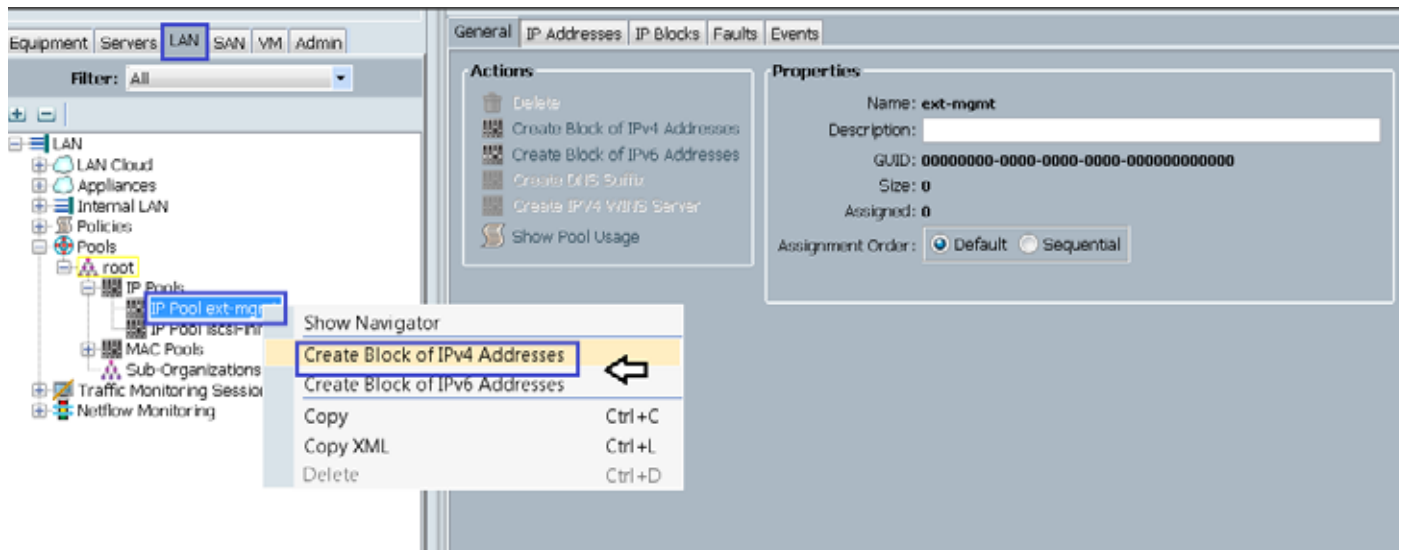
This document assumes the use of UCS 2.2(3d). Refer to [Upgrading between Cisco UCS 2.0 Releases](#) to upgrade the Cisco UCS Manager software and UCS 6296 Fabric Interconnect software to version 2.2(3d). Also, make sure the UCS C-Series version 2.2(3d) software bundles is installed on the Fabric Interconnects.

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.

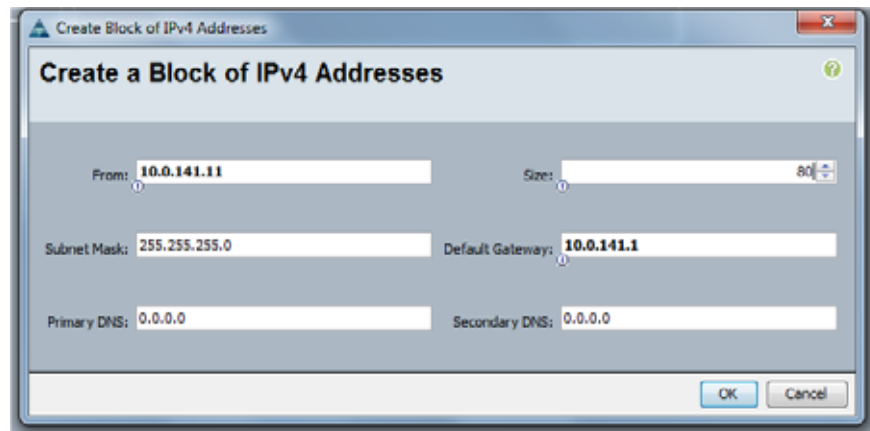
1. Select the **LAN** tab at the top of the left window.
2. Select **Pools > IP Pools > IP Pool ext-mgmt**.
3. Right-click **IP Pool ext-mgmt**
4. Select **Create Block of IPv4 Addresses**.

Figure 76 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 77 Adding Block of IPv4 Addresses for KVM Access Part 2



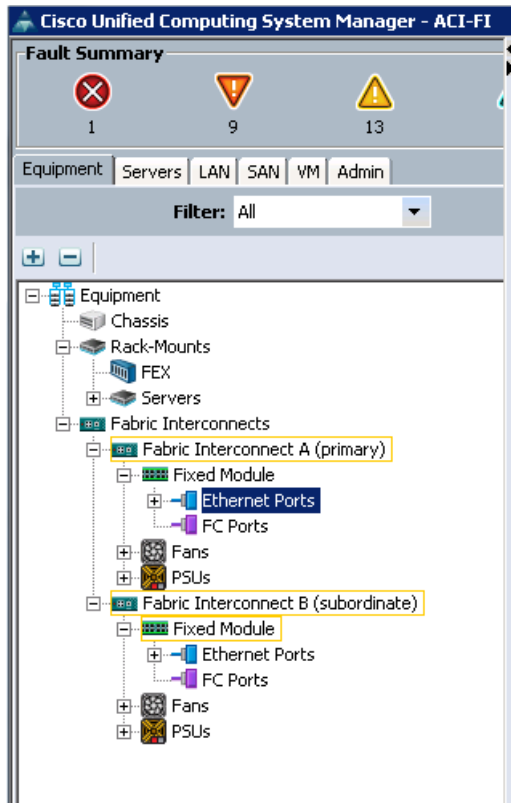
6. Click **OK** to create the IP block.
7. Click **OK** in the message box.

Enabling Uplink Port

These steps provide details for enabling uplinks ports.

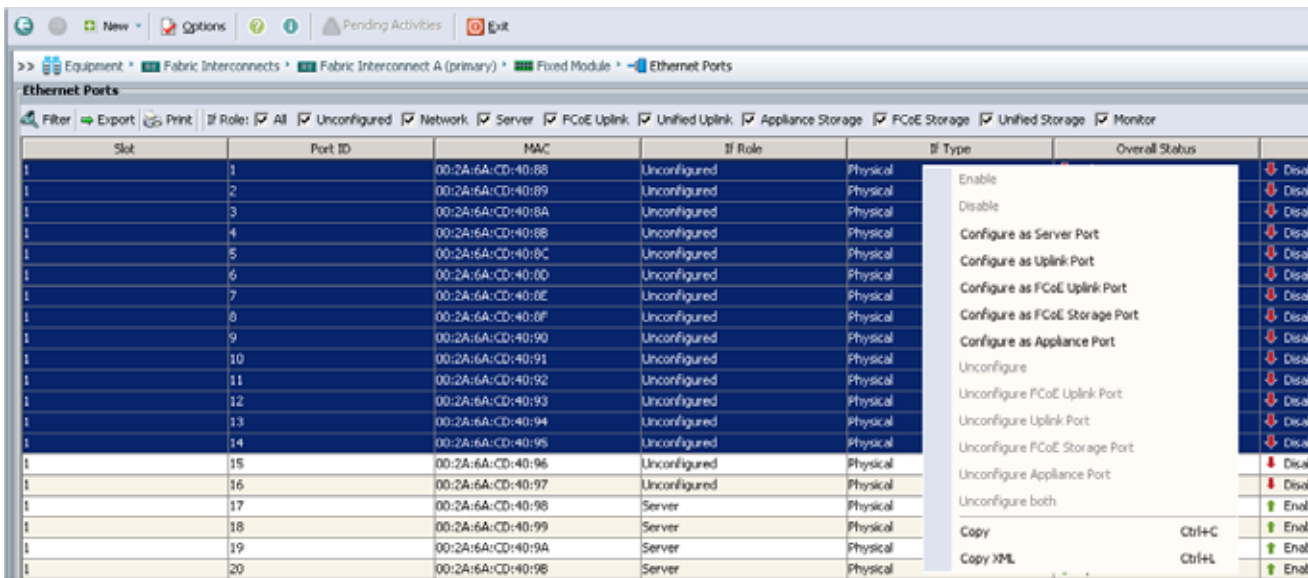
1. Select the **Equipment** tab on the top left of the window.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module > Ethernet Ports**.

Figure 78 Selecting Ethernet Ports



3. On the Right window select all the ports that are connected to the Nexus 9396 leaf switch (14 per FI), right-click them, and select **Configure as uplink Port**.

Figure 79 Enabling Uplink Ports



4. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.

5. Expand the UnConfigured Ethernet Ports section.
6. Select all the ports that are connected to the Nexus 9396 leaf switch (14 per FI), right-click them, and select **Configure as uplink Port**.



Note

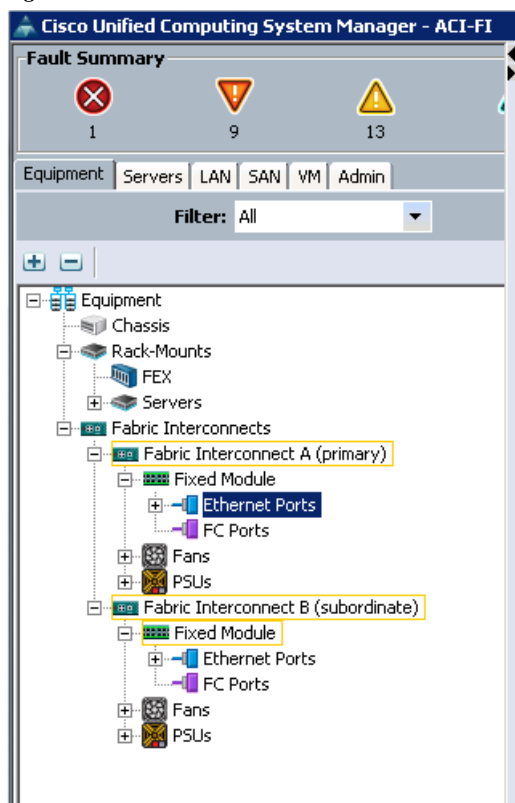
The ports that are configured as uplink port should appear as Network under IF role.

Enabling Server Ports

These steps provide details for enabling server ports.

7. Select the **Equipment** tab on the top left of the window.
8. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module > Ethernet Ports**.

Figure 80 Select Ethernet Ports



9. On the Right window select all the ports that are connected to the UCS C240 server (1 per Server), right-click them, and select **Configure as Server Port**.

Figure 81 Showing Servers Ports

1	14	00:2A:6A:CD:40:95	Unconfigured	Physical
1	15	00:2A:6A:CD:40:96	Unconfigured	Physical
1	16	00:2A:6A:CD:40:97	Unconfigured	Physical
1	17	00:2A:6A:CD:40:98	Unconfigured	Physical
1	18	00:2A:6A:CD:40:99	Unconfigured	Physical
1	19	00:2A:6A:CD:40:9A	Unconfigured	Physical
1	20	00:2A:6A:CD:40:9B	Unconfigured	Physical
1	21	00:2A:6A:CD:40:9C	Unconfigured	Physical
1	22	00:2A:6A:CD:40:9D	Unconfigured	Physical
1	23	00:2A:6A:CD:40:9E	Unconfigured	Physical
1	24	00:2A:6A:CD:40:9F	Unconfigured	Physical
1	25	00:2A:6A:CD:40:A0	Unconfigured	Physical
1	26	00:2A:6A:CD:40:A1	Unconfigured	Physical
1	27	00:2A:6A:CD:40:A2	Unconfigured	Physical
1	28	00:2A:6A:CD:40:A3	Unconfigured	Physical
1	29	00:2A:6A:CD:40:A4	Unconfigured	Physical
1	30	00:2A:6A:CD:40:A5	Unconfigured	Physical
1	31	00:2A:6A:CD:40:A6	Unconfigured	Physical
1	32	00:2A:6A:CD:40:A7	Unconfigured	Physical

Enable

Disable

Configure as Server Port

Configure as Uplink Port

Configure as FCoE Uplink Port

Configure as FCoE Storage Port

Configure as Appliance Port

Unconfigure

Unconfigure FCoE Uplink Port

Unconfigure Uplink Port

Unconfigure FCoE Storage Port

Unconfigure Appliance Port

Unconfigure both

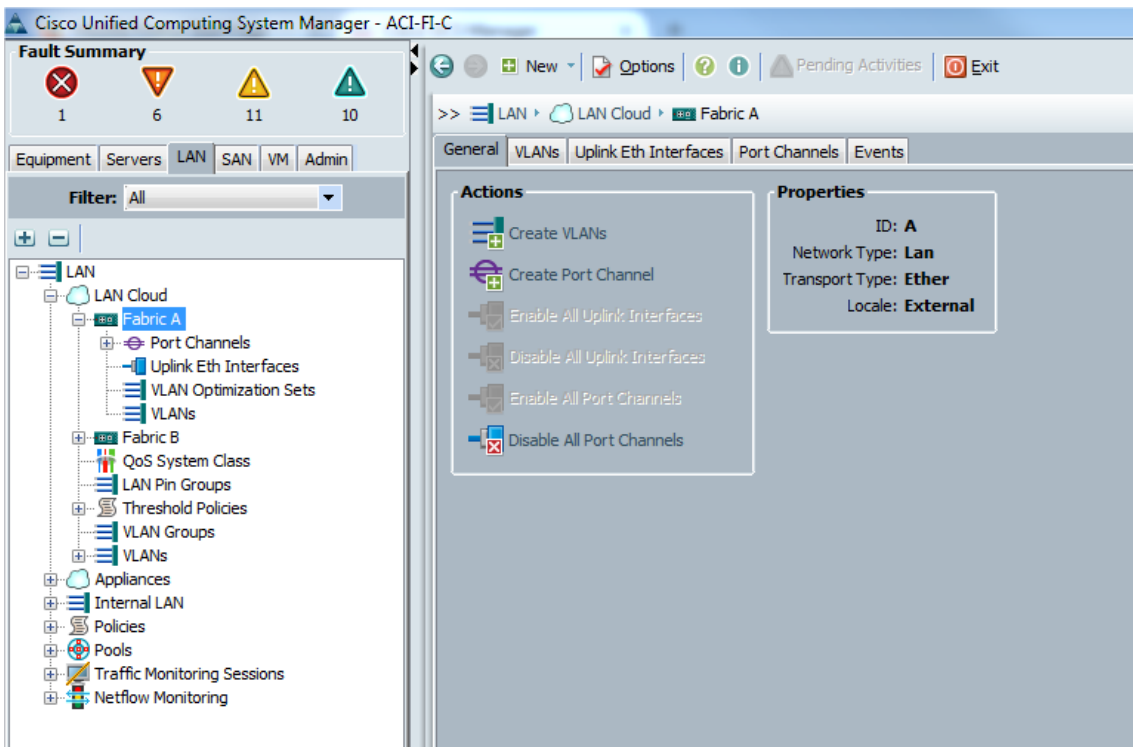
Copy XML Ctrl+C

Copy XML Ctrl+H

Configuring Port-Channels

1. Click the LAN tab on top left window.
2. Expand the LAN Cloud > Fabric A.
3. On the right window select Create Port Channel.

Figure 82 Creating Port Channel



4. On Set Port Channel Name window, perform the following actions:

- a. In the ID field, specify the ID “01” as the first port channel
- b. In Name field, type P01 as Port-channel01 and click **Next**.

Figure 83 *Setting Port-Channel ID and Name*

The screenshot shows the 'Create Port Channel' wizard in the Unified Computing System Manager. The window title is 'Create Port Channel' and the main header is 'Unified Computing System Manager'. The current step is 'Set Port Channel Name', indicated by a green checkmark and a question mark icon. On the left, a progress bar shows two steps: '1. ✓ Set Port Channel Name' and '2. □ Add Ports'. The main area contains two input fields: 'ID: 01' and 'Name: P01'. At the bottom, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

5. In Add Ports window select all the ports that is connected to the Nexus 9396 Leaf Switch and click >>. This will add all the ports in the port channel created earlier.

Figure 84 Adding Ports to the Port Channel

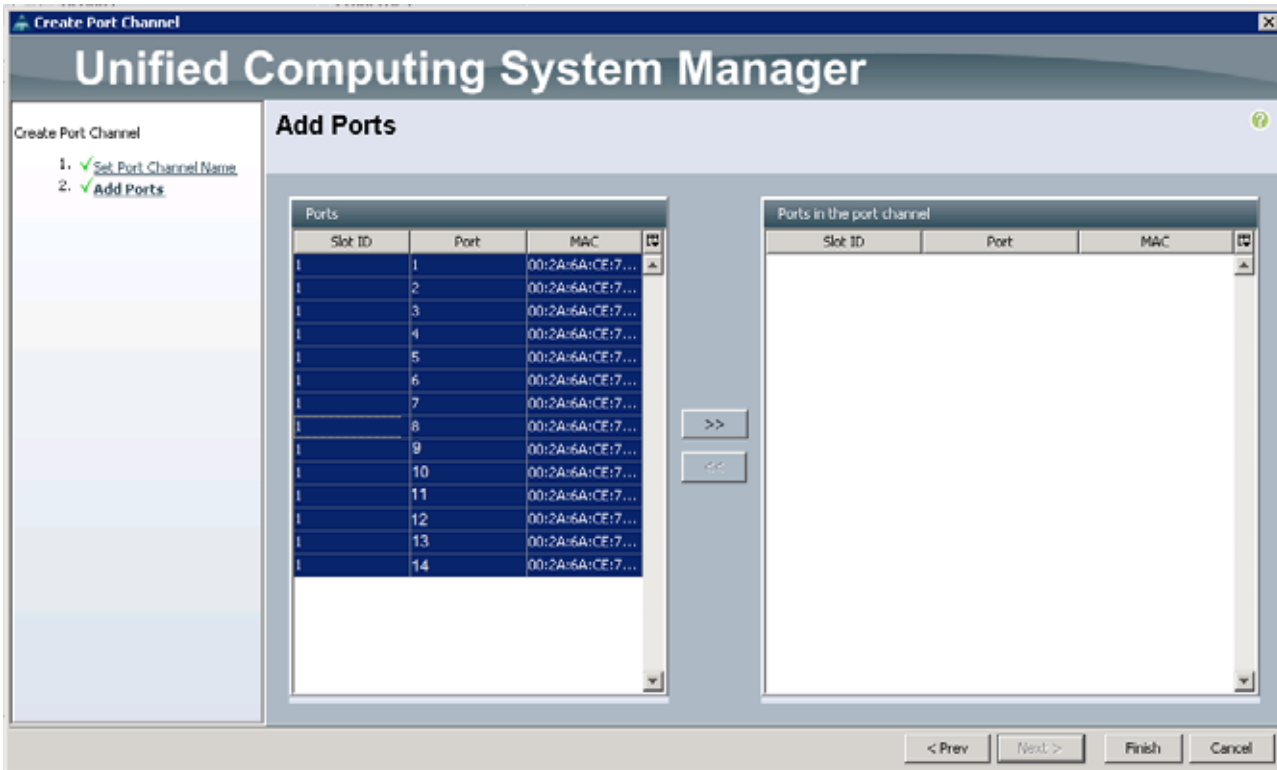
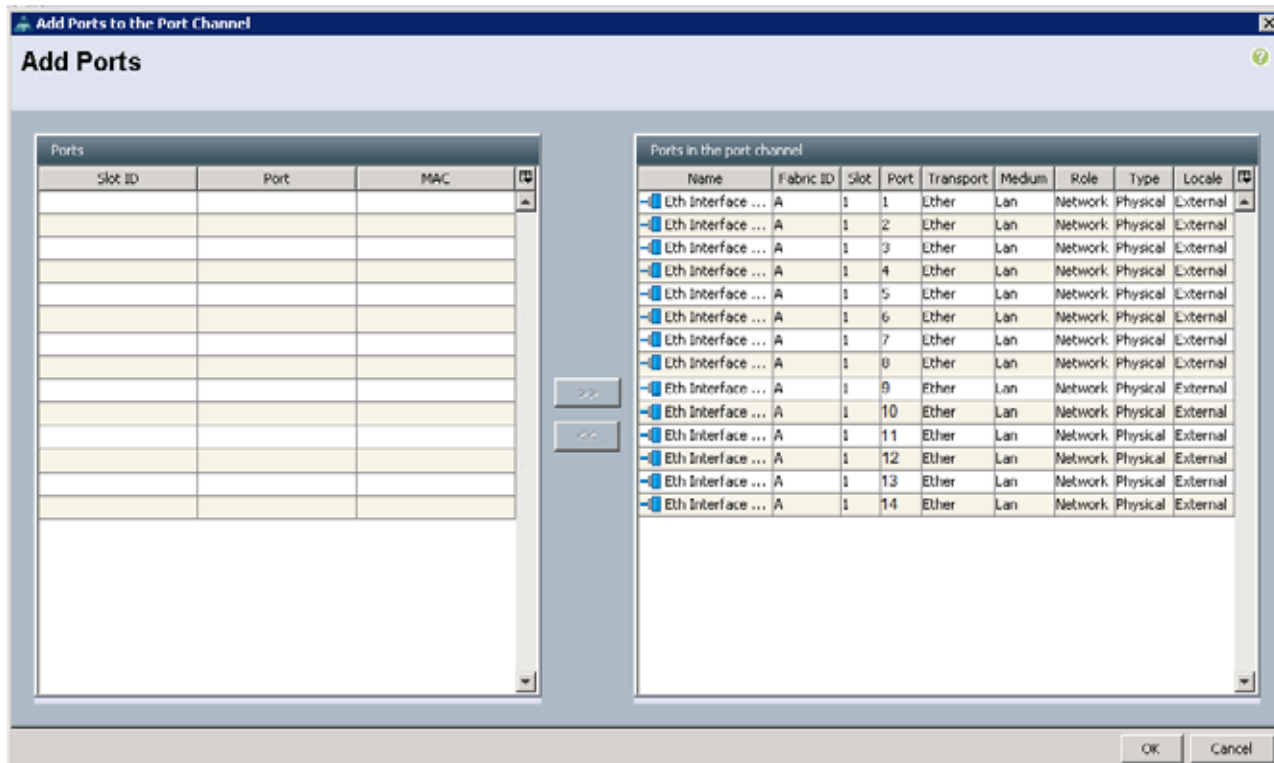
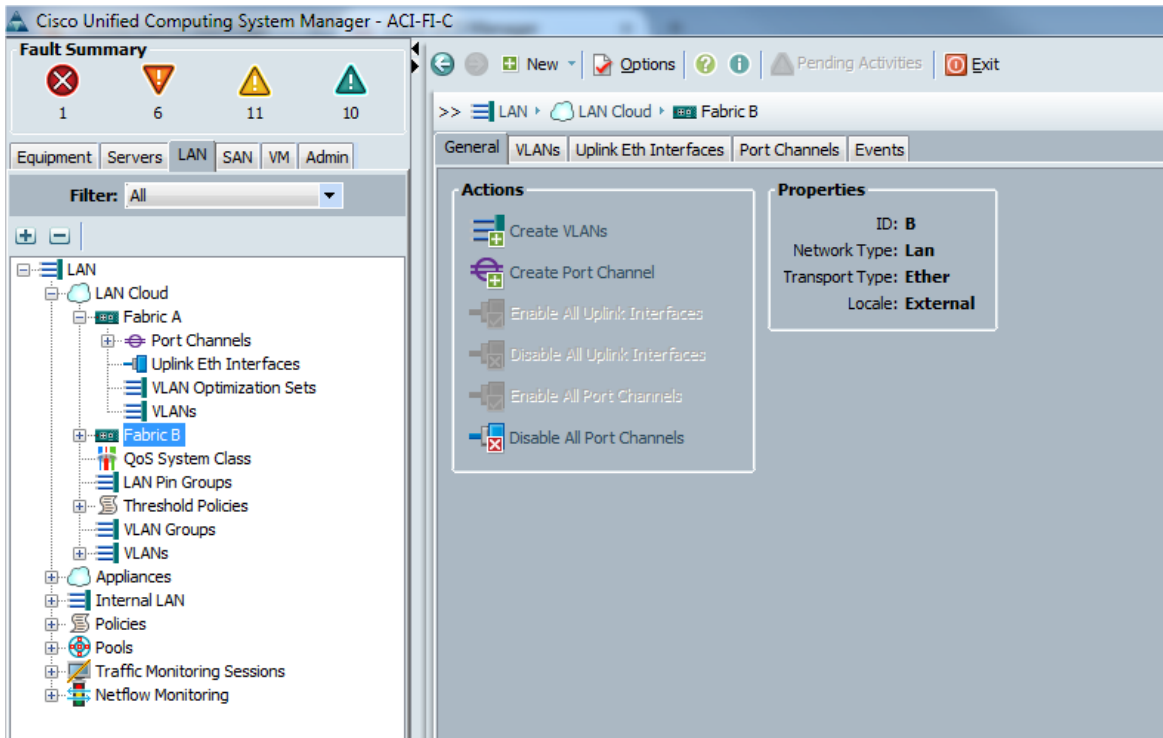


Figure 85 Added Ports to the Port Channel



6. Similarly for Fabric Interconnect B, click the **LAN** tab on top left window.
7. Expand the **LAN Cloud > Fabric B**.
8. In the right pane of the window select **Create Port Channel**.

Figure 86 Creating Port Channel



9. On Set Port Channel Name window, perform the following actions:
 - a. In the ID field, specify the ID “02” as the second port channel
 - b. In Name field, type P02 as Port-channel01 and click **Next**.

Figure 87 Setting Port-Channel ID and Name

The screenshot shows the 'Unified Computing System Manager' interface for creating a port channel. The window title is 'Create Port Channel'. The main heading is 'Set Port Channel Name'. On the left, a sidebar shows the progress: '1. ✓ Set Port Channel Name' and '2. Add Ports'. The main area has two input fields: 'ID: 02' and 'Name: P02'. At the bottom, there are navigation buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

10. In Add Ports window select all the ports that is connected to the Nexus 9396 Leaf Switch and click >>. This will add all the ports in the port channel created earlier.

Figure 88 Adding Ports to the Port Channel


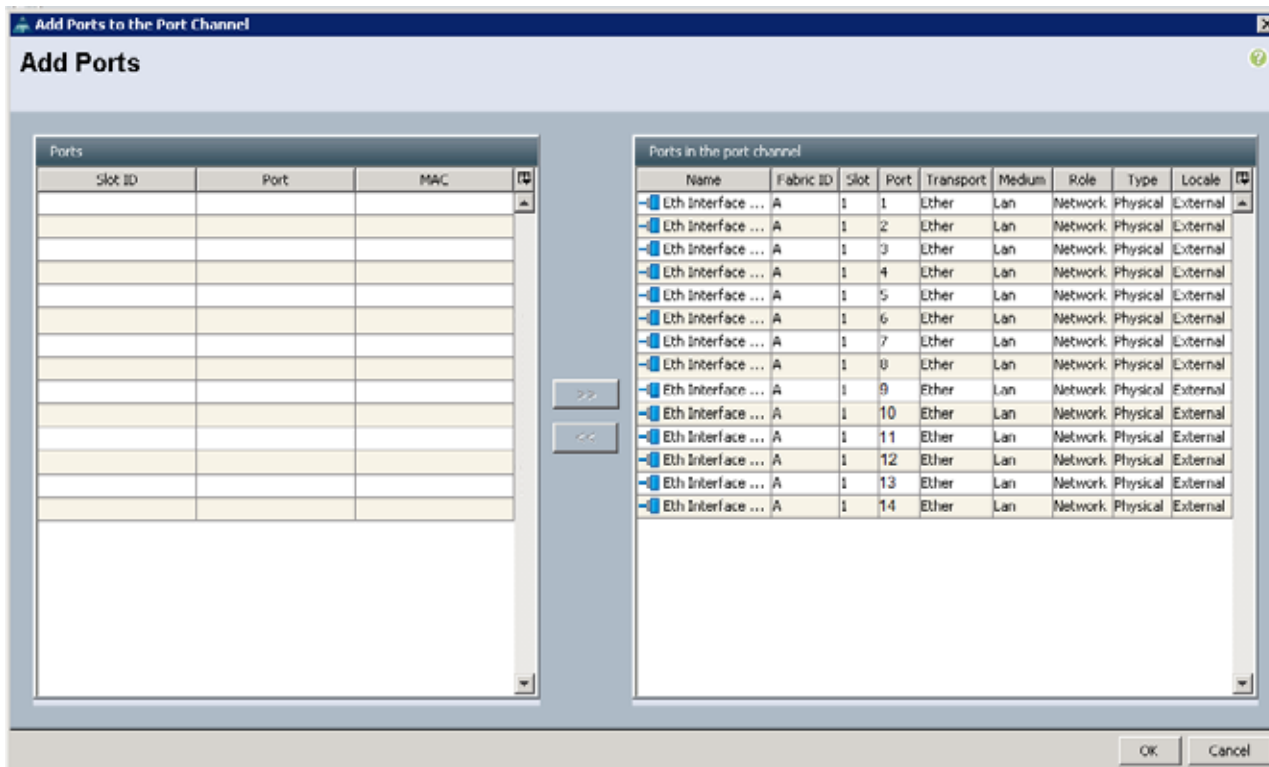


Figure 89 Added Ports to the Port Channel



11. The configured port channels and vPC can be verified by logging in to the APIC.

Figure 90 Verify Configured Port Channels and vPC

Aggregated Interfaces

INTERFACE	PROTOCOL	NATIVE VLAN	SPEED	LAYER	MODE	OPER STATE	OPER STATE REASON	ACCESS VLAN	CONFIG ACCESS VLAN	CONFIG NATIVE VLAN
po3	lACP-active	vlan-29	10 Gbps	switched	trunk	up	none	vlan-31	vlan-31	vlan-29
po5	lACP-active	vlan-29	10 Gbps	switched	trunk	up	none	vlan-31	vlan-31	vlan-29
po6	lACP-active	vlan-29	10 Gbps	switched	trunk	up	none	vlan-31	vlan-31	vlan-29
po8	lACP-active	vlan-29	10 Gbps	switched	trunk	up	none	vlan-31	vlan-31	vlan-29

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 91 Creating MAC Pool Window

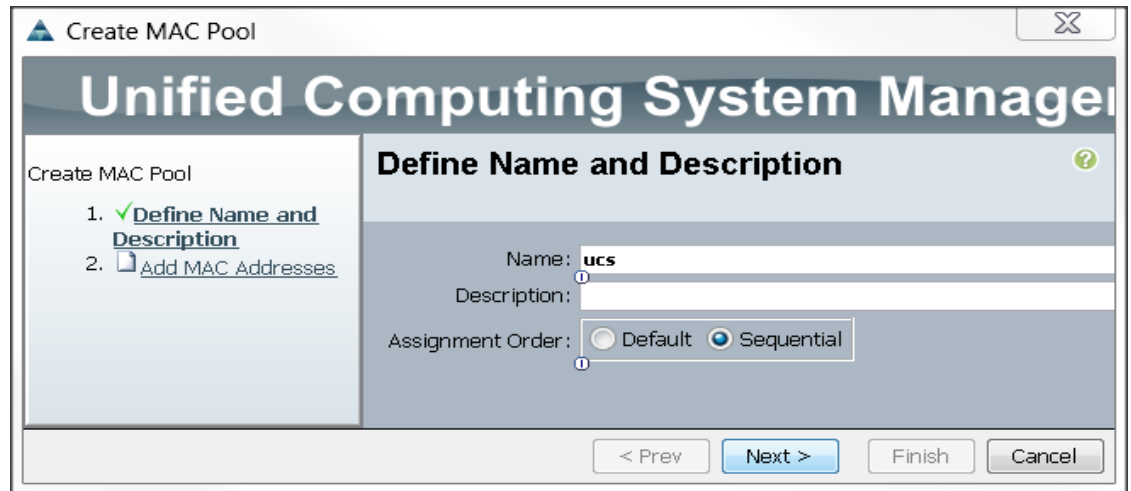
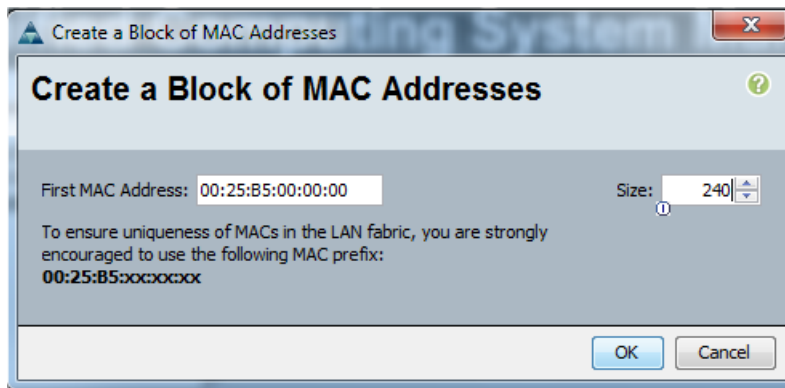
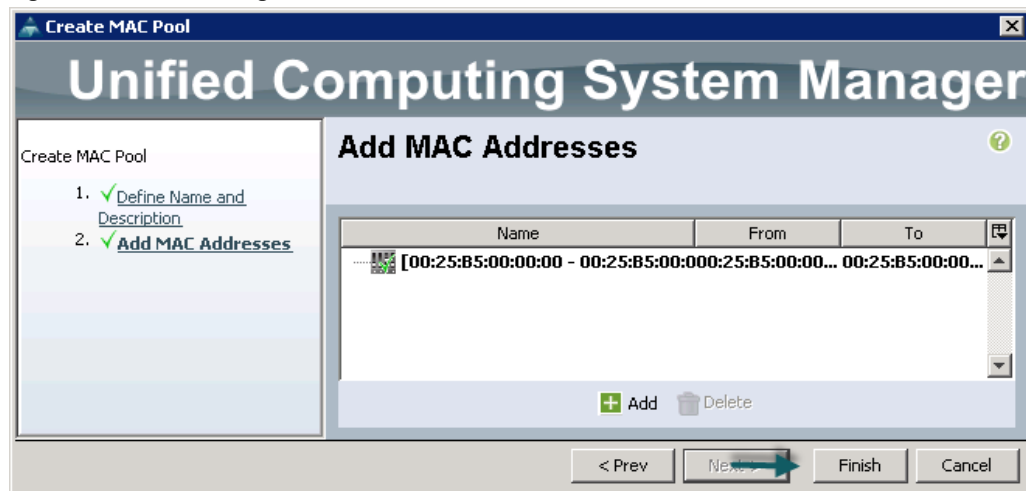


Figure 92 Specifying First MAC Address and Size



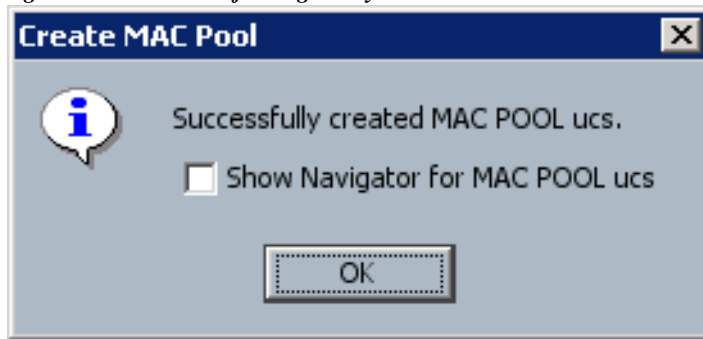
12. Click **Finish**.

Figure 93 Adding MAC Addresses



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

Figure 94 Confirming Newly Added MAC Pool



Configuring VLANs

VLANs are configured as in shown in table 6.

Table 9 VLAN Configurations

VLAN	Fabric	NIC Port	Function	Failover
vlan160_mgmt	A	eth0	Management, User connectivity	Fabric Failover to B
vlan12_HDFS	B	eth1	Hadoop	Fabric Failover to A
vlan11_DATA	A	eth2	Hadoop with multiple NICs support	Fabric Failover to B

All the VLANs created need to be trunked to the upstream distribution switch connecting the fabric interconnects. For this deployment VLAN160 is configured for management access (Installing and configuring OS, clustershell commands, setup NTP, user connectivity, etc) and vlan12_HDFS is configured for Hadoop Data traffic.

With some Hadoop distributions supporting multiple NICs, where Hadoop uses multiple IP subnets for its data traffic, vlan11_DATA 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 vlan11_DATA 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

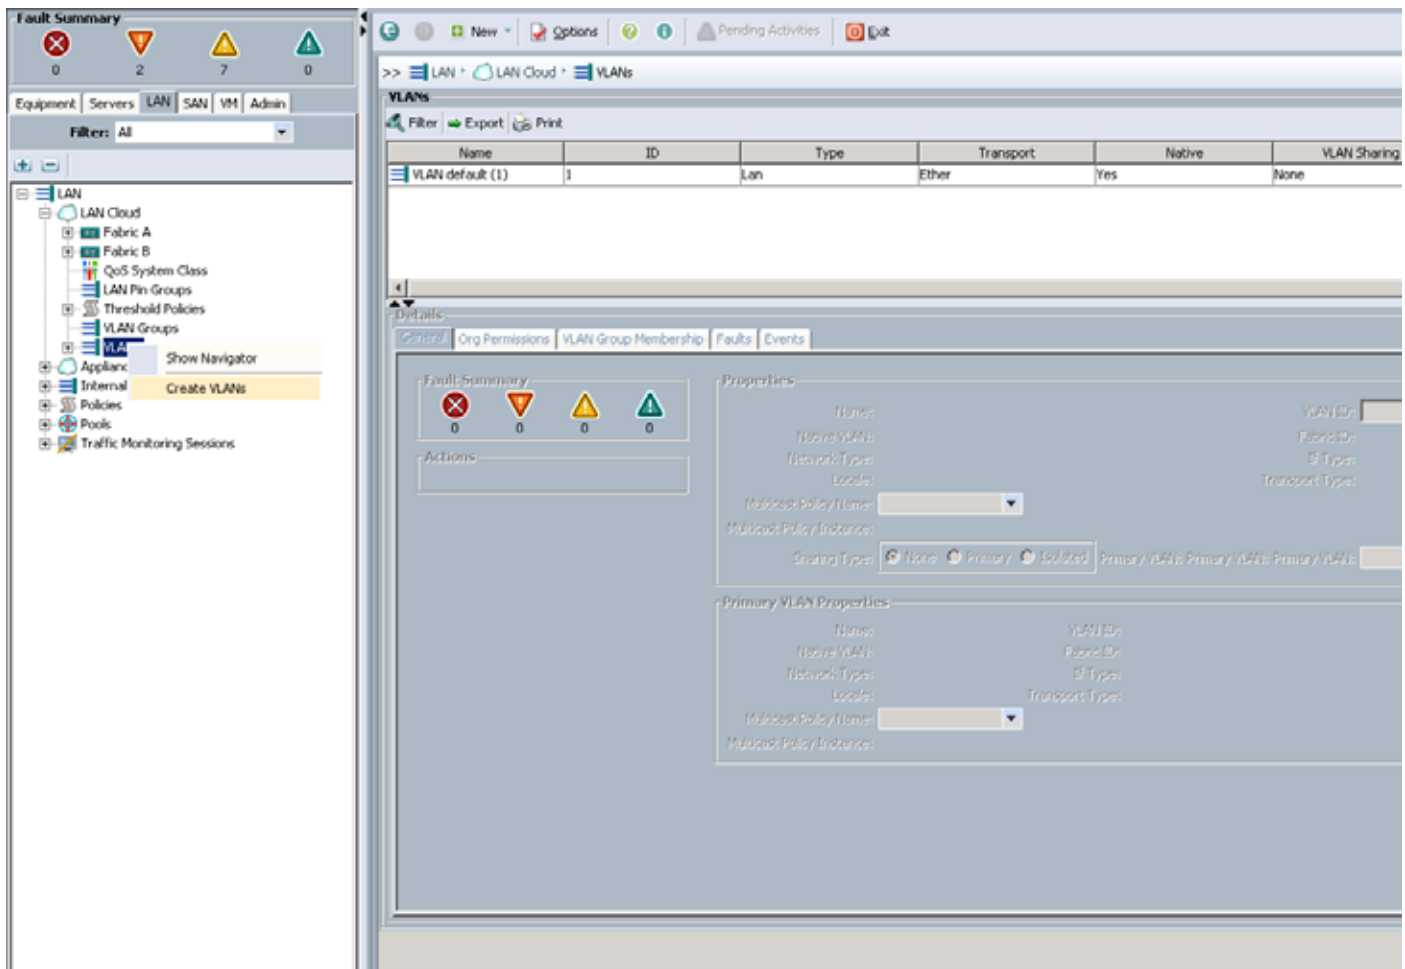
- **On Cloudera Security:** When deploying Cloudera with Security only one VLAN on one vNIC is supported. If the Cloudera install is going to have Security features enabled at a later stage, then use or create only single VLAN/vNIC in UCS Manager (could be name VLAN160_mgmt or VLAN12_HDFS) which will carry both management traffic and HDFS traffic. Ensure the MTU is set to 9000 and QoS policy is set to Platinum.
- If all the three VLANs are already created, and Cloudera Security needs to be enabled, then keep only one vNIC, VLAN160_Mgmt and delete rest of the vNICs from UCS Manager Service Profile Template. Modify the vNIC connected to VLAN160_Mgmt and update the MTU to 9000 and QoS Policy and before re-acknowledging the changes (this will lead to server reboot), on the servers

remove the configuration files for the vNICs
 /etc/sysconfig/network-scripts/<ifcfg-deleted-NICs> and re-acknowledge on the UCS Manager.
 This will restart the servers with only one vNIC/VLAN enabled.

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

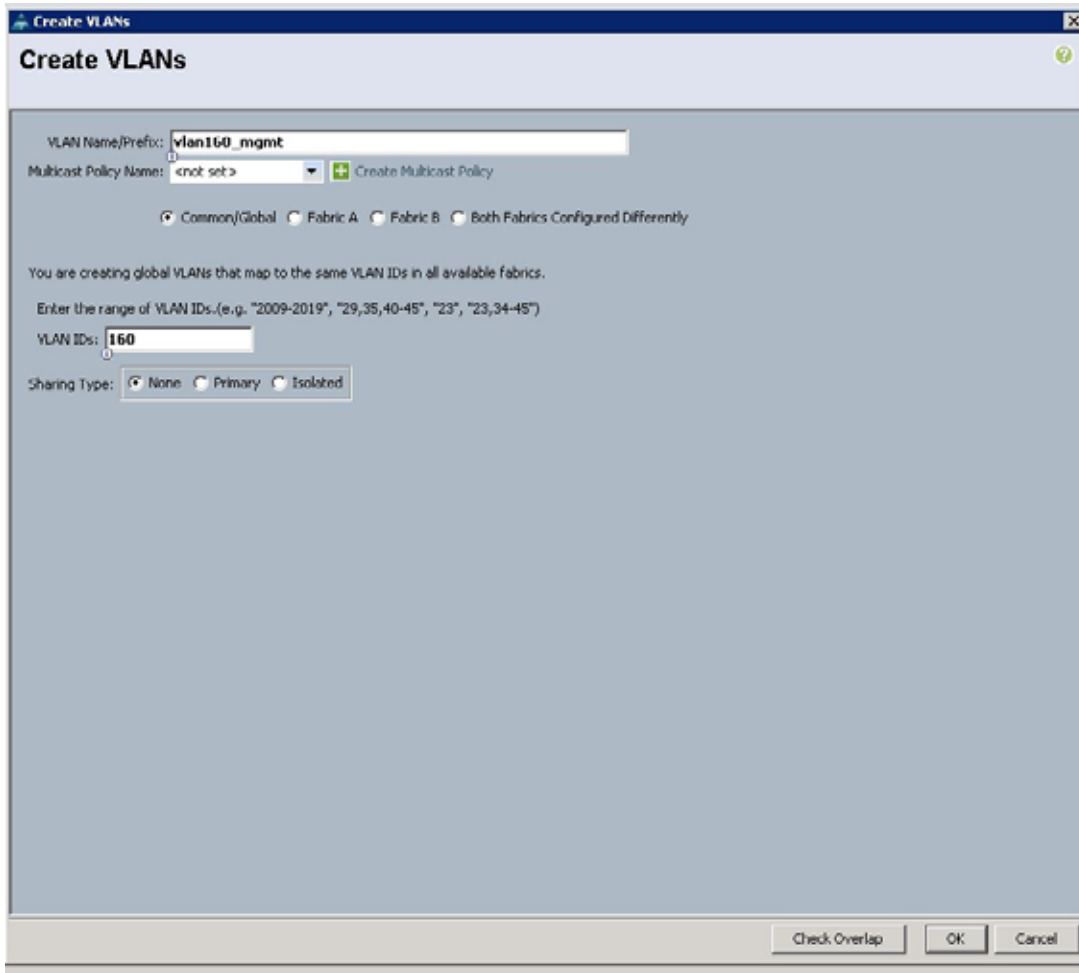
1. Select the **LAN** tab in the left pane in the UCS Manager GUI.
2. Select **LAN > VLANs**.
3. Right-click the **VLANs** under the root organization.
4. Select **Create VLANs** to create the VLAN.

Figure 95 Creating VLAN



5. Enter `vlan160_mgmt` for the VLAN Name.
6. Click the **Common/Global** radio button for the `vlan160_mgmt`.
7. Enter 160 on VLAN IDs of the Create VLAN IDs.
8. Click **OK** and then, click **Finish**.
9. Click **OK** in the success message box.

Figure 96 Creating VLAN for Management VLAN



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

Figure 97 Creating VLAN for Data

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name:

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

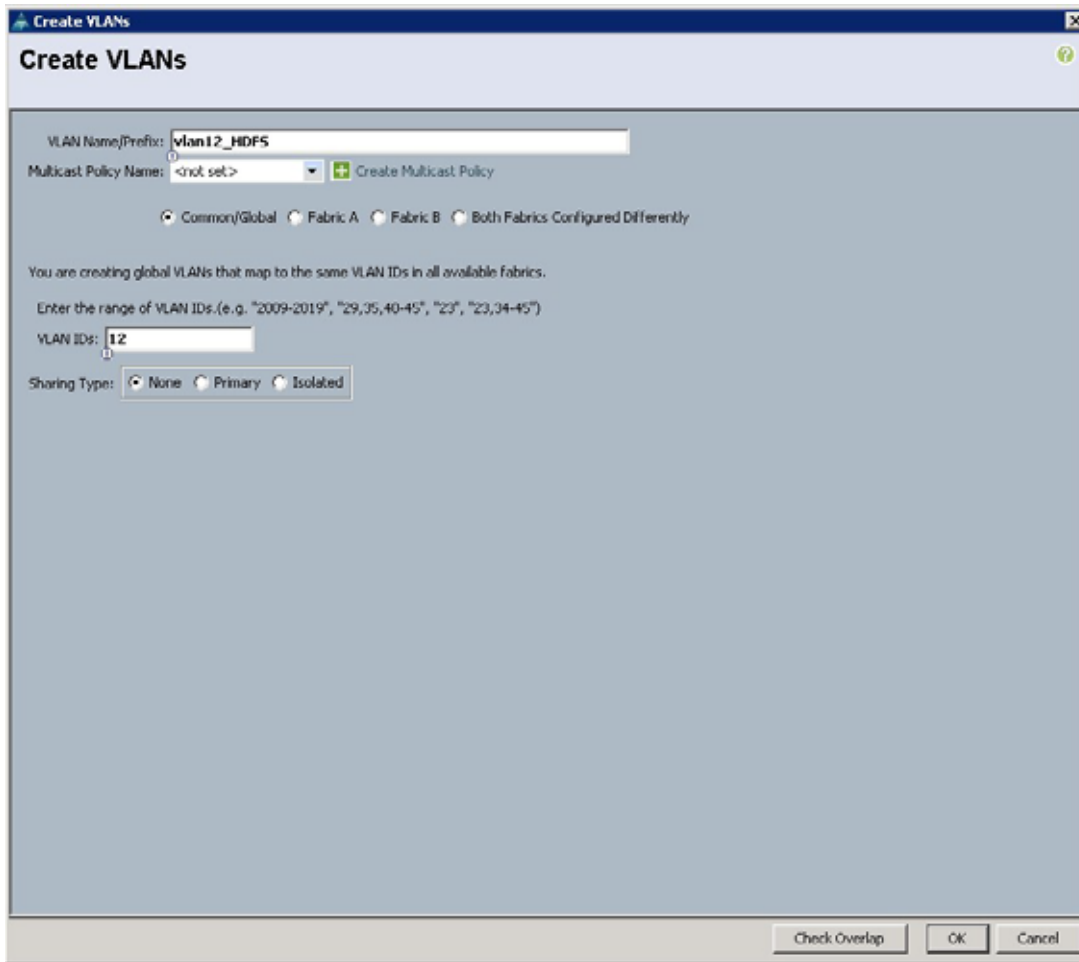
Enter the range of VLAN IDs. (e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: None Primary Isolated

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

Figure 98 Creating VLAN for Hadoop Data



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 99 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'. On the left, a progress indicator 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, 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 100 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 101 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

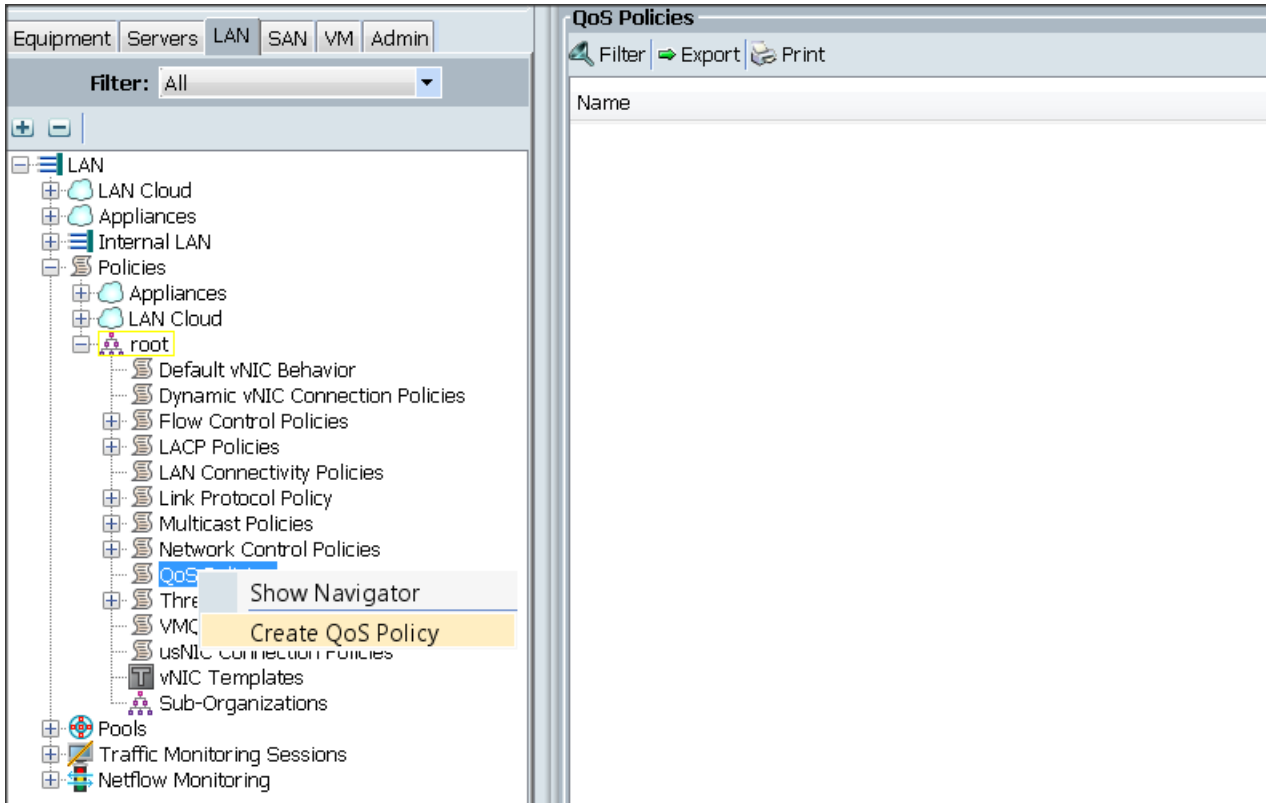
Follow these steps to create the QoS policy for a given server configuration using the Cisco UCS Manager GUI:

Best Effort Policy

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

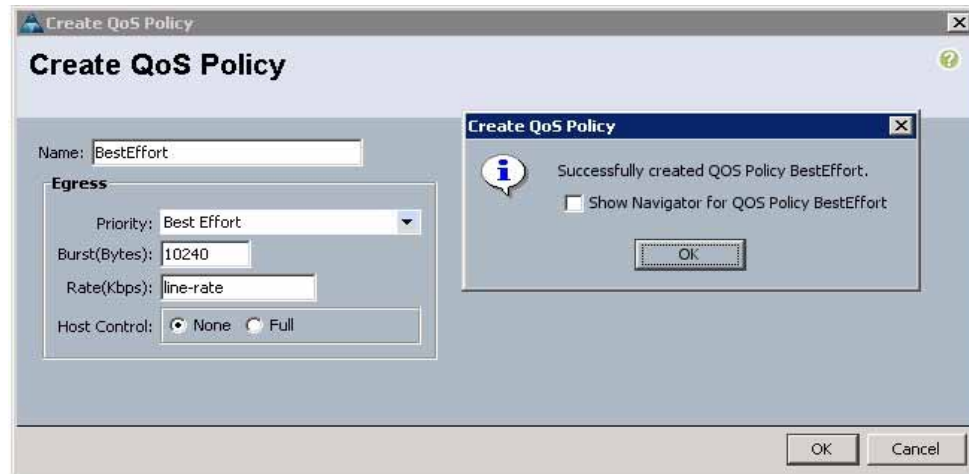
3. Right-click **QoS Policies**.
4. Select **Create QoS Policy**.

Figure 102 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.

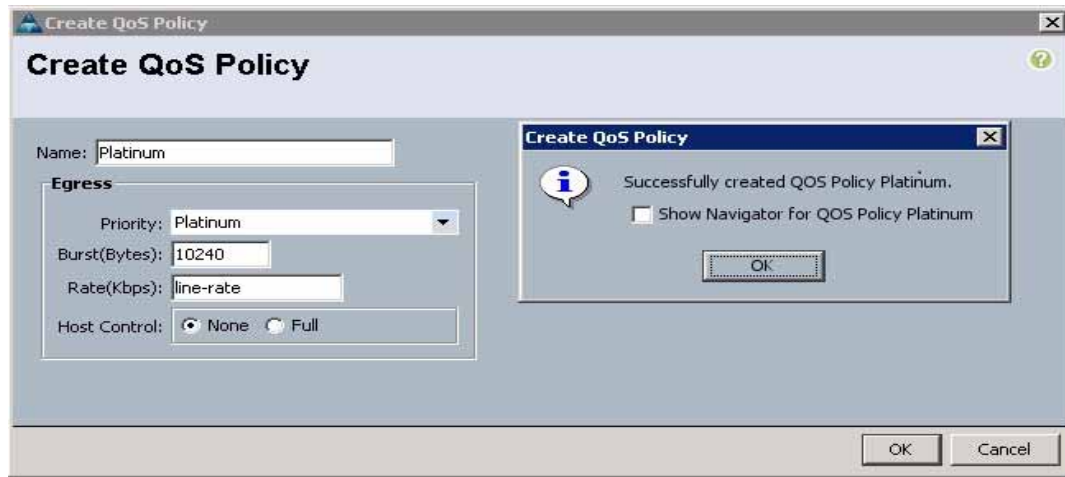
Figure 103 Creating BestEffort QoS Policy



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 104 Creating Platinum QoS Policy

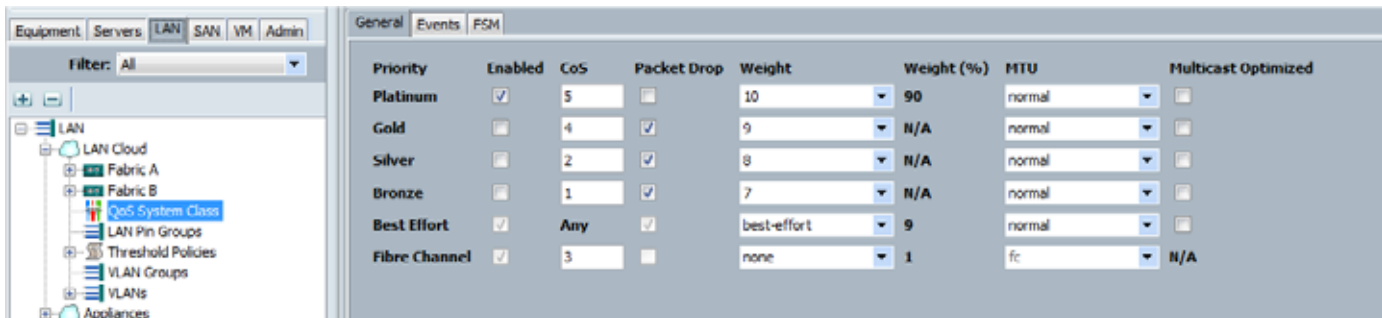


Setting Jumbo Frames

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. Check the **Enabled** Check box next to Platinum.
5. In the Best Effort row, select best-effort for weight.
6. In the Fiber Channel row, select none for weight.
7. Click **Save Changes**.
8. Click **OK**.

Figure 105 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 106 *Configuring Local Disk Policy*

Create Local Disk Configuration Policy

Name:

Description:

Mode:

Protect Configuration:

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

FlexFlash

FlexFlash State: Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State: Disable Enable

OK Cancel

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 107 Creating Server BIOS Policy

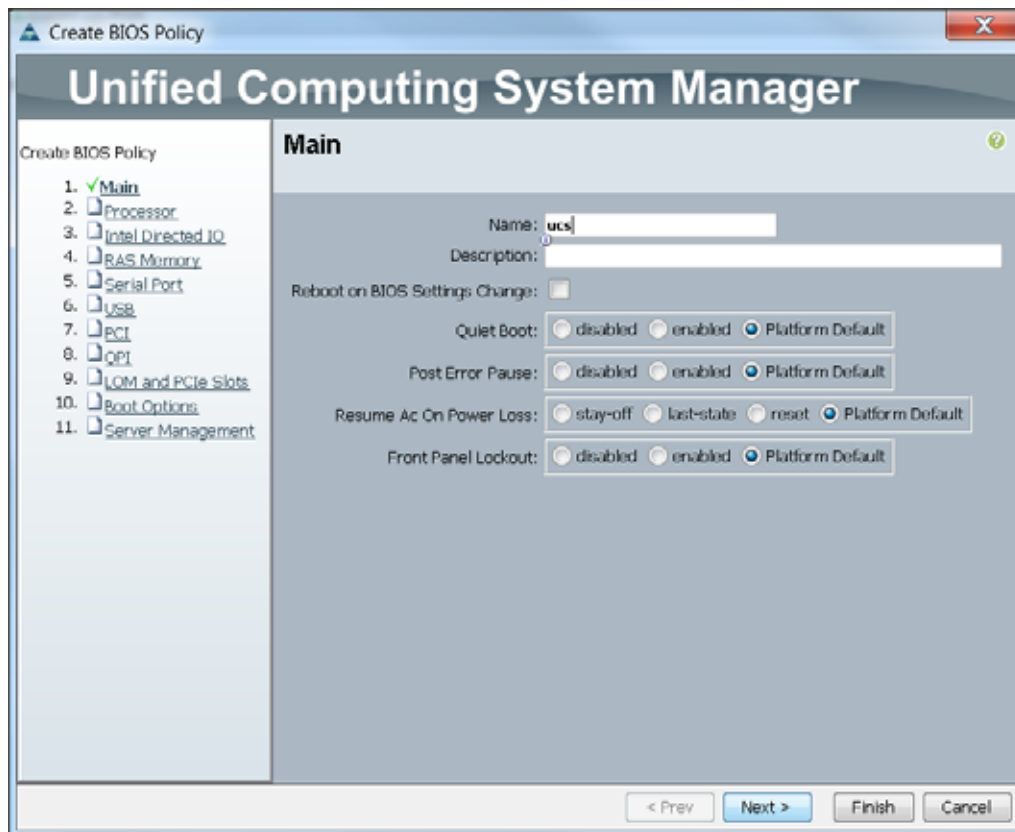


Figure 108 Creating Server BIOS Policy for Processor

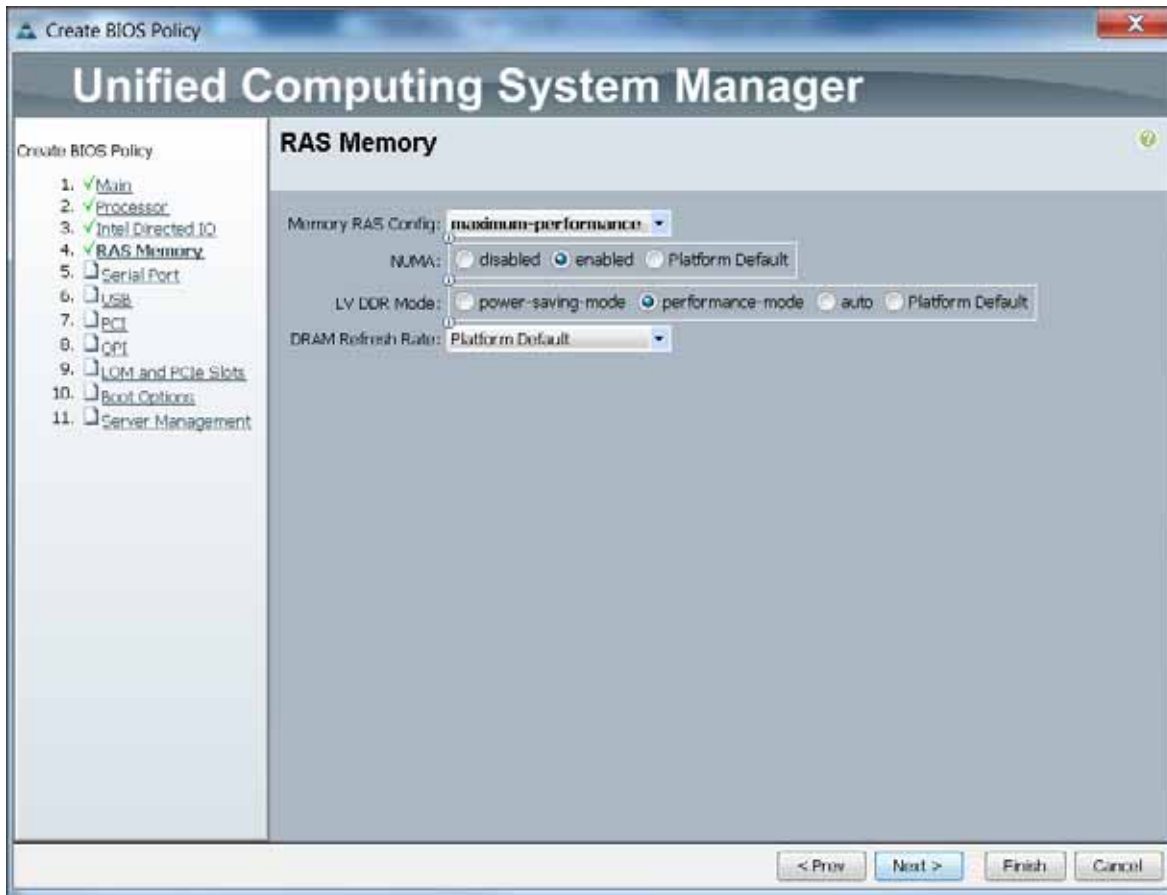


Figure 109 Creating Server BIOS Policy for Intel Directed IO



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

Figure 110 Creating Server BIOS Policy for Memory

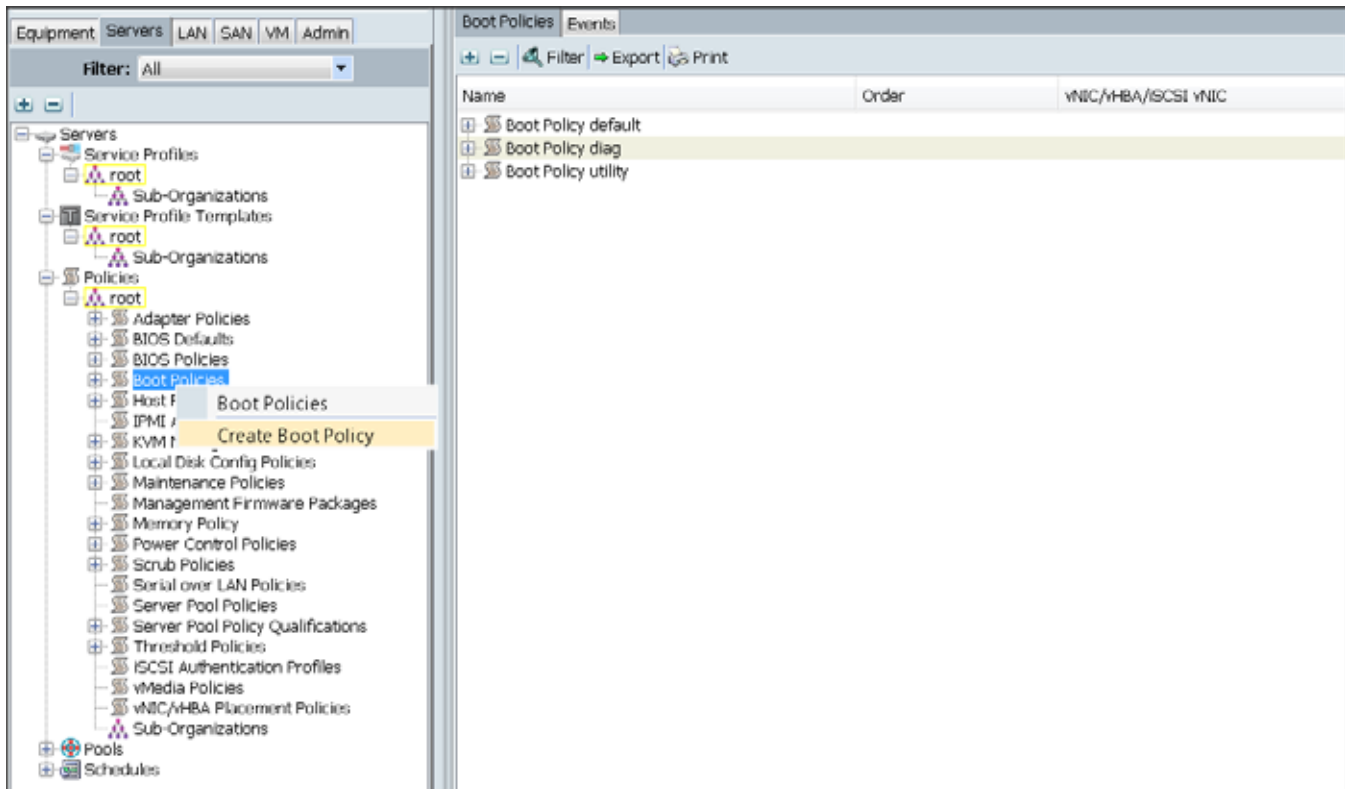


Creating Boot Policy

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

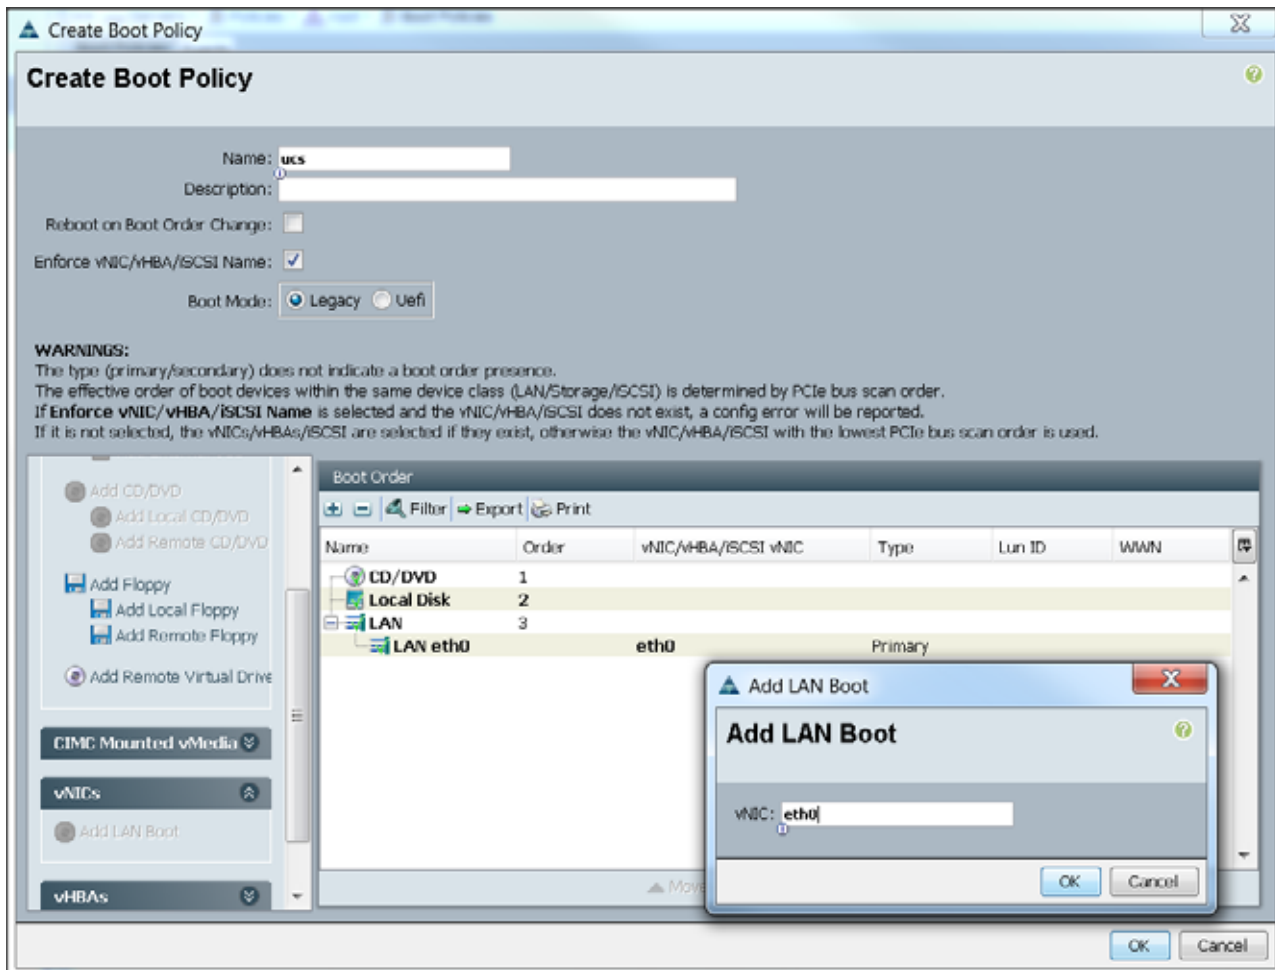
1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.
3. Right-click the **Boot Policies**.
4. Select **Create Boot Policy**.

Figure 111 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 112 Creating Boot Policy Part 2

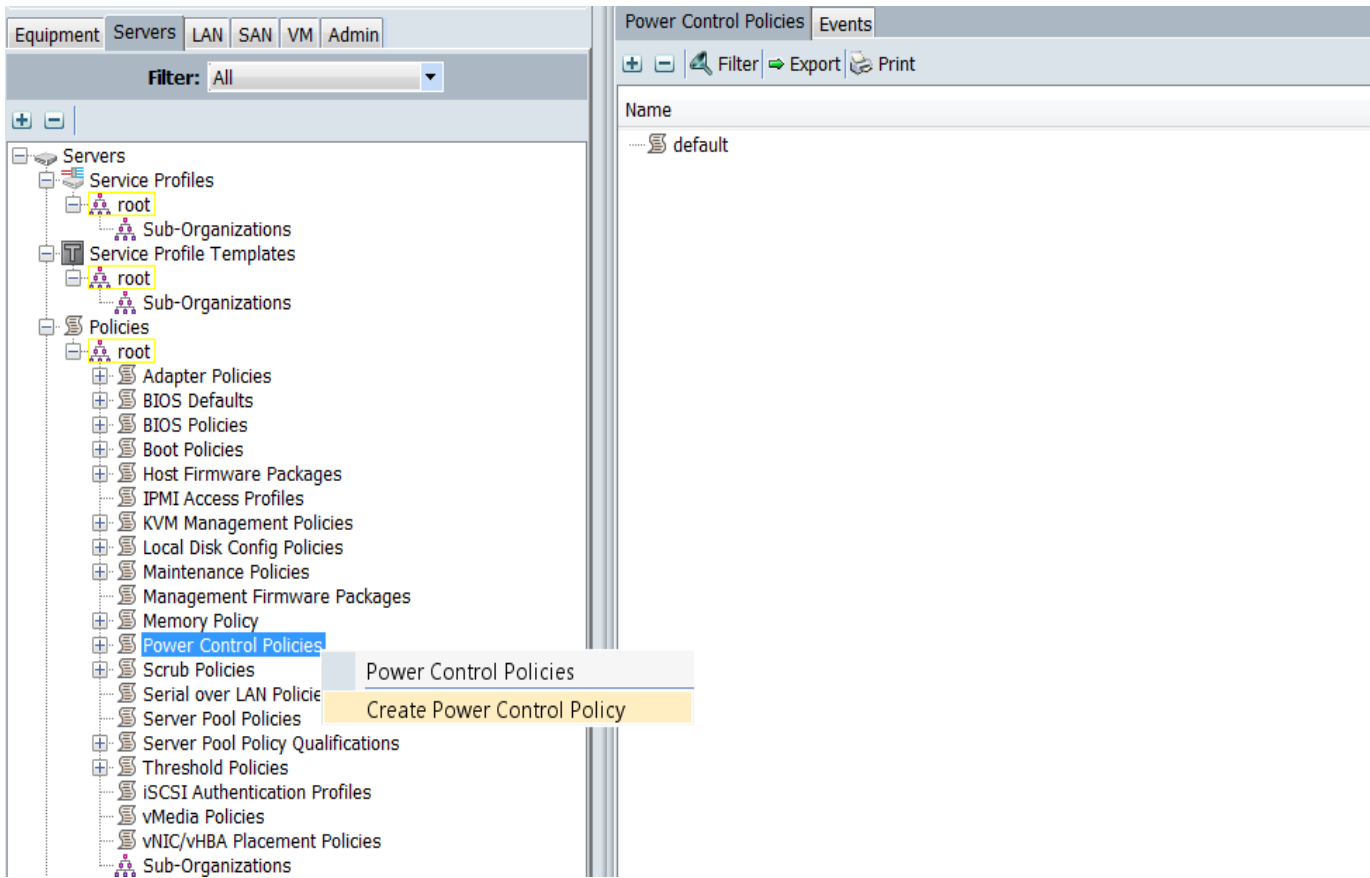


Creating Power Control Policy

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

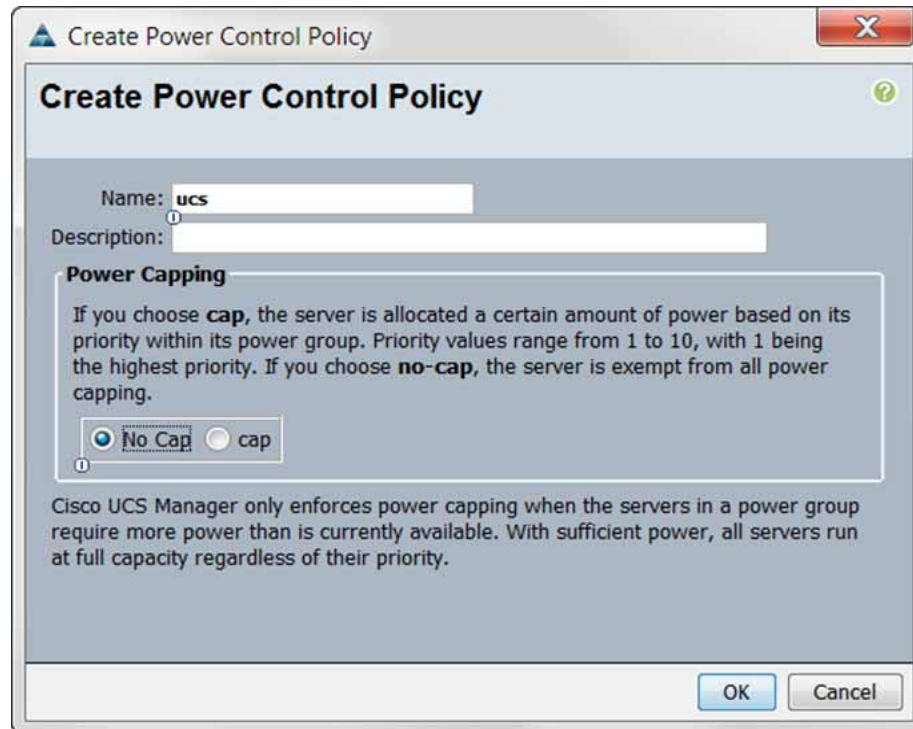
15. Select the **Servers** tab in the left pane in the UCS Manager GUI.
16. Select **Policies > root**.
17. Right-click the **Power Control Policies**.
18. Select **Create Power Control Policy**.

Figure 113 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 114 Creating Power Control Policy Part 2

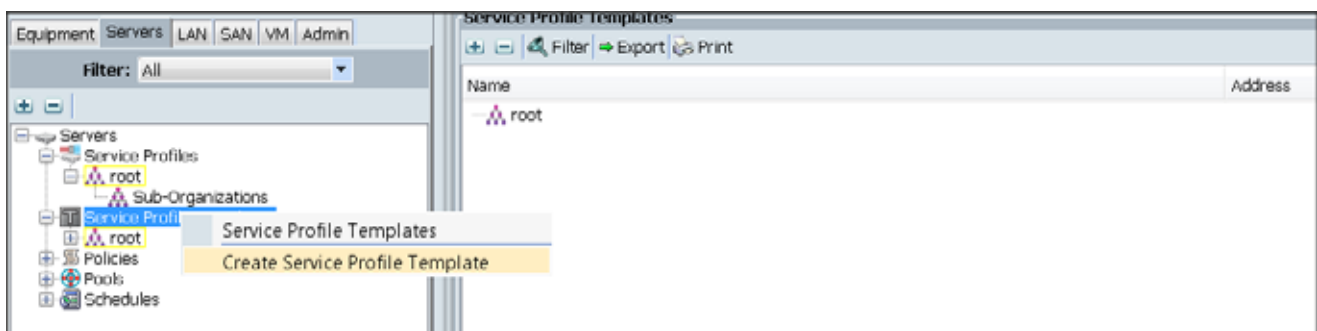


Creating Service Profile Template

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 115 Creating Service Profile Template



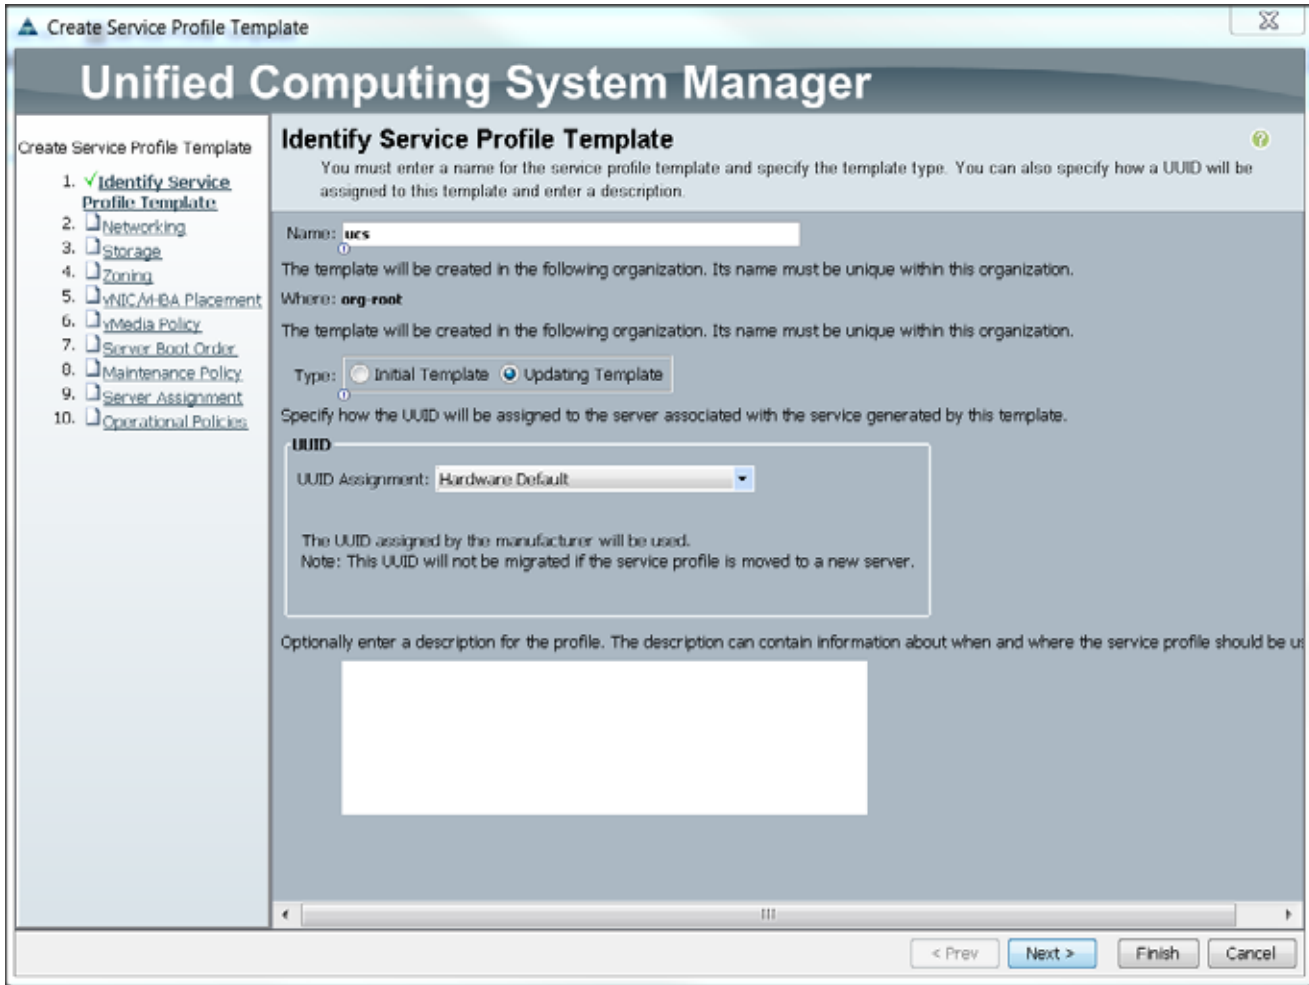
4. The Create Service Profile Template window appears.

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 116 Identify Service Profile Template



Configuring Network Settings for the Template

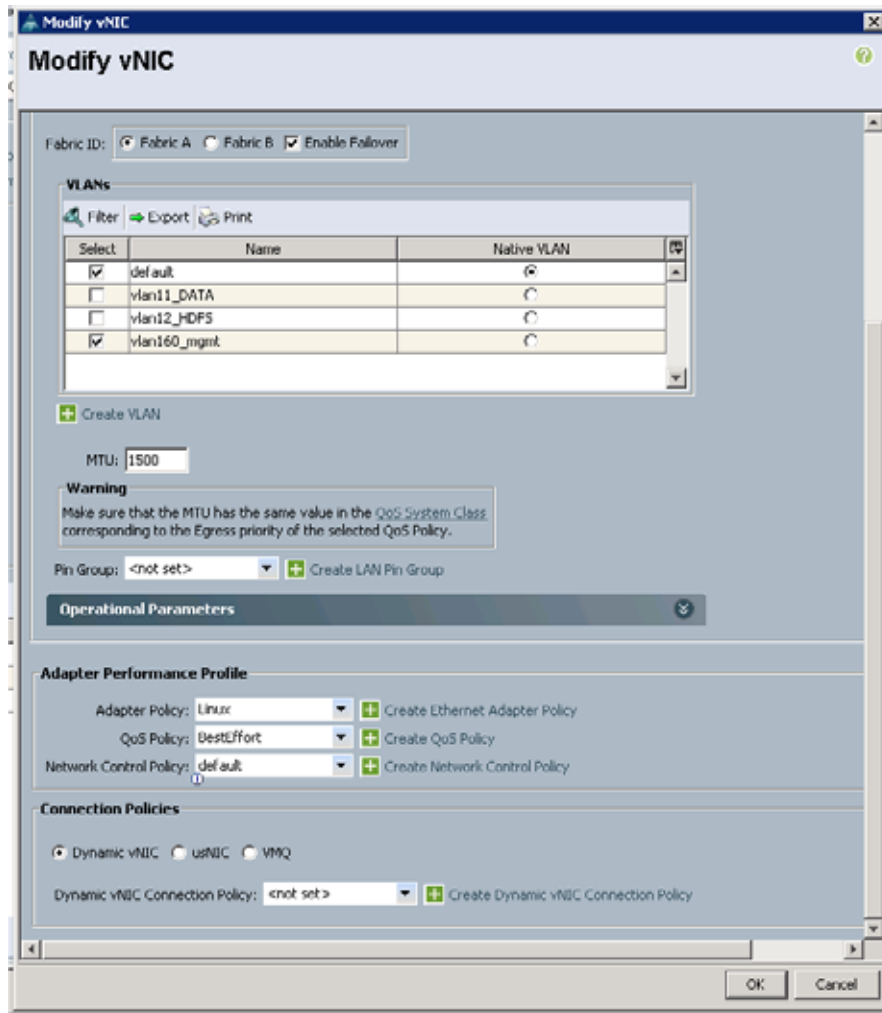
1. Keep the Dynamic vNIC Connection Policy field at the default.
2. Click the **Expert** radio button for the option, **how would you like to configure LAN connectivity?**
3. Click **Add** to add a vNIC to the template.

Figure 117 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 vlan160_mgmt check box for VLANs and select the Native VLAN default 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 118 Configuring vNIC eth0



15. The Create vNIC window appears. Name the vNIC eth1.
16. Select ucs in the Mac Address Assignment pool.
17. Click the **Fabric B** radio button and Check the **Enable failover** check box for the Fabric ID.
18. Check the **vlan12_HDFS** check box for VLANs and select the Native VLAN default radio button.
19. Select MTU size as 1500.
20. Select adapter policy as Linux.
21. Select QoS Policy as Platinum.
22. Keep the Network Control Policy as Default.
23. Keep the Connection Policies as Dynamic vNIC.
24. Keep the Dynamic vNIC Connection Policy as <not set>.
25. Click **OK**.

Figure 119 Configuring vNIC eth1

Fabric ID: Fabric A Fabric B Enable Failover

VLANs

Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input type="checkbox"/>	Vlan11_Data	<input type="radio"/>
<input checked="" type="checkbox"/>	Vlan12_HDFS	<input type="radio"/>
<input type="checkbox"/>	Vlan160_mgmt	<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

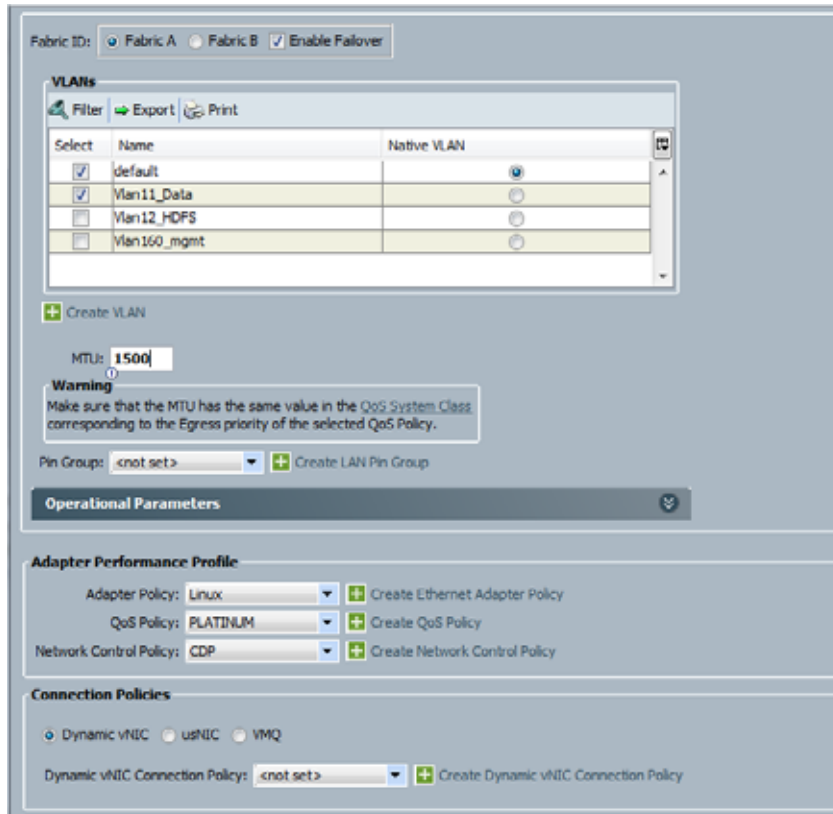
Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy: Create Dynamic vNIC Connection Policy

26. The Create vNIC window appears. Name the vNIC eth2.
27. Select ucs in the Mac Address Assignment pool.
28. Click the **Fabric A** radio button, and then Check the **Enable failover** check box for the Fabric ID.
29. Check the **vlan11_DATA** check box for VLANs and select the Native VLAN default radio button.
30. Select MTU size as 1500.
31. Select adapter policy as Linux.
32. Select QoS Policy as Platinum.
33. Keep the Network Control Policy as Default.
34. Keep the Connection Policies as Dynamic vNIC.
35. Keep the Dynamic vNIC Connection Policy as <not set>.
36. Click **OK**.

Figure 120 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 121 Configuring Storage Settings

Storage
Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage:

Mode: **Any Configuration**
Protect Configuration: **Yes**

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

FlexFlash

FlexFlash State: **Disable**

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State: **Disable**

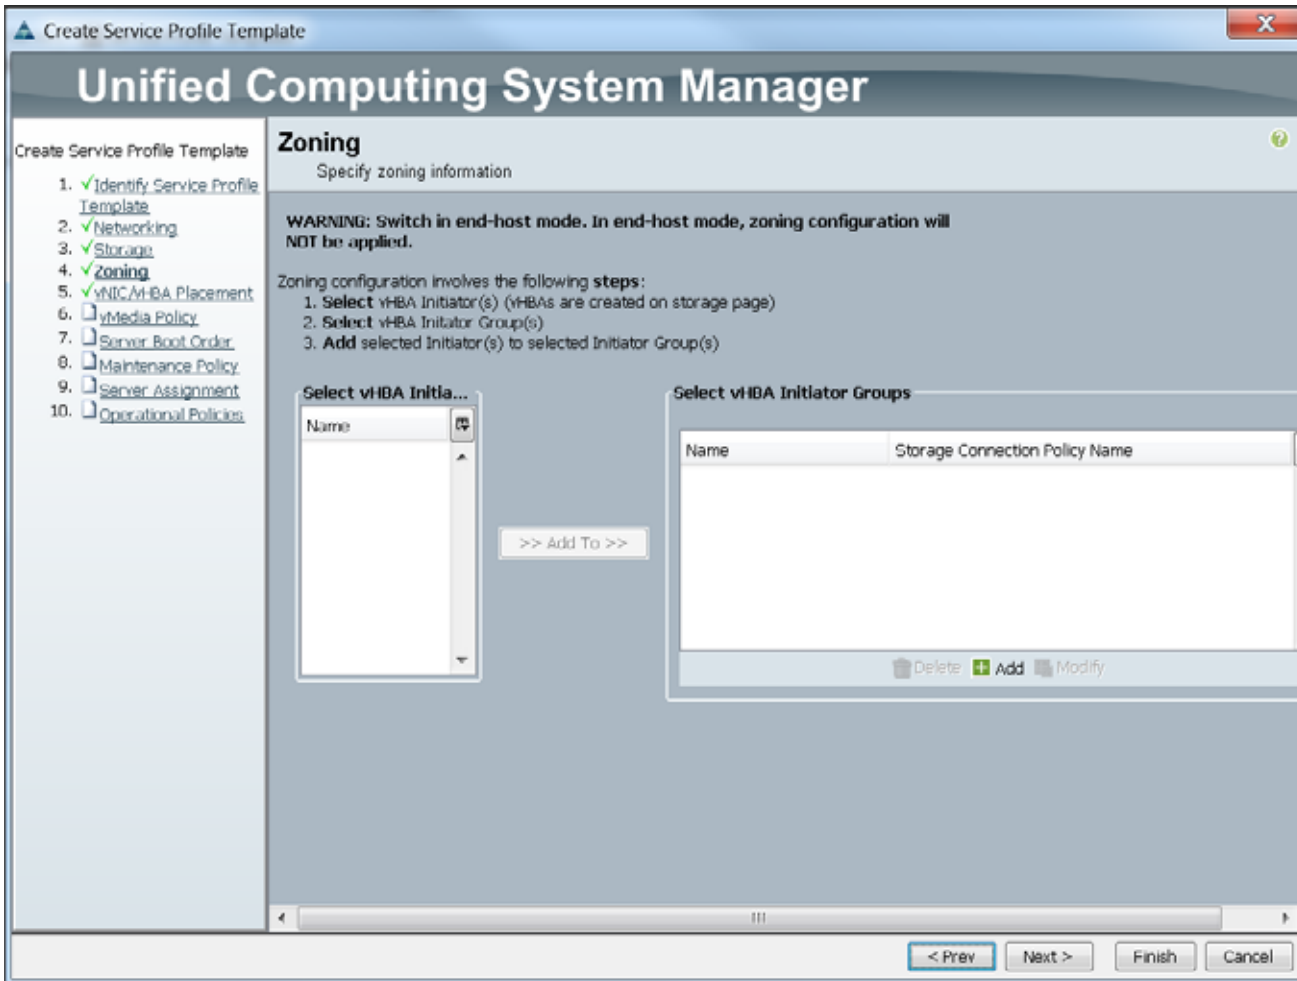
How would you like to configure SAN connectivity? Simple Expert No iHBAs Use Connectivity Policy

This server associated with this service profile will not be connected to a storage area network.

< Prev Next > Finish Cancel

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

Figure 122 Configure Zoning

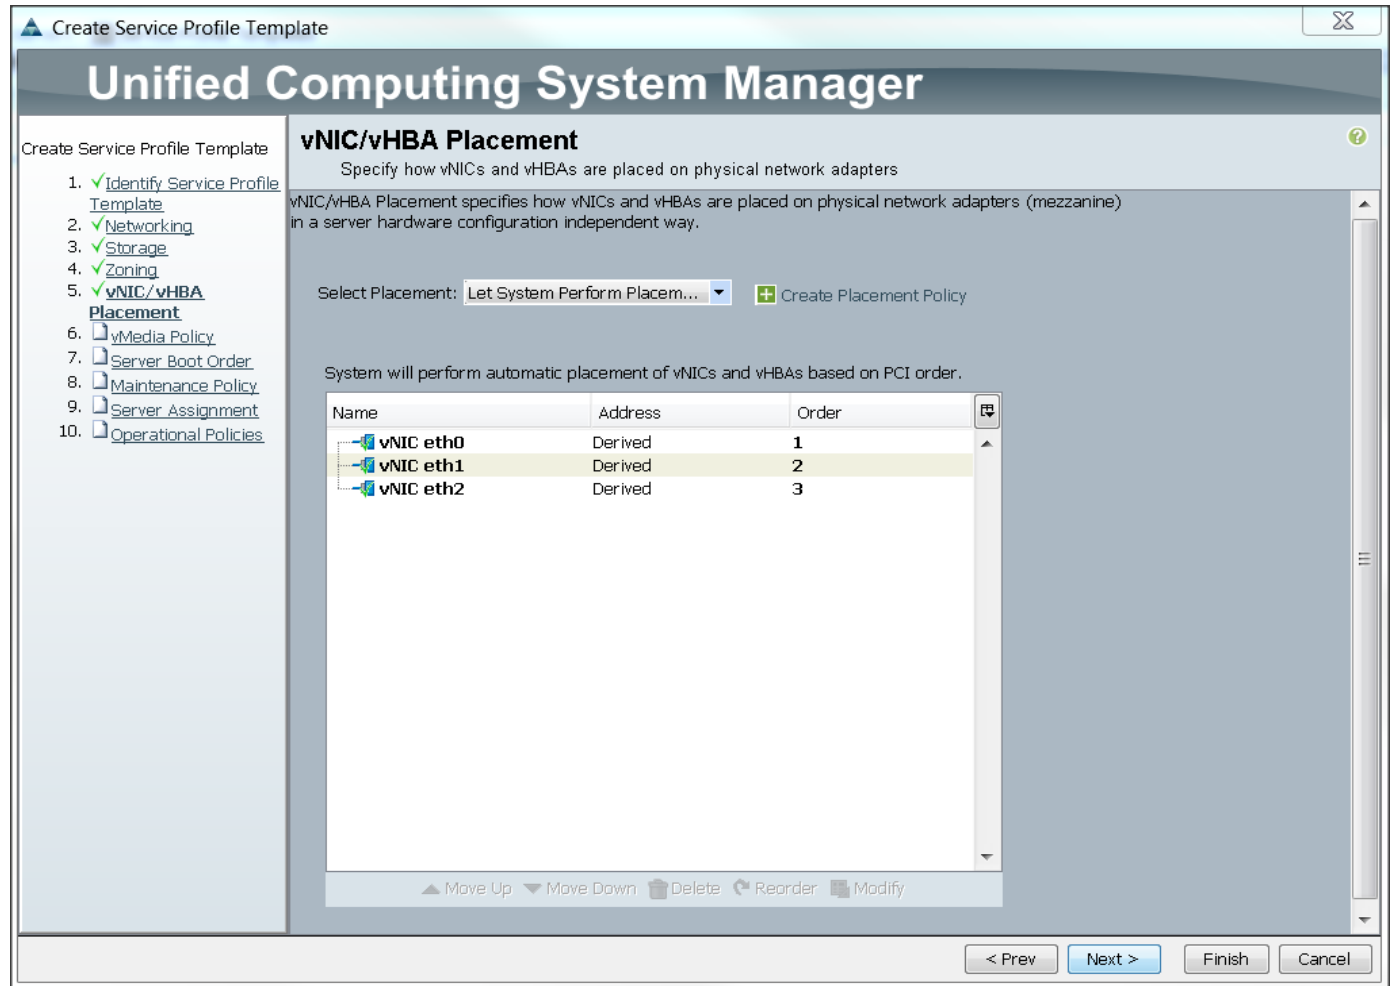


Configuring vNIC/vHBA Placement for the Template

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 123 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 124 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 125 Creating Boot Policy



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

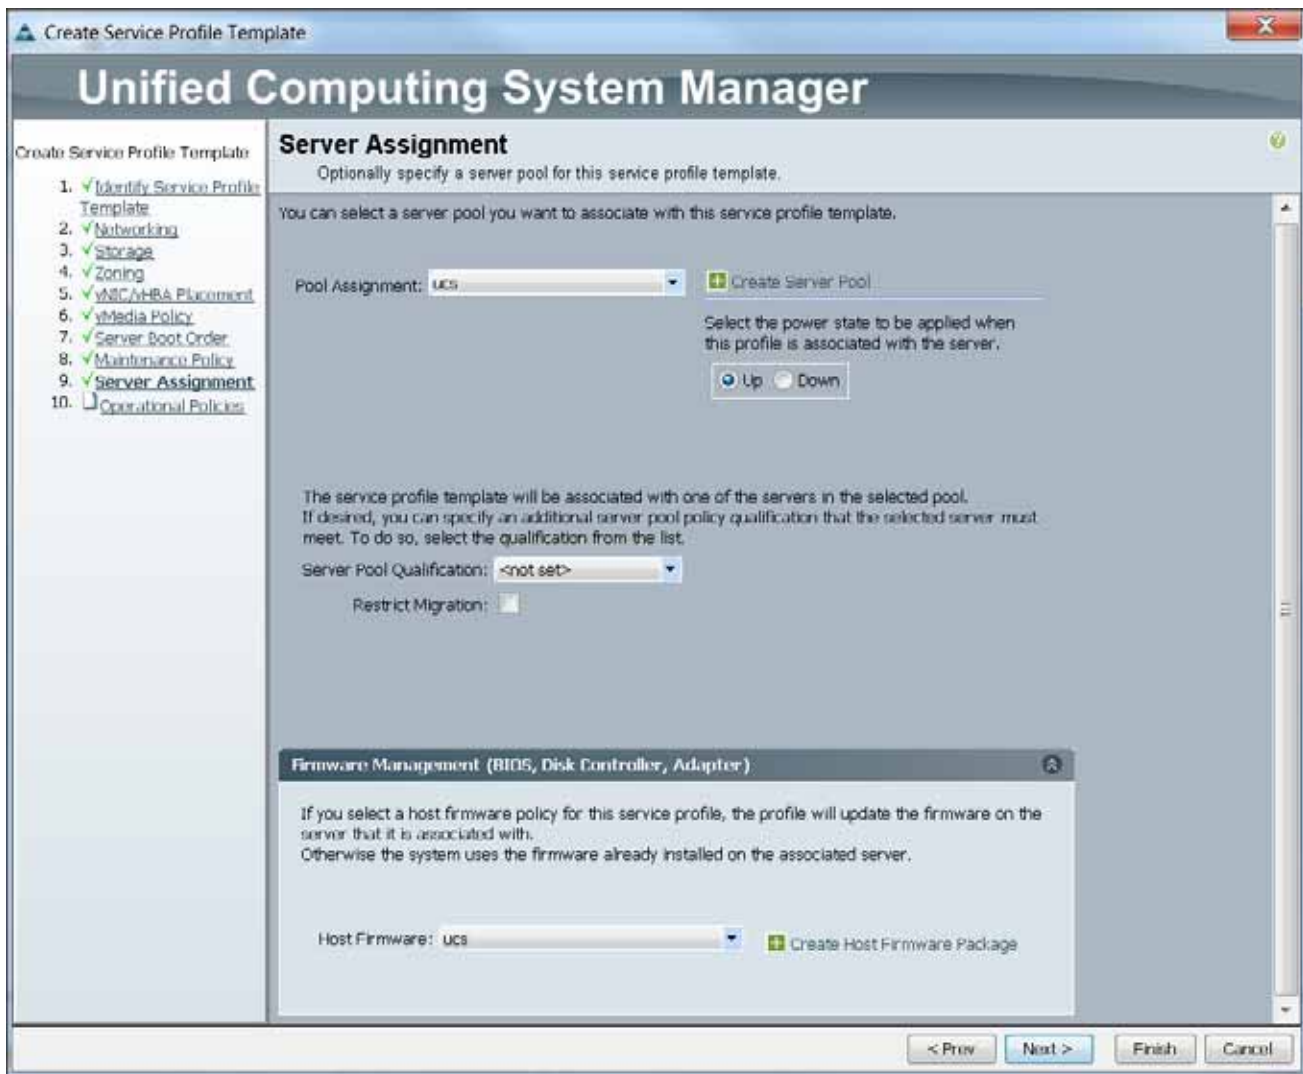
1. Keep the Maintenance policy at no policy used by default.
2. Click **Next** to continue to the next section.

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 126 Server Assignment

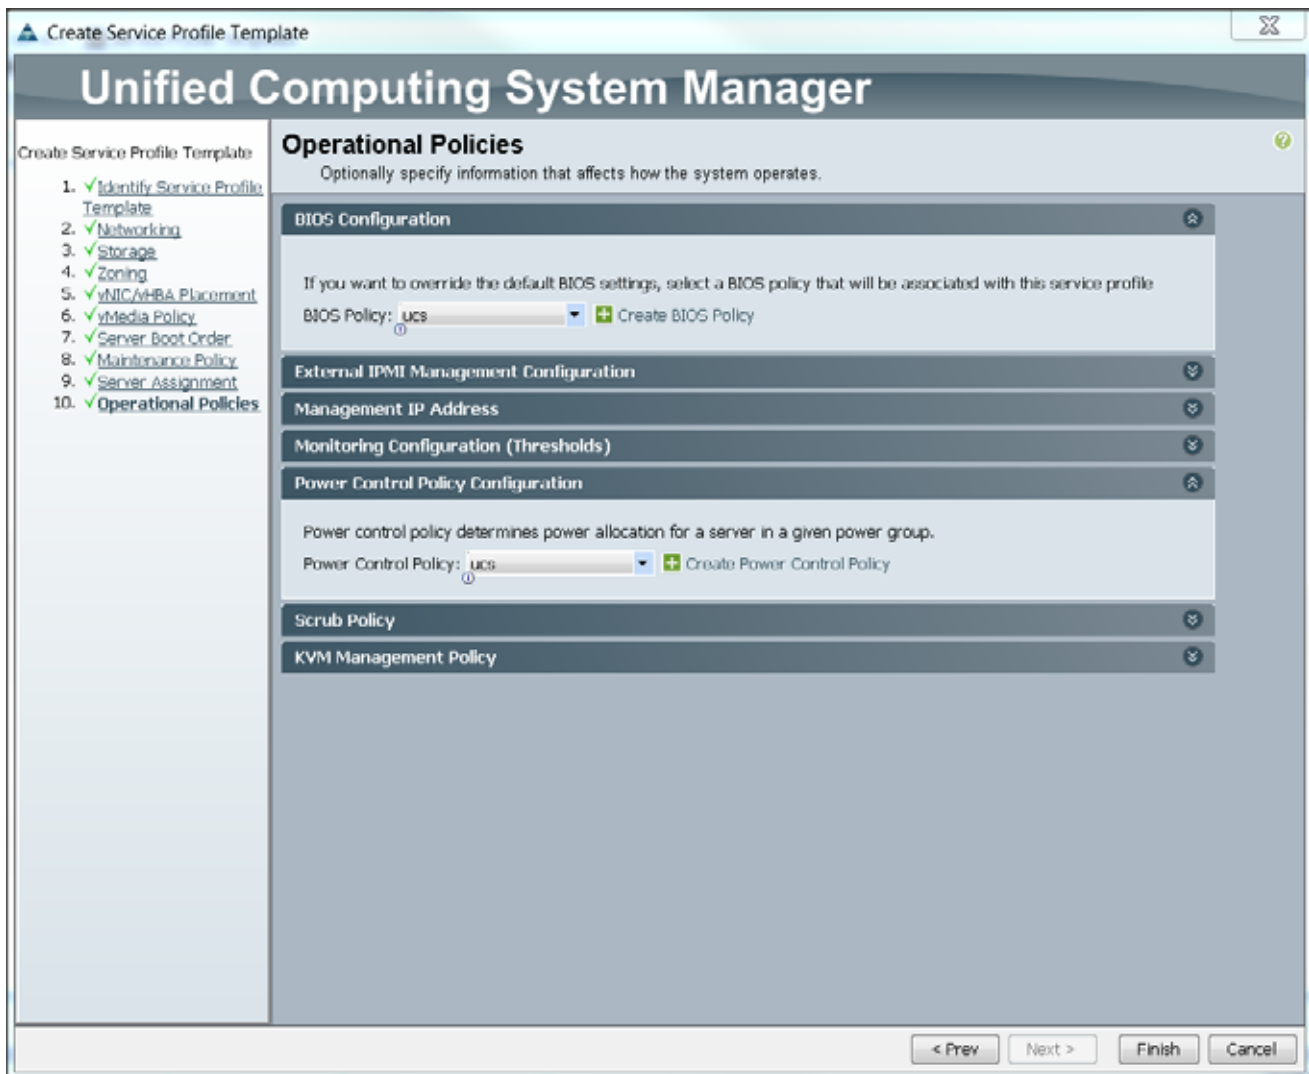


Configuring Operational Policies for the Template

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.

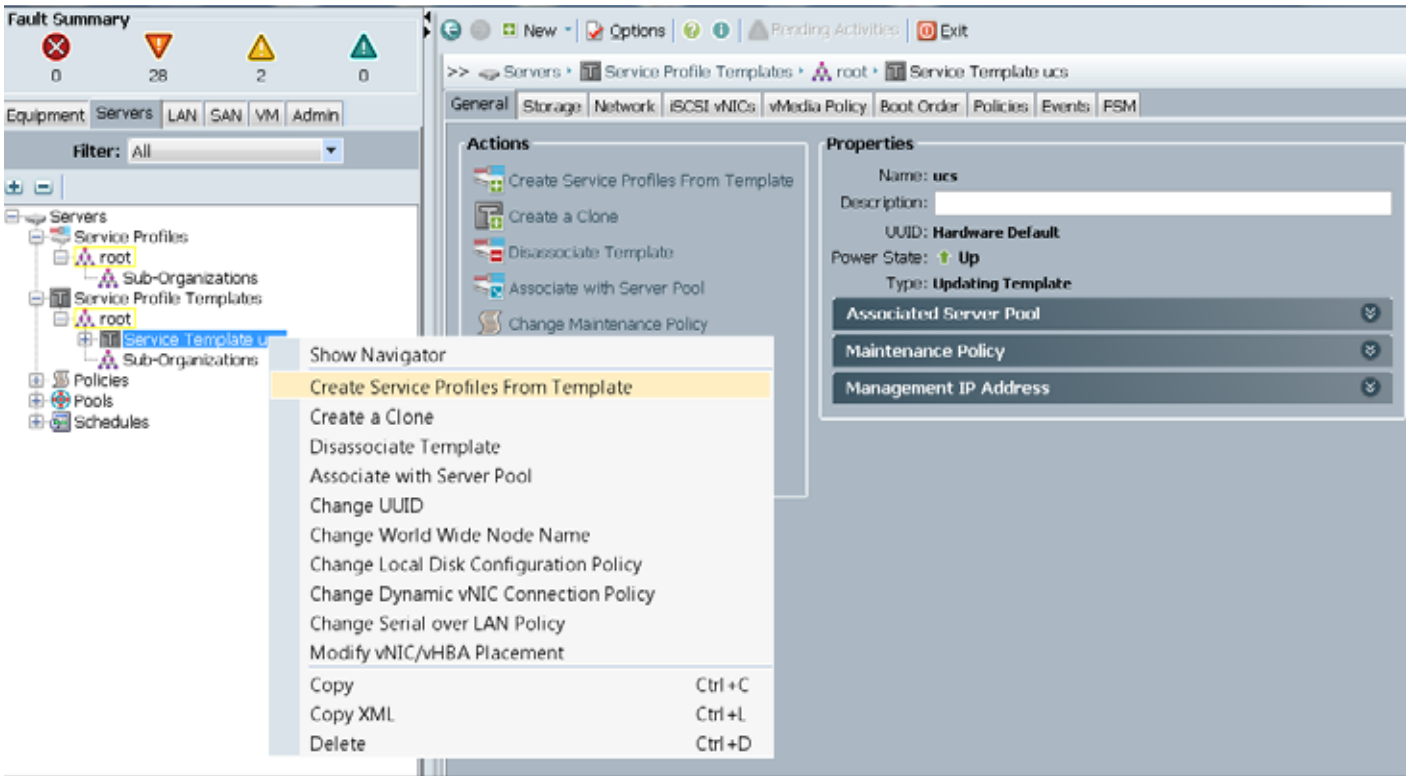
Figure 127 Selecting BIOS and Power Control Policy



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

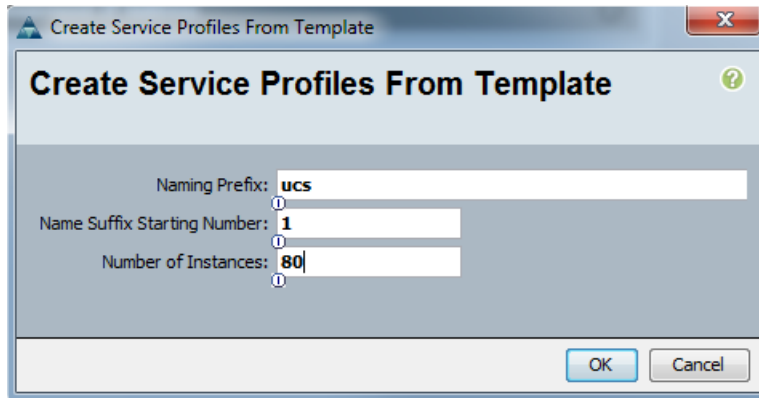
1. Go to Service Profile **Templates** > **root**.
2. Right-click Service Profile Templates ucs.
3. Select Create Service Profiles From Template.

Figure 128 Creating Service Profiles from Template



4. The Create Service Profile from Template window appears.

Figure 129 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 131.

Figure 130 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...	FOH1852V0PU		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 2	Ok	UCSC-C240...	Cisco UCS C...	FOH1850V38U		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 3	Ok	UCSC-C240...	Cisco UCS C...	FOH1844V00C		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 4	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0PY		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 5	Ok	UCSC-C240...	Cisco UCS C...	FOH1851V1Z2		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 6	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0L4		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 7	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0Q3		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 8	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V00C		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 9	Ok	UCSC-C240...	Cisco UCS C...	FOH1851V233		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 10	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0NF		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 11	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0RP		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 12	Ok	UCSC-C240...	Cisco UCS C...	FOH1851V213		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 13	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V00F		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 14	Ok	UCSC-C240...	Cisco UCS C...	FOH1851V243		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 15	Ok	UCSC-C240...	Cisco UCS C...	FOH1851V216		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A
Server 16	Ok	UCSC-C240...	Cisco UCS C...	FOH1852V0MA		24	262144	1	0	0	Operable	On	Associated	org-root/ls...	N/A

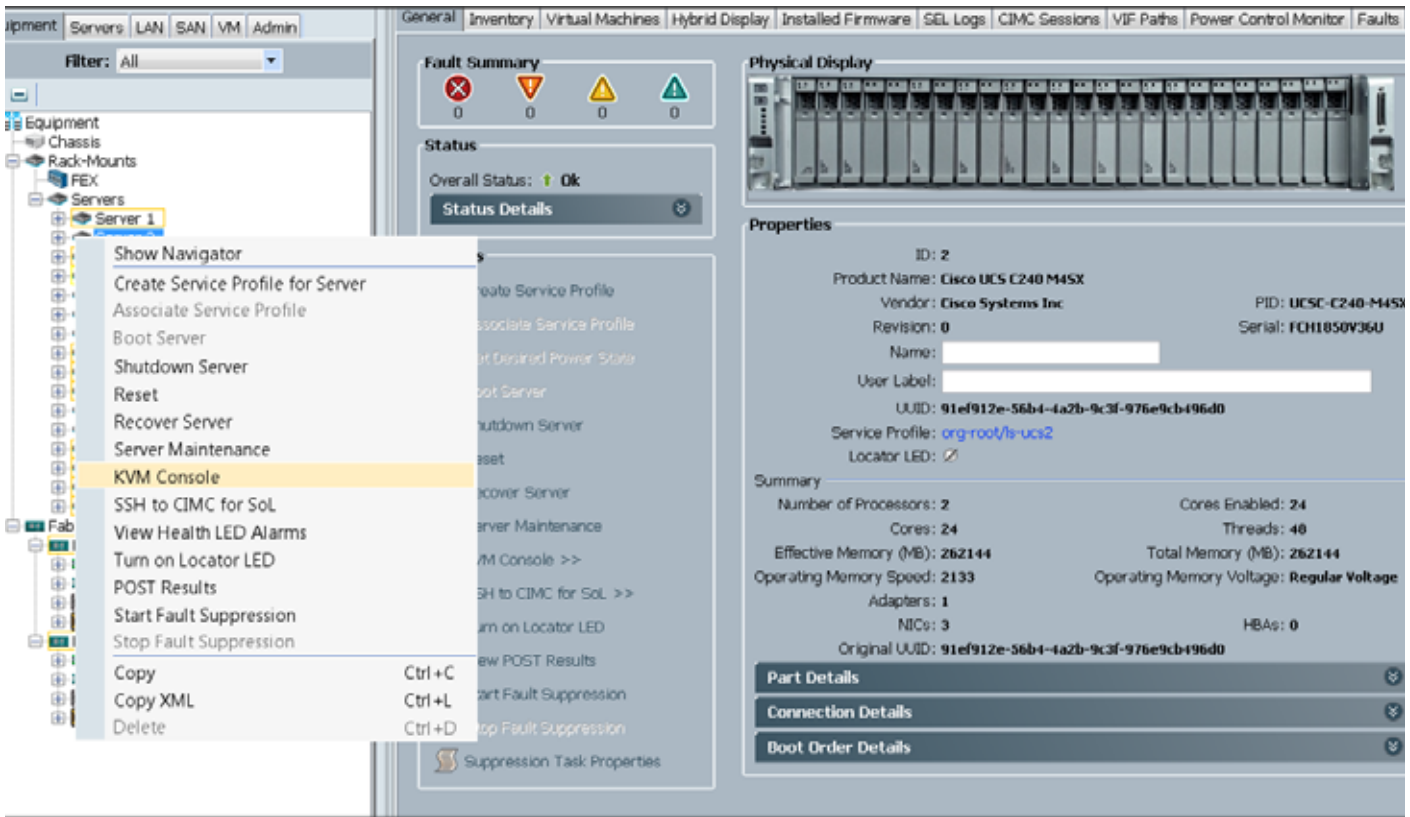
Installing Redhat Enterprise Linux 6.5 software RAID on Cisco UCS C240M4 Servers

The following section provides detailed procedure for installing Red Hat 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.


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 Servers.
4. Right-click the server and select **KVM Console**.

Figure 131 Selecting KVM Console Option



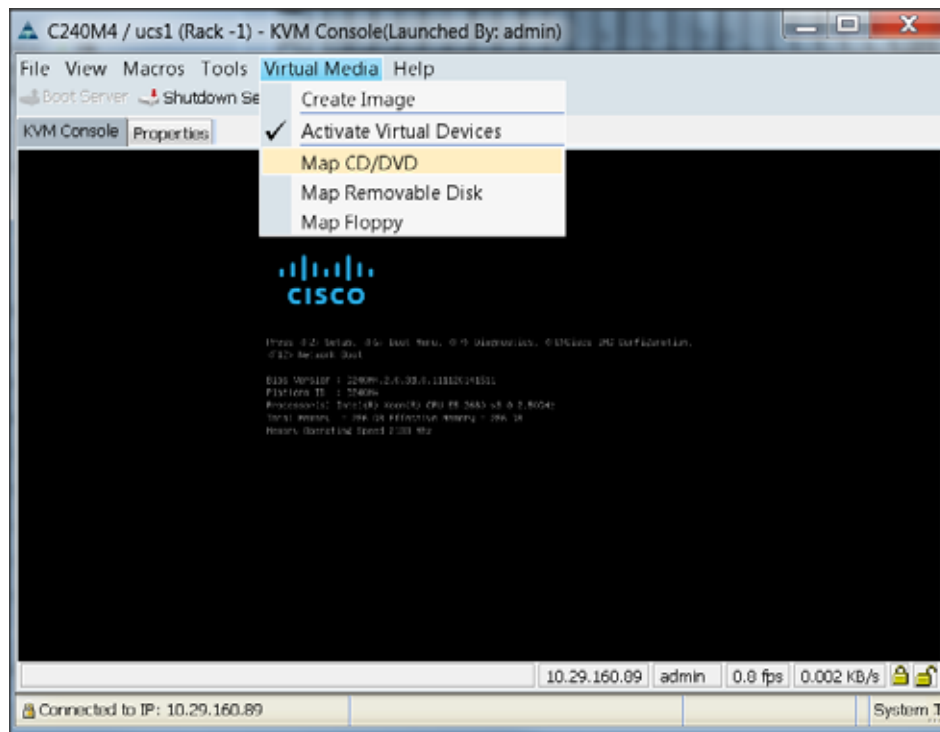
1. In the KVM window, select the **Virtual Media** tab.
2. Click the **Activate Virtual Devices** found under **Virtual Media** tab.

Figure 132 Selecting Activate Virtual Devices



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

Figure 133 Mapping ISO Image



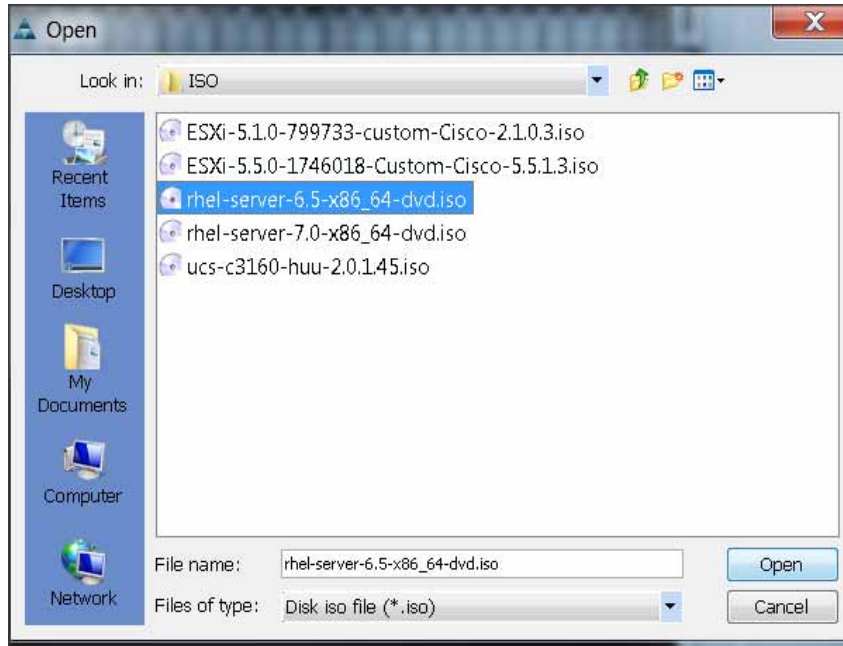
4. Browse to the Red Hat Enterprise Linux Server 6.5 installer ISO image file.



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 134 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 135 Red Hat Enterprise Linux Server 6.5 Install Media



12. Skip the Media test and start the installation

Figure 136 RHEL: Media Test and Start of Installation



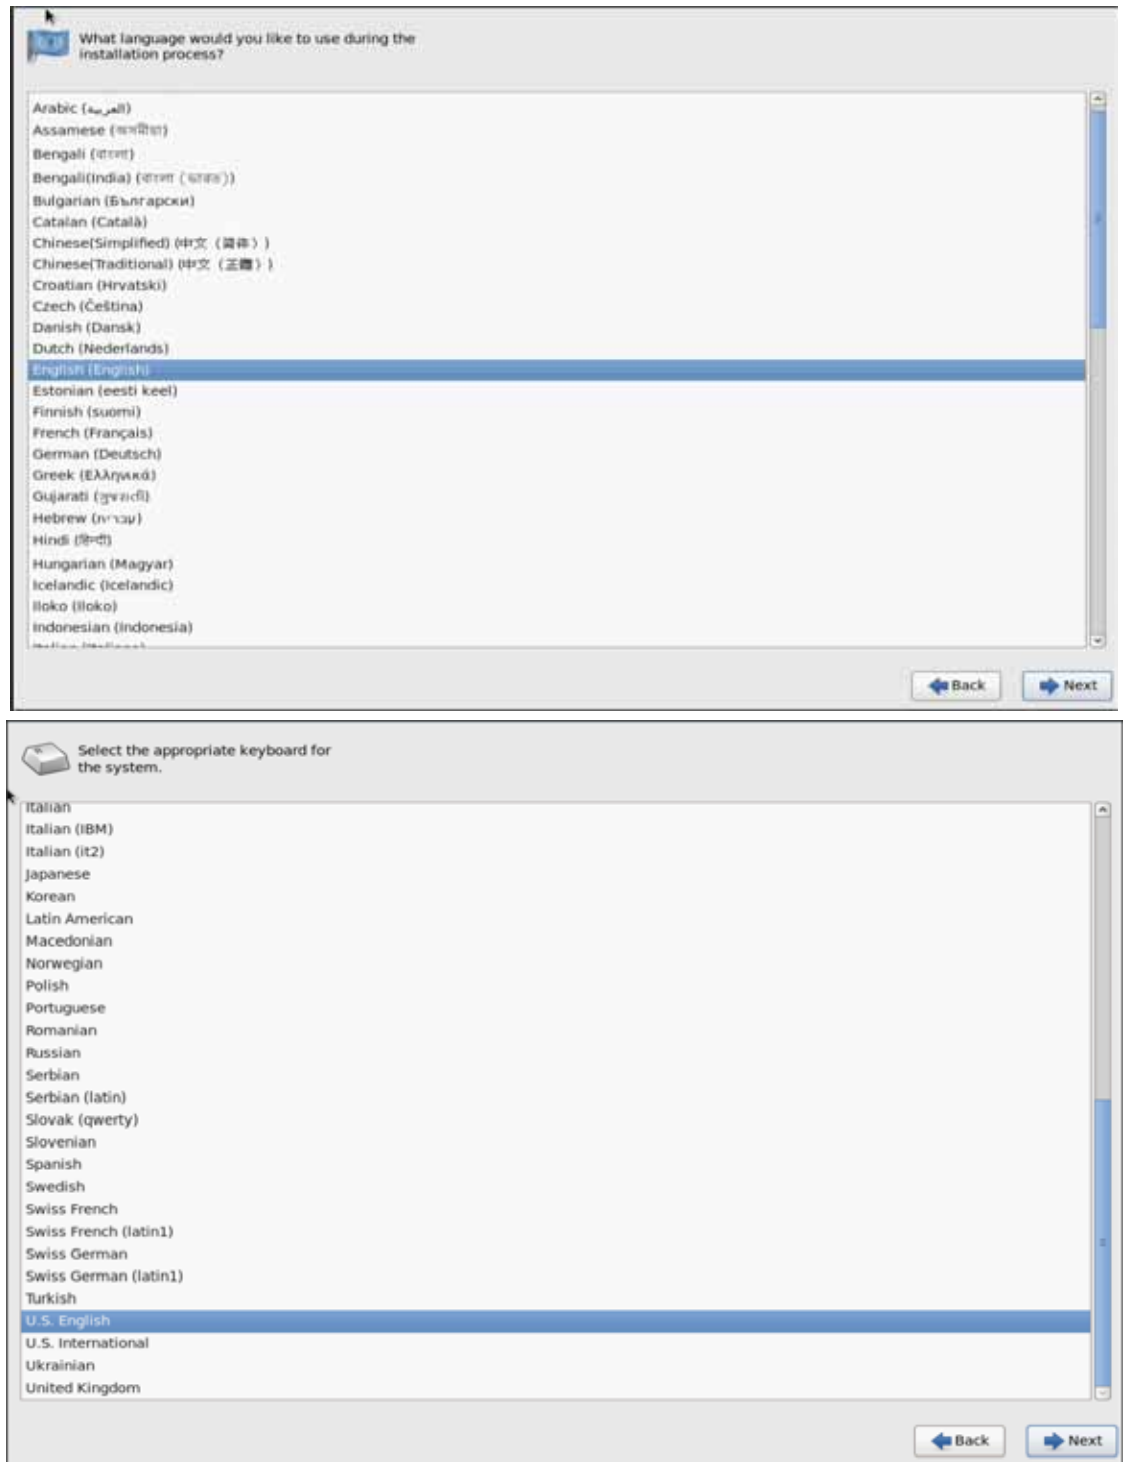
13. Click Next

Figure 137 Red Hat Enterprise Linux Server 6.5 Install Media



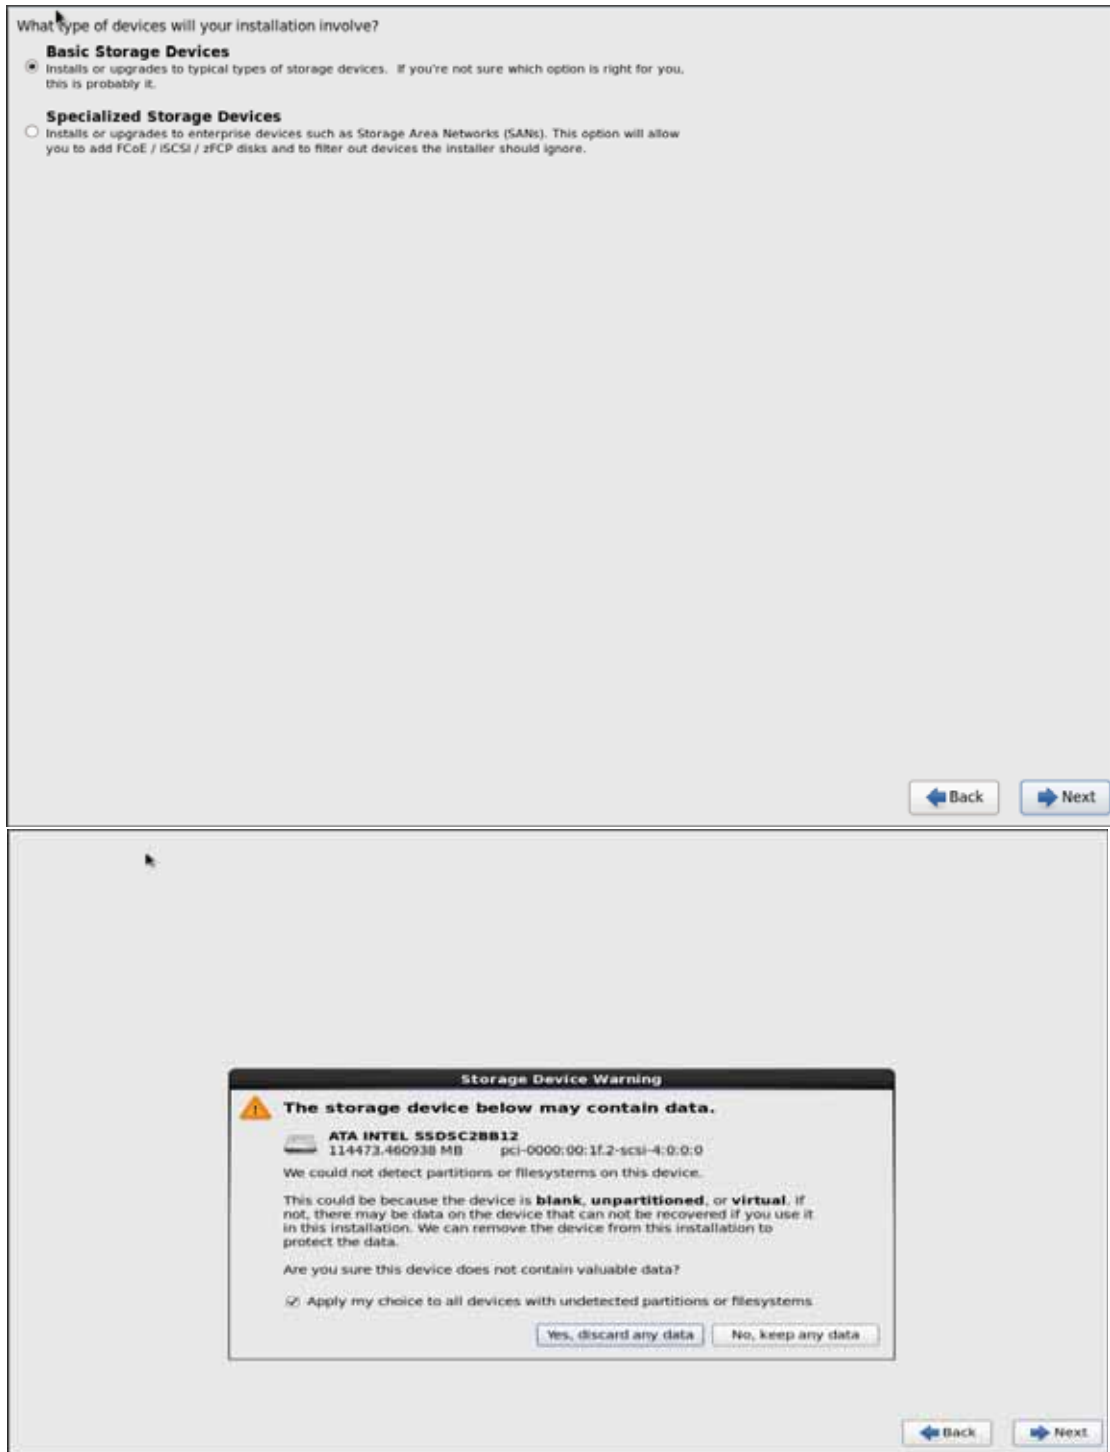
14. Select language of installation, and then Click Next

Figure 138 RHEL Installation: Language and Keyboard Selection



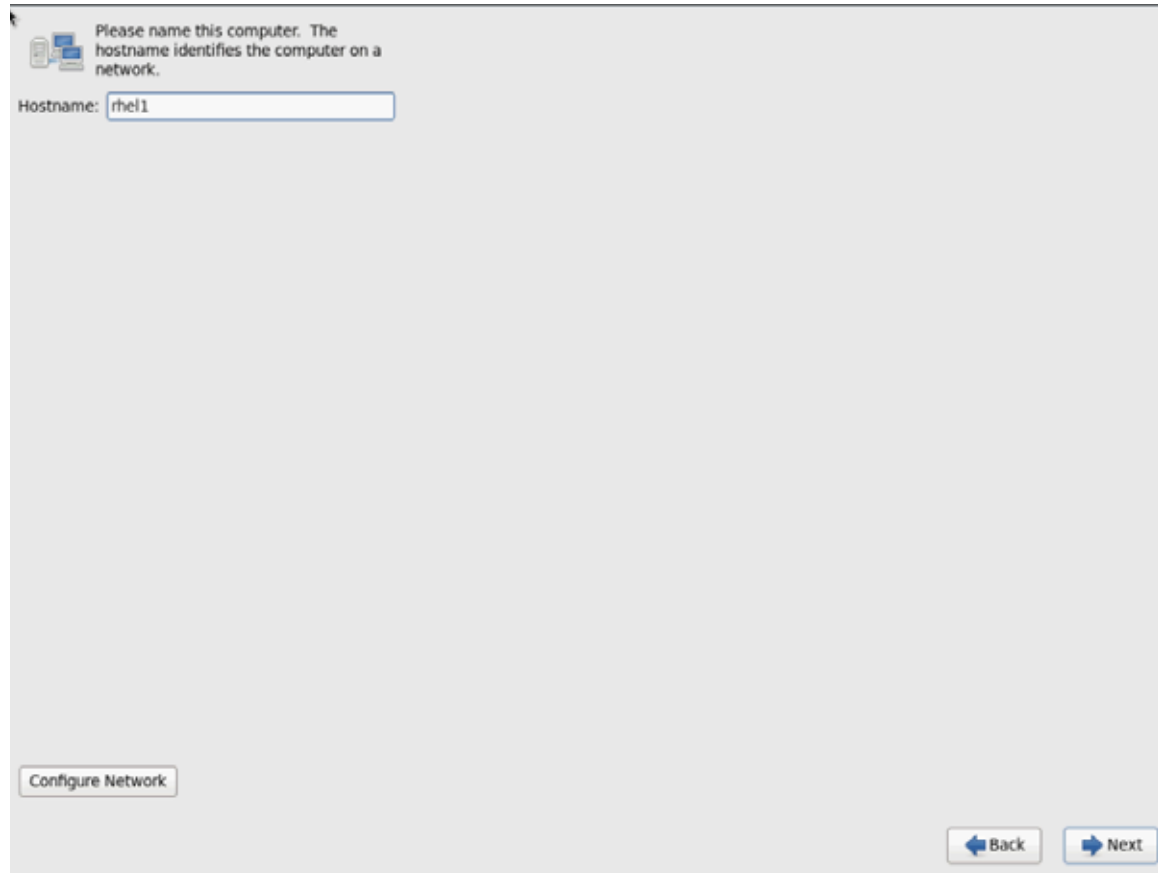
15. Select Basic Storage Devices and Click Next.

Figure 139 RHEL Installation: Storage Devices Selection



16. Provide hostname and configure Network for the host.

Figure 140 *RHEL Installation: Specify Hostname*



17. Select System eth0 and click **Edit**.
18. In the “Editing System eth0” window select the **IPv4 Settings** tab, from the Method drop-down list choose **Manual** and click **Add** to assign the IP address.

Figure 141 RHEL Installation: IPv4 Settings for eth0



Figure 142 RHEL Installation: Location Selection

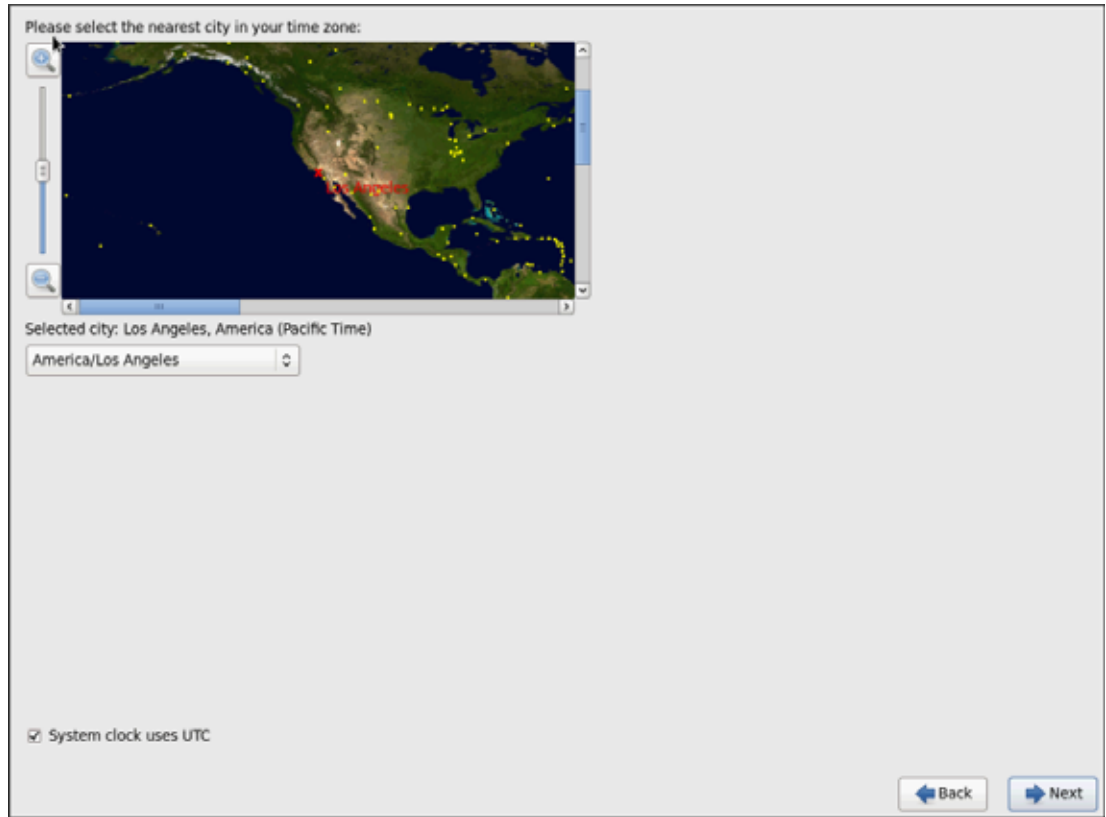



Figure 143 RHEL Installation: Enter Root Credentials



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

Figure 144 RHEL Installation: Custom Layout Creation

Which type of installation would you like?

- Use All Space**
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Replace Existing Linux System(s)**
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Shrink Current System**
Shrinks existing partitions to create free space for the default layout.
- Use Free Space**
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.
- Create Custom Layout**
Manually create your own custom layout on the selected device(s) using our partitioning tool.

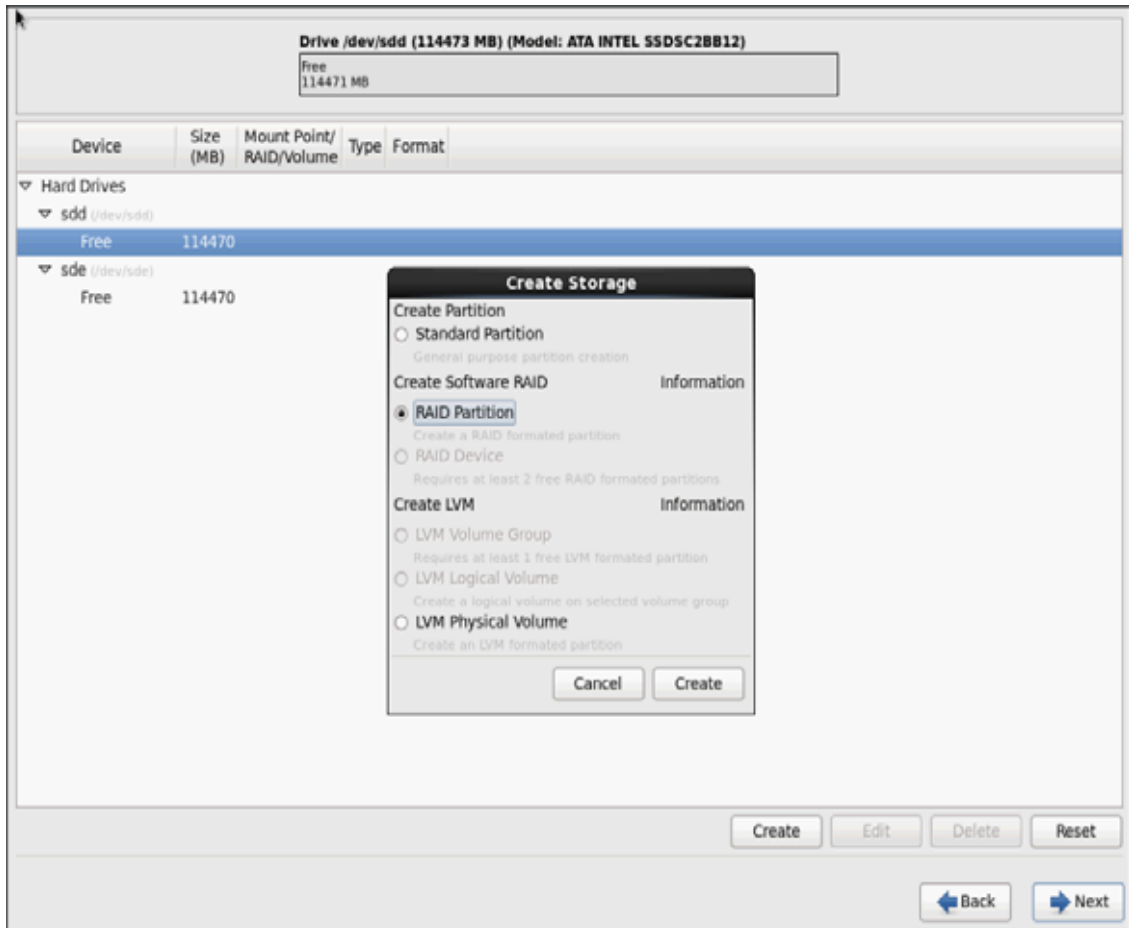
Encrypt system
 Review and modify partitioning layout

Back Next

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

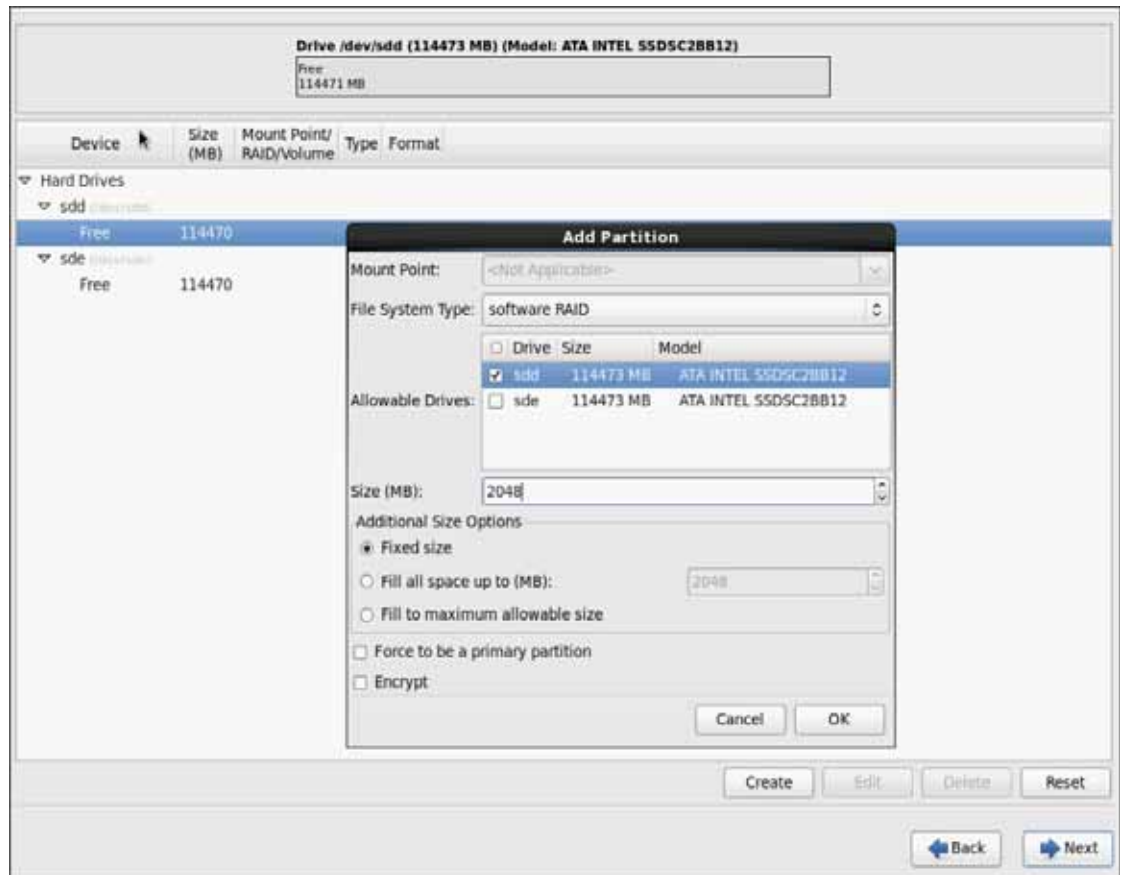
- a. Choose free volume and click on **Create** and choose **RAID Partition**.

Figure 145 RHEL Installation: Create RAID Partition



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

Figure 146 RHEL Installation: Add RAID Partition



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

Figure 147 RHEL Installation: Create RAID Partition

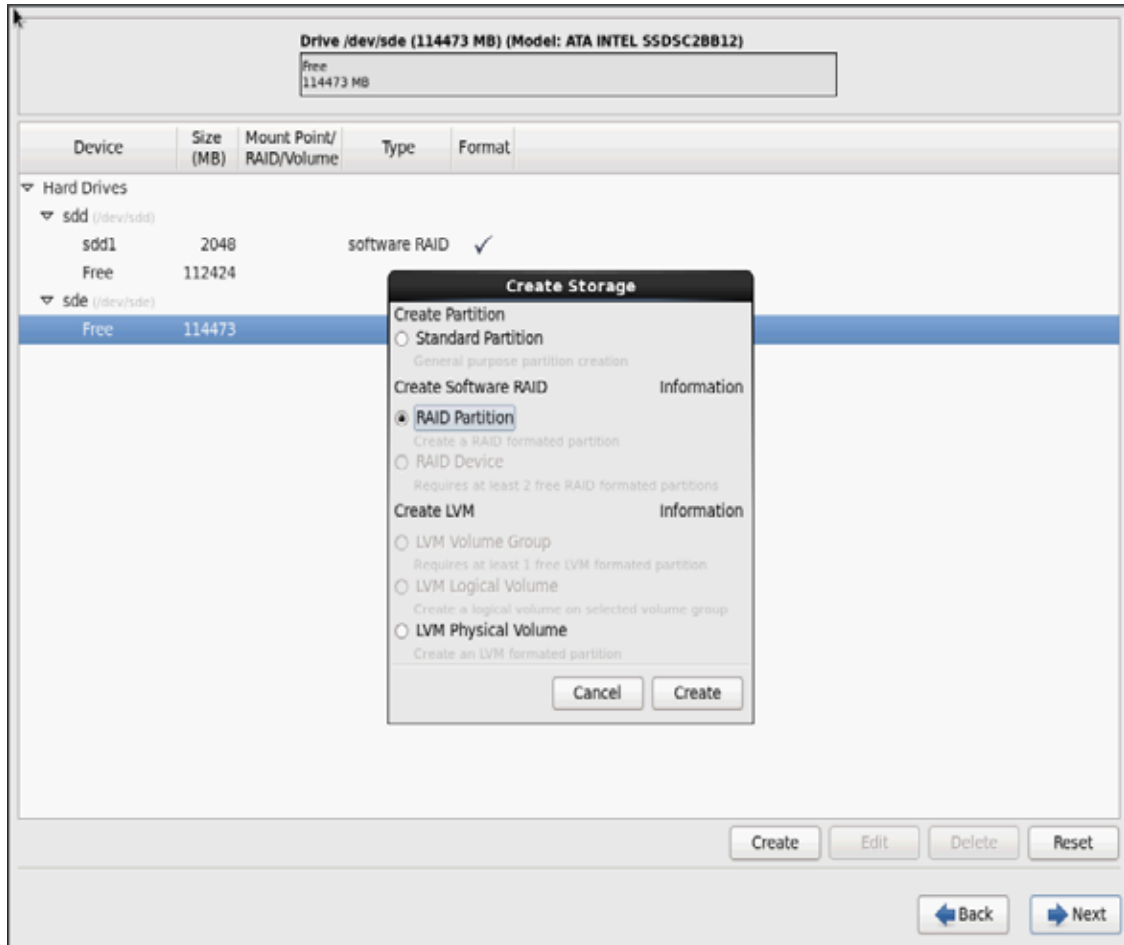
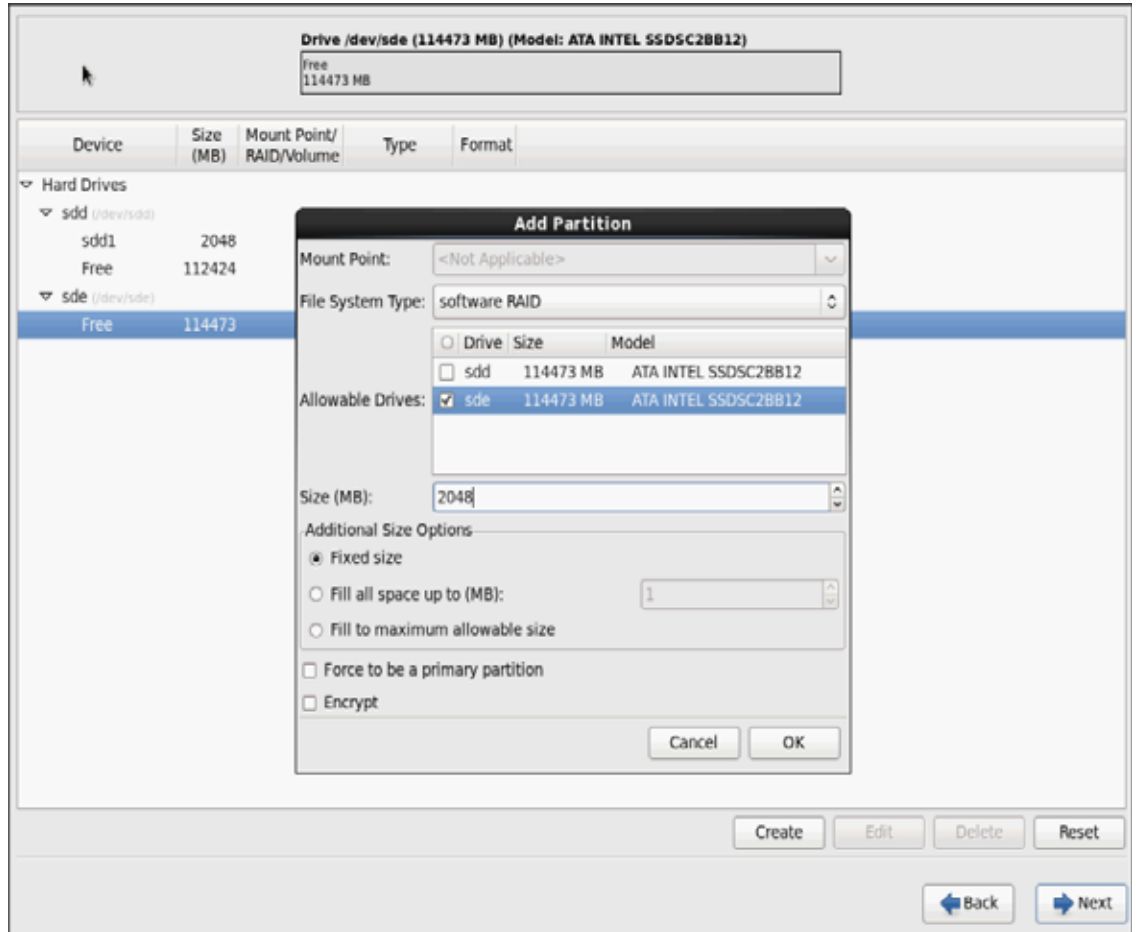


Figure 148 RHEL Installation: Add RAID Partition



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

Figure 149 RHEL Installation: Create RAID Partition

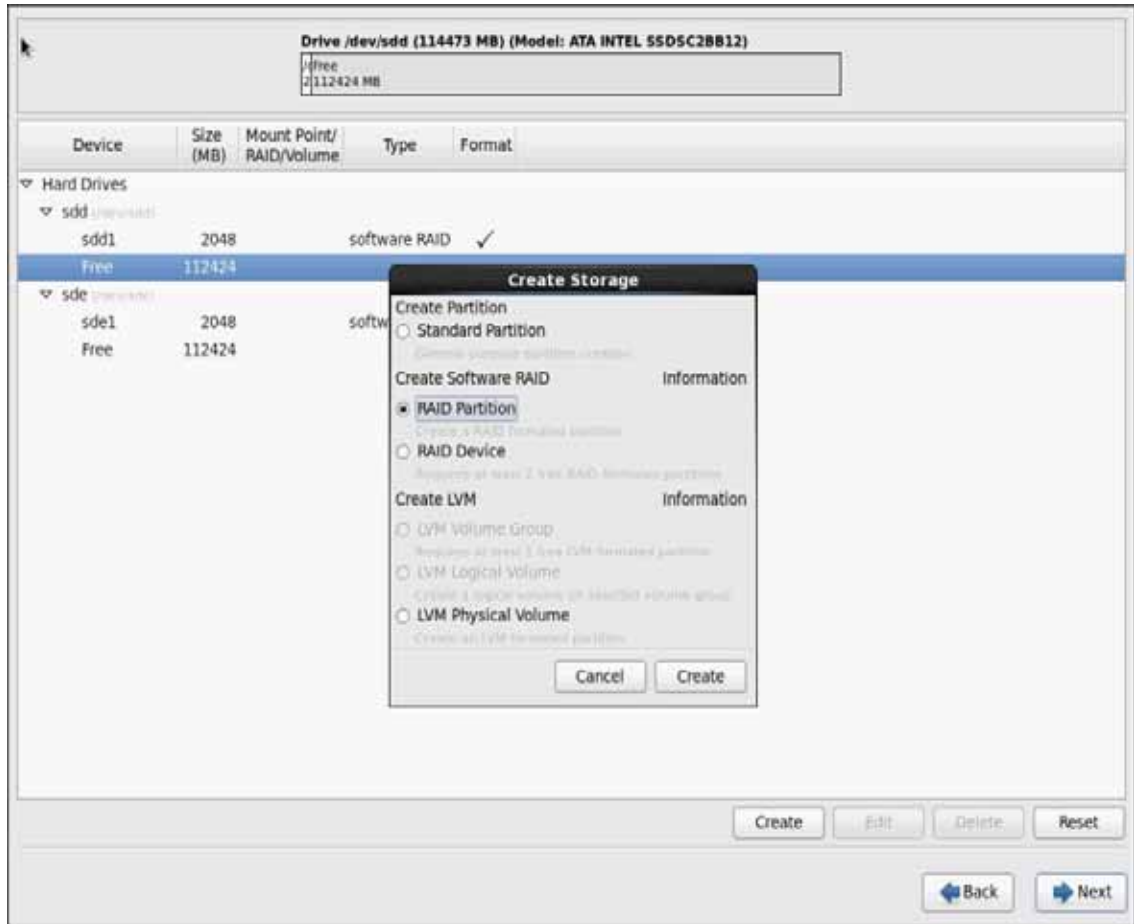


Figure 150 RHEL Installation: Add RAID Partition

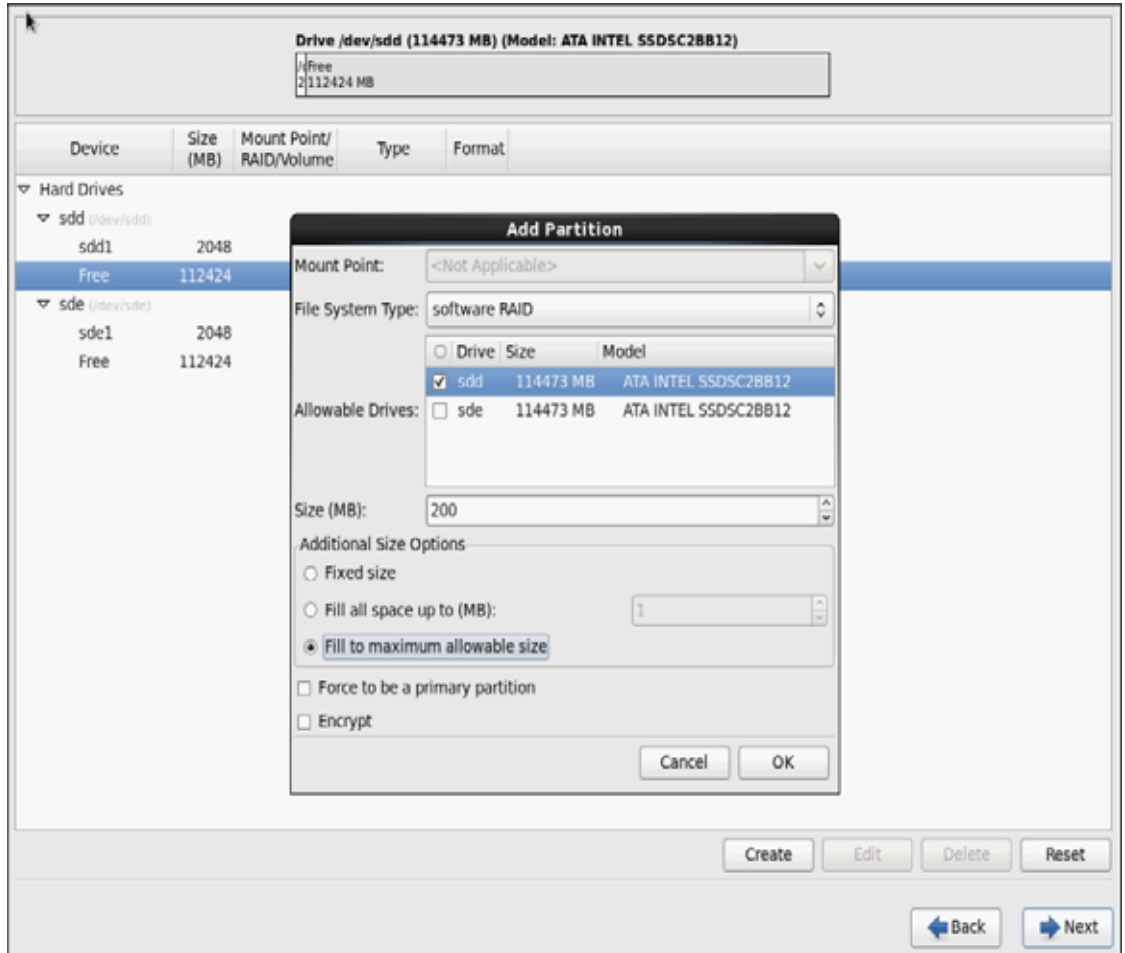


Figure 151 RHEL Installation: Create RAID Partition

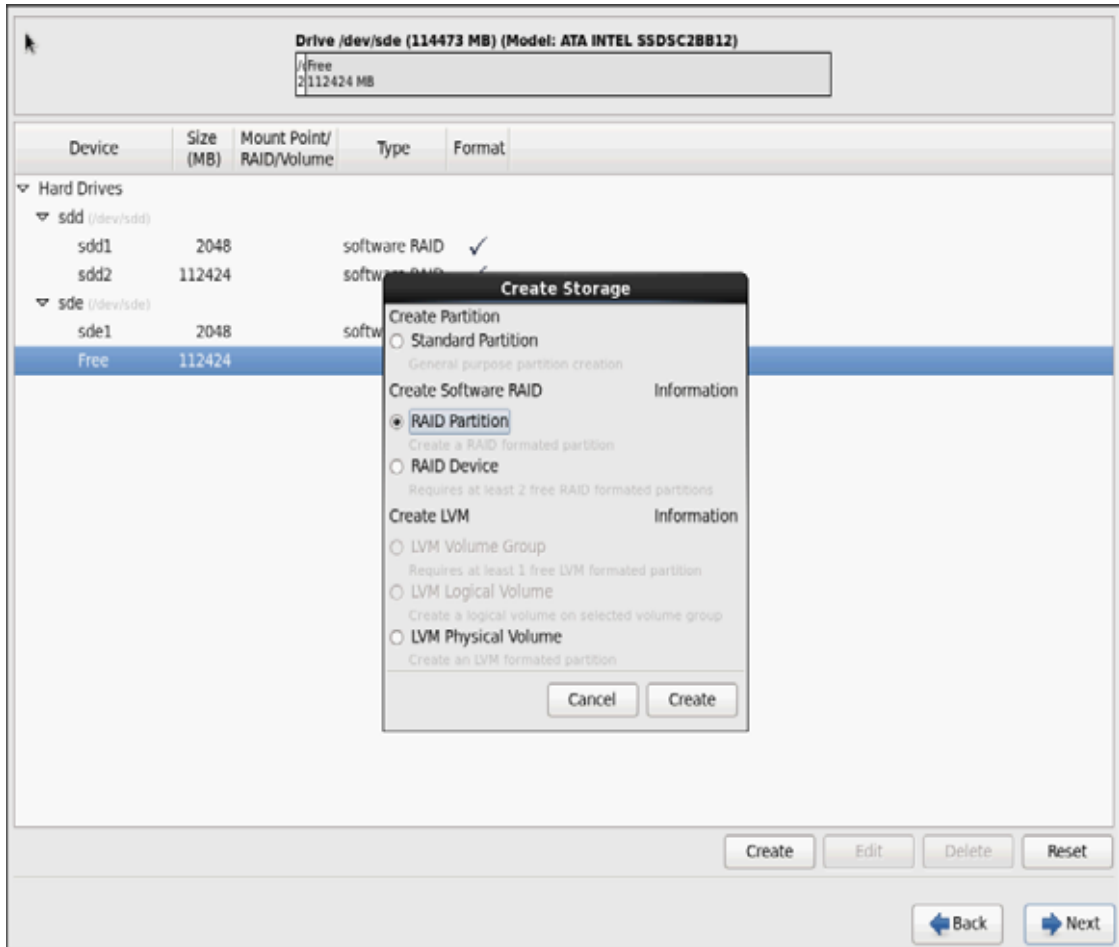
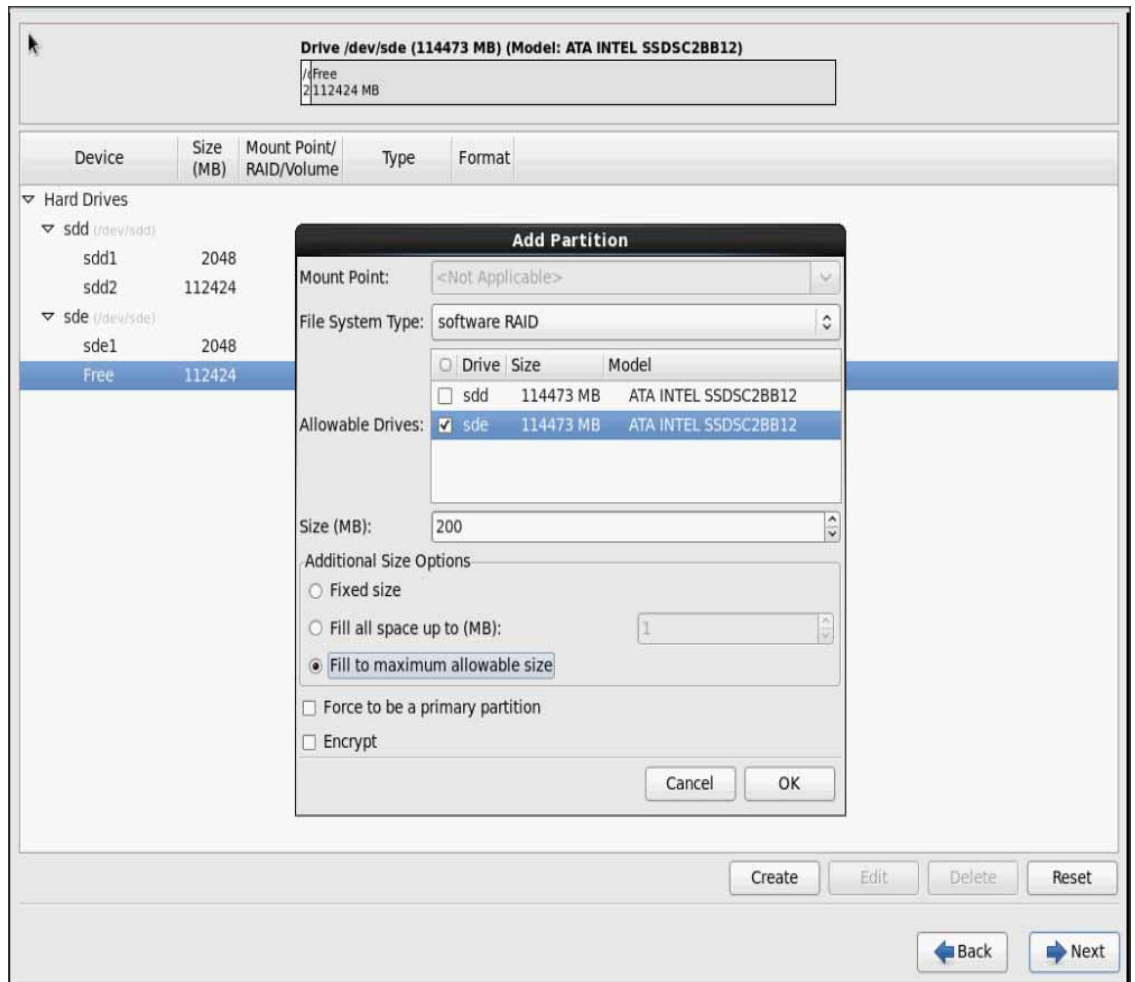


Figure 152 RHEL Installation: Add RAID Partition



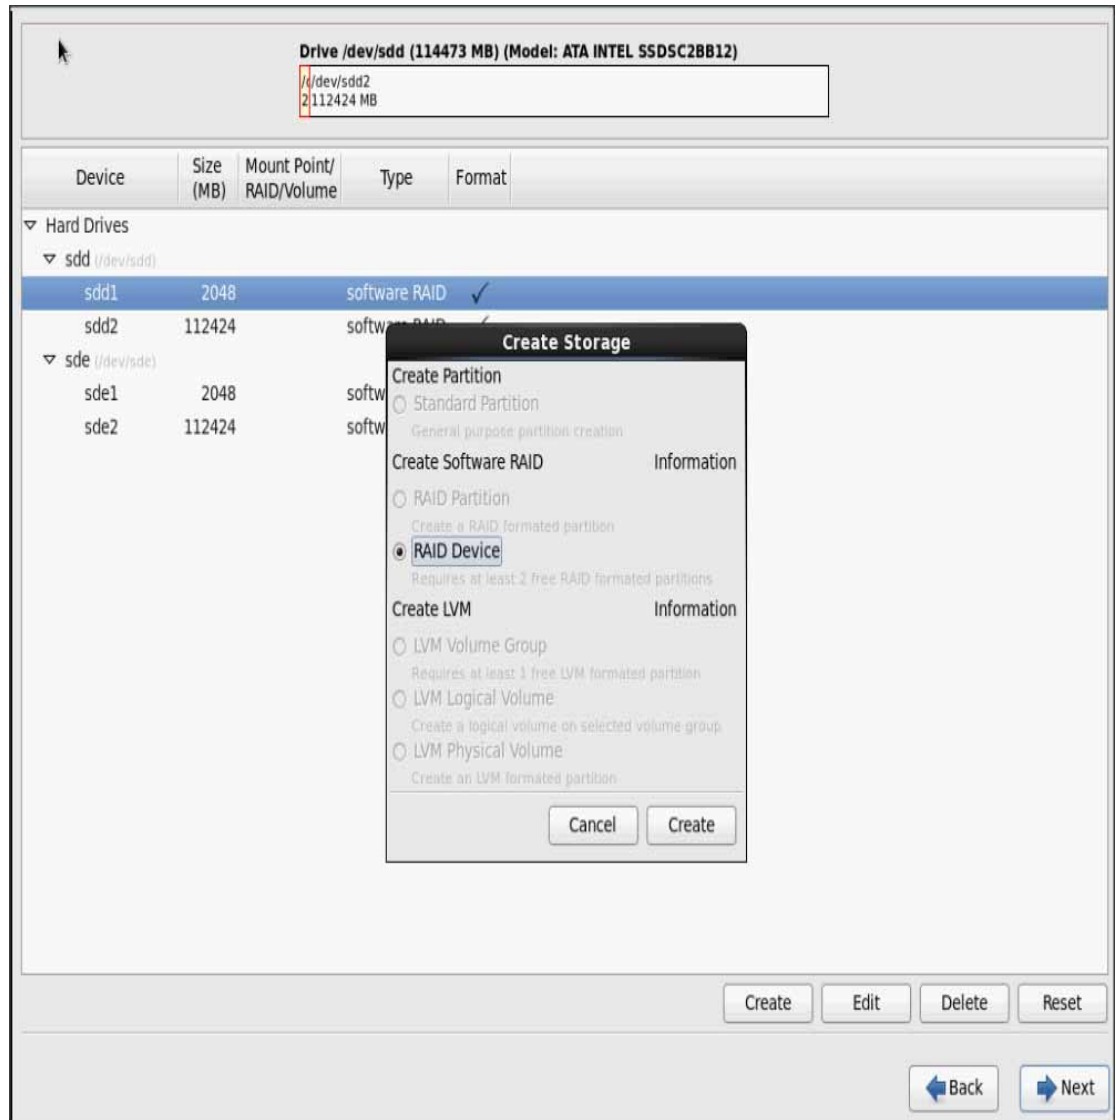
23. The above steps created 2 boot and 2 root (/) partitions. Following steps will RAID1 Devices

Figure 153 RHEL Installation: Selected RAID Devices



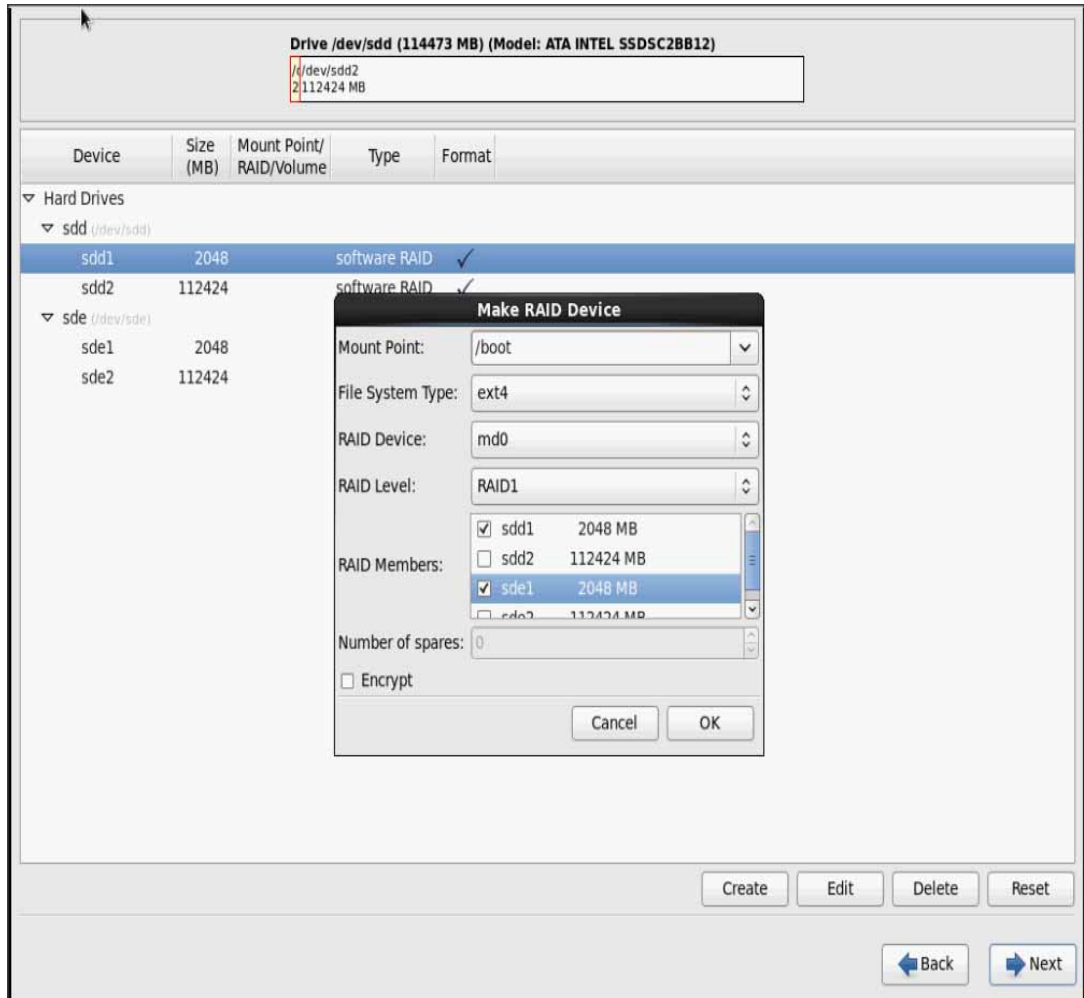
24. Choose one of the boot partitions and click on **Create > RAID Device**.

Figure 154 RHEL Installation: Select RAID Device



25. 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 155 RHEL Installation: Make RAID Device



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

Figure 156 RHEL Installation: Select RAID Device

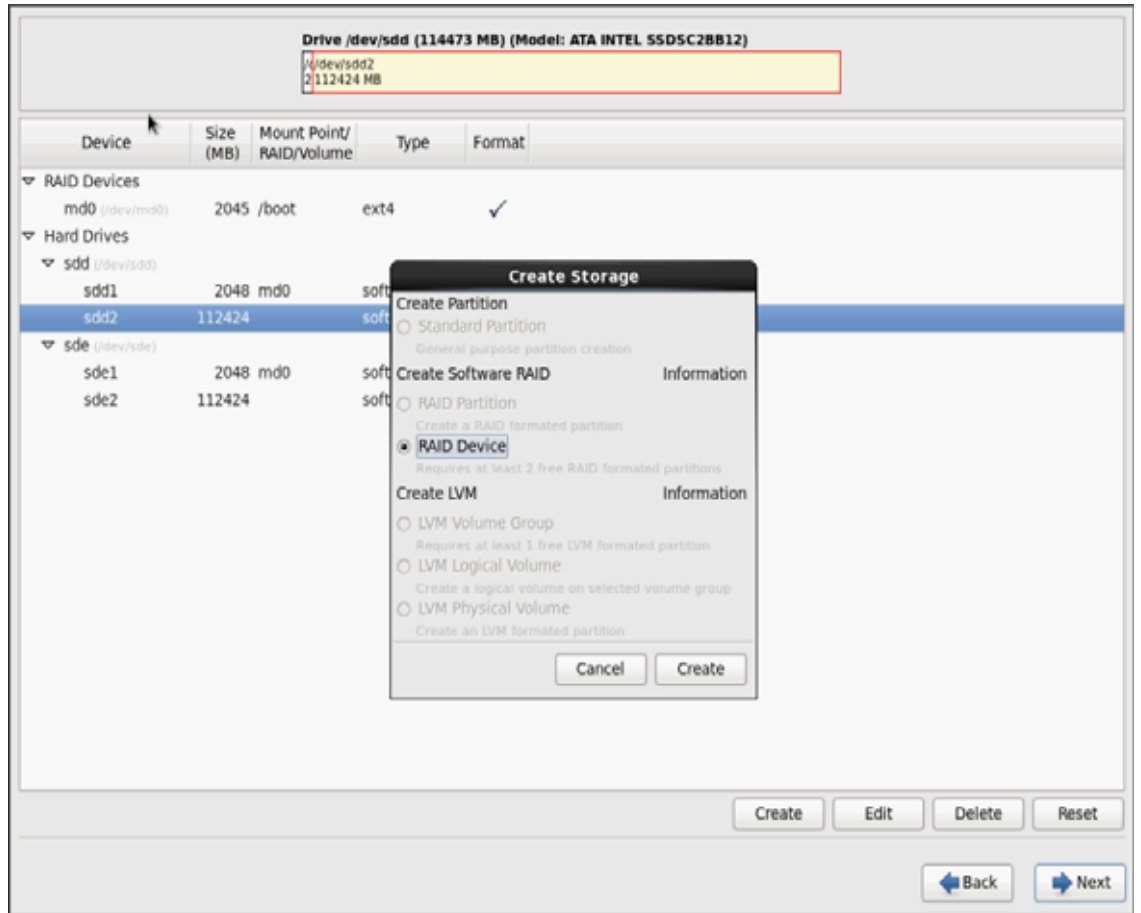


Figure 157 RHEL Installation: Make RAID Device

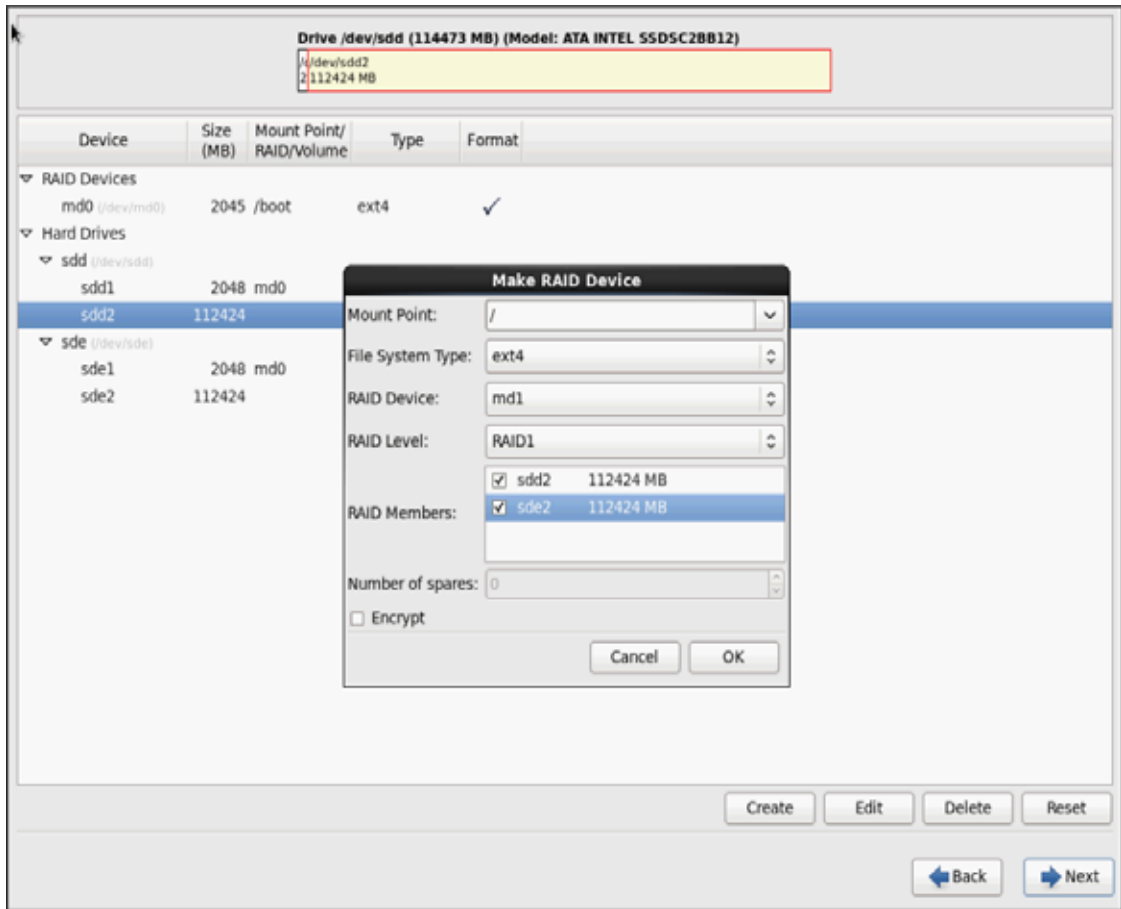
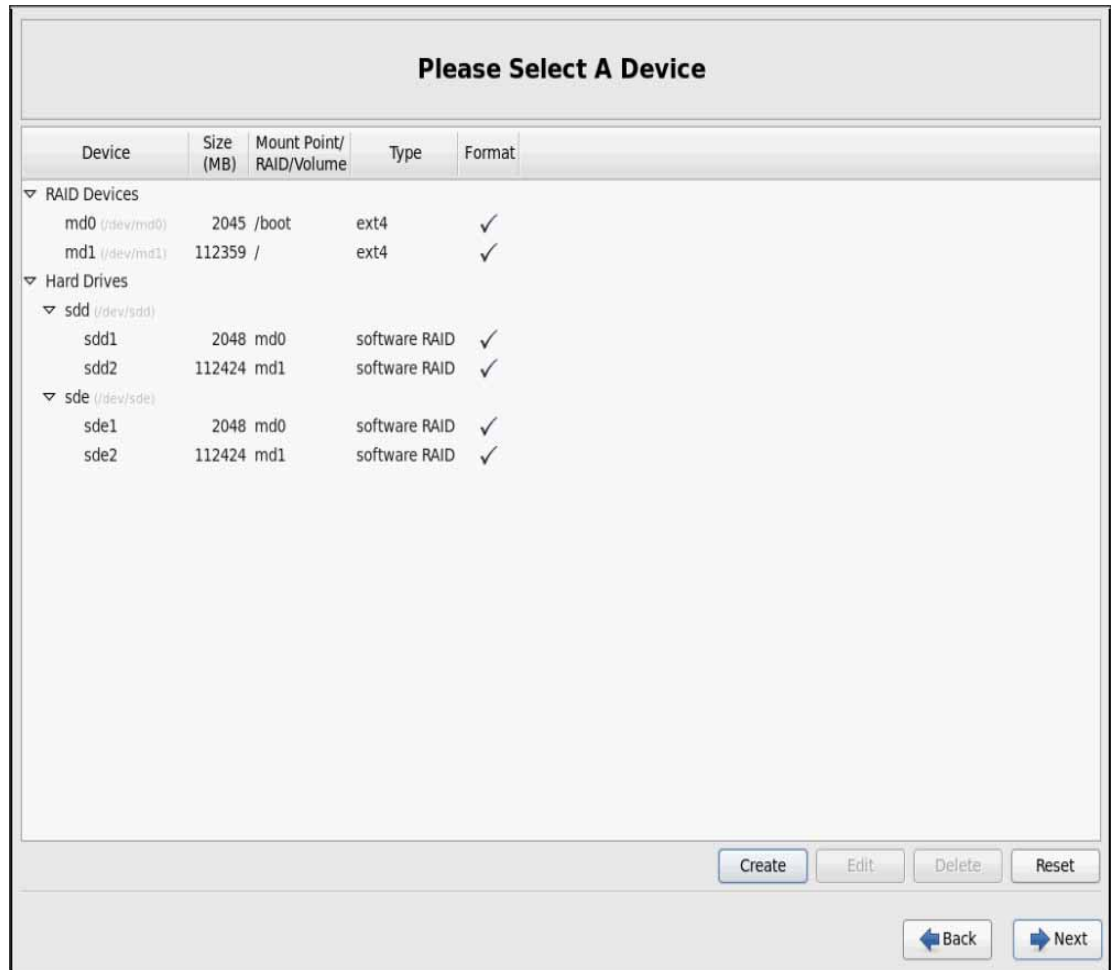
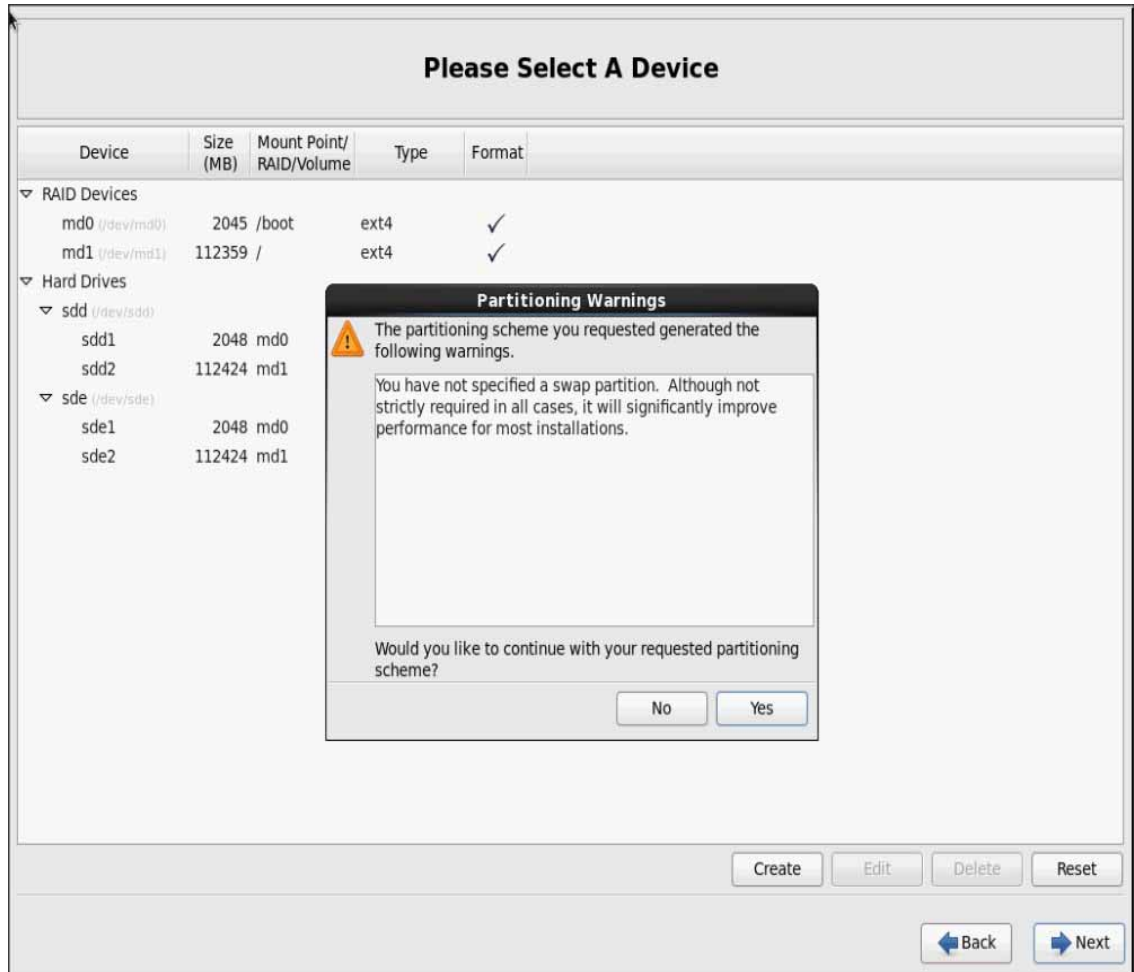


Figure 158 RHEL Installation: All the Selected Devices



27. Click on Next.

Figure 159 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**).

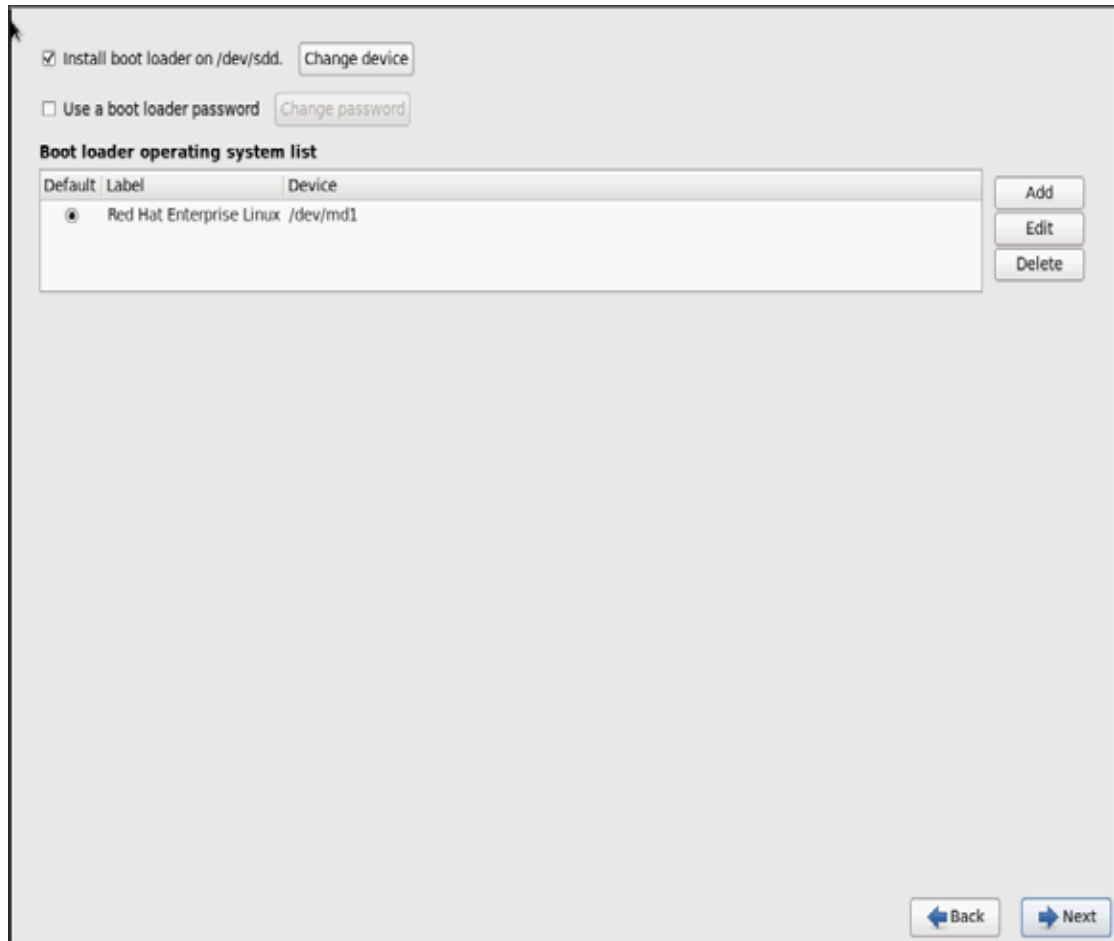
28. Click **Next**, and then click **Format**.

Figure 160 RHEL Installation: Destroy Old Devices



29. Select default settings and click **Next**.

Figure 161 RHEL Installation: Installing Boot Loader



30. Continue with RHEL Installation as shown below.

Figure 162 *RHEL Installation: Keep the Default Installation Option*

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

Basic Server
 Database Server
 Web Server
 Identity Management Server
 Virtualization Host
 Desktop
 Software Development Workstation
 Minimal

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

High Availability
 Load Balancer
 Red Hat Enterprise Linux
 Resilient Storage

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

Customize later Customize now

31. Once the installation is complete reboot the system.

Repeat the steps 1 through 40 to install Red Hat Linux 6.5 on Servers 2 through 160.



Note

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

Table 10 *Host Names and IP Addresses*

Servers	ETH 0	ETH 1	ETH 2
rhel 1	10.0.145.45	10.0.146.45	10.0.147.45
rhel 2	10.0.145.46	10.0.146.46	10.0.147.46
rhel 3	10.0.145.47	10.0.146.47	10.0.147.47

Table 10 *Host Names and IP Addresses*

rhel 4	10.0.145.48	10.0.146.48	10.0.147.48
rhel 5	10.0.145.49	10.0.146.49	10.0.147.49
rhel 6	10.0.145.50	10.0.146.50	10.0.147.50
rhel 7	10.0.145.51	10.0.146.51	10.0.147.51
rhel 8	10.0.145.52	10.0.146.52	10.0.147.52
rhel 9	10.0.145.53	10.0.146.53	10.0.147.53
rhel 10	10.0.145.54	10.0.146.54	10.0.147.54
rhel 11	10.0.145.55	10.0.146.55	10.0.147.55
rhel 12	10.0.145.56	10.0.146.56	10.0.147.56
rhel 13	10.0.145.57	10.0.146.57	10.0.147.57
rhel 14	10.0.145.58	10.0.146.58	10.0.147.58
rhel 15	10.0.145.59	10.0.146.59	10.0.147.59
rhel 16	10.0.145.60	10.0.146.60	10.0.147.60
...
rhel 160	10.0.145.204	10.0.146.204	10.0.147.204



Note

On Cloudera Security: As mentioned above in the “Configuring VLANs” Section, when deploying Cloudera with Security only one VLAN on one vNIC is supported. Please refer to the note for more details

Post OS Install Configuration

Choose one of the nodes of the cluster as Admin Node for management such as CDH installation, parallel shell, creating a local Red Hat repo and others. This CVD uses rhel1 for this purpose.



Note

rhel1 is admin node for the entire Hadoop cluster spawning across two different FI domains

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 10.0.145.45
```

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
```

```
10.0.145.45 rhel1.mgmt
10.0.145.46 rhel2.mgmt
10.0.145.47 rhel3.mgmt
10.0.145.48 rhel4.mgmt
10.0.145.49 rhel5.mgmt
10.0.145.50 rhel6.mgmt
10.0.145.51 rhel7.mgmt
10.0.145.52 rhel8.mgmt
10.0.145.53 rhel9.mgmt
10.0.145.54 rhel10.mgmt
10.0.145.55 rhel11.mgmt
```

```

10.0.145.56 rhel12.mgmt
10.0.145.57 rhel13.mgmt
10.0.145.58 rhel14.mgmt
10.0.145.59 rhel15.mgmt
. . .
10.0.145.204 rhel160.mgmt

10.0.146.45 rhel1
10.0.146.46 rhel2
10.0.146.47 rhel3
10.0.146.48 rhel4
10.0.146.49 rhel5
10.0.146.50 rhel6
10.0.146.51 rhel7
10.0.146.52 rhel8
10.0.146.53 rhel9
10.0.146.54 rhel10
10.0.146.55 rhel11
10.0.146.56 rhel12
10.0.146.57 rhel13
10.0.146.58 rhel14
10.0.146.59 rhel15
. . .
10.0.146.204 rhel160

```

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[1-160]
```

```
[root@rhel1~]# vim /etc/clustershell/groups
[root@rhel1~]# cat /etc/clustershell/groups
all:rhel[1-160].mgmt
```



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@admin ~]# ssh rhel2
The authenticity of host 'rhel2 (10.0.146.46)' can't be established.
RSA keyfingerprint is f2:0c:db:50:64:f1:ae:a6:ff:88:4a:a3:8d:9a:ee:38.
Are you sure you want to continue connecting (yes/no) ?

[root@admin ~]# ssh rhel2.mgmt
The authenticity of host 'rhel2 (10.0.145.46)' can't be established.
RSA keyfingerprint is f2:0c:db:50:64:f1:ae:a6:ff:88:4a:a3:8d:9a:ee:38.
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 rhell 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 rhell. 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 rhell.

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

5. Now on rhell 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 rhell

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



Note

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

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

- To make use of repository files on rhell 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 (rhell) 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 (rhell). Use it to regenerate the repository database(s) for the local copy of the RHEL DVD contents.

```
yum -y install createrepo
```

```
[root@rhell ~]# 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@rhell 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
```

- 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 10.0.146.45
```



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:          10.0.146.45
Address:         10.0.146.45#53
```

```
Name:   rhel1
Address: 10.0.146.45
```

```
[root@rhel2 ~]# nslookup rhel1
Server:          10.0.146.45
Address:         10.0.146.45#53
```

```
45.146.0.10.in-addr.arpa    name = rhel1.
[root@rhel2 ~]# nslookup rhel1.mgmt
Server:          10.0.146.45
Address:         10.0.146.45#53
```

```
Name:   rhel1.mgmt
Address: 10.0.145.45
```

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 10.0.145.45:80
```

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@rhel1 ~]# 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&reind=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 rhel1 (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 rhel1.

```
[root@rhel1 ~]# clush -a -b -c kmod-enic-2.1.1.66-rhel6u5.el6.x86_64.rpm
[root@rhel1 ~]# 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[1-160] 160
-----
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!
```

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 data nodes (except rhel1) by executing the following command in the admin node (rhel1)

```
for SERVER in {46..204}; do scp /root/ntp.conf
10.0.145.$SERVER:/etc/ntp.conf; done
```

```
[root@rhel1 ~]# for SERVER in {46..160}; do scp /root/ntp.conf 10.0.145.$SERVER:/etc/ntp.conf; done
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
ntp.conf      100#      136      0.1KB/s      0:00
....
...
..
.
ntp.conf      100#      136      0.1KB/s      0:00
```



Note

To run the above in clush use `-w` option: `clush -w rhel[2-160].mgmt -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@rhell ~]# 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 rhell and add the following lines:

```
vim /etc/security/limits.conf
root soft nofile 64000
root hard nofile 64000
```

```
[root@rhell ~]# 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 (rhell) to all the nodes using the following command.

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

```
[root@rhell ~]# 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@rhell ~]# 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.

VM.Swapping

Lowering `vm.swappiness` reduces anonymous paging and minimizes OOM killer invocations. Run the following on all nodes. Variable `vm.swappiness` defines how often swap should be used. 0 is No Swapping, 60 default. With `vm.swappiness` set to 1, the kernel will try to reclaim from the page cache instead of application (anonymous) pages.

```
clush -a -b " echo 'vm.swappiness=1' >> /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"
```

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 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"
```

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 rhell1:/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 -cfgeachdiskraid0 WB RA direct NoCachedBadBBU
strpsz1024 -a0

```

WB: Write back

RA: Read Ahead

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

Strpsz1024: Strip Size of 1024K



Note

The 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 the Filesystem for NameNodes and DataNodes

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
#disks_count=`lsblk -id | grep sd | wc -l`
#if [ $disks_count -eq 24 ]; then
# echo "Found 24 disks"
#else
# echo "Found $disks_count disks. Expecting 24. Exiting.."
# exit 1
#fi
[[ "-x" == "${1}" ]] && set -x && set -v && shift 1
count=1
for X in /sys/class/scsi_host/host*/scan
do
echo '- - -' > ${X}
done
for X in /dev/sd?
do
echo $X
if [[ -b ${X} && `sbin/parted -s ${X} print quit|/bin/grep -c boot` -ne 0
]]
then
echo "$X bootable - skipping."
continue
else
Y=${X##*/}1
echo "Setting up Drive => ${X}"
/sbin/parted -s ${X} mklabel gpt quit
/sbin/parted -s ${X} mkpart 1 6144s 100% quit
/sbin/mkfs.xfs -f -q -l size=65536b,lazy-count=1,su=256k -d sunit=1024,swidth=6144 -r
extsize=256k -L ${Y} ${X}1
(( $? )) && continue
#Identify UUID
UUID=`blkid ${X}1 | cut -d " " -f3 | cut -d "=" -f2 | sed 's//g'`
/bin/mkdir -p /data/disk${count}
(( $? )) && continue
echo "UUID of ${X}1 = ${UUID}, mounting ${X}1 as UUID on /data/disk${count}"
/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 ${NC}"
clush -a -B "`which ifconfig` | egrep '(^e|^p)' | awk '{print \$1}' | \ xargs -l
`which ethtool` | grep -e ^Settings -e Speed"
echo ""
clush -a -B "`which lspci` | grep -i ether"
echo ""
echo ""
# probe for disk info #
echo -e "${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 CDH to identify any discrepancies in Post OS Configuration between the servers or during troubleshooting of any cluster / Hadoop issues.

```
./cluster_verification.sh
```

Installing Cloudera

Cloudera's Distribution including Apache Hadoop (CDH) is an enterprise grade, hardened Hadoop distribution. CDH offers Apache Hadoop and several related projects into a single tested and certified product. It offers the latest innovations from the open source community with the testing and quality you expect from enterprise quality software.

Pre-Requisites for CDH Installation

This section details the pre-requisites for CDH Installation such as setting up of CDH Repo.

Cloudera Repo

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

```
mkdir -p /tmp/clouderarepo/
1. Download Cloudera Manager Repo
cd /tmp/clouderarepo/
wget http://archive.cloudera.com/cm5/redhat/6/x86_64/cm/cloudera-manager.repo
```

```

[root@redhat clouderarepo]# wget http://archive.cloudera.com/cm5/redhat/6/x86_64
--2014-02-28 13:55:06-- http://archive.cloudera.com/cm5/redhat/6/x86_64/cm/clou
Resolving archive.cloudera.com... 184.73.217.71
Connecting to archive.cloudera.com|184.73.217.71|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 289 [text/plain]
Saving to: âcloudera-manager.repoâ
100%[=====
2014-02-28 13:55:06 (24.4 MB/s) - âcloudera-manager.repoâ

```

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

```
[root@redhat clouderarepo]# reposync --config=./cloudera-manager.repo --repoid=cloudera-manager
cloudera-manager | 951 B 00:00
cloudera-manager/primary | 4.0 kB 00:00
[cloudera-manager: 1 of 7 ] Downloading RPMS/x86_64/cloudera-manager-agent-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm
cloudera-manager-agent-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm | 3.7 MB 00:05
[cloudera-manager: 2 of 7 ] Downloading RPMS/x86_64/cloudera-manager-daemons-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm
cloudera-manager-daemons-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm | 324 MB 01:52
[cloudera-manager: 3 of 7 ] Downloading RPMS/x86_64/cloudera-manager-server-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm
cloudera-manager-server-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm | 7.9 kB 00:00
[cloudera-manager: 4 of 7 ] Downloading RPMS/x86_64/cloudera-manager-server-db-2-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm
cloudera-manager-server-db-2-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm | 9.7 kB 00:00
[cloudera-manager: 5 of 7 ] Downloading RPMS/x86_64/enterprise-debuginfo-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm
enterprise-debuginfo-5.0.0-0.cm5b2.p0.119.el6.x86_64.rpm | 668 kB 00:00
[cloudera-manager: 6 of 7 ] Downloading RPMS/x86_64/jdk-6u31-linux-amd64.rpm
jdk-6u31-linux-amd64.rpm | 68 MB 00:20
[cloudera-manager: 7 of 7 ] Downloading RPMS/x86_64/oracle-j2sdk1.7-1.7.0+update25-1.x86_64.rpm
oracle-j2sdk1.7-1.7.0+update25-1.x86_64.rpm | 91 MB 00:33
```

2. Download Cloudera Manager Installer.

```
cd /tmp/clouderarepo/
wget http://archive.cloudera.com/cm5/installer/latest/cloudera-manager-installer.bin
```

```
[root@redhat clouderarepo]# wget http://archive.cloudera.com/cm5/installer/latest/cloudera-manager-installer.bin
--2014-02-28 14:11:30-- http://archive.cloudera.com/cm5/installer/latest/cloudera-manager-installer.bin
Resolving archive.cloudera.com... 184.73.217.71
Connecting to archive.cloudera.com|184.73.217.71|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 510866 (499K) [application/octet-stream]
Saving to: âcloudera-manager-installer.binâ

100%[=====>] 510,866 565K/s in 0.9s

2014-02-28 14:11:31 (565 KB/s) - âcloudera-manager-installer.binâ
```

3. Copy the repository directory to the admin (rhell) node.

```
Scp -r /tmp/clouderarepo/ rhell:/var/www/html
```

```
[root@srv1 clouderarepo]# scp -r /tmp/clouderarepo/ rhell:/var/www/html/
cloudera-manager-installer.bin
cloudera-manager.repo
cloudera-manager-agent-5.3.2-1.cm532.p0.209.el6.x86_64.rpm
enterprise-debuginfo-5.3.2-1.cm532.p0.209.el6.x86_64.rpm
cloudera-manager-server-db-2-5.3.2-1.cm532.p0.209.el6.x86_64.rpm
cloudera-manager-daemons-5.3.2-1.cm532.p0.209.el6.x86_64.rpm
cloudera-manager-daemons-5.3.2-1.cm532.p0.209 100% 476MB 68.0MB/s 00:07
oracle-j2sdk1.7-1.7.0+update67-1.x86_64.rpm 100% 135MB 67.7MB/s 00:02
cloudera-manager-server-5.3.2-1.cm532.p0.209. 100% 7852 7.7KB/s 00:00
jdk-6u31-linux-amd64.rpm 100% 68MB 67.9MB/s 00:01
```

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

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

```
createrepo --baseurl http://10.0.145.45/clouderarepo/cloudera-manager/
/var/www/html/clouderarepo/cloudera-manager
```

```
[root@rhell clouderarepo]#createrepo --baseurl http://10.0.145.45/clouderarepo/cloudera-manager/
/var/www/html/clouderarepo/cloudera-manager
Spawning worker 0 with 7 pkgs
Worker Finished
Gathering worker results

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



Note Visit <http://10.0.145.45/clouderarepo> to verify the files.

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

```
vim /var/www/html/clouderarepo/cloudera-manager/cloudera-manager.repo
[cloudera-manager]
name=Cloudera Manager
baseurl=http://10.0.145.45/clouderarepo/cloudera-manager/
gpgcheck = 0
```

Copy the file `cloudera-manager.repo` into `/etc/yum.repos.d/` on the admin node to enable it to find the packages that are locally hosted.

```
cp /var/www/html/clouderarepo/cloudera-manager/cloudera-manager.repo /etc/yum.repos.d/
```

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/cloudera-manager.repo
```

Setup the Local Parcels for CDH 5.3.2

From a host that's connected the internet, download the appropriate CDH 5.3.2 parcels that are meant for RHEL6.5 from the URL: <http://archive.cloudera.com/cdh5/parcels/> and place them in the directory `"/var/www/html/CDH5.3parcels"` of the Admin node.

The following screenshot shows the files relevant for RHEL6.5. They are,

- `CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel`
- `CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1`
- `manifest.json`.

Index of /cdh5/parcels/5.3.2

Name	Last modified	Size	Description
Parent Directory	-	-	-
CDH-5.3.2-1.cdh5.3.2.p0.10-el5.parcel	2015-02-24 23:54	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-el5.parcel.sha1	2015-02-24 23:54	41	
CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel	2015-02-24 23:55	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1	2015-02-24 23:55	41	
CDH-5.3.2-1.cdh5.3.2.p0.10-precise.parcel	2015-02-24 23:54	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-precise.parcel.sha1	2015-02-24 23:54	41	
CDH-5.3.2-1.cdh5.3.2.p0.10-sles11.parcel	2015-02-24 23:55	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-sles11.parcel.sha1	2015-02-24 23:55	41	
CDH-5.3.2-1.cdh5.3.2.p0.10-trusty.parcel	2015-02-24 23:54	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-trusty.parcel.sha1	2015-02-24 23:54	41	
CDH-5.3.2-1.cdh5.3.2.p0.10-wheezy.parcel	2015-02-24 23:55	1.5G	
CDH-5.3.2-1.cdh5.3.2.p0.10-wheezy.parcel.sha1	2015-02-24 23:55	41	
manifest.json	2015-02-24 23:55	42K	

Apache/2.4.7 (Ubuntu) Server at archive-primary.cloudera.com Port 80

1. `mkdir -p /tmp/clouderarepo/CDH5.3parcels`
2. `cd /tmp/clouderarepo/CDH5.3parcels`
3. Download Parcels

```
wget
http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel
```

```
[root@srv1 ~]# cd /tmp/clouderarepo/CDH5.3parcels
[root@srv1 CDH5.3parcels]# wget http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel
--2015-03-28 12:23:20-- http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel
Resolving archive.cloudera.com... 54.230.117.86, 54.239.132.249, 54.230.116.239, ...
Connecting to archive.cloudera.com|54.230.117.86|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1558200266 (1.5G)
Saving to: âCDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcelâ

 9% [====>] 149,019,028 11.2M/s eta 2m 5s
```

```
wget
http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1
wget http://archive.cloudera.com/cdh5/parcels/5.3.2/manifest.json
```

```
[root@Srv1 CDH5.3parcels]# wget http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1
--2015-03-28 12:24:28-- http://archive.cloudera.com/cdh5/parcels/5.3.2/CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1
Resolving archive.cloudera.com... 54.230.118.112, 54.230.118.159, 54.230.119.219, ...
Connecting to archive.cloudera.com[54.230.118.112]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 41 [application/x-sha1]
Saving to: âCDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1â

100%[=====>] 41

2015-03-28 12:24:28 (6.20 MB/s) - âCDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel.sha1â

[root@Srv1 CDH5.3parcels]# wget http://archive.cloudera.com/cdh5/parcels/5.3.2/manifest.json
--2015-03-28 12:25:19-- http://archive.cloudera.com/cdh5/parcels/5.3.2/manifest.json
Resolving archive.cloudera.com... 54.230.118.73, 54.239.132.172, 54.230.118.248, ...
Connecting to archive.cloudera.com[54.230.118.73]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 42655 (42K) [application/json]
Saving to: âmanifest.jsonâ

100%[=====>] 42,655

2015-03-28 12:25:19 (642 KB/s) - âmanifest.jsonâ
```

- Now edit the manifest.json file and remove the scripts that are not meant for RHEL6.5. Below is that script which you can copy and paste.



Note Make sure the script starts and end with initial additional braces.

```
{
  "lastUpdated": 14248220420000,
  "parcels": [
    {
      "parcelName": "CDH-5.3.2-1.cdh5.3.2.p0.10-el6.parcel",
      "components": [
        {
          "pkg_version": "0.7.0+cdh5.3.2+0",
          "pkg_release": "1.cdh5.3.2.p0.17",
          "name": "bigtop-tomcat",
          "version": "6.0.41-cdh5.3.2"
        },
        {
          "pkg_version": "0.11.0+cdh5.3.2+18",
          "pkg_release": "1.cdh5.3.2.p0.17",
          "name": "crunch",
          "version": "0.11.0-cdh5.3.2"
        },
        {
          "pkg_version": "1.5.0+cdh5.3.2+84",
          "pkg_release": "1.cdh5.3.2.p0.17",
          "name": "flume-ng",
          "version": "1.5.0-cdh5.3.2"
        },
        {
          "pkg_version": "2.5.0+cdh5.3.2+813",
          "pkg_release": "1.cdh5.3.2.p0.17",
          "name": "hadoop-0.20-mapreduce",
          "version": "2.5.0-cdh5.3.2"
        },
        {
          "pkg_version": "2.5.0+cdh5.3.2+813",
          "pkg_release": "1.cdh5.3.2.p0.17",
          "name": "hadoop-0.20-mapreduce",
          "version": "2.5.0-cdh5.3.2"
        }
      ]
    }
  ]
}
```

```

    "name": "hadoop",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "2.5.0+cdh5.3.2+813",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hadoop-hdfs",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "2.5.0+cdh5.3.2+813",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hadoop-httpfs",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "2.5.0+cdh5.3.2+813",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hadoop-kms",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "2.5.0+cdh5.3.2+813",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hadoop-mapreduce",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "2.5.0+cdh5.3.2+813",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hadoop-yarn",
    "version": "2.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "0.98.6+cdh5.3.2+83",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hbase",
    "version": "0.98.6-cdh5.3.2"
  },
  {
    "pkg_version": "1.5+cdh5.3.2+25",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hbase-solr",
    "version": "1.5-cdh5.3.2"
  },
  {
    "pkg_version": "0.13.1+cdh5.3.2+330",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hive",
    "version": "0.13.1-cdh5.3.2"
  },
  {
    "pkg_version": "0.13.1+cdh5.3.2+330",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hive-hcatalog",
    "version": "0.13.1-cdh5.3.2"
  },
  {
    "pkg_version": "3.7.0+cdh5.3.2+163",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "hue",
    "version": "3.7.0-cdh5.3.2"
  }
}

```



```

    "pkg_version": "2.1.2+cdh5.3.2+0",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "impala",
    "version": "2.1.2-cdh5.3.2"
  },
  {
    "pkg_version": "0.15.0+cdh5.3.2+193",
    "pkg_release": "1.cdh5.3.2.p0.18",
    "name": "kite",
    "version": "0.15.0-cdh5.3.2"
  },
  {
    "pkg_version": "1.0.0+cdh5.3.2+0",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "llama",
    "version": "1.0.0-cdh5.3.2"
  },
  {
    "pkg_version": "0.9+cdh5.3.2+19",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "mahout",
    "version": "0.9-cdh5.3.2"
  },
  {
    "pkg_version": "4.0.0+cdh5.3.2+339",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "oozie",
    "version": "4.0.0-cdh5.3.2"
  },
  {
    "pkg_version": "1.5.0+cdh5.3.2+62",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "parquet",
    "version": "1.5.0-cdh5.3.2"
  },
  {
    "pkg_version": "0.12.0+cdh5.3.2+51",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "pig",
    "version": "0.12.0-cdh5.3.2"
  },
  {
    "pkg_version": "1.4.0+cdh5.3.2+128",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "sentry",
    "version": "1.4.0-cdh5.3.2"
  },
  {
    "pkg_version": "4.4.0+cdh5.3.2+326",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "solr",
    "version": "4.4.0-cdh5.3.2"
  },
  {
    "pkg_version": "1.2.0+cdh5.3.2+369",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "spark",
    "version": "1.2.0-cdh5.3.2"
  },
  {
    "pkg_version": "1.99.4+cdh5.3.2+21",
    "pkg_release": "1.cdh5.3.2.p0.17",
    "name": "sqoop2",
    "version": "1.99.4-cdh5.3.2"
  }

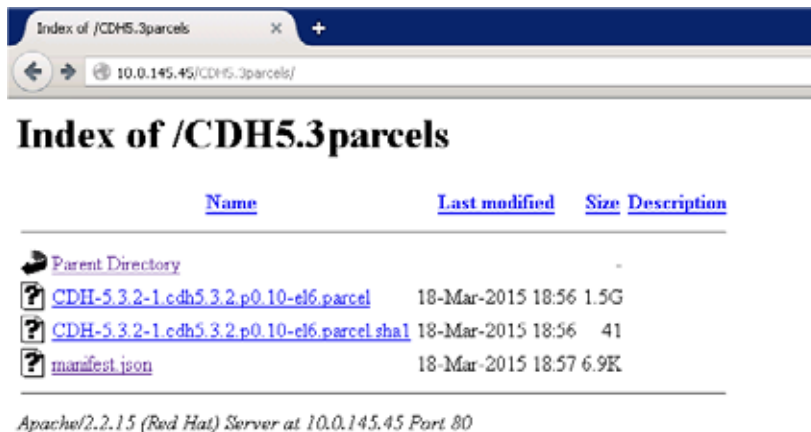
```

```

    },
    {
      "pkg_version": "1.4.5+cdh5.3.2+64",
      "pkg_release": "1.cdh5.3.2.p0.17",
      "name": "sqoop",
      "version": "1.4.5-cdh5.3.2"
    },
    {
      "pkg_version": "0.9.0+cdh5.3.2+13",
      "pkg_release": "1.cdh5.3.2.p0.17",
      "name": "whirr",
      "version": "0.9.0-cdh5.3.2"
    },
    {
      "pkg_version": "3.4.5+cdh5.3.2+83",
      "pkg_release": "1.cdh5.3.2.p0.17",
      "name": "zookeeper",
      "version": "3.4.5-cdh5.3.2"
    }
  ],
  "replaces": "IMPALA, SOLR, SPARK",
  "hash": "a1722a9c033d33ca4ed4558eaf6c10c803b06a16"
}
]
}

```

5. scp /tmp/clouderarepo/CDH5.3parcels to the /var/www/html directory of Admin node (rhel1).
`scp -r /tmp/clouderarepo/CDH5.3parcels/ rhel1:/var/www/html/`
6. Verify that these files are accessible by visiting the URL <http://10.0.145.45/CDH5.3parcels/> in admin node.



Setting MySQL Database for Cloudera manager

To use a MySQL database, follow these procedures:

1. Installing the MySQL Server
2. Configuring and Starting the MySQL Server
3. Installing the MySQL JDBC Driver

4. Creating Databases for Activity Monitor, Reports Manager, Hive Metastore Server, Sentry Server, Cloudera Navigator Audit Server, and Cloudera Navigator Metadata Server

Following steps provide details of the above procedure for setting MySQL database for Cloudera Manager:

1. Installing the mysql server

In the admin node where Cloudera manager will be installed, use the following command to install mysql server.

```
[root@rhell ]# yum -y install mysql-server
```

2. Configuring and starting the MySQL Server

- a. Stop the MySQL server if it is running.

```
[root@rhell ]# service mysqld stop
```

- b. Move old InnoDB log if exists.

Move files /var/lib/mysql/ib_logfile0 and /var/lib/mysql/ib_logfile1 out of /var/lib/mysql/ to a backup location.

```
mv /var/lib/mysql/ib_logfile0 /root/ib_logfile0.bkp
```

```
mv /var/lib/mysql/ib_logfile1 /root/ib_logfile1.bkp
```

- c. Determine the location of the option file, my.cnf and edit/add following lines

```
vim /etc/my.cnf
[mysqld]
transaction-isolation = READ-COMMITTED

# InnoDB settings
innodb_flush_method = O_DIRECT

max_connections = 550
```

```
[root@rhell ~]# vi /etc/my.cnf
[root@rhell ~]# cat /etc/my.cnf
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
user=mysql
transaction-isolation = READ-COMMITTED
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0

[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid

# InnoDB settings
innodb_flush_method = O_DIRECT

max_connections = 550
```



Note max_connections need to be increased based on number of nodes and applications. Please follow the recommendations as mention in Cloudera document:

http://www.cloudera.com/content/cloudera/en/documentation/core/v5-3-x/topics/cm_ig_mysql.html

- d. Ensure MySQL Server starts at boot

```
chkconfig mysqld on
```

```
[root@rhell ~]# chkconfig --list mysqld
mysqld          0:off  1:off  2:on   3:on   4:on   5:on   6:off
```

- e. Start the MySQL Server

```
service mysqld start
```

```
[root@rhell ~]# chkconfig --list mysqld
mysqld          0:off  1:off  2:on   3:on   4:on   5:on   6:off
[root@rhell ~]# service mysqld start
Initializing MySQL database: Installing MySQL system tables...
OK
Filling help tables...
OK

To start mysqld at boot time you have to copy
support-files/mysql.server to the right place for your system

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER !
To do so, start the server, then issue the following commands:

/usr/bin/mysqladmin -u root password 'new-password'
/usr/bin/mysqladmin -u root -h rhell password 'new-password'

Alternatively you can run:
/usr/bin/mysql_secure_installation

which will also give you the option of removing the test
databases and anonymous user created by default. This is
strongly recommended for production servers.

See the manual for more instructions.

You can start the MySQL daemon with:
cd /usr ; /usr/bin/mysqld_safe &

You can test the MySQL daemon with mysql-test-run.pl
cd /usr/mysql-test ; perl mysql-test-run.pl

Please report any problems with the /usr/bin/mysqlbug script!

Starting mysqld: [ OK ]
```

- f. set the MySQL root password

In the admin node (rhell) run the mysql_secure_installation to set MySQL root password. The file is located at /usr/bin directory.

```

[root@rhell ]# cd /usr/bin/
[root@rhell bin]# mysql_secure_installation

*****OUTPUT*****
[root@rhell bin]# mysql_secure_installation

NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR ALL MySQL
      SERVERS IN PRODUCTION USE!  PLEASE READ EACH STEP CAREFULLY!

In order to log into MySQL to secure it, we'll need the current
password for the root user.  If you've just installed MySQL, and
you haven't set the root password yet, the password will be blank,
so you should just press enter here.

Enter current password for root (enter for none):
OK, successfully used password, moving on...

Setting the root password ensures that nobody can log into the MySQL
root user without the proper authorization.

Set root password? [Y/n] y
New password:
Re-enter new password:
Password updated successfully!
Reloading privilege tables..
... Success!

By default, a MySQL installation has an anonymous user, allowing anyone
to log into MySQL without having to have a user account created for
them.  This is intended only for testing, and to make the installation
go a bit smoother.  You should remove them before moving into a
production environment.

Remove anonymous users? [Y/n] y
... Success!

Normally, root should only be allowed to connect from 'localhost'.  This
ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] n
... skipping.

By default, MySQL comes with a database named 'test' that anyone can
access.  This is also intended only for testing, and should be removed
before moving into a production environment.

Remove test database and access to it? [Y/n] Y
- Dropping test database...
... Success!
- Removing privileges on test database...
... Success!

Reloading the privilege tables will ensure that all changes made so far
will take effect immediately.

Reload privilege tables now? [Y/n] Y
... Success!

Cleaning up...

All done!  If you've completed all of the above steps, your MySQL

```

```
installation should now be secure.
```

```
Thanks for using MySQL!
```

```
*****end*****
```

3. Installing the MySQL JDBC Driver

Install the JDBC driver on the Cloudera Manager Server host, as well as hosts which run the Activity Monitor, Reports Manager, Hive Metastore Server, Sentry Server, Cloudera Navigator Audit Server, and Cloudera Navigator Metadata Server roles.

- a. From the host connected to the internet, download the MySQL JDBC driver from: <http://www.mysql.com/downloads/connector/j/5.1.html>.



Note select the platform independent from the drop down list and download mysql-connector-java-5.1.34.tar.gz file

- b. Copy the downloaded file to the admin (rhel1) node, log in to the admin node and extract the file.

```
scp mysql-connector-java-5.1.34.tar.gz rhel1:/root/
```

- c. Login to the admin (rhel1) node and extract the file.

```
tar xzvf mysql-connector-java-5.1.34.tar.gz
```

- d. Create /usr/share/java/ directory in the admin node.

```
mkdir -p /usr/share/java/
```

- e. Go the mysql-connector-java-5.1.34 directory and copy the mysql-connector-java-5.1.35-bin.jar and rename it to the folder created above as shown in the command below:

```
cd mysql-connector-java-5.1.34
cp mysql-connector-java-5.1.34/mysql-connector-java-5.1.34-bin.jar
/usr/share/java/mysql-connector-java.jar
```

4. Creating Databases for Activity Monitor, Reports Manager and Hive Metastore Server.

- a. In the admin node (rhel1) Log into MySQL as the root user:

```
mysql -u root -p
```

Enter the password that was supplied in step 2.f above

```
Enter password:
```

- b. Create databases for the Activity Monitor, Reports Manager and Hive Metastore Server using the command below:

```
Mysql> create database amon DEFAULT CHARACTER SET utf8;
Mysql> create database rman DEFAULT CHARACTER SET utf8;
Mysql> create database metastore DEFAULT CHARACTER SET utf8;
```

```
mysql> create database amon DEFAULT CHARACTER SET utf8;
Query OK, 1 row affected (0.00 sec)

mysql> create database rman DEFAULT CHARACTER SET utf8;
Query OK, 1 row affected (0.00 sec)

mysql> create database metastore DEFAULT CHARACTER SET utf8;
Query OK, 1 row affected (0.00 sec)

mysql> create database sentry DEFAULT CHARACTER SET utf8;
Query OK, 1 row affected (0.00 sec)

mysql>
```

```
mysql> grant all on rman.*TO 'root'@'%' IDENTIFIED BY 'password';
mysql> grant all on metastore.*TO 'root'@'%' IDENTIFIED BY 'password';
mysql> grant all on amon.*TO 'root'@'%' IDENTIFIED BY 'password';
```

```
mysql> grant all on rman.*TO 'root'@'%' IDENTIFIED BY 'password';
Query OK, 0 rows affected (0.00 sec)

mysql> grant all on metastore.*TO 'root'@'%' IDENTIFIED BY 'password';
Query OK, 0 rows affected (0.00 sec)

mysql> grant all on amon.*TO 'root'@'%' IDENTIFIED BY 'password';
Query OK, 0 rows affected (0.00 sec)

mysql>
```

Cloudera Installation

The following section describes installation of Cloudera Manager first and then using Cloudera Manager to install CDH 5.3

Installing Cloudera Manager

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

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

Follow the steps below to install Cloudera Manager.

Update the repo files to point to local repository.

```
rm -f /var/www/html/clouderarepo/*.repo
cp /etc/yum.repos.d/c*.repo /var/www/html/clouderarepo/
```

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

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

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

```
cd /var/www/html/clouderarepo/
./cloudera-manager-installer.bin --skip_repo_package=1
```

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

```

Cloudera Manager 5

The Cloudera Manager Installer enables you to install Cloudera Manager and
bootstrap an entire CDH cluster, requiring only that you have SSH access to
your cluster's machines, and that those machines have Internet access.

This installer is only recommended for demonstration and proof of concept
deployments, but is not recommended for production deployments because it is
not intended to scale and may require database migration as your cluster
grows.

The Cloudera Manager Installer will automatically:

* Detect the operating system on the Cloudera Manager host
* Install the package repository for Cloudera Manager and the Java Runtime
Environment (JRE)
* Install the JRE if it's not already installed
* Install and configure an embedded PostgreSQL database
* Install and run the Cloudera Manager Server

Once server installation is complete, you can browse to Cloudera Manager's
web interface and use the cluster installation wizard to set up your CDH
cluster.

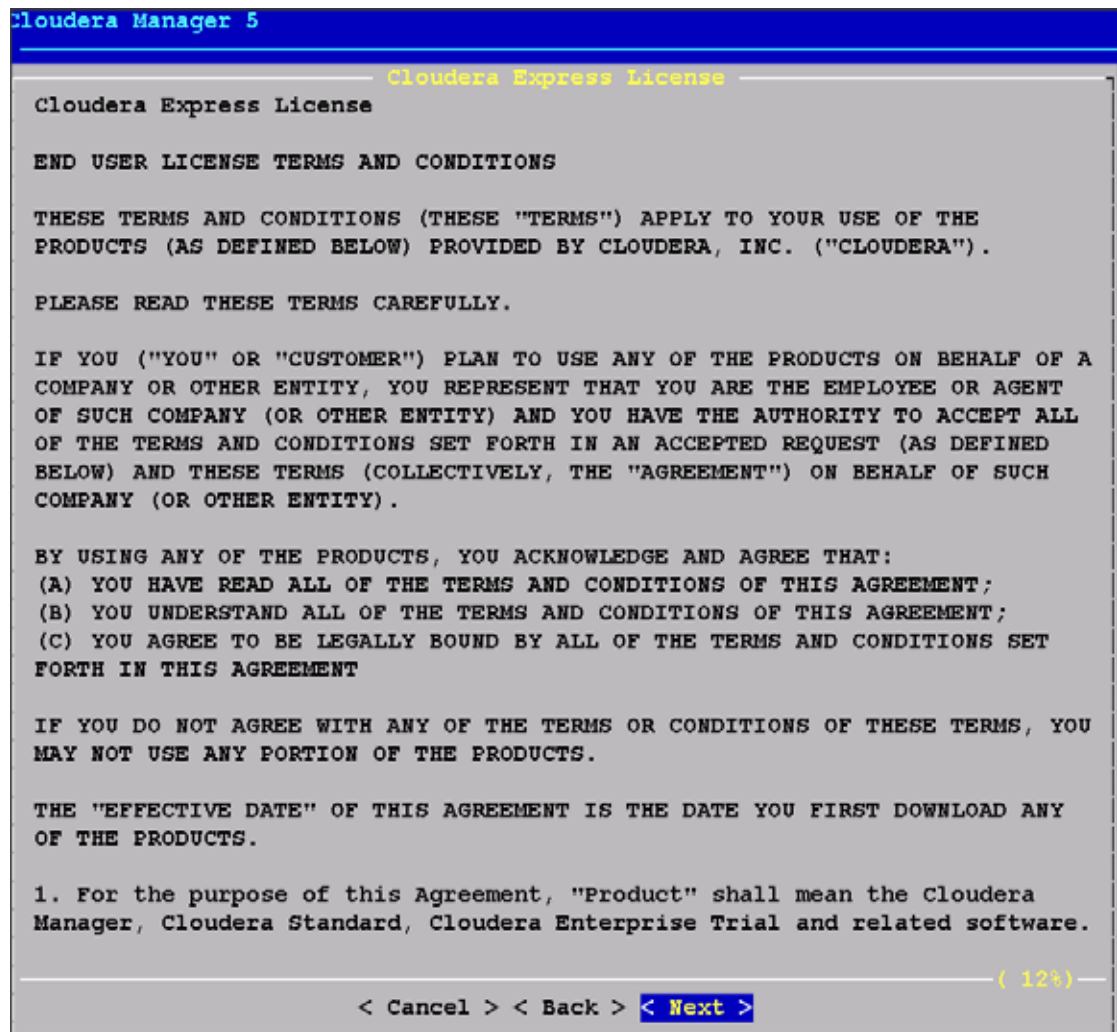
Cloudera Manager supports the following 64-bit operating systems:

* Red Hat Enterprise Linux 5 (Update 7 or later recommended)
* Red Hat Enterprise Linux 6 (Update 4 or later recommended)
* Oracle Enterprise Linux 5 (Update 6 or later recommended)
* Oracle Enterprise Linux 6 (Update 4 or later recommended)
* CentOS 5 (Update 7 or later recommended)
* CentOS 6 (Update 4 or later recommended)
* SUSE Linux Enterprise Server 11 (Service Pack 2 or later recommended)

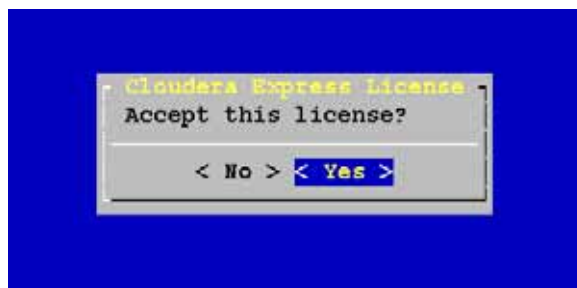
< Cancel > < Next >

```


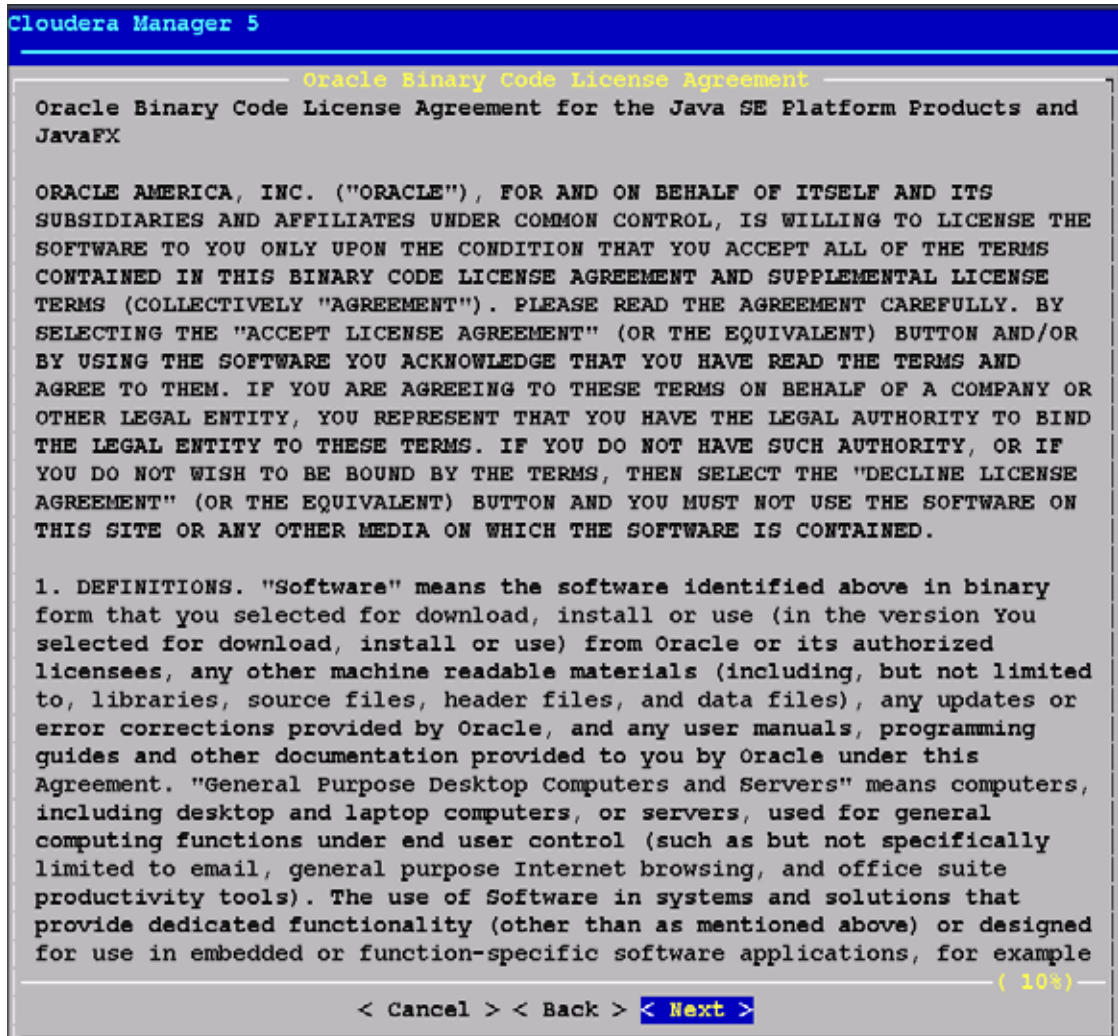
4. Click **Next** in the End User License agreement page.




5. Click **Yes** in the license agreement confirmation page.




6. Click **Next** in Oracle Binary Code License Agreement and **Yes** in the Oracle Binary Code License Agreement for the Java SE Platform Products page.




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



- Save the url displayed `http://10.0.145.45:7180`. You will need this url to access Cloudera Manager. If you are unable to connect to the server, check to see if iptables and SELinux are disabled.



- Click **OK**.



- Once the installation of Cloudera Manager is complete. Install CDH5 using the Cloudera Manager web interface.

Setting up the Cloudera Manager Server Database

The Cloudera Manager Server database stores information about service and host configurations.

Preparing a Cloudera Manager Server External Database

- Run the `scm_prepare_database.sh` script on the host where the Cloudera Manager Server package is installed:

```
[root@rhell ~]# cd /usr/share/cm/s/schema
[root@rhell schema]# ./scm_prepare_database.sh mysql amon root <password>
[root@rhell schema]# ./scm_prepare_database.sh mysql rman root <password>
[root@rhell schema]# ./scm_prepare_database.sh mysql metastore root <password>
```

- Verify the database connectivity using the following command.

```
[root@rhell ~]# mysql -u root -p
Mysql> connect amon
Mysql> connect rman
Mysql> connect metastore
```

```
[root@rhel1 schema]# mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 21
Server version: 5.1.71 Source distribution

Copyright (c) 2000, 2013, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql> connect amon
Connection id:    25
Current database: amon

mysql> connect metastore
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Connection id:    27
Current database: metastore

mysql> █
```

The MySQL External database setup is complete.

Installing Cloudera Enterprise Data Hub Edition (CDH5)

Role Assignment

This is one of the critical consideration for the installation. Inspect and customize the role assignments of all the nodes based on your requirements.

Table 11 Service Assignment

Service Name	Host
NameNode	rhel1, rhel3 (HA)
HistoryServer	rhel1
ResouceManager	rhel2, rhel3 (HA)
Hue Server	rhel2
HiveMetastore Server	rhel1
HiveServer2	rhel2
HBase Master	rhel2
Oozie Server	rhel1
Zookeeper	rhel1, rhel2, rhel3

Table 11 Service Assignment

JournalNodes	rhel1, rhel2, rhel3
DataNode	rhel4 to rhel160
NodeManager	rhel4 to rhel160
RegionServer	rhel4 to rhel160
Sqoop Server	rhel1
Impala Catalog Server Daemon	rhel2
Solr Server	rhel1
Spark Server	rhel1
Spark Worker	rhel4 to rhel160

Scaling the Cluster

The role assignment recommendation above is for clusters larger than 64 servers and in High Availability (HA). For smaller cluster running without HA the recommendation is to dedicate one server for name node and a second server for secondary name node and YARN Resource Manager. For larger clusters larger than 64 nodes the recommendation is to dedicate one server each for name node, YARN Resource Manager and one more for running both NameNode (HA) and ResourceManager(HA) as in the table (no Secondary Namenode when in HA).

HDFS High Availability (HA)

The HDFS HA feature provides the option of running two NameNodes in the same cluster, in an Active/Passive configuration. These are referred to as the Active NameNode and the Standby NameNode. Unlike the Secondary NameNode, the Standby NameNode is hot standby, allowing a fast failover to a new NameNode in the case that a machine crashes, or a graceful administrator-initiated failover for the purpose of planned maintenance. There cannot be more than two NameNodes.

For more information, see:

http://www.cloudera.com/content/cloudera/en/documentation/core/v5-3-x/topics/cdh_hag_hdfs_ha_intro.html



Note

Setting up HDFS HA is done after Cloudera Install.

Map-Reduce HA (YARN/MRv2)

The YARN ResourceManager (RM) is responsible for tracking the resources in a cluster and scheduling applications (for example, MapReduce jobs). Before CDH 5, the RM was a single point of failure in a YARN cluster. The RM high availability (HA) feature adds redundancy in the form of an Active/Standby RM pair to remove this single point of failure. Furthermore, upon failover from the Standby RM to the Active, the applications can resume from their last check-pointed state; for example, completed map tasks in a MapReduce job are not re-run on a subsequent attempt. This allows events such the following to be handled without any significant performance effect on running applications.

- Unplanned events such as machine crashes.
- Planned maintenance events such as software or hardware upgrades on the machine running the ResourceManager.

For more information, see:

http://www.cloudera.com/content/cloudera/en/documentation/core/v5-3-x/topics/cdh_hag_rm_ha_config.html



Note

Setting up YARN HA is done after Cloudera Install.

To install Cloudera Enterprise Data Hub, follow these steps:

1. Access the Cloudera Manager using the URL displayed by the Installer, [http:// 10.0.145.45:7180](http://10.0.145.45:7180).
2. Login to the Cloudera Manager. Enter "admin" for both the Username and Password fields.

Figure 163 *Login to Cloudera Manager*

Login

Username:
admin

Password:

Remember me on this computer.

Login

3. If you do not have a Cloudera license, click Cloudera Enterprise Data Hub Trial Edition. If you do have a Cloudera license, Click "Upload License" and select your license.
4. Based on requirement Choose appropriate Cloudera Editions for Installation.

Figure 164 Installing Cloudera Enterprise

Welcome to Cloudera Manager. Which edition do you want to deploy?

Upgrading to **Cloudera Enterprise Data Hub Edition** provides important features that help you manage and monitor your Hadoop clusters in mission-critical environments.

	Cloudera Express	Cloudera Enterprise Data Hub Edition Trial	Cloudera Enterprise
License	Free	60 Days After the trial period, the product will continue to function as Cloudera Express . Your cluster and your data will remain unaffected.	Annual Subscription cisco_dev_cloudera_e <input type="button" value="Upload"/> Success Cloudera Enterprise is available in three editions: • Basic Edition • Flex Edition • Data Hub Edition
Node Limit	Unlimited	Unlimited	Unlimited
CDH	✓	✓	✓
Core Cloudera Manager Features	✓	✓	✓
Advanced Cloudera Manager Features		✓	✓
Cloudera Navigator		✓	✓
Cloudera Support			✓

For full list of features available in **Cloudera Express** and **Cloudera Enterprise**, [click here](#).

- Click **Continue** in the confirmation page.

Figure 165 Confirmation Page

Thank you for choosing Cloudera Manager and CDH.

This installer will install **Cloudera Enterprise 5.3.2** and enable you to later choose packages for the services below (there may be some license implications).

- Apache Hadoop (Common, HDFS, MapReduce, YARN)
- Apache HBase
- Apache ZooKeeper
- Apache Oozie
- Apache Hive
- Hue (Apache licensed)
- Apache Flume
- Cloudera Impala (Apache licensed)
- Apache Sentry
- Apache Sqoop
- Cloudera Search (Apache licensed)
- Apache Spark

You are using Cloudera Manager to install and configure your system. You can learn more about Cloudera Manager by clicking on the **Support** menu above.

[» Continue](#)

Edit the Cloudera Enterprise Parcel Settings to Use the CDH 5.3.2 Parcels

1. At this point, open another tab in the same browser window and visit the URL:
<http://10.0.145.45:7180/cmf/parcel/status> for modifying the parcel settings.
2. Click **Edit Settings** on this page:
3. Click – to remove all the remote repository URLs, and add the URL to the location where we kept the CDH 5.3.2 parcels i.e. <http://10.0.145.45/CDH5.3parcels/>

Figure 166 Edit Cloudera Enterprise Parcel Settings

Settings

Search X Settings marked will not take effect until

Category	Property	Value
Performance	Local Parcel Repository Path	/opt/cloudera/parcel-repo
Advanced	Parcel Update Frequency	1 hour(s)
Monitoring	Remote Parcel Repository URLs	<input type="text" value="http://10.0.145.45/CDH5.3parcels/"/> + Reset to the default value: http://archive.cloudera.com/cdh5/parcels/latest/ http://archive.cloudera.com/cdh4/parcels/latest/ http://archive.cloudera.com/impala/parcels/latest/ http://archive.cloudera.com/search/parcels/latest/
Security	Create System-Wide Symlinks for Active Parcels	<input checked="" type="checkbox"/>
Ports and Addresses	Cloudera Manager Manages Parcels	<input checked="" type="checkbox"/>
Kerberos	Create Users and Groups, and Apply File Permissions for Parcels	<input checked="" type="checkbox"/>
Other	Automatically Download New Parcels	<input type="checkbox"/>
Support		
External Authentication		
Parcels		
Network		
Custom Service Descriptors		

4. Click **Save Changes** to finish the configuration.

Now navigate back to the Cloudera installation home page i.e. <http://10.0.145.45:7180>

5. Click **Continue** in the confirmation page.

Figure 167 Cloudera Confirmation Page

Thank you for choosing Cloudera Manager and CDH.

This installer will install **Cloudera Enterprise 5.3.2** and enable you to later choose packages for the services below (there may be some license implications).

- Apache Hadoop (Common, HDFS, MapReduce, YARN)
- Apache HBase
- Apache ZooKeeper
- Apache Oozie
- Apache Hive
- Hue (Apache licensed)
- Apache Flume
- Cloudera Impala (Apache licensed)
- Apache Sentry
- Apache Sqoop
- Cloudera Search (Apache licensed)
- Apache Spark

You are using Cloudera Manager to install and configure your system. You can learn more about Cloudera Manager by clicking on the **Support** menu above.

» Continue

- Specify the hosts that are part of the cluster using their IP addresses or hostname. The figure below shows use of a pattern to specify IP addresses range.

10.0.146.[45-204] or rhel[1-160]



Note Here, eth1 IP is provided as this is provided with better QoS policy than management VLAN.

- After the IP addresses or hostnames are entered, click **Search**.

Figure 168 Searching for Cluster Nodes

Specify hosts for your CDH cluster installation.

Hint: Search for hostnames and/or IP addresses using [patterns](#) IP.

rhel[1-160]

SSH Port:

8. Cloudera Manager will "discover" the nodes in the cluster. Verify that all desired nodes have been found and selected for installation.
9. For the method of installation, select the **Use Parcels (Recommended)** radio button.
10. For the CDH version, select the **CDH5** radio button.
11. For the specific release of Cloudera Manager, select the **Custom Repository** radio button.
12. Enter the URL for the repository within the admin node.
<http://10.0.145.45/clouderarepo/cloudera-manager> and click **Continue**.

Figure 169 Cluster Installation: Selecting Repository

Select Repository

Cloudera recommends the use of parcels for installation over packages, because parcels enable Cloudera Manager to easily manage the software on your cluster, deployment and upgrade of service binaries. Electing not to use parcels will require you to manually upgrade packages on all hosts in your cluster when software is available, and will prevent you from using Cloudera Manager's rolling upgrade capabilities.

Choose Method

Use Packages

Use Parcels (Recommended) [More Options](#)

Select the version of CDH

CDH-5.3.2-1.cdh5.3.2.p0.10

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

Matched release for this Cloudera Manager Server

Custom Repository

Example for SLES, Redhat or other RPM based distributions:
http://archive.cloudera.com/cm5/redhat/6/x86_64/cm/5/

Example for Ubuntu or other Debian based distributions:
`deb http://archive.cloudera.com/cm5/ubuntu/lucid/amd64/cm/ lucid-cm5 contrib`

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

[Back](#) 1 2 3 4 5 6 7 8

13. Check "install Oracle java SE Development kit (JDK)" and click **Continue**.

Figure 170 Cluster Installation: JDK Installation

Cluster Installation

JDK Installation Options

Oracle Binary Code License Agreement for the Java SE Platform Products and JavaFX

ORACLE AMERICA, INC. ("ORACLE"), FOR AND ON BEHALF OF ITSELF AND ITS SUBSIDIARIES AND AFFILIATES UNDER COMMON CONTROL, IS WILLING TO LICENSE THE SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS (COLLECTIVELY "AGREEMENT"). PLEASE READ THE AGREEMENT CAREFULLY. BY SELECTING THE "ACCEPT LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND/OR BY USING THE SOFTWARE YOU ACKNOWLEDGE THAT YOU HAVE READ THE TERMS AND AGREE TO THEM. IF YOU ARE AGREEING TO THESE TERMS ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO BE BOUND BY THE TERMS, THEN SELECT THE "DECLINE LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND YOU MUST NOT USE THE SOFTWARE ON THIS SITE OR ANY OTHER MEDIA ON WHICH THE SOFTWARE IS CONTAINED.

1. DEFINITIONS. "Software" means the software identified above in binary form that you selected for download, install or use (in the version you selected for download, install or use) from Oracle or its authorized licensees, any other machine readable materials (including, but not limited to, libraries, source files, header files, and data files), any updates or error corrections provided by Oracle, and any user manuals, programming guides and other documentation provided to you by Oracle under this Agreement. "General Purpose Desktop Computers and Servers" means computers, including desktop and laptop computers, or servers, used for general computing functions under end user control (such as but not specifically limited to email, general purpose Internet browsing, and office suite productivity tools). The use of Software in systems and solutions that provide dedicated functionality (other than as mentioned above) or designed for use in embedded or function-specific software applications, for example but not limited to: Software embedded in or bundled with industrial control systems, wireless mobile telephones, wireless handheld devices, kiosks, TV/STB, Blu-ray Disc devices, telematics and network control switching equipment, printers and storage management systems, and other related systems are excluded from this definition and not licensed under this Agreement. "Programs" means (a) Java technology applets and applications intended to run on the Java Platform, Standard Edition platform on Java-enabled General Purpose Desktop Computers and Servers; and (b) JavaFX technology applications intended to run on the JavaFX Runtime on JavaFX-enabled General Purpose Desktop Computers and Servers. "Commercial Features" means those features identified in Table 1-1 (Commercial Features In Java SE Product Editions) of the Java SE documentation accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>. "README File" means the README file for the Software accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

2. LICENSE TO USE. Subject to the terms and conditions of this Agreement including, but not limited to, the Java Technology Restrictions of the Supplemental License Terms.

Install Oracle Java SE Development Kit (JDK)
 Check this box to accept the Oracle Binary Code License Agreement and install the JDK. Leave it unchecked to use a currently installed JDK.

Install Java Unlimited Strength Encryption Policy Files
 Check this checkbox if local laws permit you to deploy unlimited strength encryption and you are running a secure cluster.

⏪ Back

1
2
3
4
5
6
7
8

⏩ Continue

14. Click **Continue** again.

Cluster Installation

Enable Single User Mode

Only supported for CDH 5.2 and above.

By default, service processes run as distinct users on the system. For example, HDFS DataNodes run as user "hdfs" and HBase RegionServers run as user "hbase." Enabling "single user mode" configures Cloudera Manager to run service processes as a single user, by default "cloudera-scm", thereby prioritizing isolation between managed services and the rest of the system over isolation between the managed services.

The **major benefit** of this option is that the Agent does not run as root. However, this mode complicates installation, which is described fully in the [documentation](#). Most notably, directories which in the regular mode are created automatically by the Agent, must be created manually on every host with appropriate permissions, and sudo (or equivalent) access must be set up for the configured user.

Switching back and forth between single user mode and regular mode is not supported.

Single User Mode

Configure all clusters to run in single user mode where the Cloudera Manager agent and all service processes run as the same system user. Only supported for CDH 5.2 and above.

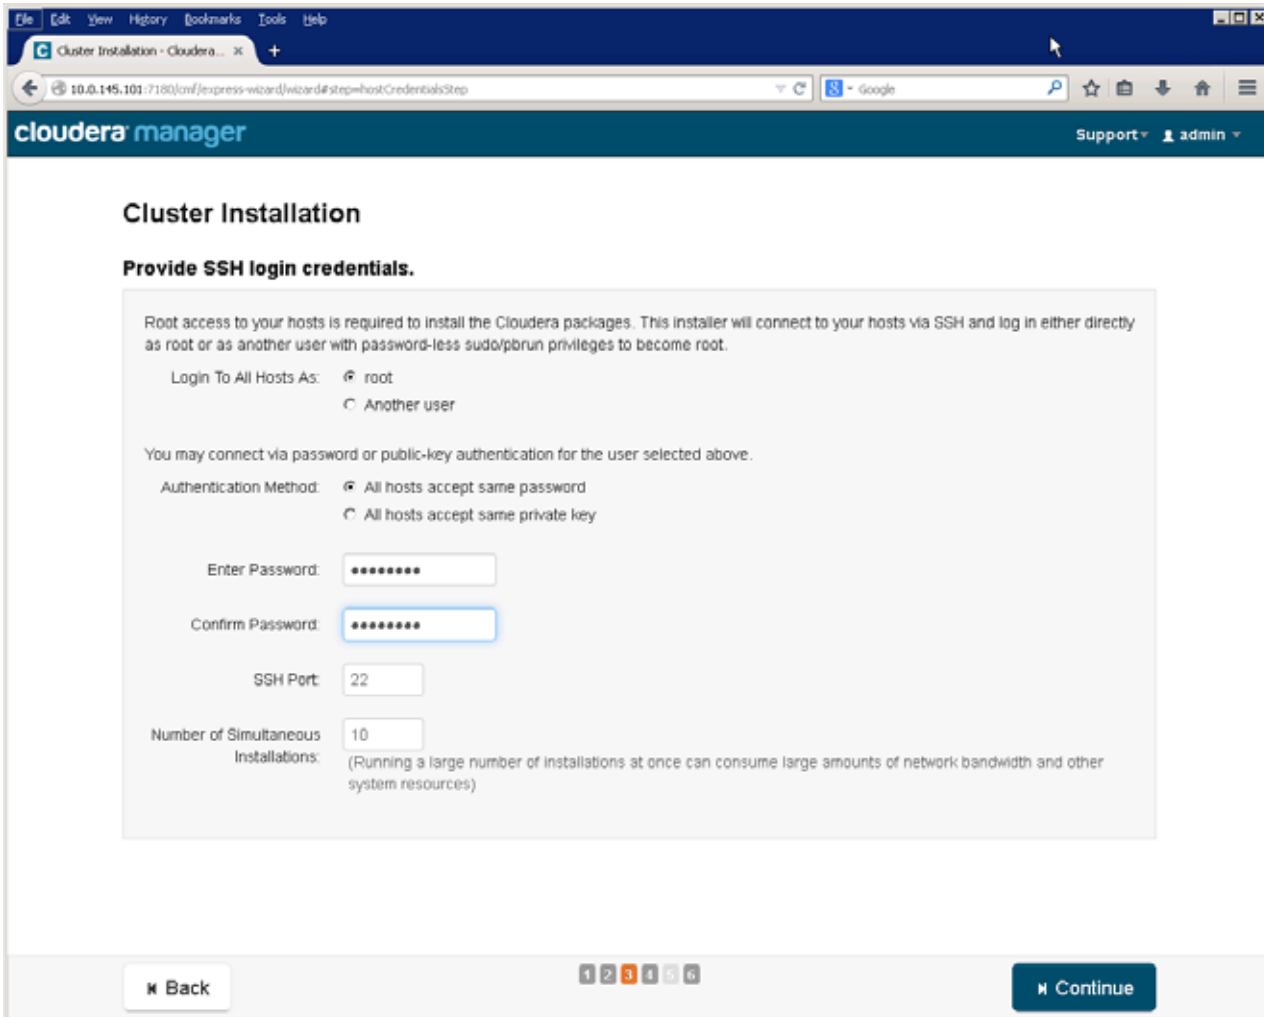
◀ Back

1 2 3 4 5 6 7 8

Continue ▶

15. Provide SSH login credentials for the cluster and click **Continue**.

Figure 171 Login Credentials to Start CDH Installation



The screenshot shows a web browser window displaying the Cloudera Manager interface. The page title is "Cluster Installation" and the sub-header is "Provide SSH login credentials." The main content area contains the following information:

- Root access to your hosts is required to install the Cloudera packages. This installer will connect to your hosts via SSH and log in either directly as root or as another user with password-less sudo/pbrun privileges to become root.**
- Login To All Hosts As:** root, Another user
- You may connect via password or public-key authentication for the user selected above.**
- Authentication Method:** All hosts accept same password, All hosts accept same private key
- Enter Password:** [password field]
- Confirm Password:** [password field]
- SSH Port:** [22]
- Number of Simultaneous Installations:** [10] (Running a large number of installations at once can consume large amounts of network bandwidth and other system resources)

At the bottom of the page, there is a "Back" button on the left, a progress indicator with six steps (step 3 is highlighted), and a "Continue" button on the right.









16. After the installation is successful click **Continue** to begin the parcel installation.

Figure 172 Cluster Installation: Completed

Cluster Installation

Installation completed successfully.

8 of 8 host(s) completed successfully.

Hostname	IP Address	Progress	Status	
rhel1	10.0.146.45		✓ Installation completed successfully.	Details ⓘ
rhel2	10.0.146.46		✓ Installation completed successfully.	Details ⓘ
rhel3	10.0.146.47		✓ Installation completed successfully.	Details ⓘ
rhel4	10.0.146.48		✓ Installation completed successfully.	Details ⓘ
rhel5	10.0.146.49		✓ Installation completed successfully.	Details ⓘ
rhel6	10.0.146.50		✓ Installation completed successfully.	Details ⓘ
rhel7	10.0.146.51		✓ Installation completed successfully.	Details ⓘ
rhel8	10.0.146.52		✓ Installation completed successfully.	Details ⓘ

17. Installation using parcels begins.

Figure 173 Cluster Installation: Installation Selected Parcels

Cluster Installation

Installing Selected Parcels

The selected parcels are being downloaded and installed on all the hosts in the cluster.

CDH 5.3.2-1.cdh5.3.2.p0.10

	Downloading 100%
	Distributing 0%
	Activating 0%

18. Once the installation is completed successfully click **Continue** to select the required services.

Figure 174 Cluster Installation: Selected Parcels Installation Complete

Cluster Installation

Installing Selected Parcels

The selected parcels are being downloaded and installed on all the hosts in the cluster.

CDH 5.3.2-1.cdh5.3.2.p0.10



19. Wait for Cloudera Manager to inspect the hosts on which it has just performed the installation.
20. Review and verify the summary. Click **Finish**.

Figure 175 Inspecting Hosts for Correctness

Cluster Installation

Inspect hosts for correctness ⌂ Run Again

Validations

- ✓ Inspector ran on all 8 hosts.
- ✓ The following failures were observed in checking hostnames...
- ✓ No errors were found while looking for conflicting init scripts.
- ✓ No errors were found while checking /etc/hosts.
- ✓ All hosts resolved localhost to 127.0.0.1.
- ✓ All hosts checked resolved each other's hostnames correctly and in a timely manner.
- ✓ Host clocks are approximately in sync (within ten minutes).
- ✓ Host time zones are consistent across the cluster.
- ✓ No users or groups are missing.
- ✓ No conflicts detected between packages and parcels.
- ✓ No kernel versions that are known to be bad are running.
- ✓ No performance concerns with Transparent Huge Pages settings.
- ✓ CDH 5 Hue Python version dependency is satisfied.
- ✓ 0 hosts are running CDH 4 and 8 hosts are running CDH5.
- ✓ All checked hosts in each cluster are running the same version of components.
- ✓ All managed hosts have consistent versions of Java.
- ✓ All checked Cloudera Management Daemons versions are consistent with the server.

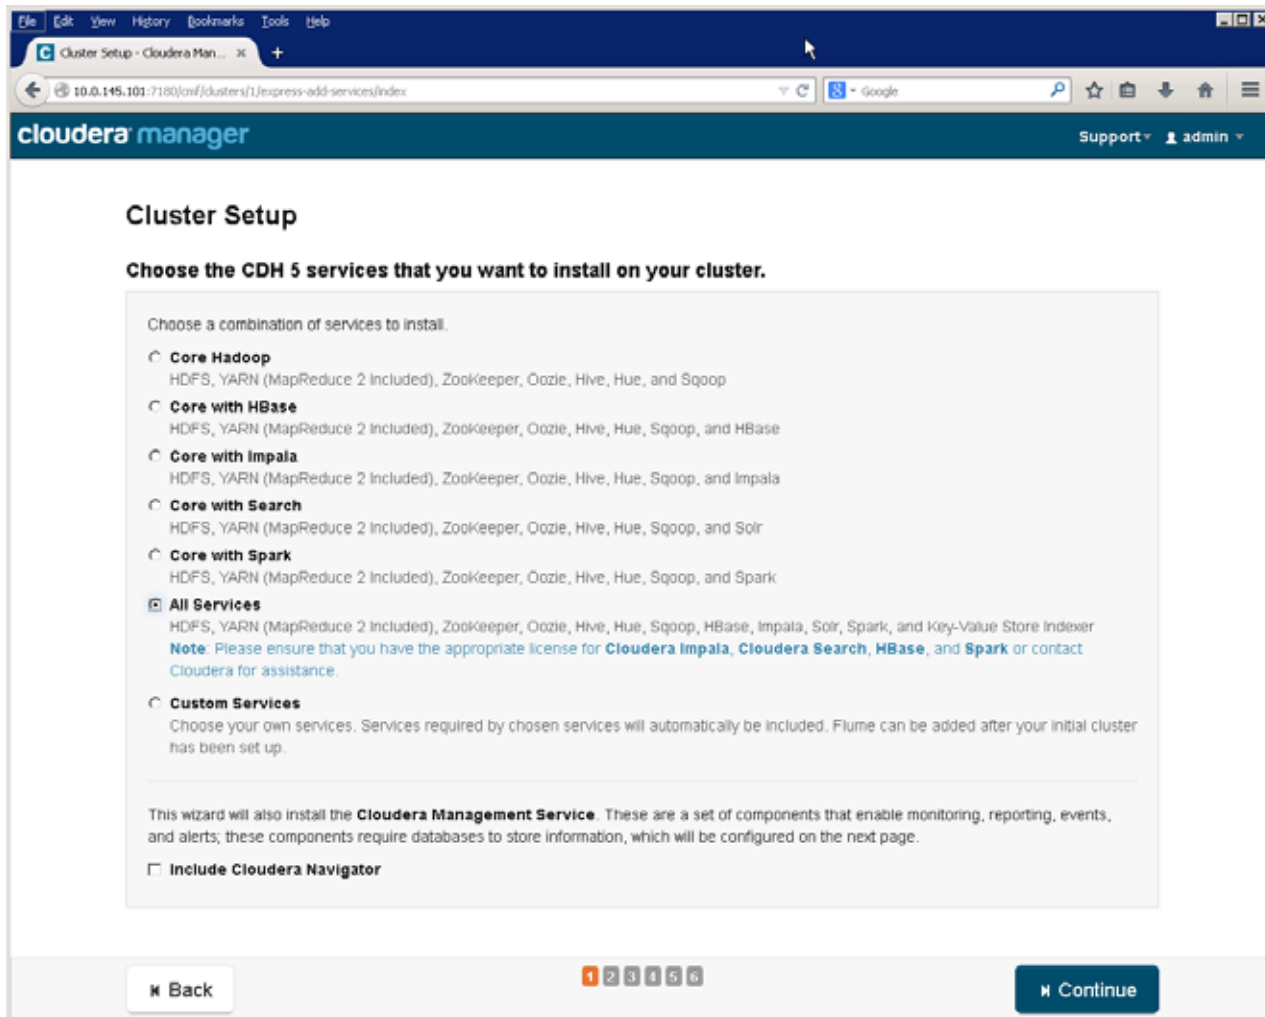
⌂ Back

1 2 3 4 5 6 7 8

⌂ Finish

21. Select services that need to be started on the cluster.

Figure 176 Selecting CDH Version and Services



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

Reconfigure the service assignment to match the table in “Role Assignment” section on page 188.

Figure 177 Cluster Setup: Role Assignment - Part1

Cluster Setup

Customize Role Assignments

You can customize the role assignments for your new cluster here, but if assignments are made incorrectly, such as assigning too many roles to a single host, this can impact the performance of your services. Cloudera does not recommend altering assignments unless you have specific requirements, such as having pre-selected a specific host for a specific role.

You can also view the role assignments by host. [View By Host](#)

HBase

- Master** × 1 New:
- HBase REST Server**:
- HBase Thrift Server**:
- RegionServer** × 158 New:

HDFS

- NameNode** × 1 New:
- SecondaryNameNode** × 1 New:
- Balancer** × 1 New:
- HttpFS**:
- NFS Gateway**:
- DataNode** × 158 New:

Hive

- Gateway** × 160 New:
- Hive Metastore Server** × 1 New:
- WebHCat Server**:
- HiveServer2** × 1 New:

Figure 178 Cluster Setup: Role Assignment - Part2

Hue

- Hue Server** × 1 New:

Impala

- Impala Catalog Server** × 1 New:
- Impala StateStore** × 1 New:
- Impala Daemon** × 159 New:

Key-Value Store Indexer

- Lily HBase Indexer** × 1 New:

Cloudera Management Service

- Service Monitor** × 1 New:
- Activity Monitor** × 1 New:
- Host Monitor** × 1 New:
- Reports Manager** × 1 New:
- Event Server** × 1 New:
- Alert Publisher** × 1 New:
- Navigator Audit Server** × 1 New:
- Navigator Metadata Server** × 1 New:

Figure 179 Cluster Setup: Role Assignment - Part3

The screenshot displays the role assignment configuration for a Cloudera cluster. The roles and their configurations are as follows:

- Oozie:** 1 Oozie Server (New) on host rhel1.
- Solr:** 1 Solr Server (New) on host rhel1.
- Spark:** 1 History Server (New) on host rhel1 and 1 Gateway on host Select hosts.
- Sqoop 2:** 1 Sqoop 2 Server (New) on host rhel1.
- YARN (MR2 Included):** 1 ResourceManager (New) on host rhel2, 1 JobHistory Server (New) on host rhel2, and 158 NodeManagers (New) on host Same As DataNode.
- ZooKeeper:** 3 Servers (New) on host rhel[1-3].

At the bottom of the screen, there is a 'Back' button on the left, a progress indicator in the center showing 6 steps, and a 'Continue' button on the right.

Using Custom Database

The role assignment recommendation above is for clusters of up to 160 servers. For clusters larger than 160 nodes the recommendation is to dedicate one server each for name node, secondary name node and YARN Resource Manager.

1. Select the **Use Custom Database** radio button.
2. In Database Host Name sections use port 3306 for TCP/IP because connection to the remote server always uses TCP/IP.
3. Enter the Database Name. Username and password that was used during the database creation stage. (Please refer Setting MySQL Database for Cloudera Manager section).
4. Click **Test Connection** to verify the connection, once the connection is successful click **Continue**.

Figure 180 Database Setup

Cluster Setup

Database Setup

Configure and test database connections. Create the databases first according to the [Installing and Configuring an External Database](#) section of the [Installation Guide](#).

Use Custom Databases
 Use Embedded Database

Hive ✔ Successful

Database Host Name: Database Type: Database Name: Username: Password:

Activity Monitor ✔ Successful

Currently assigned to run on **rhel1**.

Database Host Name: Database Type: Database Name: Username: Password:

Reports Manager ✔ Successful

Currently assigned to run on **rhel1**.

Database Host Name: Database Type: Database Name: Username: Password:

Show Password

1 2 3 4 5 6

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

Configuring Yarn (MR2 Included) and HDFS Services

The following parameters are used for Cisco UCS Integrated Infrastructure for Big Data Performance Optimized cluster configuration described in this document. These parameters are to be changed based on the cluster configuration, number of nodes and specific workload.

Table 12 Yarn-MR2 Included

Service	Value
mapreduce.map.memory.mb	3 GiB
mapreduce.reduce.memory.mb	3 GiB
mapreduce.map.java.opts.max.heap	2560 MiB
yarn.nodemanager.resource.memorymb	180 GiB
yarn.nodemanager.resource.cpu-vcores	32
yarn.scheduler.minimum-allocation-mb	4 GiB
yarn.scheduler.maximum-allocation-mb	180 GiB

Table 12 *Yarn-MR2 Included*

yarn.scheduler.maximum-allocation-vcores	40
mapreduce.task.io.sort.mb	256 MiB

Table 13 *HDFS*

Service	Value
dfs.datanode.failed.volumes.tolerated	11
dfs.datanode.du.reserved	10 GiB
dfs.datanode.data.dir.perm	755
Java Heap Size of Namenode in Bytes	2628 MiB
dfs.namenode.handler.count	54
dfs.namenode.service.handler.count	54
Java Heap Size of Secondary namenode in Bytes	2628 MiB

Figure 181

Cluster Setup

Review Changes

HDFS Root Directory The HDFS directory shared by HBase RegionServers.

hbase.rootdir HBase (Service-Wide)

Enable Replication Allow HBase tables to be replicated.

hbase.replication HBase (Service-Wide)

Enable Indexing Allow indexing of tables in HBase by Lily HBase Indexer.
Note: Replication must be enabled for indexing to work.

HBase (Service-Wide)

DataNode Data Directory Comma-delimited list of directories on the local file system where the DataNode stores HDFS block data. Typical values are /data/N/dfs/dn for N = 1, 2, 3... These directories should be mounted using the noatime option and the disks should be configured using JBOD. RAID is not recommended.

dfs.data.dir, dfs.datanode.data.dir DataNode Default Group

⏪ Back

1 2 3 4 5 6

Continue ⏩

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



Note

In case Sqoop2 service doesn't startup due to error "Unable to create database", while cluster setup stage, please go to troubleshooting section in the Appendix A.

Figure 182 Starting the Cluster Services

Command	Context	Status	Started at	Ended at
✓ First Run		Finished	Mar 29, 2015 1:50:16 PM PDT	Mar 29, 2015 2:00:31 PM PDT
Finished First Run of all services successfully.				

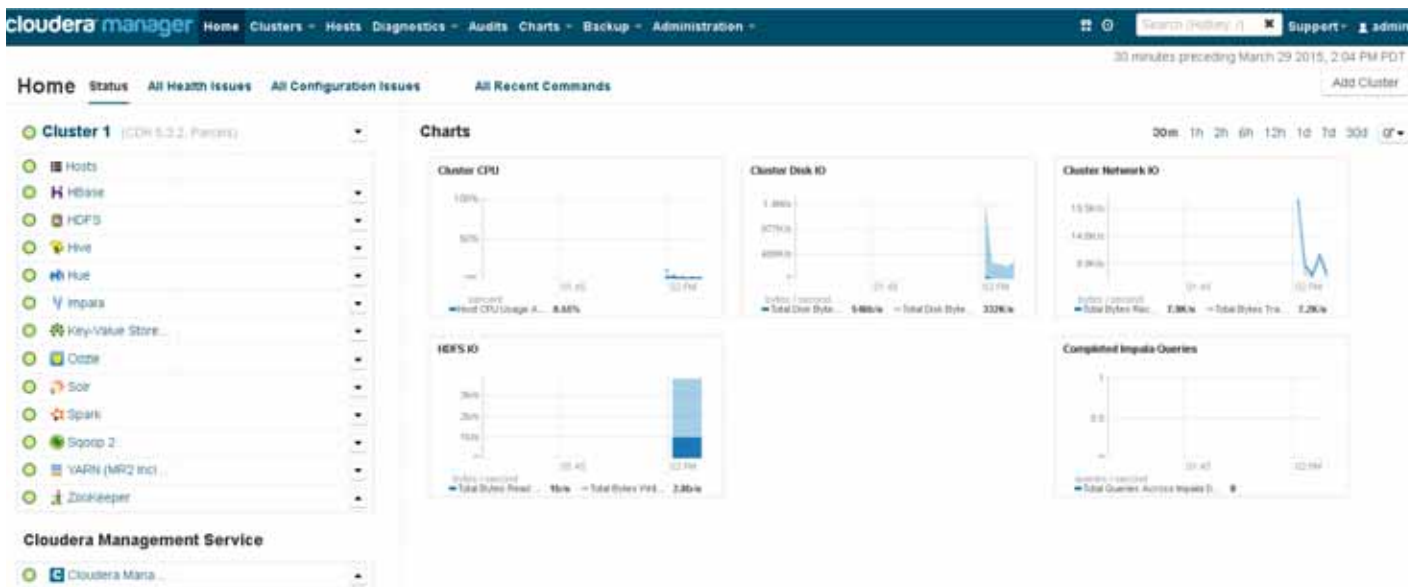
Command Progress	
Completed 33 of 33 steps	
✓	Initializing Zookeeper Service Completed 3 steps successfully.
✓	Starting Zookeeper Service Completed 3 steps successfully. Details
✓	Checking if the name directories of the Hadoop are empty. Formatting HDFS only if empty. Successfully formatted hadoop.
✓	Starting HDFS Service Successfully started HDFS service. Details
✓	Creating HDFS tmp directory Successfully created HDFS directory /tmp.
✓	Creating HBase root directory Successfully created HDFS directory.
✓	Starting HBase Service Service started successfully. Details
✓	Initializing Solr in Zookeeper Successfully initialized solr service. Details
✓	Creating HDFS home directory for Solr Successfully created HDFS directory. Details
✓	Starting Solr Service Service started successfully. Details
✓	Creating MR2 job history directory Successfully created HDFS directory. Details
✓	Creating NodeManager remote application log directory Successfully created HDFS directory. Details
✓	Starting YARN (MR2 included) Service Successfully started service. Details
✓	Creating Hive Metastore Database Tables Created Hive Metastore Database Tables successfully. Details
✓	Creating Hive user directory Successfully created HDFS directory. Details
✓	Creating Hive warehouse directory Successfully created HDFS directory. Details
✓	Starting Hive Service Service started successfully. Details
✓	Starting Key-Value Store Indexer Service Service started successfully. Details
✓	Execute command CreateSparkUserDirCommand on service Spark Successfully created HDFS directory. Details
✓	Execute command CreateSparkHistoryDirCommand on service Spark Successfully created HDFS directory. Details
✓	Execute command SparkUploaderServiceCommand on service Spark Command upload spark jar finished successfully on service spark_on_yarn. Details
✓	Starting Spark Service Service started successfully. Details
✓	Creating Sqoop 2 user directory Successfully created HDFS directory. Details
✓	Creating Sqoop 2 Database Created sqoop database successfully.

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

Figure 183 Cluster Setup Completion



Figure 184 Service Status of the Cluster



Setting up HDFS HA

The **Enable High Availability** workflow leads through adding a second (standby) NameNode and configuring JournalNodes. During the workflow, Cloudera Manager creates a **federated namespace**.

- Log in to the admin node (rhel1) and create the Edit directory for the JournalNode hosts.

```
clush -w rhel[1-3] mkdir -p /data/disk1/namenode-edits
clush -w rhel[1-3] chmod 777 /data/disk1/namenode-edits
```



```

root@rhell1 ~]# clush -w rhel[1-3] mkdir -p /data/disk1/namenode-edits
root@rhell1 ~]# clush -w rhel[1-3] chmod 777 /data/disk1/namenode-edits
root@rhell1 ~]# █

```

2. Log in to the Cloudera manager and go to the HDFS service.
3. In the top right corner Select **Actions** > **Enable High Availability**. A screen showing the hosts that are eligible to run a standby NameNode and the JournalNodes displays.

Enable High Availability for HDFS

Getting Started

This wizard leads you through adding a standby NameNode, restarting the HDFS service and any dependent services, and then re-deploying client configurations.

Nameservice Name

Enabling High Availability creates a new nameservice. Accept the default name `nameservice1` or provide another name in Nameservice Name.

◀ Back ▶ Continue

4. Specify a name for the nameservice or accept the default name `nameservice1` and click **Continue**.
5. In the **NameNode Hosts** field, click **Select a host**. The host selection dialog displays.

Enable High Availability for HDFS

Assign Roles

NameNode Hosts

[Select a host](#)

JournalNode Hosts

[Select hosts](#)

We recommend that JournalNodes be hosted on machines of similar hardware specifications as the NameNodes. The hosts of NameNodes and the JobTracker are generally good options. You must have a minimum of three and an odd number of JournalNodes

⏪ Back
🔍 🔍 🔍 🔍
Continue ⏩

6. Check the checkbox next to the hosts (rhe13) where you want the standby NameNode to be set up and click **OK**.



Note The standby NameNode cannot be on the same host as the active NameNode, and the host that is chosen should have the same hardware configuration (RAM, disk space, number of cores, and so on) as the active NameNode.

7. In the **JournalNode Hosts** field, click **Select hosts**. The host selection dialog displays.
8. Check the checkboxes next to an odd number of hosts (a minimum of three) to act as JournalNodes and click **OK**. Here we are using the same nodes as Zookeeper nodes.

Enable High Availability for HDFS

Assign Roles

NameNode Hosts

JournalNode Hosts

We recommend that JournalNodes be hosted on machines of similar hardware specifications as the NameNodes. The hosts of NameNodes and the JobTracker are generally good options. You must have a minimum of three and an odd number of JournalNodes.



Note JournalNodes should be hosted on hosts with similar hardware specification as the NameNodes. It is recommended that you put a JournalNode each on the same hosts as the active and standby NameNodes, and the third JournalNode on ResourceManager node.

9. Click **Continue**.

10. In the **JournalNode Edits Directory** property, enter a directory location created earlier in step 1 for the JournalNode edits directory into the fields for each JournalNode host.



Note The directories you specify should be empty, and must have the appropriate permissions.

cloudera manager Support | Admin

Enable High Availability for HDFS

Review Changes

Set the following configuration values for your new hosts. Required values are marked with *

Parameter	Group	Value	Description
Service: HDFS			
NameNode Data Directories* (fs.name.node.name.dir)	HA1	/data/ha1/dfsnn	Determines where on the local file system the NameNode should store the name table (storage). For redundancy, enter a comma-separated list of directories to replicate the name table in all of the directories. Typical values are /data/dfsnn where 1-3.
	HA2	/data/ha2/dfsnn	
	HA3	/data/ha3/dfsnn	
JournalNode Edits Directory* (fs.journalnode.edits.dir)	HA1	/data/ha1/journalnode-edits	Directory on the local file system where NameNode edits are written.
	HA2	/data/ha2/journalnode-edits	
	HA3	/data/ha3/journalnode-edits	

Extra Options

- First initialize the ZooKeeper ZNode for subcluster. Any previous ZNode used for this NameService will be overwritten.
- Clear any existing data present in name directories of Standby NameNode.
Make sure you have backed up any existing data in the name directories of Standby NameNode.
- Clear any existing data present in the JournalNode edits directory for this NameService.
Make sure you have backed up any existing data in the edits directory on all hosts (active, standby, JournalNodes).

Back Continue

11. **Extra Options:** Decide whether Cloudera Manager should clear existing data in ZooKeeper, standby NameNode, and JournalNodes. If the directories are not empty (for example, re-enabling a previous HA configuration), Cloudera Manager will not automatically delete the contents—select to delete the contents by keeping the default checkbox selection. The recommended default is to clear the directories.



Note If chosen not to do so, the data should be in sync across the edits directories of the JournalNodes and should have the same version data as the NameNodes.

12. Click **Continue**.

Cloudera Manager executes a set of commands that will stop the dependent services, delete, create, and configure roles and directories as appropriate, create a nameservice and failover controller, and restart the dependent services and deploy the new client configuration.



Note Formatting of name directory is expected to fail as the directories are not empty as is the case here.

Enable High Availability for HDFS

Progress

Command	Contact	Status	Started at	Ended at
Enable High Availability	ICE2	Finished	May 13, 2015 10:46:36 PM EDT	May 13, 2015 10:56:20 PM EDT

Successfully enabled high availability and automatic failover.

Command Progress

Completed 20 of 20 steps.

- Stop HDFS and its dependent services. Successfully executed command stop on service HDFS.
- Creating roles to enable high availability. Successfully added new NameNode to HDFS as role.
- Deleting the SecondaryNameNode role. The checkpoints directories of the SecondaryNameNode will not be deleted. Successfully deleted role.
- Configuring NameNodes and the HDFS service to enable high availability. Successfully updated config value.
- Check that name directories for the new Standby NameNode either do not exist or are writable and empty. Can optionally clear directories. Confirmed that directory is empty and writable by an HDFS service. [Details](#)
- Check that edit directories for the nameservice either do not exist or are writable and empty. Can optionally clear directories. Confirmed that directory is empty and writable by an HDFS service. [Details](#)
- Initializing high availability state in Zookeeper. Successfully initialized high availability state in Zookeeper from failovercontroller (role). [Details](#)
- Starting the NameNodes. Supervisor returned RUNNING. [Details](#)
- Formatting the name directories of the current NameNode. If the name directories are not empty, this is expected to fail. Command (101) has failed.** [Details](#)
- Initializing shared edit directory of NameNodes. Successfully initialized shared edit directory of the NameNode. [Details](#)
- Starting the NameNode that will be transitioned to active mode (role). Supervisor returned RUNNING. [Details](#)
- Waiting for the Active NameNode to start up. NameNode started responding to NFS successfully. [Details](#)
- Snapshotting Standby NameNode by initializing its name directories. Successfully snapshotting Standby NameNode (role). [Details](#)
- Starting Standby NameNode. Supervisor returned RUNNING. [Details](#)
- Starting the Failover Controller on the host of the Active NameNode. Supervisor returned RUNNING. [Details](#)
- Starting the Failover Controller on the host of the Standby NameNode. Supervisor returned RUNNING. [Details](#)
- Waiting for the Standby NameNode to start up. NameNode started responding to NFS successfully. [Details](#)
- Creating HDFS tmp directory if not already created. HDFS directory /tmp already exists. [Details](#)
- Start HDFS and its dependent services. Successfully executed command start on service HDFS.
- Deploying configurations for clients of services in this cluster. Successfully deployed all client configurations. [Details](#)

Back Continue

13. In the next screen additional steps are suggested by the Cloudera Manager to update the Hue and Hive metastore. Click **finish** for the screen shown below.



Note The following subsections will cover configuring Hue and Hive for HA as needed.

Enable High Availability for HDFS

Congratulations!
Successfully enabled High Availability.

The following manual steps must be performed after completing this wizard:

- Configure the HDFS Web Interface Role of Hue service(s) **Hue** to be an HTTPS role instead of a NameNode. [Documentation](#)
- For each of the Hive service(s) **Hive**, stop the Hive service, back up the Hive Metastore Database to a persistent store, run the service command "update Hive Metastore NameNodes", then restart the Hive services.

14. In the Cloudera Manager, Click on **Home > HDFS > Instances** to see Namenode in High Availability.

Federation and High Availability

[+ Add NameService](#)

Name	Highly Available	Automatic Failover	NameNode	SecondaryNameNode
nameservice1	<input checked="" type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes	<ul style="list-style-type: none"> namenode_m1 (Active) namenode_m2 (Standby) 	

Role Instances [Migrate Roles](#) [Add Role Instances](#) [Role Groups](#) [Roll Back](#)

Configuring Hive Metastore to Use HDFS HA

The Hive metastore can be configured to use HDFS high availability.

1. Go the Hive service.
2. Select **Actions > Stop**.
3. Click **Stop** to confirm the command.
4. Back up the Hive metastore database.
5. Select **Actions > Update Hive Metastore NameNodes** and confirm the command.
6. Select **Actions > Start**.
7. Restart the Hue and Impala services if stopped prior to updating the metastore.

Configuring Hue to Work with HDFS HA

1. Go to the HDFS service.
2. Click the **Instances** tab.
3. Click **Add Role Instances**.
4. Select the text box below the HttpFS field. The Select Hosts dialog displays.
5. Select the host on which to run the role and click **OK**.
6. Click **Continue**.
7. Check the checkbox next to the HttpFS role and select **Actions for Selected > Start**.

8. After the command has completed, go to the Hue service.
9. Click the **Configuration** tab.
10. Locate the HDFS Web Interface Role property or search for it by typing its name in the Search box.
11. Select the HttpFS role you just created instead of the NameNode role, and save your changes.



12. Restart the Hue service.

Configuring Impala to Work with HDFS HA

1. Complete the steps to reconfigure the Hive metastore database, as described in the preceding section. Impala shares the same underlying database with Hive, to manage metadata for databases, tables, and so on.
2. Log in to the admin node (rhel1) and ssh to rhel2. Run command `impala-shell`.
3. Issue the **INVALIDATE METADATA** statement from an Impala shell. This one-time operation makes all Impala daemons across the cluster aware of the latest settings for the Hive metastore database. Alternatively, restart the Impala service.

```
[root@rhel2 ~]# impala-shell
Starting Impala Shell without Kerberos authentication
Connected to rhel2:21000
Server version: impalad version 2.1.2-cdh5 RELEASE (build 92438b7fb62a875041ffd2ab324f408ecf8129e5)
Welcome to the Impala shell. Press TAB twice to see a list of available commands.

Copyright (c) 2012 Cloudera, Inc. All rights reserved.

(Shell build version: Impala Shell v2.1.2-cdh5 (92438b7) built on Tue Feb 24 12:36:33 PST 2015)
[rhel2:21000] > |
```

```
[rhel2:21000] > INVALIDATE METADATA;
Query: invalidate METADATA

Fetched 0 row(s) in 4.63s
```

For more information, see:

http://www.cloudera.com/content/cloudera/en/documentation/core/v5-3-x/topics/impala_invalidate_metadata.html

Configuring Oozie to Use HDFS HA

To configure an Oozie workflow to use HDFS HA, use the HDFS nameservice (nameservice1) instead of the NameNode URI in the <name-node> element of the workflow.

Setting up MapReduce v2 (YARN) HA

1. Log in to the **Cloudera manager** and go to the **YARN service**.
2. Select **Actions > Enable High Availability**. A screen showing the hosts that are eligible to run a standby ResourceManager displays. The host where the current ResourceManager is running is not available as a choice.
3. Select the host (rhe13) where the standby ResourceManager is to be installed, and click **Continue**.

Enable High Availability for YARN (MR2 Included)

Getting Started

This wizard leads you through adding a standby ResourceManager, restarting this YARN (MR2 Included) service and any dependent services, and then re-deploying client configurations.

The screenshot shows a web interface for configuring YARN High Availability. It features a section titled 'ResourceManager Hosts' with two input fields. The first field is labeled 'rhe12 (Current)' and the second field is labeled 'rhe13'. The interface is light gray with a white background for the text.

4. Cloudera Manager proceeds to execute a set of commands that stop the YARN service, add a standby ResourceManager, initialize the ResourceManager high availability state in ZooKeeper, restart YARN, and redeploy the relevant client configurations.

Enable High Availability for YARN (MR2 Included)

Progress

Command	Context	Status	Started at	Ended at
✓ Enable ResourceManager HA	YARN (MR2 Included)	Finished	May 13, 2015 7:00:34 PM EDT	May 13, 2015 7:04:40 PM EDT

Successfully enabled ResourceManager HA.

Command Progress

Completed 4 of 4 steps.

- ✓ Stop yarn and its dependent services
Successfully executed command stop on service YARN (MR2 Included)
- ✓ Add Standby ResourceManager
Successfully added new ResourceManager to YARN (MR2 Included) on rhel3.
- ✓ Start yarn and its dependent services
Successfully executed command start on service Hue
- ✓ Deploy client config for Cluster 1
Successfully deployed all client configurations.
[Details](#)

⏪ Back ⏩ Finish

5. Click **Finish** once the installation is completed successfully.

Changing the log directory

To change the default log from the /var prefix to /data/disk1, follow these steps:

1. Log into the cloudera home page and click **Clusters**.
2. From the configuration drop-down menu select “All Log Directories”.

Figure 185 Changing Log Directory

The screenshot shows the Cloudera Manager Administration interface. The top navigation bar includes 'Home', 'Clusters', 'Hosts', 'Diagnostics', 'Audits', 'Charts', 'Backup', and 'Administration'. The main header is 'Cloudera Security' with tabs for 'Status' and 'Configuration'. The 'Configuration' tab is active, and a dropdown menu is open over the 'All Log Directories' option. The 'Status' table lists various services with their status indicators. To the right, there are four charts: 'Host CPU Usage' showing 0.5% usage, 'Cluster Disk IO' showing disk activity, 'HDFS IO' showing read and write rates, and 'Running MapReduce Jobs' showing zero jobs.

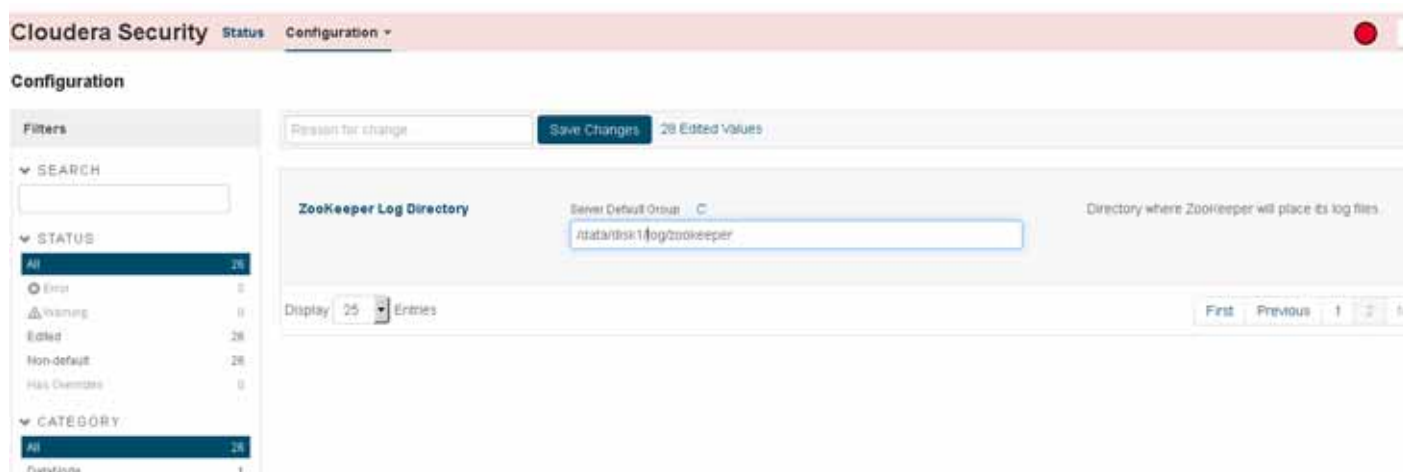
3. Change the path of the log directory to `/data/disk1/log/<service-name>` as shown in the fig below.

Figure 186 Change Path of the Log Directory - Part1

The screenshot displays the Cloudera Security Configuration interface. On the left is a navigation sidebar with 'Filters' and 'GROUP' sections. The main area shows a list of configurations for various services, each with a 'Reason for Change' field, a 'Save Changes' button, and a count of '27 Edited Values'. The configurations are as follows:

Service	Default Group	Current Path	Description
DataNode Log Directory	DataNode Default Group	/data/nfs1/hadoop-hdfs	Directory where DataNode will place its log files.
Falover Controller Log Directory	Falover Controller Default Group	/data/nfs1/hadoop-hdfs	Directory where Falover Controller will place its log files.
HttpFS Log Directory	HttpFS Default Group	/data/nfs1/hadoop-ftp	Directory where HttpFS will place its log files.
JournalNode Log Directory	JournalNode Default Group	/data/nfs1/hadoop-hdfs	Directory where JournalNode will place its log files.
NFS Gateway Log Directory	NFS Gateway Default Group	/data/nfs1/hadoop-hdfs	Directory where NFS Gateway will place its log files.
NameNode Log Directory	NameNode Default Group	/data/nfs1/hadoop-hdfs	Directory where NameNode will place its log files.
SecondaryNameNode Log Directory	SecondaryNameNode Default Group	/data/nfs1/hadoop-hdfs	Directory where SecondaryNameNode will place its log files.
Hive Metastore Server Log Directory	Hive Metastore Server Default Group	/data/nfs1/hog/hive	Directory where Hive Metastore Server will place its log files.
HiveServer2 Log Directory	HiveServer2 Default Group	/data/nfs1/hog/hive	Directory where HiveServer2 will place its log files.
WebHCat Server Log Directory	WebHCat Server Default Group	/data/nfs1/hog/hc-staging	Directory where WebHCat Server will place its log files.
Hue Server Log Directory	Hue Server Default Group	/data/nfs1/hog/hue	Directory where Hue Server will place its log files.
Kerberos Ticket Renewer Log Directory	Kerberos Ticket Renewer Default Group	/data/nfs1/hog/hue	Directory where Kerberos Ticket Renewer will place its log files.
Catalog Server Log Directory	Impala Catalog Server Default Group	/data/nfs1/hog/impala-catalog	Directory where Catalog Server will place its log files.
Impala Daemon Log Directory	Impala Daemon Default Group	/data/nfs1/hog/impalad	Directory where Impala Daemon will place its log files.
Liama Log Directory	Impala Liama ApplicationMaster Default Group	/data/nfs1/hog/impala-liama	Directory where Liama will place its log files.
StateStore Log Directory	Impala StateStore Default Group	/data/nfs1/hog/statestore	Directory where StateStore will place its log files.
JobTracker Log Directory	JobTracker Default Group	/data/nfs1/hog/hadoop-0-20-mapreduce	Directory where JobTracker will place its log files.
TaskTracker Log Directory	TaskTracker Default Group	/data/nfs1/hog/hadoop-0-20-mapreduce	Directory where TaskTracker will place its log files.
Oozie Server Log Directory	Oozie Server Default Group	/data/nfs1/hog/oozie	Directory where Oozie Server will place its log files.
Soir Server Log Directory	Soir Server Default Group	/data/nfs1/hog/soir	Directory where Soir Server will place its log files.
History Server Log Directory	History Server Default Group		The log directory for log files of the role History Server.

Figure 187 Directory Path for Log File



4. Click **Save Changes**.

Conclusion

Hadoop has become a popular data management across all verticals. Cisco UCS Integrated Infrastructure for Big Data and Cisco Application Centric Infrastructure (ACI) along with Cloudera offers a dependable deployment model for enterprise Hadoop that offer a fast and predictable path for businesses to unlock value in big data. This architecture allows using the UCS Manager capabilities in FI for provisioning the servers within a single domain while interconnecting multiple Fabric Interconnect domains with ACI.

The configuration detailed in the document can be extended to clusters of various sizes depending on what application demands as discussed in the Scalability section. Next generation Big Data Infrastructure needs to cater to the emerging trends in Big Data Applications to meet multiple Lines of Business (LOB) SLAs. Cisco UCS Integrated Infrastructure for Big Data and Cisco ACI brings numerous advantages to a Big Data cluster – fewer point of management for the network, enhanced performance, superior failure handling characteristics, unprecedented scalability. Further, ACI paves way to the next generation data center network accelerating innovation with its SDN capabilities in the Big Data space.

Bill of Materials

This section gives the BOM for the 160 node Performance optimized Cluster. See Table 14 for BOM for the master rack, Table 15 for the expansion rack, Table 16 and 17 for software components and Table 18 for Nexus 9k and APIC.

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

Table 14 Bill of Materials for C240M4SX Base Rack

Part Number	Description	Quantity
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-FI6296UP	SMARTNET 24X7X4 UCS 6296UP 2RU Fabric Int/2 PSU/4 Fans	2
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	60
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	108
UCS-BLKE-6200	UCS 6200 Series Expansion Module Blank	6
UCS 6296UP Fan Module	UCS 6296UP Fan Module	8
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	6
RACK-UCS2	Cisco R42610 standard rack w/side panels	1

Table 14 *Bill of Materials for C240M4SX Base Rack*

Part Number	Description	Quantity
RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	2
CON-UCW3-RPDUX	UC PLUS 24X7X4 Cisco RP208-30-U-X Single Phase PDU 2x (Country Specific)	6

Table 15 *Bill of Materials for C240M4SX Expansion Rack*

Part Number	Description	Quantity
UCSC-C240-M4SX	UCS C240 M4 SFF 24 HD w/o CPU, mem, HD, PCIe, PS, railkit w/expndr	64
UCSC-MRAID12G	Cisco 12G SAS Modular Raid Controller	64
UCSC-MRAID12G-2GB	Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	64
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	64
CAB-9K12A-NA	Power Cord 125VAC 13A NEMA 5-15 Plug North America	128
UCSC-PSU2V2-1200W	1200W V2 AC Power Supply for 2U C-Series Servers	128
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	64
UCSC-HS-C240M4	Heat Sink for UCS C240 M4 Rack Server	128
UCSC-SCCBL240	Supercap cable 250mm	64
UCS-CPU-E52680D	2.50 GHz E5-2680 v3/120W 12C/30MB Cache/DDR4 2133MHz	128
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	1024
UCS-HD12T10KS2-E	1.2 TB 6G SAS 10K rpm SFF HDD	1536
UCS-SD120G0KSB-EV	120 GB 2.5 inch Enterprise Value 6G SATA SSD (BOOT)	128
UCSC-PCI-1C-240M4	Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	64



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

SFP-H10GB-CU5M=	10GBASE-CU SFP+ Cable 5 Meter	128
RACK-UCS2	Cisco R42610 standard rack w/side panels	4

RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	8
CON-UCW3-RPDUX	UC PLUS 24X7X4 Cisco RP208-30-U-X Single Phase PDU 2x (Country Specific)	24

Table 16 *Red Hat Enterprise Linux License*

Red Hat Enterprise Linux		
RHEL-2S-1G-3A	Red Hat Enterprise Linux	160
CON-ISV1-RH2S1G3A	3 year Support for Red Hat Enterprise Linux	160

Table 17 *Cloudera License*



Note Choose one of the part numbers.

Part Number	Description	Quantity
UCS-BD-CEBN=	Cloudera Enterprise Basic Edition	160
UCS-BD-CEFN=	Cloudera Enterprise Flex Edition	160
UCS-BD-CEDN=	Cloudera Enterprise Data Hub Edition	160

Table 18 *Bill of Materials for Nexus Device and APIC*

Part Number	Description	Quantity
N9K-C9508-B2	Nexus 9508 Chassis Bundle with 1 Sup, 3 PS, 2 SC, 6 FM, 3 FT	2
N9K-C9396PX	Nexus 9300 with 48p 1/10G SFP+ and 1 uplink module slot	2
N9k-X9736PQ	Spine Line-Card	2
APIC-L1	APIC Appliance	3
N9K POWERCABLES	Power Cables	3
CAB-C13-C14-AC	Power cord, C13 to C14 (recessed receptacle), 10A	4
QSFP-H40G-CU3M	40GBASE-CR4 Passive Copper Cable, 3m	24
N9K-M12PQ	ACI Uplink Module for Nexus 9300, 12p 40G QSFP	3
N9K-C9500-RMK	Nexus 9500 Rack Mount Kit	2
CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	6
N9K-C9500-LC-CV	Nexus 9500 Linecard slot cover	16
N9K-C9500-SUP-C V	Nexus 9500 Supervisor slot cover	2

Table 18 *Bill of Materials for Nexus Device and APIC*

N9K-PAC-3000W-B	Nexus 9500 3000W AC PS, Port-side Intake	6
N9K-SUP-A	Supervisor for Nexus 9500	2
N9K-SC-A	System Controller for Nexus 9500	4
N9K FABRIC	Fabric Module	2
N9300 RACK	Rack Mount Kit	3
N9K-C9300-RMK	Nexus 9300 Rack Mount Kit	3

Appendix A

Troubleshooting

Troubleshooting Sqoop2 process startup

In case Sqoop2 service doesn't startup due to error "Unable to create database", while cluster setup stage, do the following:



Note

This step is not needed if Sqoop2 service comes up fine.

1. Form the node connected to the internet download and Copy derby from Apache Derby.

```
wget
http://apache.mirrors.pair.com//db/derby/db-derby-10.11.1.1/db-derby-10.11.1.1-bin.zip
```

2. Copy to Admin node (rhel1) and unzip the file.

```
scp db-derby-10.11.1.1-bin.zip rhel1:/root/
unzip db-derby-10.11.1.1-bin.zip
```

3. Copy derby.jar and derbyclient.jar to /var/lib/sqoop2/ on Admin Node (rhel1).

```
cd db-derby-10.11.1.1-bin/lib/
cp derby.jar /var/lib/sqoop2/
cp derbyclient.jar /var/lib/sqoop2/
```

4. Copy derby.jar and derbyclient.jar to /var/lib/sqoop2/ on all Nodes.

```
cd /var/lib/sqoop2/

# ls
derbyclient.jar derby.jar mysql-connector-java.jar postgresql-9.0-801.jdbc4.jar
tomcat-deployment

clush -a -b -c derby*
clush -a -b ls /var/lib/sqoop2/
```

5. Change the link to derby.jar and derbyclient.jar in parcels.

```
clush -a -b rm -f
/opt/cloudera/parcels/CDH-5.3.2-1.cd5.3.2.p0.10/lib/sqoop2/webapps/sqoop/WEB-INF/lib/
derby-10.8.2.2.jar
```

```
clush -a -b ln -s /var/lib/sqoop2/derby.jar
/opt/cloudera/parcels/CDH-5.3.2-1.cdh5.3.2.p0.10/lib/sqoop2/webapps/sqoop/WEB-INF/lib/
derby-10.8.2.2.jar
```

```
clush -a -b ls -l
/opt/cloudera/parcels/CDH-5.3.2-1.cdh5.3.2.p0.10/lib/sqoop2/webapps/sqoop/WEB-INF/lib/
derby-10.8.2.2.jar
```

6. Retry cluster setup operation.

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.

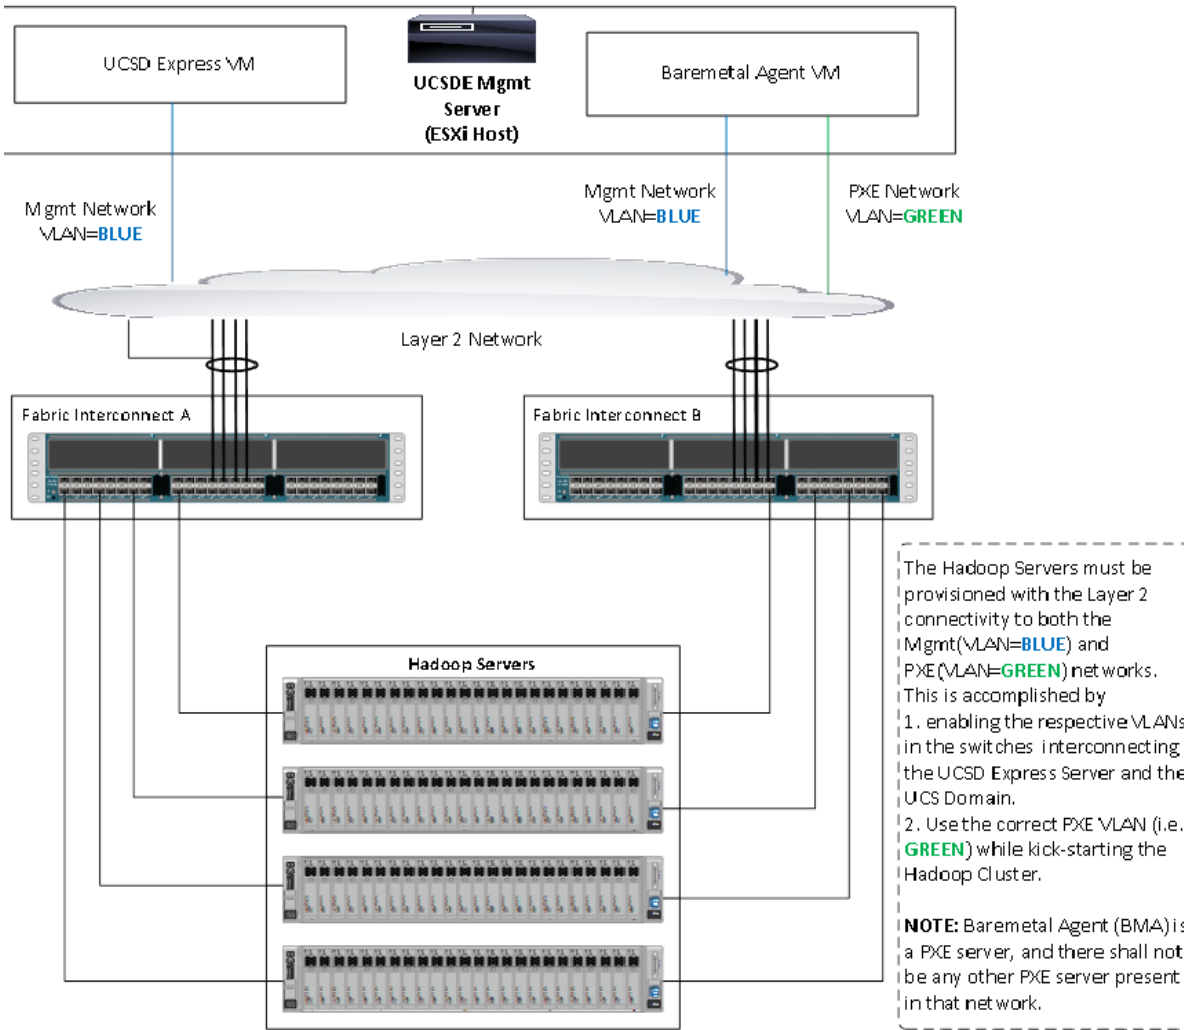
UCSD Express Management Server Configuration

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 188 UCSD Express Management Server that lives outside the UCS Domain

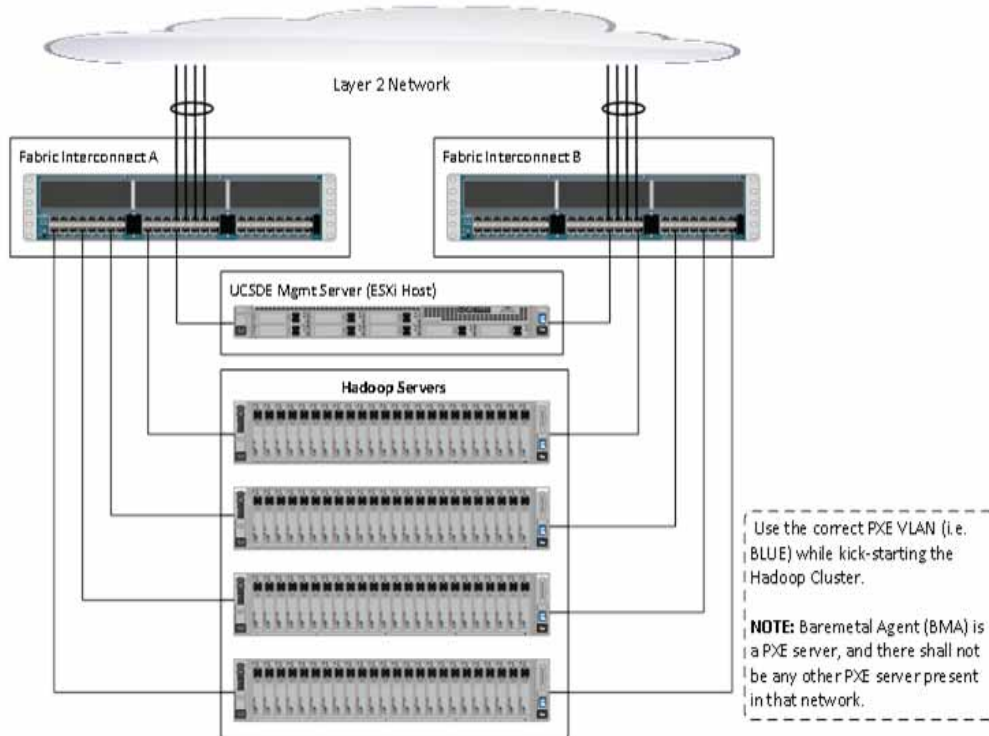


2. The UCSD Express Management server is hosted on a C220 M4 rack server that is connected to and managed by the same UCS Domain. This is the method used in this document.

Figure 189 UCSD 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.

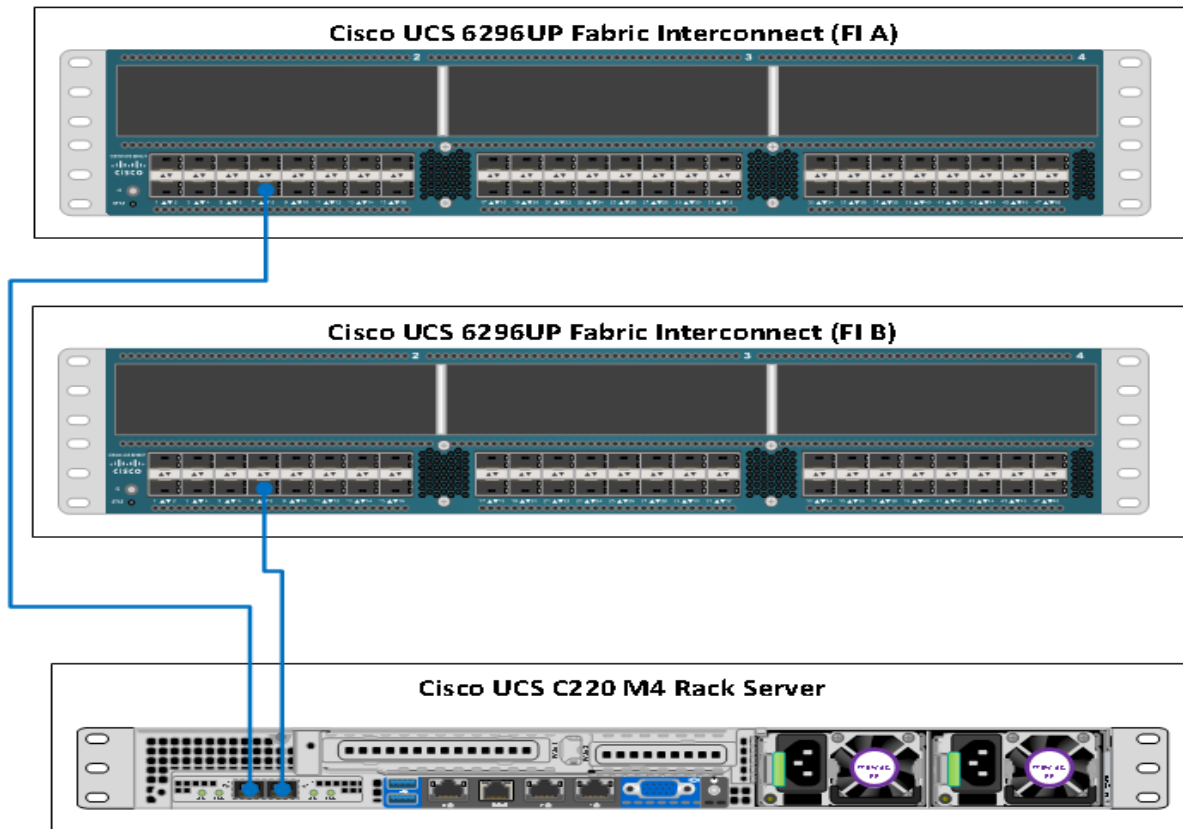


UCSD Management Server Cabling

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 190 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. Its service profile need to be

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

Table 19 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.



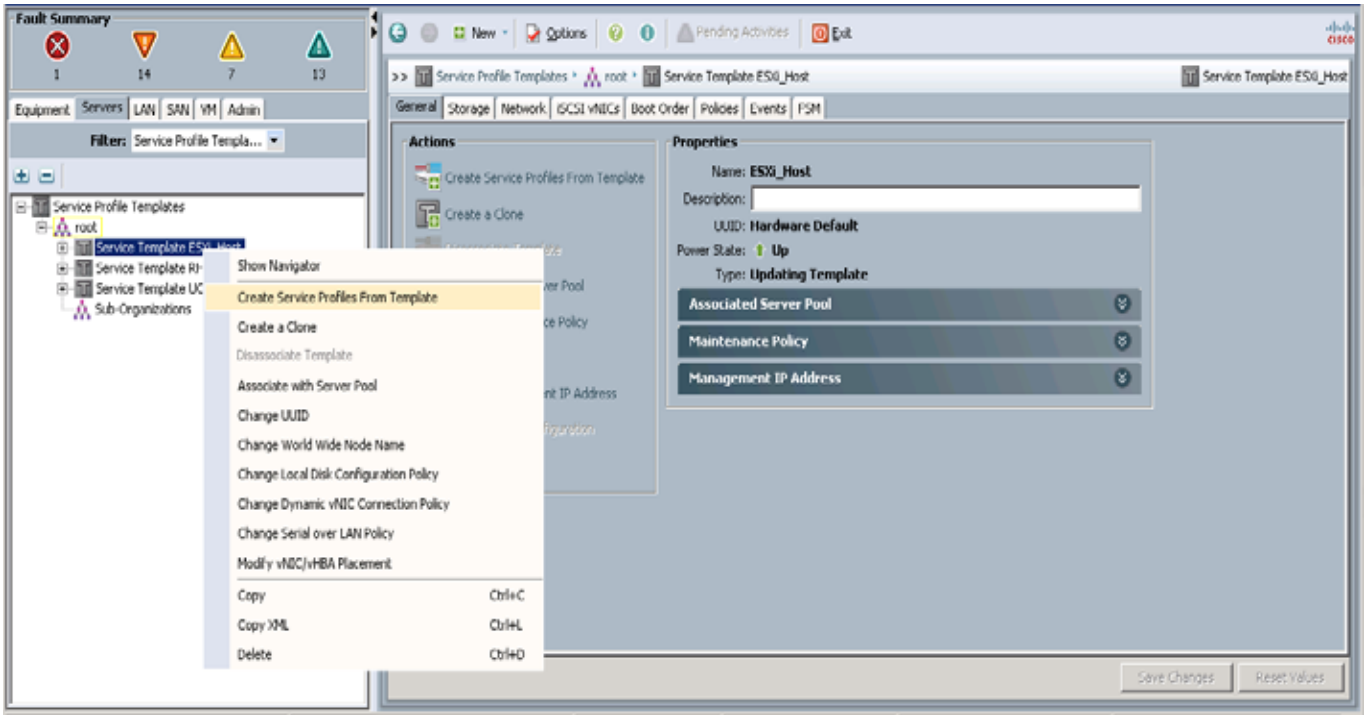
Note Create the service profile template named as ESXi_Host with two vNICs as shown in the above table. For vNIC eth0, select default VLAN as the native VLAN, and for vNIC eth1, select PXE VLAN (vlan85_PXE) as the native VLAN.

Creating Service Profile from the Template

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

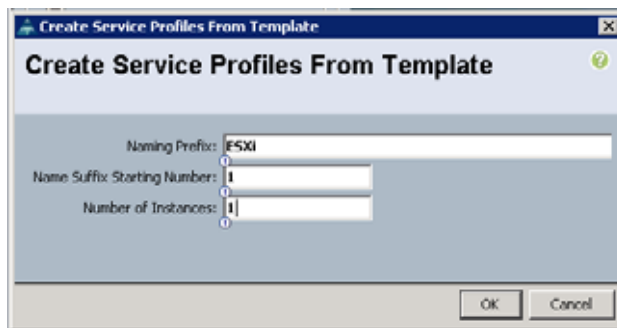
1. Go to Service Profile **Templates** > **root**.
2. Right-click **Service Profile Templates ESXi_Host**.
3. Select **Create Service Profiles From Template**.

Figure 191 Creating Service Profiles from Template



4. The Create Service Profile from Template window appears.

Figure 192 Selecting Name and Total number of Service Profiles



Association of the Service Profiles will take place automatically.

Installing VMware vSphere ESXi 5.5

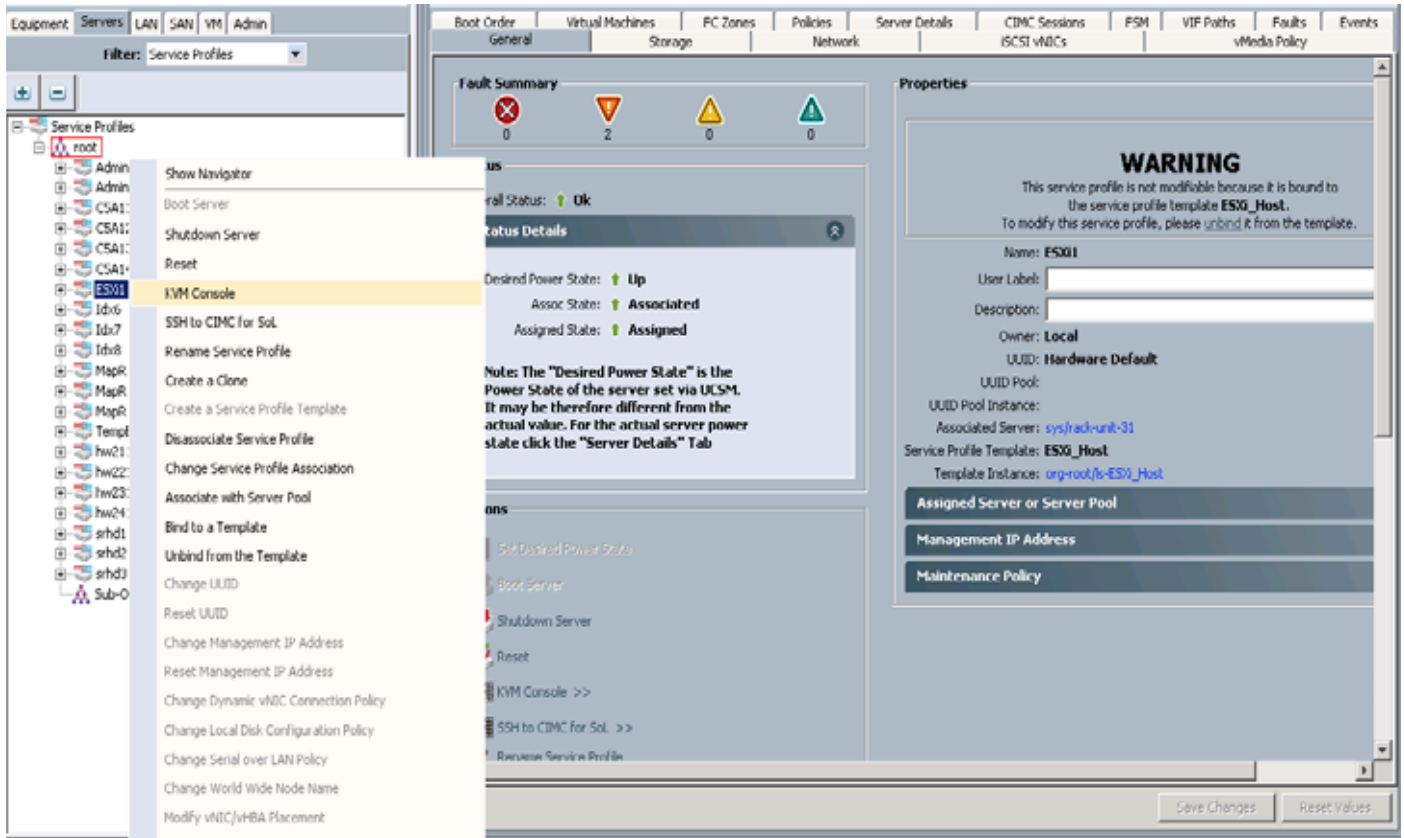
The following section provides detailed procedures for installing VMware vSphere ESXi 5.5.

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.

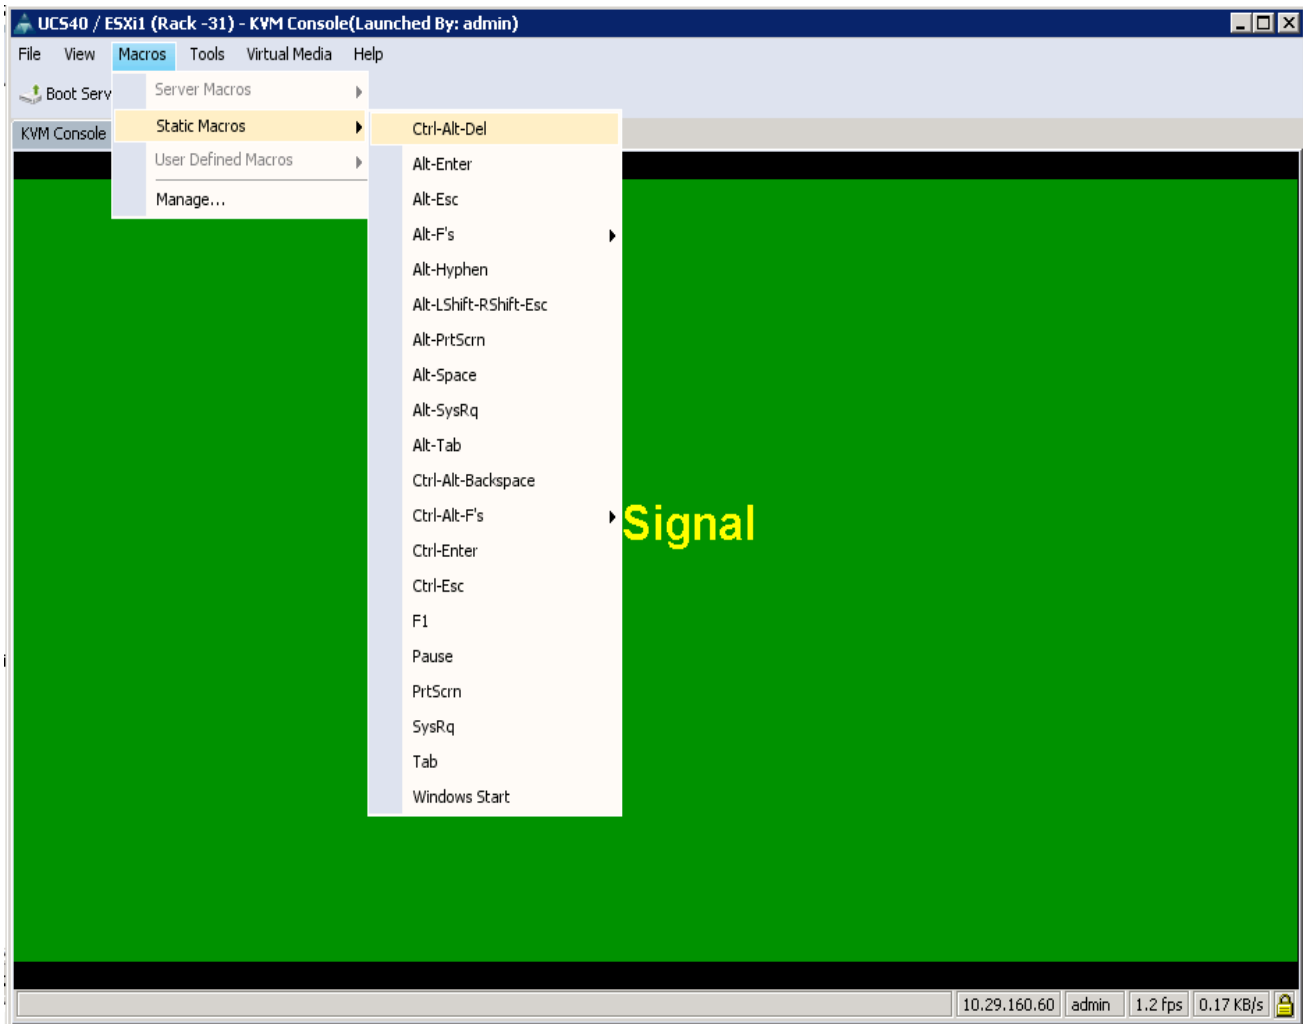
2. Select the **Servers** tab.
3. In the navigation pane expand Service Profiles.
4. Right click on the newly created service profile ESXi1 and select KVM Console.

Figure 193 *Selecting KVM Console*



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

Figure 194 Sending Ctrl-Alt-Del to Reset the Server



- 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 195 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

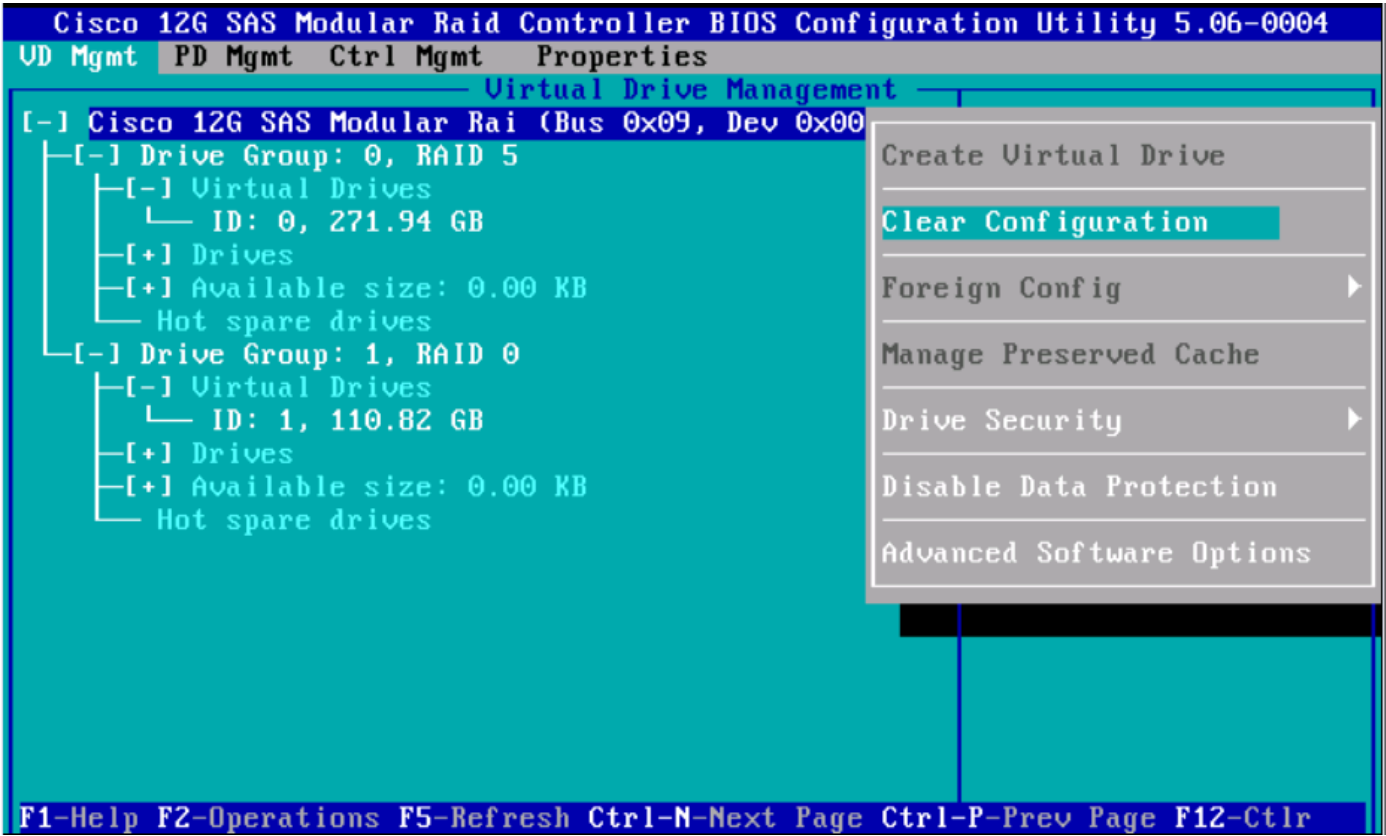
ID LUN VENDOR      PRODUCT                REVISION      CAPACITY
-- -- -
      LSI          Cisco 12G SAS Modular Raid 4.250.00-3632 2048MB
20 0 SEAGATE        ST9146803SS           0004          140014MB
21 0 SEAGATE        ST9146803SS           0004          140014MB
22 0 TOSHIBA       MBF2300RC  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

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

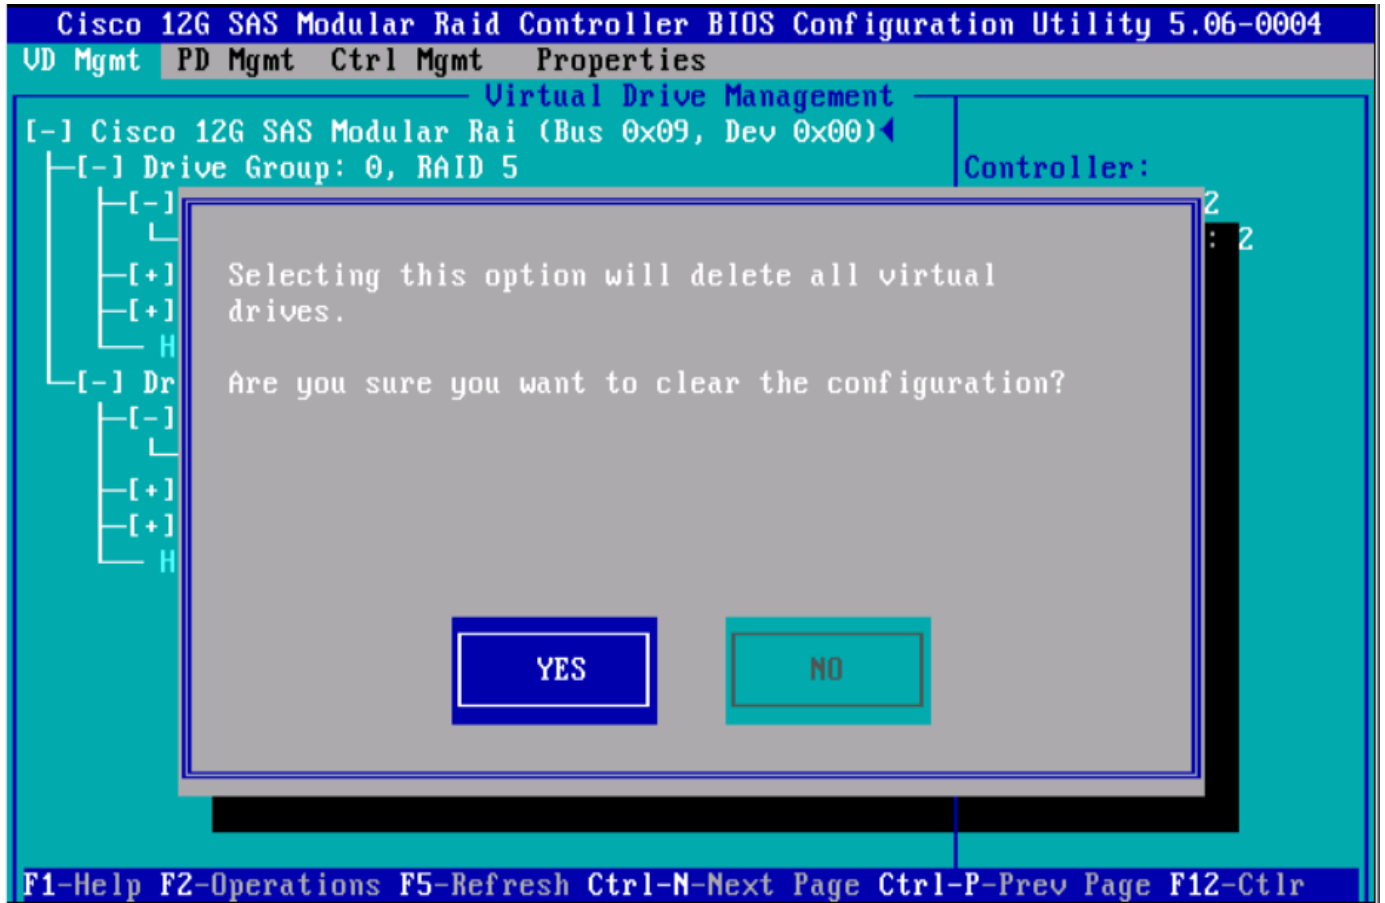
0 Virtual Drive(s) found on the host adapter.
0 Virtual Drive(s) handled by BIOS
Press <Ctrl><R> to Run MegaRAID Configuration Utility
10.29.160.60 admin 0.4 fps 0.001 KB/s

```

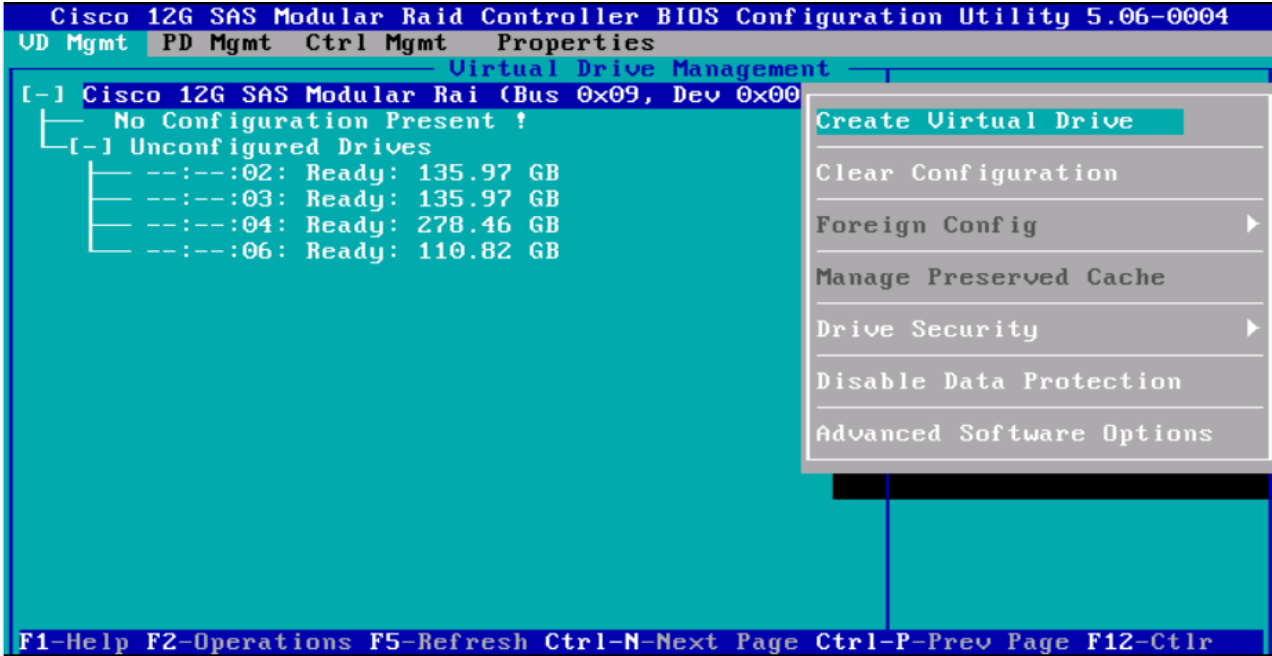
7. In the MegaRAID configuration utility, under VD Mgmt section, use the arrow keys to select the Cisco 12G SAS Modular RAID (Bus 0xNN, Dev 0xNN) line item.
8. Press the function key **F2**.
9. Select the option Clear Configuration, and press **ENTER**.



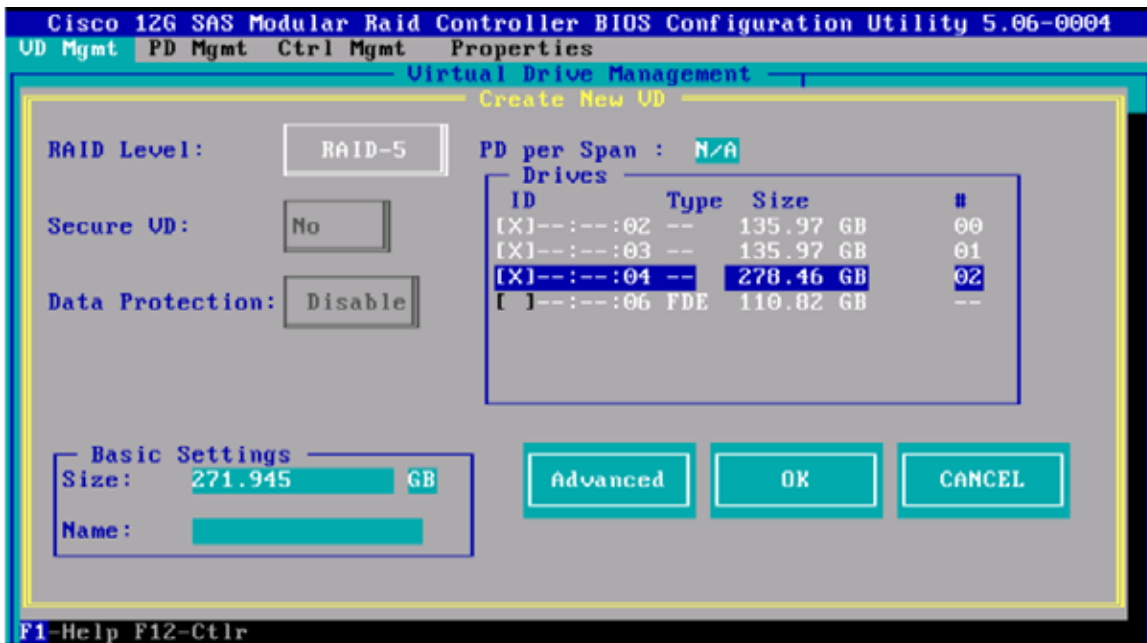
10. To the question Are you sure you want to clear the configuration? click **YES** and press **ENTER** key.



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**.



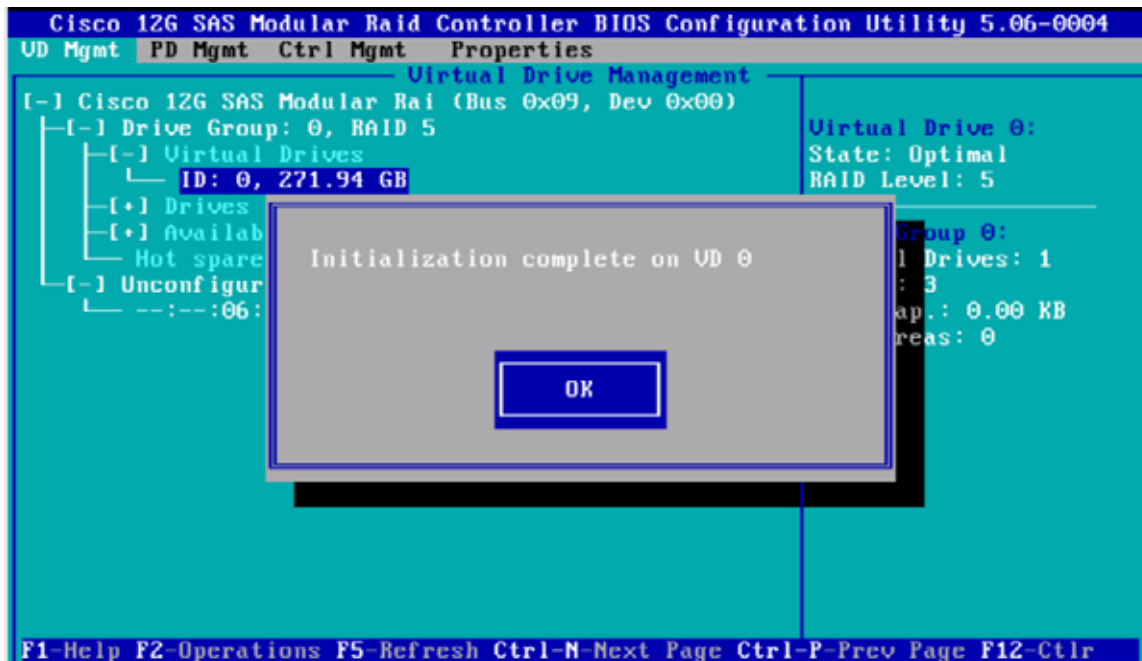
13. In the RAID Level: press **ENTER** and choose **RAID-5**.
14. In the Drives section, press **SPACE** on the desired number of drives to select them to be part of the RAID group. Use the Up and Down arrow keys to navigate.



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

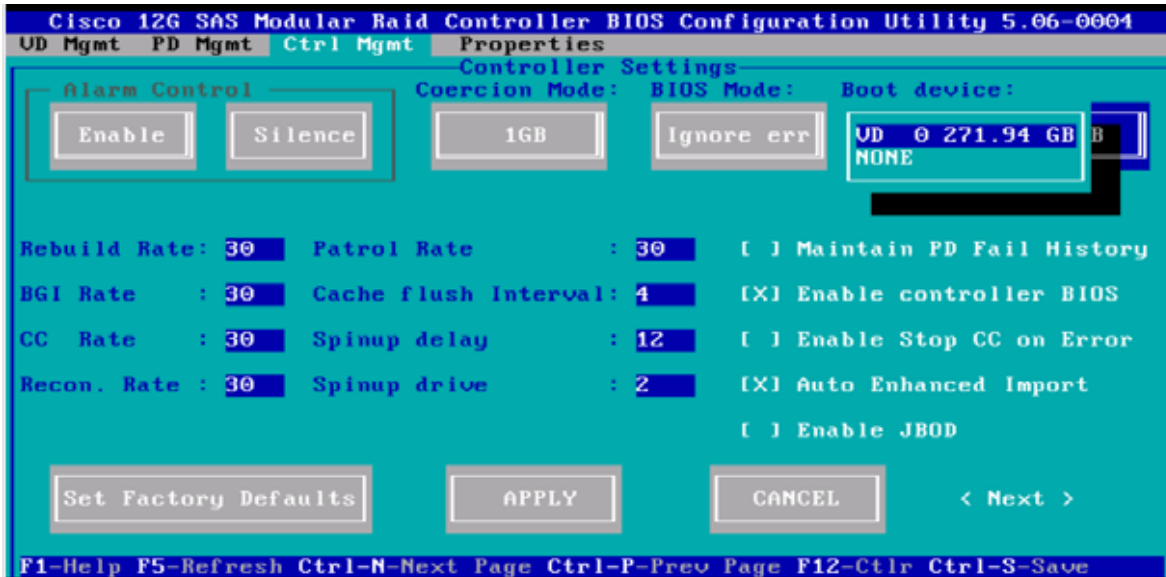


17. After the initialization is complete, the following message appears. Press **OK** to continue.

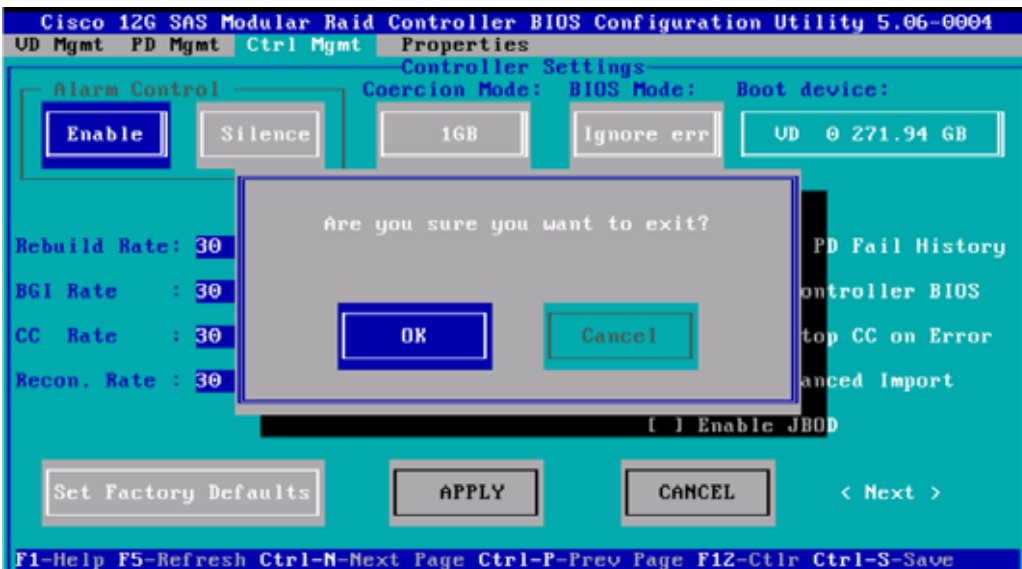


18. Press **Ctrl-N** twice to navigate to the Ctrl Mgmt screen.


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.

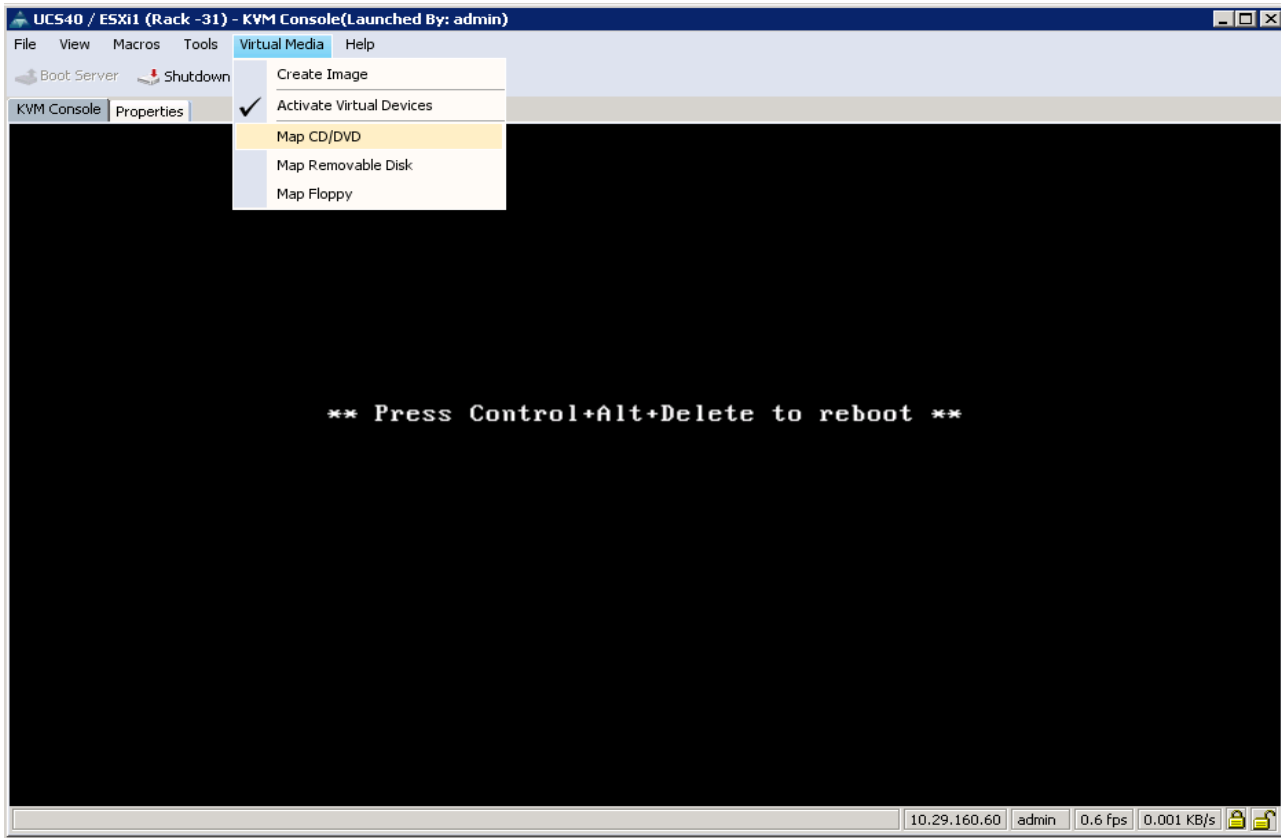


23. In the KVM window, select the Virtual Media menu.
24. Click the Activate Virtual Devices found in the right hand corner of the Virtual Media selection menu.



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

Figure 196 Mapping the CD/DVD Virtual Media



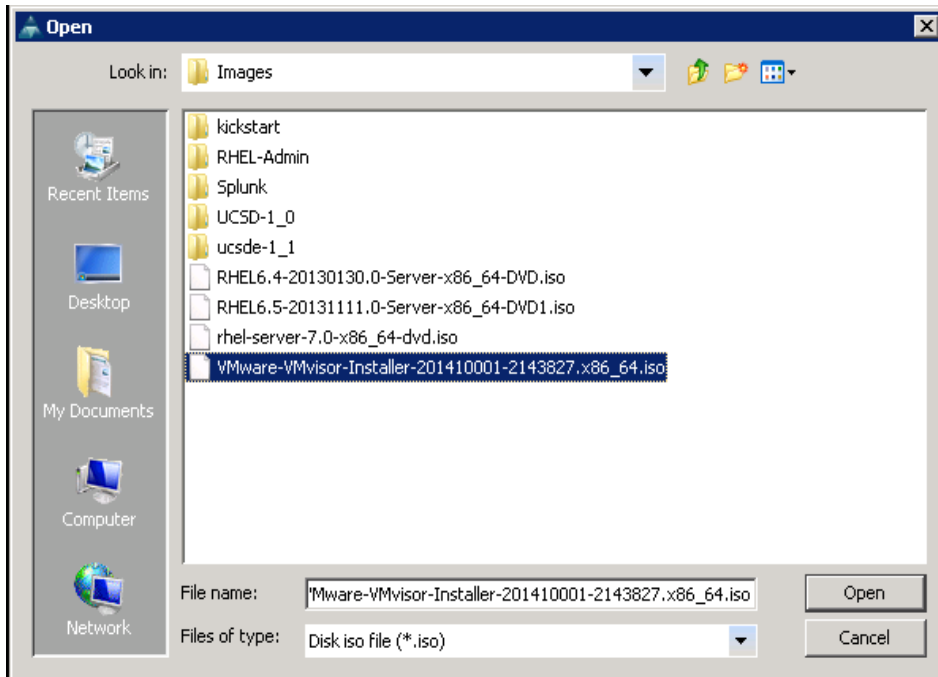
26. Browse to the VMware vSphere ESXi 5.5 installer ISO image file.



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 197 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 198 ESXi Standard Boot Menu



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

Figure 199 Loading the ESXi Installer




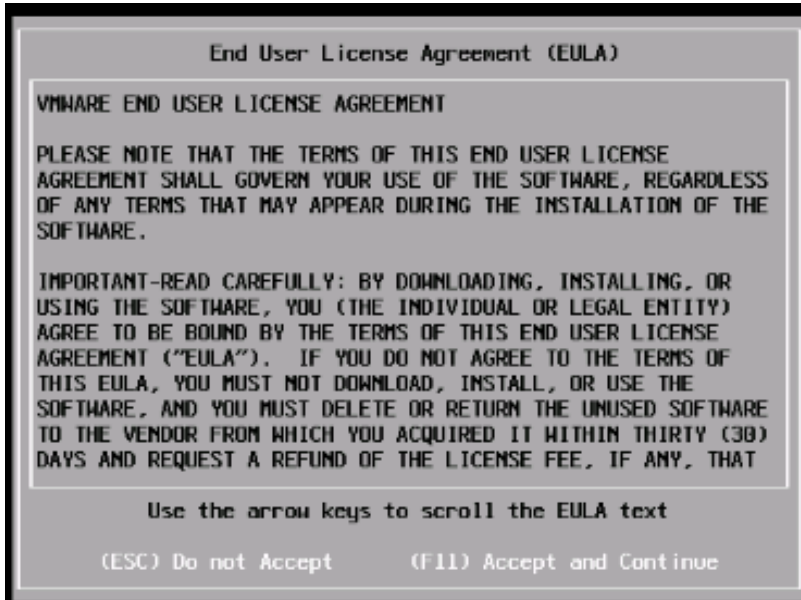
Figure 200 VMWare ESXi Installation screen



33. Press ENTER to continue.

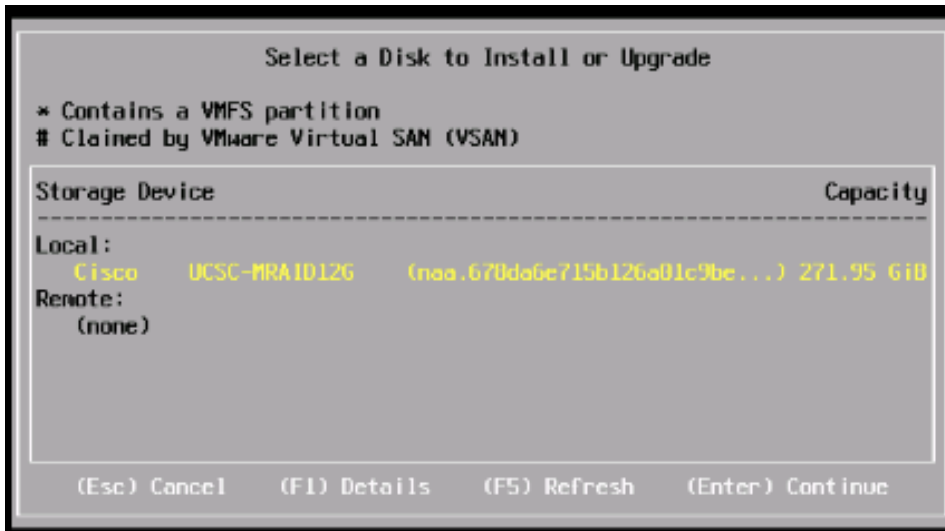
34. Press **F11** to accept End user License Agreement (EULA) and continue.

Figure 201 Accept End User License Agreement (EULA)



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

Figure 202 Selecting the Storage Device for installing the ESXi operating system.



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

Figure 203 Choose the Keyboard layout



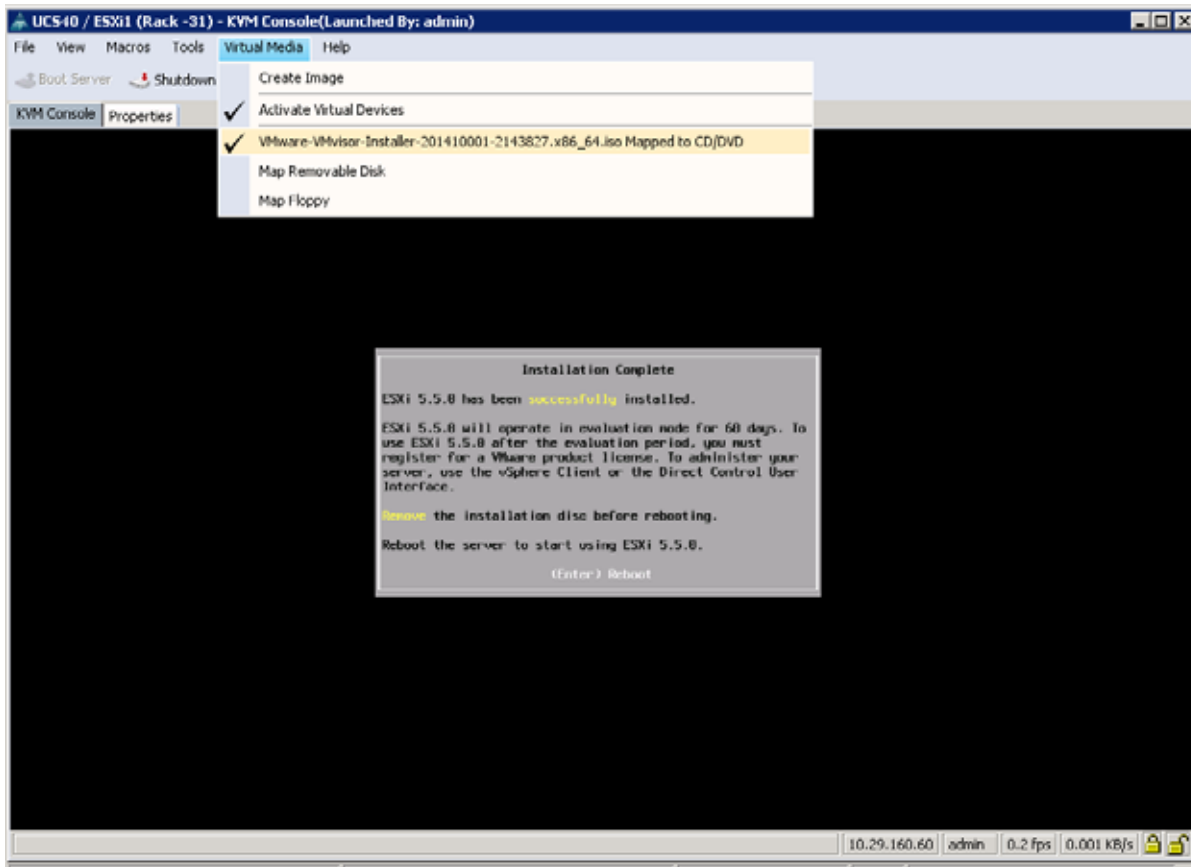
37. Choose the root password and confirm it. Press **ENTER** to continue.

Figure 204 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 205 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

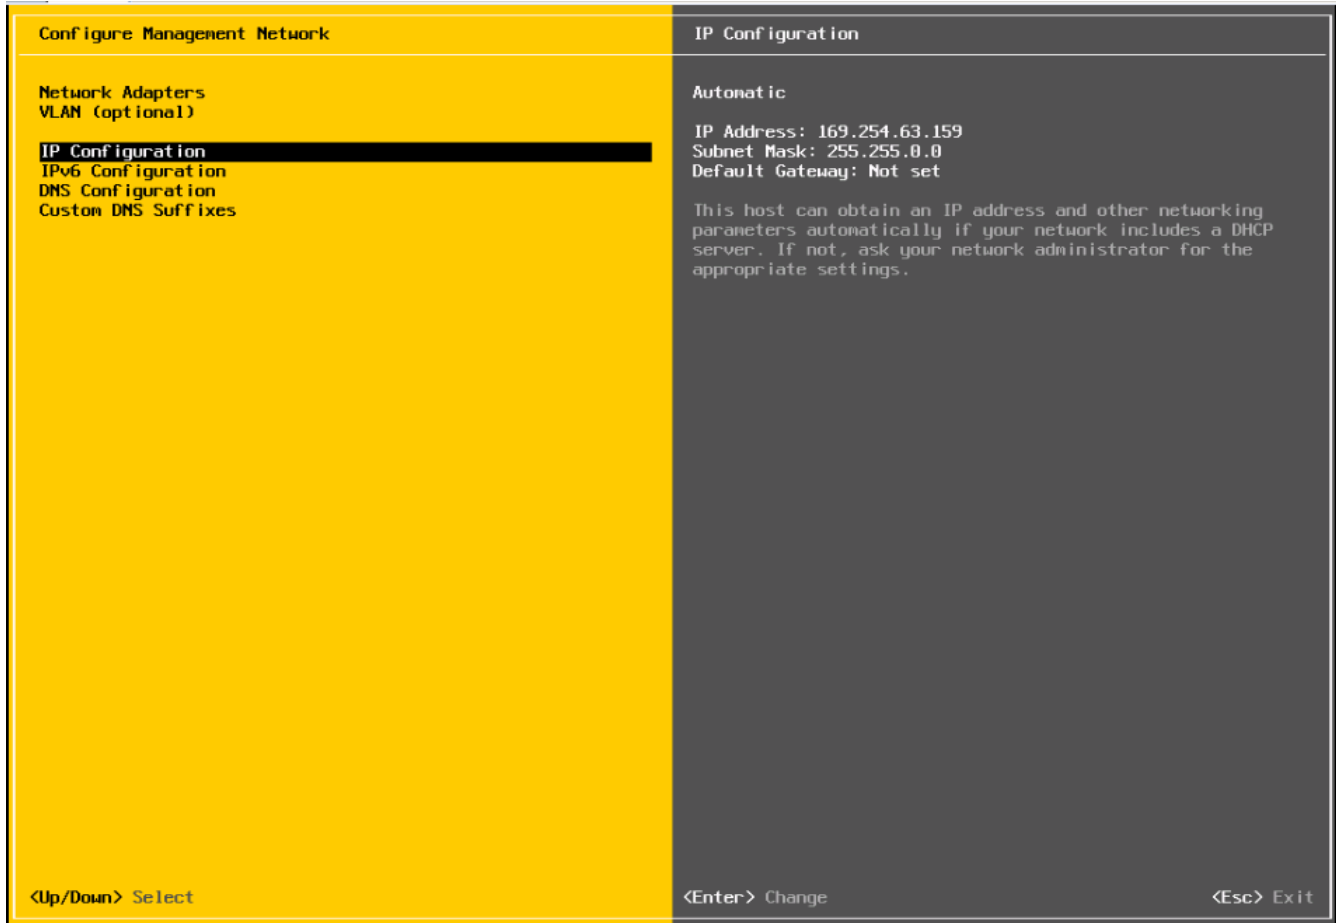
1. Once the server reboots, press **F2** to log on.
2. Enter username as root, and the password chosen above.

Figure 206 VMWare ESXi initial screen as seen via the KVM Console



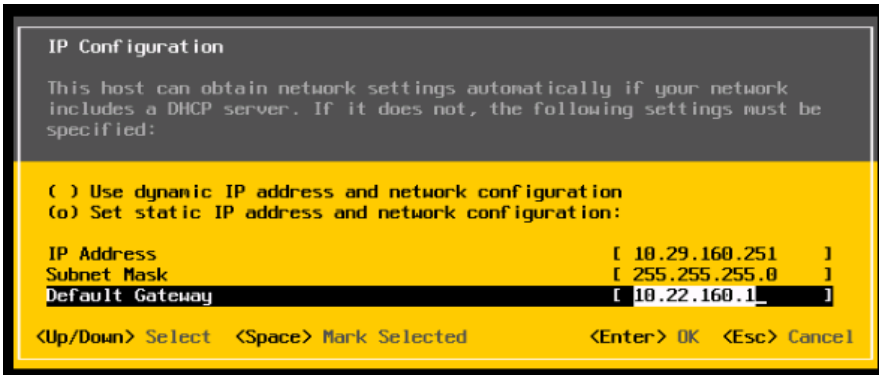
3. Press **F2** to continue
4. Select Configure Management Network, and press **ENTER**.
5. Select **IP Configuration** option.

Figure 207 Enter the IP configuration option of the Management Network



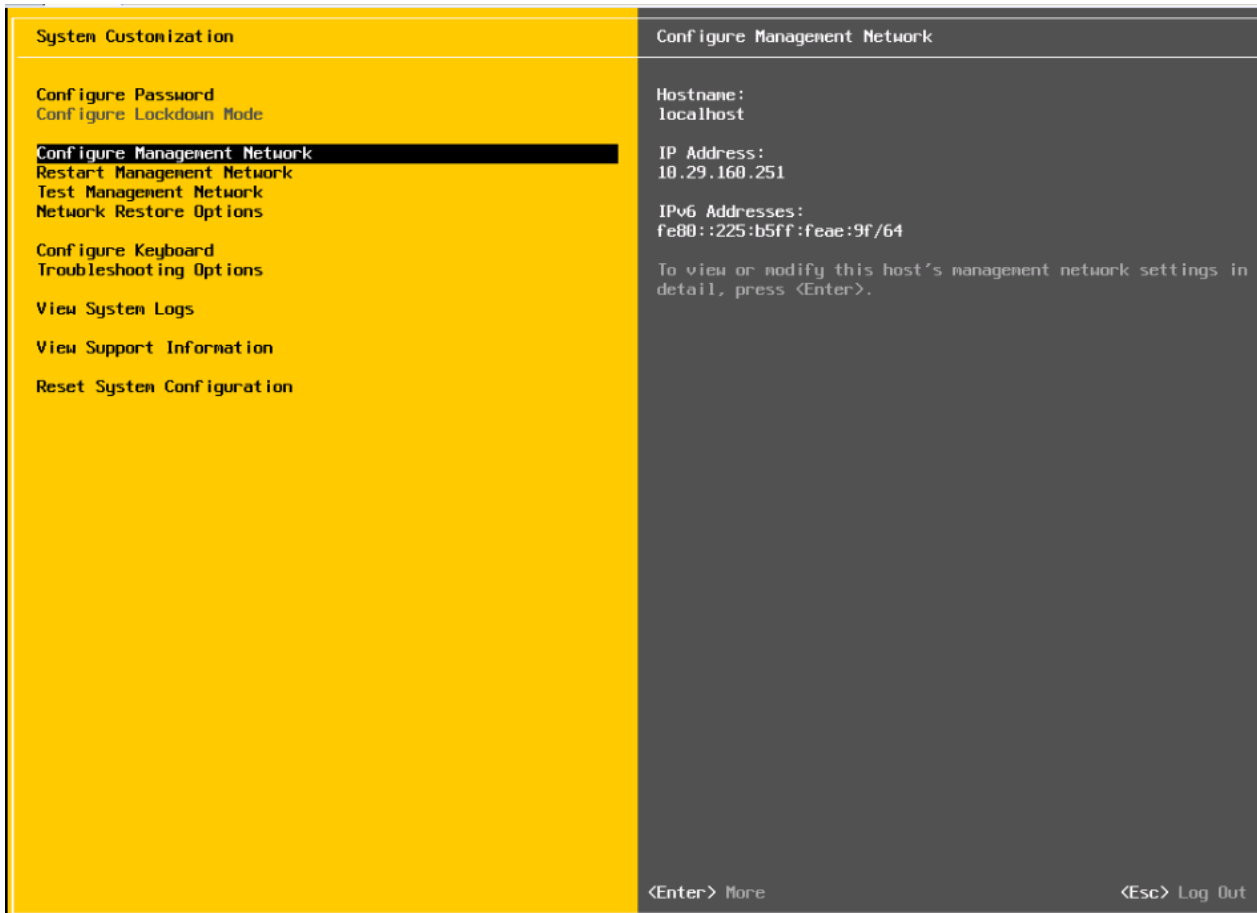
6. Press **ENTER** to continue.
7. Use the Up/Down arrow keys to highlight the Set Static IP address and network configuration option, and press **SPACE** key to select it.
8. Enter the static IP address, Subnet Mask and Default Gateway.

Figure 208 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 209 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 210 Accessing the ESXi web interface



Figure 211 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.

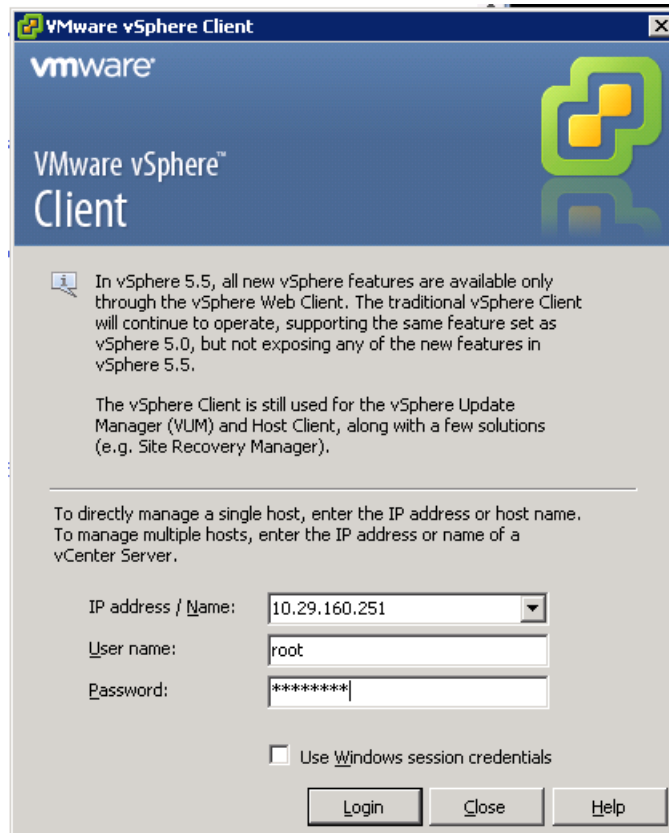
Figure 212 Installing the vSphere Client software



Configuring the vSphere ESXi hypervisor

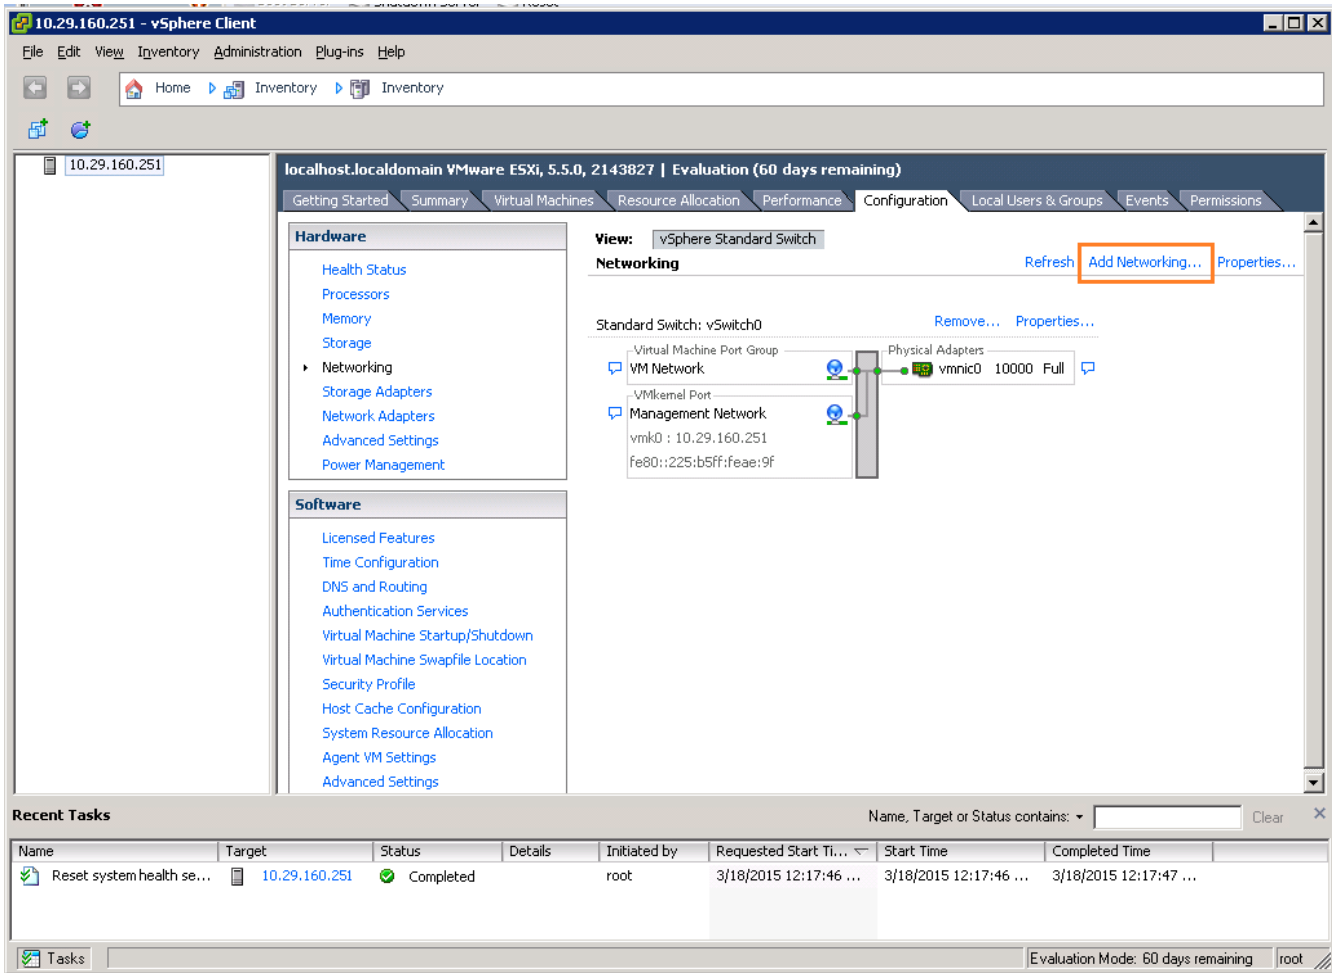
1. After the installation is complete, launch the VMWare vSphere client.
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click on **Login** to continue.

Figure 213 Logging into the ESXi using vSphere Client



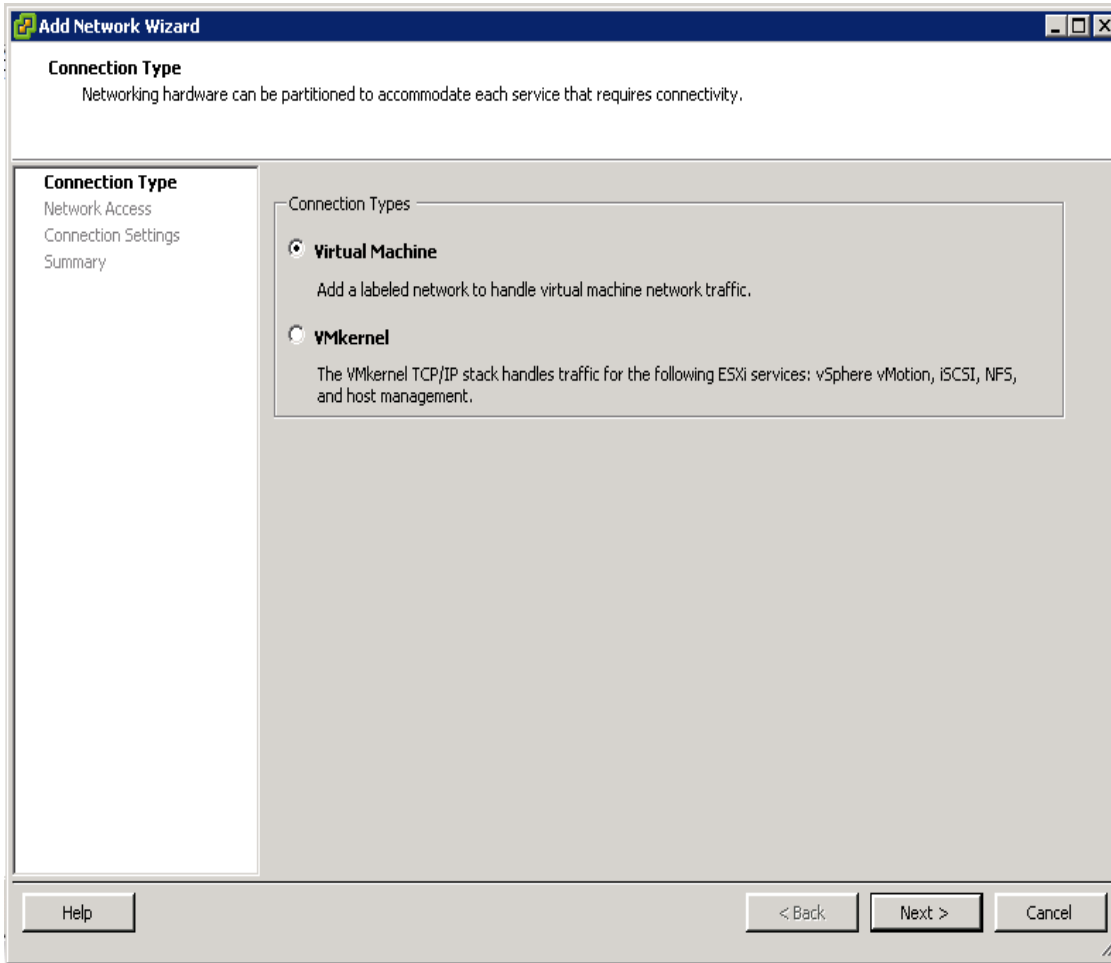
4. In the vSphere Client, click on the Configuration tab on the right, and within the Hardware section, click on Networking.
5. Click on Add Networking link on the upper right hand side.

Figure 214 vSphere Client Networking screen



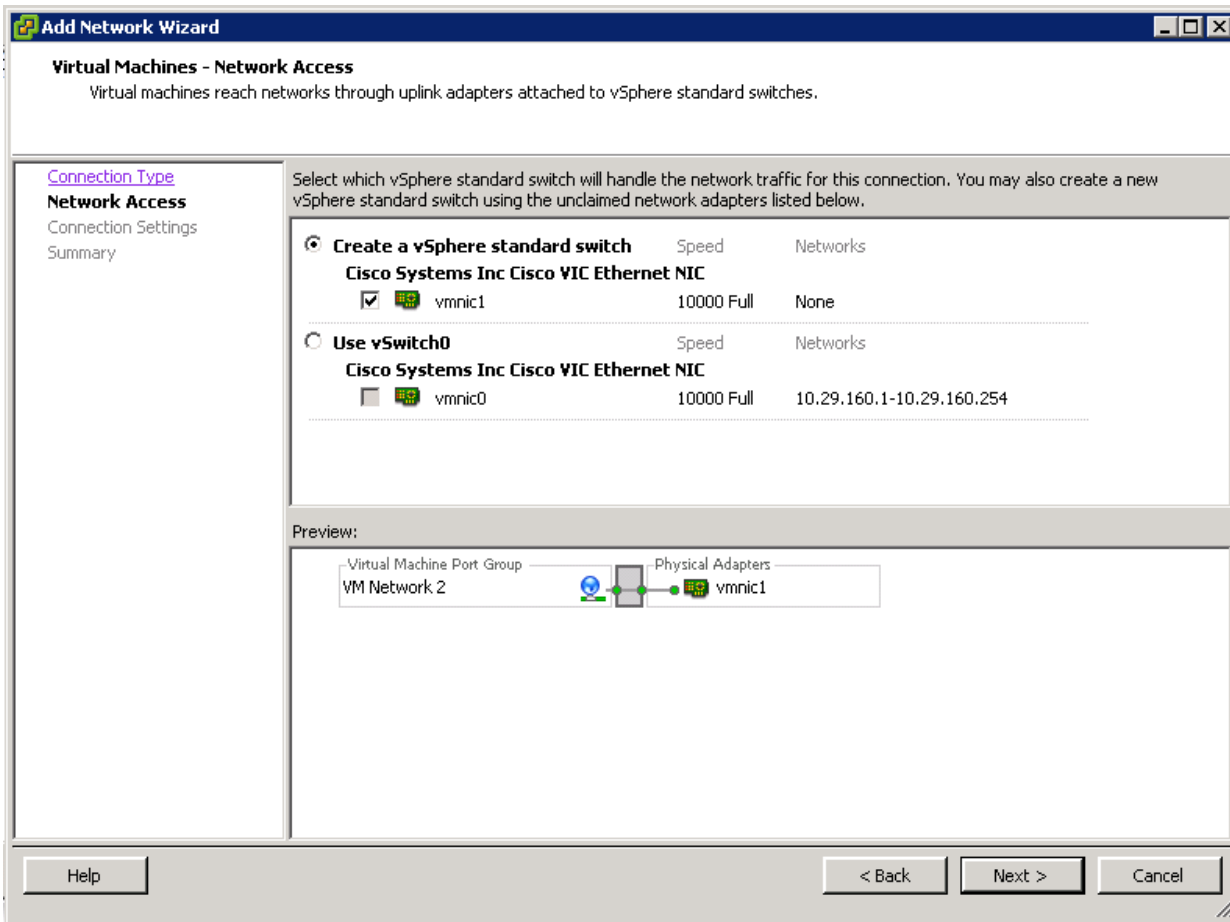
6. In the Add Networking dialog box, click the **Virtual Machine** radio button and click **Next**.

Figure 215 Adding a new Virtual Machine Network



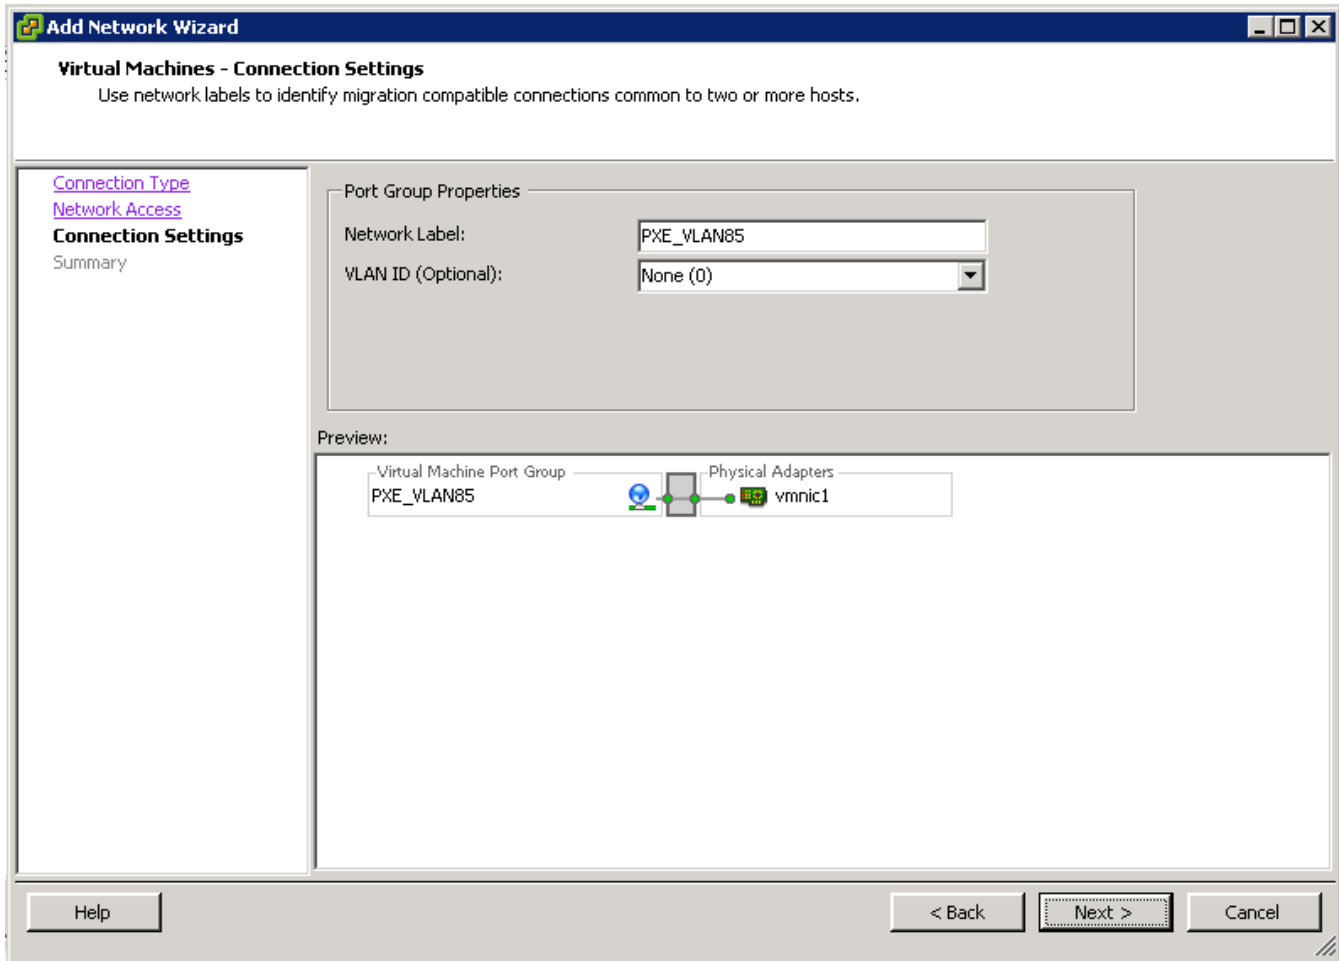
7. Click the **Create a vSphere standard switch** radio button and make sure that the checkbox next to `vmnic1` is checked.
8. Click **Next**.

Figure 216 Creating a new vSphere Standard Switch



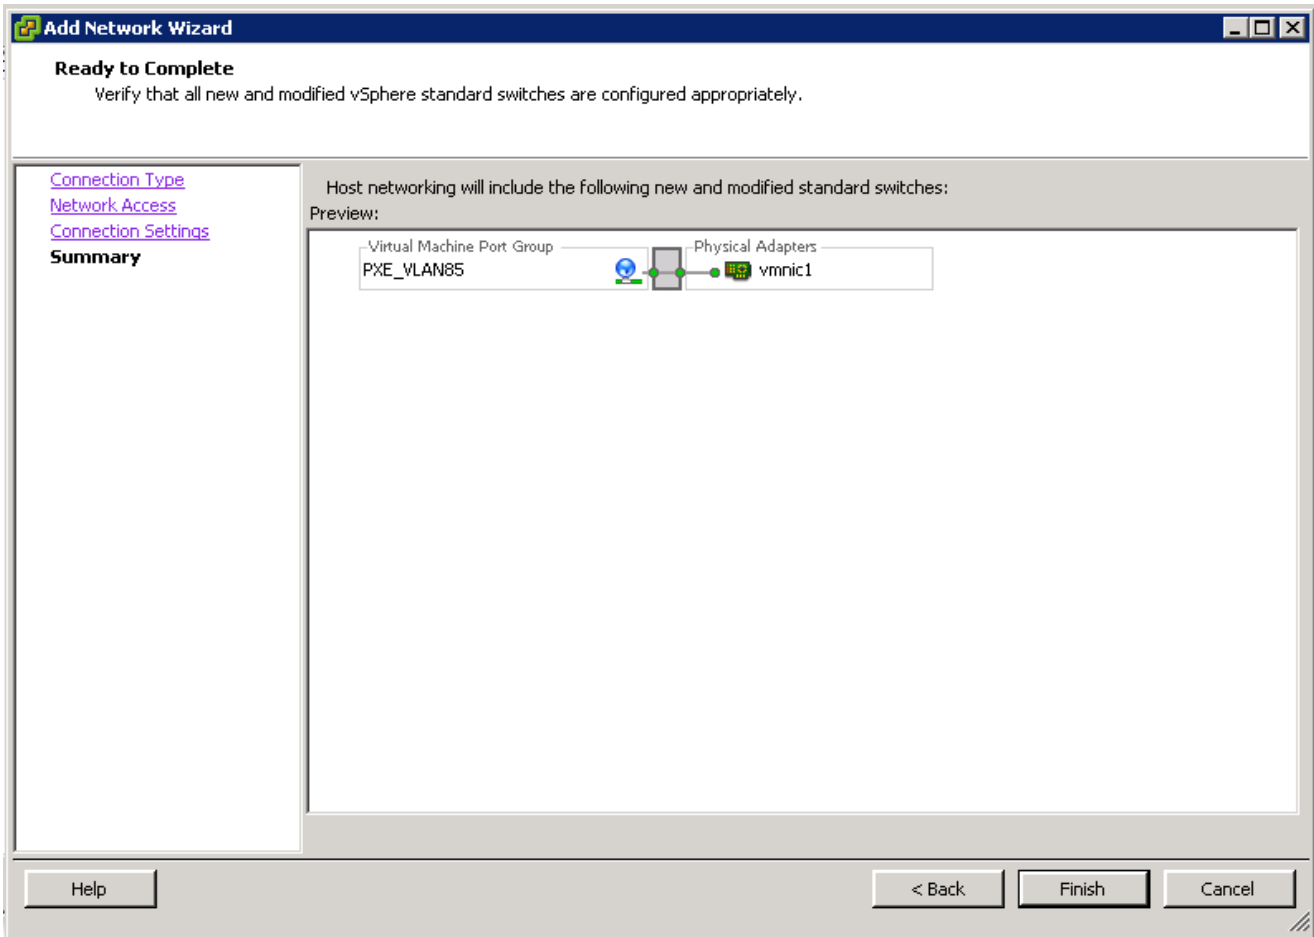
9. In the Port Group Properties, change the Network Label field to PXE_VLAN85.
10. Leave the VLAN ID(Optional) field as None(0).
11. Click **Next**.

Figure 217 Creating the Port Group for the PXE VLAN



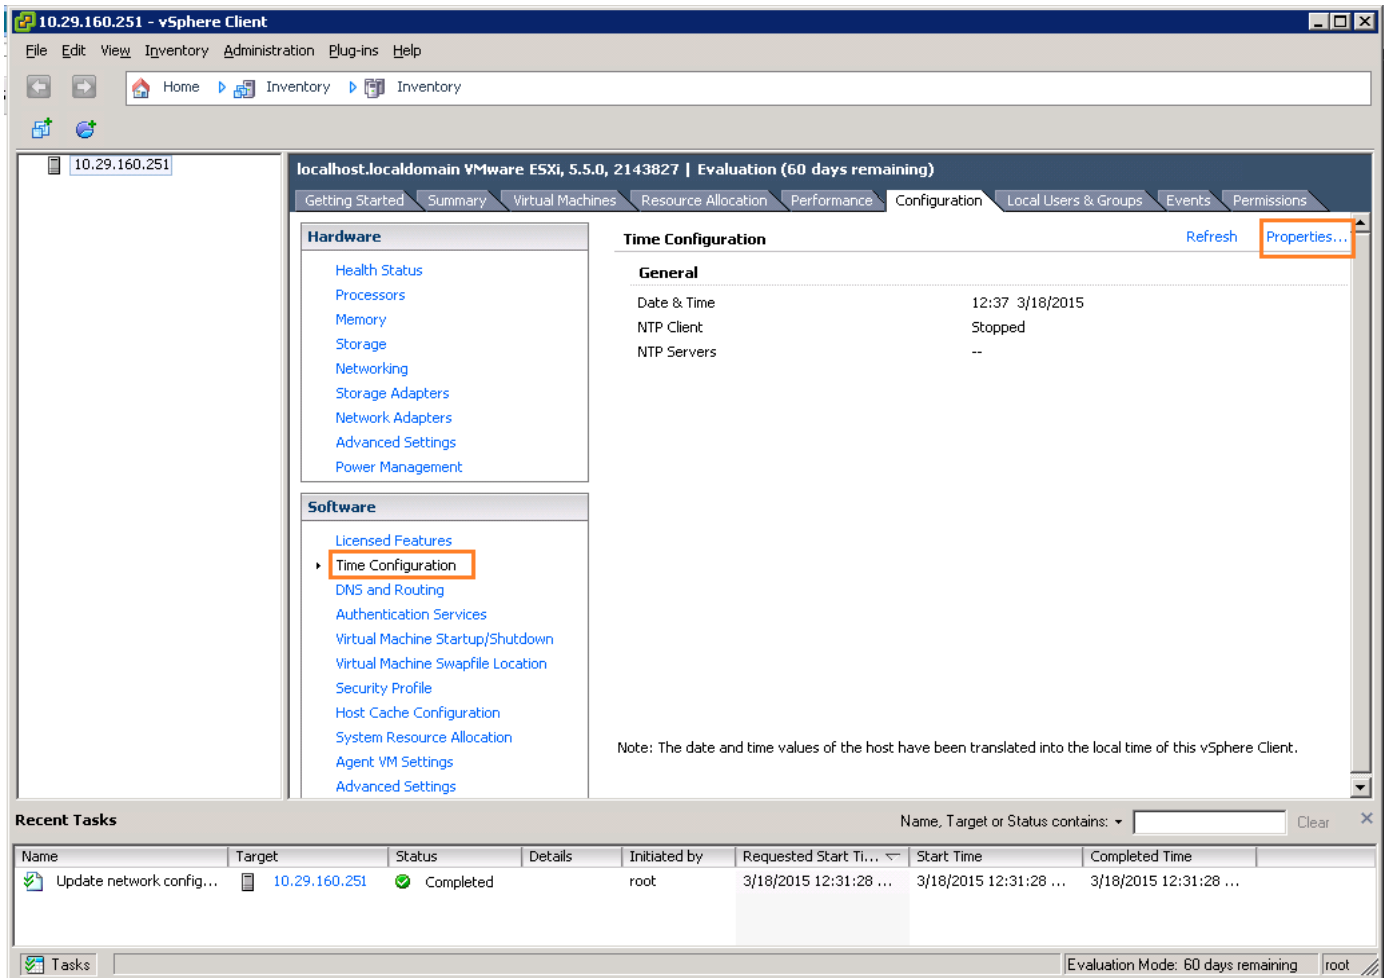
12. Click **Finish** to complete adding the Network.

Figure 218 Verify the Created vSphere Standard Switches



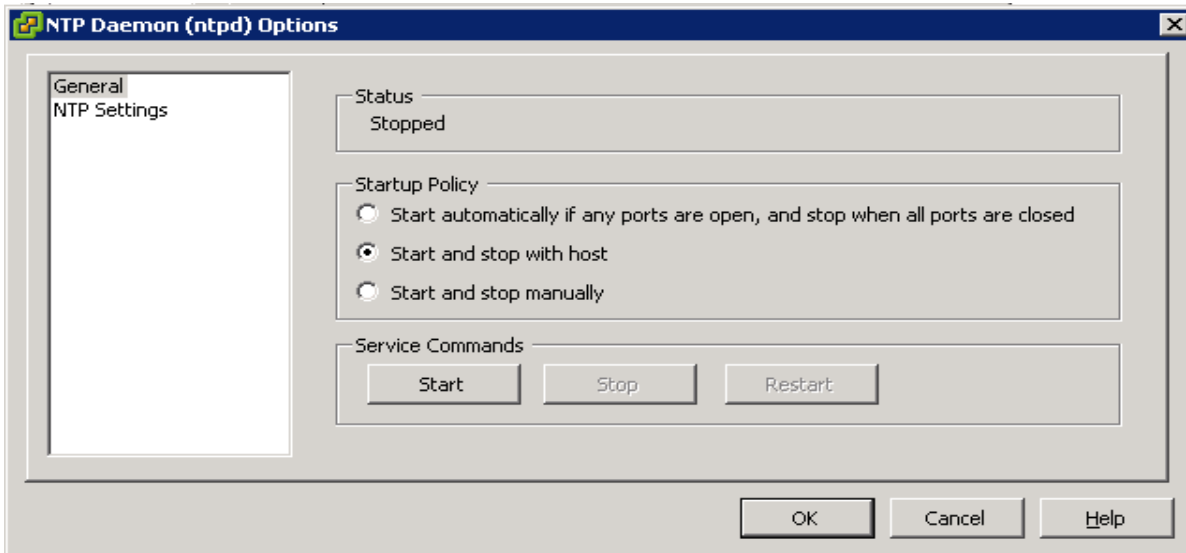
13. Click on the **Time Configuration** under the Software section.
14. Click on **Properties** at the upper right hand corner.

Figure 219 Enabling the NTP Client on the ESXi



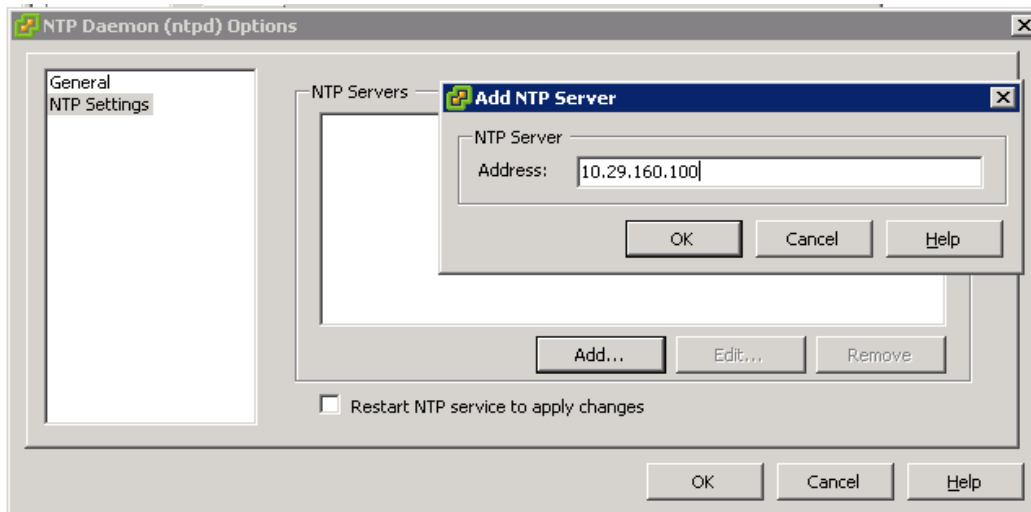
15. In the NTP Daemon (ntpd) Options dialog box, click **Options**.
16. Click on the **General** options.
17. Click to select the start and stop with **host** radio button.

Figure 220 NTP Daemon



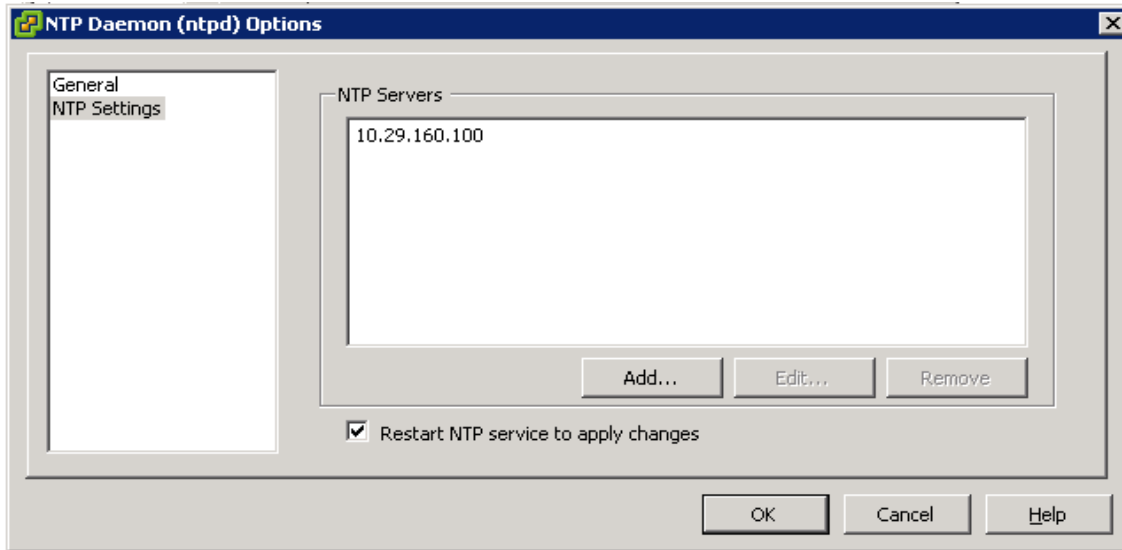
18. Click on NTP Settings option.
19. Click on **Add** button to add the NTP server's IP address.
20. Press **OK** to continue.

Figure 221 Adding a new NTP Server to the ESXi NTP Settings



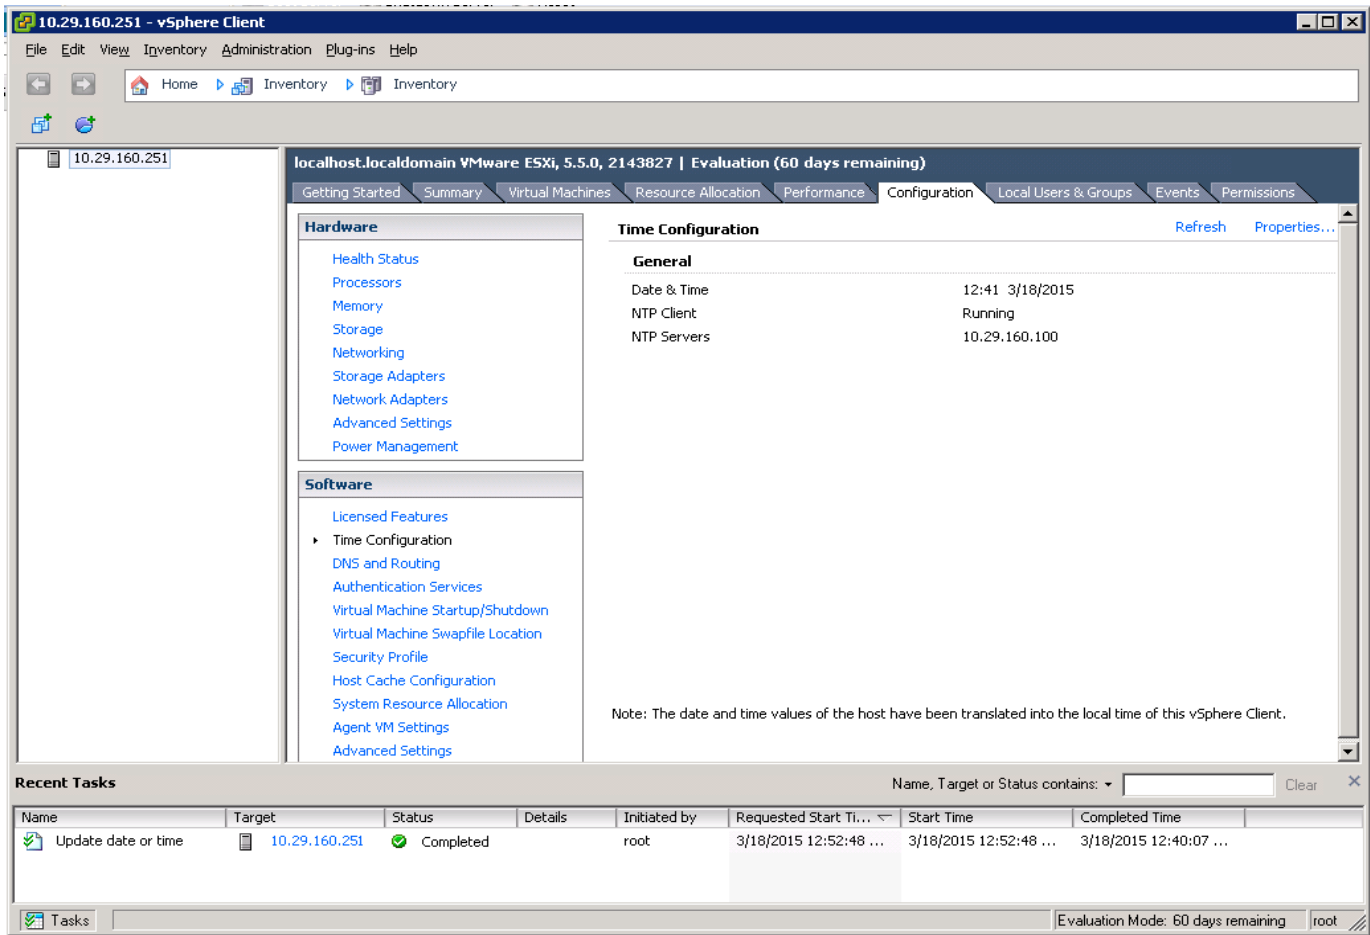
21. In the next screen, verify the IP-address in the NTP Servers list.
22. Click on the checkbox **Restart NTP service to apply changes**.
23. Press the button **OK** twice to complete the time configurations.

Figure 222 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 223 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 20 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&rellifecycle=&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&rellifecycle=&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&rellifecycle=&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&rellifecycle=&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&rellifecycle=&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 224 Cisco UCS Director Express for Big Data 1.0 Download Page

The screenshot displays the Cisco software download page for UCS Director Express for Big Data 1.0. The page includes a navigation menu, a search bar, and a list of software releases. The releases are organized into a table with columns for File Information, Release Date, and Size. The fourth release, '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' dropdown menu.

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	Download Add to cart
Cisco UCSD Express For Big Data Patch (Patch needs to be applied on 1.0. MD5 Checksum - ca44a9a25057af5072acafaf7fc7d933)	19-NOV-2014	1.76 MB	Download Add to cart
Cisco UCSD Express Hotfix for Bash Code Injection Vulnerability (Bash Shells hock - CVE-2014-6271, CVE-2014-7169) Note: Patch has README that explains h ow to apply this patch	06-OCT-2014	1.82 MB	Download Add to cart
Cisco UCS Director Express for Big Data 1.0 (OVF Appliance) MD5 Checksum (8d6cb7dc36107ca5c1f93a9faf69d49c)	05-SEP-2014	2663.09 MB	Download Add to cart
Cisco UCS Director Express for Big Data BMA Update Package MD5 Checksu m - 517fa2a881b8cab6dff0c3ad17a1cc9b	05-SEP-2014	343.95 MB	Download Add to cart

- 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 225 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

Add Devices
 Add Notification

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"/>

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 226 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.

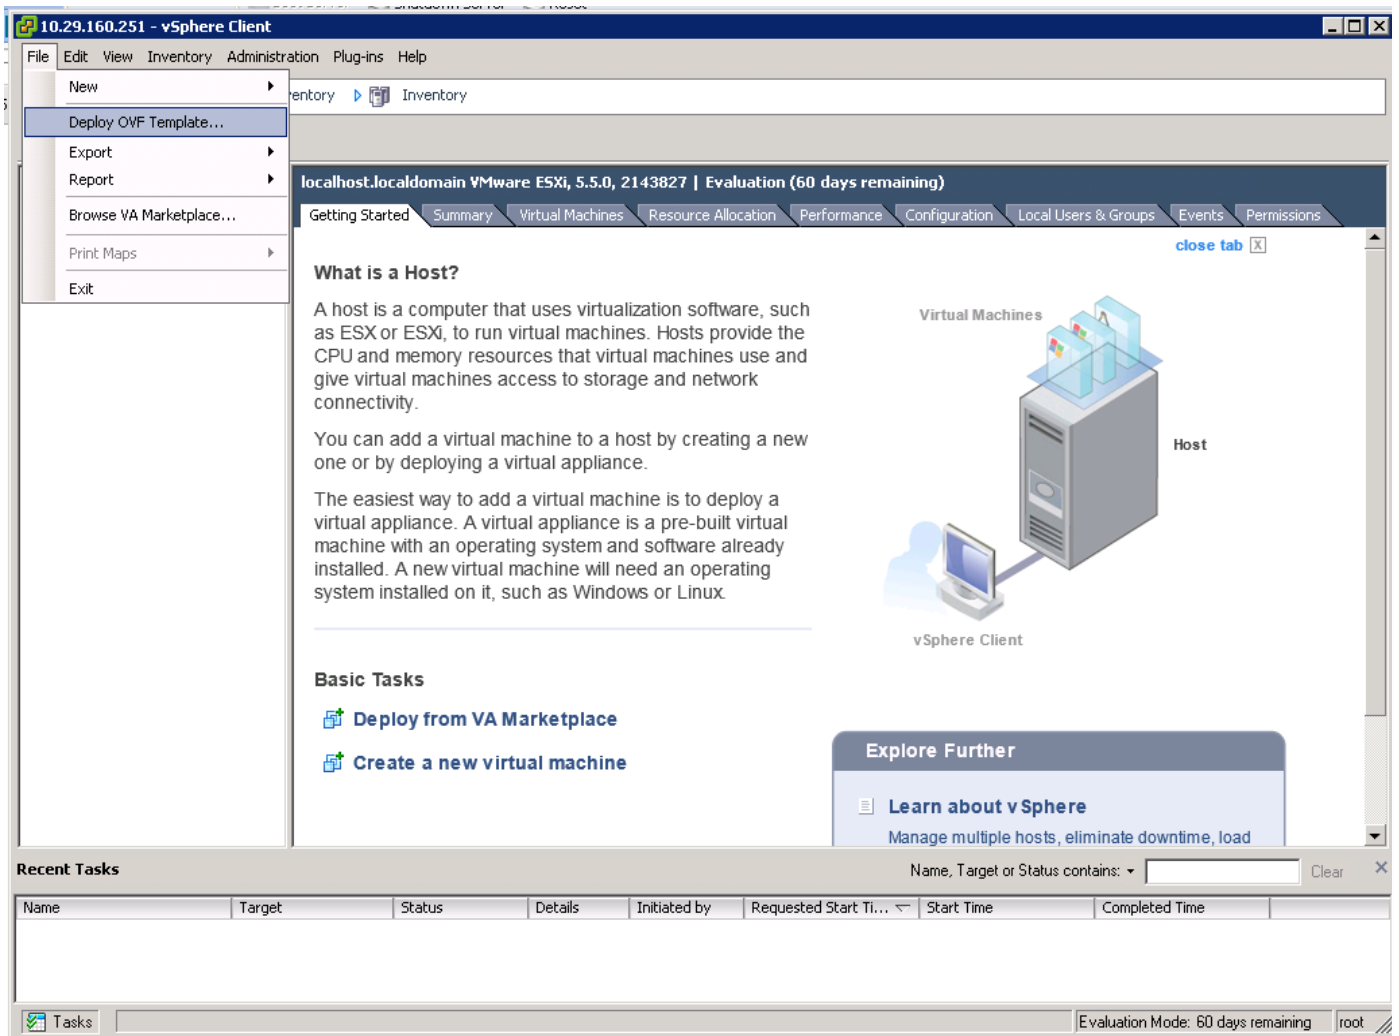
Installing Cisco UCS Director Express for Big Data

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

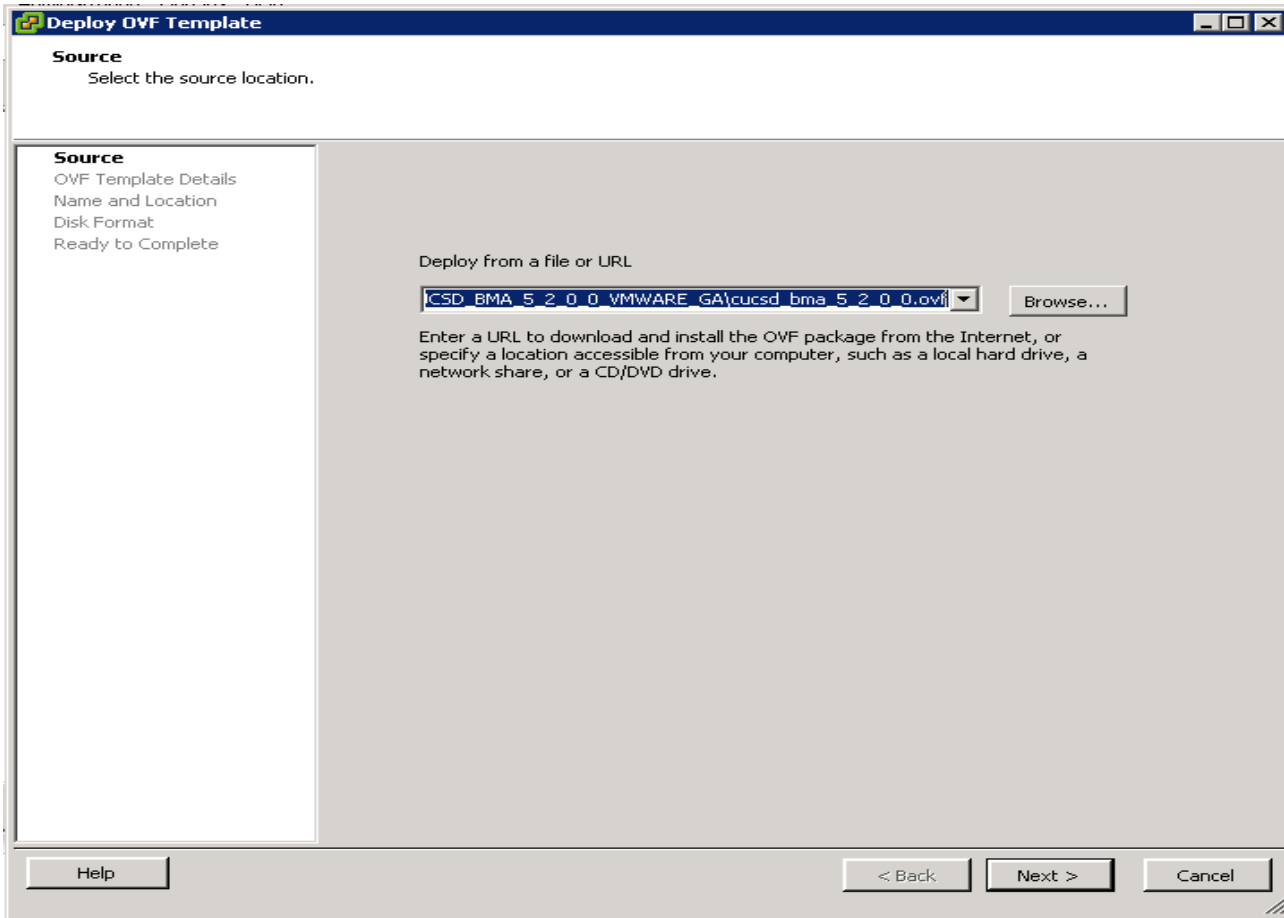
1. Launch the VMWare vSphere client software
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click on **Login** to continue.
4. From the **File** menu, Select **Deploy OVF Template**.

Figure 227 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 228 Select the Cisco UCS Director Baremetal Agent OVF file



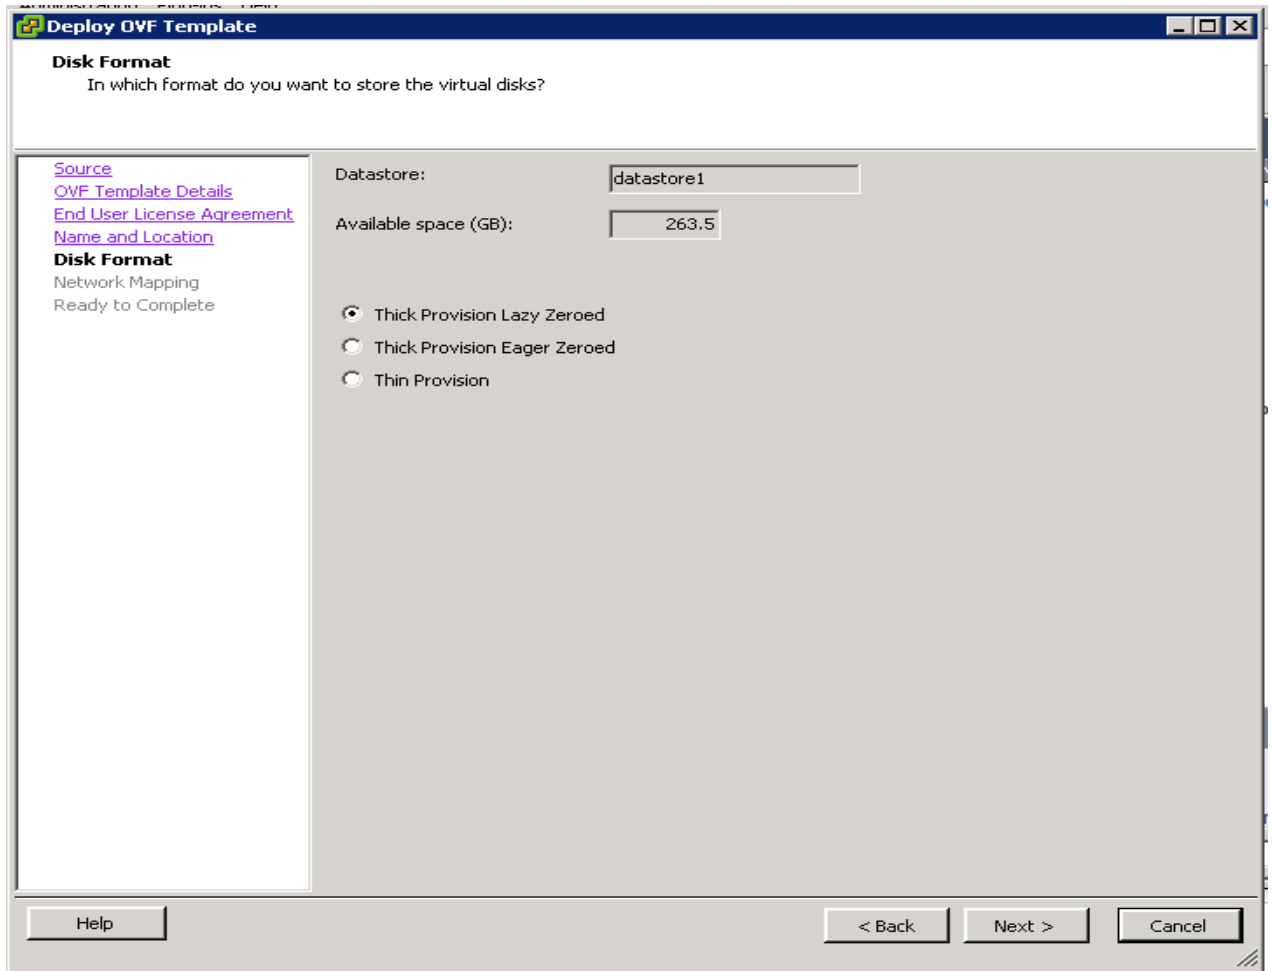
7. Review the details of the OVF template, Click **Next**.
8. Accept the End User License Agreement. Click **Next** to continue.
9. In the **Name and Location** option, Enter the name of the VM. Click **Next** to continue.

Figure 229 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".

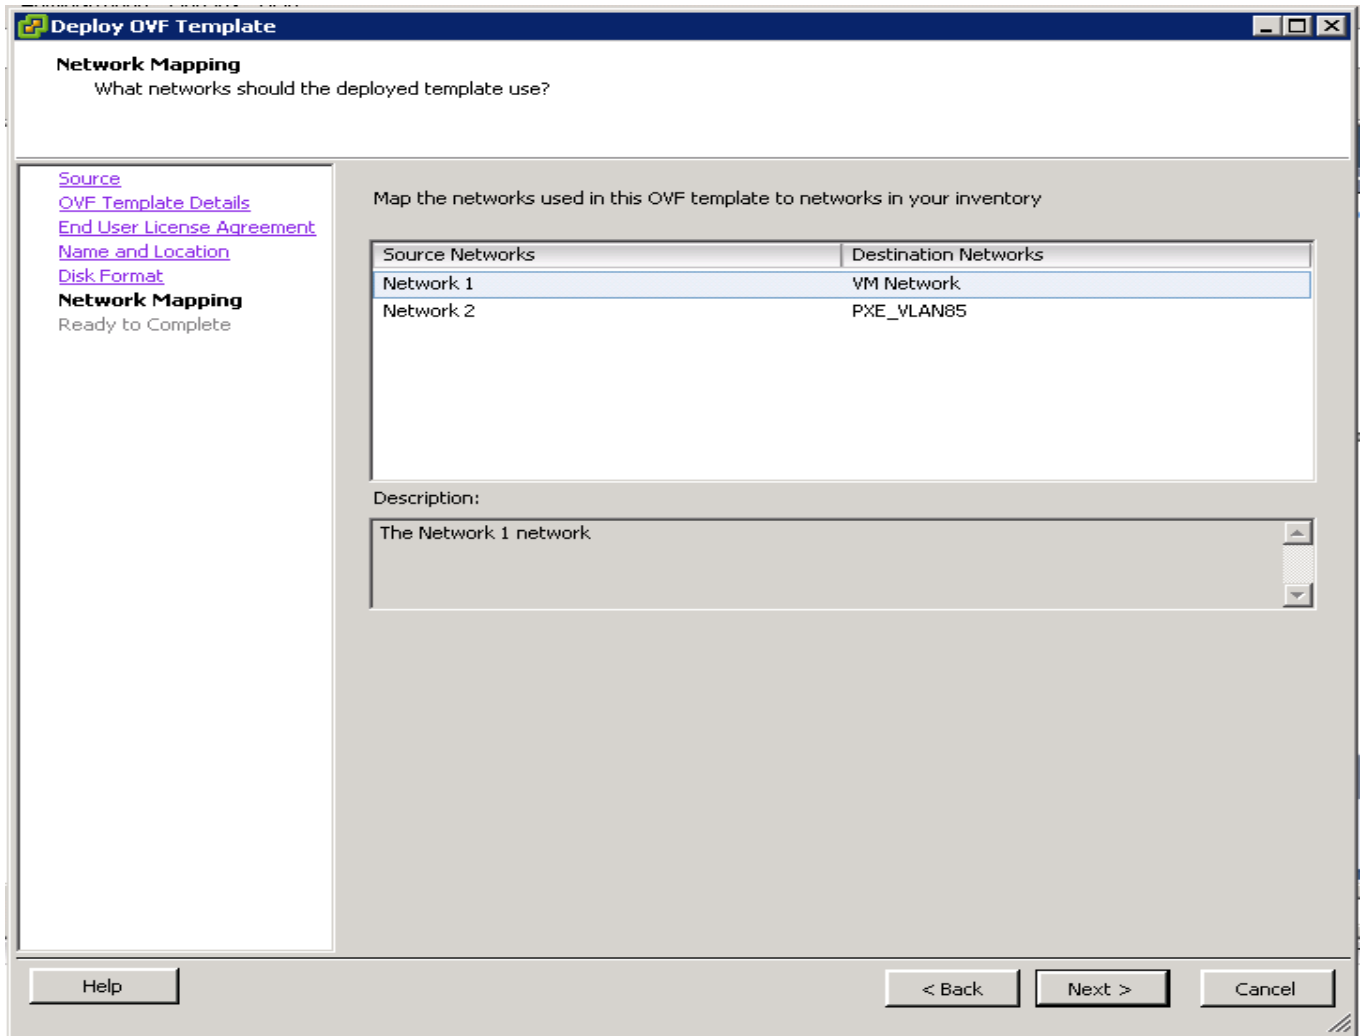
10. In the Disk Format option, click the **Thick Provision Lazy Zeroed** radio button. Click **Next** to continue.

Figure 230 Select the Disk Format for the VM



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 231 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 232 Deploy the Cisco UCS Director Baremetal Agent VM

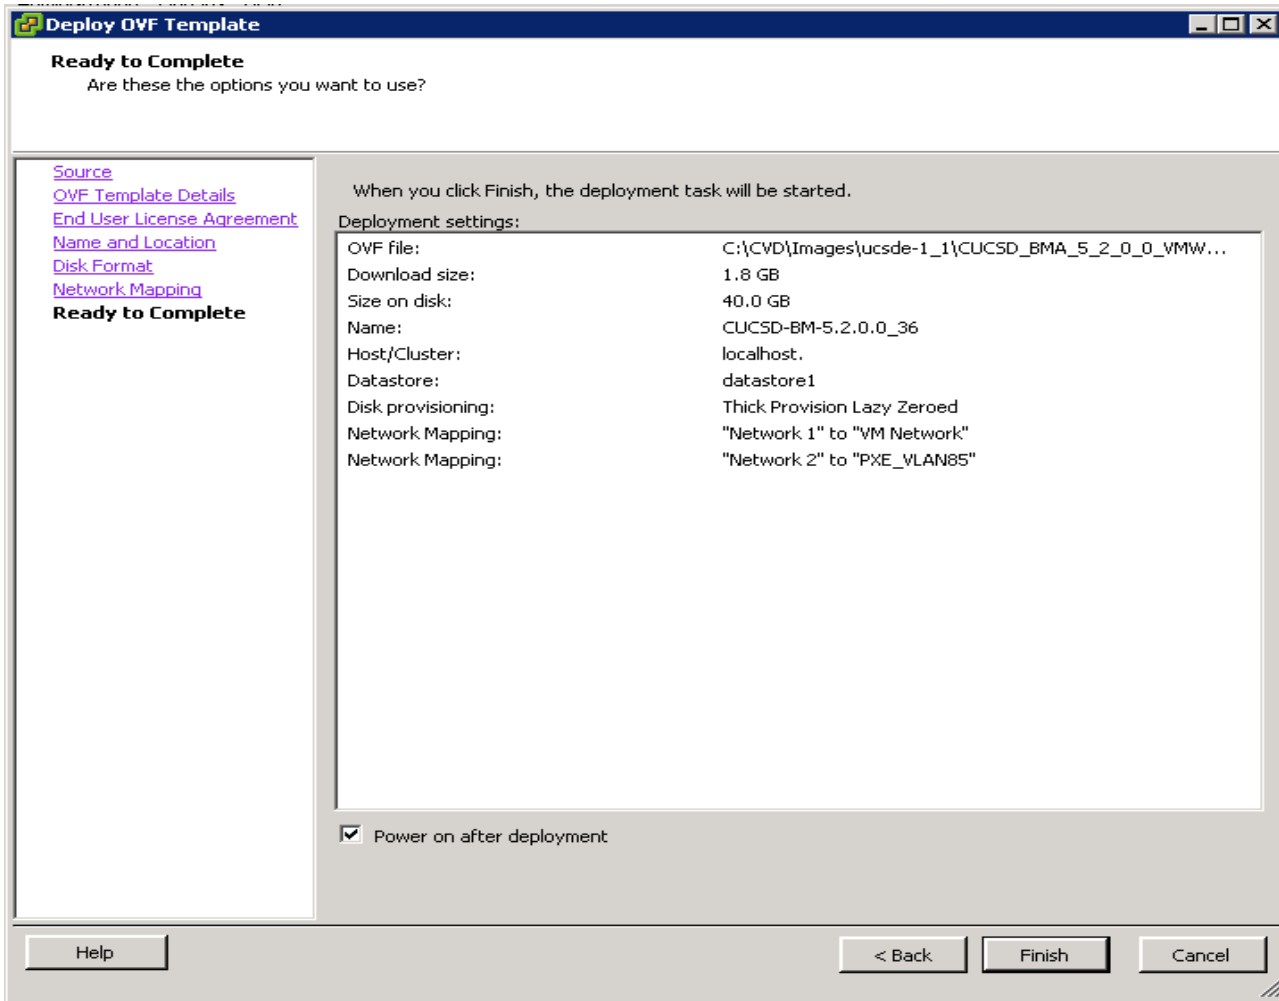
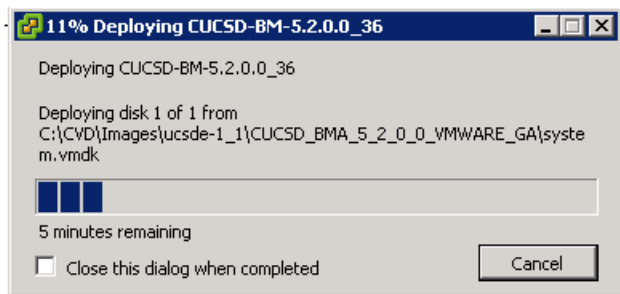


Figure 233 Cisco UCS Director Baremetal Agent VM Deployment in Progress

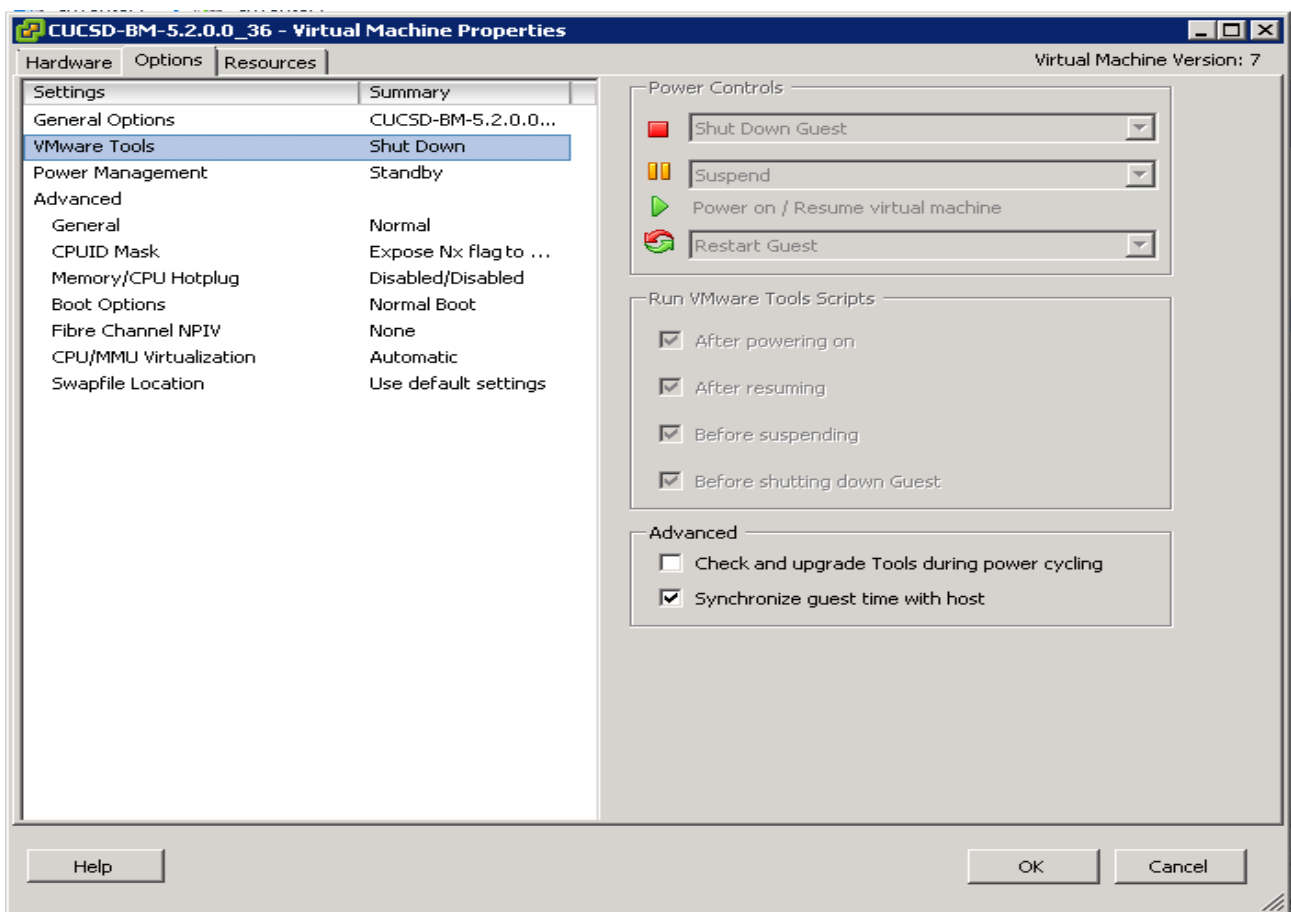


Configuring the Cisco UCS Director Baremetal Agent VM (BMA-VM)

The Cisco UCS Director Baremetal Agent VM named as CUCSD-BM-5.2.0.0_36 shall be known as BMA-VM here onwards.

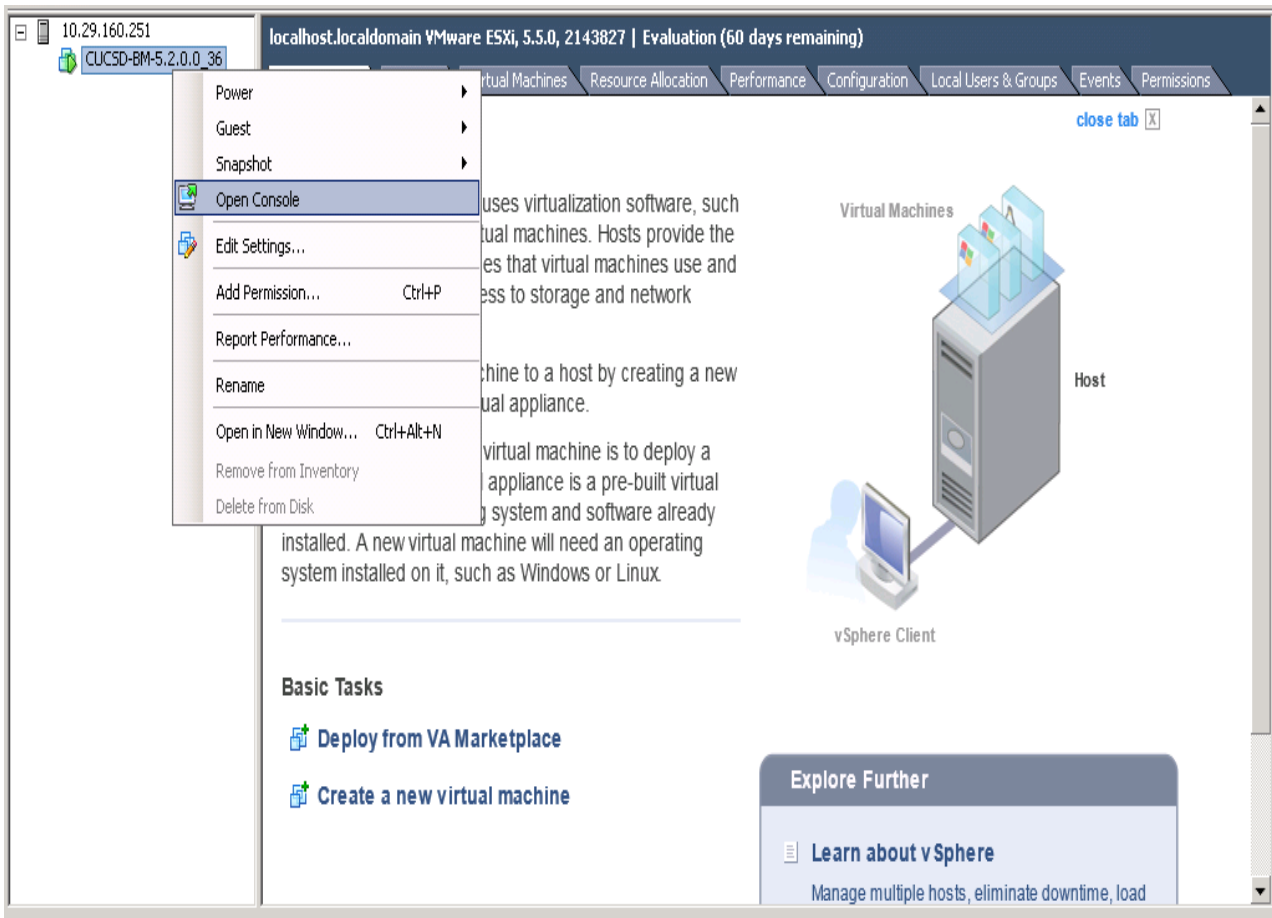
1. Right click on the BMA-VM, and select **Edit Settings**.
2. In the Virtual Machine Properties dialog box, click on the Options Tab.
3. Click on the VMWare **Tools**, Click on the **Synchronize guest time with host** option in the Advanced **section**.
4. Click on **OK** button to accept the changes.

Figure 234 Edit VM Settings to Synchronize the Guest Time with the ESXi Host



5. Right click on the BMA-VM, and select **Open Console**.

Figure 235 Access the VM Console of the BMA-VM



- In the console accept the End User License Agreement by typing **yes** and press **ENTER**.

Figure 236 Accept the EULA

```
Do you agree with the terms of the End User License Agreement?
yes/no [no]: yes_
```

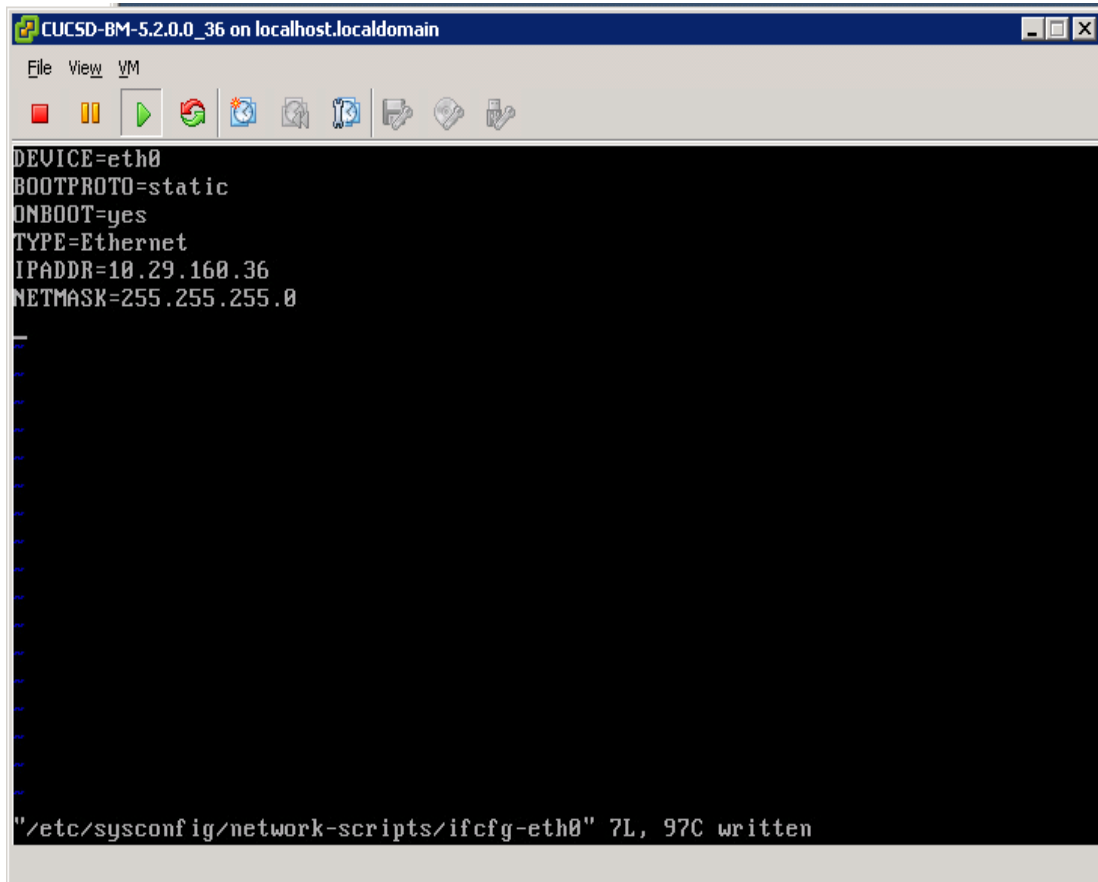
- Login as **root** user using the default password **pxeboot**.
- Configure the network interfaces by editing the `ifcfg-eth0` and `ifcfg-eth1` files located at `/etc/sysconfig/network-scripts/` directory, as follows:

Table 21 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 237 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 238 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 239 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 240 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 241 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 242 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 243 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 244 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 245 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 246 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	

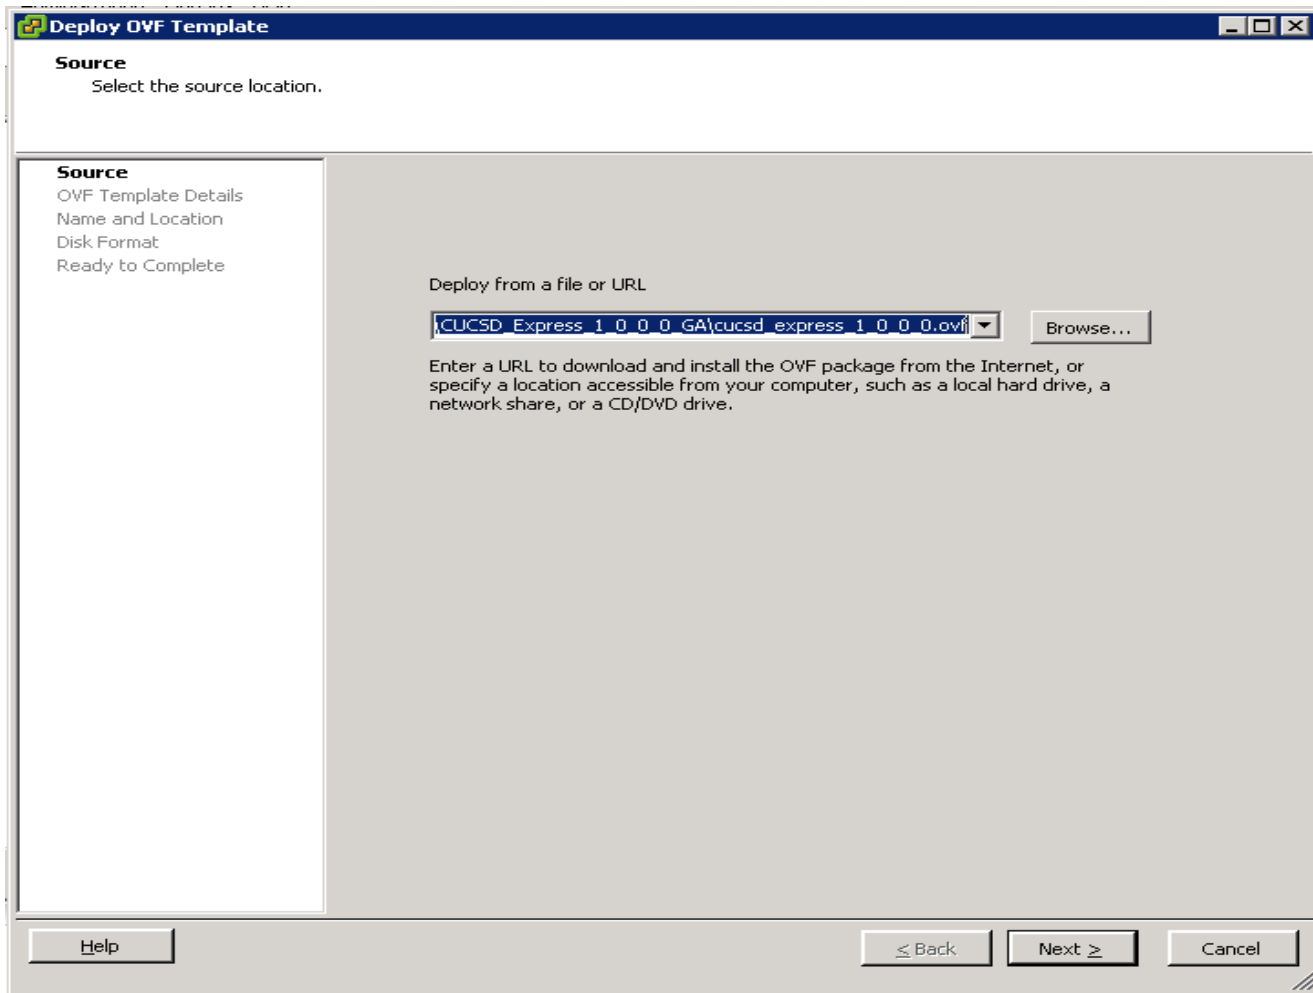
Apache/2.2.3 (CentOS) Server at 10.29.160.36 Port 80

BMA-VM configurations are complete.

Deploying the Cisco UCS Director Express OVF

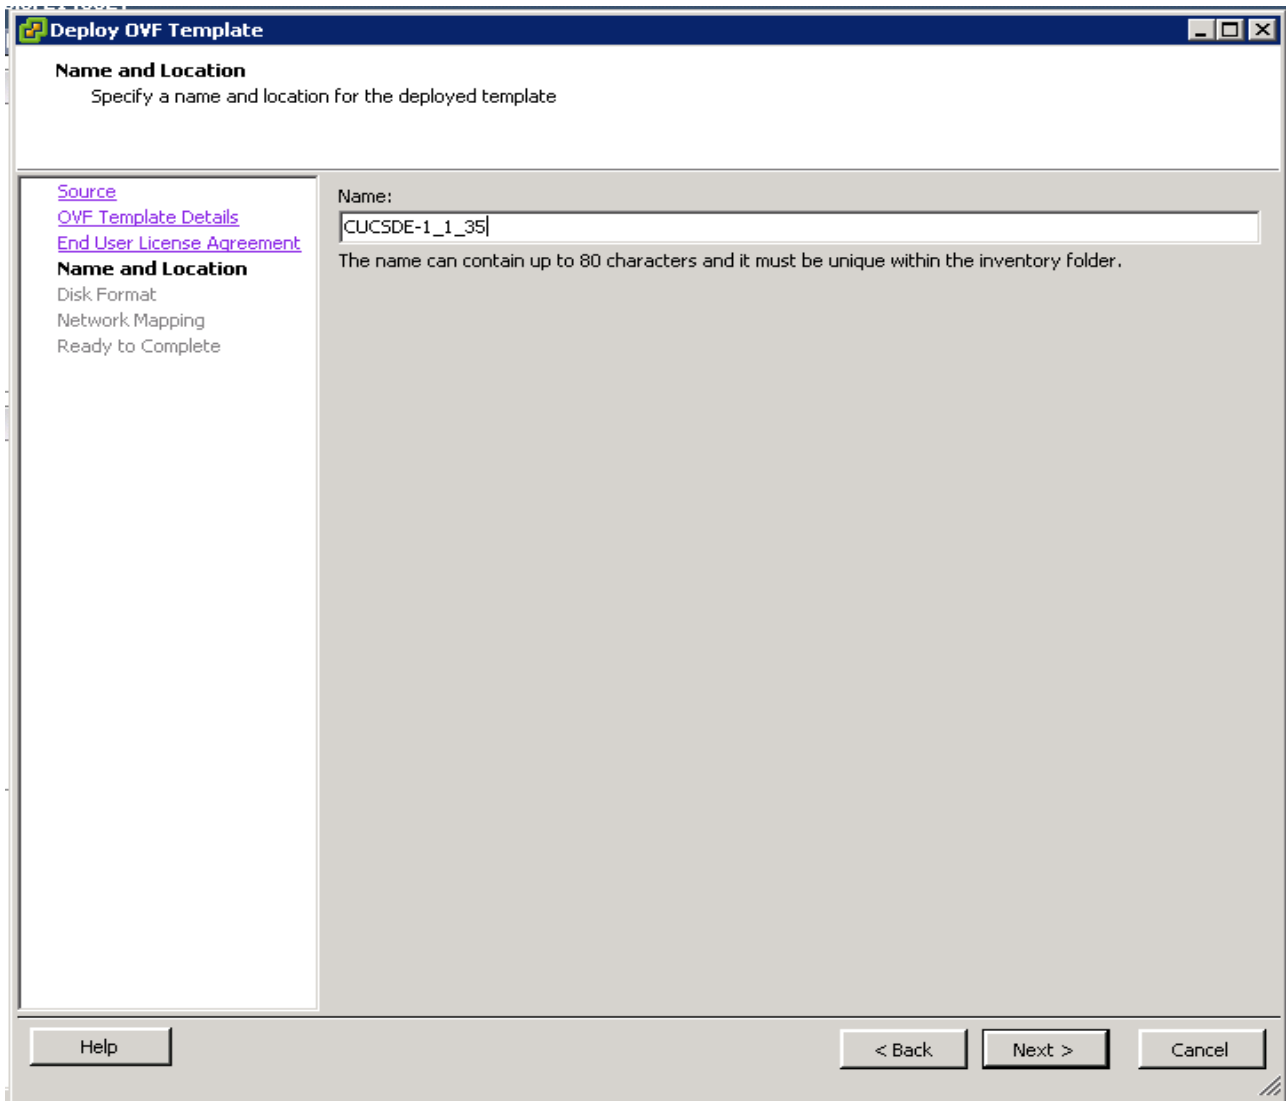
1. Launch the VMWare vSphere client software
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click **Login** to continue.
4. From the **File** menu, Select **Deploy OVF Template**.
5. Choose the Cisco UCS Director Express for Big Data 1.0 OVF template. Click **Open**.

Figure 247 Deploy the Cisco UCSD Express 1.0 OVF



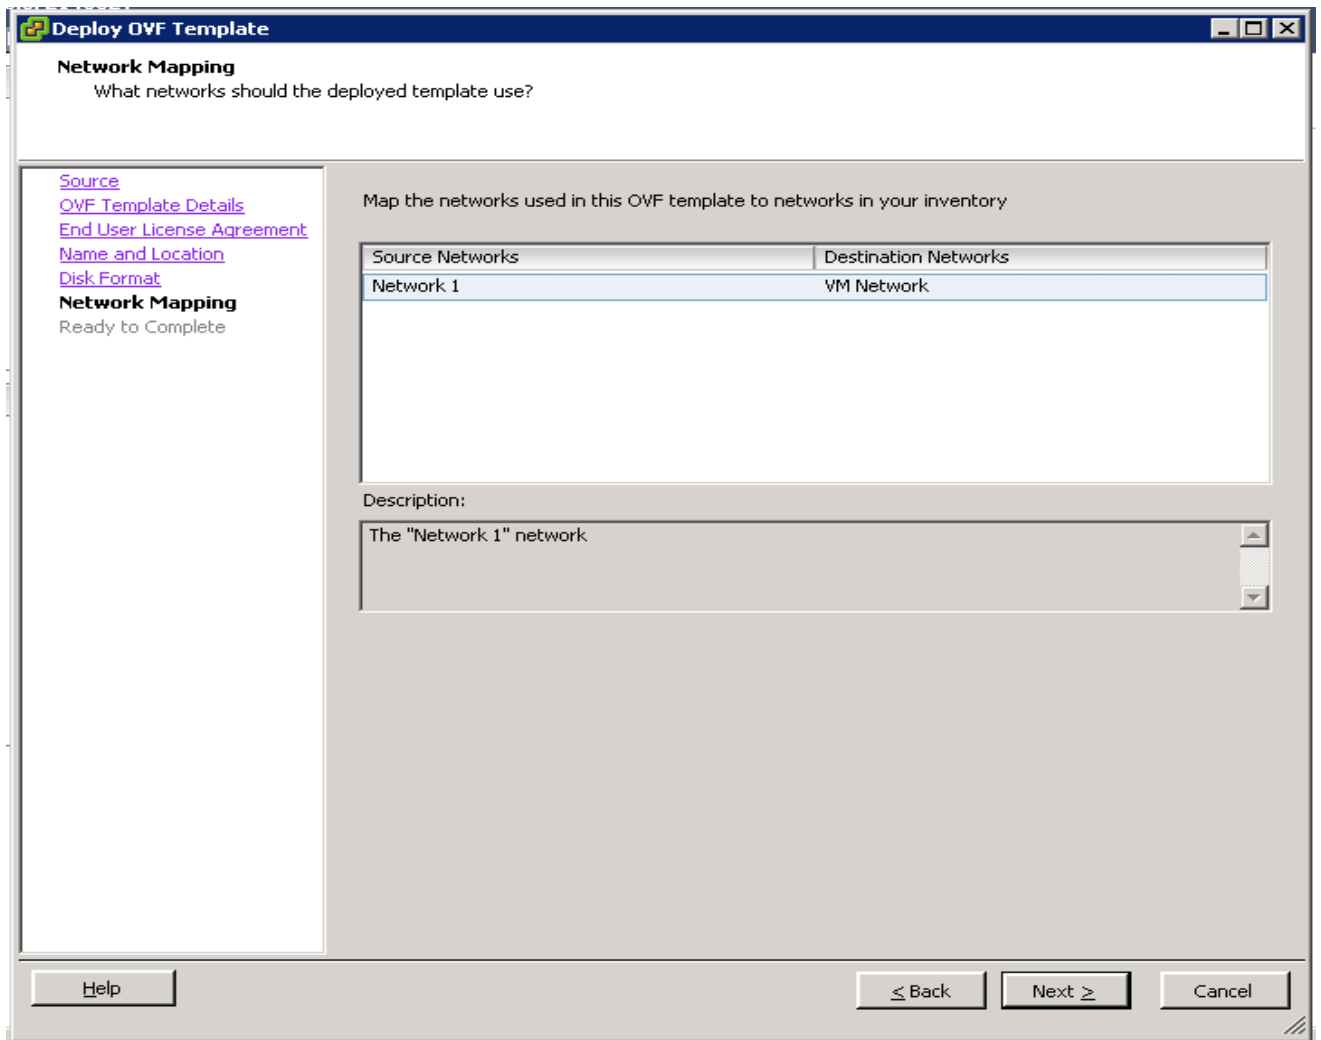
6. Review the details of the OVF, and Click **Next** to continue.
7. Accept the EULA, Click **Next** to continue.
8. Name the VM, Click **Next** to continue.

Figure 248 Name the Cisco UCS Director Express VM



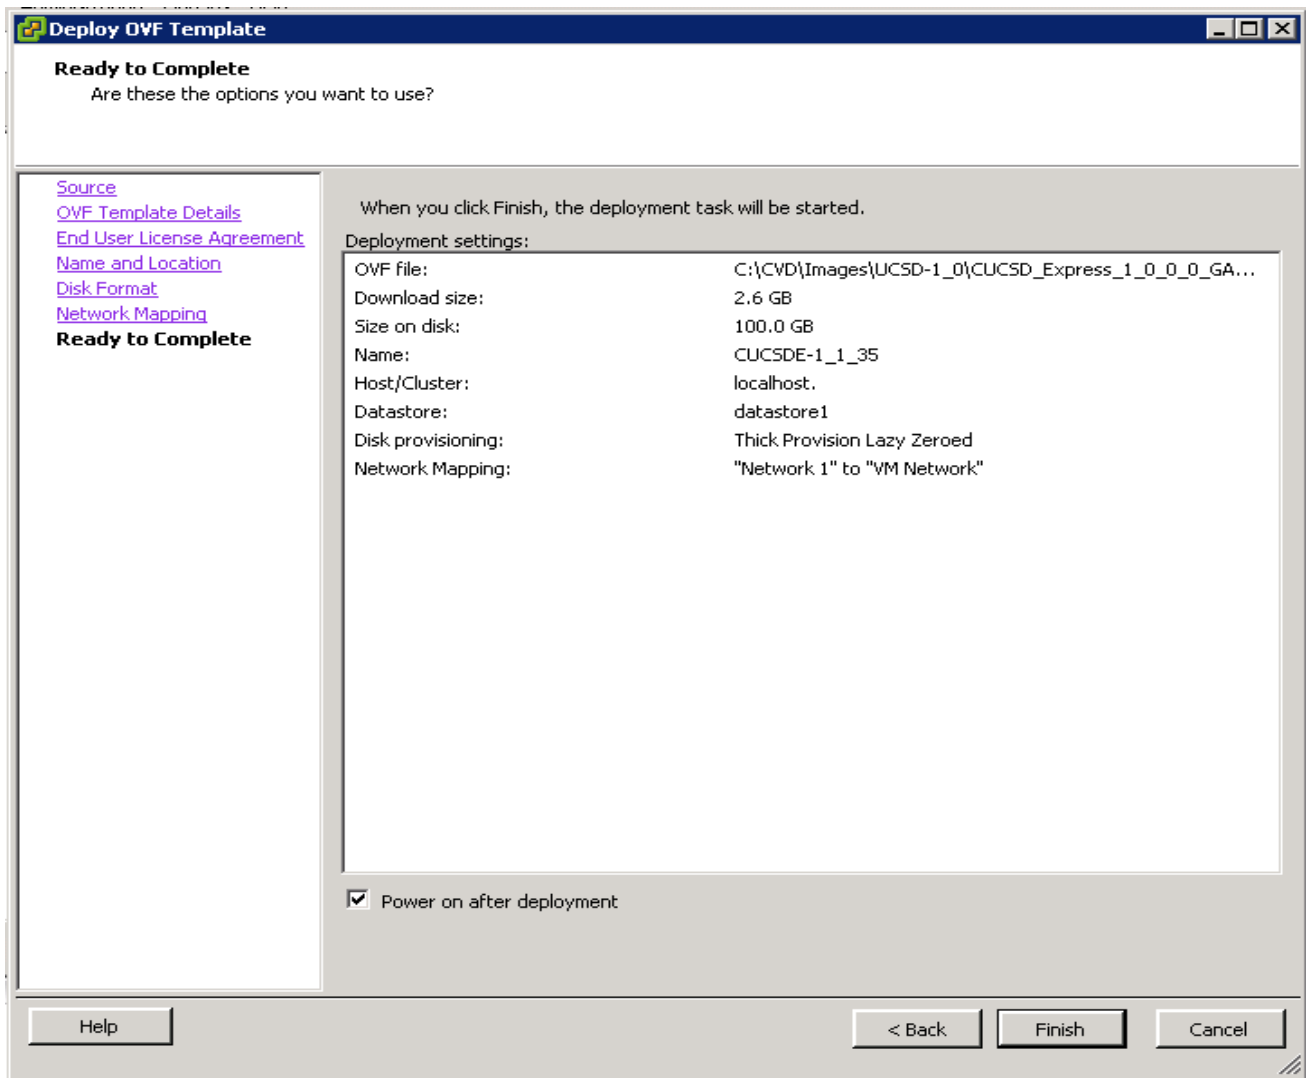
9. Choose the destination network **VM Network** for the source network **Network 1**. Click **Next** to continue.

Figure 249 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 250 Deploy the Cisco UCS Director Express VM

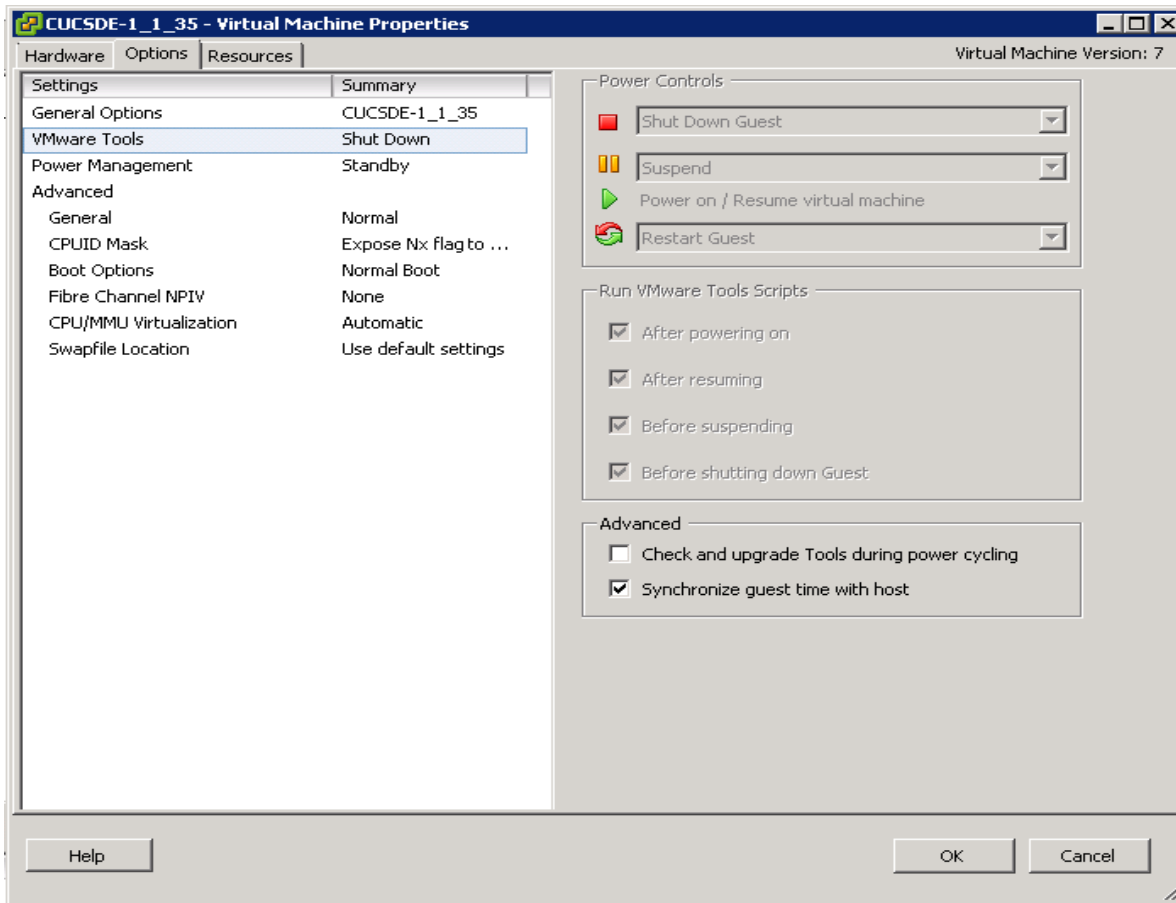


Configuring the Cisco UCS Director Express VM (UCSD-VM)

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 251 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 252 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 253 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 254 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 255 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 256 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 257 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 258 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 259 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 260 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 exist, 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 thread:
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 261 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)

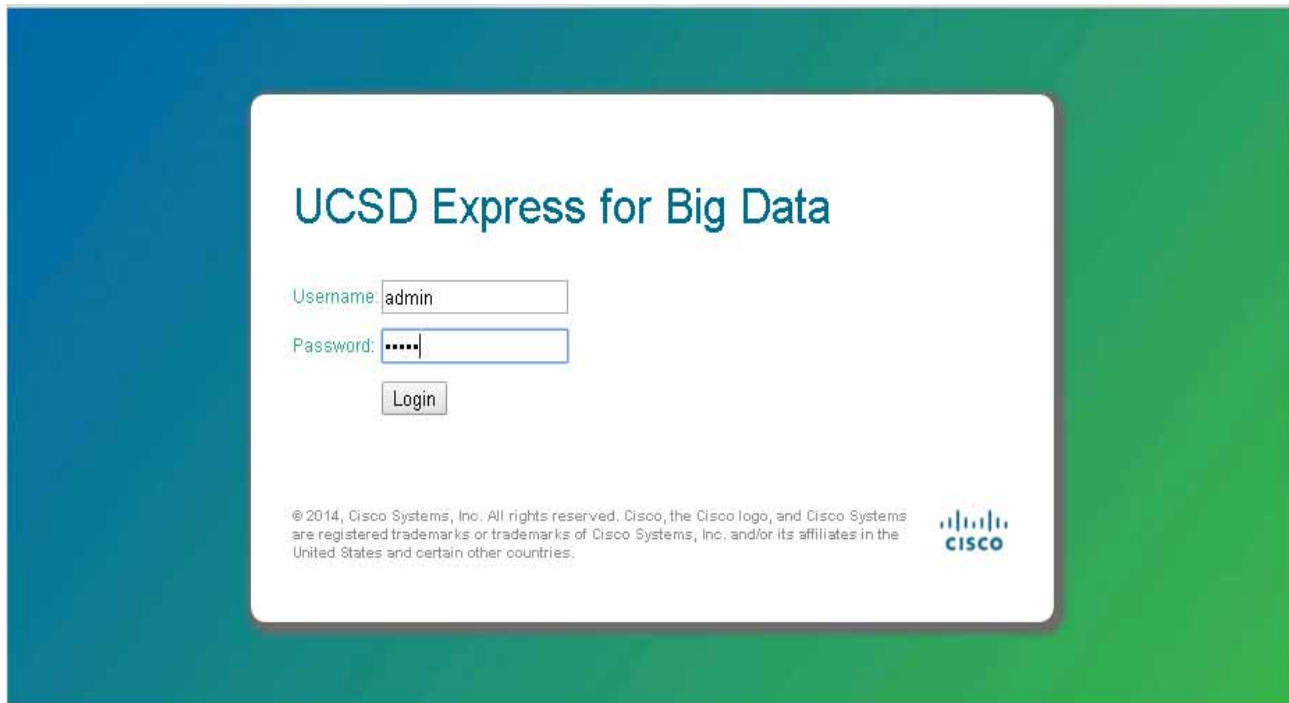
The Cisco UCS Director Express for Big Data, henceforth known as UCSD-Express, needs to be configured with the IP address to the UCS domain (i.e. UCS Manager's) physical account. This allows the UCSD-Express to query the UCS Manager and perform inventory collection.

The UCSD-Express will also need to be configured with the BMA's physical account and configure its services such as DHCP.

Add the licenses to UCSD-Express

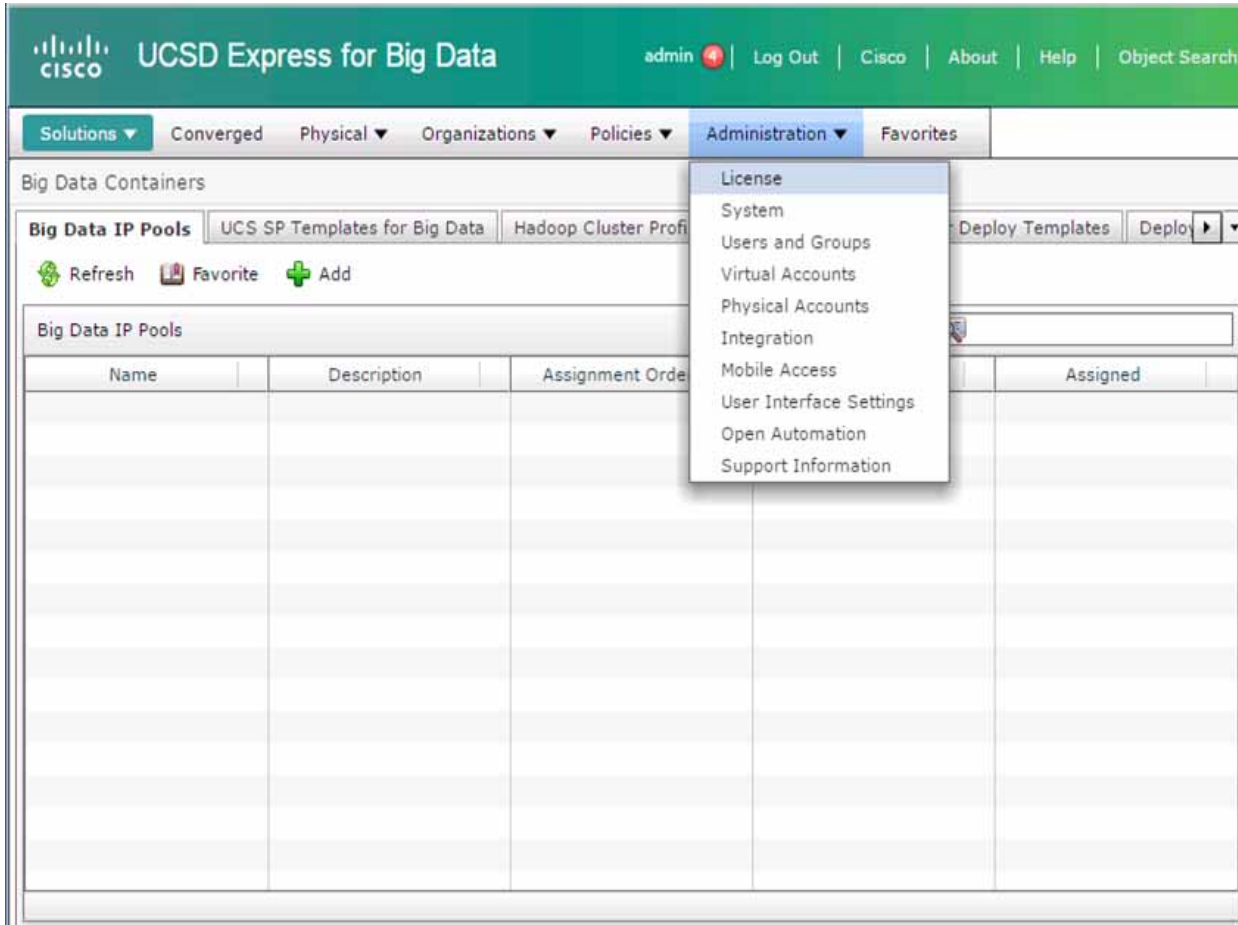
1. Using a web browser, visit the URL <http://<UCSD-VM's IP>/>.
2. Login as user **admin** with the default password **admin**.

Figure 262 Logging onto the Cisco UCS Director Express for Big Data



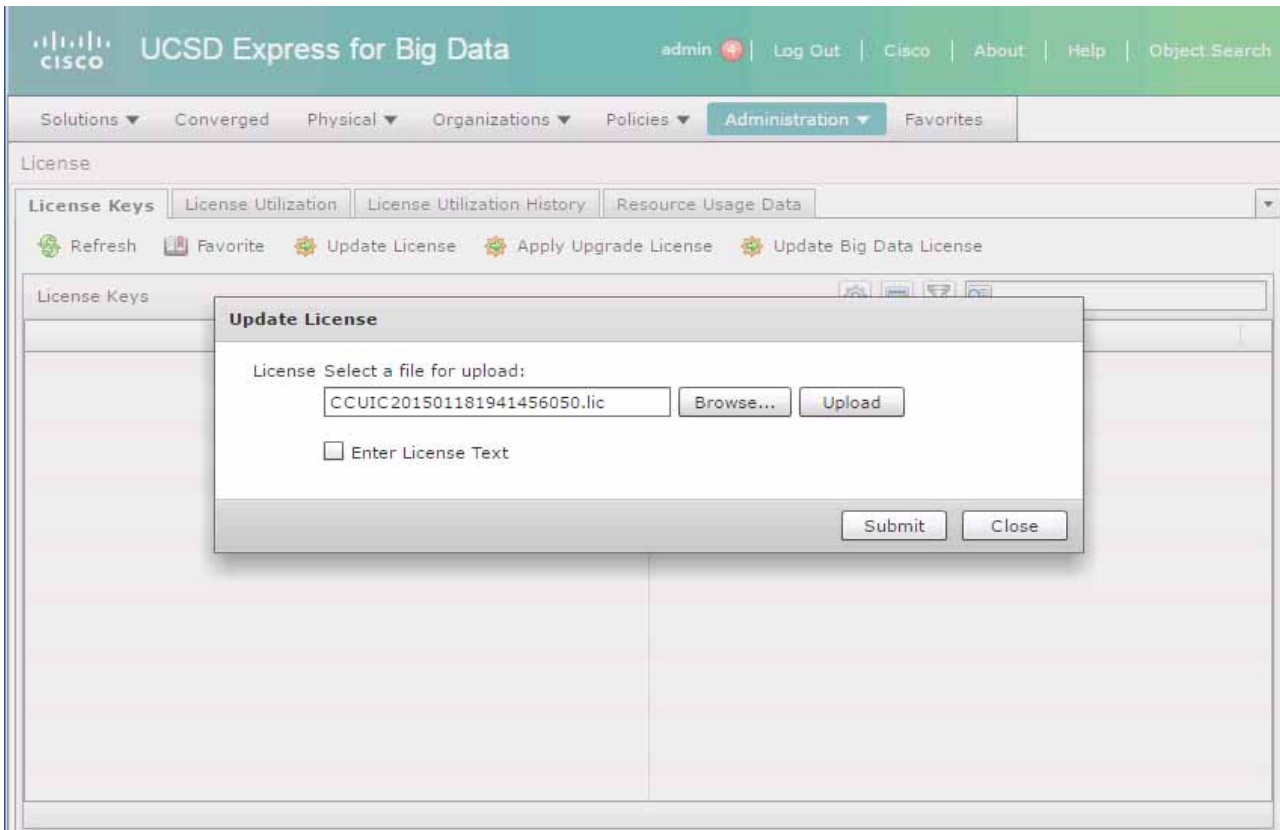
3. Navigate to **Administration > License** screen.

Figure 263 Accessing the License Administration Page



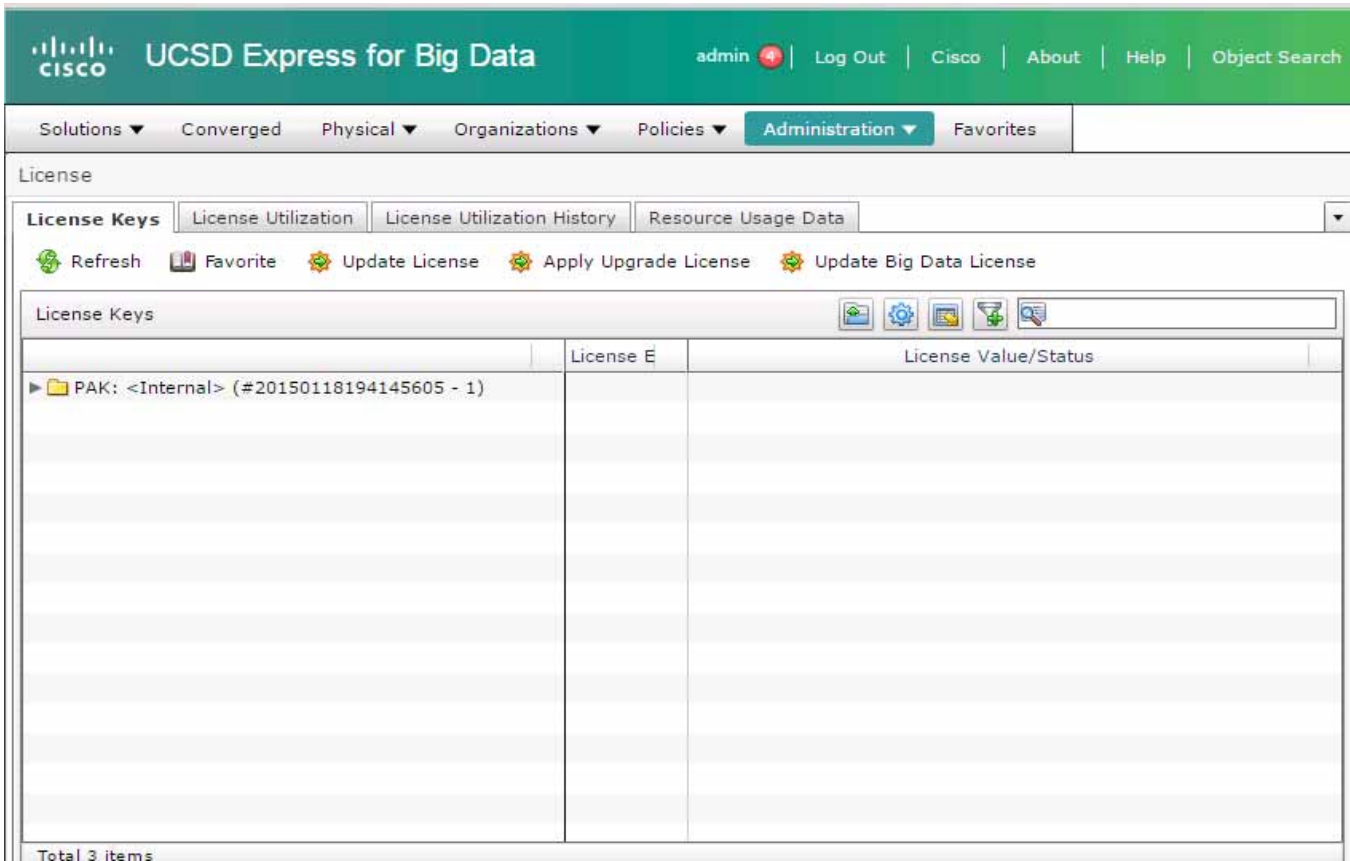
4. Click on **License Keys** tab.
5. Click on **Update License**.
6. In the **Update License** dialog box, click **Browse** to select the license file.
7. Click **Upload**.
8. After the license file gets uploaded, Click **Submit** to apply the license.

Figure 264 Applying the Base Cisco UCS Director License.



9. The license keys are displayed as shown below.

Figure 265 Cisco UCS Director Base Licenses got Applied Successfully



10. Click on **Update Big Data License**.
11. In the **Update Big Data Subscription** dialog box, click **Browse** to select the Big Data specific license file.
12. Click **Upload**.
13. After the license file gets uploaded, Click **Submit**.

Figure 266 Applying the Cisco UCS Director Express Big Data Subscription License

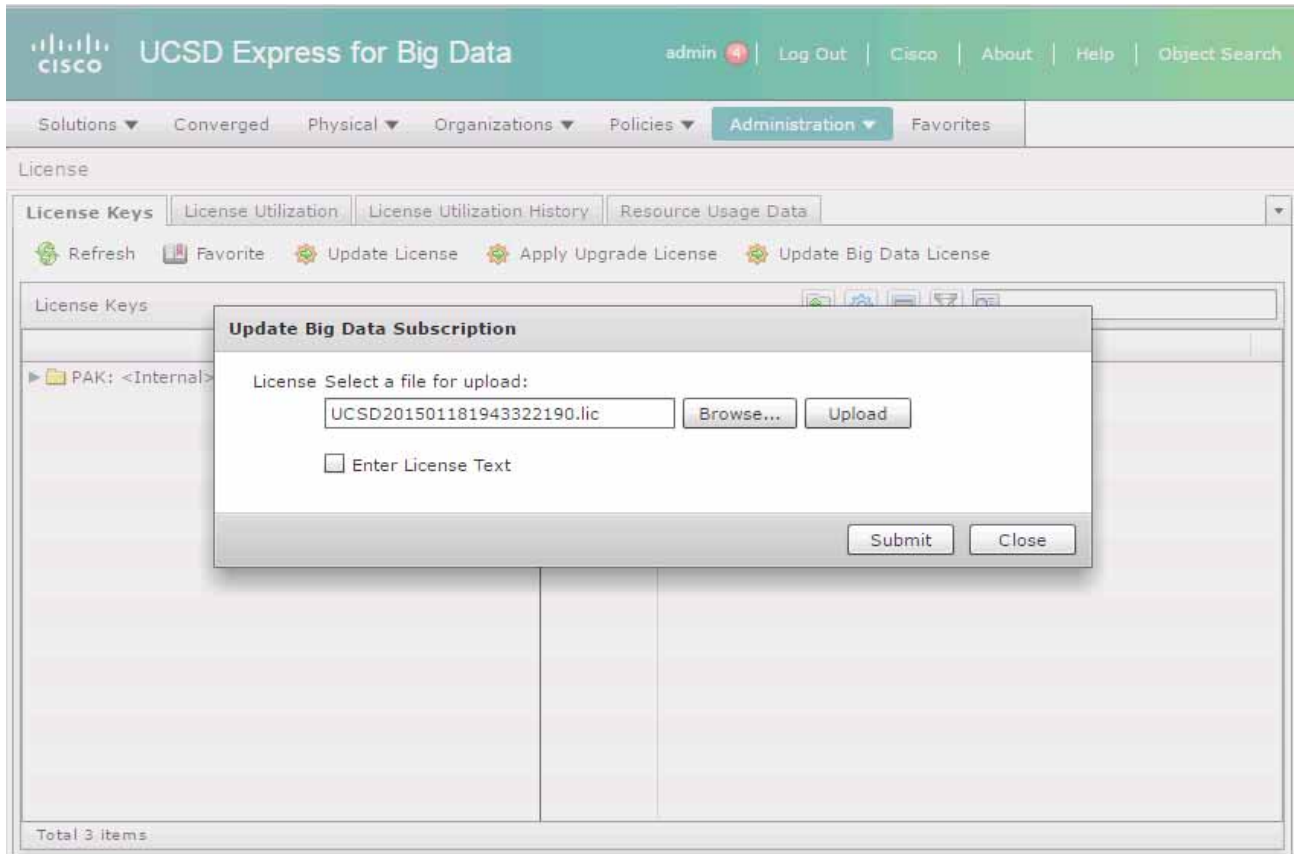


Figure 267 Completion of the License Application.

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 268 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.

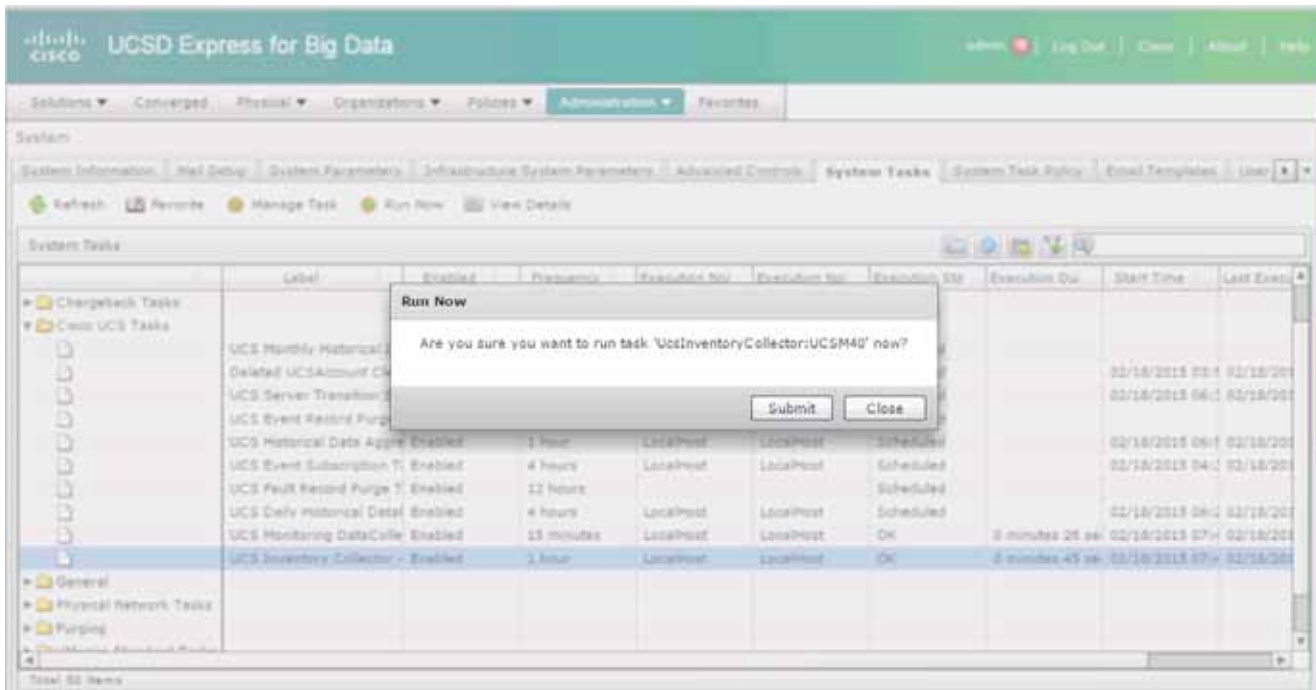


Note

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 269 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 270 Adding the Bare Metal Agent Appliance Information

Add Bare Metal Agent Appliance

BMA Name: BMA36 *

BMA Management Address: 10.29.160.36 *

NOTE: This address must be reachable from the Cisco UCS Director appliance

Login ID: root *

Password: ***** *

BMA Uses Different Interfaces for Management and PXE Traffic

BMA PXE Interface Address: 192.168.85.36 *

Description: [Empty]

Location: [Empty]

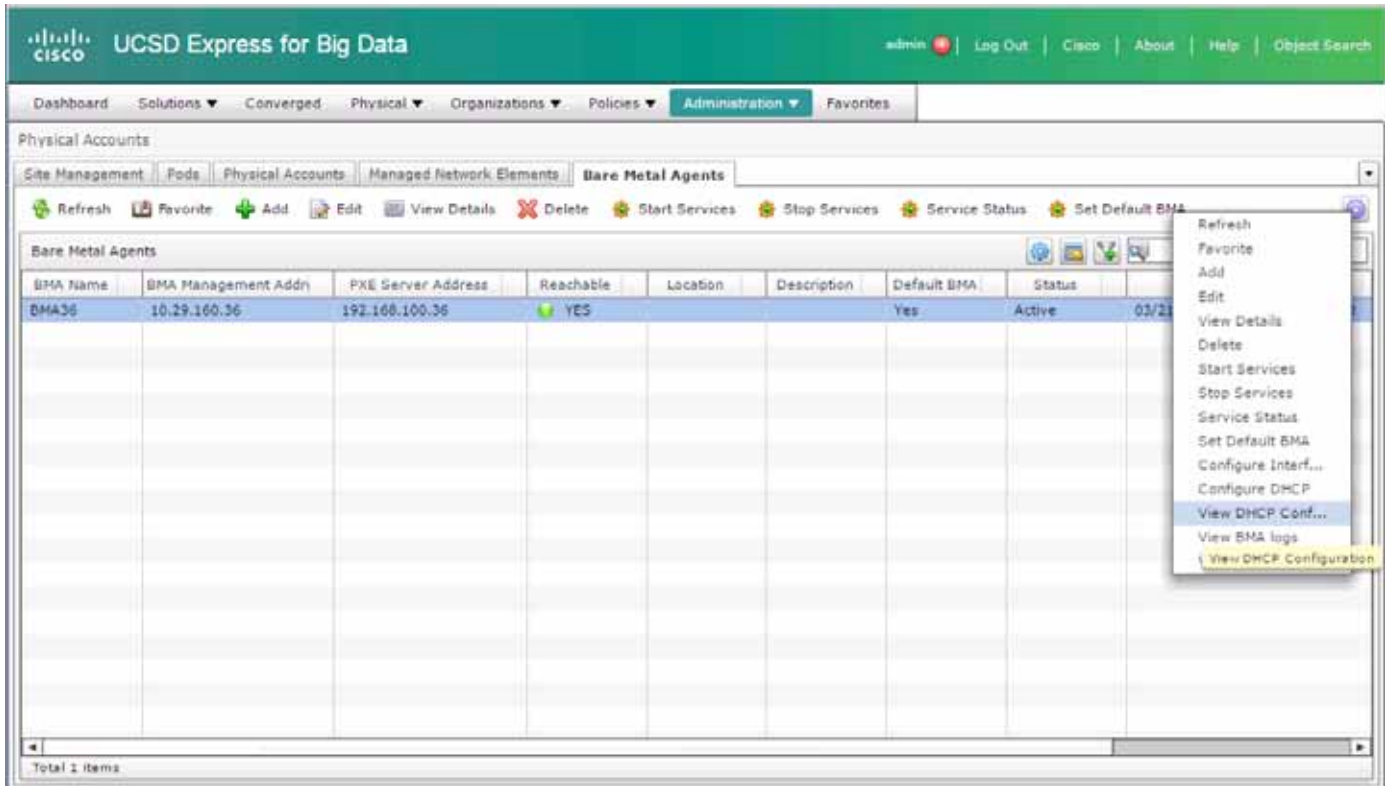
UCSD Database Address: 10.29.160.35 *

Submit Close

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 271 Configuring the DHCP



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 272 Configuring the DHCP services on the BMA.

Configure DHCP

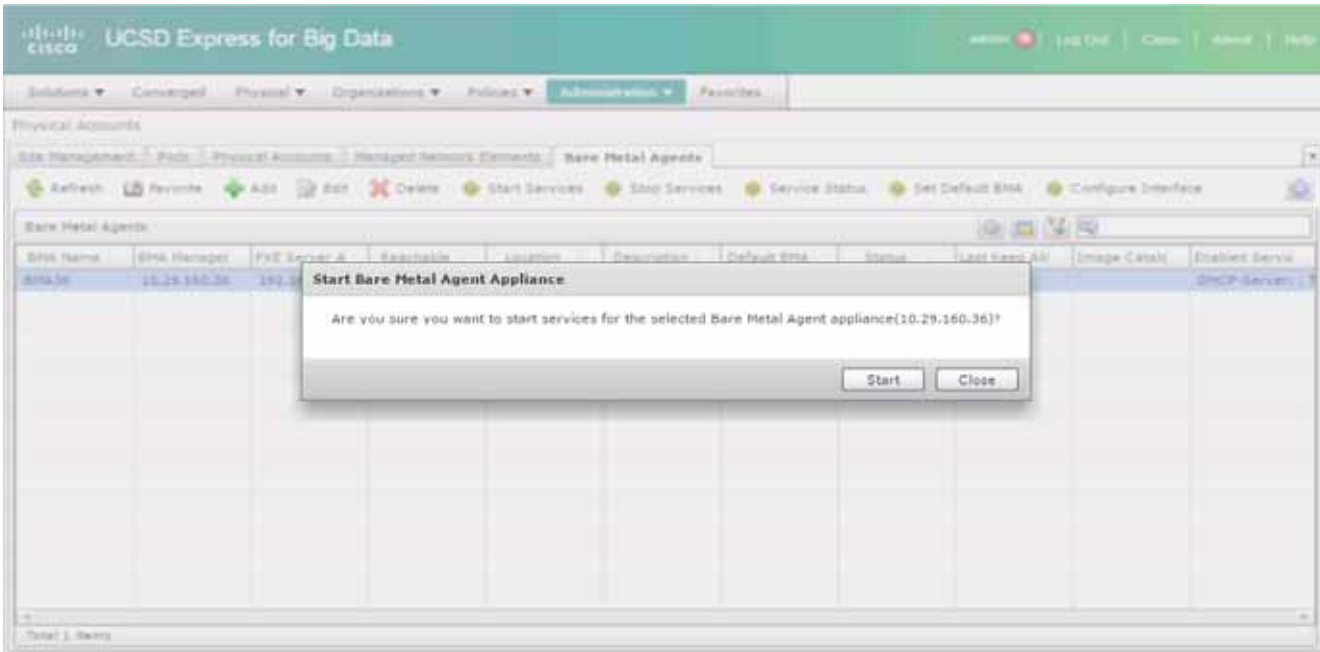
DHCP Subnet	<input type="text" value="192.168.85.0"/>	*
DHCP Netmask	<input type="text" value="255.255.255.0"/>	*
DHCP Start IP	<input type="text" value="192.168.85.160"/>	*
DHCP End IP	<input type="text" value="192.168.85.254"/>	*
Router IP Address	<input type="text" value="192.168.85.36"/>	

Submit Close

Start the BMA services

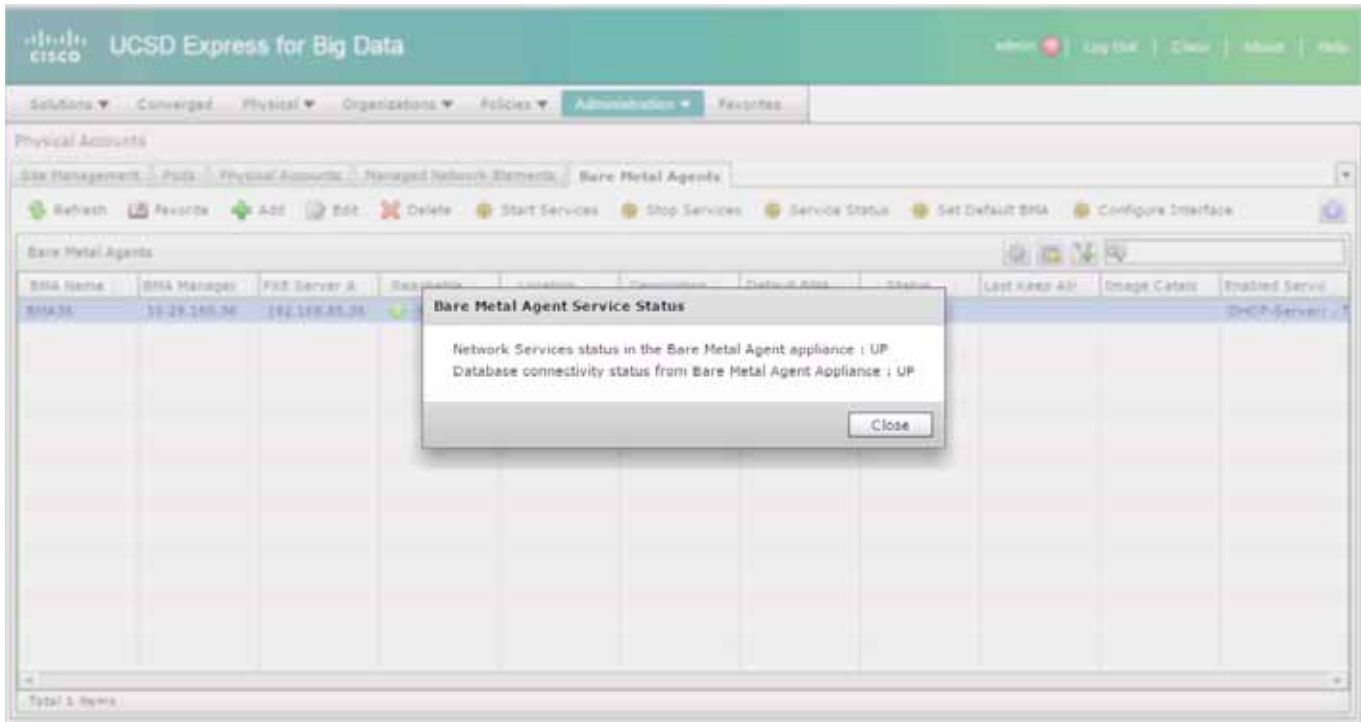
1. Navigate to **Administration >Physical Accounts >Bare Metal Agents**.
2. Select the BMA entry.
3. Click **Start Services**.
4. In the **Start Bare Metal Agent Appliance** dialog box, click **Start** to start the services.

Figure 273 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 274 Verifying the Bare Metal Agent Services Status

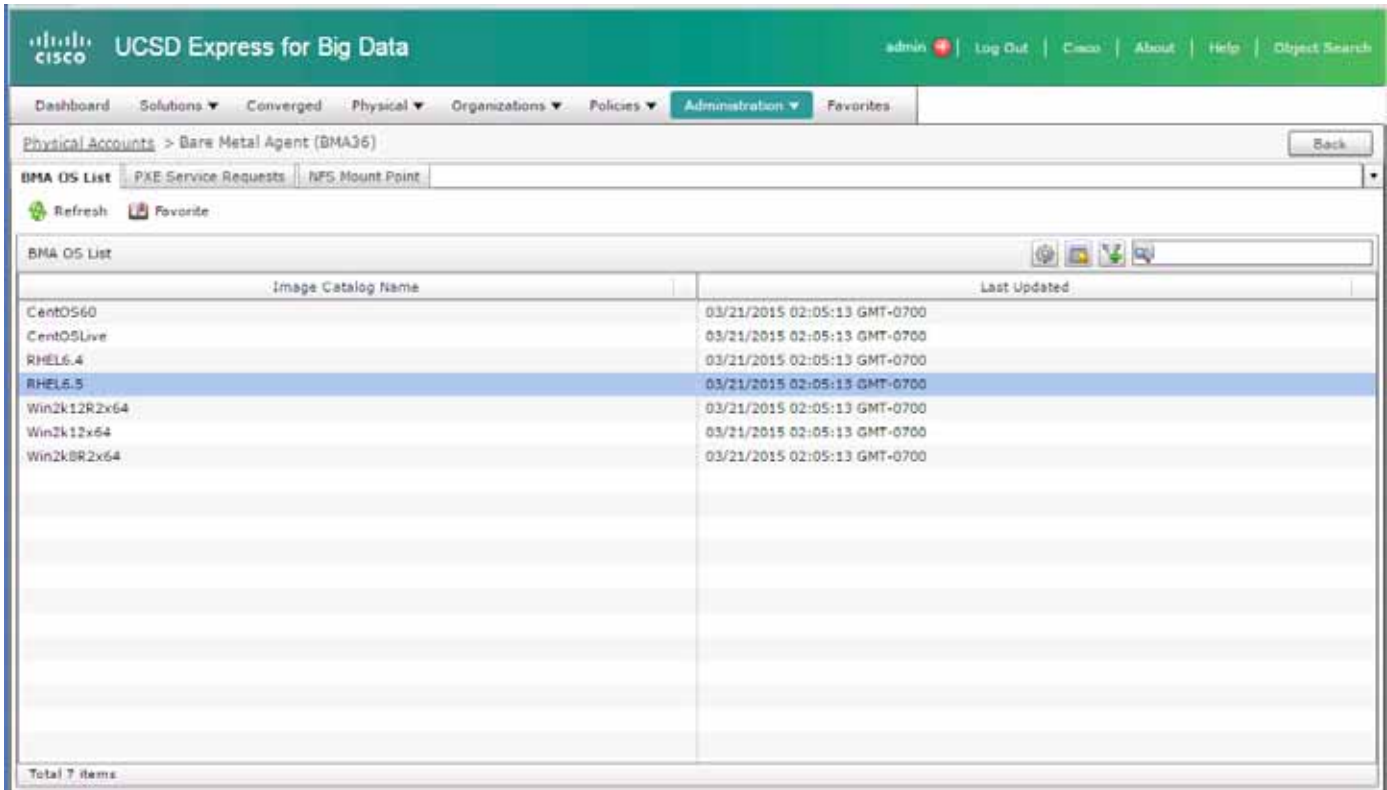


Note

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 275 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.

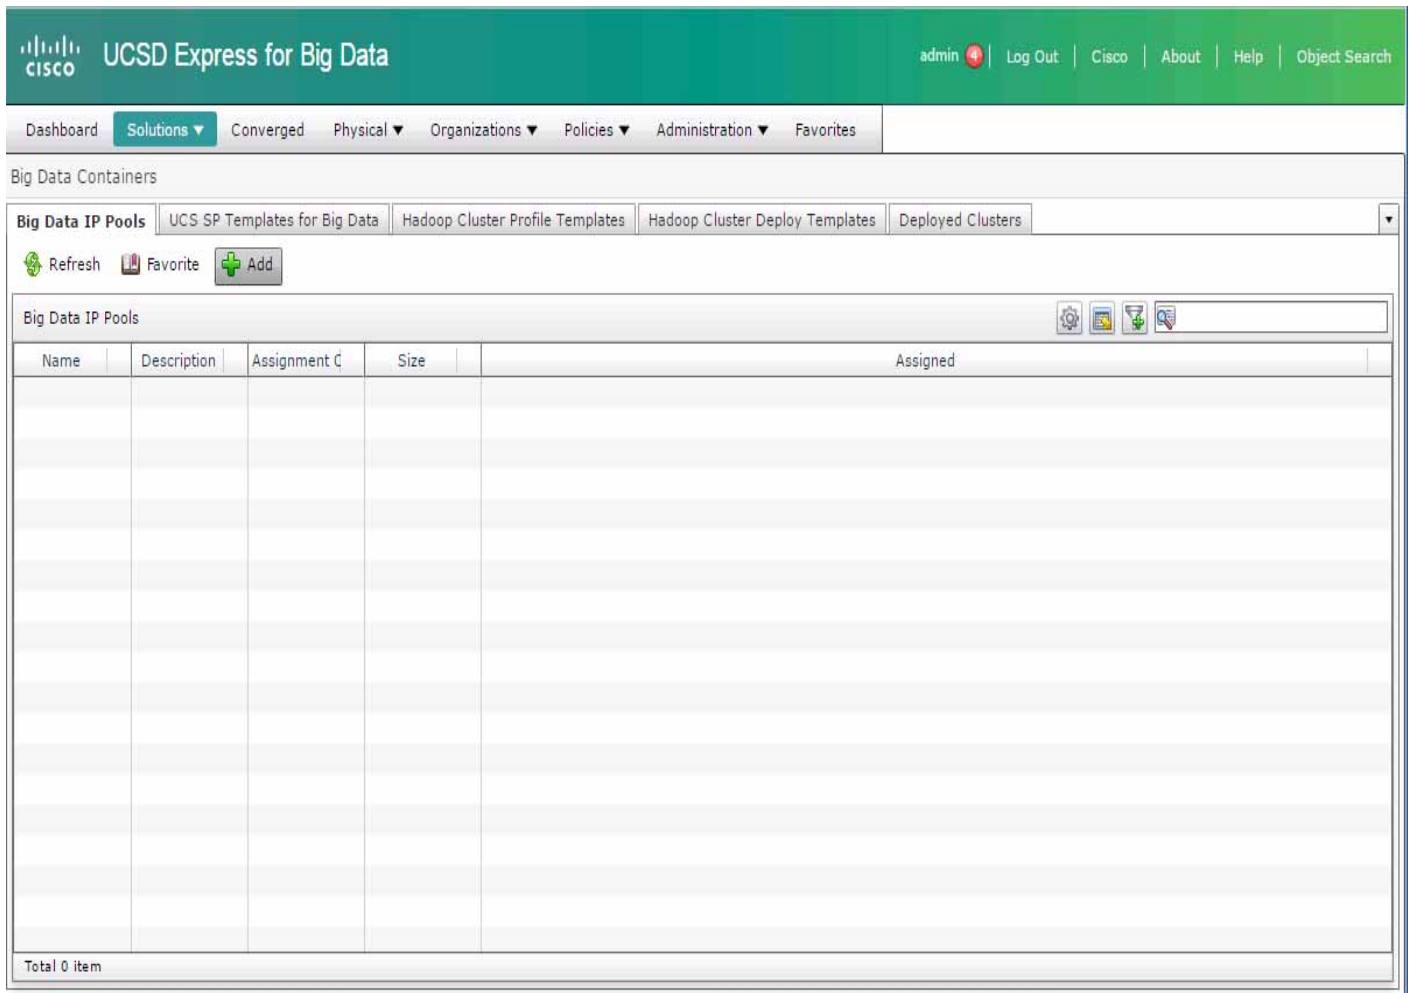
 **Note**

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 276 Creating the IP Address Pools



6. In the **Create an IP Pool** dialog box.
7. Enter the name **MGMT**. Click **Next** to continue.

Figure 277 *Creating the IP Address pool for MGMT VLAN*

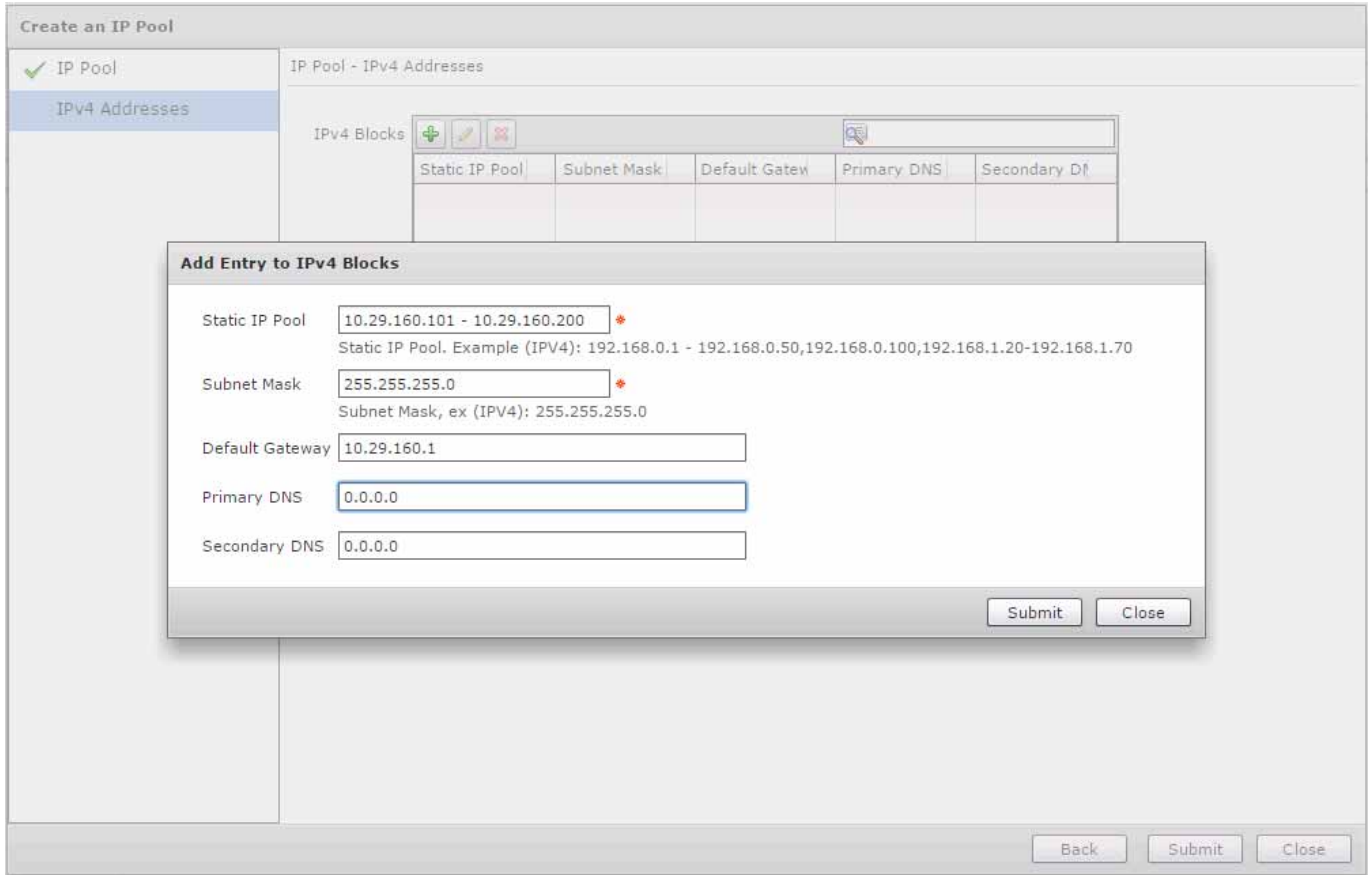
The screenshot shows a web-based configuration interface for creating an IP pool. The title bar reads "Create an IP Pool". On the left, a sidebar contains two menu items: "IP Pool" (highlighted) and "IPv4 Addresses". The main content area is titled "IP Pool Management" and contains the following fields:

- IP Pool Name:** A text input field containing the text "MGMT". A small red asterisk is visible to the right of the field, indicating a required field.
- Description:** An empty text input field.
- Assignment Order:** A dropdown menu with "Default" selected.

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 278 Adding a Block of IP Address to the MGMT IP Address Pool

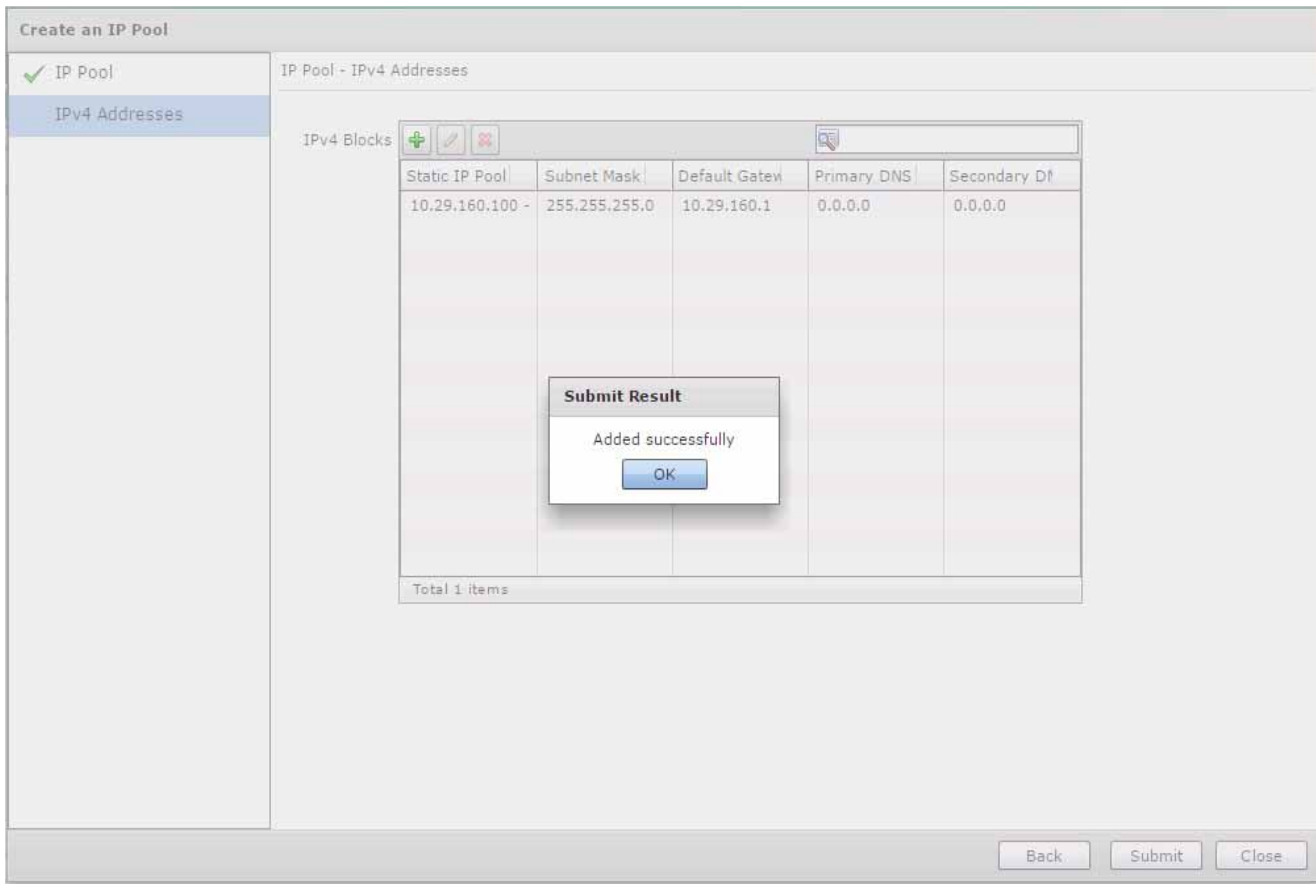



Note

The Default Gateway, Primary and Secondary DNS fields are optional.

11. Click **Submit** again to create the Big Data IP Pool.

Figure 279 *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 280 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 lists the configured IP pools:

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.

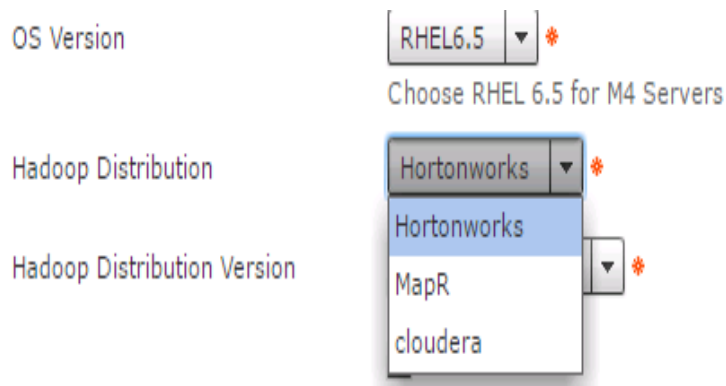
11. In the password fields, enter the preferred passwords and confirm them.
12. Choose the OS Version from the drop-down box. For C220 M4/C240 M4 rack servers, only OS supported is RHEL 6.5.


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 **Hadoop** from the drop-down list.

Figure 281 Selecting the Desired Hadoop Distribution



14. In the Hadoop Distribution Version field, select **Cloudera-5.3.0** from the drop down list.

Figure 282 Selecting the Hadoop Distribution Version

Cloudera	Hadoop Distribution	cloudera
	Hadoop Distribution Version	cloudera-5.2.0 cloudera-5.2.0 cloudera-5.3.0 cloudera-5.0.1 cloudera-5.2.1
	UCS Manager Account	
Hortonworks	Hadoop Distribution	Hortonworks
	Hadoop Distribution Version	Hortonworks-2.1 Hortonworks-2.1 Hortonworks-2.2
	UCS Manager Account	
MapR	Hadoop Distribution	MapR
	Hadoop Distribution Version	MapR-4.0.1 MapR-4.0.1 MapR-3.1.1 MapR-4.0.2
	UCS Manager Account	

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 283 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)

- Double-click on **eth1** and select appropriate IP-pool, MAC Address Pool and enter the DATA1 VLAN ID. Click **Submit**.

Figure 284 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

- Double-click on **eth2** and select appropriate IP-pool, MAC Address Pool and enter the DATA VLAN ID. Click **Submit**.

Figure 285 Editing the vNIC Template to Provide the DATA2 Network Configurations

Edit Entry

vNIC Name: eth2 *

IP Pool: DATA(192.168.12.101 - 192.168.12.200) *

MAC Address Pool: mac_pool1 (1978) *

VLAN ID: 12 *
[4048-4093],[1-3967]
(DATA2 VLAN)

Submit Close



Note

The following figure show the expanded version of the Instant Hadoop Cluster Creation dialog box with all the fields filled in.

Figure 286 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-root	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 287 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 23:38:19 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 14:28:38 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 288 Viewing a Completed Master Workflow

Workflow Status | Log | Objects Created and Modified | Input/Output

Service Request

Status Refresh

▼ Overview Current status for the service request.

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

Close



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 289 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.

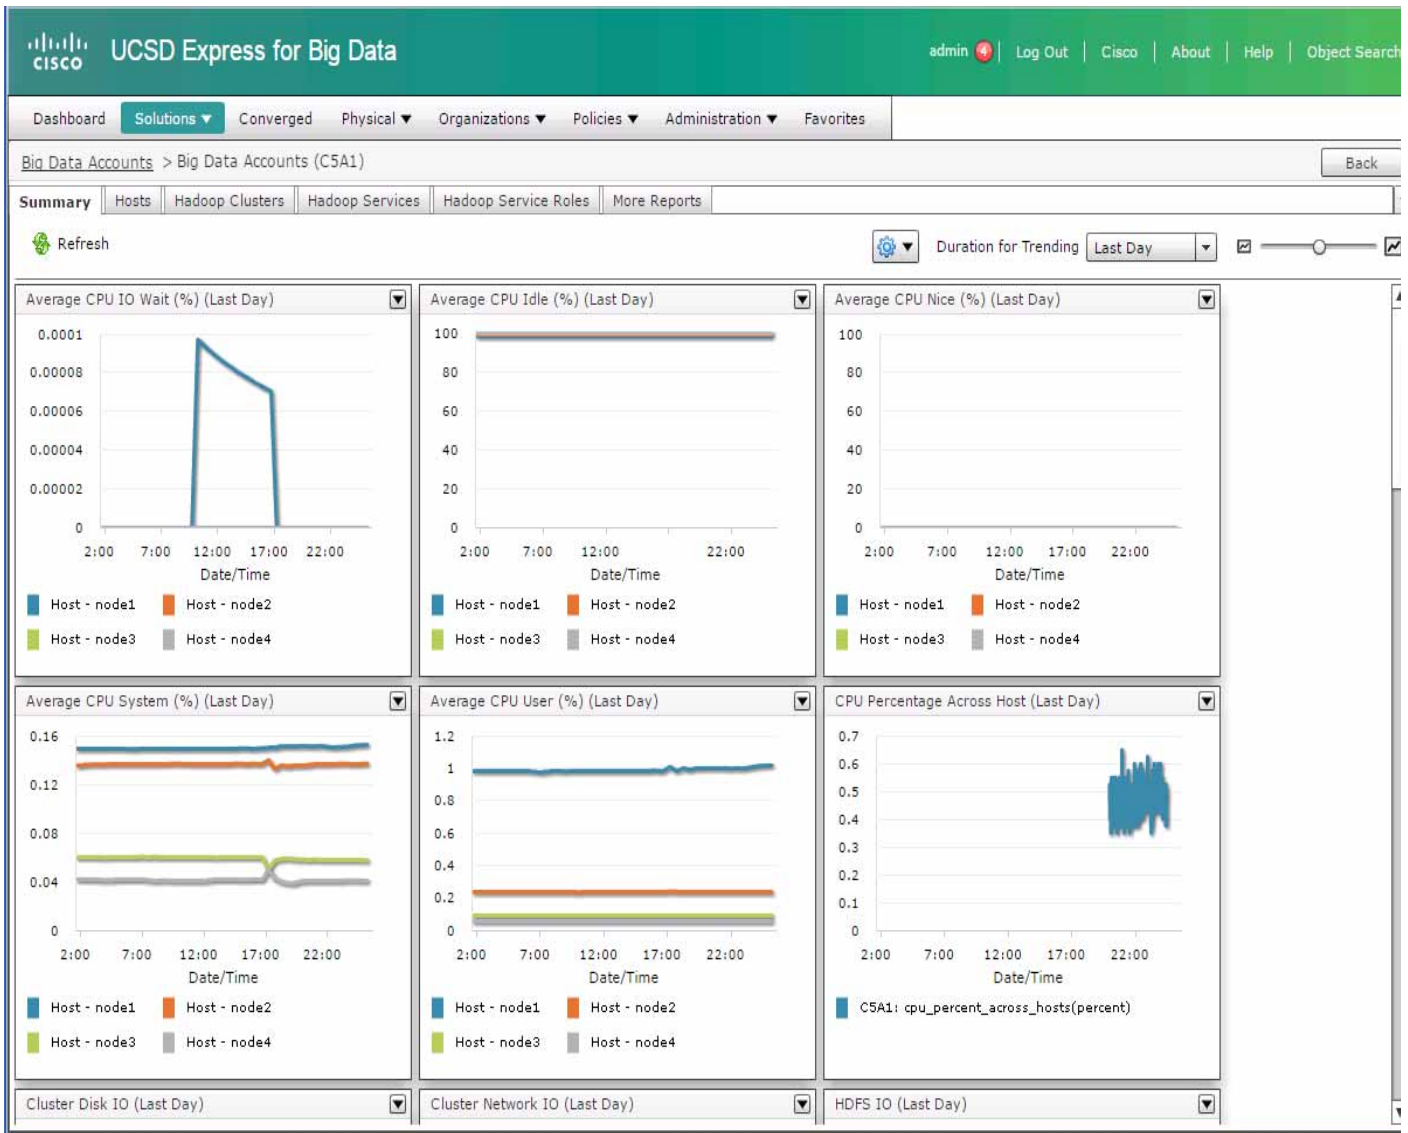
Figure 290 Big Data Accounts Summary Screen

Account Name	Account Type	Data Center	Management Console IP	Login
CSA1	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

Total 3 Items

2. Double-click on one of the accounts to view the cluster-wide performance charts.

Figure 291 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 292 *Big Data Accounts – Viewing the List of Hosts of a Particular Hadoop Cluster*

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

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 293 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 and the text 'UCSD Express for Big Data'. The user is logged in as 'admin'. The breadcrumb trail is 'Big Data Accounts > Big Data Accounts (CSA1)'. The 'Hadoop Services' tab is selected, showing a table of services. The table has columns for Status, Health, Service Type, and Service Name. The services listed are: flume, sqoop, ks_indexer, hue, Sentry, zookeeper, oozie, impala, hdfs, solr, spark, hbase, yarn, and hive. All services are in a 'STARTED' state with a 'GOOD' health. The bottom of the table indicates 'Total 14 items'.

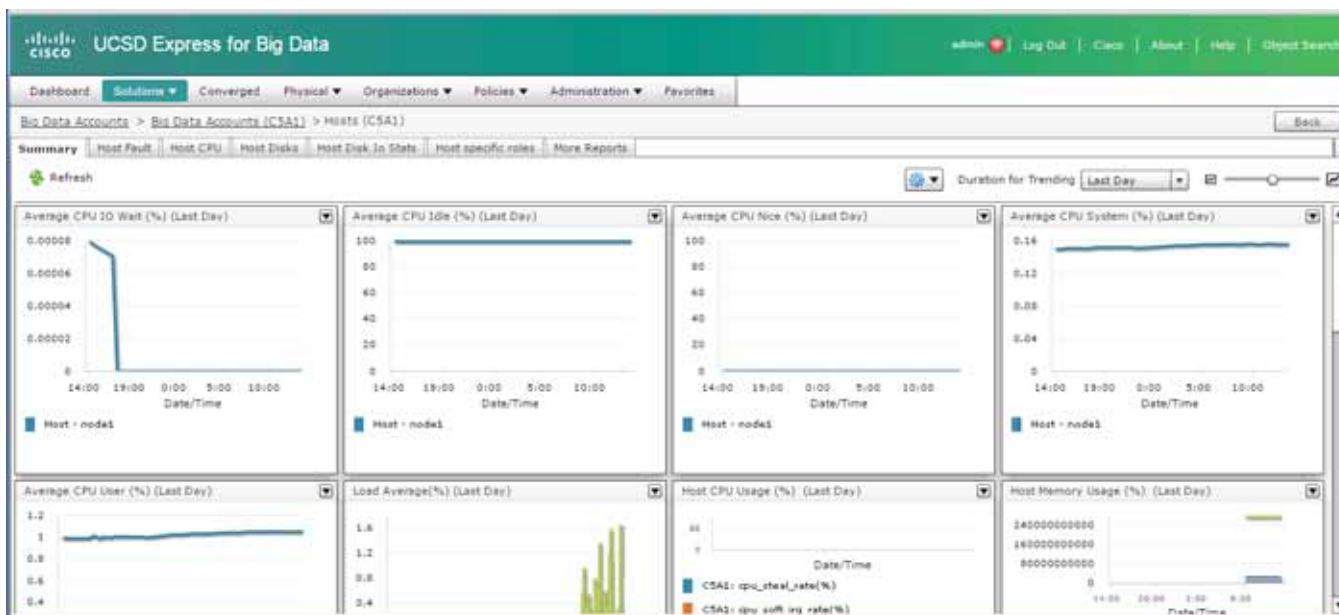
Status	Health	Service Type	Service Name
STARTED	GOOD	FLUME	flume
STARTED	GOOD	SQOOP	sqoop
STARTED	GOOD	KS_INDEXER	ks_indexer
STARTED	GOOD	HUE	hue
STARTED	GOOD	SENTRY	Sentry
STARTED	GOOD	ZOOKEEPER	zookeeper
STARTED	GOOD	OOZIE	oozie
STARTED	GOOD	IMPALA	impala
STARTED	GOOD	HDFS	hdfs
STARTED	GOOD	SOLR	solr
STARTED	GOOD	SPARK	spark
STARTED	GOOD	HBASE	hbase
STARTED	GOOD	YARN	yarn
STARTED	GOOD	HIVE	hive

Total 14 items

Host level Monitoring

In the **Hosts** tab, double-click on one of the hosts to view the host's statistics.

Figure 294 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 22 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 22 *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 23 *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