



VMware vSphere 5.0 Built On FlexPod


Last Updated: November 1, 2012



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors

John Kennedy, Technical Leader, Cisco

John Kennedy is a technical marketing engineer in the Server Access and Virtualization Technology group. Currently, John is focused on the validation of FlexPod architecture while contributing to future SAVTG products. John spent two years in the Systems Development unit at Cisco, researching methods of implementing long-distance vMotion for use in the Data Center Interconnect Cisco Validated Designs. Previously, John worked at VMware for eight and a half years as a senior systems engineer supporting channel partners outside the United States and serving on the HP Alliance team. He is a VMware Certified Professional on every version of VMware ESX and ESXi, vCenter, and Virtual Infrastructure, including vSphere 5. He has presented at various industry conferences.

John George, Reference Architect, Infrastructure and Cloud Engineering, NetApp

John George is a reference architect in the NetApp Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting cloud infrastructure solutions that include NetApp products. Before assuming his current role, he supported and administered Nortel's worldwide training network and VPN infrastructure. John holds a Master's Degree in Computer Engineering from Clemson University.

Ganesh Kamath, Technical Marketing Engineer, NetApp

Ganesh Kamath is a technical architect in the NetApp TSP Solutions Engineering team focused on architecting and validating solutions for TSPs based on NetApp products. Ganesh's diverse experiences at NetApp include working as a technical marketing engineer as well as a member of the NetApp Rapid Response Engineering team, qualifying specialized solutions for our most demanding customers.

Lindsey Street, Systems Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a systems architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelor of Science degree in computer networking and her Master of Science degree in Information Security from East Carolina University.

Chris Reno, Reference Architect, Infrastructure and Cloud Engineering, NetApp

Chris Reno is a reference architect in the NetApp Infrastructure and Cloud Engineering team and is focused on creating, validating, supporting, and evangelizing solutions based on NetApp products. Chris has his Bachelor of Science degree in International Business and Finance and his Bachelor of Arts degree in Spanish from the University of North Carolina-Wilmington while also holding numerous industry certifications.

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 of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <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. (1005R)

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.

© 2012 Cisco Systems, Inc. All rights reserved.



VMware vSphere 5.0 Built On FlexPod

Overview

Industry trends indicate a vast data center transformation toward shared infrastructures. By using virtualization, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod®, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements.

Audience

This document describes the architecture and deployment procedures of an infrastructure composed of Cisco Unified Computing System™, NetApp®, and VMware® virtualization that uses FCoE-based storage serving NAS and SAN protocols. The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the core FlexPod architecture.

Architecture

The FlexPod architecture is highly modular or “podlike.” Although each customer's FlexPod unit varies in its exact configuration, once a FlexPod unit is built, it can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units).

Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for all virtualization solutions. VMware vSphere 5.0 Built On FlexPod includes NetApp storage, Cisco networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere® software in a



Corporate Headquarters:

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

Copyright 2012

single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. This is why the reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an FCoE-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the VMware vSphere 5.0 Built On FlexPod components and the network connections for a configuration with FCoE-based storage. This design uses the Cisco Nexus® 5548UP, Cisco Nexus 2232PP FEX, and Cisco UCS C-Series and B-Series with the Cisco UCS virtual interface card (VIC) and the NetApp FAS family of storage controllers connected in a highly available design using Cisco Virtual Port Channels (vPCs). This infrastructure is deployed to provide FCoE-booted hosts with file- and block-level access to shared storage datastores. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture—be it FC, FCoE, or 10GbE—no recabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 1 *VMware vSphere Built on FlexPod Components with FCoE Boot Topology*


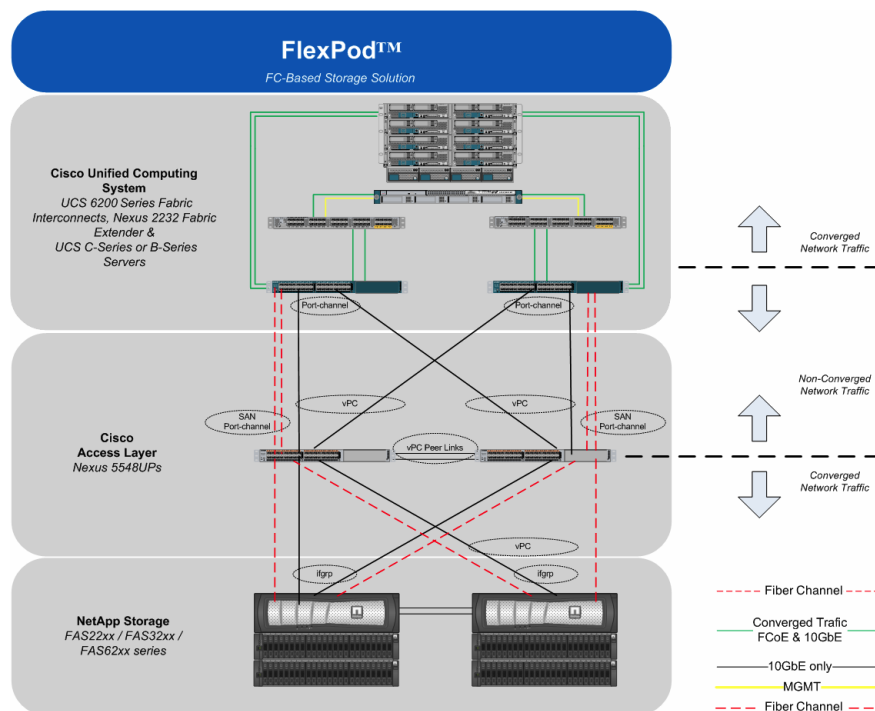


Figure 2 VMware vSphere Built on FlexPod with FC Boot Topology

The reference configuration includes:

- Two Cisco Nexus 5548UP switches
- Two Cisco Nexus 2232PP fabric extenders
- Two Cisco UCS 6248UP fabric interconnects
- Support for 16 Cisco UCS C-Series servers without any additional networking components
- Support for 8 Cisco UCS B-Series servers without any additional blade server chassis
- Support for hundreds of Cisco UCS C-Series and B-Series servers by way of additional fabric extenders and blade server chassis
- One NetApp FAS3240-A (HA pair)

Storage is provided by a NetApp FAS3240-A (HA configuration in a single chassis). All system and network links feature redundancy, providing end-to-end high availability (HA). For server virtualization, the deployment includes VMware vSphere. Although this is the base design, each of the components can be scaled flexibly to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capacity and throughput, and special hardware or software features can be added to introduce new capabilities.

This document guides you through the low-level steps for deploying the base architecture, as shown in [Figure 1](#). These procedures cover everything from physical cabling to compute and storage configuration to configuring virtualization with VMware vSphere.

FlexPod Benefits

One of the founding design principles of the FlexPod architecture is flexibility. Previous FlexPod architectures have highlighted FCoE- or FC-based storage solutions as well as IP-based storage solutions in addition to showcasing a variety of application workloads. This particular FlexPod architecture is a predesigned configuration that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of data center switches, NetApp FAS storage components, and VMware virtualization software. FlexPod is a base configuration, but it can scale up for greater performance and capacity, and it can scale out for environments that require consistent, multiple deployments. FlexPod has the flexibility to be sized and optimized to accommodate many different use cases. These use cases can be layered on an infrastructure that is architected based on performance, availability, and cost requirements.

FlexPod is a platform that can address current virtualization needs and simplify the evolution to an IT as a service (ITaaS) infrastructure. The VMware vSphere 5.0 Built On FlexPod solution can help improve agility and responsiveness, reduce total cost of ownership (TCO), and increase business alignment and focus.

This document focuses on deploying an infrastructure that is capable of supporting VMware vSphere, VMware vCenter™ with NetApp plug-ins, and NetApp OnCommand® as the foundation for virtualized infrastructure. Additionally, this document details a use case for those who want to design an architecture with FCoE-based storage using storage protocols such as iSCSI, FCoE, CIFS, and NFS. For a detailed study of several practical solutions deployed on FlexPod, refer to the [NetApp Technical Report 3884: FlexPod Solutions Guide](#).

Benefits of the Cisco Unified Computing System

Cisco Unified Computing System is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud computing environments.

The system's x86-architecture rack-mount and blade servers are powered by Intel® Xeon® processors. These industry-standard servers deliver world-record performance to power mission-critical workloads. Cisco servers, combined with a simplified, converged architecture, drive better IT productivity and superior price/performance for lower total cost of ownership (TCO). Building on Cisco's strength in enterprise networking, Cisco's Unified Computing System is integrated with a standards-based, high-bandwidth, low-latency, virtualization-aware unified fabric. The system is wired once to support the desired bandwidth and carries all Internet protocol, storage, inter-process communication, and virtual machine traffic with security isolation, visibility, and control equivalent to physical networks. The system meets the bandwidth demands of today's multicore processors, eliminates costly redundancy, and increases workload agility, reliability, and performance.

Cisco Unified Computing System is designed from the ground up to be programmable and self integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards, even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time. With model-based management, administrators manipulate a model of a desired system configuration, associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender technology reduces the number of system components to purchase, configure, manage, and maintain by condensing three network layers into one. This represents a radical simplification over traditional systems, reducing capital and operating costs while increasing business agility, simplifying and speeding deployment, and improving performance.

Cisco Unified Computing System helps organizations go beyond efficiency: it helps them become more effective through technologies that breed simplicity rather than complexity. The result is flexible, agile, high-performance, self-integrating information technology, reduced staff costs with increased uptime through automation, and more rapid return on investment.

This reference architecture highlights the use of the Cisco UCS C200-M2, Cisco UCS C220-M3, Cisco UCS B200-M2, Cisco UCS B200-M3, the Cisco UCS 6248UP, and the Nexus 2232PP FEX to provide a resilient server platform balancing simplicity, performance and density for production-level virtualization. Also highlighted in this architecture, is the use of Cisco UCS service profiles that enable FCoE boot of the native operating system. Coupling service profiles with unified storage delivers on demand stateless computing resources in a highly scalable architecture.

Recommended support documents include:

- Cisco Unified Computing System:
<http://www.cisco.com/en/US/products/ps10265/index.html>
- Cisco Unified Computing System C-Series Servers:
<http://www.cisco.com/en/US/products/ps10493/index.html>
- Cisco Unified Computing System B-Series Servers:
<http://www.cisco.com/en/US/products/ps10280/index.html>

Benefits of Cisco Nexus 5548UP

The Cisco Nexus 5548UP Switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages, including:

Architectural Flexibility

- Unified ports that support traditional Ethernet, Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE)
- Synchronizes system clocks with accuracy of less than one microsecond, based on IEEE 1588
- Offers converged Fabric extensibility, based on emerging standard IEEE 802.1BR, with Fabric Extender (FEX) Technology portfolio, including:
 - Cisco Nexus 2000 FEX
 - Adapter FEX
 - VM-FEX

Infrastructure Simplicity

- Common high-density, high-performance, data-center-class, fixed-form-factor platform
- Consolidates LAN and storage
- Supports any transport over an Ethernet-based fabric, including Layer 2 and Layer 3 traffic
- Supports storage traffic, including iSCSI, NAS, FC, RoE, and IB over Ethernet

- Reduces management points with FEX Technology

Business Agility

- Meets diverse data center deployments on one platform
- Provides rapid migration and transition for traditional and evolving technologies
- Offers performance and scalability to meet growing business needs

Specifications At-a Glance

- A 1 -rack-unit, 1/10 Gigabit Ethernet switch
- 32 fixed Unified Ports on base chassis and one expansion slot totaling 48 ports
- The slot can support any of the three modules: Unified Ports, 1/2/4/8 native Fibre Channel, and Ethernet or FCoE
- Throughput of up to 960 Gbps

This reference architecture highlights the use of the Cisco Nexus 5548UP. As mentioned, this platform is capable of serving as the foundation for wire-once, unified fabric architectures. This document provides guidance for an architecture capable of delivering storage protocols including iSCSI, FC, FCoE, CIFS, and NFS. With the storage protocols license enabled on the Nexus 5548UP, end-users can take advantage of a full feature SAN switch as well as traditional Ethernet networking.

Recommended support documents include:

- Cisco Nexus 5000 Family of switches: <http://www.cisco.com/en/US/products/ps9670/index.html>

Benefits of the NetApp FAS Family of Storage Controllers

The NetApp Unified Storage Architecture offers customers an agile and scalable storage platform. All NetApp storage systems use the Data ONTAP® operating system to provide SAN (FCoE, FC, iSCSI), NAS (CIFS, NFS), and primary and secondary storage in a single unified platform so that all workloads can be hosted on the same storage array. A single process for activities such as installation, provisioning, mirroring, backup, and upgrading is used throughout the entire product line, from the entry level to enterprise-class controllers. Having a single set of software and processes simplifies even the most complex enterprise data management challenges. Unifying storage and data management software and processes streamlines data ownership, enables companies to adapt to their changing business needs without interruption, and reduces total cost of ownership.

This reference architecture focuses on the use case of using FCoE-based storage to solve customers' challenges and meet their needs in the data center. Specifically this entails FCoE boot of Cisco UCS hosts; provisioning of virtual machine datastores by using NFS; and application access through FCoE, iSCSI, CIFS, or NFS, all while using NetApp unified storage.

In a shared infrastructure, the availability and performance of the storage infrastructure are critical because storage outages or performance issues can affect thousands of users. The storage architecture must provide a high level of availability and performance. For detailed documentation about best practices, NetApp and its technology partners have developed a variety of best practice documents.

This reference architecture highlights the use of the NetApp FAS3200 product line, specifically the FAS3240-A with an FCoE adapter card, Flash Cache, and SAS storage.

Recommended support documents include:

- NetApp storage systems: www.netapp.com/us/products/storage-systems/
- NetApp FAS3200 storage systems: www.netapp.com/us/products/storage-systems/fas3200/

- NetApp TR-3437: Storage Best Practices and Resiliency Guide
- NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines
- NetApp TR-3749: NetApp and VMware vSphere Storage Best Practices
- NetApp TR-3884: FlexPod Solutions Guide
- NetApp TR-3824: MS Exchange 2010 Best Practices Guide

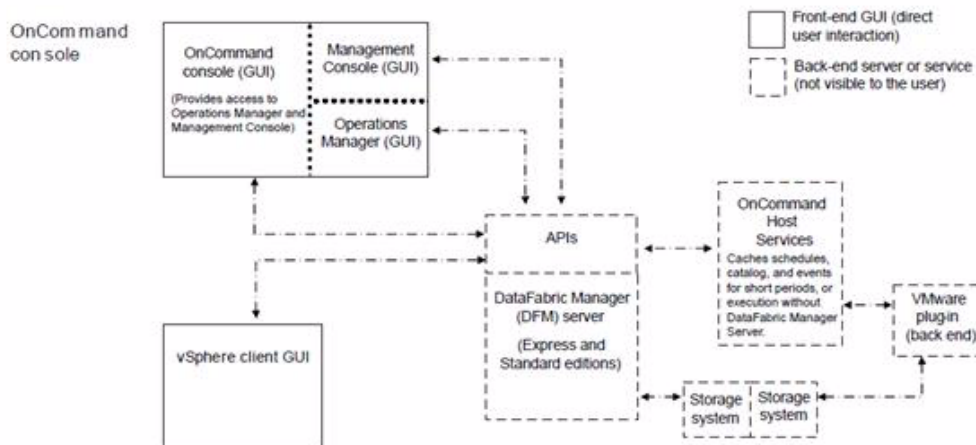
Benefits of NetApp OnCommand Unified Manager Software

NetApp OnCommand management software delivers efficiency savings by unifying storage operations, provisioning, and protection for both physical and virtual resources. The key product benefits that creates this value include:

- **Simplicity.** A single unified approach and a single set of tools to manage both the physical world and the virtual world as you move to a services model to manage your service delivery. This makes NetApp the most effective storage for the virtualized data center. It has a single configuration repository for reporting, event logs, and audit logs.
- **Efficiency.** Automation and analytics capabilities deliver storage and service efficiency, reducing IT capex and opex spend by up to 50%.
- **Flexibility.** With tools that let you gain visibility and insight into your complex multiprotocol, multivendor environments and open APIs that let you integrate with third-party orchestration frameworks and hypervisors, OnCommand offers a flexible solution that helps you rapidly respond to changing demands.

OnCommand gives you visibility across your storage environment by continuously monitoring and analyzing its health. You get a view of what is deployed and how it is being used, enabling you to improve your storage capacity utilization and increase the productivity and efficiency of your IT administrators. And this unified dashboard gives at-a-glance status and metrics, making it far more efficient than having to use multiple resource management tools.

Figure 3 OnCommand Architecture



OnCommand Host Package

You can discover, manage, and protect virtual objects after installing the NetApp OnCommand Host Package software. The components that make up the OnCommand Host Package are:

- **OnCommand host service VMware plug-in.** A plug-in that receives and processes events in a VMware environment, including discovering, restoring, and backing up virtual objects such as virtual machines and datastores. This plug-in executes the events received from the host service.
- **Host service.** The host service software includes plug-ins that enable the NetApp DataFabric® Manager server to discover, back up, and restore virtual objects, such as virtual machines and datastores. The host service also enables you to view virtual objects in the OnCommand console. It enables the DataFabric Manager server to forward requests, such as the request for a restore operation, to the appropriate plug-in, and to send the final results of the specified job to that plug-in. When you make changes to the virtual infrastructure, automatic notification is sent from the host service to the DataFabric Manager server. You must register at least one host service with the DataFabric Manager server before you can back up or restore data.
- **Host service Windows PowerShell™ cmdlets.** Cmdlets that perform virtual object discovery, local restore operations, and host configuration when the DataFabric Manager server is unavailable.

Management tasks performed in the virtual environment by using the OnCommand console include:

- Create a dataset and then add virtual machines or datastores to the dataset for data protection.
- Assign local protection and, optionally, remote protection policies to the dataset.
- View storage details and space details for a virtual object.
- Perform an on-demand backup of a dataset.
- Mount existing backups onto an ESX® server to support tasks such as backup verification, single file restore, and restoration of a virtual machine to an alternate location.
- Restore data from local and remote backups as well as restoring data from backups made before the introduction of OnCommand management software.
- View storage details and space details for a virtual object.

Storage Service Catalog

The Storage Service Catalog, a component of OnCommand, is a key NetApp differentiator for service automation. It lets you integrate storage provisioning policies, data protection policies, and storage resource pools into a single service offering that administrators can choose when provisioning storage. This automates much of the provisioning process, and it also automates a variety of storage management tasks associated with the policies.

The Storage Service Catalog provides a layer of abstraction between the storage consumer and the details of the storage configuration, creating "storage as a service." The service levels defined with the Storage Service Catalog automatically specify and map policies to the attributes of your pooled storage infrastructure. This higher level of abstraction between service levels and physical storage lets you eliminate complex, manual work, encapsulating storage and operational processes together for optimal, flexible, dynamic allocation of storage.

The service catalog approach also incorporates the use of open APIs into other management suites, which leads to a strong ecosystem integration.

FlexPod Management Solutions

The FlexPod platform provides open APIs for easy integration with a broad range of management tools. NetApp and Cisco work with trusted partners to provide a variety of management solutions. Products designated as Validated FlexPod Management Solutions must pass extensive testing in Cisco and NetApp labs against a broad set of functional and design requirements. Validated solutions for automation and orchestration provide unified, turnkey functionality. Now you can deploy IT services in minutes instead of weeks by reducing complex, multiadministrator processes to repeatable workflows that are easily adaptable. The following list details the current vendors for these solutions.



Note

Some of the following links are available only to partners and customers.

- CA
<http://solutionconnection.netapp.com/CA-Infrastructure-Provisioning-for-FlexPod.aspx>
<http://www.youtube.com/watch?v=mmkNUvVZY94>
- Cloupia
<http://solutionconnection.netapp.com/cloupia-unified-infrastructure-controller.aspx>
<http://www.cloupia.com/en/flexpodtoclouds/videos/Cloupia-FlexPod-Solution-Overview.html>
- Gale Technologies
<http://solutionconnection.netapp.com/galeforce-turnkey-cloud-solution.aspx>
<http://www.youtube.com/watch?v=y1f81zjFF0>

Products designated as FlexPod Management Solutions have demonstrated the basic ability to interact with all components of the FlexPod platform. Vendors for these solutions currently include BMC Software Business Service Management, Cisco Intelligent Automation for Cloud, DynamicOps, FireScope, Nimsoft, and Zenoss. Recommended documents include:

- <https://solutionconnection.netapp.com/flexpod.aspx>
- <http://www.netapp.com/us/communities/tech-ontap/tot-building-a-cloud-on-flexpod-1203.html>

Benefits of VMware vSphere with the NetApp Virtual Storage Console

VMware vSphere, coupled with the NetApp Virtual Storage Console (VSC), serves as the foundation for VMware virtualized infrastructures. vSphere 5.0 offers significant enhancements that can be employed to solve real customer problems. Virtualization reduces costs and maximizes IT efficiency, increases application availability and control, and empowers IT organizations with choice. VMware vSphere delivers these benefits as the trusted platform for virtualization as demonstrated by its contingent of more than 300,000 customers worldwide.

VMware vCenter Server is the best way to manage and use the power of virtualization. A vCenter domain manages and provisions resources for all the ESX hosts in the given data center. The ability to license various features in vCenter at differing price points allows customers to choose the package that best serves their infrastructure needs.

The VSC is a vCenter plug-in that provides end-to-end virtual machine (VM) management and awareness for VMware vSphere environments running on top of NetApp storage. The following core capabilities make up the plug-in:

- Storage and ESXi™ host configuration and monitoring by using Monitoring and Host Configuration

- Datastore provisioning and VM cloning by using Provisioning and Cloning
- Backup and recovery of VMs and datastores by using Backup and Recovery
- Online alignment and single and group migrations of VMs into new or existing VMFS datastores by using Optimization and Migration

Because the VSC is a vCenter plug-in, all vSphere clients that connect to vCenter can access VSC. This availability is different from a client-side plug-in that must be installed on every vSphere client.

Software Revisions

It is important to note the software versions used in this document. [Table 1](#) details the software revisions used throughout this document.

Table 1 *Software Revisions*

| Layer | Compute | Version or Release | Details |
|----------|--------------------------------------|---------------------------|--|
| Compute | Cisco UCS Fabric Interconnect | 2.0(4a) | Embedded management |
| | Cisco UCS C 220 M3 | 2.0(4a) | Software bundle release |
| | Cisco UCS B 200 M3 | 2.0(4a) | Software bundle release |
| | Cisco E-Nic | 2.1.2.22 | Ethernet driver for Cisco VIC |
| | Cisco F-Nic | 1.5.0.8 | FCoE driver for Cisco VIC |
| Network | Cisco Nexus Fabric Switch | 5.1(3)N2(1b) | Operating system version |
| Storage | NetApp FAS3240-A | Data ONTAP 8.1 | Operating system version |
| Software | Cisco UCS Hosts | VMware vSphere ESXi 5.0 | Operating system version |
| | Microsoft® .NET Framework | 3.5.1 | Feature enabled within Windows® operating system |
| | Microsoft SQL Server® | MS SQL Server 2008 R2 SP1 | VM (1 each): SQL Server DB |
| | VMware vCenter | 5.0 | VM (1 each): VMware vCenter |
| | NetApp OnCommand | 5.0 | VM (1 each): OnCommand |
| | NetApp Virtual Storage Console (VSC) | 4.0 | Plug-in within VMware vCenter |
| | Cisco Nexus 1010-x | 4.2.1.SP1.4 | Virtual Services Appliance |
| | Cisco Nexus 1000v | 4.2.1.SV1.5.1a | Virtual Services Blade within the 1010-x |

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with IP-based storage. Therefore, reference is made to which component is being configured with each step, either A or B. For example, controller A and controller B are used to identify the two NetApp storage controllers that are provisioned with this document, and Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, **<text>** appears as part of the command structure. See the following example for the vlan create command:

```
controller A> vlan create
```

Usage:

```
vlan create [-g {on|off}] <ifname> <vlanid_list>
vlan add <ifname> <vlanid_list>
vlan delete -g <ifname> [<vlanid_list>]
vlan modify -g {on|off} <ifname>
vlan stat <ifname> [<vlanid_list>]
```

Example:

```
controller A> vlan create vif0 <management VLAN ID>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. [Table 2](#) describes the VLANs necessary for deployment as outlined in this guide. The VM-Mgmt VLAN is used for management interfaces of the VMware vSphere hosts. [Table 3](#) lists the VSANs necessary for deployment as outlined in this guide.

[Table 4](#) lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.



Note

If you use separate in-band and out-of-band management VLANs, you must create a Layer 3 route between these VLANs. For this validation, a common management VLAN was used.

Table 2 **Necessary VLANs**

| VLAN Name | VLAN Purpose | ID Used in Validating This Document |
|------------------|---|-------------------------------------|
| Mgmt in band | VLAN for in-band management interfaces | 3175 |
| Mgmt out of band | VLAN for out-of-band management interfaces | 3175 |
| Native | VLAN to which untagged frames are assigned | 2 |
| NFS | VLAN for NFS traffic | 3170 |
| FCoE-A | VLAN for FCoE traffic for fabric A | 101 |
| FCoE-B | VLAN for FCoE traffic for fabric B | 102 |
| vMotion | VLAN designated for the movement of VMs from one physical host to another | 3173 |
| VM Traffic | VLAN for VM application traffic | 3174 |
| Packet Control | VLAN for Packet Control traffic | 3176 |

Table 3 **Necessary VSANs**

| VSAN Name | VSAN Purpose | ID Used in Validating This Document |
|-----------|---|-------------------------------------|
| VSAN A | VSAN for fabric A traffic. ID matches FCoE-A VLAN | 101 |
| VSAN B | VSAN for fabric B traffic. ID matches FCoE-B VLAN | 102 |

Table 4 **Configuration Variables**

| Variable | Description | Variable Value in Our Reference Implementation |
|------------------------|---|--|
| <# of disks> | Number of disks to assign to a controller | 24 |
| <necessary licenses> | License needed for a controller | cf, fcp, nfs, Flex_clone |
| <NTP server IP> | NTP server IP | 192.168.175.4 |
| <# of disks for aggr1> | Number of disks to assign to aggr1 | 19 |

| | | |
|---|---|------------------------|
| <ntap SNMP request role> | Creates the SNMP request role | snmpv3role |
| <ntap SNMP managers> | SNMP management group | snmpv3group |
| <ntap SNMP users> | Creates an SNMP user | snmpv3user |
| <ntap SNMP community> | SNMP community | icefxp7-cmtty |
| <ntap admin email address> | SNMP admin e-mail address | JaneDoe@Netapp.com |
| <ntap SNMP site name> | SNMP site name | RTP, Building 1, Lab 3 |
| <NFS VLAN ID> | VLAN for NFS traffic | 3170 |
| <Controller A NFS IP> | Controller A NFS IP | 192.168.170.96 |
| <NFS Netmask> | NFS netmask | 255.255.255.0 |
| <Controller B NFS IP> | Controller B NFS IP | 192.168.170.97 |
| <ESXi Host 1 NFS IP> | ESXi host 1 NFS IP | 192.168.170.98 |
| <ESXi Host 2 NFS IP> | ESXi host 2 NFS IP | 192.168.170.99 |
| <Nexus A Switch name> | Cisco Nexus switch A host name | ice5548-1 |
| <Nexus A mgmt0 IP> | Cisco Nexus switch A mgmt IP | 192.168.175.69 |
| <Nexus A mgmt0 netmask> | Cisco Nexus switch A MGMT 0 netmask | 255.255.255.0 |
| <Nexus A mgmt0 gateway> | Cisco Nexus switch A MGMT 0 gateway | 192.168.175.1 |
| <Nexus B Switch name> | Cisco Nexus switch B host name | ice5548-2 |
| <Nexus B mgmt0 IP> | Cisco Nexus switch B mgmt IP | 192.168.175.70 |
| <Nexus B mgmt0 netmask> | Cisco Nexus switch B MGMT 0 netmask | 255.255.255.0 |
| <Nexus B mgmt0 gateway> | Cisco Nexus switch B MGMT 0 gateway | 192.168.175.1 |
| <MGMT VLAN ID>. | VLAN ID for in-band and out-of-band management interfaces | 3175 |
| <Native VLAN ID> | VLAN ID for native VLAN | 2 |
| <Packet Control VLAN ID> | VLAN ID for the packet control VLAN | 3176 |
| <vMotion VLAN ID> | VLAN for vMotion® traffic | 3173 |
| <VM-Traffic VLAN ID> | VLAN for VM application traffic | 3174 |

| | | |
|---|---|--------------------------------|
| <Nexus vPC domain ID> | Cisco Nexus vPC domain ID | 23 |
| <Fabric A FCoE VLAN ID> | Fabric A FCoE VLAN ID | 101 |
| <VSAN A ID> | VSAN A ID | 101 |
| <Fabric B FCoE VLAN ID> | Fabric B FCoE VLAN ID | 102 |
| <VSAN B ID> | VSAN B ID | 102 |
| <Controller A 2a WWPN> | World Wide Port Name (WWPN) for controller A_2a | Queried from storage |
| <Controller B 2a WWPN> | World Wide Port Name (WWPN) for controller B_2a | Queried from storage |
| <Controller A 2b WWPN> | World Wide Port Name (WWPN) for controller A_2b | Queried from storage |
| <Controller B 2b WWPN> | World Wide Port Name (WWPN) for controller B_2b | Queried from storage |
| <VM-Host-Infra-01 vHBA_A WWPN> | World Wide Port Name (WWPN) for vHBA A | Queried from Cisco UCS Manager |
| <VM-Host-Infra-01 vHBA_B WWPN> | World Wide Port Name (WWPN) for vHBA B | Queried from Cisco UCS Manager |
| <VM-Host-Infra-01 IP address>. | Infra 01 VM IP | 192.168.175.99 |
| <VM-Host-Infra-02 IP address>. | Infra 02 VM IP | 192.168.175.100 |
| <root password> | root password for your ESXi environment | ***** |
| <password> | vpxuser password | ***** |
| <Primary VSM IP Address> | VSM primary IP address | 192.165.175.193 |
| <license filename> | File name of bootflash license key | Will come from Cisco |
| <Storage Controller A> | Name of storage controller A | ice3240-1a |
| <Storage Controller B> | Name of storage controller B | ice3240-1b |
| <global ssl country> | SSL country for DFM install | US |
| <global ssl state> | SSL state for DFM install | "North Carolina" |
| <global ssl locality> | SSL locality for DFM install | RTP |

| | | |
|---|--|---|
| <global ssl org> | SSL organization for DFM install | NetApp |
| <global ssl org unit> | SSL organization unit for DFM install | ICE |
| <global ntap dfm hostname> | Global NetApp DFM host name for DFM install | icefxp1-vsc-oc |
| <ntap admin email address> | NetApp admin e-mail address for DFM install | JaneDoe@netapp.com |
| <ntap SNMP password> | DFM SNMP password | ***** |
| <ntap autosupport mailhost> | Local mail server | mail@mailhost.com |
| <ntap A hostname> | DFM host name of controller A | Ice3240-1a |
| <ntap B hostname> | DFM host name of controller B | Ice3240-1b |
| <global default password> | Global password | ***** |
| <ntap SNMP traphosts> | SNMP traphost name | icefxp1-vsc-oc |
| <OC_install_directory> | Installation folder for the OnCommand host package | "C:\Program Files\NetApp\OnCommand Host Package" |
| <VSC installation directory> | Installation location of VSC | "C:\Program Files\NetApp\Virtual Storage Console" |
| <DFM_Install_dir> | Installation location of DFM | "C:\Program Files\NetApp\DFM" |

**Note**

In this document, management IPs and host names must be assigned for the following components:

- NetApp storage controllers A and B
- Cisco UCS fabric Interconnects A and B and the Cisco UCS cluster
- Cisco Nexus 5548s A and B
- VMware ESXi hosts
- VMware vCenter SQL Server virtual machine
- VMware vCenter virtual machine
- NetApp Virtual Storage Console or OnCommand virtual machine

For all host names except the virtual machine host names, the IP addresses must be preconfigured in the local DNS server. Additionally, the NFS IP addresses of the NetApp storage systems are used to monitor the storage systems from OnCommand DataFabric Manager. In this validation, a management host name was assigned to each storage controller (that is, ice3240-1a-m) and provisioned in DNS. A host name was also assigned for each controller in the NFS VLAN (that is, ice3240-1a) and provisioned in DNS. This NFS VLAN host name was then used when the storage system was added to OnCommand Data Fabric Manager.

Deployment

This document describes the steps to deploy base infrastructure components as well to provision VMware vSphere as the foundation for virtualized workloads. When you finish these deployment steps, you will be prepared to provision applications on top of a VMware virtualized infrastructure. The outlined procedure contains the following steps:

- Initial NetApp controller configuration
- Initial Cisco UCS configuration
- Initial Cisco Nexus configuration
- Creation of necessary VLANs for management, basic functionality, and virtualized infrastructure specific to VMware
- Creation of necessary VSANs for booting of the Cisco UCS hosts
- Creation of necessary vPCs to provide high availability among devices
- Creation of necessary service profile pools: MAC, UUID, server, and so forth
- Creation of necessary service profile policies: adapter, boot, and so forth
- Creation of two service profile templates from the created pools and policies: one each for fabric A and B
- Provisioning of two servers from the created service profiles in preparation for OS installation
- Initial configuration of the infrastructure components residing on the NetApp controller
- Installation of VMware vSphere 5.0
- Installation and configuration of VMware vCenter
- Enabling of NetApp Virtual Storage Console (VSC)
- Configuration of NetApp OnCommand
- Configuration of NetApp vStorage APIs for Storage Awareness (VASA) Provider

The VMware vSphere built on FlexPod architecture is flexible; therefore the configuration detailed in this section can vary for customer implementations, depending on specific requirements. Although customer implementations can deviate from the following information, the best practices, features, and configurations described in this section should be used as a reference for building a customized VMware vSphere built on FlexPod solution.

Cabling Information

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain details for the prescribed and supported configuration of the NetApp FAS3240-A running Data ONTAP 8.1. This configuration uses a dual-port FCoE adapter, built-in FC ports, and external SAS disk shelves. For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps

Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

It is possible to order a FAS3240-A system in a different configuration from what is prescribed in the tables in this section. Before starting, be sure that the configuration matches the descriptions in the tables and diagrams in this section.

Figure 4 shows a FlexPod cabling diagram. The labels indicate connections to endpoints rather than port numbers on the physical device. For example, connection 1 is an FCoE port connected from NetApp controller A to Cisco Nexus 5548 A. SAS connections 27, 28, 29, and 30 as well as ACP connections 31 and 32 should be connected to the NetApp storage controller and disk shelves according to best practices for the specific storage controller and disk shelf quantity. Additionally, this paper assumes the FCoE adapter is installed in the second PCI slot of the controller. If the FCoE card is not installed in slot 2 but is installed in a slot according to best practices based on the quantity of add-on cards, modify the card name as appropriate.

**Note**

Cables Necessary for FC configuration are optional when implementing FCoE, and are identified by **green text**. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide at https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

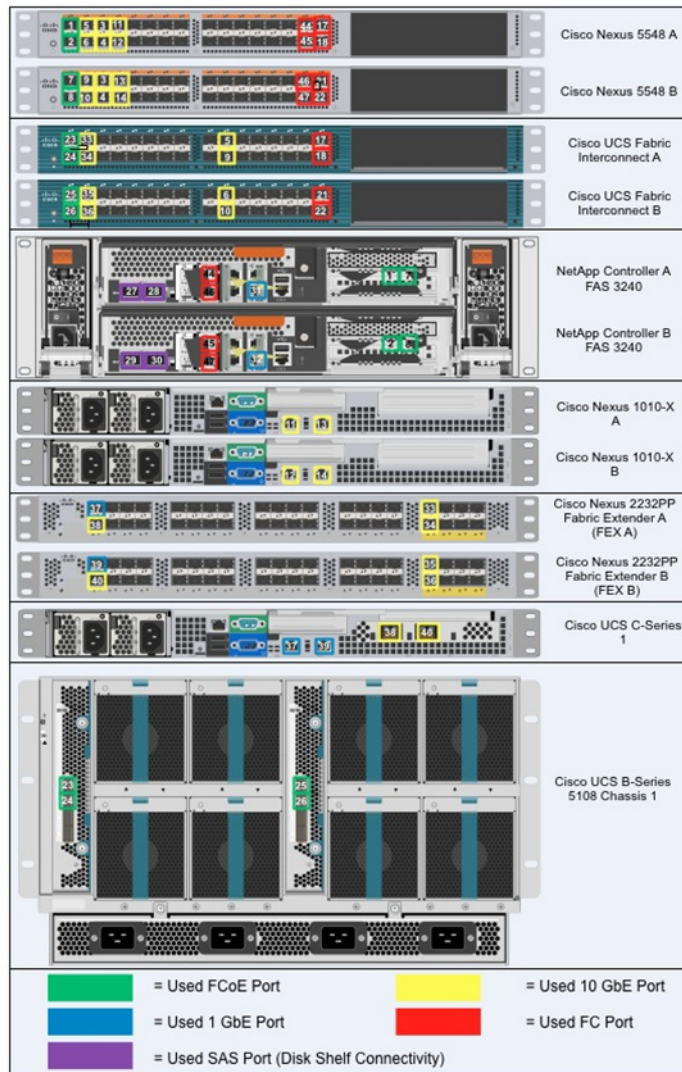
Figure 4 **Cabling**

Table 5 *Cisco Nexus 5548 A Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 5548 A | Eth1/1 | FCoE | NetApp controller A | e2a |
| | Eth1/2 | FCoE | NetApp controller B | e2a |
| | Eth1/3 | 10GbE | Cisco UCS fabric interconnect A | Eth1/19 |
| | Eth1/4 | 10GbE | Cisco UCS fabric interconnect B | Eth1/19 |
| | Eth1/5 | 10GbE | Cisco Nexus 5548 B | Eth1/5 |
| | Eth1/6 | 10GbE | Cisco Nexus 5548 B | Eth1/6 |
| | Eth1/7 | 10GbE | Cisco Nexus 1010-X A | LOM 1 |
| | Eth1/8 | 10GbE | Cisco Nexus 1010-X B | LOM 1 |
| | MGMT0 | 100MbE | 100MbE management switch | Any |

**Note**

For devices requiring GbE connectivity, use the GbE copper SFP+s (GLC-T=).

Table 6 *Cisco Nexus 5548 B Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 5548 B | Eth1/1 | FCoE | NetApp controller A | e2b |
| | Eth1/2 | FCoE | NetApp controller B | e2b |
| | Eth1/3 | 10GbE | Cisco UCS fabric interconnect A | Eth1/20 |
| | Eth1/4 | 10GbE | Cisco UCS fabric interconnect B | Eth1/20 |
| | Eth1/5 | 10GbE | Cisco Nexus 5548 A | Eth1/5 |
| | Eth1/6 | 10GbE | Cisco Nexus 5548 A | Eth1/6 |
| | Eth1/7 | 10GbE | Cisco Nexus 1010-X A | LOM2 |
| | Eth1/8 | 10GbE | Cisco Nexus 1010-X B | LOM2 |
| | MGMT0 | 100MbE | 100MbE management switch | Any |

**Note**

For devices requiring GbE connectivity, use the GbE copper SFP+s (GLC-T=).

Table 7 NetApp Controller A Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------|-------------|--------------|--------------------------|-------------|
| NetApp controller A | wrench port | 100MbE | 100MbE management switch | Any |
| | e0P | GbE | SAS shelves | ACP port |
| | 0a | SAS | SAS shelves | SAS |
| | 0b | SAS | SAS shelves | SAS |
| | e2a | 10GbE / FCoE | Cisco Nexus 5548 A | Eth1/1 |
| | e2b | 10GbE / FCoE | Cisco Nexus 5548 B | Eth1/1 |
| | 0c* | FC | Cisco Nexus 5548 A | FC1/29 |
| | 0d* | FC | Cisco Nexus 5548 B | FC1/29 |

Table 8 NetApp Controller B Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------|-------------|--------------|--------------------------|-------------|
| NetApp controller B | wrench port | 100MbE | 100MbE management switch | Any |
| | e0P | GbE | SAS shelves | ACP port |
| | 0a | SAS | SAS shelves | SAS |
| | 0b | SAS | SAS shelves | SAS |
| | e2a | 10GbE / FCoE | Cisco Nexus 5548 A | Eth1/2 |
| | e2b | 10GbE / FCoE | Cisco Nexus 5548 B | Eth1/2 |
| | 0c* | FC | Cisco Nexus 5548 A | FC1/30 |
| | 0d* | FC | Cisco Nexus 5548 B | FC1/30 |

**Note**

Connections denoted with * necessary for FC configuration only.

Table 9 Cisco UCS Fabric Interconnect A Ethernet Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------------------|------------|------------|------------------------------------|-------------|
| Cisco UCS fabric interconnect A | Eth1/19 | 10GbE | Cisco Nexus 5548 A | Eth1/3 |
| | Eth1/20 | 10GbE | Cisco Nexus 5548 B | Eth1/3 |
| | Eth1/1 | FCoE/10GbE | Cisco UCS B-Series chassis 1 FEX A | Port 1 |
| | Eth1/2 | FCoE/10GbE | Cisco UCS B-Series chassis 1 FEX A | Port 2 |
| | Eth1/3 | 10GbE | Cisco Nexus 2232PP FEX A | Eth2/1 |
| | Eth1/4 | 10GbE | Cisco Nexus 2232PP FEX A | Eth2/2 |
| | MGMT0 | 100MbE | 100MbE management switch | Any |
| | L1 | GbE | Cisco UCS fabric interconnect B | L1 |
| | L2 | GbE | Cisco UCS fabric interconnect B | L2 |

Table 10 Cisco UCS Fabric interconnect B Ethernet Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------------------|------------|------------|------------------------------------|-------------|
| Cisco UCS fabric interconnect B | Eth1/19 | 10GbE | Cisco Nexus 5548 A | Eth1/4 |
| | Eth1/20 | 10GbE | Cisco Nexus 5548 B | Eth1/4 |
| | Eth1/1 | 10GbE/FCoE | Cisco UCS B-Series chassis 1 FEX B | Port 1 |
| | Eth1/2 | 10GbE/FCoE | Cisco UCS B-Series chassis 1 FEX B | Port 2 |
| | Eth1/3 | 10GbE | Cisco Nexus 2232PP FEX B | Eth2/1 |
| | Eth1/4 | 10GbE | Cisco Nexus 2232PP FEX B | Eth2/2 |
| | MGMT0 | 100MbE | 100MbE management switch | Any |
| | L1 | GbE | Cisco UCS fabric interconnect A | L1 |
| | L2 | GbE | Cisco UCS fabric interconnect A | L2 |

Table 11 Cisco Nexus 1010-X A Ethernet Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|----------------------|------------|------------|--------------------|-------------|
| Cisco Nexus 1010-X A | LOM1 | 10GbE | Cisco Nexus 5548 A | Eth1/7 |
| | LOM2 | 10GbE | Cisco Nexus 5548 B | Eth1/7 |

Table 12 *Cisco Nexus 1010-X B Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|----------------------|------------|------------|--------------------|-------------|
| Cisco Nexus 1010-X B | LOM 1 | 10GbE | Cisco Nexus 5548 A | Eth1/8 |
| | LOM2 | 10GbE | Cisco Nexus 5548 B | Eth1/8 |

Table 13 *Cisco Nexus 2232PP Fabric Extender A (FEX A) Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 2232PP FEX A | Eth1/1 | GbE | Cisco UCS C-Series 1 | M1 |
| | Eth1/2 | 10GbE | Cisco UCS C-Series 1 | Port 0 |
| | Eth2/1 | 10GbE | Cisco UCS fabric interconnect A | Eth1/3 |
| | Eth2/2 | 10GbE | Cisco UCS fabric interconnect A | Eth1/4 |

Table 14 *Cisco Nexus 2232PP Fabric Extender B (FEX B) Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 2232PP FEX B | Eth1/1 | GbE | Cisco UCS C-Series 1 | M2 |
| | Eth1/2 | 10GbE | Cisco UCS C-Series 1 | Port 1 |
| | Eth2/1 | 10GbE | Cisco UCS fabric interconnect B | Eth1/3 |
| | Eth2/2 | 10GbE | Cisco UCS fabric interconnect B | Eth1/4 |

Table 15 *Cisco UCS C-Series 1 Ethernet Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|----------------------|------------|------------|--------------------------|-------------|
| Cisco UCS C-Series 1 | M1 | GbE | Cisco Nexus 2232PP FEX A | Eth1/1 |
| | M2 | GbE | Cisco Nexus 2232PP FEX B | Eth1/1 |
| | Port 0 | 10GbE | Cisco Nexus 2232PP FEX A | Eth1/2 |
| | Port 1 | 10GbE | Cisco Nexus 2232PP FEX B | Eth1/2 |

Table 16 *Cisco Nexus 5548 A Fibre Channel Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 5548 A | FC1/29* | FC | Controller_A | 0c |
| | FC1/30* | FC | Controller_B | 0c |
| | FC1/31 | FC | Cisco UCS fabric interconnect A | Port 31 |
| | FC1/32 | FC | Cisco UCS fabric interconnect A | Port 32 |

Table 17 *Cisco Nexus 5548 B Fibre Channel Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|--------------------|------------|------------|---------------------------------|-------------|
| Cisco Nexus 5548 B | FC1/29* | FC | Controller_A | 0d |
| | FC1/30* | FC | Controller_B | 0d |
| | FC1/31 | FC | Cisco UCS fabric interconnect B | Port 31 |
| | FC1/32 | FC | Cisco UCS fabric interconnect B | Port 32 |

**Note**

Connections denoted with * necessary for FC configuration only.

Table 18 *Cisco UCS Fabric Interconnect A Fibre Channel Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------------------|------------|------------|--------------------|-------------|
| Cisco UCS fabric interconnect A | Port 31 | FC | Cisco Nexus 5548 A | FC1/31 |
| | Port 32 | FC | Cisco Nexus 5548 A | FC1/32 |

Table 19 *Cisco UCS Fabric Interconnect B Fibre Channel Cabling Information*

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---------------------------------|------------|------------|--------------------|-------------|
| Cisco UCS fabric interconnect B | Port 31 | FC | Cisco Nexus 5548 B | FC1/31 |
| | Port 32 | FC | Cisco Nexus 5548 B | FC1/32 |

NetApp FAS3240 Deployment Procedure: Part 1

This section provides a detailed procedure for configuring the NetApp FAS3240-A for use in a VMware vSphere built on FlexPod solution. These steps should be followed precisely. Failure to do so could result in an improper configuration.



Note

The configuration steps described in this section provide guidance for configuring the FAS3240-A running Data ONTAP 8.1.

Assign Controller Disk Ownership

These steps provide details for assigning disk ownership and disk initialization and verification.



Note

Typical best practices should be followed when determining the number of disks to assign to each controller head. You may choose to assign a disproportionate number of disks to a given storage controller in an HA pair, depending on the intended workload. In this reference architecture, half the total number of disks in the environment is assigned to one controller and the remainder to its partner. Divide the number of disks in half and use the result in the following command for <# of disks>.

Controller A

1. If the controller is at a **LOADER-A>** prompt, enter **autoboot** to start Data ONTAP. During controller boot, when prompted for the Boot Menu, press **Ctrl-C**.
2. At the menu prompt, select option **5** for Maintenance mode boot.
3. If prompted with **Continue to boot?** enter **Yes**.
4. Enter **ha-config show** to verify that the controllers and chassis configuration is **ha**.



Note

If either component is not in HA mode, use the ha-config modify command to put the components in HA mode.

5. Enter **disk show**. No disks should be assigned to the controller.
6. To determine the total number of disks connected to the storage system, enter **disk show -a**.
7. Enter **disk assign -n <# of disks>**.



Note

This reference architecture allocates half the disks to each controller. Workload design could dictate different percentages, however.

8. Enter **halt** to reboot the controller.
9. If the controller stops at a **LOADER-A>** prompt, enter **autoboot** to start Data ONTAP.
10. During controller boot, when prompted, press **Ctrl-C**.
11. At the menu prompt, select option **4** for “Clean configuration and initialize all disks.”
12. The installer asks if you want to zero the disks and install a new file system. Enter **y**.
13. A warning is displayed that this will erase all of the data on the disks. Enter **y** to confirm that this is what you want to do.

**Note**

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Controller B

1. If the controller is at a **LOADER-B>** prompt, enter **autoboot** to start Data ONTAP. During controller boot, when prompted, press **Ctrl-C** for the special boot menu.

**Note**

If you are using two controllers in separate chassis, the prompt will be at the **LOADER-A>** prompt.

2. At the menu prompt, select option **5** for Maintenance mode boot.
3. If prompted with Continue to boot? enter **Yes**.
4. Type **ha-config show** to verify that the controllers and chassis configuration is **ha**.

**Note**

If either controller is not in HA mode, use the **ha-config modify** command to put the components in HA mode.

5. Enter **disk show**. No disks should be assigned to this controller.
6. To determine the total number of disks connected to the storage system, enter **disk show -a**. This will now show the number of remaining unassigned disks connected to the controller.
7. Enter **disk assign -n <# of disks>** to assign the remaining disks to the controller.
8. Enter **halt** to reboot the controller.
9. If the controller stops at a **LOADER-B>** prompt, enter **autoboot** to start Data ONTAP.
10. During controller boot, when prompted, press **Ctrl-C** for the boot menu.
11. At the menu prompt, select option **4** for “Clean configuration and initialize all disks.”
12. The installer asks if you want to zero the disks and install a new file system. Enter **y**.
13. A warning is displayed that this will erase all of the data on the disks. Enter **y** to confirm that this is what you want to do.

**Note**

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Set Up Data ONTAP 8.1

These steps provide details for setting up Data ONTAP 8.1.

Controller A and Controller B

1. After the disk initialization and the creation of the root volume, Data ONTAP setup begins.
2. Enter the host name of the storage system.
3. Enter **n** for enabling IPv6.
4. Enter **y** for configuring interface groups.
5. Enter **1** for the number of interface groups to configure.

6. Name the interface **vif0**.
7. Enter **l** to specify the interface as LACP.
8. Enter **i** to specify IP load balancing.
9. Enter **2** for the number of links for **vif0**.
10. Enter **e2a** for the name of the first link.
11. Enter **e2b** for the name of the second link.
12. Press **Enter** to accept the blank IP address for **vif0**.
13. Enter **n** for interface group **vif0** taking over a partner interface.
14. Press **Enter** to accept the blank IP address for **e0a**.
15. Enter **n** for interface **e0a** taking over a partner interface.
16. Press **Enter** to accept the blank IP address for **e0b**.
17. Enter **n** for interface **e0b** taking over a partner interface.
18. Enter the IP address of the out-of-band management interface, **e0M**.
19. Enter the net mask for **e0M**.
20. Enter **y** for interface **e0M** taking over a partner IP address during failover.
21. Enter **e0M** for the name of the interface to be taken over.
22. Press **Enter** to accept the flow controller as **full**.
23. Enter **n** to continue setup through the Web interface.
24. Enter the IP address for the default gateway for the storage system.
25. Enter the IP address for the administration host.
26. Enter the local time zone (such as PST, MST, CST, or EST or Linux® time zone format; for example, **America/New_York**).
27. Enter the location for the storage system.
28. Press **Enter** to accept the default root directory for HTTP files [/home/http].
29. Enter **y** to enable DNS resolution.
30. Enter the DNS domain name.
31. Enter the IP address for the first name server.
32. Enter **n** to finish entering DNS servers, or select **y** to add up to two more DNS servers.
33. Enter **n** for running the NIS client.
34. Press **Enter** to acknowledge the AutoSupport™ message.
35. Enter **y** to configure the SP LAN interface.
36. Enter **n** to setting up DHCP on the SP LAN interface.
37. Enter the IP address for the SP LAN interface.
38. Enter the net mask for the SP LAN interface.
39. Enter the IP address for the default gateway for the SP LAN interface.
40. Enter the fully qualified domain name for the mail host to receive SP messages and AutoSupport.
41. Enter the IP address for the mail host to receive SP messages and AutoSupport.

**Note**

If you make a mistake during setup, press **Ctrl+C** to get a command prompt. Enter **setup** and run the setup script again. Or, you can complete the setup script and at the end enter setup to redo the setup script. If you redo the setup script, you must use the passwd command at the end setup to input the administrative password. You will not be automatically prompted to do so. At the end of the setup script, the storage system must be rebooted for changes to take effect.

42. Enter the new administrative (root) password.
43. Enter the new administrative (root) password again to confirm.
44. Log in to the storage system with the new administrative password.

Install Data ONTAP to Onboard Flash Storage

The following steps describe installing Data ONTAP to the onboard flash storage.

Controller A and Controller B

1. To install the Data ONTAP image to the onboard flash device, enter **software install** and indicate the HTTP or HTTPS Web address of the NetApp Data ONTAP 8.1 flash image; for example, **http://192.168.175.5/81_q_image.tgz**.
2. Enter **download** and press **Enter** to download the software to the flash device.

Harden Storage System Logins and Security

The following steps describe hardening the storage system logins and security.

Controller A and Controller B

1. Enter **secureadmin disable ssh**.
2. Enter **secureadmin setup -f ssh** to enable SSH on the storage controller.
3. If prompted, enter **yes** to rerun SSH setup.
4. Accept the default values for SSH1.x protocol.
5. Enter **1024** for SSH2 protocol.
6. If the information specified is correct, enter **yes** to create the SSH keys.
7. Enter **options telnet.enable off** to disable telnet on the storage controller.
8. Enter **secureadmin setup ssl** to enable SSL on the storage controller.
9. If prompted, enter **yes** to rerun SSL setup.
10. Enter the country name code, state or province name, locality name, organization name, and organization unit name.
11. Enter the fully qualified domain name of the storage system.
12. Enter the administrator's e-mail address.
13. Accept the **default for days until the certificate expires**.
14. Enter **1024** for the SSL key length.
15. Enter **options httpd.admin.enable off** to disable HTTP access to the storage system.
16. Enter **options httpd.admin.ssl.enable on** to enable secure access to the storage system.

Install the Required Licenses

The following steps provide details about storage licenses that are used to enable features in this reference architecture. A variety of licenses come installed with the Data ONTAP 8.1 software.



Note

The following licenses are needed to deploy this reference architecture:

- `cluster (cf)`: To configure storage controllers into an HA pair
- `FCP`: To enable the FCP protocol
- `nfs`: To enable the NFS protocol
- `flex_clone`: To enable the provisioning of NetApp FlexClone® volumes and files

Controller A and Controller B

1. Enter **license add <necessary licenses>** to add licenses to the storage system.
2. Enter **license** to double-check the installed licenses.
3. Enter **reboot** to reboot the storage controller.
4. Log back in to the storage controller with the root password.

Enable Licensed Features

The following steps provide details for enabling licensed features.

Controller A and Controller B

1. Enter **options licensed_feature.multistore.enable on**.
2. Enter **options licensed_feature.nearstore_option.enable on**.

Enable Active-Active Controller Configuration Between Two Storage Systems

This step provides details for enabling active-active controller configuration between the two storage systems.

Controller A Only

1. Enter **cf enable** and press **Enter** to enable active-active controller configuration.

Start FCP

This step provides details for enabling the FCP protocol.

Controller A and Controller B

1. Enter **fcv start**.
2. Record the WWPN or FC port name for later use by typing `fcv show adapters`.



Note

If using FC instead of FCoE between storage and the network and there are no available target ports, reconfiguration is necessary.

3. Type **fcadmin config**.

**Note**

Only FC ports that are configured as targets can be used to connect to initiator hosts on the SAN.

4. Type **fcadmin config -t target 0c**.
5. Type **fcadmin config -t target 0d**.

**Note**

If an initiator port is made into a target port, a reboot is required. NetApp recommends rebooting after completing the entire configuration because other configuration steps might also require a reboot.

Set Up Storage System NTP Time Synchronization and CDP Enablement

The following steps provide details for setting up storage system NTP time synchronization and enabling Cisco Discovery Protocol (CDP).

Controller A and Controller B

1. Enter date **CCyymmddhhmm**, where CCyy is the four-digit year, mm is the two-digit month, dd is the two-digit day of the month, hh is the two-digit hour, and the second mm is the two-digit minute to set the storage system time to the actual time.
2. Enter **options timed.proto ntp** to synchronize with an NTP server.
3. Enter **options timed.servers <NTP server IP>** to add the NTP server to the storage system list.
4. Enter **options timed.enable on** to enable NTP synchronization on the storage system.
5. Enter **options cdpd.enable on**.

Create Data Aggregate aggr1

This step provides details for creating the data **aggregate aggr1**.

**Note**

In most cases, the following command finishes quickly, but depending on the state of each disk, it might be necessary to zero some or all of the disks in order to add them to the aggregate. This could take up to 60 minutes to complete.

Controller A

1. Enter **aggr create aggr1 -B 64 <# of disks for aggr1>** to create **aggr1** on the storage controller.

Controller B

1. Enter **aggr create aggr1 -B 64 <# of disks for aggr1>** to create **aggr1** on the storage controller.

Create an SNMP Requests Role and Assign SNMP Login Privileges

This step provides details for creating the SNMP request role and for assigning SNMP login privileges to it.

Controller A and Controller B

1. Run the following command: **useradmin role add <ntap SNMP request role> -a login-snmp**.

Create an SNMP Management Group and Assign an SNMP Request Role

This step provides details for creating an SNMP management group and assigning an SNMP request role to it.

Controller A and Controller B

1. Run the following command: **useradmin group add <ntap SNMP managers> -r <ntap SNMP request role>.**

Create an SNMP User and Assign It to an SNMP Management Group

This step provides details for creating an SNMP user and assigning it to an SNMP management group.

Controller A and Controller B

1. Run the following command: **useradmin user add <ntap SNMP user> -g <ntap SNMP managers>.**



Note

After the user is created, the system prompts for a password. Enter the SNMP password.

Set Up SNMP v1 Communities on Storage Controllers

These steps provide details for setting up SNMP v1 communities on the storage controllers so that OnCommand System Manager can be used.

Controller A and Controller B

1. Run the following command: **snmp community delete all.**
2. Run the following command: **snmp community add ro <ntap SNMP community>.**

Set Up SNMP Contact Information for Each Storage Controller

This step provides details for setting SNMP contact information for each of the storage controllers.

Controller A and Controller B

1. Run the following command: **snmp contact <ntap admin email address>.**

Set SNMP Location Information for Each Storage Controller

This step provides details for setting SNMP location information for each of the storage controllers.

Controller A and Controller B

1. Run the following command: **snmp location <ntap SNMP site name>.**

Reinitialize SNMP on Storage Controllers

This step provides details for reinitializing SNMP on the storage controllers.

Controller A and Controller B

1. Run the following command: **snmp init 1**.

Initialize NDMP on the Storage Controllers

This step provides details for initializing NDMP.

Controller A and Controller B

1. Run the following command: **ndmpd on**.

Enable Flash Cache

This step provides details for enabling the NetApp Flash Cache module, if installed.

Controller A and Controller B

1. Enter the following command to enable Flash Cache on each controller: **options flexscale.enable on**.

Add VLAN Interfaces

The following steps provide details for adding VLAN interfaces on the storage controllers.

Controller A

1. Run the following command: **vlan create vif0 <NFS VLAN ID>**.
2. Run the following command: **wrfile -a /etc/rc vlan create vif0 <NFS VLAN ID>**.
3. Run the following command: **ifconfig vif0-<NFS VLAN ID> <Controller A NFS IP> netmask <NFS Netmask> mtusize 9000 partner vif0-<NFS VLAN ID>**.
4. Run the following command: **wrfile -a /etc/rc ifconfig vif0-<NFS VLAN ID> <Controller A NFS IP> netmask <NFS Netmask> mtusize 9000 partner vif0-<NFS VLAN ID>**.
5. Run the following command to verify additions to the /etc/rc file: **rdfile /etc/rc**.

Controller B

1. Run the following command: **vlan create vif0 <NFS VLAN ID>**.
2. Run the following command: **wrfile -a /etc/rc vlan create vif0 <NFS VLAN ID>**.
3. Run the following command: **ifconfig vif0-<NFS VLAN ID> <Controller B NFS IP> netmask <NFS Netmask>.mtusize 9000 partner vif0-<NFS VLAN ID>**.
4. Run the following command: **wrfile -a /etc/rc ifconfig vif0-<NFS VLAN ID> <Controller B NFS IP> netmask <NFS Netmask>.mtusize 9000 partner vif0-<NFS VLAN ID>**.
5. Run the following command to verify additions to the /etc/rc file: **rdfile /etc/rc**.

Add Infrastructure Volumes

The following steps describe adding volumes on the storage controller for SAN boot of the Cisco UCS hosts as well as virtual machine provisioning.

**Note**

In this reference architecture, controller A houses the boot LUNs for the VMware hypervisor in addition to the swap files, while controller B houses the first datastore for virtual machines.

Controller A

1. Run the following command: **vol create esxi_boot -s none aggr1 100g.**
2. Run the following command: **sis on /vol/esxi_boot.**
3. Run the following command: **vol create infra_swap -s none aggr1 100g.**
4. Run the following command: **snap sched infra_swap 0 0 0.**
5. Run the following command: **snap reserve infra_swap 0.**

Controller B

1. Run the following command: **vol create infra_datastore_1 -s none aggr1 500g.**
2. Run the following command: **sis on /vol/infra_datastore_1.**

Export NFS Infrastructure Volumes to ESXi Servers

These steps provide details for setting up NFS exports of the infrastructure volumes to the VMware ESXi servers.

Controller A

1. Run the following command: **exportfs -p rw=<ESXi Host 1 NFS IP>:<ESXi Host 2 NFS IP>,root=<ESXi Host 1 NFS IP>:<ESXi Host 2 NFS IP>,nosuid /vol/infra_swap.**
2. Run the following command: **exportfs.** Verify that the NFS exports are set up correctly.

Controller B

1. Run the following command: **exportfs -p rw=<ESXi Host 1 NFS IP>:<ESXi Host 2 NFS IP>,root=<ESXi Host 1 NFS IP>:<ESXi Host 2 NFS IP>,nosuid /vol/infra_datastore_1.**
2. Run the following command: **exportfs.** Verify that the NFS exports are set up correctly.

Cisco Unified Computing System Deployment Procedure

The following section provides a detailed procedure for configuring the Cisco Unified Computing System for use in a FlexPod environment. These steps should be followed precisely because a failure to do so could result in an improper configuration.

**Note**

The following sections document the steps necessary to provision Cisco UCS C-Series and B-Series servers as part of a FlexPod environment.

Perform Initial Setup of the Cisco UCS C-Series Blade Servers

These steps provide details for initial setup of the Cisco UCS C-Series M2 Servers with Server BIOS and CIMC software at a 1.4.3c or higher level. This procedure is not necessary for Cisco UCS C-Series M3 servers because they shipped with a minimum of 1.4.3c firmware. This procedure is also not

necessary for Cisco UCS C-Series M2 servers that are already at a 1.4.3c or higher level. It is important to get the systems to a known state with the appropriate firmware package, so that they can be discovered by the Cisco UCS Manager.

**Note**

If the Cisco UCS C-Series blade servers are not part of the architecture to be deployed, this section may be skipped.

All Cisco UCS C-Series M2 Blade Servers

1. From a system connected to the internet, download the latest Cisco UCS Host Upgrade Utility release for your C-Series servers from www.cisco.com. Navigate to **Downloads Home > Products Unified Computing and Servers > Cisco UCS C-Series Rack-Mount Standalone Server Software**.
2. After downloading the **Host Upgrade Utility package**, install the contents to a recordable CD / DVD.
3. Connect a monitor and keyboard to the Front Console port of the server, power on the server and insert the CD media into the C-Series server.
4. Monitor the system as it proceeds through Power On Self Test (POST).
5. Enter **F8** to enter the **CIMC Config**.
6. Upon entering the CIMC Configuration Utility, select the box to **return the CIMC to its factory defaults**.
7. Enter **F10** to save.
8. Enter **F10** to confirm the configuration and the system will reset the CIMC to its factory defaults and automatically reboot.
9. Monitor the system as it proceeds through POST.
10. Enter **F6** to enter the **Boot Selection Menu**.
11. Select the **SATA DVD Drive** when prompted and the server will boot into the upgrade utility.
12. Press **Y** in the C-Series Host Based Upgrade utility screen to acknowledge the Cisco EULA.
13. Select option **8** to upgrade all of the upgradeable items installed.
14. The system will then begin updating the various components, a process that can take 10 -15 minutes.
15. If the system prompts a note that the current version of the LOM is equal to the upgrade version and if the system should continue, enter **Y**.
16. Press any key to acknowledge completion of the updates.
17. Select option **10** to reboot the server.
18. Eject the upgrade media from the server.

Perform Initial Setup of the Cisco UCS 6248 Fabric Interconnects

These steps provide details for initial setup of the Cisco UCS 6248 fabric Interconnects.

Cisco UCS 6248 A

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect.
2. At the prompt to enter the configuration method, enter **console** to continue.
3. If asked to either do a new setup or restore from backup, enter **setup** to continue.

4. Enter **y** to continue to set up a new fabric interconnect.
5. Enter **y** to enforce strong passwords.
6. Enter the password for the admin user.
7. Enter the same password again to confirm the password for the admin user.
8. When asked if this fabric interconnect is part of a cluster, 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.

Cisco UCS 6248 B

1. Connect to the console port on the second Cisco UCS 6248 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 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.

Log into Cisco UCS Manager

Cisco UCS Manager

These steps provide details for logging into the Cisco UCS environment.

1. Open a Web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.
2. Select the **Launch UCS Manager** 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 username and enter the administrative password and click **Login** to log in to the Cisco UCS Manager software.

Upgrade the Cisco UCS Manager Software to Version 2.0(4a)

This document assumes the use of the Cisco UCS Manager 2.0(4a). Refer to [Upgrading Between Cisco UCS 2.0 Releases](#) to upgrade the Cisco UCS Manager software and UCS 6248 Fabric Interconnect software to version 2.0(4a). Do not load the Cisco UCS C-Series version 2.0(4a) software bundle on the Fabric Interconnects.

Add a 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. This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

Cisco UCS Manager

1. Select the **Admin** tab at the top of the left window.
2. Select **All > Communication Management**.
3. Right-click **Management IP Pool**.
4. Select **Create Block of IP Addresses**.
5. Enter the starting IP address of the block and number of IPs needed as well as the subnet and gateway information.
6. Click **OK** to create the IP block.
7. Click **OK** in the message box.

Synchronize Cisco Unified Computing System to NTP

These steps provide details for synchronizing the Cisco Unified Computing System environment to the NTP server.

Cisco UCS Manager

1. Select the **Admin** tab at the top of the left window.
2. Select **All > Timezone Management**.
3. In the right pane, select the appropriate timezone in the Timezone drop-down menu.
4. Click **Save Changes** and then **OK**.
5. Click **Add NTP Server**.
6. Input the NTP server IP and click **OK**.
7. Click **OK**.

Configure Unified Ports

These steps provide details for modifying an unconfigured Ethernet port into an FC uplink port in the Cisco UCS environment.



Note

Modification of the unified ports leads to a reboot of the fabric interconnect in question. This reboot can take up to 10 minutes.

Cisco UCS Manager

1. Navigate to the **Equipment** tab in the left pane.
2. Select **Fabric Interconnect A**.
3. In the right pane, click the **General** tab.
4. Select **Configure Unified Ports**.
5. Select **Yes** to launch the wizard.
6. Use the slider tool and move one position to the left to configure the last two ports (31 and 32) as **FC uplink ports**.
7. Ports 31 and 32 now have the "B" indicator indicating their reconfiguration as FC uplink ports.
8. Click **Finish**.
9. Click **OK**.
10. The Cisco UCS Manager GUI will close as the primary fabric interconnect reboots.
11. Upon successful reboot, open a Web browser, navigate to the **Cisco UCS 6248 fabric interconnect cluster address**, and Launch the **Cisco UCS Manager**.
12. When prompted, enter **admin** for the username and enter the administrative password and click **Login** to log in to the Cisco UCS Manager software.
13. Navigate to the **Equipment** tab in the left pane.
14. Select **Fabric Interconnect B**.
15. In the right pane, click the **General** tab.
16. Select **Configure Unified Ports**.
17. Select **Yes** to launch the wizard.
18. Use the slider tool and move one position to the left to configure the last two ports (31 and 32) as **FC uplink ports**.
19. Ports 31 and 32 now have the "B" indicator indicating their reconfiguration as FC uplink ports.
20. Click **Finish**.
21. Click **OK**.

Edit the Chassis Discovery Policy

These steps provide details for modifying the chassis discovery policy. Setting the discovery policy will simplify the addition of B-Series UCS Chassis and additional fabric extenders for further C-Series connectivity.

Cisco UCS Manager


1. Navigate to the **Equipment** tab in the left pane and select Equipment in the list on the left.
2. In the right pane, click the **Policies** tab.
3. Under **Global Policies**, change the **Chassis Discovery Policy** to **2-link** or to match the number of uplink ports cabled between your chassis or fexes and the Fabric Interconnects.
4. Leave the **Link Grouping Preference** set to **None**.
5. Click **Save Changes** in the bottom right corner.
6. Click **OK**.

Enable Server and Uplink Ports


These steps provide details for enabling server and uplinks ports.

Cisco UCS Manager

1. Select the **Equipment** tab on the top left of the window.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand the **Unconfigured Ethernet Ports** section.
4. Select the ports that are connected to the Chassis or Cisco 2232 FEX (2 per FEX), right-click them, and select **Configure as Server Port**.



5. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.



6. Select ports **19** and **20** that are connected to the Cisco Nexus 5548 switches, right-click them, and select **Configure as Uplink Port**.
7. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
8. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
9. Expand the **Unconfigured Ethernet Ports** section.
10. Select the ports that are connected to the Chassis or Cisco 2232 FEX (2 per FEX), right-click them, and select **Configure as Server Port**.
11. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
12. Select ports **19** and **20** that are connected to the Cisco Nexus 5548 switches, right-click them, and select **Configure as Uplink Port**.
13. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
14. In the case of using the 2208 or 2204 FEX or the external 2232 FEX, navigate to each device by selecting **Equipment -> Chassis ->** or **Equipment -> Rack-Mounts -> FEX -> <FEX #>** and select the **Connectivity Policy** tab in the right-pane and modify the **Admin States** to be **Port Channel**.
15. Click **Save Changes** and then **OK**.


Acknowledge Cisco UCS Chassis and FEX

These steps provide details for acknowledging all UCS Chassis and External 2232 FEX Modules.


Cisco UCS Manager

1. Select the **Equipment** tab on the top left of the window.

2. Expand **Chassis** and select each Chassis listed.
3. For each Chassis, select **Acknowledge Chassis**.



4. Click **Yes** and **OK** to complete acknowledging the Chassis.
5. If you have C-Series servers in your configuration, expand **Rack Mounts** and **FEX**.
6. For each FEX listed, select it and select **Acknowledge Fex**.



7. Click **Yes** and **OK** to complete acknowledging the FEX.

Create Uplink PortChannels to the Cisco Nexus 5548 Switches

These steps provide details for configuring the necessary PortChannels out of the Cisco UCS environment.

Cisco UCS Manager

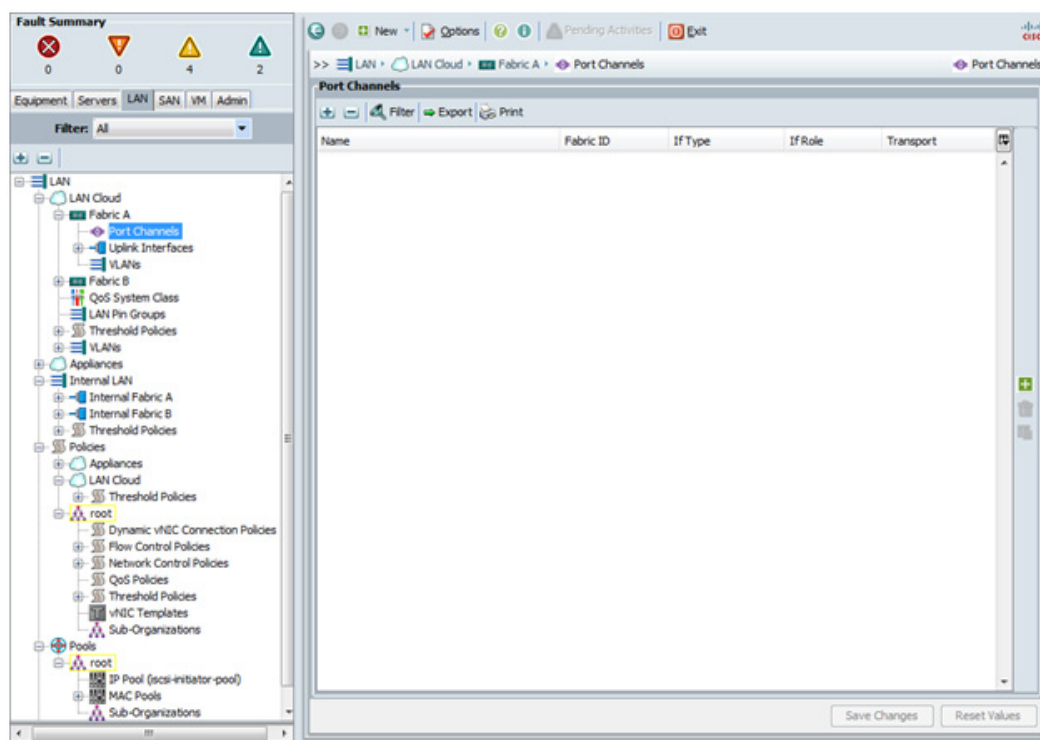
1. Select the **LAN** tab on the left of the window.



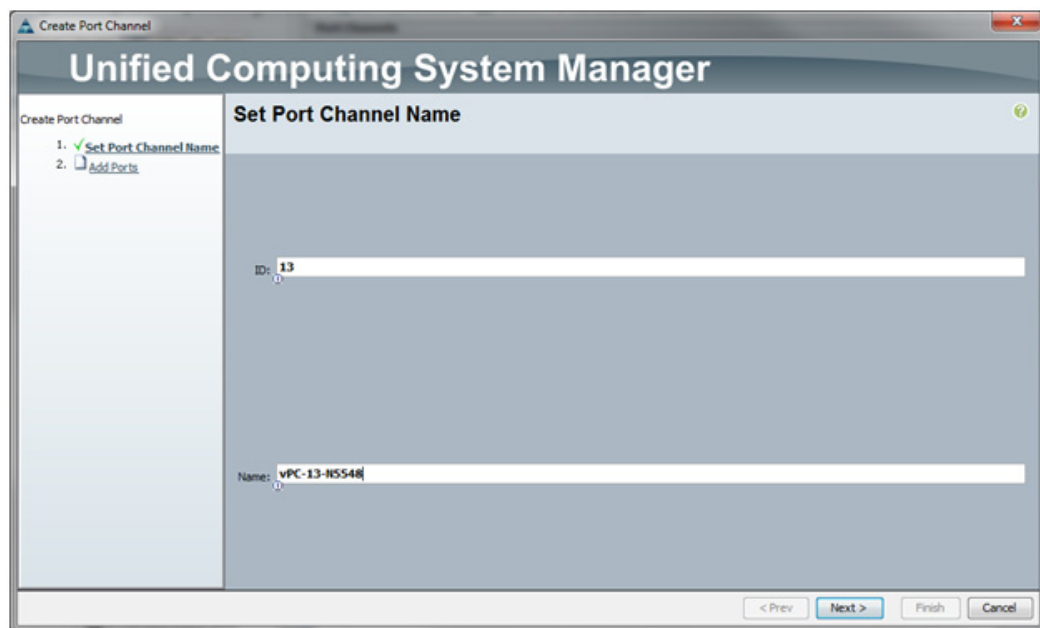
Note

Two PortChannels are created, one from fabric A to both Cisco Nexus 5548 switches and one from fabric B to both Cisco Nexus 5548 switches.

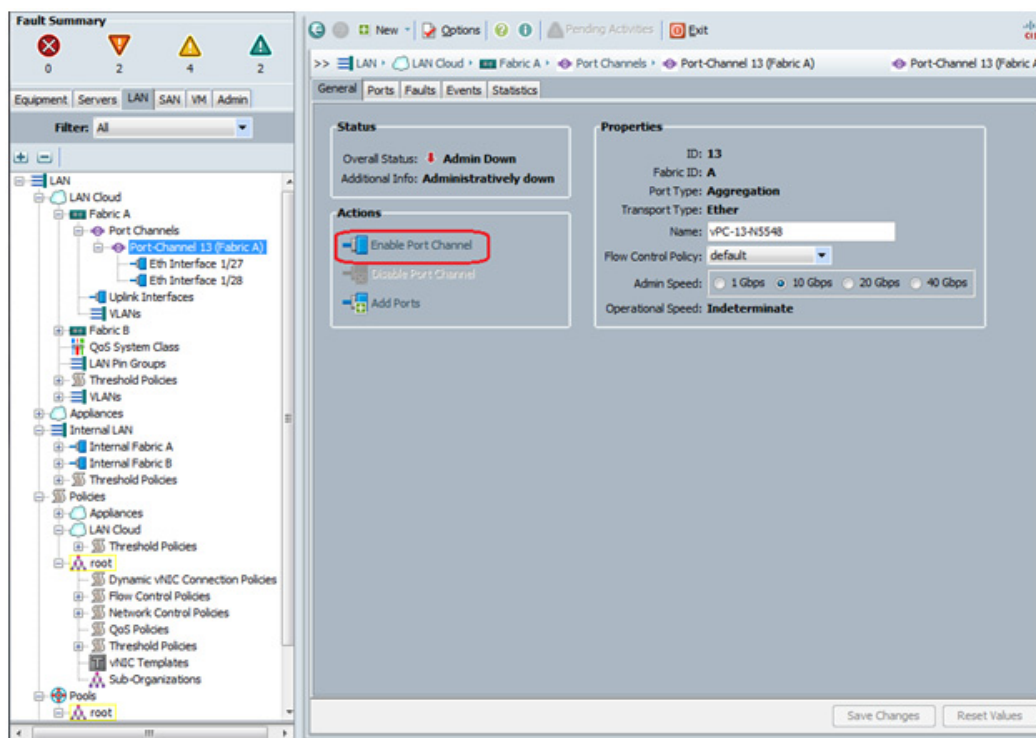
2. Under **LAN Cloud**, expand the **Fabric A** tree.
3. Right-click **Port Channels**.



4. Select **Create Port Channel**.
5. Enter **13** as the unique ID of the PortChannel.
6. Enter **vPC-13-N5548** as the name of the PortChannel.
7. Click **Next**.



8. Select the port with slot ID: **1** and port: **19** and also the port with slot ID: **1** and port **20** to be added to the PortChannel.
9. Click >> to add the ports to the Port Channel.
10. Click **Finish** to create the Port Channel.
11. Select the check box for **Show navigator for Port-Channel 13 (Fabric A)**.
12. Click **OK** to continue.
13. Under **Actions**, select **Enable Port Channel**.
14. In the pop-up box, click **Yes**, then **OK** to enable.



15. Click **OK** to close the Navigator.
16. Under **LAN Cloud**, expand the **Fabric B** tree.
17. Right-click **Port Channels**.
18. Select **Create Port Channel**.
19. Enter **14** as the unique ID of the Port Channel.
20. Enter **vPC-14-N5548** as the name of the Port Channel.
21. Click **Next**.
22. Select the port with slot ID: **1** and port: **19** and also the port with slot ID: **1** and port **20** to be added to the Port Channel.
23. Click >> to add the ports to the Port Channel.
24. Click **Finish** to create the Port Channel.
25. Select Check box for **Show navigator for Port-Channel 14 (Fabric B)**.

26. Click **OK** to continue.
27. Under **Actions**, select **Enable Port Channel**.
28. In the pop-up box, click **Yes**, then **OK** to enable.
29. Click **OK** to close the Navigator.

Create an Organization

These steps provide details for configuring an organization in the Cisco UCS environment. Organizations are used as a means to organize 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 included below.

Cisco UCS Manager

1. From the **New...** menu at the top of the window, select **Create Organization**.
2. Enter a name for the organization.
3. Enter a description for the organization (optional).
4. Click **OK**.
5. In the message box that displays, click **OK**.

Create MAC Address Pools

These steps provide details for configuring the necessary MAC address pools for the Cisco UCS environment.

Cisco UCS Manager


1. Select the **LAN** tab on the left of the window.
2. Select **Pools > root**.




Note

Two MAC address pool are created, one for each switching fabric.

3. Right-click **MAC Pools** under the root organization.
4. Select **Create MAC Pool** to create the MAC address pool.
5. Enter **MAC_Pool_A** for the name of the MAC pool.
6. (Optional) Enter a description of the MAC pool.
7. Click **Next**.
8. Click **Add**.
9. Specify a starting MAC address. It is recommend to place 0A in the next to last octet of the starting MAC address to differentiate the MAC addresses as Fabric A addresses.
10. Specify a size of the MAC address pool sufficient to support the available blade or server resources.



11. Click **OK**.
12. Click **Finish**.
13. In the message box that displays, click **OK**.
14. Right-click **MAC Pools** under the root organization.
15. Select **Create MAC Pool** to create the MAC address pool.
16. Enter **MAC_Pool_B** for the name of the MAC pool.
17. (Optional) Enter a description of the MAC pool.
18. Click **Next**.
19. Click **Add**.
20. Specify a starting MAC address. It is recommend to place 0B in the next to last octet of the starting MAC address to differentiate the MAC addresses as Fabric B addresses.
21. Specify a size of the MAC address pool sufficient to support the available blade or server resources.



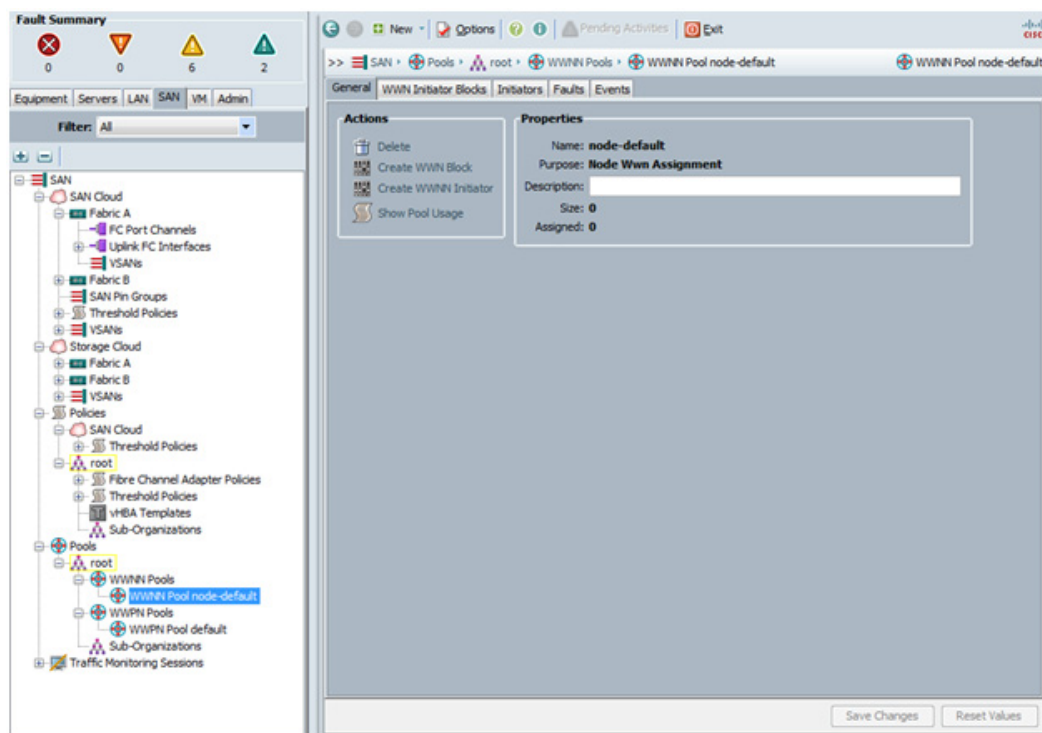
22. Click **OK**.
23. Click **Finish**.
24. In the message box that displays, click **OK**.

Create WWNN Pools

These steps provide details for configuring the necessary WWNN pools for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **SAN** tab at the top left of the window.
2. Select **Pools > root**.
3. Right-click **WWNN Pools**.
4. Select **Create WWNN Pool**.



5. Enter **WWNN_Pool** as the name of the WWNN pool.
6. (Optional) Add a description for the WWNN pool.
7. Click **Next** to continue.
8. Click **Add** to add a block of WWNNs.
9. The default is fine, modify if necessary.
10. Specify a size of the WWNN block sufficient to support the available blade or server resources.



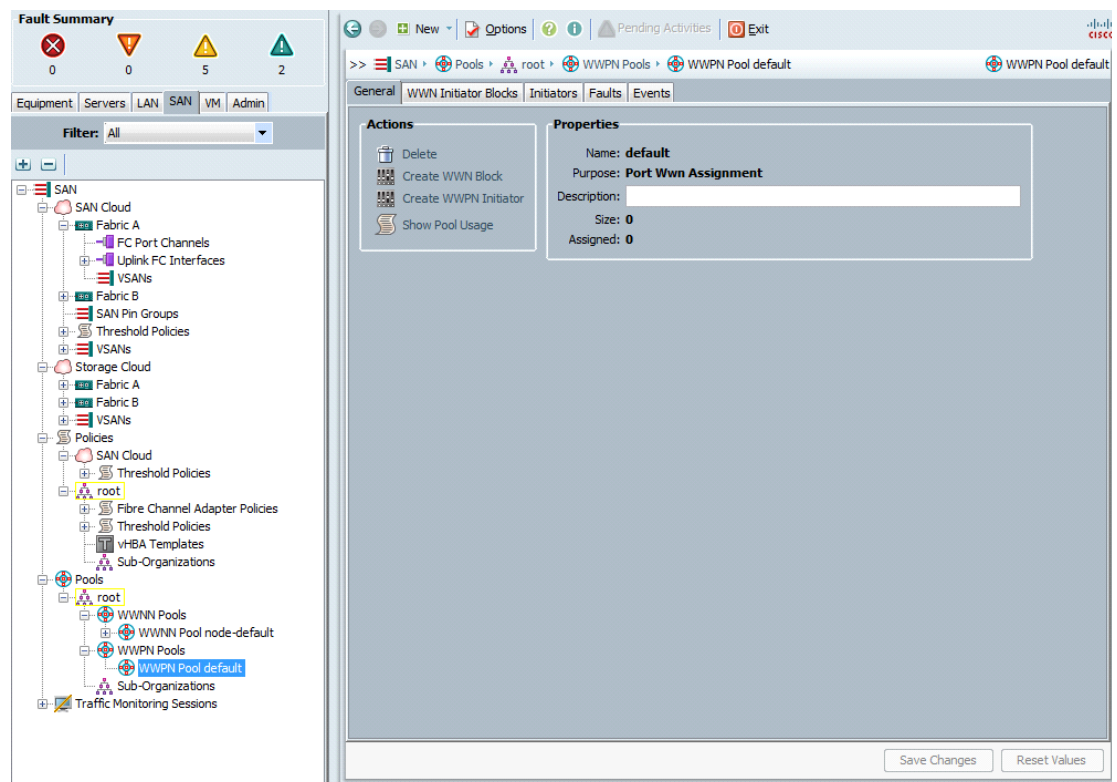
11. Click **OK** to proceed.
12. Click **Finish** to proceed.
13. Click **OK** to finish.

Create WWPN Pools

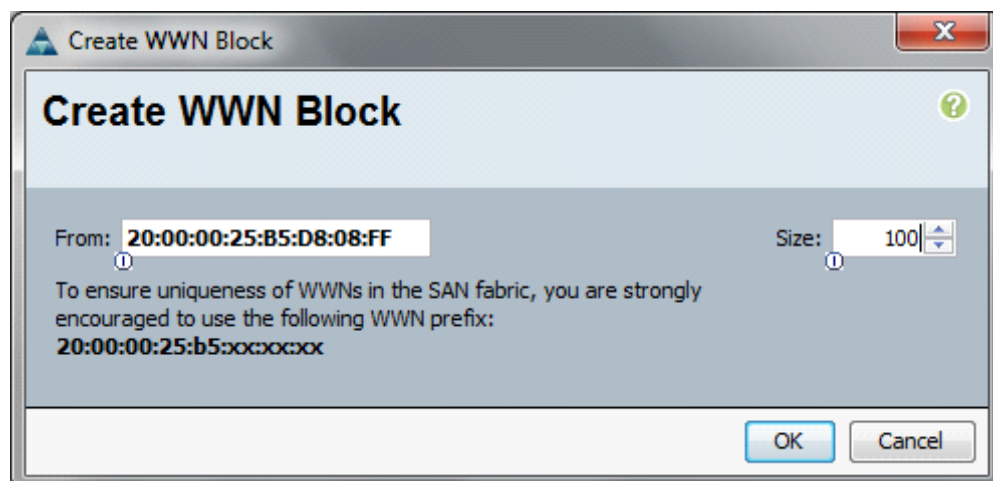
These steps provide details for configuring the necessary WWPN pools for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **SAN** tab at the top left of the window.
2. Select **Pools > root**.
3. Two WWPN pools are created, one for fabric A and one for fabric B.
4. Right-click **WWPN Pools**.
5. Select **Create WWPN Pool**.



6. Enter **WWPN_Pool_A** as the name for the WWPN pool for fabric A.
7. (Optional). Give the WWPN pool a description.
8. Click **Next**.
9. Click **Add** to add a block of WWPNs.
10. Enter the starting WWPN in the block for fabric A.
11. Specify a size of the WWPN block sufficient to support the available blade or server resources.



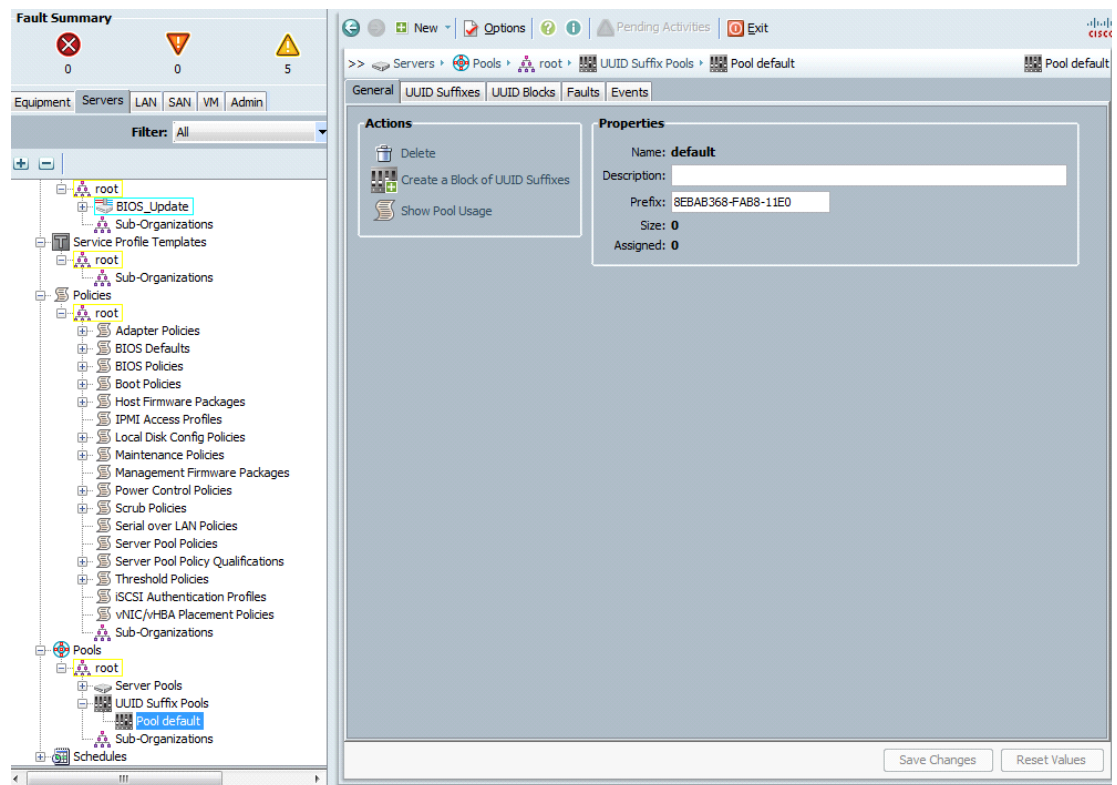
12. Click **OK**.
13. Click **Finish** to create the WWPN pool.
14. Click **OK**.
15. Right-click **WWPN Pools**.
16. Select **Create WWPN Pool**.
17. Enter **WWPN_Pool_B** as the name for the WWPN pool for fabric B.
18. (Optional) Give the WWPN pool a description.
19. Click **Next**.
20. Click **Add** to add a block of WWPNs.
21. Enter the starting WWPN in the block for fabric B.
22. Specify a size of the WWPN block sufficient to support the available blade or server resources.
23. Click **OK**.
24. Click **Finish**.
25. Click **OK** to finish.

Create UUID Suffix Pool

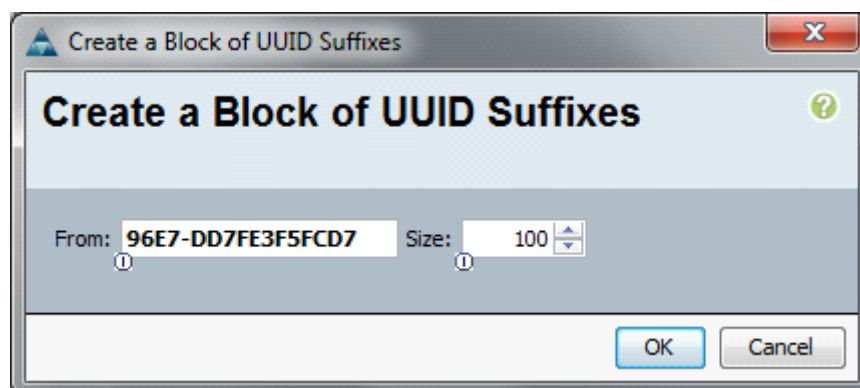
These steps provide details for configuring the necessary UUID suffix pool for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **Servers** tab on the top left of the window.
2. Select **Pools > root**.
3. Right-click **UUID Suffix Pools**.
4. Select **Create UUID Suffix Pool**.



5. Name the UUID suffix pool **UUID_Pool**.
6. (Optional) Give the UUID suffix pool a description.
7. Leave the prefix at the derived option.
8. Click **Next** to continue.
9. The From field is fine at the default setting.
10. Specify a size of the UUID block sufficient to support the available blade or server resources.



11. Click **OK**.
12. Click **Finish** to proceed.

13. Click **OK** to finish.

Create Server Pool

These steps provide details for configuring the necessary server pool for the Cisco UCS environment.



Note

Consider creating unique server pools to the granularity that is required in your environment.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Select **Pools > root**.
3. Right-click **Server Pools**.
4. Select **Create Server Pool**.
5. Name the server pool **Infra_Pool**.
6. (Optional) Give the server pool a description.
7. Click **Next** to continue to add servers.
8. Select two servers to be used for the VMware management cluster and add to the Infra_Pool server pool. Click >> to add them to the pool.
9. Click **Finish**.
10. Select **OK** to finish.

Create VLANs

These steps provide details for configuring the necessary VLANs for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **LAN** tab on the left of the window.



Note

Six VLANs are created.

2. Select **LAN Cloud**.
3. Right-click **VLANs**.
4. Select **Create VLANs**.
5. Enter **MGMT-VLAN** as the name of the VLAN to be used for management traffic.
6. Keep the **Common/Global** option selected for the scope of the VLAN.
7. Enter the VLAN ID for the management VLAN. Keep the sharing type as **none**.
8. Click **OK**, then **OK**.

Create VLANs

VLAN Name/Prefix:

☒ 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

9. Right-click **VLANs**.
10. Select **Create VLANs**.
11. Enter **NFS-VLAN** as the name of the VLAN to be used for the NFS VLAN.
12. Keep the **Common/Global** option selected for the scope of the VLAN.
13. Enter the VLAN ID for the NFS VLAN.
14. Click **OK**, then **OK**.
15. Right-click **VLANs**.
16. Select **Create VLANs**.
17. Enter **Packet-Control-VLAN** as the name of the VLAN to be used for the Nexus 1000v.
18. Keep the **Common/Global** option selected for the scope of the VLAN.
19. Enter the VLAN ID for the Packet Control VLAN.
20. Click **OK**, then **OK**.

Create VLANs

VLAN Name/Prefix:

☒ 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

21. Right-click **VLANs**.
22. Select **Create VLANs**.
23. Enter **vMotion-VLAN** as the name of the VLAN to be used for the vMotion VLAN.
24. Keep the **Common/Global** option selected for the scope of the VLAN.
25. Enter the VLAN ID for the vMotion VLAN.
26. Click **OK**, then **OK**.
27. Right-click **VLANs**.
28. Select **Create VLANs**.
29. Enter **VM-Traffic-VLAN** as the name of the VLAN to be used for the VM Traffic VLAN.
30. Keep the **Common/Global** option selected for the scope of the VLAN.
31. Enter the VLAN ID for the VM Traffic VLAN.
32. Click **OK**, then **OK**.

Create VLANs

VLAN Name/Prefix:

☒ 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

33. Right-click **VLANs**.
34. Select **Create VLANs**.
35. Enter **Native-VLAN** as the name of the VLAN to be used for the Native VLAN.
36. Keep the **Common/Global** option selected for the scope of the VLAN.
37. Enter the VLAN ID for the Native VLAN.
38. Click **OK**, then **OK**.

39. In the expanded list of VLANs in the left pane, right-click the newly created Native-VLAN and select **Set as Native VLAN**.
40. Click **Yes** and **OK**.

Create VSANs and SAN PortChannels

These steps provide details for configuring the necessary VSANs and SAN PortChannels for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **SAN** tab at the top left of the window.
2. Expand the **SAN Cloud** tree.
3. Right-click **VSANs**.
4. Select **Create VSAN**.
5. Enter **VSAN_A** as the VSAN name for fabric A.
6. Keep the **Disabled** option selected for the Default Zoning.
7. Select **Fabric A**.
8. Enter the VSAN ID for fabric A.
9. Enter the FCoE VLAN ID for fabric A.

**Note**

It is recommended to use the same ID for the VSAN and FCoE VLAN ID.

10. Click **OK** and then **OK** to create the VSAN.

Create Storage VSAN

Name:

Default Zoning: ☒ Disabled ☐ Enabled

☐ Common/Global ☒ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID:

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN:

OK Cancel

11. Right-click **VSANs**.
12. Select **Create VSAN**.
13. Enter **VSAN_B** as the VSAN name for fabric B.
14. Keep the **Disabled** option selected for the Default Zoning
15. Select **Fabric B**.
16. Enter the VSAN ID for fabric B.
17. Enter the FCoE VLAN ID for fabric B.

**Note**

It is recommended to use the same ID for the VSAN and FCoE VLAN ID.

18. Click **OK** and then **OK** to create the VSAN.

Create VSAN

Name:

Default Zoning: ☒ Disabled ☐ Enabled

☐ Common/Global ☐ Fabric A ☒ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.
Enter the VSAN ID that maps to this VSAN.

VSAN ID:

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN:

OK Cancel

19. Under **SAN Cloud**, expand the **Fabric A** tree.
20. Right-click **FC Port Channels**.
21. Select **Create Port Channel**.
22. Enter **1** for the Port Channel ID and **SPo1** for the PortChannel name.
23. Click **Next**.
24. Select ports **31** and **32** and click **>>** to add the ports to the Port Channel.
25. Click **Finish**.
26. Select the Check box for **Show navigator for FC Port-Channel 1 (Fabric A)**.
27. Click **OK** to complete creating the Port Channel.
28. In the VSAN pull-down under Properties select the **vsan VSAN_A for fabric A**.
29. Click **Apply**, then click **OK**.
30. Under **Actions**, click **Enable Port Channel**.
31. Click **Yes** and then **OK** to enable the Port Channel. This action also enables the two FC ports in the Port Channel.
32. Click **OK** to Close the Navigator.
33. Under **SAN Cloud**, expand the **Fabric B** tree.


34. Right-click **FC Port Channels**.
35. Select **Create Port Channel**.
36. Enter **2** for the Port Channel ID and **SPo2** for the Port Channel name.
37. Click **Next**.
38. Select ports **31** and **32** and click **>>** to add the ports to the PortChannel.
39. Click **Finish**.
40. Select Check box for **Show navigator for FC Port-Channel 2 (Fabric B)**.
41. Click **OK** to complete creating the Port Channel.
42. In the VSAN pull-down under Properties select **VSAN_B for fabric B**.
43. Click **Apply** and then click **OK**.
44. Under **Actions**, click **Enable Port Channel**.
45. Click **Yes**, then **OK** to enable the Port Channel. This action also enables the two FC ports in the Port Channel.
46. Click **OK** to Close the Navigator.

Create a Firmware Management Package

These steps provide details for a firmware management policy for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Select **Policies > root**.
3. Right-click **Management Firmware Packages**.
4. Select **Create Management Firmware Package**.
5. Enter **VM-Host-Infra** as the management firmware package name.
6. Select the appropriate packages and latest versions of the Server Management Firmware for the servers that you have.
7. Click **OK** to complete creating the firmware management package.
8. Click **OK**.

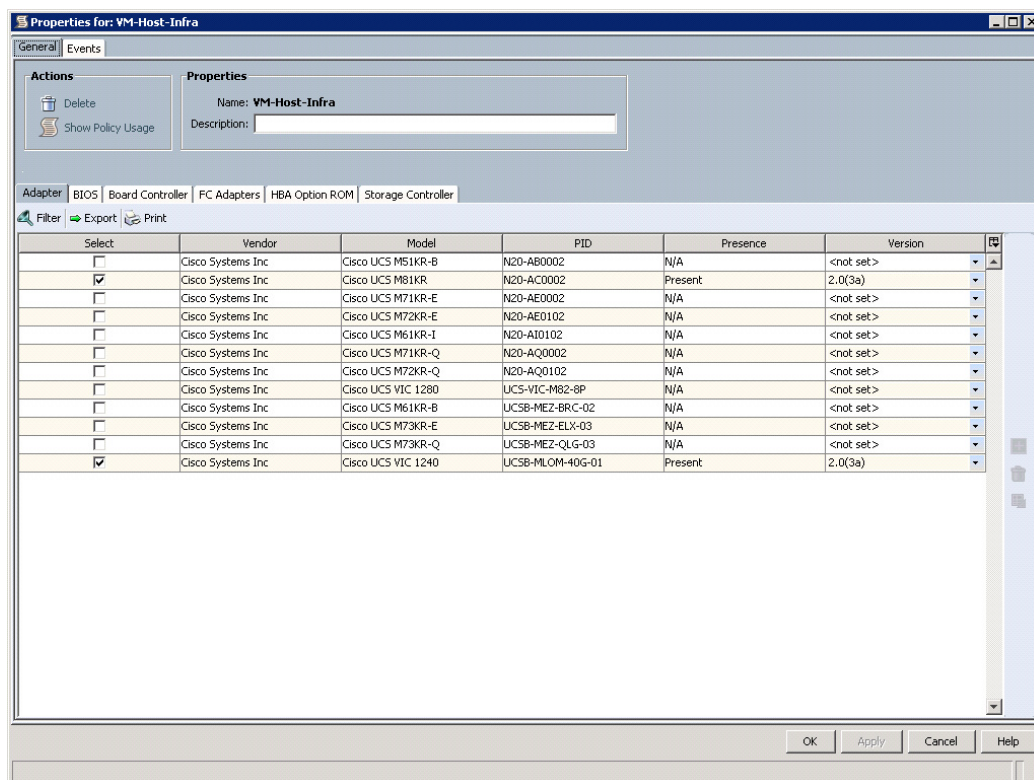


Create Host Firmware Package

These steps provide details for creating a firmware management policy for a given server configuration in the Cisco UCS environment. Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These often include adapter, BIOS, board controller, FC adapters, HBA option ROM, and storage controller properties.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Select **Policies > root**.
3. Right-click **Host Firmware Packages**.
4. Select **Create Host Firmware Package**.
5. Enter **VM-Host-Infra** as the name of the host firmware package.
6. Navigate the tabs of the **Create Host Firmware Package** Navigator and select the appropriate packages and versions for the server configuration.
7. Click **OK** to complete creating the host firmware package.
8. Click **OK**.

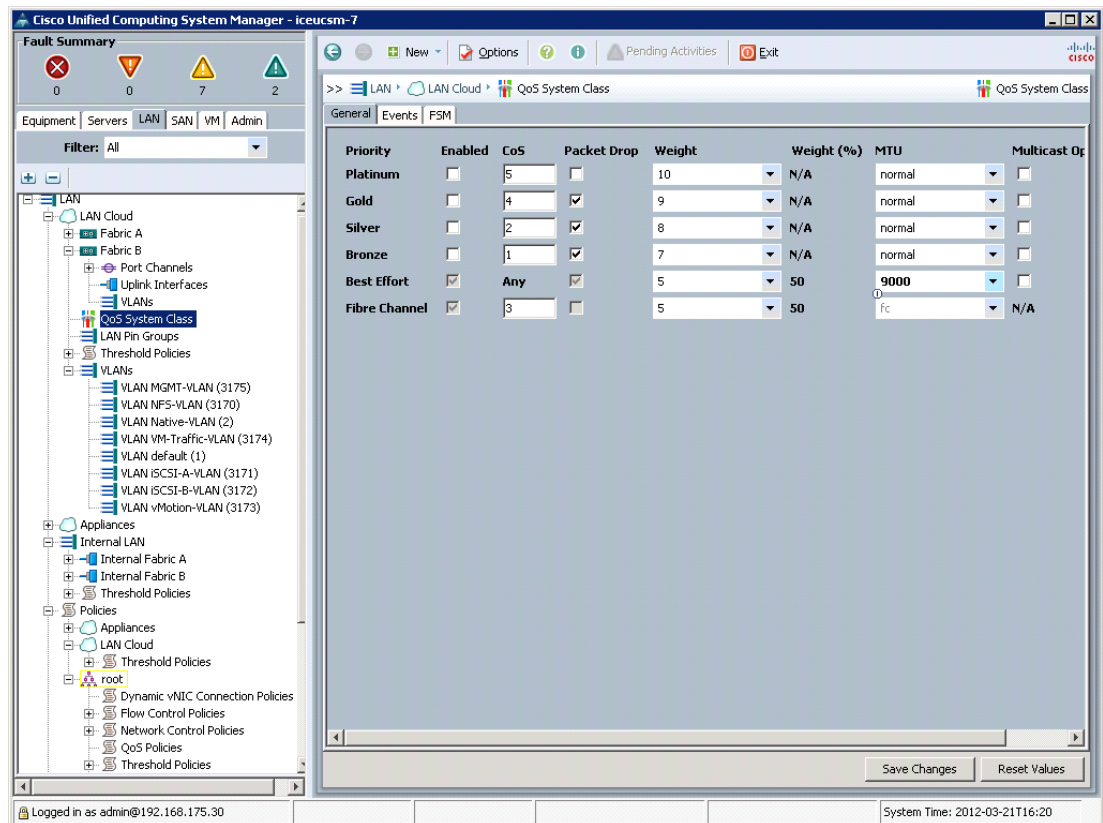


Set Jumbo Frames in Cisco UCS Fabric

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

Cisco UCS Manager

1. Select the **LAN** tab at the top left of the window.
2. Go to **LAN Cloud > QoS System Class**.
3. In the right pane, click the **General** tab
4. On the Best Effort row, type **9000** in the MTU box.
5. Click **Save Changes** in the bottom right corner.
6. Click **OK** to continue.



Create a Local Disk Configuration Policy (Optional)

These steps provide details for creating a local disk configuration for the Cisco UCS environment, which is necessary if the servers in question do not have a local disk.

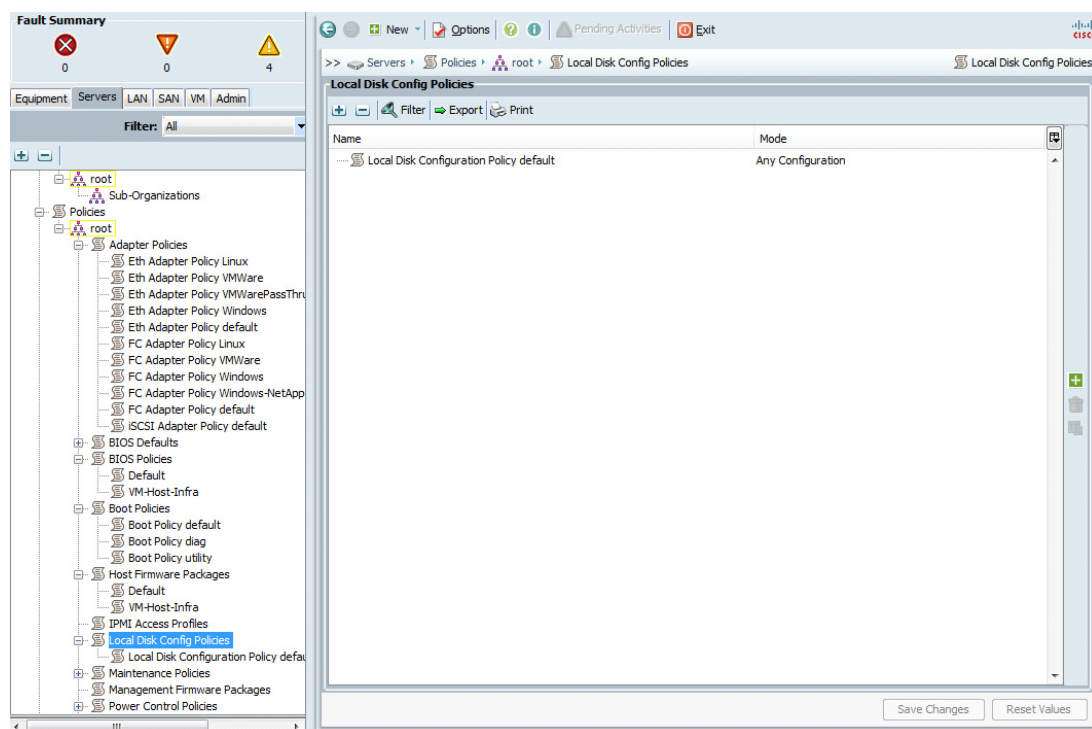


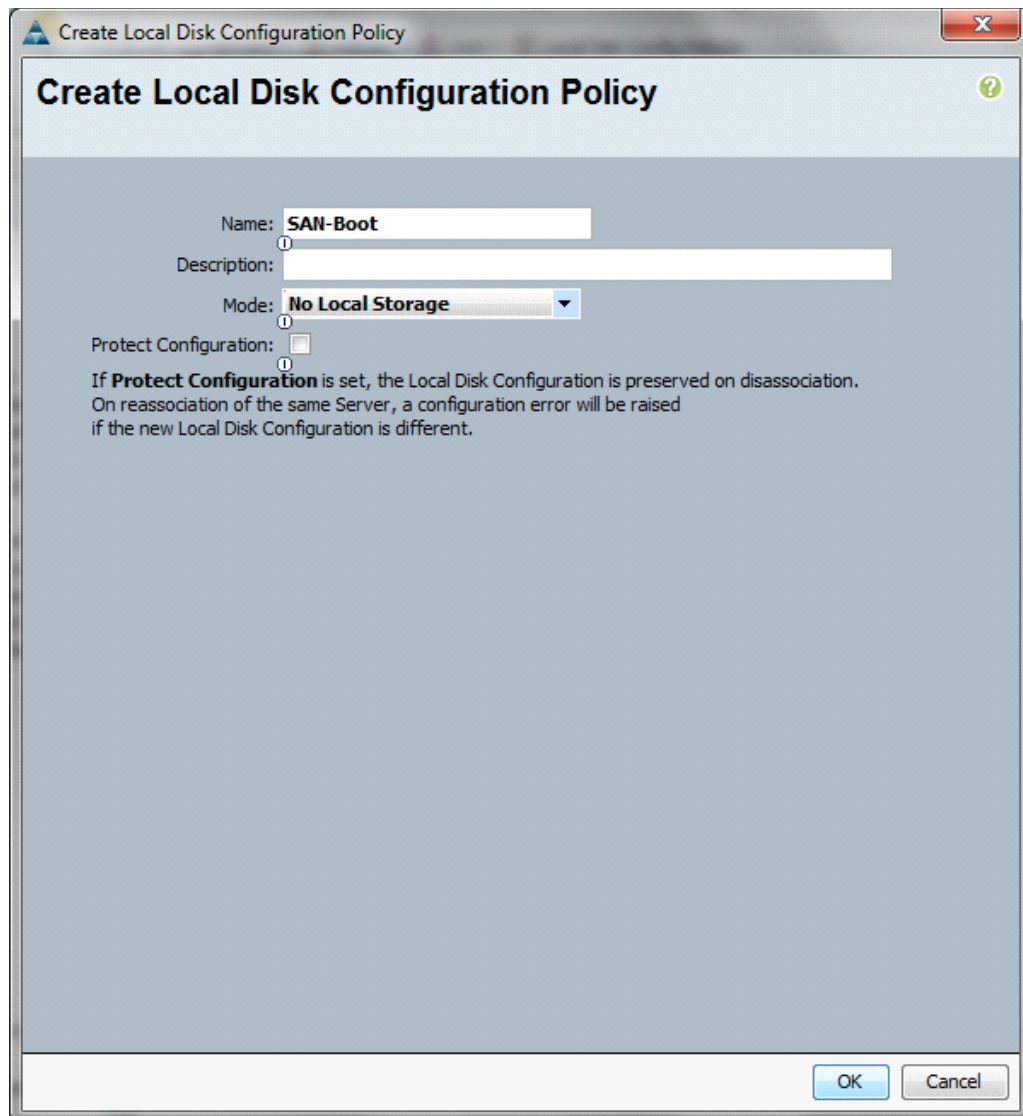
Note

This policy should not be used on servers that contain local disks.

Cisco UCS Manager

1. Select the **Servers** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **Local Disk Config Policies**.
4. Select **Create Local Disk Configuration Policy**.
5. Enter **SAN-Boot** as the local disk configuration policy name.
6. Change the Mode to **No Local Storage**. Uncheck the **Protect Configuration** box.
7. Click **OK** to complete creating the Local Disk Configuration Policy.
8. Click **OK**.





Create a Network Control Policy for Cisco Discovery Protocol (CDP)

These steps provide details for creating a network control policy that enables CDP on virtual network ports.

Cisco UCS Manager

1. Select the **LAN** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **Network Control Policies**.
4. Select **Create Network Control Policy**.
5. Enter **Enable_CDP** as the policy name.
6. Click the **Enabled** radio button for **CDP**.

7. Click **OK** to complete creating the Network Control Policy.
8. Click **OK**.

Create Network Control Policy

Name:

CDP: ☐ Disabled ☒ Enabled

MAC Register Mode: ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail: ☒ Link Down ☐ Warning

MAC Security

Forge: ☒ Allow ☐ Deny

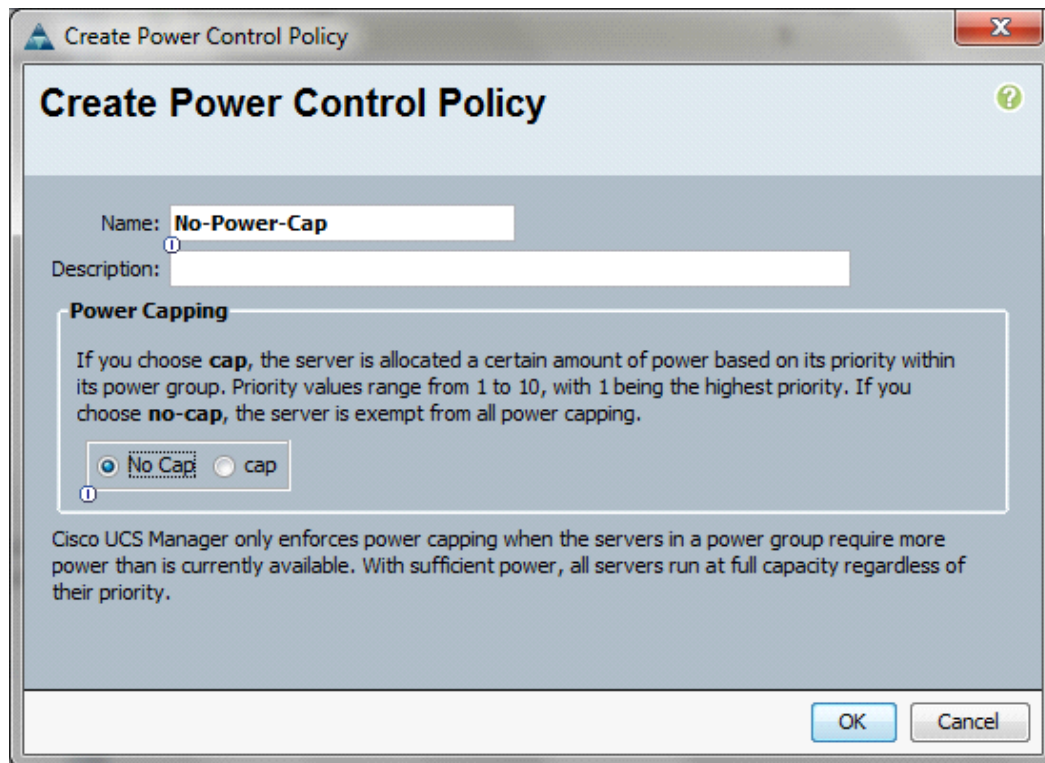
OK Cancel

Create a Power Control Policy

These steps provide details for creating a Power Control Policy for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Go to **Policies > root**.
3. Right-click **Power Control Policies**.
4. Select **Create Power Control Policy**.
5. Enter **No-Power-Cap** as the power control policy name.
6. Change the Power Capping to **No Cap**.
7. Click **OK** to complete creating the Power Control Policy.
8. Click **OK**.



Create a Server Pool Qualification Policy (Optional)

These steps provide details for creating an optional server pool qualification policy for the Cisco UCS environment.

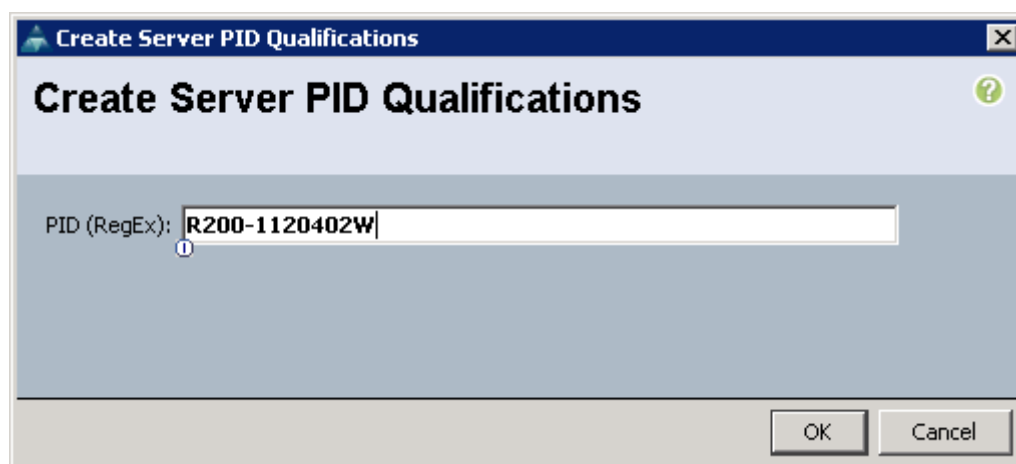


Note

This example details a policy for a C200-M2.

Cisco UCS Manager

1. Select the **Servers** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **Server Pool Policy Qualifications**.
4. Select **Create Server Pool Policy Qualification**.
5. Type **C200-M2** as the name for the Policy.
6. Select **Create Server PID Qualifications**.
7. Enter **R200-1120402W** as the PID.
8. Click **OK** to complete creating the Server Pool Qualification Policy.
9. Click **OK**.
10. Click **OK**.




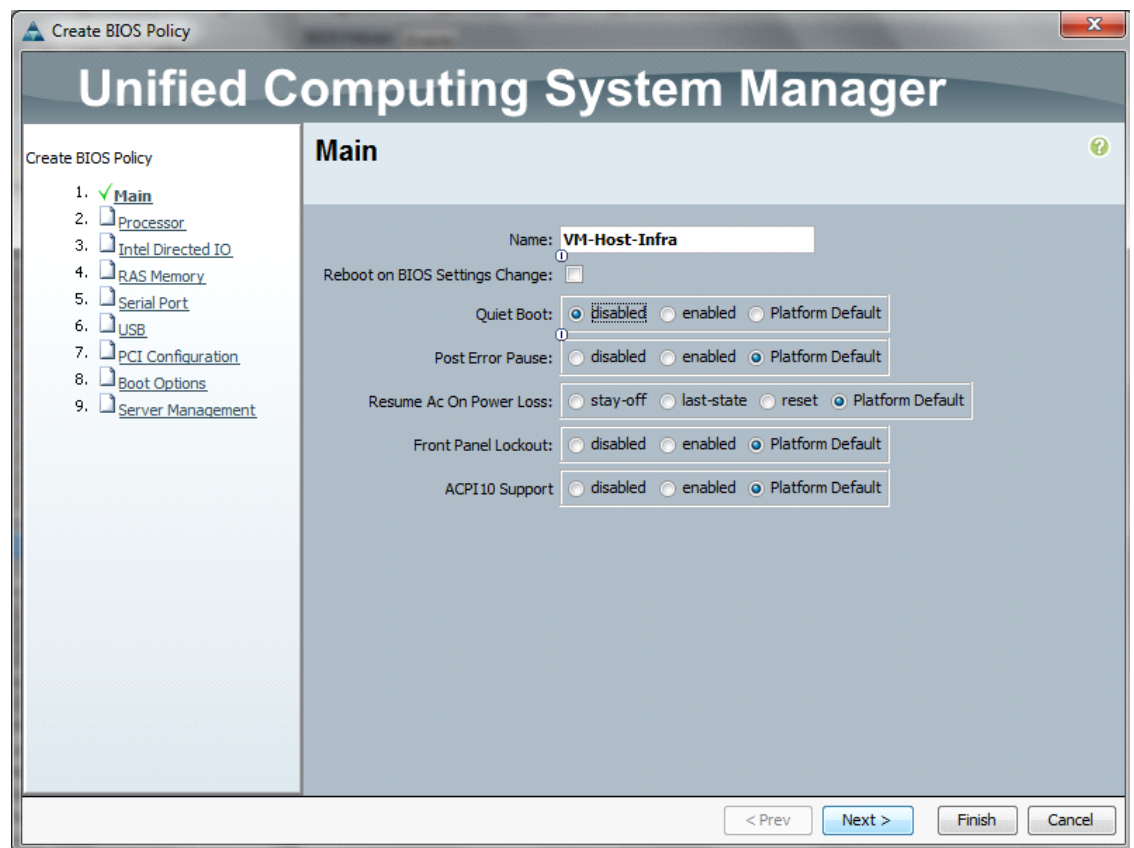
Create a Server BIOS Policy

These steps provide details for creating a server BIOS policy for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **Servers** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **BIOS Policies**.
4. Select **Create BIOS Policy**.
5. Enter **VM-Host-Infra** as the BIOS policy name.
6. Change the Quiet Boot property to **Disabled**.
7. Click **Finish** to complete creating the BIOS policy.
8. Click **OK**.






Create vNIC / vHBA Placement Policy for Virtual Machine Infrastructure Hosts

These steps provide details for creating a vNIC placement policy for infrastructure hosts.

Cisco UCS Manager

1. Select the **Servers** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **vNIC / vHBA Placement policy** and select **Create Placement Policy**.
4. Enter the name **VM-Host-Infra**.
5. Click **1** and select **Assigned Only**.
6. Click **OK**.
7. Click **OK**.




Create vNIC Templates

These steps provide details for creating multiple vNIC templates for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **LAN** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **vNIC Templates**.



4. Select **Create vNIC Template**.
5. Enter **vNIC_Template_A** as the vNIC template name.
6. Leave **Fabric A** selected. Do not check the **Enable Failover** box. Under target, make sure the **VM box** is not selected. Select **Updating Template** as the Template Type. Under VLANs, select **MGMT-VLAN**, **NFS-VLAN**, **Native-VLAN**, **VM-Traffic-VLAN**, **Packet-Control-VLAN**, and **vMotion-VLAN**. Set **Native-VLAN** as the Native VLAN. Under MTU, enter 9000. Under MAC Pool, select **MAC_Pool_A**. Under Network Control Policy: select **Enable_CDP**.
7. Click **OK** to complete creating the vNIC template.
8. Click **OK**.

Create vNIC Template

Name:

Description:

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

Target

☒ Adapter
☐ VM

Warning
If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☐ Initial Template ☒ Updating Template

VLANs

| Select | Name | Native VLAN |
|-------------------------------------|-------------|----------------------------------|
| <input type="checkbox"/> | default | <input type="radio"/> |
| <input checked="" type="checkbox"/> | MGMT-VLAN | <input type="radio"/> |
| <input checked="" type="checkbox"/> | NFS-VLAN | <input type="radio"/> |
| <input checked="" type="checkbox"/> | Native-VLAN | <input checked="" type="radio"/> |

Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

OK Cancel

9. Select the **LAN** tab on the left of the window.
10. Go to **Policies > root**.
11. Right-click **vNIC Templates**.
12. Select **Create vNIC Template**.
13. Enter **vNIC_Template_B** as the vNIC template name.

14. Select **Fabric B**. Do not check the **Enable Failover** box. Under target, make sure the **VM box** is not selected. Select **Updating Template** as the Template Type. Under VLANs, select **MGMT-VLAN**, **NFS-VLAN**, **Native-VLAN**, **VM-Traffic-VLAN**, **Packet-Control-VLAN**, and **vMotion-VLAN**. Set **Native-VLAN** as the Native VLAN. Under MTU, enter **9000**. Under MAC Pool: select **MAC_Pool_B**. Under Network Control Policy: select **Enable_CDP**.
15. Click **OK** to complete creating the vNIC template.
16. Click **OK**.

Create vNIC Template

Name:

Description:

Fabric ID: ☐ Fabric A ☒ Fabric B ☐ Enable Failover

Target

☒ Adapter ☐ VM

Warning

If VM is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☐ Initial Template ☒ Updating Template

VLANs

| Select | Name | Native VLAN |
|-------------------------------------|-------------|----------------------------------|
| <input type="checkbox"/> | default | <input type="radio"/> |
| <input checked="" type="checkbox"/> | MGMT-VLAN | <input type="radio"/> |
| <input checked="" type="checkbox"/> | NFS-VLAN | <input type="radio"/> |
| <input checked="" type="checkbox"/> | Native-VLAN | <input checked="" type="radio"/> |

+ Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:


OK Cancel

Create vHBA Templates for Fabric A and B

These steps provide details for creating multiple vHBA templates for the Cisco UCS environment.

Cisco UCS Manager

1. Select the **SAN** tab on the left of the window.
2. Go to **Policies > root**.
3. Right-click **vHBA Templates**.



4. Select **Create vHBA Template**.
5. Enter **vHBA_Template_A** as the vHBA template name.
6. Select **Fabric A**. Under Select VSAN, select **VSAN_A**. Under WWN Pool, select **WWPN_Pool A**.
7. Click **OK** to complete creating the vHBA template.
8. Click **OK**.

Create vHBA Template

Name:

Description:

Fabric ID: ☒ A ☐ B

Select VSAN: + Create VSAN

Template Type: ☒ Initial Template ☐ Updating Template

Max Data Field Size:

WWN Pool:

QoS Policy:

Pin Group:

Stats Threshold Policy:

9. Select the **SAN** tab on the left of the window.
10. Go to **Policies > root**.
11. Right-click **vHBA Templates**.
12. Select **Create vHBA Template**.
13. Enter **vHBA_Template_B** as the vHBA template name.
14. Select **Fabric B**. Under Select VSAN, select **VSAN_B**. Under WWN Pool, select **WWPN_Pool B**.
15. Click **OK** to complete creating the vHBA template.
16. Click **OK**.

Create Boot Policies

These steps provide details for creating boot policies for the Cisco UCS environment. In the proceeding steps, 2 boot policies will be configured. The first policy will configure the primary target as the first target port of controller A while the second boot policy will configure the primary target as the second target port of controller A. The subsequent screenshots apply to an environment in which each storage controller's 2a port is connected to fabric A and each storage controller's 2b port is connected to fabric B.




Note


If you are using FC between the Nexus 5548 and the NetApp Storage systems substitute port 0c for port 2a and port 0d for port 2b in this procedure.

Cisco UCS Manager

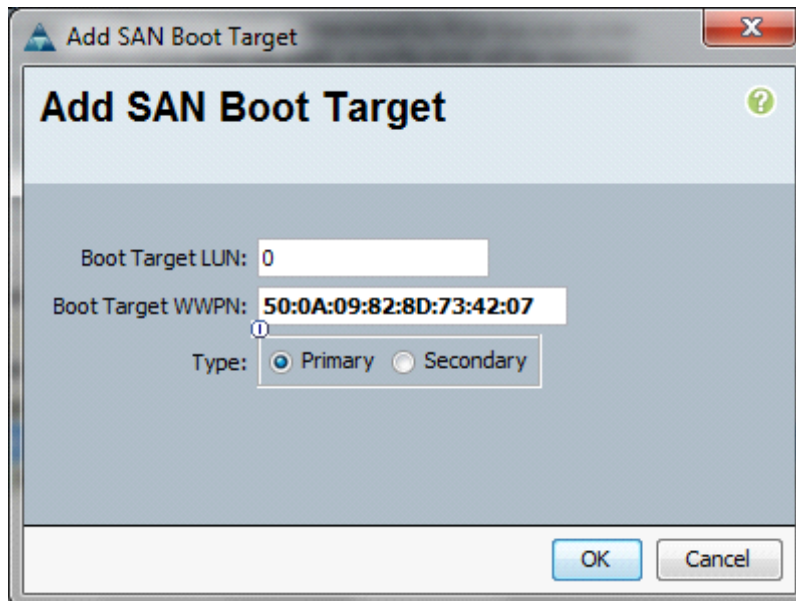
1. Select the **Servers** tab at the top left of the window.
2. Go to **Policies > root**.
3. Right-click **Boot Policies**.



4. Select **Create Boot Policy**.
5. Name the boot policy **Boot-Fabric-A**.
6. (Optional) Give the boot policy a description.
7. Leave **Reboot on Boot Order Change** unchecked.
8. Expand the **Local Devices** drop-down menu and select **Add CD-ROM**.
9. Expand the **vHBAs** drop-down menu and select **Add SAN Boot**.
10. Enter **Fabric-A** in the vHBA field in the Add SAN Boot window that displays.
11. Make sure that **Primary** is selected as the type.
12. Click **OK** to add the SAN boot initiator.

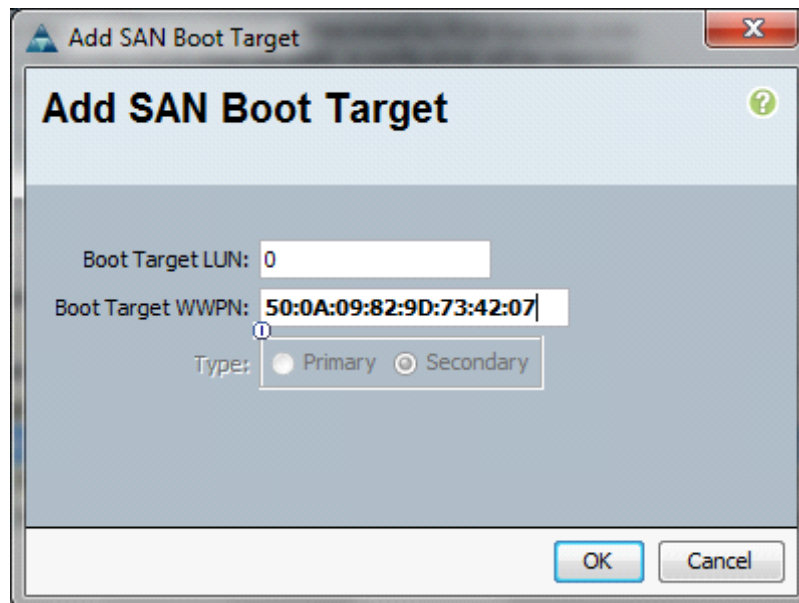


13. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.
14. Enter the WWPN for the primary FC adapter interface **2a**, or **0c** of controller A. To obtain this information, log in to controller A and run the **fcp show adapter** command.
15. Be sure to use the FC portname and not the FC node name.
16. Keep the type as **Primary**.
17. Click **OK** to add the SAN boot target.

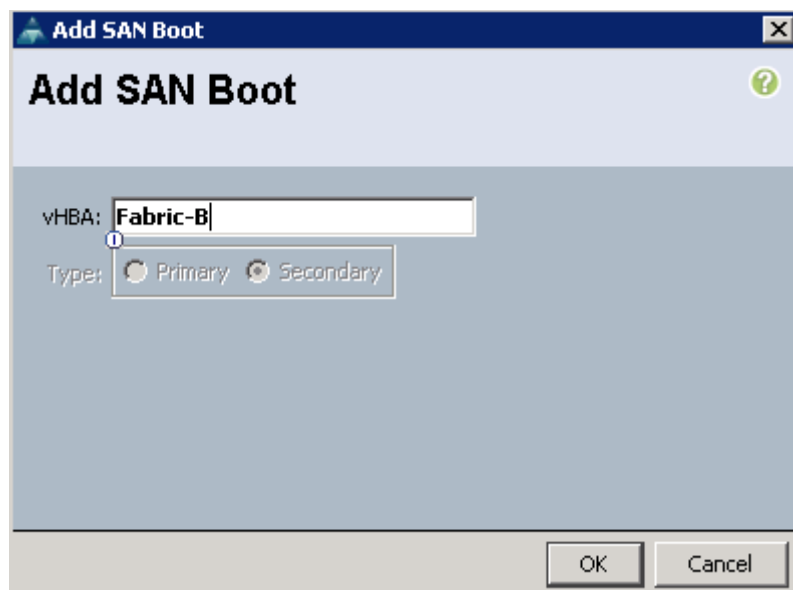


18. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.


19. Enter the WWPN for the primary FC adapter interface **2a** or **0c** of controller B. To obtain this information, log in to the controller B and run the **fcp show adapters** command.
20. Be sure to use the FC portname and not the FC node name.
21. Click **OK** to add the SAN boot target.




22. Select **Add SAN Boot** under the vHBA drop-down menu.
23. Enter **Fabric-B** in the vHBA field in the Add SAN Boot window that displays.
24. The type should automatically be set to **Secondary** and it should be grayed out.
25. Click **OK** to add the SAN boot target.




26. Select **Add SAN Boot Target** under the vHBA drop-down menu.
27. The Add SAN Boot Target window displays. Keep the value for **Boot Target LUN** as **0**.
28. Enter the WWPN for the primary FC adapter interface **2b**, or **0d** of the controller A. To obtain this information, log in to controller A and run the **fcp show adapters** command.
29. Be sure to use the FC portname and not the FC node name.
30. Keep the type as **Primary**.
31. Click **OK** to add the SAN boot target.



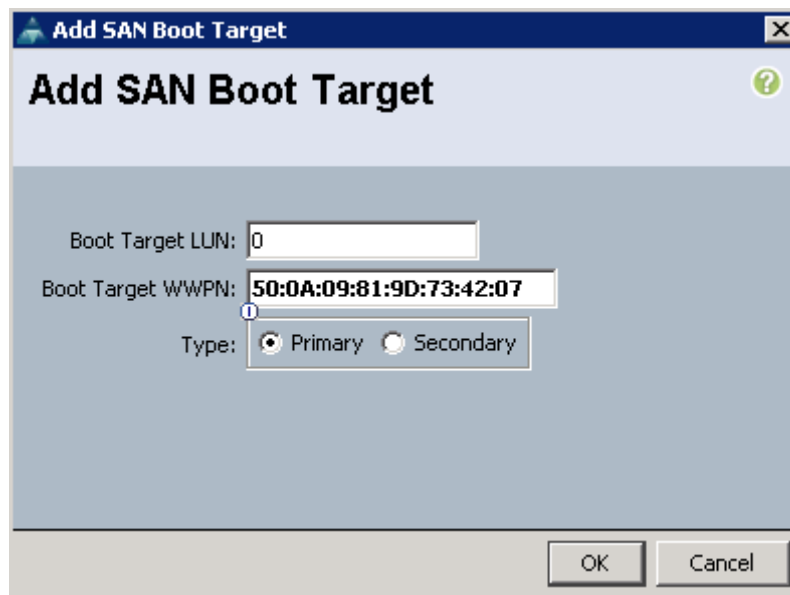
32. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.
33. Enter the WWPN for the primary FC adapter interface **2b**, or adapter **0d** of controller B. To obtain this information, log in to controller B and run the **fcp show adapters** command.
34. Be sure to use the FC portname and not the FC node name.
35. Click **OK** to add the SAN boot target.



36. Click **OK** then **OK** to complete adding the Boot Policy.
37. Right-click **Boot Policies** again.
38. Select **Create Boot Policy**.
39. Name the boot policy **Boot-Fabric-B**.
40. (Optional) Give the boot policy a description.
41. Leave **Reboot on Boot Order Change** unchecked.
42. Expand the **Local Devices** drop-down menu and select **Add CD-ROM**.
43. Click the **vHBA** drop-down menu and select **Add SAN Boot**.
44. Enter **Fabric-B** in the vHBA field in the Add SAN Boot window that displays.
45. Make sure that **Primary** is selected as the type.
46. Click **OK** to add the SAN boot target.



47. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.
48. Enter the WWPN for the primary FC adapter interface **2b**, or adapter **0d** of controller A. To obtain this information, log in to controller B and run the **fcv show adapters** command.
49. Be sure to use the FC portname and not the FC node name.
50. Keep the type as **Primary**.
51. Click **OK** to add the SAN boot target.



52. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.

53. Enter the WWPN for the primary FC adapter interface **2b**, or adapter **0d** of controller B. To obtain this information, log in to controller B and run the **fcv show adapters** command.
54. Be sure to use the FC portname and not the FC node name.
55. Click **OK** to add the SAN boot target.

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 50:0A:09:81:8D:73:42:07

Type: ☐ Primary ☒ Secondary

OK Cancel

56. Select **Add SAN Boot** under the vHBA drop-down menu.
57. Enter **Fabric-A** in the vHBA field in the Add SAN Boot window that displays.
58. The type should automatically be set to **Secondary** and it should be grayed out. This is fine.
59. Click **OK** to add the SAN boot target.

Add SAN Boot

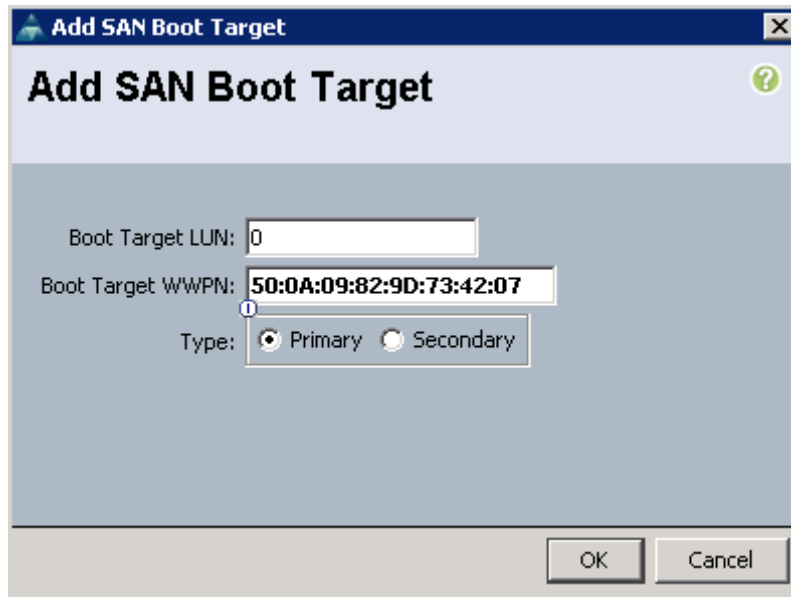
vHBA: Fabric-A

Type: ☐ Primary ☒ Secondary

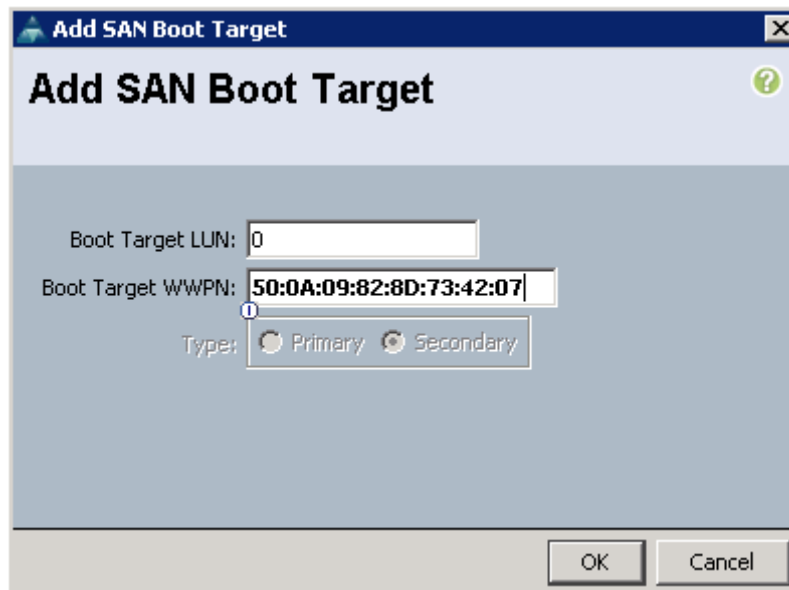
OK Cancel

60. Select **Add SAN Boot Target** under the vHBA drop-down menu.

61. The Add SAN Boot Target window displays. Keep the value for **Boot Target LUN** as **0**.
62. Enter the WWPN for the primary FC adapter interface **2a**, or adapter **0c** of controller A. To obtain this information, log in to controller A and run the **fcp show adapters** command.
63. Be sure to use the FC portname and not the FC node name.
64. Keep the type as **Primary**.
65. Click **OK** to add the SAN boot target.



66. Under the vHBA drop-down menu, select **Add SAN Boot Target**. Keep the value for **Boot Target LUN** as **0**.
67. Enter the WWPN for the primary FC adapter interface **2a**, or adapter **0c** of controller B. To obtain this information, log in to controller B and run the **fcp show adapters** command.
68. Be sure to use the FC portname and not the FC node name.
69. Click **OK** to add the SAN boot target.



70. Click **OK** to create the boot policy in the Create Boot Policy pop-up window.
71. Click **OK**.

Create Service Profile Templates

This section details the creation of two service profile templates: one for fabric A boot and one for fabric B boot. The first profile is created and then cloned and modified for the second host.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Go to **Service Profile Templates > root**.
3. Right-click **root**.
4. Select **Create Service Profile Template**.
5. The Create Service Profile Template window displays.
 - a. These steps detail configuration information for the Identify the Service Profile Template Section.
 - b. Name the service profile template **VM-Host-Infra-Fabric-A**. This service profile template is configured to boot from controller A on Fabric A.
 - c. Select **Updating Template**.
 - d. In the UUID section, select **UUID_Pool** as the UUID pool.
 - e. Click **Next** to continue to the next section.

6. Storage section

- a. Select **Default** for the Local Storage field if the server in question has local disks.
- b. Select the **SAN-Bootlocal storage policy** if the server in question does not have local disk.
- c. Select the **Expert** option for the How would you like to configure SAN connectivity field.
- d. In the WWNN Assignment field, select **WWNN_Pool**.
- e. Click the **Add** button at the bottom of the window to add vHBAs to the template.
- f. The Create vHBA window displays. Name the vHBA **Fabric-A**.
- g. Check the box for **Use SAN Connectivity Template**.
- h. Select **vHBA_Template_A** in the vHBA Template field.
- i. Select **VMware** in the Adapter Policy field.
- j. Click **OK** to add the vHBA to the template.

Create vHBA

Name:

Use SAN Connectivity Template: ☒

[+ Create vHBA Template](#)

vHBA Template:

Adapter Performance Profile

Adapter Policy: [+ Create Fibre Channel Adapter Policy](#)

- k. Click the **Add** button at the bottom of the window to add vHBAs to the template.
- l. The Create vHBA window displays. Name the vHBA **Fabric-B**.
- m. Check the box for **Use SAN Connectivity Template**.
- n. Select **vHBA_Template_B** in the vHBA Template field.
- o. Select **VMware** in the Adapter Policy field.
- p. Click **OK** to add the vHBA to the template.

Create vHBA

Name:

Use SAN Connectivity Template: ☒

[+ Create vHBA Template](#)

vHBA Template:

Adapter Performance Profile

Adapter Policy: [+ Create Fibre Channel Adapter Policy](#)

- Verify - Review the table to make sure that both of the vHBAs were created.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. [Identify Service Profile Template](#)
2. **Storage**
3. [Networking](#)
4. [vNIC/vHBA Placement](#)
5. [Server Boot Order](#)
6. [Maintenance Policy](#)
7. [Server Assignment](#)
8. [Operational Policies](#)

Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: Mode: **Any Configuration**

[+ Create Local Disk Configuration Policy](#)

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

How would you like to configure SAN connectivity? ☐ Simple ☒ Expert ☐ No vHBAs

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

World Wide Node Name

WWNN Assignment:

The WWNN will be assigned from the selected pool.
 The available/total WWNNs are displayed after the pool name.

| Name | WWPN |
|---------------|---------|
| vHBA Fabric-A | Derived |
| vHBA Fabric-B | Derived |


[Delete](#) [+ Add](#) [Modify](#)

< Prev Next > Finish Cancel

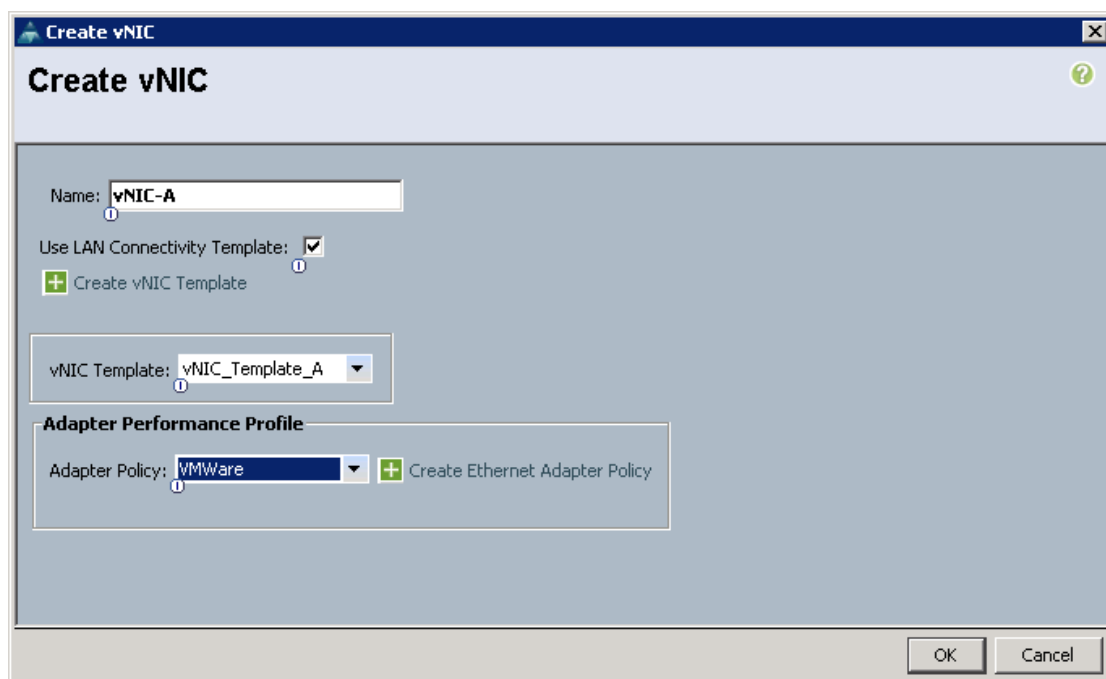
q. Click **Next** to continue to the next section.

7. Networking Section.

- Leave the **Dynamic vNIC Connection Policy** field at the default.
- Select **Expert** for the How would you like to configure LAN connectivity? option.



- c. Click **Add** to add a vNIC to the template.
- d. The Create vNIC window displays. Name the vNIC **vNIC-A**.
- e. Check the **Use LAN Connectivity Template** checkbox.
- f. Select **vNIC_Template_A** for the vNIC Template field.
- g. Select **VMWare** in the Adapter Policy field.
- h. Click **OK** to add the vNIC to the template.



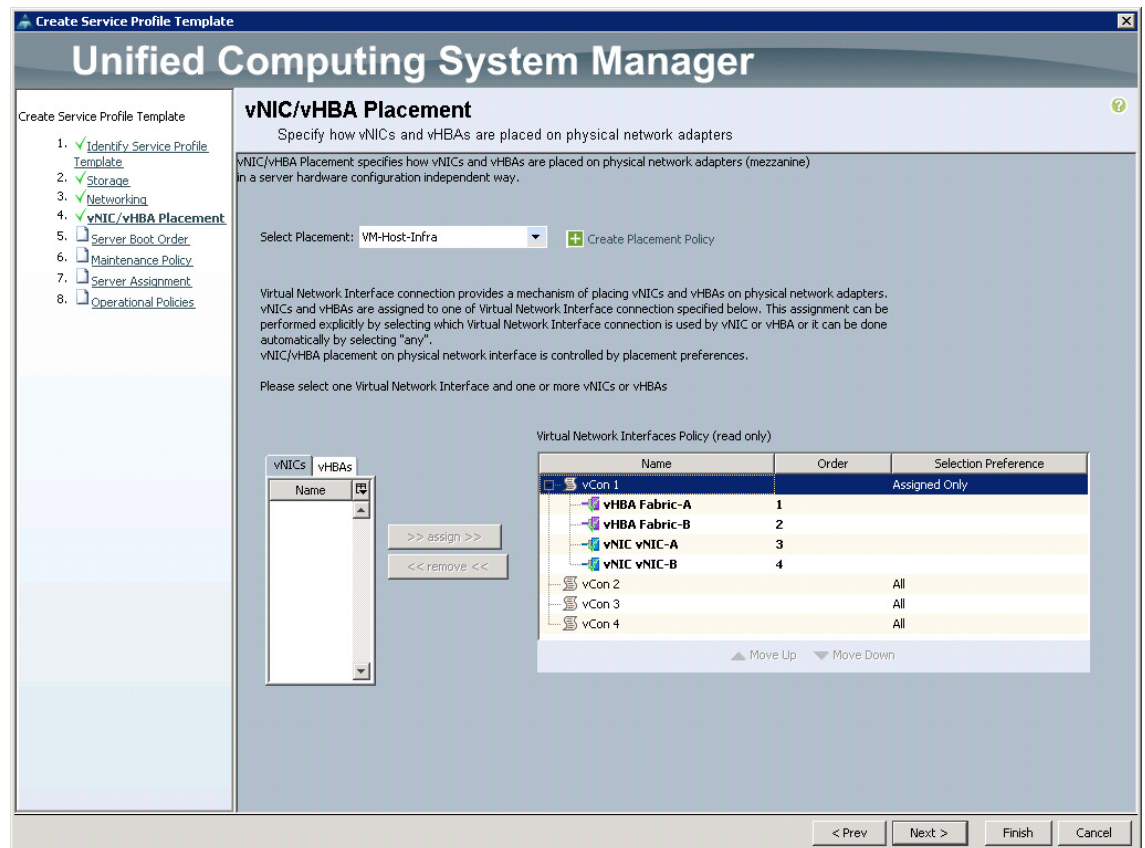
- i. Click **Add** to add a vNIC to the template.
- j. The Create vNIC window displays. Name the vNIC **vNIC-B**.
- k. Check the **Use LAN Connectivity Template** checkbox.
- l. Select **vNIC_Template_B** for the vNIC Template field.
- m. Select **VMWare** in the Adapter Policy field.
- n. Click **OK** to add the vNIC to the template.
- Verify: Review the table to make sure that all of the vNICs were created.

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The 'Networking' section is active, showing options for LAN configuration. The left sidebar lists the steps: 1. Identify Service Profile Template, 2. Storage, 3. Networking (selected), 4. vNIC/vHBA Placement, 5. Server Boot Order, 6. Maintenance Policy, 7. Server Assignment, and 8. Operational Policies. The main area has a title bar 'Networking' and a subtitle 'Optionally specify LAN configuration information.' Below this is a dropdown for 'Dynamic vNIC Connection Policy' set to 'Select a Policy to use (no Dynamic vNIC Policy by default)' and a '+ Create Dynamic vNIC Connection Policy' button. A section titled 'How would you like to configure LAN connectivity?' has three radio buttons: 'Simple' (selected), 'Expert', and 'No vNICs'. Below this is a note: 'Click Add to specify one or more vNICs that the server should use to connect to the LAN.' A table lists vNICs:

| Name | MAC Address | Fabric ID | Native VLAN |
|-------------|-------------|-----------|-------------|
| vNIC vNIC-A | Derived | derived | |
| vNIC vNIC-B | Derived | derived | |

Below the table are buttons for 'Delete', '+ Add', and 'Modify'. At the bottom of the main area is a section for 'iSCSI vNICs'. The bottom of the wizard has navigation buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

- o. Click **Next** to continue to the next section.
8. vNIC/vHBA Placement Section.
 - a. Select the **VM-Host-Infra Placement Policy** in the Select Placement field.
 - b. Select **vCon1** and assign the vHBAs / vNICs in the following order:
 - vHBA Fabric-A
 - vHBA Fabric-B
 - vNIC-A
 - vNIC-B



- Verify: Review the table to make sure that all of the vNICs were assigned in the appropriate order.
 - c. Click **Next** to continue to the next section.
9. Server Boot Order Section
- a. Select **Boot-Fabric-A** in the Boot Policy field.
 - b. Verify: Review the table to make sure that all of the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
 - c. Click **Next** to continue to the next section.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. [Identify Service Profile Template](#)
2. [Storage](#)
3. [Networking](#)
4. [vNIC/vHBA Placement](#)
5. **[Server Boot Order](#)**
6. [Maintenance Policy](#)
7. [Server Assignment](#)
8. [Operational Policies](#)

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Boot-Fabric-A** [+ Create Boot Policy](#)

Name: **Boot-Fabric-A**

Description:

Reboot on Boot Order Change: **no**

Enforce vNIC/vHBA/iSCSI Name: **yes**

WARNINGS:

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

Boot Order

| Name | Order | vNIC/vHBA/iSCSI vNIC | Type | Lun ID | WWN |
|----------------------|-------|----------------------|-----------|--------|-------------------------|
| CD-ROM | 1 | | | | |
| Storage | 2 | | | | |
| SAN primary | | Fabric-A | Primary | | |
| SAN Target primary | | | Primary | 0 | 50:0A:09:81:8D:CD:92:BC |
| SAN Target secondary | | | Secondary | 0 | 50:0A:09:81:9D:CD:92:BC |
| SAN secondary | | Fabric-B | Secondary | | |
| SAN Target primary | | | Primary | 0 | 50:0A:09:82:8D:CD:92:BC |
| SAN Target secondary | | | Secondary | 0 | 50:0A:09:82:9D:CD:92:BC |

[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#)

< Prev Next > Finish Cancel

10. Maintenance Policy Section

- Keep the default of **no policy used by default**.
- Click **Next** to continue to the next section.

11. Server Assignment Section

- Select **Infra_Pool** in the Pool Assignment field.
- Optionally, select a Server Pool Qualification policy.
- Select **Down** for the power state.
- Select **VM-Host-Infra** in the Host Firmware field.
- Select **VM-Host-Infra** in the Management Firmware field.
- Click **Next** to continue to the next section.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Server Assignment

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment:

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

☐ Up ☒ Down

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

Server Pool Qualification:

Restrict Migration: ☐

Firmware Management (BIOS, Disk Controller, Adapter)

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

Host Firmware:

Management Firmware:

< Prev Next > Finish Cancel

12. Operational Policies Section

- Select **VM-Host-Infra** in the BIOS Policy field.
- Select **No-Power-Cap** for the Power Control Policy.
- Click **Finish** to create the Service Profile template.
- Click **OK** in the pop-up window to proceed.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Operational Policies

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy: VM-Host-Infra + Create BIOS Policy

External IPMI Management Configuration

Management IP Address

Monitoring Configuration (Thresholds)

Power Control Policy Configuration

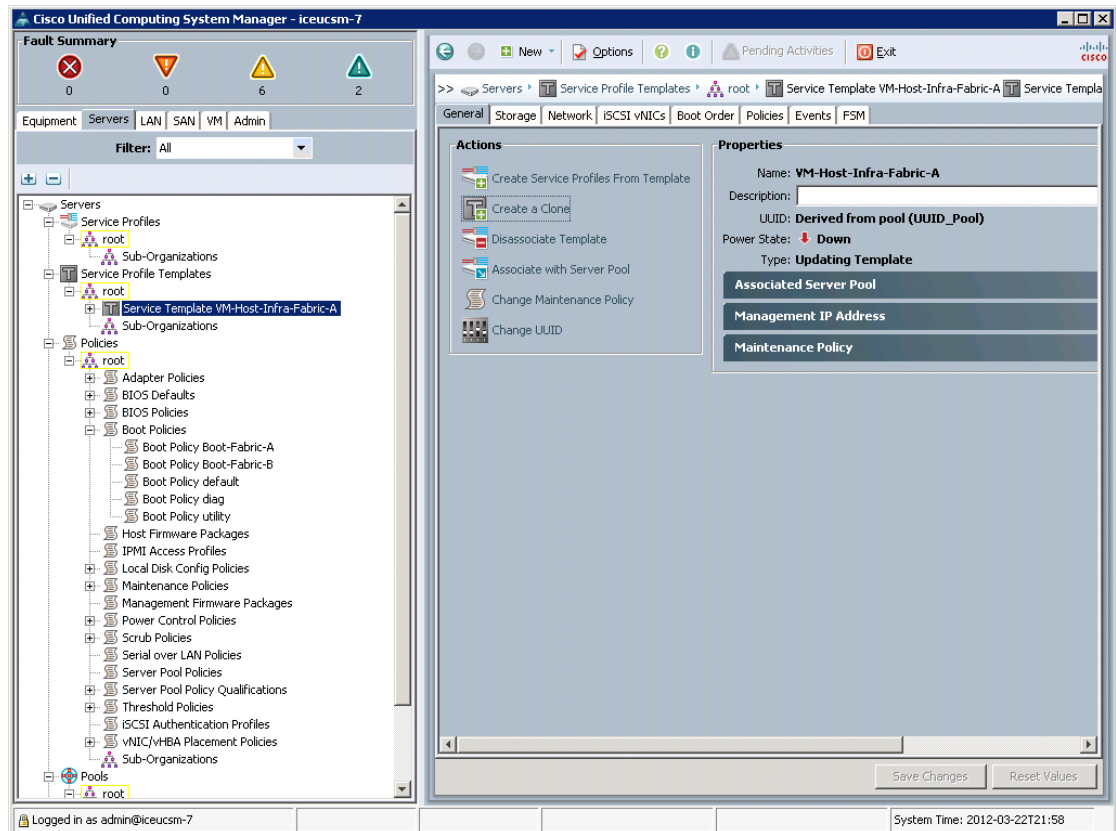
Power control policy determines power allocation for a server in a given power group.

Power Control Policy: No-Power-Cap + Create Power Control Policy

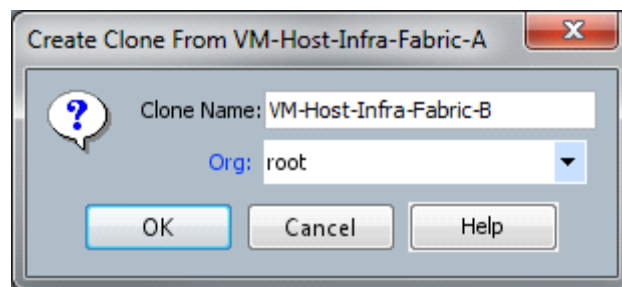
Scrub Policy

< Prev Next > Finish Cancel

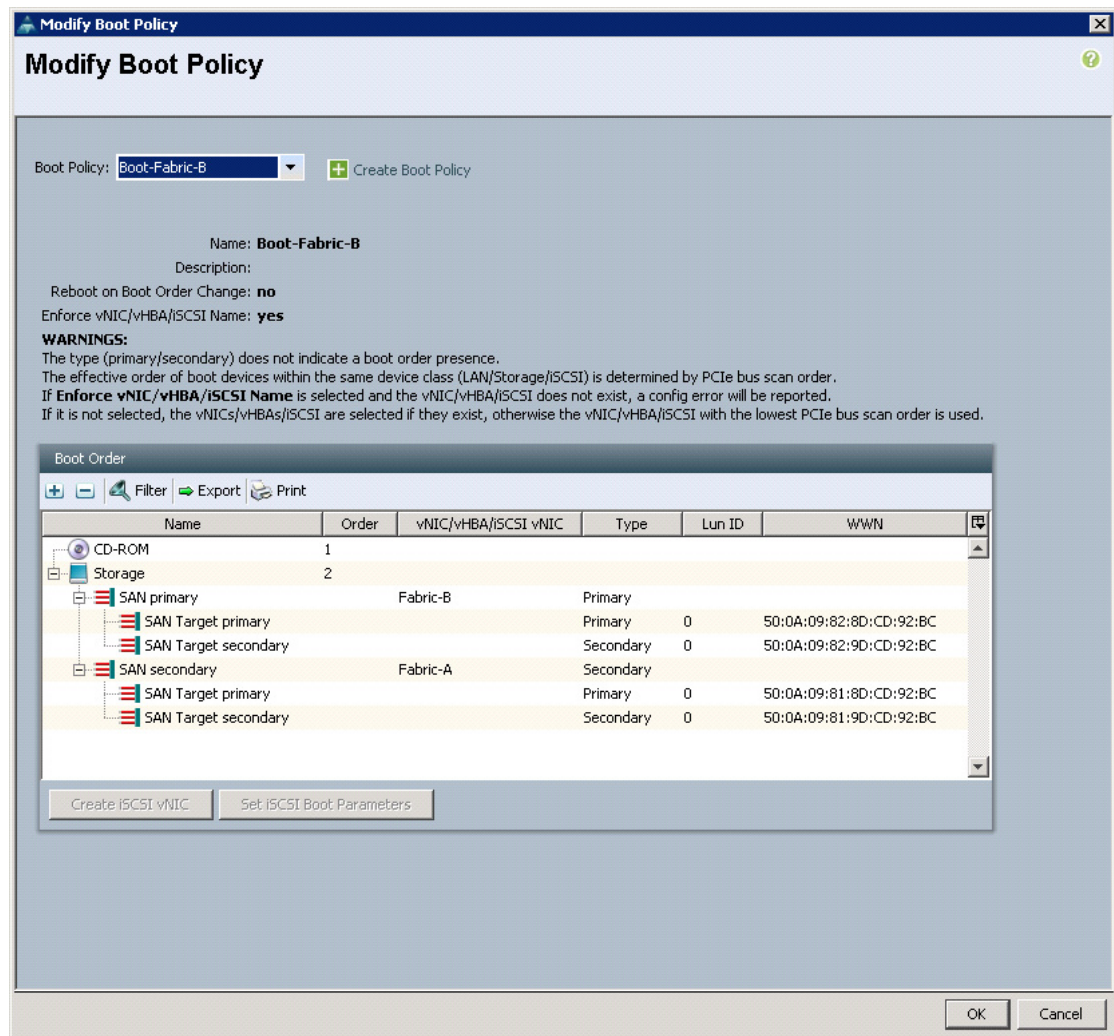
13. Select the **Servers** tab at the top left of the window.
14. Go to **Service Profile Templates > root**.
15. Right-click the previously created **VM-Host-Infra-Fabric-A** template.
16. Click **Create a Clone**.



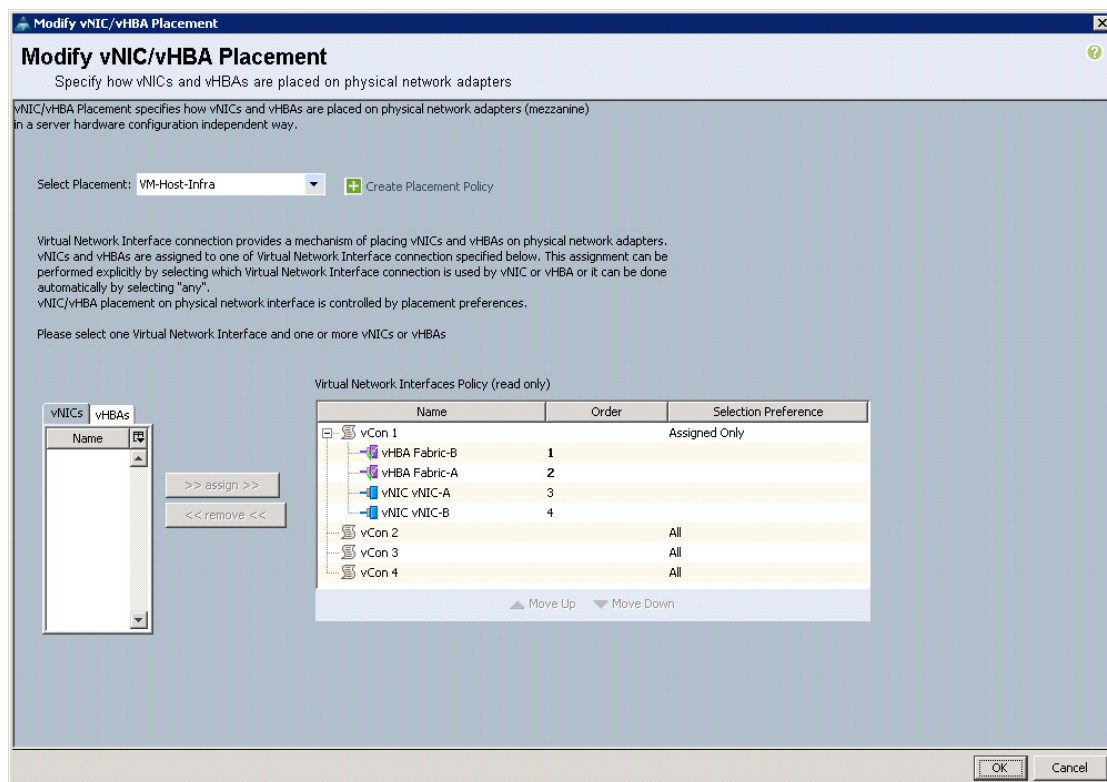
17. Enter **VM-Host-Infra-Fabric-B** in the Clone Name field and click OK.



18. Click **OK**.
19. Select the newly created service profile template and select the **Boot Order** tab.
20. Click **Modify Boot Policy**.



21. Select **Boot-Fabric-B Boot Policy** and click **OK**, then **OK**.
22. Select the **Network** tab and click **Modify vNIC/HBA Placement**.
23. Move **vHBA Fabric-B** ahead of **vHBA Fabric-A** in the placement order and click **OK**, then **OK**.

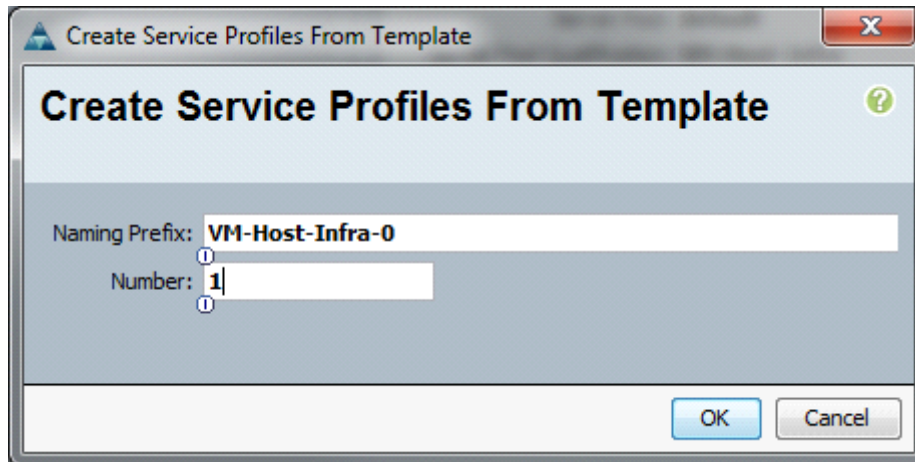


Create Service Profiles

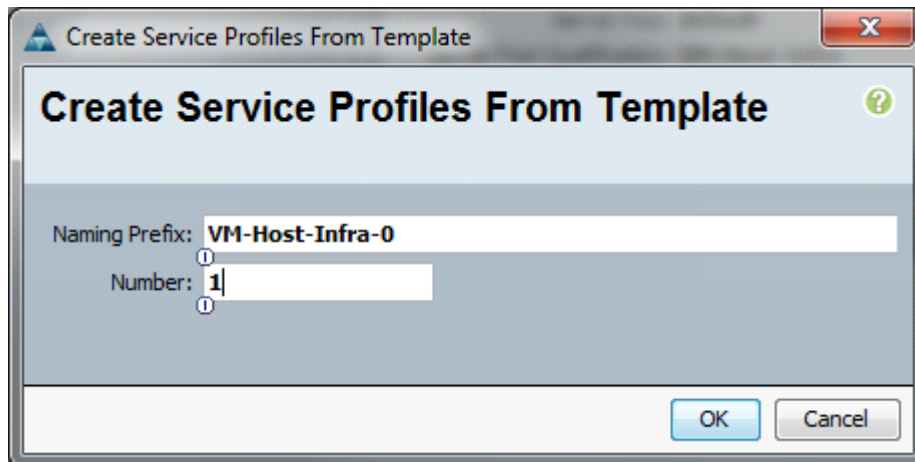
These steps provide details for creating service profiles from template.

Cisco UCS Manager

1. Select the **Servers** tab at the top left of the window.
2. Select **Service Profile Templates > VM-Host-Infra-Fabric-A**.
3. Right-click and select **Create Service Profiles From Template**.
4. Enter **VM-Host-Infra-0** for the service profile prefix.
5. Enter **1** for the number of service profiles to create.
6. Click **OK** to create the service profile.



7. Select **Service Profile Templates > VM-Host-Infra-Fabric-B**.
8. Right-click and select **Create Service Profiles From Template**.
9. Enter **VM-Host-Infra-0** for the service profile prefix.
10. Enter **1** for the number of service profiles to create.
11. Click **OK** to create the service profile.



12. Click **OK** in the message box.
13. Verify that Service Profiles **VM-Host-Infra-01** and **VM-Host-Infra-02** are created. The service profiles will automatically be associated with the servers in their assigned server pools.

Add More Servers to the FlexPod Unit

Add server pools, service profile templates, and service profiles in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created (in the previous steps), the infrastructure blades in the environment each have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information in the tables below.

| NetApp Controller | FC Port | FC Portname |
|-------------------|----------|-------------|
| ControllerA | 2a or 0c | |
| | 2b or 0d | |
| ControllerB | 2a or 0c | |
| | 2b or 0d | |



Note

On each NetApp controller use `fc show adapter` to gather the information above. If using FC between storage and the Cisco Nexus 5548s, substitute 0c for 2a and 0d for 2b.

| Cisco UCS Service Profile Name | Fabric-A WWPN | Fabric-B WWPN |
|--------------------------------|---------------|---------------|
| VM-Host-Infra-01 | | |
| VM-Host-Infra-02 | | |

To gather the information in the table above, launch the Cisco UCS Manager GUI, and in the left pane select the Servers tab. From there, expand Servers > Service Profiles > root > . Click each service profile and then click the Storage tab on the right. While doing so, record the WWPN information in the right display window for both vHBA_A and vHBA_B for each service profile in the table above.

Cisco Nexus 5548 Deployment Procedure

The following section provides a detailed procedure for configuring the Cisco Nexus 5548 switches for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.



Note

The configuration steps detailed in this section provides guidance for configuring the Nexus 5548 UP running release 5.1(3)N2(1b). This configuration also uses the native VLAN on the trunk ports to discard untagged packets, by setting the native VLAN on the PortChannel, but not including this VLAN in the allowed VLANs on the PortChannel.



Note

If you are using Fibre Channel between the storage systems and the Nexus 5548 instead of FCoE, the 0c (assuming FAS32xx) FC ports from the 2 storage systems should be cabled to ports 1/29 and 1/30 on Nexus switch A and the 0d ports should be cabled to ports 1/29 and 1/30 on Nexus switch B.

Set up Initial Cisco Nexus 5548 Switch

These steps provide details for the initial Cisco Nexus 5548 Switch setup.

Cisco Nexus 5548 A

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempts to enter Power on Auto Provisioning.

1. Enter **yes** to abort Power on Auto Provisioning.
2. Enter **yes** to enforce secure password standards.
3. Enter the password for the admin user.
4. Enter the password a second time to commit the password.
5. Enter **yes** to enter the basic configuration dialog.
6. Create another login account (yes/no) [**n**]: **Enter**.
7. Configure read-only SNMP community string (yes/no) [**n**]: **Enter**.
8. Configure read-write SNMP community string (yes/no) [**n**]: **Enter**.
9. Enter the switch name: <**Nexus A Switch name**> **Enter**.
10. Continue with out-of-band (mgmt0) management configuration? (yes/no) [**y**]: **Enter**.
11. Mgmt0 IPv4 address: <**Nexus A mgmt0 IP**> **Enter**.
12. Mgmt0 IPv4 netmask: <**Nexus A mgmt0 netmask**> **Enter**.
13. Configure the default gateway? (yes/no) [**y**]: **Enter**.
14. IPv4 address of the default gateway: <**Nexus A mgmt0 gateway**> **Enter**.
15. Enable the telnet service? (yes/no) [**n**]: **Enter**.
16. Enable the ssh service? (yes/no) [**y**]: **Enter**.
17. Type of ssh key you would like to generate (dsa/rsa):**rsa**.
18. Number of key bits <**768-2048**> :**1024 Enter**.
19. Configure the ntp server? (yes/no) [**y**]: **Enter**.
20. NTP server IPv4 address: <**NTP Server IP**> **Enter**.
21. Enter basic FC configurations (yes/no) [**n**]: **Enter**.
22. Would you like to edit the configuration? (yes/no) [**n**]: **Enter**.
23. Be sure to review the configuration summary before enabling it.
24. Use this configuration and save it? (yes/no) [**y**]: **Enter**.
25. Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus A. It is recommended to continue setup via the console or serial port.
26. Log in as user **admin** with the password previously entered.

Cisco Nexus 5548 B

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and enter Power on Auto Provisioning.

1. Enter **yes** to abort Power on Auto Provisioning.
2. Enter **yes** to enforce secure password standards.
3. Enter the password for the admin user.
4. Enter the password a second time to commit the password.
5. Enter **yes** to enter the basic configuration dialog.

6. Create another login account (yes/no) [n]: **Enter**.
7. Configure read-only SNMP community string (yes/no) [n]: **Enter**.
8. Configure read-write SNMP community string (yes/no) [n]: **Enter**.
9. Enter the switch name: <Nexus B Switch name> **Enter**.
10. Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: **Enter**.
11. Mgmt0 IPv4 address: <Nexus B mgmt0 IP> **Enter**.
12. Mgmt0 IPv4 netmask: <Nexus B mgmt0 netmask> **Enter**.
13. Configure the default gateway? (yes/no) [y]: **Enter**.
14. IPv4 address of the default gateway: <Nexus B mgmt0 gateway> **Enter**.
15. Enable the telnet service? (yes/no) [n]: **Enter**.
16. Enable the ssh service? (yes/no) [y]: **Enter**.
17. Type of ssh key you would like to generate (dsa/rsa):**rsa**.
18. Number of key bits <768-2048>:**1024** **Enter**.
19. Configure the ntp server? (yes/no) [y]: **Enter**.
20. NTP server IPv4 address: <NTP Server IP> **Enter**.
21. Enter basic FC configurations (yes/no) [n]: **Enter**.
22. Would you like to edit the configuration? (yes/no) [n]: **Enter**.
23. Be sure to review the configuration summary before enabling it.
24. Use this configuration and save it? (yes/no) [y]: **Enter**.
25. Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus B. It is recommended to continue setup via the console or serial port.
26. Log in as user **admin** with the password previously entered.

Enable Appropriate Cisco Nexus Features

These steps provide details for enabling the appropriate Cisco Nexus features.

Nexus A and Nexus B

1. Type **config t** to enter the global configuration mode.
2. Type **feature lacp**.
3. Type **feature vpc**.
4. Type **feature fcoe**.
5. Type **feature npiv**.
6. Type **feature fport-channel-trunk**.

Configure FC Ports

These steps provide details for configuring the necessary FCoE ports on the Nexus devices.

Nexus A and Nexus B

1. Type **slot 1**.

2. Type **port 31-32 type fc**.



Note

If you are using FC between the Nexus 5548 and storage, change this to: Type **port 29-32 type fc**.

3. Type **copy run start**.
4. Type **reload**.

Set Global Configurations

These steps provide details for setting global configurations.

Nexus A and Nexus B

1. From the global configuration mode, type **spanning-tree port type network default** to make sure that, by default, the ports are considered as network ports in regards to spanning-tree.
2. Type **spanning-tree port type edge bpduguard default** to enable bpduguard on all edge ports by default.
3. Type **spanning-tree port type edge bpdufilter default** to enable bpdufilter on all edge ports by default.
4. Type **policy-map type network-qos jumbo**.
5. Type **class type network-qos class-default**.
6. Type **mtu 9000**.
7. Type **exit**.
8. Type **class type network-qos class-fcoe**.
9. Type **pause no-drop**.
10. Type **mtu 2158**.
11. Type **exit**.
12. Type **exit**.
13. Type **system qos**.
14. Type **service-policy type queuing input fcoe-default-in-policy**.
15. Type **service-policy type queuing output fcoe-default-out-policy**.
16. Type **service-policy type qos input fcoe-default-in-policy**.
17. Type **service-policy type network-qos jumbo**.
18. Type **exit**.
19. Type **copy run start**.

Create Necessary VLANs

These steps provide details for creating the necessary VLANs.

Nexus A and Nexus B

1. From the global configuration mode, type **vlan <MGMT VLAN ID>**.
2. Type name **MGMT-VLAN**.

3. Type **exit**.
4. Type **vlan <Native VLAN ID>**.
5. Type **name Native-VLAN**.
6. Type **exit**.
7. Type **vlan <NFS VLAN ID>**.
8. Type **name NFS-VLAN**.
9. Type **exit**.
10. Type **vlan <Packet Control VLAN ID>**.
11. Type **name Packet-Control-VLAN**.
12. Type **exit**.
13. Type **vlan <vMotion VLAN ID>**.
14. Type **name vMotion-VLAN**.
15. Type **exit**.
16. Type **vlan <VM-Traffic VLAN ID>**.
17. Type **name VM-Traffic-VLAN**.
18. Type **exit**.

Add Individual Port Descriptions for Troubleshooting

These steps provide details for adding individual port descriptions for troubleshooting activity and verification.

Cisco Nexus 5548 A

1. From the global configuration mode, type **interface Eth1/1**.
2. Type **description <Controller A:e2a>**.
3. Type **exit**.
4. Type **interface Eth1/2**.
5. Type **description <Controller B:e2a>**.
6. Type **exit**.
7. Type **interface Eth1/5**.
8. Type **description <Nexus B:Eth1/5>**.
9. Type **exit**.
10. Type **interface Eth1/6**.
11. Type **description <Nexus B:Eth1/6>**.
12. Type **exit**.
13. Type **interface Eth1/3**.
14. Type **description <UCSM A:Eth1/19>**.
15. Type **exit**.
16. Type **interface Eth1/4**.

17. Type **description** <UCSM B:Eth1/19>.
18. Type **exit**.
19. If using native FC, type **interface fc1/29**.
20. If using native FC, type **switchport description** <Controller A:0c>
21. Type **exit**.
22. If using native FC, type **interface fc1/30**.
23. If using native FC, type **switchport description** <Controller B:0c>
24. Type **exit**.
25. Type **interface fc1/31**.
26. Type **switchport description** <UCSM A:fc1/31>.
27. Type **exit**.
28. Type **interface fc1/32**.
29. Type **switchport description** <UCSM A:fc1/32>.
30. Type **exit**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **interface Eth1/1**.
2. Type **description** <Controller A:e2b>.
3. Type **exit**.
4. Type **interface Eth1/2**.
5. Type **description** <Controller B:e2b>.
6. Type **exit**.
7. Type **interface Eth1/5**.
8. Type **description** <Nexus A:Eth1/5>.
9. Type **exit**.
10. Type **interface Eth1/6**.
11. Type **description** <Nexus A:Eth1/6>.
12. Type **exit**.
13. Type **interface Eth1/3**.
14. Type **description** <UCSM A:Eth1/20>.
15. Type **exit**.
16. Type **interface Eth1/4**.
17. Type **description** <UCSM B:Eth1/20>.
18. Type **exit**.
19. If using native FC to storage, type **interface fc1/29**.
20. If using native FC to storage, type **switchport description** <Controller A:0d>
21. Type **exit**.
22. If using native FC to storage, type **interface fc1/30**.

23. If using native FC to storage, type **switchport description <Controller B:0d>**
24. Type **exit**.
25. Type **interface fc1/31**.
26. Type **switchport description <UCSM B:fc1/31>**.
27. Type **exit**.
28. Type **interface fc1/32**.
29. Type **switchport description <UCSM B:fc1/32>**.
30. Type **exit**.

Create Necessary PortChannels

These steps provide details for creating the necessary PortChannels between devices.

Cisco Nexus 5548 A and Nexus 5548 B

1. From the global configuration mode, type **interface Po10**.
2. Type **description vPC peer-link**.
3. Type **exit**.
4. Type **interface Eth1/5-6**.
5. Type **channel-group 10 mode active**.
6. Type **no shutdown**.
7. Type **exit**.
8. Type **interface Po11**.
9. Type **description <Controller A>**.
10. Type **exit**.
11. Type **interface Eth1/1**.
12. Type **channel-group 11 mode active**.
13. Type **no shutdown**.
14. Type **exit**.
15. Type **interface Po12**.
16. Type **description <Controller B>**.
17. Type **exit**.
18. Type **interface Eth1/2**.
19. Type **channel-group 12 mode active**.
20. Type **no shutdown**.
21. Type **exit**.
22. Type **interface Po13**.
23. Type **description <UCSM A>**.
24. Type **exit**.
25. Type **interface Eth1/3**.

26. Type **channel-group 13 mode active**.
27. Type **no shutdown**.
28. Type **exit**.
29. Type **interface Po14**.
30. Type **description <UCSM B>**.
31. Type **exit**.
32. Type **interface Eth1/4**.
33. Type **channel-group 14 mode active**.
34. Type **no shutdown**.
35. Type **exit**.
36. Type **copy run start**.

Add Port Channel Configurations

These steps provide details for adding Port Channel configurations.

Cisco Nexus 5548 A and Nexus 5548 B

1. From the global configuration mode, type **interface Po10**.
2. Type **switchport mode trunk**.
3. Type **switchport trunk native vlan <Native VLAN ID>**.
4. Type **switchport trunk allowed vlan <MGMT VLAN ID, NFS VLAN ID, Packet Control VLAN ID, vMotion VLAN ID, VM-Traffic VLAN ID>**.
5. Type **spanning-tree port type network**.
6. Type **no shutdown**.
7. Type **exit**.
8. Type **interface Po11**.
9. Type **switchport mode trunk**.
10. Type **switchport trunk native vlan <Native VLAN ID>**.
11. Type **switchport trunk allowed vlan <NFS VLAN ID>**.
12. Type **spanning-tree port type edge trunk**.
13. Type **no shutdown**.
14. Type **exit**.
15. Type **interface Po12**.
16. Type **switchport mode trunk**.
17. Type **switchport trunk native vlan <Native VLAN ID>**.
18. Type **switchport trunk allowed vlan <NFS VLAN ID>**.
19. Type **spanning-tree port type edge trunk**.
20. Type **no shutdown**.
21. Type **exit**.

22. Type **interface Po13**.
23. Type **switchport mode trunk**.
24. Type **switchport trunk native vlan <Native VLAN ID>**.
25. Type **switchport trunk allowed vlan <MGMT VLAN ID, NFS VLAN ID, Packet Control VLAN ID, vMotion VLAN ID, VM-Traffic VLAN ID>**.
26. Type **spanning-tree port type edge trunk**.
27. Type **no shutdown**.
28. Type **exit**.
29. Type **interface Po14**.
30. Type **switchport mode trunk**.
31. Type **switchport trunk native vlan <Native VLAN ID>**.
32. Type **switchport trunk allowed vlan <MGMT VLAN ID, NFS VLAN ID, Packet Control VLAN ID, vMotion VLAN ID, VM-Traffic VLAN ID>**.
33. Type **spanning-tree port type edge trunk**.
34. Type **no shutdown**.
35. Type **exit**.
36. Type **copy run start**.

Configure Virtual Port Channels

These steps provide details for configuring virtual Port Channels (vPCs).

Cisco Nexus 5548 A

1. From the global configuration mode, type **vpc domain <Nexus vPC domain ID>**.
2. Type **role priority 10**.
3. Type **peer-keepalive destination <Nexus B mgmt0 IP> source <Nexus A mgmt0 IP>**.
4. Type **exit**.
5. Type **interface Po10**.
6. Type **vpc peer-link**.
7. Type **exit**.
8. Type **interface Po11**.
9. Type **vpc 11**.
10. Type **exit**.
11. Type **interface Po12**.
12. Type **vpc 12**.
13. Type **exit**.
14. Type **interface Po13**.
15. Type **vpc 13**.
16. Type **exit**.
17. Type **interface Po14**.

18. Type **vpc 14**.
19. Type **exit**.
20. Type **copy run start**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **vpc domain <Nexus vPC domain ID>**.
2. Type **role priority 20**.
3. Type **peer-keepalive destination <Nexus A mgmt0 IP> source <Nexus B mgmt0 IP>**.
4. Type **exit**.
5. Type **interface Po10**.
6. Type **vpc peer-link**.
7. Type **exit**.
8. Type **interface Po11**.
9. Type **vpc 11**.
10. Type **exit**.
11. Type **interface Po12**.
12. Type **vpc 12**.
13. Type **exit**.
14. Type **interface Po13**.
15. Type **vpc 13**.
16. Type **exit**.
17. Type **interface Po14**.
18. Type **vpc 14**.
19. Type **exit**.
20. Type **copy run start**.

Configure Ports for the Cisco Nexus 1010 Virtual Appliances

These steps provide details for configuring the ports connected to the Cisco Nexus 1010.

Cisco Nexus 5548 A

1. From the global configuration mode, type **interface Eth 1/7**.
2. Type **description <Nexus 1010 A:Eth1>**.
3. Type **switchport mode trunk**.
4. Type **switchport trunk allowed vlan <MGMT VLAN ID, Packet Control VLAN ID>**.
5. Type **speed 1000**.
6. Type **spanning-tree port type edge trunk**.
7. Type **no shut**.
8. Type **exit**.
9. Type **interface Eth1/8**.

10. Type **description** <Nexus 1010 B:Eth1>.
11. Type **switchport mode trunk**.
12. Type **switchport trunk allowed vlan** <MGMT VLAN ID, Packet Control VLAN ID>.
13. Type **speed 1000**.
14. Type **spanning-tree port type edge trunk**.
15. Type **no shut**.
16. Type **exit**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **interface Eth 1/7**.
2. Type **switchport mode trunk**.
3. Type **switchport trunk allowed vlan** <MGMT VLAN ID, Packet Control VLAN ID>.
4. Type **speed 1000**.
5. Type **spanning-tree port type edge trunk**.
6. Type **no shut**.
7. Type **exit**.
8. Type **interface Eth1/8**.
9. Type **description** <Nexus 1010 B:Eth2>.
10. Type **switchport mode trunk**.
11. Type **switchport trunk allowed vlan** <MGMT VLAN ID, Packet Control VLAN ID>.
12. Type **speed 1000**.
13. Type **spanning-tree port type edge trunk**.
14. Type **no shut**.
15. Type **exit**.

FCoE Boot Option



Note

Perform this procedure ONLY if implementing the FCoE configuration. For FC, proceed to the “FC Boot Option” section below.

Create VSANs, Assign FCoE Ports, Turn on FCoE Ports

These steps provide details for configuring VSANs, assigning FC ports and enabling FC ports.



Note

This procedure sets up FCoE connections between the Nexus 5548s and the NetApp Storage Systems.

Cisco Nexus 5548 A

1. From the global configuration mode, type **vlan** <Fabric A FCoE VLAN ID>.
2. Type **name FCoE_Fabric_A**.

3. Type **fcoe vsan <VSAN A ID>**.
4. Type **exit**.
5. Type **interface po11**.
6. Type **switchport trunk allowed vlan add <Fabric A FCoE VLAN ID>**.
7. Type **exit**.
8. Type **interface vfc11**.
9. Type **switchport description <Controller A:2a>**.
10. Type **bind interface po11**.
11. Type **no shutdown**.
12. Type **exit**.
13. Type **interface po12**.
14. Type **switchport trunk allowed vlan add <Fabric A FCoE VLAN ID>**.
15. Type **exit**.
16. Type **interface vfc12**.
17. Type **switchport description <Controller B:2a>**.
18. Type **bind interface po12**.
19. Type **no shutdown**.
20. Type **exit**.
21. Type **interface san-port-channel 1**.
22. Type **channel mode active**.
23. Type **exit**.
24. Type **vsan database**.
25. Type **vsan <VSAN A ID> name Fabric_A**.
26. Type **vsan <VSAN A ID> interface fc1/31-32**.
27. Type **vsan <VSAN A ID> interface san-port-channel 1**.
28. Type **vsan <VSAN A ID> interface vfc11**.
29. Type **vsan <VSAN A ID> interface vfc12**.
30. Type **exit**.
31. Type **interface fc1/31-32**.
32. Type **channel-group 1 force**.
33. Type **no shutdown**.
34. Type **exit**.
35. Type **show int san-port-channel 1 to confirm connectivity**.
36. Type **exit**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **vlan <Fabric B FCoE VLAN ID>**.
2. Type **name FCoE_Fabric_B**.

3. Type **fcoe vsan <VSAN B ID>**.
4. Type **exit**.
5. Type **interface po11**.
6. Type **switchport trunk allowed vlan add <Fabric B FCoE VLAN ID>**.
7. Type **exit**.
8. Type **interface vfc11**.
9. Type **switchport description <Controller A:2b>**.
10. Type **bind interface po11**.
11. Type **no shutdown**.
12. Type **exit**.
13. Type **interface po12**.
14. Type **switchport trunk allowed vlan add <Fabric B FCoE VLAN ID>**.
15. Type **exit**.
16. Type **interface vfc12**.
17. Type **switchport description <Controller B:2b>**.
18. Type **bind interface po12**.
19. Type **no shutdown**.
20. Type **exit**.
21. Type **interface san-port-channel 2**.
22. Type **channel mode active**.
23. Type **exit**.
24. Type **vsan database**.
25. Type **vsan <VSAN B ID> name Fabric_B**.
26. Type **vsan <VSAN B ID> interface fc1/31-32**.
27. Type **vsan <VSAN B ID> interface san-port-channel 2**.
28. Type **vsan <VSAN B ID> interface vfc11**.
29. Type **vsan <VSAN B ID> interface vfc12**.
30. Type **exit**.
31. Type **interface fc1/31-32**.
32. Type **channel-group 2 force**.
33. Type **no shutdown**.
34. Type **exit**.
35. Type **show int san-port-channel 2 to confirm connectivity**.
36. Type **exit**.

Create Device Aliases and Create Zones for FCoE Devices

These steps provide details for configuring device aliases and zones for the primary boot path. Instructions are given for all target ports, however, the redundant path is enabled following operating system installation.

Cisco Nexus 5548 A

1. From the global configuration mode, type **device-alias database**.
2. Type **device-alias name VM-Host-Infra-01_A pwn <Fabric-A WWPN>**.
3. Type **device-alias name VM-Host-Infra-02_A pwn <Fabric-A WWPN>**.
4. Type **device-alias name controller_A_2a pwn <Controller A 2a WWPN>**.
5. Type **device-alias name controller_B_2a pwn <Controller B 2a WWPN>**.



Note

Get this information from the table in the [Gather Necessary Information](#) section.

6. After all of the necessary device-alias are created, type **exit**.
7. Type **device-alias commit**.
8. Create the zone for each service profile.
 - a. Type **zone name VM-Host-Infra-01_A vsan <VSAN A ID>**.
 - b. Type **member device-alias VM-Host-Infra-01_A**.
 - c. Type **member device-alias controller_A_2a**.
 - d. Type **member device-alias controller_B_2a**.
 - e. Type **exit**.
 - f. Type **zone name VM-Host-Infra-02_A vsan <VSAN A ID>**.
 - g. Type **member device-alias VM-Host-Infra-02_A**.
 - h. Type **member device-alias controller_A_2a**.
 - i. Type **member device-alias controller_B_2a**.
 - j. Type **exit**.
9. After the zoning for each Cisco UCS service profiles has been created, create a zoneset to organize and manage them.
10. Create the zoneset and add the necessary members.
 - a. Type **zoneset name flexpod vsan <VSAN A ID>**.
 - b. Type **member VM-Host-Infra-01_A**.
 - c. Type **member VM-Host-Infra-02_A**.
 - d. Type **exit**.
11. Activate the zoneset.
 - a. Type **zoneset activate name flexpod vsan <VSAN A ID>**.
 - b. Type **exit**.
12. Type **copy run start**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **device-alias database**.
2. Type **device-alias name VM-Host-Infra-01_B pwwn <Fabric-B WWPN>**.
3. Type **device-alias name VM-Host-Infra-02_B pwwn <Fabric-B WWPN>**.
4. Type **device-alias name controller_A_2b pwwn <Controller A 2b WWPN>**.
5. Type **device-alias name controller_B_2b pwwn <Controller B 2b WWPN>**.

**Note**

Get this information from the table in the [Gather Necessary Information](#) section.

6. After all of the necessary device-alias are created, type **exit**.
7. Type **device-alias commit**.
8. Create the zones for each service profile.
 - a. Type **zone name VM-Host-Infra-01_B vsan <VSAN B ID>**.
 - b. Type **member device-alias VM-Host-Infra-01_B**.
 - c. Type **member device-alias controller_A_2b**.
 - d. Type **member device-alias controller_B_2b**.
 - e. Type **exit**.
 - f. Type **zone name VM-Host-Infra-02_B vsan <VSAN B ID>**.
 - g. Type **member device-alias VM-Host-Infra-02_B**.
 - h. Type **member device-alias controller_A_2b**.
 - i. Type **member device-alias controller_B_2b**.
 - j. Type **exit**.
9. After all of the zones for the Cisco UCS service profiles have been created, create a zoneset to organize and manage them.
10. Create the zoneset and add the necessary members.
 - a. Type **zoneset name flexpod vsan <VSAN B ID>**.
 - b. Type **member VM-Host-Infra-01_B**.
 - c. Type **member VM-Host-Infra-02_B**.
 - d. Type **exit**.
11. Activate the zoneset.
 - a. Type **zoneset activate name flexpod vsan <VSAN B ID>**.
 - b. Type **exit**.
12. Type **copy run start**.

FC Boot Option

**Note**

Perform this procedure **ONLY** if implementing the FC connectivity option.

Create VSANs, Assign FC Ports, Turn on FC Ports

These steps provide details for configuring VSANs, assigning FC ports and enabling FC ports.



Note

This procedure sets up FC connections between the Nexus 5548s and the NetApp Storage Systems

Cisco Nexus 5548 A

1. From the global configuration mode, type **interface san-port-channel 1**.
2. Type **channel mode active**.
3. Type **exit**.
4. Type **vsan database**.
5. Type **vsan <VSAN A ID> name Fabric_A**.
6. Type **vsan <VSAN A ID> interface fc1/29-32**.
7. Type **vsan <VSAN A ID> interface san-port-channel 1**.
8. Type **exit**.
9. Type **interface fc1/29-30**.
10. Type **channel-group 1 force**.
11. Type **no shutdown**.
12. Type **exit**.
13. Type **show int san-port-channel 1** to confirm connectivity.

Cisco Nexus 5548 B

1. From the global configuration mode, type **interface san-port-channel 2**.
2. Type **channel mode active**.
3. Type **exit**.
4. Type **vsan database**.
5. Type **vsan <VSAN B ID> name Fabric_B**.
6. Type **vsan <VSAN B ID> interface fc1/29-32**.
7. Type **vsan <VSAN B ID> interface san-port-channel 2**.
8. Type **exit**.
9. Type **interface fc1/29-30**.
10. Type **channel-group 2 force**.
11. Type **no shutdown**.
12. Type **exit**.
13. Type **show int san-port-channel 2** to confirm connectivity.

Create Device Aliases and Create Zones for FC Devices

These steps provide details for configuring device aliases and zones for the primary boot path. Instructions are given for all target ports, however, the redundant path is enabled following operating system installation.

Cisco Nexus 5548 A

1. From the global configuration mode, type **device-alias database**.
2. Type **device-alias name VM-Host-Infra-01_A pwwn <Fabric-A WWPN>**.
3. Type **device-alias name VM-Host-Infra-02_A pwwn <Fabric-A WWPN>**.
4. Type **device-alias name controller_A_0c pwwn <Controller A 0c WWPN>**.
5. Type **device-alias name controller_B_0c pwwn <Controller B 0c WWPN>**.



Note

Get this information from the table in the [Gather Necessary Information](#) section.

6. After all of the necessary device-alias are created, type **exit**.
7. Type **device-alias commit**.
8. Create the zone for each service profile.
 - a. Type **zone name VM-Host-Infra-01_A vsan <VSAN A ID>**.
 - b. Type **member device-alias VM-Host-Infra-01_A**.
 - c. Type **member device-alias controller_A_0c**.
 - d. Type **member device-alias controller_B_0c**.
 - e. Type **exit**.
 - f. Type **zone name VM-Host-Infra-02_A vsan <VSAN A ID>**.
 - g. Type **member device-alias VM-Host-Infra-02_A**.
 - h. Type **member device-alias controller_A_0c**.
 - i. Type **member device-alias controller_B_0c**.
 - j. Type **exit**.
9. After the zoning for each Cisco UCS service profiles has been created, create a zoneset to organize and manage them.
10. Create the zoneset and add the necessary members.
 - a. Type **zoneset name flexpod vsan <VSAN A ID>**.
 - b. Type **member VM-Host-Infra-01_A**.
 - c. Type **member VM-Host-Infra-02_A**.
 - d. Type **exit**.
11. Activate the zoneset.
 - a. Type **zoneset activate name flexpod vsan <VSAN A ID>**.
 - b. Type **exit**.
12. Type **copy run start**.

Cisco Nexus 5548 B

1. From the global configuration mode, type **device-alias database**.
2. Type **device-alias name VM-Host-Infra-01_B pwwn <Fabric-B WWPN>**.
3. Type **device-alias name VM-Host-Infra-02_B pwwn <Fabric-B WWPN>**.
4. Type **device-alias name controller_A_0d pwwn <Controller A 0d WWPN>**.
5. Type **device-alias name controller_B_0d pwwn <Controller B 0d WWPN>**.



Note

Get this information from the table in the [Gather Necessary Information](#) section.

6. After all of the necessary device-alias are created, type **exit**.
7. Type **device-alias commit**.
8. Create the zones for each service profile.
 - a. Type **zone name VM-Host-Infra-01_B vsan <VSAN B ID>**.
 - b. Type **member device-alias VM-Host-Infra-01_B**.
 - c. Type **member device-alias controller_A_0d**.
 - d. Type **member device-alias controller_B_0d**.
 - e. Type **exit**.
 - f. Type **zone name VM-Host-Infra-02_B vsan <VSAN B ID>**.
 - g. Type **member device-alias VM-Host-Infra-02_B**.
 - h. Type **member device-alias controller_A_0d**.
 - i. Type **member device-alias controller_B_0d**.
 - j. Type **exit**.
9. After all of the zones for the Cisco UCS service profiles have been created, create a zoneset to organize and manage them.
10. Create the zoneset and add the necessary members.
 - a. Type **zoneset name flexpod vsan <VSAN B ID>**.
 - b. Type **member VM-Host-Infra-01_B**.
 - c. Type **member VM-Host-Infra-02_B**.
 - d. Type **exit**.
11. Activate the zoneset.
 - a. Type **zoneset activate name flexpod vsan <VSAN B ID>**.
 - b. Type **exit**.
12. Type **copy run start**.

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, it is recommended to use virtual port channels to uplink the Cisco Nexus 5548 switches included in the FlexPod environment into

the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to type copy run start to save the configuration on each switch once configuration is completed.

NetApp FAS3240 Deployment Procedure: Part 2

Add Infrastructure Host Boot LUNs

The following steps describe adding the necessary boot LUNs on the storage controller for SAN boot of the Cisco UCS hosts.

Controller A

1. Run the following command: **lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-01.**
2. Run the following command: **lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-02.**

Create FCP Igroups

These steps describe creating the Fibre Channel igroups on the storage controller for SAN boot of the Cisco UCS hosts.

Controller A

1. Run the following command: **igroup create -f -t vmware VM-Host-Infra-01 <VM-Host-Infra-01 vHBA_A WWPN> <VM-Host-Infra-01 vHBA_B WWPN>.**
2. Run the following command: **igroup create -f -t vmware VM-Host-Infra-02 <VM-Host-Infra-02 vHBA_A WWPN> <VM-Host-Infra-02 vHBA_B WWPN>.**
3. Run the following command: **igroup set VM-Host-Infra-01 alua yes.**
4. Run the following command: **igroup set VM-Host-Infra-02 alua yes.**

Map LUNs to Igroups

These steps describe mapping the boot LUNs to the Fibre Channel igroups on the storage controller for SAN boot of the Cisco UCS hosts.

Controller A

1. Run the following command: **lun map /vol/esxi_boot/VM-Host-Infra-01 VM-Host-Infra-01 0.**
2. Run the following command: **lun map /vol/esxi_boot/VM-Host-Infra-02 VM-Host-Infra-02 0.**
3. Run the following command: **lun show -m.**
4. Verify that the created LUNs are mapped correctly.

VMware ESXi 5.0 Deployment Procedure

The following subsections (through “Move the VM Swap File Location”) provide detailed procedures for installing VMware ESXi 5.0 in this environment. The deployment procedures that follow are customized to include the environment variables described in previous sections. By the end of this section, two FCP booted ESX hosts will be provisioned.



Note

Multiple methods exist for installing ESXi in such an environment. This procedure highlights using the built-in KVM console and virtual media features in Cisco UCS Manager to map remote installation media to each individual server and connect to their FCP boot LUNs.

Log in to the Cisco UCS 6200 Fabric Interconnects

Cisco UCS Manager

1. Log in to the Cisco UCS 6200 fabric interconnects and launch the Cisco UCS Manager application.
2. In the main window, select the **Servers** tab.
3. Select **Servers > Service Profiles > root > VM-Host-Infra-01**.
4. Navigate to the Actions section and select the **KVM Console** link.
5. Select **Servers > Service Profiles > root > VM-Host-Infra-02**.
6. Navigate to the Actions section and select the **KVM Console** link.

Set Up the ESXi Install

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. In the KVM window, select the **Virtual Media** tab.
2. Click the **Add Image...** button in the window that displays.
3. Browse to the **ESXi installer ISO image file**.
4. Click **Open** to add the image to the list of virtual media.
5. Click the checkbox for **Mapped** next to the entry corresponding to the image you just added.
6. In the KVM window, select the **KVM** tab to monitor during boot.
7. In the KVM window, select the **Boot Server** button in the upper left corner.
8. Click **OK**.
9. Click **OK**.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. On reboot, the machine detects the presence of the ESXi install media.
2. Select the **ESXi Installer** from the menu that displays.
3. After the installer is finished loading, press **Enter** to continue with the install.
4. Read through the EULA and press **F11** to accept and continue with the install.

5. Select the **NetApp LUN** that you set up previously as the install disk for ESXi and then press **Enter** to continue.
6. Select the appropriate keyboard layout and press **Enter** to continue.
7. Enter and confirm the **Root password** and press **Enter** to continue.
8. The installer warns you that existing partitions will be removed on the volume. After you are sure this is what you want, press **F11** to install ESXi.
9. After the install is complete, be sure to unmap the ESXi install image by unchecking the **Mapped** checkbox in the Virtual Media window. This is so that the server reboots into ESXi and not into the installer.
10. The Virtual Media window might warn you that it is preferable to eject the media from the guest. Because we cannot do this (and the media is read-only), click **Yes** and unmap it anyway.
11. Back in the KVM tab, press **Enter** to reboot the server.

Set Up the ESXi Hosts' Management Networking

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. After the server is done rebooting, press **F2** (the Customize System option).
2. Log in with **root** as the login name and the root password entered above.
3. Select the **Configure Management Network** menu option.
4. Select the **VLAN (optional)** menu option.
5. Enter the <MGMT VLAN ID>, and press **Enter**.

Set Up Each ESXi Host's Management Networking

ESXi Host VM-Host-Infra-01

1. From the Configure Management Network menu, select the **IP Configuration menu** option.
2. Select the **Set static IP address and network configuration:** option using the space bar to manually set up the management networking.
3. Enter the IP address for managing the first ESXi host.
4. Enter the subnet mask for the first ESXi host.
5. Enter the default gateway for the first ESXi host.
6. Press **Enter** to accept the changes to the management networking.
7. Select the **DNS Configuration** menu option.
8. Because we manually specified the IP configuration for the ESXi host, we also must specify the DNS information manually.
9. Enter the primary DNS server's IP address.
10. (Optional) Enter the secondary DNS server's IP address.
11. Enter the Fully Qualified Domain Name (FQDN) for the first ESXi host.
12. Press **Enter** to accept the changes to the DNS configuration.
13. Press **Esc** to exit the Configure Management Network submenu.
14. Press **y** to confirm the changes made and return to the main menu.

15. Select **Test Management Network** and verify that the management network is set up correctly.
16. Press **Esc** to log out of the VMware Console.

ESXi Host VM-Host-Infra-02

1. From the Configure Management Network menu, select the **IP Configuration menu** option.
2. Select the **Set static IP address and network configuration:** option using the space bar to manually setup the management networking.
3. Enter the IP address for managing the second ESXi host.
4. Enter the subnet mask for the second ESXi host.
5. Enter the default gateway for the second ESXi host.
6. Press **Enter** to accept the changes to the management networking.
7. Select the DNS Configuration menu option.
8. Because we manually specified the IP configuration for the ESXi host, we also must specify the DNS information manually.
9. Enter the primary DNS server's IP address.
10. (Optional) Enter the secondary DNS server's IP address.
11. Enter the Fully Qualified Domain Name (FQDN) for the second ESXi host.
12. Press **Enter** to accept the changes to the DNS configuration.
13. Press **Esc** to exit the Configure Management Network submenu.
14. Press **y** to confirm the changes made and return to the main menu.
15. Select **Test Management Network** and verify that the management network is set up correctly.
16. Press **Esc** to log out of the VMware Console.

Download VMware vSphere Client and vSphere Remote Command Line

1. Open a Web browser and navigate to **http://<VM-Host-Infra-01 IP address>**.
2. Download and install both the **vSphere Client** and the Windows version of the **vSphere Remote Command Line**. Note that these downloads come from the VMware web site and Internet access is required on the management workstation.

Log in to VMware ESXi Host Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

1. Open the vSphere client and enter **<VM-Host-Infra-01 IP address>** as the host you are trying to connect to.
2. Enter **root** for the username.
3. Enter the root password.
4. Click the **Login** button to connect.

ESXi Host VM-Host-Infra-02

1. Open the **vSphere client** and enter **<VM-Host-Infra-02 IP address>** as the host you are trying to connect to.

2. Enter **root** for the username.
3. Enter the root password.
4. Click the **Login** button to connect.

Load Updated Cisco VIC enic Driver Version

Download and expand the zip files for the latest Cisco enic and fnic drivers, available for UCS servers. The enic version used for this reference architecture is 2.1.2.22, and the fnic version is 1.5.0.8. Be sure to download the latest drivers. For reference the links are below:

- <https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESX50-Cisco-enic-21222&productId=229>
- <https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI50-CISCO-fnic-1508&productId=229>

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. In the vSphere client, select the host on the left panel.
2. Go to the **Summary** tab.
3. On the right, under **Resources > Storage**, right-click **datastore1** and select **Browse Datastore...**
4. Click the fourth button and select **Upload File...**
5. Navigate to the folder where the expanded downloaded zip file is located and select the **net-enic-2.1.2.22-IOEM.500.0.0.441683.x86_64.vib** file.
6. Click **Open**.
7. Click **Yes**. The .vib file is uploaded to datastore1.
8. Click the fourth button and select **Upload File...**
9. Navigate to the folder where the expanded downloaded zip file is located and select the **scsi-fnic-1.5.0.8-IOEM.500.0.0.472560.x86_64.vib** file.
10. Click **Open**.
11. Click **Yes**. The .vib file is uploaded to datastore1.
12. Open the **VMware vSphere CLI Command Prompt** that was installed earlier.
13. For each ESXi Host in the Command Prompt, type **esxcli -s <ESXi Host IP> -u root -p <root password> software vib install -v /vmfs/volumes/datastore1/net-enic-2.1.2.22-IOEM.500.0.0.441683.x86_64.vib** as shown below.

```

Administrator: Command Prompt
VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.99 -u root -p NetApp123 software vib install -v /vmfs/volumes/datastore1/net-enic-2.1.2.2-10EM.500.0.0.441683.x86_64.vib
Installation Result
Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
Reboot Required: true
VIBs Installed: Cisco_bootbank_net-enic_2.1.2.2-10EM.500.0.0.441683
VIBs Removed: VMware_bootbank_net-enic_1.4.2.15a-1vmw.500.0.0.469512
VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp123 software vib install -v /vmfs/volumes/datastore1/scsi-fnic-1.5.0.8-10EM.500.0.0.472560.x86_64.vib
Installation Result
Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
Reboot Required: true
VIBs Installed: Cisco_bootbank_scsi-fnic_1.5.0.8-10EM.500.0.0.472560
VIBs Removed: VMware_bootbank_scsi-fnic_1.5.0.3-1vmw.500.0.0.469512
VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.99 -u root -p NetApp123 software vib install -v /vmfs/volumes/datastore1/scsi-fnic-1.5.0.8-10EM.500.0.0.472560.x86_64.vib
Installation Result
Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
Reboot Required: true
VIBs Installed: Cisco_bootbank_scsi-fnic_1.5.0.8-10EM.500.0.0.472560
VIBs Removed: VMware_bootbank_scsi-fnic_1.5.0.3-1vmw.500.0.0.469512
VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>

```

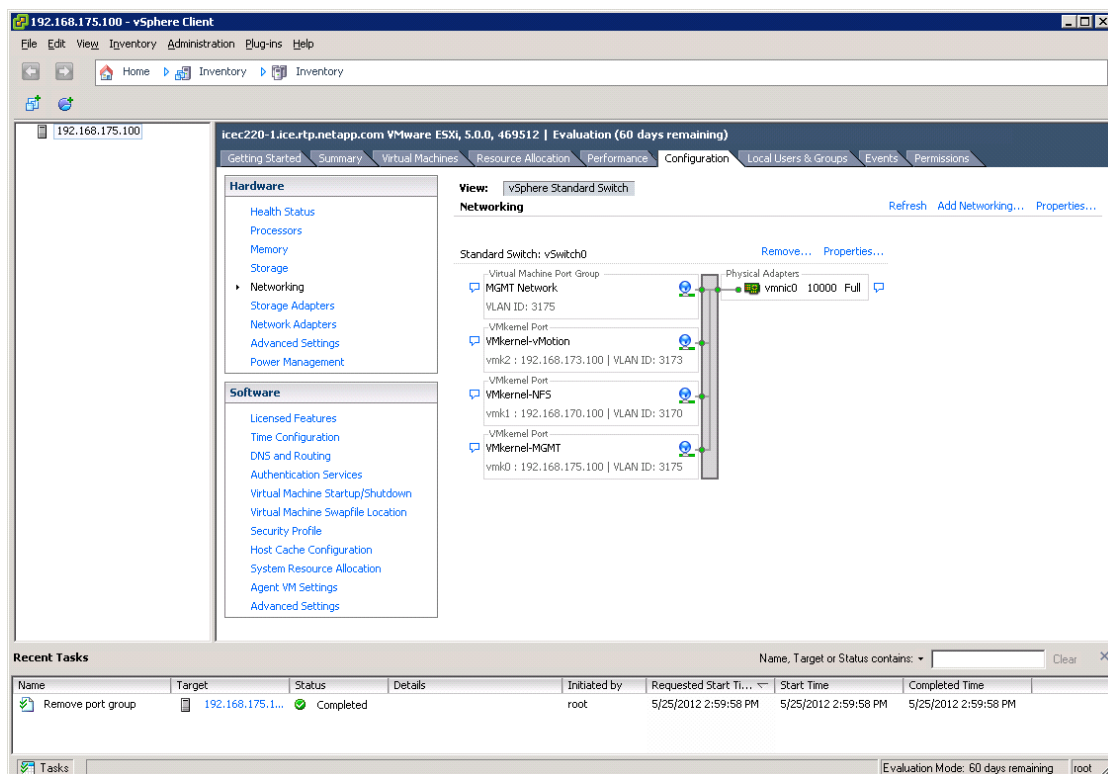
14. For each ESXi Host in the Command Prompt, type `esxcli -s <ESXi Host IP> -u root -p <root password> software vib install -v /vmfs/volumes/datastore1/scsi-fnic-1.5.0.8-10EM.500.0.0.472560.x86_64.vib` as shown above.
15. Back in the vSphere client, right-click the host on the left panel and select **Reboot**.
16. Click **Yes** to continue.
17. Enter a reason for the reboot and click **OK**.
18. After the reboot is complete log back into both hosts using the vSphere client.

Set up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

1. In vSphere client, select the host on the left panel.
2. Go to the **Configuration** tab.
3. Click the **Networking** link in the Hardware box.
4. Click the **Properties...** link in the right field on vSwitch0.
5. Select the **vSwitch configuration** and click **Edit...**
6. Under the **General** tab, change the MTU: to **9000**.
7. Click **OK**.
8. Select the **Management Network configuration** and click **Edit...**
9. Change the **Network Label:** to **VMkernel-MGMT** and select the **Management Traffic:** checkbox.
10. Click **OK**.


11. Select the **VM Network configuration** and click **Edit...**
12. Change the **Network Label:** to **MGMT Network** and enter **<MGMT VLAN ID>** for the VLAN ID (Optional): field.
13. Click **OK** to continue.
14. Click **Add...**
15. Select the **VMkernel** radio button and click **Next**.
16. Change the **Network Label:** to **VMkernel-NFS** and enter **<NFS VLAN ID>** for the VLAN ID (Optional): field.
17. Click **Next**.
18. Enter the IP Address and Subnet Mask for the NFS VLAN interface for **VM-Host-Infra-01**.
19. Click **Next**.
20. Click **Finish**.
21. Select the **VMkernel-NFS** configuration and click **Edit...**
22. Change the **MTU:** field to **9000** and click **OK**.
23. Click **Add...**
24. Select the **VMkernel** radio button and click **Next**.
25. Change the **Network Label:** to **VMkernel-vMotion**, enter **<vMotion VLAN ID>** for the VLAN ID (Optional): field, and select the **Use this port group for vMotion** checkbox.
26. Click **Next**.
27. Enter the IP address and subnet mask for the vMotion VLAN interface for **VM-Host-Infra-01**.
28. Click **Next**.
29. Click **Finish**.
30. Select the **VMkernel-vMotion** configuration and click **Edit...**
31. Change the **MTU:** field to **9000** and click **OK**.
32. Click **Close** to close the dialog box. The ESXi host networking setup should be similar to the one shown below.



ESXi Host VM-Host-Infra-02

1. In vSphere client, select the host on the left panel.
2. Go to the **Configuration** tab.
3. Click the **Networking** link in the Hardware box.
4. Click the **Properties...** link in the right field on vSwitch0.
5. Select the **vSwitch** configuration and click **Edit...**
6. Under the **General** tab, change the MTU: to **9000**.
7. Click **OK**.
8. Select the **Management Network** configuration and click **Edit...**
9. Change the **Network Label:** to **VMkernel-MGMT** and select the **Management Traffic:** checkbox.
10. Click **OK**.
11. Select the **VM Network** configuration and click **Edit...**
12. Change the **Network Label:** to **MGMT Network** and enter **<MGMT VLAN ID>** for the VLAN ID (Optional): field.
13. Click **OK** to continue.
14. Click **Add...**
15. Select the **VMkernel** radio button and click **Next**.
16. Change the **Network Label:** to **VMkernel-NFS** and enter **<NFS VLAN ID>** for the VLAN ID (Optional): field.

17. Click **Next**.
18. Enter the IP address and subnet mask for the NFS VLAN interface for **VM-Host-Infra-02**.
19. Click **Next**.
20. Click **Finish**.
21. Select the **VMkernel-NFS** configuration and click **Edit...**
22. Change the **MTU: field** to **9000** and click **OK**.
23. Click **Add...**
24. Select the **VMkernel** radio button and click **Next**.
25. Change the **Network Label: to VMkernel-vMotion**, enter **<vMotion VLAN ID>** for the VLAN ID (Optional): field, and select the **Use this port group for vMotion** checkbox.
26. Click **Next**.
27. Enter the IP address and subnet mask for the vMotion VLAN interface for **VM-Host-Infra-02**.
28. Click **Next**.
29. Click **Finish**.
30. Select the **VMkernel-vMotion** configuration and click **Edit...**
31. Change the **MTU: field** to **9000** and click **OK**.
32. Click **Close** to close the dialog box. The ESXi host networking setup should be similar to the one shown below.



Mount Required Datastores

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. In the vSphere client, select the host on the left panel.
2. Select the **Configuration** tab.
3. Select **Storage** in the Hardware box.
4. Select **Add Storage** in the top ribbon in the Datastore section on the right panel.
5. The Add Storage wizard displays. Select the **Network File System** radio button, and click **Next** to continue.
6. The wizard prompts for the location of the NFS export. Enter the **<Controller B NFS IP>** for the Server:.
7. Enter **/vol/infra_datastore_1** as the Folder: path to the NFS export.
8. Make sure the **Mount NFS read only** checkbox is not checked.
9. Enter **infra_datastore_1** as the Datastore Name.
10. Click **Next** to continue.
11. Review the information you entered, and click **Finish** to add the datastore.
12. Select **Add Storage** in the top ribbon in the Datastore section on the right panel.
13. The Add Storage wizard displays. Select the **Network File System** radio button, and click **Next** to continue.
14. The wizard prompts for the location of the NFS export. Enter the **<Controller A NFS IP>** for the Server:.
15. Enter **/vol/infra_swap** as the Folder: path to the NFS export.
16. Make sure the **Mount NFS read only** checkbox is not checked.
17. Enter **infra_swap** as the Datastore Name.
18. Click **Next** to continue.
19. Review the information you entered, and click **Finish** to add the datastore.

NTP Time Configuration

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. In the vSphere client, select the host on the left panel.
2. Select the **Configuration** tab.
3. Click the **Time Configuration** link in the Software box.
4. Click the **Properties...** link on the right panel.
5. A Time Configuration window displays. Click **Options...** at the bottom.
6. An NTP Daemon Options window displays. Select the **Start and stop with host** radio button. Select **NTP Settings** in the left box, then click **Add...**
7. Another pop-up window displays. Enter the **IP address of the NTP server**, and click **OK** to continue.
8. On the original NTP Daemon Options window, check the **Restart NTP service** to apply changes checkbox. Click **OK**.

9. Check the **NTP Client Enabled** checkbox. Click **OK** at the bottom of the window to continue and close the window.
10. On the Time Configuration window, verify that the clock is now set to the correct time.

Move the VM Swap File Location

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

1. In the vSphere client, select the host on the left panel.
2. Select the **Configuration** tab.
3. In the Software box, select **Virtual Machine Swapfile Location**.
4. On the right panel, click **Edit...**
5. Select the radio button for **Store the swapfile in a swap file datastore** selected below if it is not already selected.
6. Select **infra_swap** as the datastore you want to store the swapfile on.
7. Click **OK** at the bottom of the page to finish.

VMware vCenter 5.0 Deployment Procedure

The following sections provide detailed procedures for installing VMware vCenter 5.0 within a FlexPod environment. The deployment procedures that follow are customized to include the specific environment variables that have been noted in previous sections. By the end of this section, a VMware vCenter server will be configured along with a Microsoft SQL Server providing the database to support vCenter. Although this procedure walks through the installation and configuration of an external Microsoft SQL Server 2008 R2 database, other types of external databases are supported by vCenter. If you choose to use an alternate database, please refer to VMware vSphere 5.0 documentation for how to setup this database and integrate it into vCenter.

Build a Microsoft SQL Server VM

1. Log in to host **VM-Host-Infra-01** with the VMware vSphere Client.
2. In vSphere Client, select the host on the left panel.
3. Right-click the host and select **New Virtual Machine...**
4. Select the **Custom** radio button and click **Next**.
5. Name the Virtual Machine and click **Next**.
6. Select **infra_datastore_1** and click **Next**.
7. Select the **Virtual Machine Version: 8** radio button and click **Next**.
8. Make sure the **Windows** radio button and **Microsoft Windows Server 2008 R2 (64-bit) Version:** is selected and click **Next**.
9. Select **2** virtual sockets and **1** core per virtual socket and click **Next**.
10. Make sure **4 GB of memory** is selected and click **Next**.
11. Select **1** NIC total.
12. For NIC 1:, select **MGMT Network** and the **VMXNET 3 Adapter**. Click **Next**.


13. Leave the **LSI Logic SAS SCSI controller** selected and click **Next**.
14. Leave **Create a new virtual disk** selected and click **Next**.
15. Make the disk size at least 40 GB and click **Next**.
16. Click **Next**.
17. Select the checkbox next to **Edit the virtual machine settings before completion** and click **Continue**.
18. Select **Options** tab.
19. Select **Boot Options**.
20. On the right, select the **Force BIOS Setup** checkbox.
21. Click **Finish**.
22. On the left panel, expand the host field by clicking the **plus sign**.
23. Right-click the just created SQL Server Virtual Machine and click **Open Console**.
24. Click the third button (green right-arrow) to Power On the VM.
25. Click the ninth button (CD with a Wrench) to map the Windows Server 2008 R2 SP1 iso and select **Connect to ISO image on local disk...**
26. Navigate to the Windows Server 2008 R2 SP1 iso, select it, and click **Open**.
27. Back in the BIOS Setup Utility Window, click in the window and use the Right Arrow key to move to the **Boot menu**. Use the Down Arrow key to **highlight CD-ROM drive**. Use the + key two times to move CD-ROM Drive to the top of the list. Press **F10** and **Enter** to Save and Exit the BIOS Setup Utility.
28. The Windows Installer will boot. Select the appropriate Language, Time and currency format, and Keyboard and click **Next**. Click **Install** now. Make sure **Windows Server 2008 R2 Standard (Full Installation)** is selected and click **Next**. Accept the license terms and click **Next**. Select **Custom (advanced)**. Make sure **Disk 0 Unallocated Space** is selected and click **Next**. Windows Installation will complete.
29. After Windows Installation is complete and the VM has rebooted, click **OK** to enter the Administrator password. Enter and confirm the Administrator password and click the **Blue Arrow** to log in. Click **OK** to confirm the Password Change.
30. When you are logged in to the VM desktop, in the VM console window, select the **VM menu**. Under Guest, select **Install/Upgrade VMware Tools**. Click **OK**.
31. If prompted to eject the windows install media prior to running setup for VMware tools, click **OK**.
32. In the popup window, select **Run setup64.exe**.
33. In the VMware Tools installer window, click **Next**.
34. Make sure **Typical** is selected and click **Next**.
35. Click **Install**.
36. Click **Finish**.
37. Click **Yes** to restart the VM.
38. After the reboot is complete select the **VM menu**. Under Guest, select **Send Ctrl + Alt + del**. Then enter the password to log back into the VM.
39. In Server Manager, click the + sign next to Diagnostics in the left hand pane.
40. Click **Device Manager**. In the center pane, double-click **Display Adapters**.

41. Right-click on **Standard VGA Graphics Adapter** and, from the pop-up menu, click on **Update Driver Software...**
42. Click on **Browse my computer for driver software.**
43. In the drop down text box, type **C:\Program Files\Common Files\VMware\Drivers\wddm_video**. Make sure the **Include subfolders** check box is checked.
44. Click **Next**.
45. Verify Windows has successfully installed the VMware SVGA 3D video driver. Click **Close**.
46. Click **Yes** to restart the Guest OS.
47. After logging in, set the VM's timezone, IP address and gateway, and hostname. A reboot will be required.
48. Log back into the virtual machine and download and install all required Windows Updates. This will require several reboots.

Install Microsoft SQL Server 2008 R2

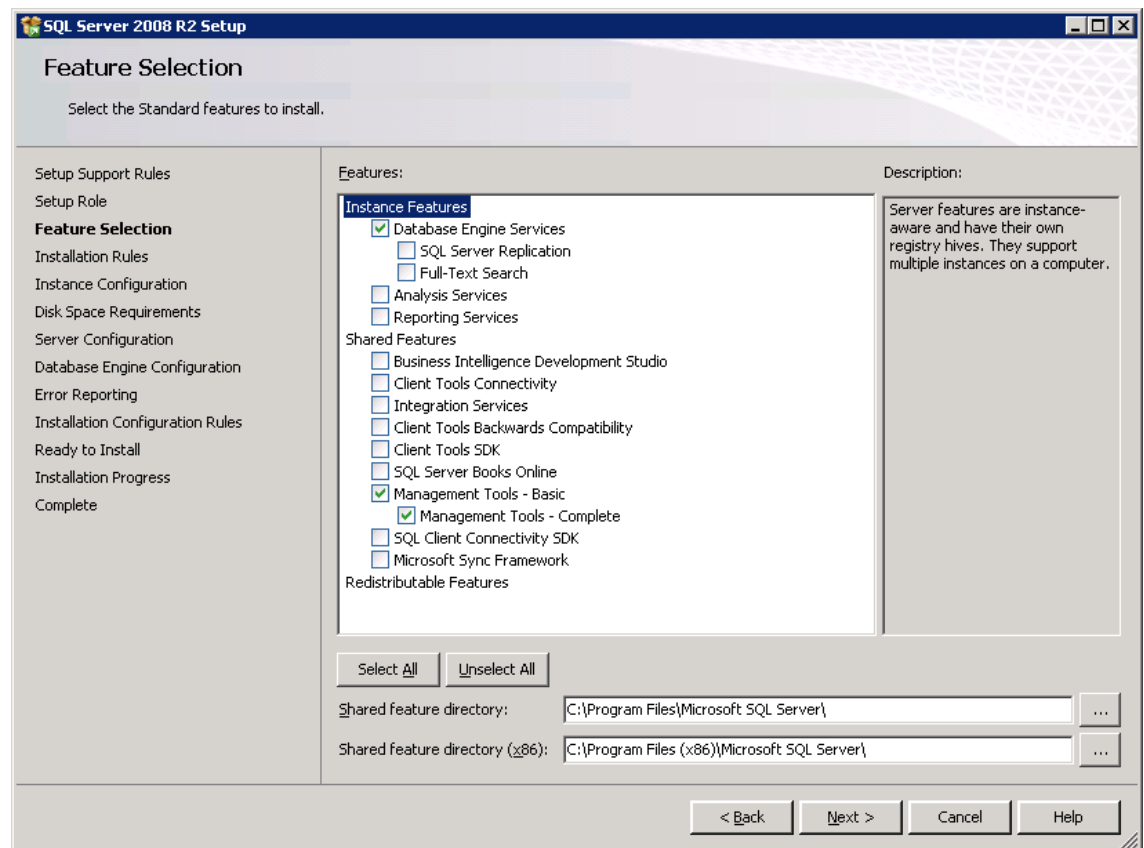
vCenter SQL Server VM

1. Log into the vCenter SQL server VM as the local administrator and open Server Manager.
2. Expand Features and click **Add Features**.
3. Expand .NET Framework 3.5.1 Features and select only **.NET Framework 3.5.1**.




4. Click **Next**.

5. Click **Install**.
6. Click **Close**.
7. Open Windows Firewall with Advanced Security by clicking **Start > Administrative Tools > Windows Firewall with Advanced Security**.
8. Highlight **Inbound Rules** and click **New Rule...**
9. Select **Port** and click **Next**.
10. Select **TCP** and enter the specific local port **1433**. Click **Next**.
11. Select **Allow the connection** and click **Next**, then **Next** again.
12. Name the rule **SQL Server** and click **Finish**.
13. Close **Windows Firewall with Advanced Security**.
14. In the vCenter SQL Server VMware console, click the ninth button (CD with a Wrench) to map the Microsoft SQL Server 2008 R2 iso and select **Connect to ISO image on local disk...**
15. Navigate to the **SQL Server 2008 R2 iso**, select it, and click **Open**.
16. In the popup, click **Run SETUP.EXE**.
17. In the SQL Server Installation Center window, click **Installation** on the left.
18. Click **New installation or add features to an existing installation** on the right.
19. Click **OK**.
20. Select **Enter the product key:**, enter a product key, and click **Next**.
21. Select the checkbox to accept the license terms and decide whether to select the second checkbox. Click **Next**.
22. Click **Install** to install the Setup Support Files.
23. Address any warnings except the Windows Firewall warning. The Windows Firewall issue was addressed earlier. Click **Next**.
24. Select **SQL Server Feature Installation** and click **Next**.
25. Under Instance Features, select **only Database Engine Services, Management Tools - Basic, and Management Tools - Complete**.
26. Click **Next**.

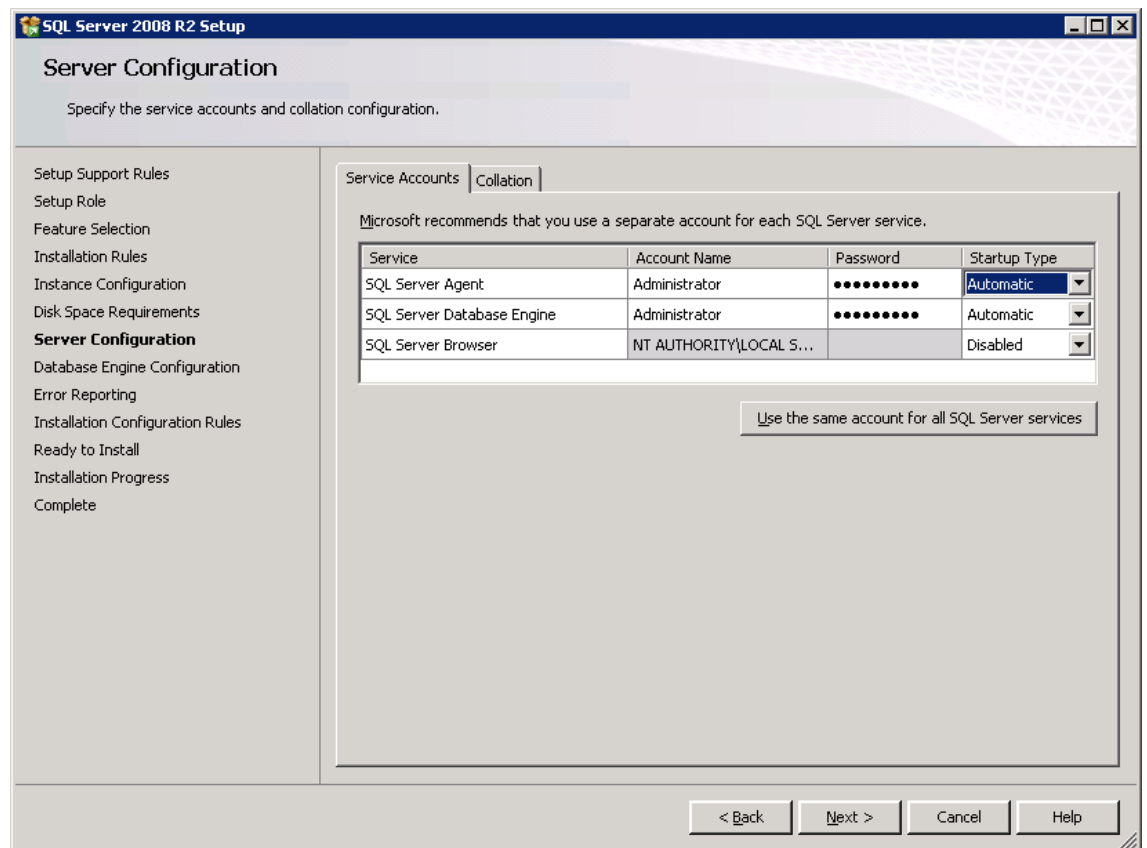


27. Click **Next**.

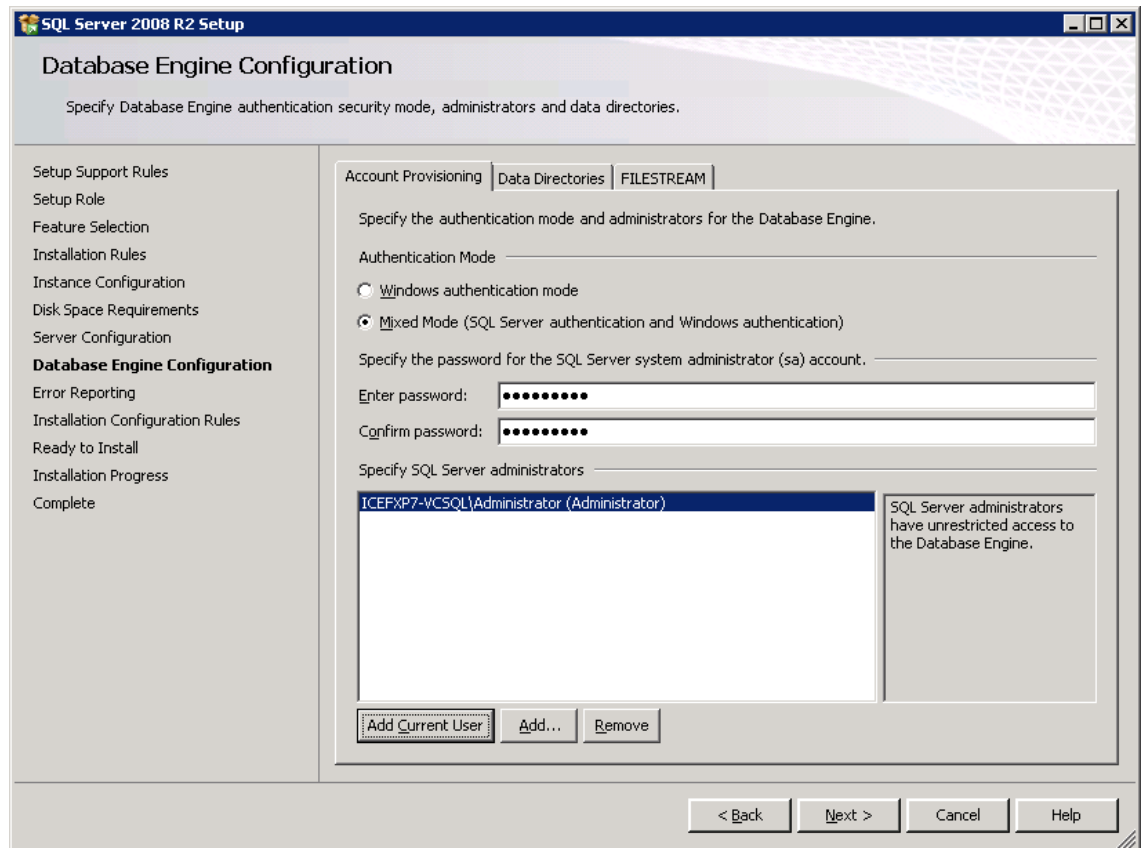
28. Leave **Default instance** selected and click **Next**.



29. Click **Next**.
30. Click in the first **Account Name** field next to SQL Server Agent and click <<**Browse...**>>. Enter the local machine Administrator name (example: system name\Administrator), check it and click **OK**. Enter the Administrator Password. Change the SQL Server Agent Startup Type to **Automatic**. Next to SQL Server Database Engine, select **Administrator** under Account Name and again enter the Administrator Password.
31. Click **Next**.



32. Select **Mixed Mode (SQL Server authentication and Windows authentication)**. Enter and confirm the password for the **sa** account.
33. Click **Add Current User**.
34. Click **Next**.



35. Decide whether to send Error Reports to Microsoft and click **Next**.
36. Click **Next**.
37. Click **Install**.
38. When the installation is complete, click **Close** to close the SQL Server Installer.
39. Close the **SQL Server Installation Center**.
40. Install all available Microsoft Windows updates by navigating to **Start -> All Programs -> Windows Updates**.
41. Open the SQL Server Management Studio by selecting **Start > All Programs > Microsoft SQL Server 2008 R2 > SQL Server Management Studio**.
42. Under Server name: select the **local machine** name. Under Authentication: select **SQL Server Authentication**. Enter **sa** for the Login: and the sa password. Click **Connect**.
43. On the left, click **New Query**.
44. Input the following script, substituting a vpxuser password for the **<Password>** variable:

```

use [master]

go

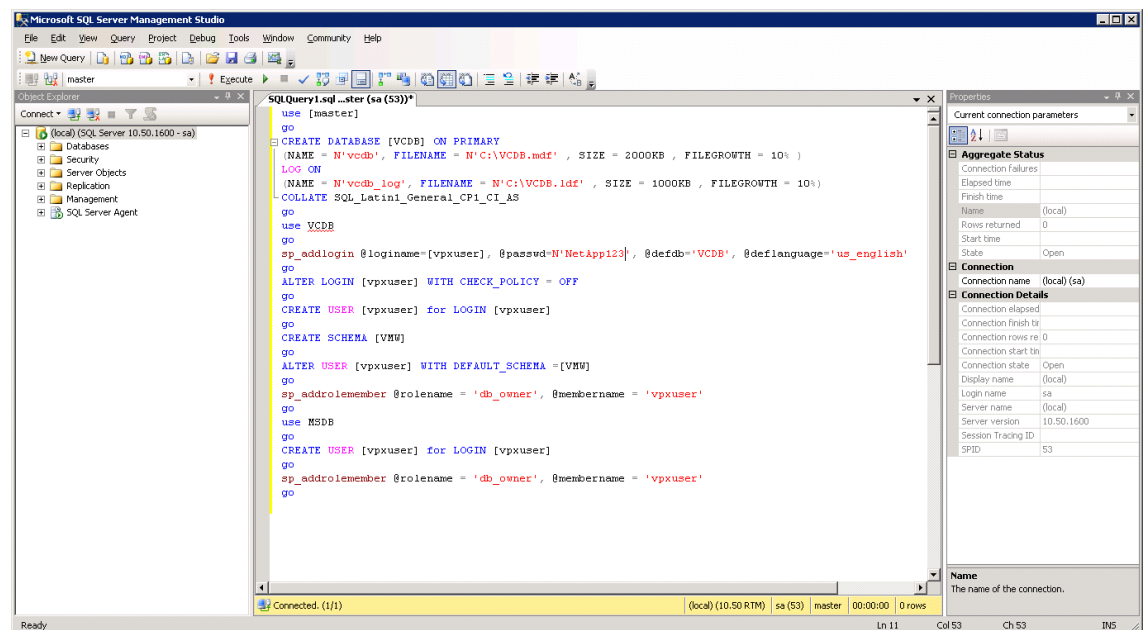
CREATE DATABASE [VCDB] ON PRIMARY
(NAME = N'vcdb', FILENAME = N'C:\VCDB.mdf' , SIZE = 2000KB , FILEGROWTH = 10% )
LOG ON

```

```

(NAME = N'vcdb_log', FILENAME = N'C:\VCDB.ldf' , SIZE = 1000KB , FILEGROWTH =
10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
sp_addlogin @loginame=[vpxuser], @passwd=N'<Password>', @defdb='VCDB',
@deflanguage='us_english'
go
ALTER LOGIN [vpxuser] WITH CHECK_POLICY = OFF
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
CREATE SCHEMA [VMW]
go
ALTER USER [vpxuser] WITH DEFAULT_SCHEMA =[VMW]
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go

```




45. In the upper-middle of the window, click **Execute**. The Query should execute successfully.

46. Close **Microsoft SQL Server Management Studio**.
47. Disconnect the Microsoft SQL Server 2008 R2 iso from the SQL Server virtual machine.

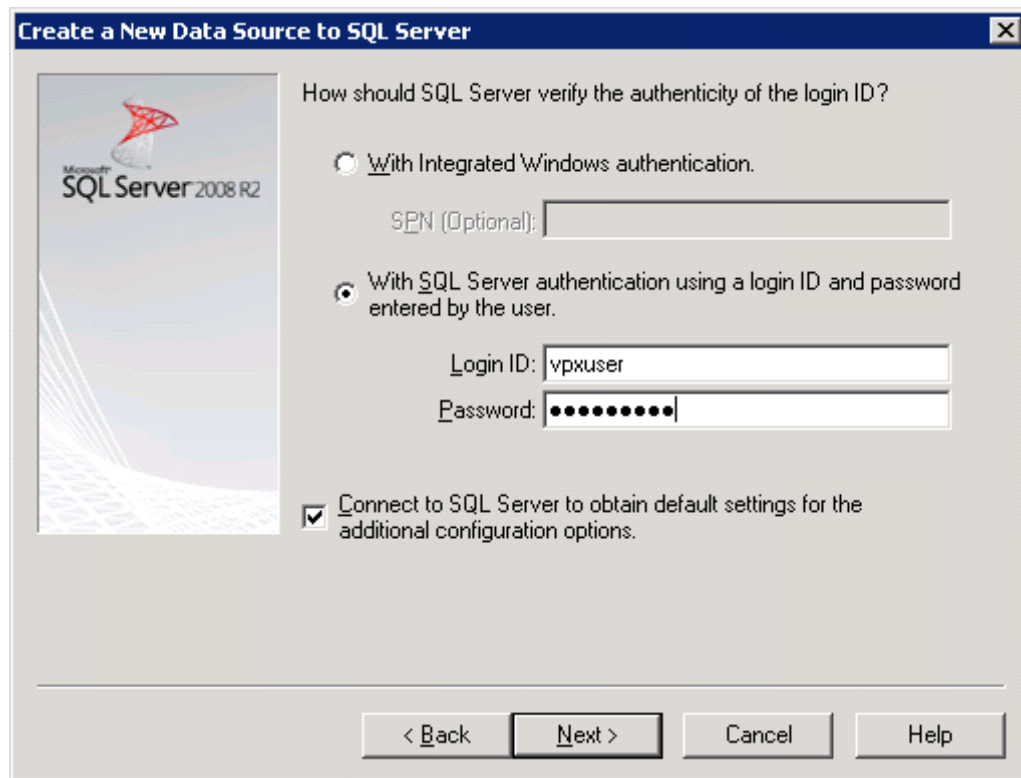
Build a VMware vCenter Virtual Machine

Using the instructions above that were used to build the SQL Server VM, build a VMware vCenter VM with 4 GB RAM, 2 CPUs, and one virtual network interface in the <**MGMT VLAN ID**> **VLAN**. Bring up the VM, install VMware Tools, and assign an IP address and host name in the Active Directory Domain.

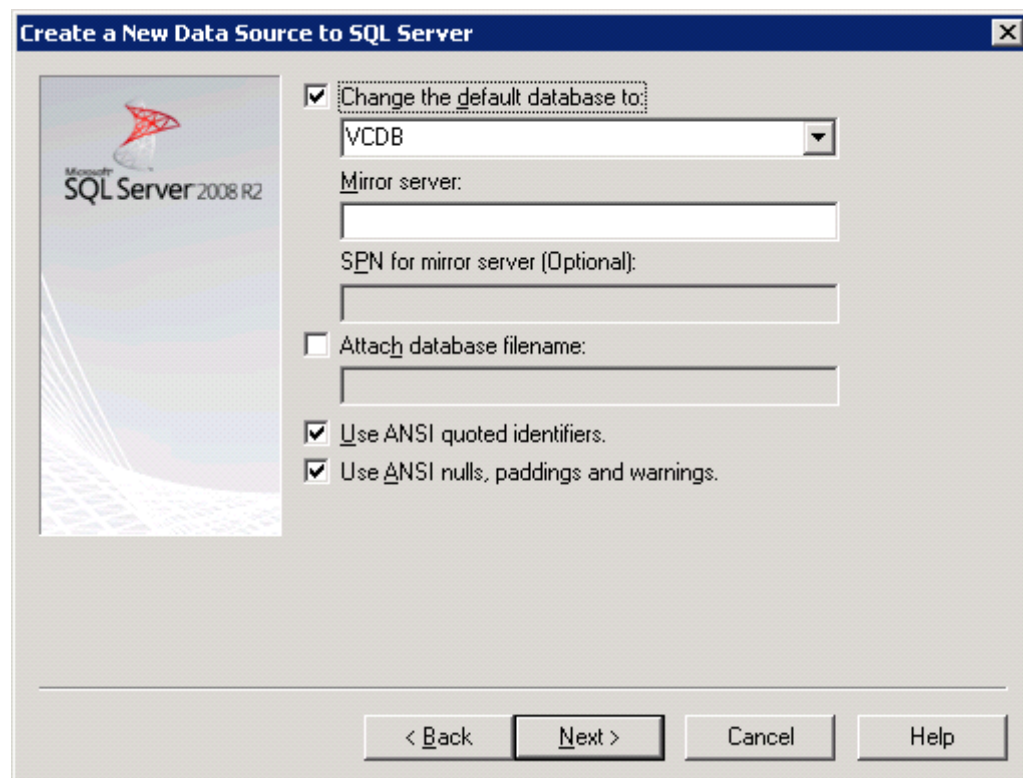
1. Log into the vCenter VM as the local Administrator and open **Server Manager**.
2. Expand **Features** and click **Add Features**.
3. Expand **.NET Framework 3.5.1 Features** and select only **.NET Framework 3.5.1**.
4. Click **Next**.
5. Click **Install**.
6. Click **Close** to close the Add Features Wizard.
7. Close **Server Manager**.
8. Download and install the Client Components of the Microsoft Server 2008 R2 Native Client from <http://go.microsoft.com/fwlink/?LinkID=188401&clid=0x409>.
9. Create the vCenter Database Data Source Name (DSN). Open **Data Sources (ODBC)** by selecting **Start > Administrative Tools > Data Sources (ODBC)**.
10. Select the **System DSN** tab.
11. Click **Add...**
12. Select **SQL Server Native Client 10.0** and click **Finish**.
13. Name the Data Source **VCDB**. In the Server: field, enter the IP address of the vCenter SQL server.
14. Click **Next**.



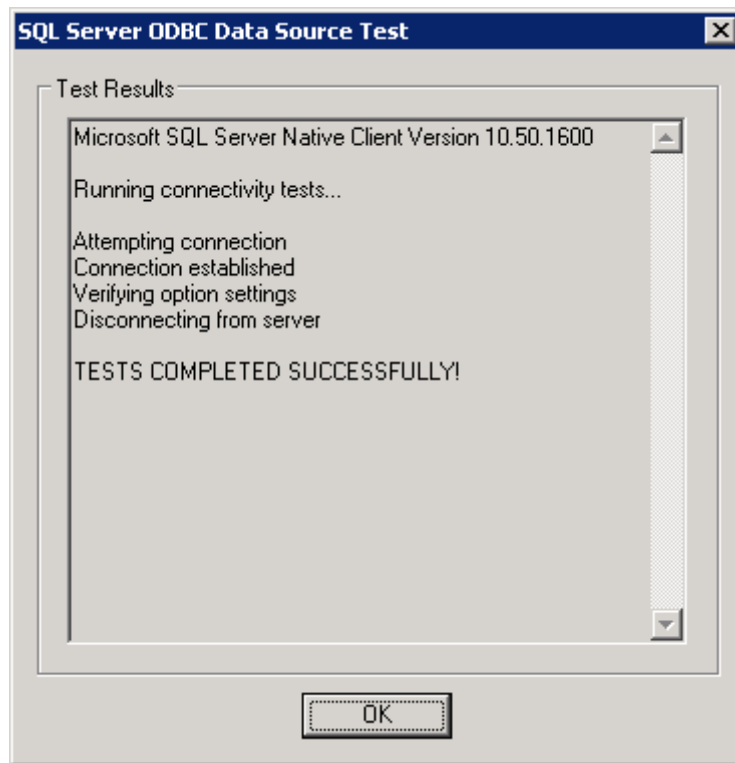
15. Select **With SQL Server authentication...** and input **vpuser** for the Login ID: and the vpuser password.
16. Click **Next**.



17. Select the checkbox next to **Change the default database to:**, select **VCDB** from the pulldown, and click **Next**.



18. Click **Finish**.
19. Click **Test DataSource...** The test should complete successfully.



20. Click **OK** then **OK**.
21. Install all available Microsoft Windows updates by navigating to **Start -> All Programs -> Windows Updates**.


**Note**

A restart may be required.


Install VMware vCenter Server

vCenter Server VM


1. In the vCenter Server VMware console, click the ninth button (CD with a Wrench) to map the VMware vCenter iso and select **Connect to ISO image on local disk...**
2. Navigate to the VMware vCenter 5.0 (VIMSetup) iso, select it, and click **Open**.
3. In the popup, click **Run autorun.exe**.
4. In the VMware vCenter Installer window, make sure **vCenter Server** is selected and click **Install**.



5. Select the appropriate language and click **OK** to continue.
6. Click **Next**.
7. Click **Next**.
8. Agree to the license terms and click **Next**.
9. Enter a User Name, Organization, and vCenter License key. Click **Next**.
10. Select **Use an existing supported database**, select **VCDB** using the pulldown, and click **Next**.



11. Enter the vpxuser Password and click Next.



12. Click **OK**.
13. Click **Next**.
14. Click **Next**.
15. Make sure **Create a standalone VMware vCenter Server instance** is selected and click **Next**.
16. Click **Next**.
17. Click **Next**.
18. Select the appropriate Inventory Size and click **Next**.
19. Click **Install**.
20. Click **Finish**.
21. Click **Exit** in the VMware vCenter Installer window.
22. Disconnect the **VMware vCenter iso** from the vCenter VM.
23. Install all available Microsoft Windows updates by navigating to **Start -> All Programs -> Windows Updates**.

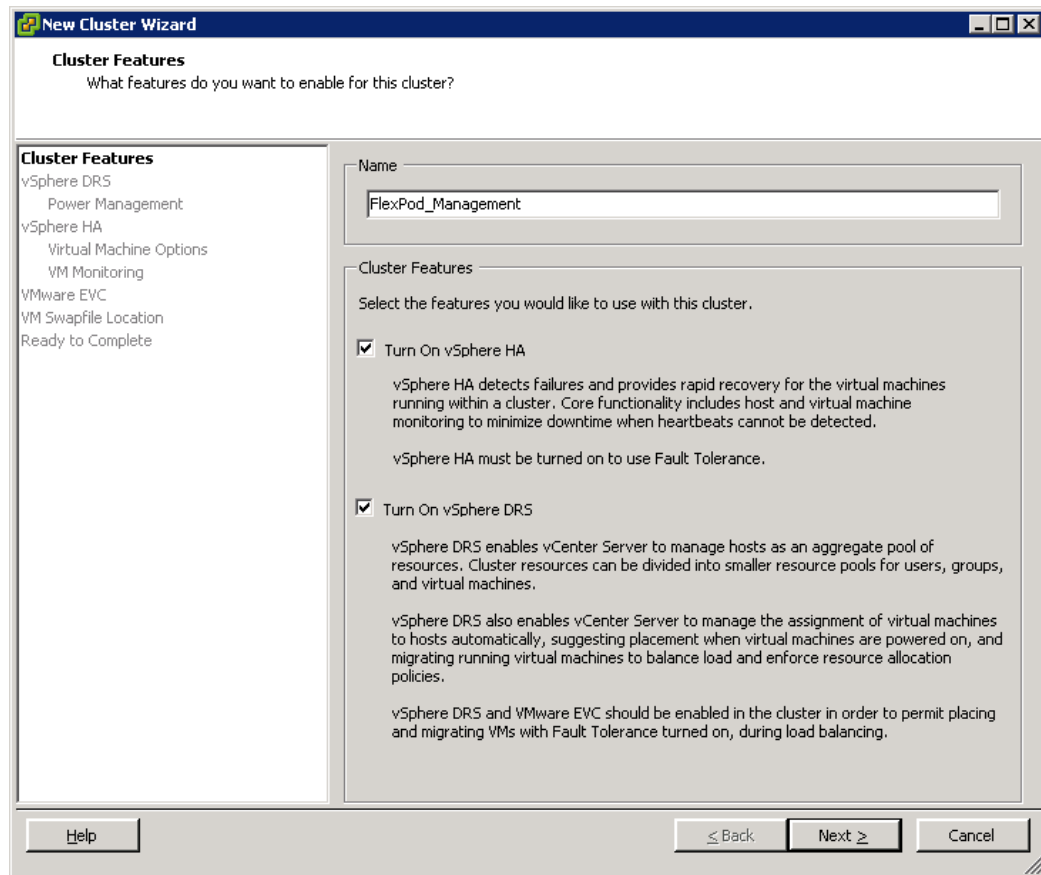
**Note**

A restart may be required.

vCenter Setup

vCenter Server Virtual Machine

1. Using the vSphere Client, log into the vCenter Server just created as Administrator.
2. Near the center of the window, click **Create a datacenter**.
3. Enter **FlexPod_DC_1** as the datacenter name.
4. Right-click the newly created **FlexPod_DC_1**, and select **New Cluster...**
5. Name the Cluster **FlexPod_Management** and select the checkboxes next to **Turn On vSphere HA** and **Turn on vSphere DRS**.
6. Click **Next**.



7. Accept the defaults for **vSphere DRS** and click **Next**.
8. Accept the defaults for **Power Management** and click **Next**.
9. Accept the defaults for **vSphere HA** and click **Next**.
10. Accept the defaults for **Virtual Machine Options** and click **Next**.
11. Accept the defaults for **VM Monitoring** and click **Next**.
12. Accept the defaults for **VMware EVC** and click **Next**.

**Note**

If mixing Cisco UCS B or C-Series M2 and M3 servers within a vCenter Cluster, it will be necessary to turn on VMware EVC mode. Please refer to VMware KB Article 1003212 for more information on setting up EVC mode.

13. Select **Store the swapfile in the datastore specified by the host** and click **Next**.
14. Click **Finish**.
15. Right-click the newly created **FlexPod_Management** cluster and select **Add Host...**
16. In the Host: field, enter either the IP address or hostname of the VM-Host-Infra_01 host. Enter **root** for the Username: and the root password for this host for the Password:. Click **Next**.
17. Click **Yes**.
18. Click **Next**.

19. Select **Assign a new license key to the host**. Click **Enter Key...** Enter a vSphere license key and click **OK**. Click **Next**.
20. Click **Next**.
21. Click **Next**.
22. Click **Finish**. VM-Host-Infra-01 is added to the cluster.
23. Using the instructions above, add **VM-Host-Infra-02** to the cluster.

Cisco Nexus 1010-X and 1000v Deployment Procedure

The following sections provide detailed procedures for installing a pair high-availability Cisco Nexus 1010-X virtual appliances in a FlexPod configuration. The primary and standby Nexus 1000v Virtual Supervisor Modules (vsm) will be installed on the 1010-Xs. The deployment procedures that follow are customized to include the specific environment variables that have been noted in previous sections. By the end of this section, a Nexus 1000v distributed virtual switch (dvs) will be provisioned. This procedure assumes that the Cisco Nexus 1000v software version 4.2.1.SV1.5.1a has been downloaded from www.cisco.com and expanded. This procedure also assumes that VMware vSphere 5.0 Enterprise Plus licensing is installed.

Configure the CIMC Interface on Both Nexus 1010-Xs

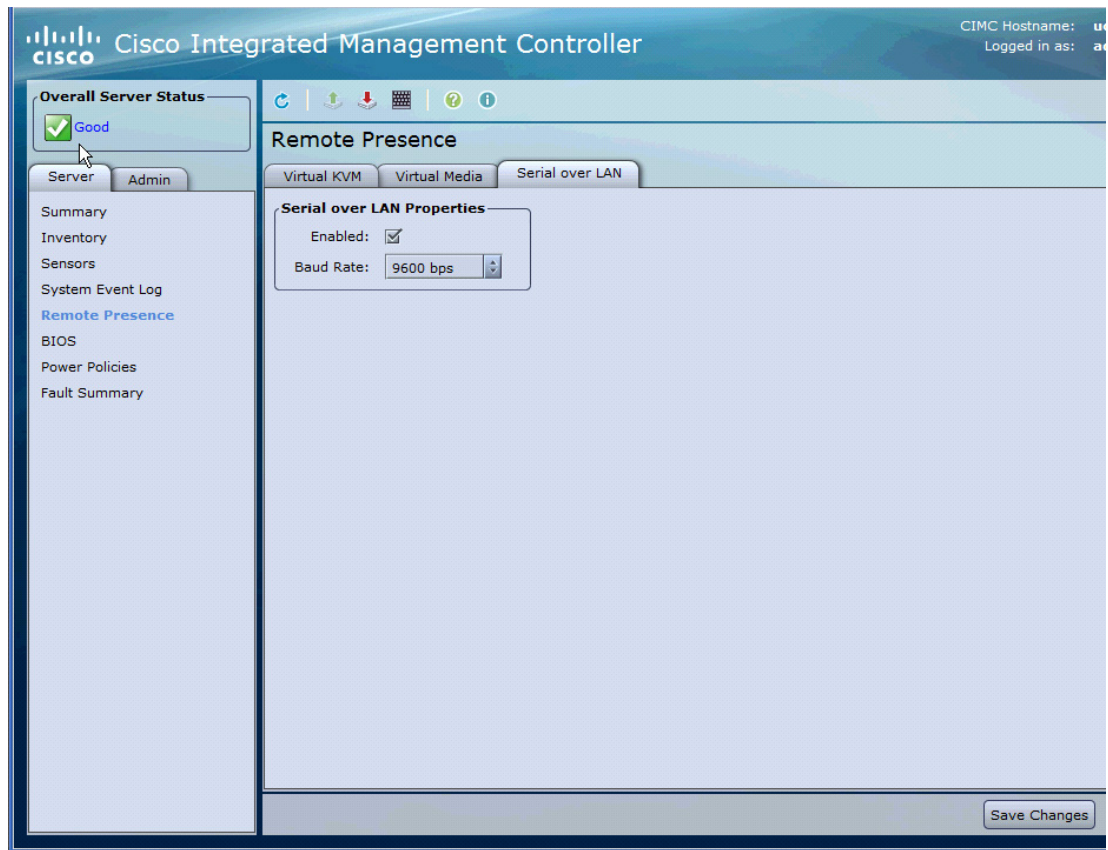
Nexus 1010-X A and Nexus 1010-X B

1. Connect to the console port on the back of the Cisco Nexus 1010-X virtual appliance.
2. Reboot the virtual appliance.
3. Press **F8** when prompted to configure the CIMC interface.
4. Unselect the box next to **DHCP enabled:**.
5. Set the CIMC IP address in the management VLAN.
6. Set the CIMC subnetmask.
7. Set the CIMC gateway.
8. Set and reenter the CIMC Default password.
9. Press **F10** to save the configuration.
10. Press **ESC** to reboot the virtual appliance.

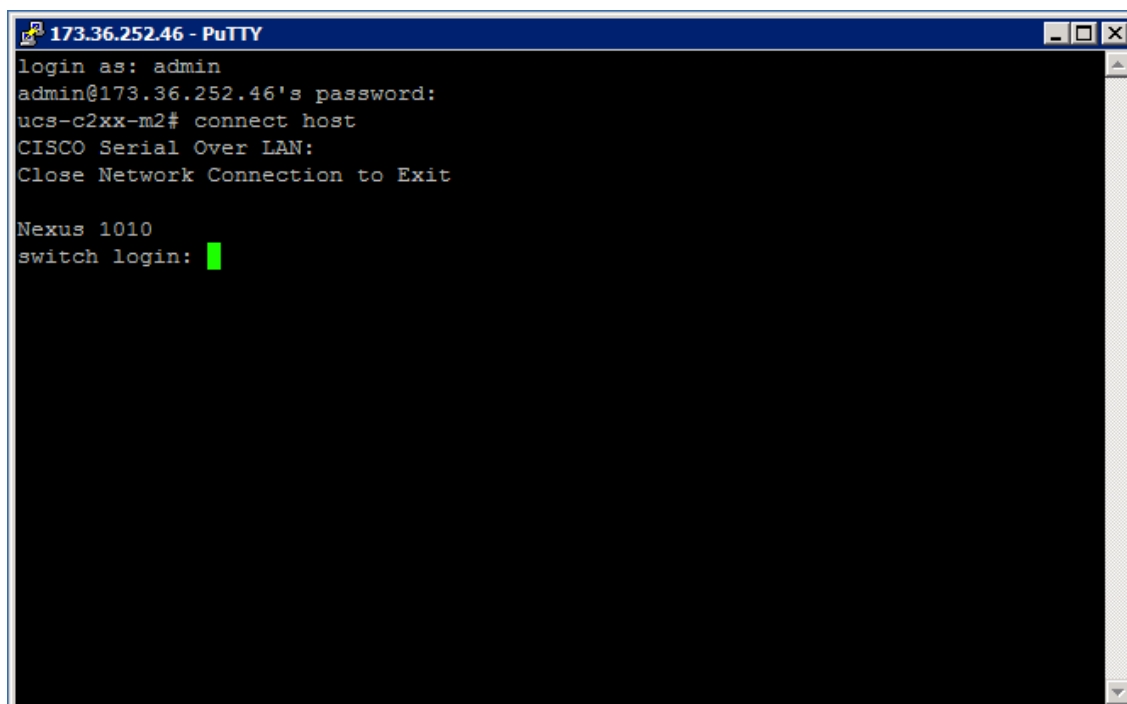
Configure Serial over LAN for both Nexus 1010-Xs

Nexus 1010-X A and Nexus 1010-X B

1. Use a web browser to open the url **https://<CIMC IP address>**.
2. Log in to the CIMC with the CIMC Default password.
3. In the left hand column, click on **Remote Presence**.
4. Click on the **Serial over LAN** tab.
5. Check the **Enabled: check box for Serial over LAN Properties**.
6. Using the pulldown, set the Baud Rate to **9600** bps.
7. Click the **Save Changes** button.



8. Log out of the CIMC web interface.
9. Using an ssh client, connect to <CIMC IP address>, using the default CIMC username and password.
10. Type **connect host**.



```

173.36.252.46 - PuTTY
login as: admin
admin@173.36.252.46's password:
ucs-c2xx-m2# connect host
CISCO Serial Over LAN:
Close Network Connection to Exit

Nexus 1010
switch login: █

```

Configure the Nexus 1010-X Virtual Appliances

Nexus 1010-X A

1. Reboot the virtual appliance. The appliance should boot into a setup mode.
2. Enter the password for the admin user.
3. Enter the default password again to confirm.
4. Enter **primary** for the HA role.
5. Enter **1** for the network uplink type.
6. Enter **<Packet Control VLAN ID>** for the control VLAN ID.
7. Enter a unique Nexus 1010 domain ID.
8. Enter **<MGMT VLAN ID>** for the management VLAN ID.
9. Enter **yes** to enter the basic system configuration dialogue.
10. Enter **no** to skip creating another login account.
11. Enter **no** to skip configuring a read-only SNMP community string.
12. Enter **no** to skip configuring a read-write SNMP community string.
13. Enter the hostname for the Nexus 1010-X VSA.
14. Enter **yes** to continue with the out-of-band management configuration.
15. Enter the Nexus 1010-X VSA management IP address.
16. Enter the Nexus 1010-X VSA management netmask.
17. Enter **yes** to configure the default gateway.
18. Enter the Nexus 1010-X VSA management default gateway.

19. Enter **no** to skip configuring the advanced IP options.
20. Enter **no** to skip enabling telnet service.
21. Enter **yes** to enable ssh service.
22. Enter **rsa** as the type of ssh key to be generated.
23. Enter **1024** for the number of key bits.
24. Enter **yes** to enable the http server.
25. Enter **yes** to enable the ntp server.
26. Enter the NTP server IP address.
27. Review the configuration summary. If everything is correct, enter **no** to skip editing the configuration.
28. Enter **yes** to use this configuration and save it.
29. The Nexus 1010-X will save the configuration and reboot. After reboot, log back in as **admin**.

Nexus 1010-X B

1. Reboot the virtual appliance. The appliance should boot into a setup mode.
2. Enter the password for the admin user that you entered on the Primary Nexus 1010-X.
3. Enter the default password again to confirm.
4. Enter **secondary** for the HA role.
5. Enter **1** for the network uplink type.
6. Enter **<Packet Control VLAN ID>** for the control VLAN ID.
7. Enter the same unique Nexus 1010 domain ID entered on Nexus 1010-X A.
8. Enter **<MGMT VLAN ID>** for the management VLAN ID.
9. The Nexus 1010-X will save the configuration and reboot.

Setup the Primary Cisco Nexus 1000v VSM

Nexus 1010-X A

1. Type **show module**. Continue periodically typing this command until module 2 (Nexus 1010-X B) has a status of ha-standby.
2. Enter the global configuration mode by typing **config t**.
3. Type **virtual-service-blade VSM-1** to create a VSB named VSM-1.
4. Type **dir /repository** to list the available Cisco Nexus 1000v ISO images.
5. If the desired Nexus 1000v ISO file (nexus-1000v.4.2.1.SV1.5.1a.iso), is not present on the Nexus 1010-X, use the copy command to copy it to the Nexus 1010-X disk. You will need to place the file either on an FTP server or on a Unix or Linux machine (scp) that is accessible from the Nexus 1010-X's management interface. An example copy command from an FTP server is **copy ftp://192.168.175.5/nexus-1000v.4.2.1.SV1.5.1a.iso /repository/**.
6. Type **virtual-service-blade-type new nexus-1000v.4.2.1.SV1.5.1a.iso**.
7. Type **interface control vlan <Packet Control VLAN ID>** to add the packet and control vlan to this VSB for control traffic.

8. Type **interface packet vlan <Packet Control VLAN ID>** to add the packet and control vlan to this VSB for packet traffic.
9. Type **enable primary** to enable this VSB as the primary.
10. Press **Enter** to verify the ISO image.
11. Enter a unique domain ID for the VSM. This domain ID should be different than the VSA domain ID.
12. Enter **V4** as the management IP version.
13. Enter the Nexus 1000v VSM management IP address.
14. Enter the Nexus 1000v VSM management netmask.
15. Enter the Nexus 1000v VSM default gateway IP address.
16. Enter the Nexus 1000v VSM's hostname or switch name.
17. Enter the password for the VSM's admin user.
18. Type **copy run start**.
19. Type **show virtual-service-blade summary**. Continue periodically typing this command until the PRIMARY VSM-1 has the state VSB POWERED ON.

Set Up the Secondary Cisco Nexus 1000v VSM

Nexus 1010-X A

1. Type **system switchover** to activate Nexus 1010-X B.

Nexus 1010-X B

1. Login to Nexus 1010-X B as the admin user.
2. Enter the global configuration mode by typing **config t**.
3. Type **virtual-service-blade VSM-1**.
4. Type **dir /repository** to list the available Cisco Nexu 1000v ISO images.
5. If the desired Nexus 1000v ISO file (nexus-1000v.4.2.1.SV1.5.1a.iso), is not present on the Nexus 1010-X, use the copy command to copy it to the Nexus 1010-X disk. You will need to place the file either on an FTP server or on a Unix or Linux machine (scp) that is accessible from the Nexus 1010-X's management interface. An example copy command from an FTP server is **copy ftp://192.168.175.5/nexus-1000v.4.2.1.SV1.5.1a.iso /repository/**.
6. Type **enable secondary** to enable this VSB as the secondary.
7. Press **Enter** to verify the ISO image.
8. Enter the unique domain ID that was entered for the primary VSM.
9. Enter **V4** as the management IP version.
10. Enter the Nexus 1000v VSM management IP address entered on the Primary VSM.
11. Enter the Nexus 1000v VSM management netmask.
12. Enter the Nexus 1000v VSM default gateway IP address.
13. Enter the Nexus 1000v VSM's hostname or switch name entered on the Primary VSM.
14. Enter the password for the VSM's admin user entered on the Primary VSM.

15. Type **show virtual-service-blade summary**. Continue periodically typing this command until both the PRIMARY and SECONDARY VSM-1s have the state VSB POWERED ON.
16. Type **copy run start**.
17. Type **system switchover on Nexus 1010-X B** to activate Nexus 1010-X A. Nexus 1010-X B will reboot.

Install the Virtual Ethernet Module (VEM) on each ESXi Host

vCenter Server Virtual Machine

1. From the main window in the vSphere Client connected to vCenter, select the first server in the list under the FlexPod Management cluster.
2. Select the **Summary** tab.
3. Under Storage on the right, right-click **infra_datastore_1** and select **Browse Datastore...**
4. Select the root folder (/) and click the 3rd button at the top to add a folder.
5. Name the folder **VEM** and click **OK**.
6. On the left, select the **VEM** folder.
7. Click the 4th button at the top and select **Upload File...**
8. Navigate to the **cross_cisco-vem-v142-4.2.1.1.5.1a.0-3.0.1.vib** file and click **Open**.
9. Click **Yes**. The VEM file should now appear in the VEM folder in the datastore.
10. Open the VMware vSphere CLI Command Prompt.
11. In the VMware vSphere CLI, for each ESXi Host, enter the following: **esxcli -s <Host Server IP> -u root -p <Root Password> software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v142-4.2.1.1.5.1a.0-3.0.1.vib**.

```

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.100 -u root -p NetApp123 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v142-4.2.1.1.5.1a.0-3.0.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: Cisco_bootbank_cisco-vem-v142-esx_4.2.1.1.5.1a.0-3.0.1
  VIBs Removed:
  VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp123 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v142-4.2.1.1.5.1a.0-3.0.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: Cisco_bootbank_cisco-vem-v142-esx_4.2.1.1.5.1a.0-3.0.1
  VIBs Removed:
  VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>_

```

Register the Nexus 1000v as a vCenter Plugin

1. Using a web browser, navigate to the **<Primary VSM IP Address>** using **https://**.

2. Right-click the **cisco_nexus_1000v_extension.xml** hyperlink and select **Save Target As...**
3. Save the xml document to the local Desktop.
4. In the vSphere Client connected to vCenter, select **Plug-ins > Manage Plug-ins...**
5. Right-click in the white space in the window and select **New Plug-in...**
6. Browse to the Desktop and select the **cisco_nexus_1000v_extension.xml** document saved earlier. Click **Open**.
7. Click **Register Plug-in**.
8. Click **Ignore**.
9. Click **OK**.
10. The Cisco_Nexus_1000v should now appear in the list of available plug-ins.
11. Click **Close** to close the Plug-in Manager.

Base Configuration of the Primary VSM

1. Using an ssh client, log into the Primary Nexus 1000v VSM as **admin**.
2. If you have your Nexus 1000v license product authorization key (PAK), the license must be installed.
3. Type **show license host-id**. The command will output the license VDH number.
4. From your Cisco Nexus 1000v software license claim certificate, locate the product authorization key.
5. On the Web, go to the Product License Registration site on the Cisco Software Download Web site.
6. From the Product License Registration Web site, follow the instructions for registering your VSM license. The license key file is sent to you in an e-mail. The license key authorizes use on only the host ID device. You must obtain separate license key file(s) for each of your Primary VSMs.
7. Copy your license to either an FTP server or a UNIX or Linux machine.
8. Using "**copy scp://**", or "**copy ftp://**" copy your license to the bootflash on the VSM.
9. Type **install license bootflash:<license filename>**.
10. Type **show license usage** and verify the license is installed.
11. Type **copy run start** to save the configuration.
12. Enter the global configuration mode by typing **config t**.
13. Type **svs connection vCenter**.
14. Type **protocol vmware-vim**.
15. Type **remote ip address <vCenter Server IP> port 80**.
16. Type **vmware dvs datacenter-name FlexPod_DC_1**.
17. Type **connect**.
18. Type **exit**.
19. Type **ntp server <NTP Server IP> use-vrf management**.
20. Type **vlan <MGMT VLAN ID>**.
21. Type **name MGMT-VLAN**.
22. Type **vlan <NFS VLAN ID>**.

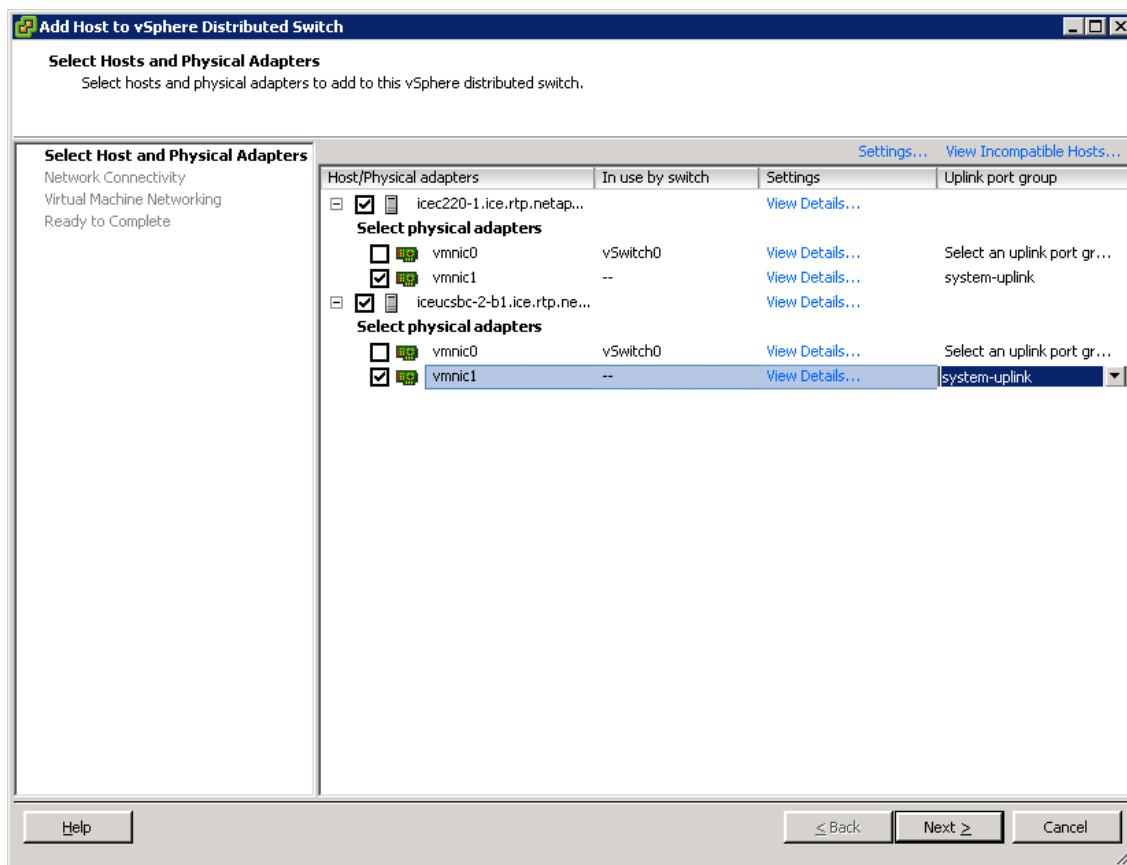
23. Type **name NFS-VLAN**.
24. Type **vlan <vMotion VLAN ID>**.
25. Type **name vMotion-VLAN**.
26. Type **vlan <Packet-Control VLAN ID>**.
27. Type **name Packet-Control-VLAN**.
28. Type **vlan <VM-Traffic VLAN ID>**.
29. Type **name VM-Traffic-VLAN**.
30. Type **vlan <Native VLAN ID>**.
31. Type **name Native-VLAN**.
32. Type **exit**.
33. Type **port-profile type ethernet system-uplink**.
34. Type **vmware port-group**.
35. Type **switchport mode trunk**.
36. Type **switchport trunk native vlan <Native VLAN ID>**.
37. Type **switchport trunk allowed vlan <MGMT VLAN ID>, <NFS VLAN ID>, <vMotion VLAN ID>, <Packet-Control VLAN ID>, <VM-Traffic VLAN ID>**.
38. Type **channel-group auto mode on mac-pinning**.
39. Type **no shutdown**.
40. Type **system vlan <MGMT VLAN ID>, <NFS VLAN ID>, <vMotion VLAN ID>, <Packet-Control VLAN ID>, <VM-Traffic VLAN ID>**.
41. Type **system mtu 9000**.
42. Type **state enabled**.
43. Type **port-profile type vethernet MGMT-VLAN**.
44. Type **vmware port-group**.
45. Type **switchport mode access**.
46. Type **switchport access vlan <MGMT VLAN ID>**.
47. Type **no shutdown**.
48. Type **system vlan <MGMT VLAN ID>**.
49. Type **state enabled**.
50. Type **port-profile type vethernet NFS-VLAN**.
51. Type **vmware port-group**.
52. Type **switchport mode access**.
53. Type **switchport access vlan <NFS VLAN ID>**.
54. Type **no shutdown**.
55. Type **system vlan <NFS VLAN ID>**.
56. Type **state enabled**.
57. Type **port-profile type vethernet vMotion-VLAN**.
58. Type **vmware port-group**.

59. Type **switchport mode access**.
60. Type **switchport access vlan <vMotion VLAN ID>**.
61. Type **no shutdown**.
62. Type **system vlan <vMotion VLAN ID>**.
63. Type **state enabled**.
64. Type **port-profile type vethernet VM-Traffic-VLAN**.
65. Type **vmware port-group**.
66. Type **switchport mode access**.
67. Type **switchport access vlan <VM-Traffic VLAN ID>**.
68. Type **no shutdown**.
69. Type **system vlan <VM-Traffic VLAN ID>**.
70. Type **state enabled**.
71. Type **exit**.
72. Type **copy run start**.


Migrate the ESXi Hosts' Networking to the Nexus 1000v

vCenter Server Virtual Machine


1. In the VMware vSphere Client connected to vCenter, select **Home > Networking**.
2. Expand the vCenter, DataCenter, and Nexus 1000v Folder. Select the **Nexus 1000v** switch.
3. Under Basic Tasks for the vSphere Distributed Switch, select **Add a host**.
4. For both hosts, select **vmnic1** and use the pulldown to select the **system-uplink Uplink port group**. Click **Next**.



- For all VMkernel ports, select the appropriate Destination port group from the Nexus 1000v. Click **Next**.



6. Select the **Migrate virtual machine networking** checkbox. Expand each **VM** and select the port groups for migration individually. Click **Next**.



7. Click **Finish**. Wait for the migration process to complete.
8. In the vSphere Client vCenter window, select **Home > Hosts and Clusters**.
9. Select the first ESXi host and select the **Configuration** tab. In the Hardware box, select **Networking**.
10. Make sure **vSphere Standard Switch** is selected at the top next to View:. vSwitch0 should have no active VMkernel or VM network ports on it. On the upper-right of vSwitch0, click **Remove...**
11. Click **Yes**.
12. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View:.
13. Click **Manage Physical Adapters...**
14. Scroll down to the system-uplink box and click **<Click to Add NIC>**.
15. Select **vmnic0** and click **OK**.
16. Click **OK** to close the Manage Physical Adapters window.
17. Two system-uplinks should now be present.
18. Select the second ESXi host and select the **Configuration** tab. In the Hardware box, select **Networking**.
19. Make sure **vSphere Standard Switch** is selected at the top next to View:. vSwitch0 should have no active VMkernel or VM network ports on it. On the upper-right of vSwitch0, click **Remove...**
20. Click **Yes**.

21. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View:.
22. Click **Manage Physical Adapters...**
23. Scroll down to the system-uplink box and click **<Click to Add NIC>**.
24. Select **vmnic0** and click OK.
25. Click **OK** to close the Manage Physical Adapters window.
26. Two system-uplinks should now be present.
27. Back in the ssh client connected to the Nexus 1000v, type **show interface status** to verify that all interfaces and port channels have been correctly configured.

```

192.168.175.193 - PuTTY
2012 Jun  1 22:14:45 icefxp1-vsm %VEM_MGR-2-MOD_ONLINE: Module 3 is online

icefxp1-vsm# show interface status

-----
Port          Name                Status  Vlan    Duplex  Speed  Type
-----
mgmt0         --                  up      routed  full    1000   --
Eth3/1        --                  up      trunk   full    10G    --
Eth3/2        --                  up      trunk   full    10G    --
Eth4/1        --                  up      trunk   full    10G    --
Eth4/2        --                  up      trunk   full    10G    --
Po1           --                  up      trunk   full    10G    --
Po2           --                  up      trunk   full    10G    --
Veth1         VMware VMkernel, v up    3175   auto    auto    --
Veth2         VMware VMkernel, v up    3170   auto    auto    --
Veth3         VMware VMkernel, v up    3173   auto    auto    --
Veth4         VMware VMkernel, v up    3175   auto    auto    --
Veth5         VMware VMkernel, v up    3170   auto    auto    --
Veth6         VMware VMkernel, v up    3173   auto    auto    --
Veth7         icefxp1-vc, Networ up    3175   auto    auto    --
Veth8         icefxp1-vcsql, Net up    3175   auto    auto    --
control0      --                  up      routed  full    1000   --
icefxp1-vsm#

```

28. Type **show module** and verify that the two ESXi hosts are present as modules.

```

192.168.175.193 - PuTTY
icefxp1-vsm# show module
Mod  Ports  Module-Type                Model                Status
---  -
1    0      Virtual Supervisor Module  Nexus1000V           active *
2    0      Virtual Supervisor Module  Nexus1000V           ha-standby
3    248    Virtual Ethernet Module    NA                    ok
4    248    Virtual Ethernet Module    NA                    ok

Mod  Sw                Hw
---  -
1    4.2(1)SV1(5.1a)    0.0
2    4.2(1)SV1(5.1a)    0.0
3    4.2(1)SV1(5.1a)    VMware ESXi 5.0.0 Releasebuild-469512 (3.0)
4    4.2(1)SV1(5.1a)    VMware ESXi 5.0.0 Releasebuild-469512 (3.0)

Mod  MAC-Address(es)                Serial-Num
---  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3    02-00-0c-00-03-00 to 02-00-0c-00-03-80  NA
4    02-00-0c-00-04-00 to 02-00-0c-00-04-80  NA

Mod  Server-IP          Server-UUID          Server-Name
--More--

```

29. Type **copy run start**.
30. Use **exit** to log out of the Nexus 1000v.

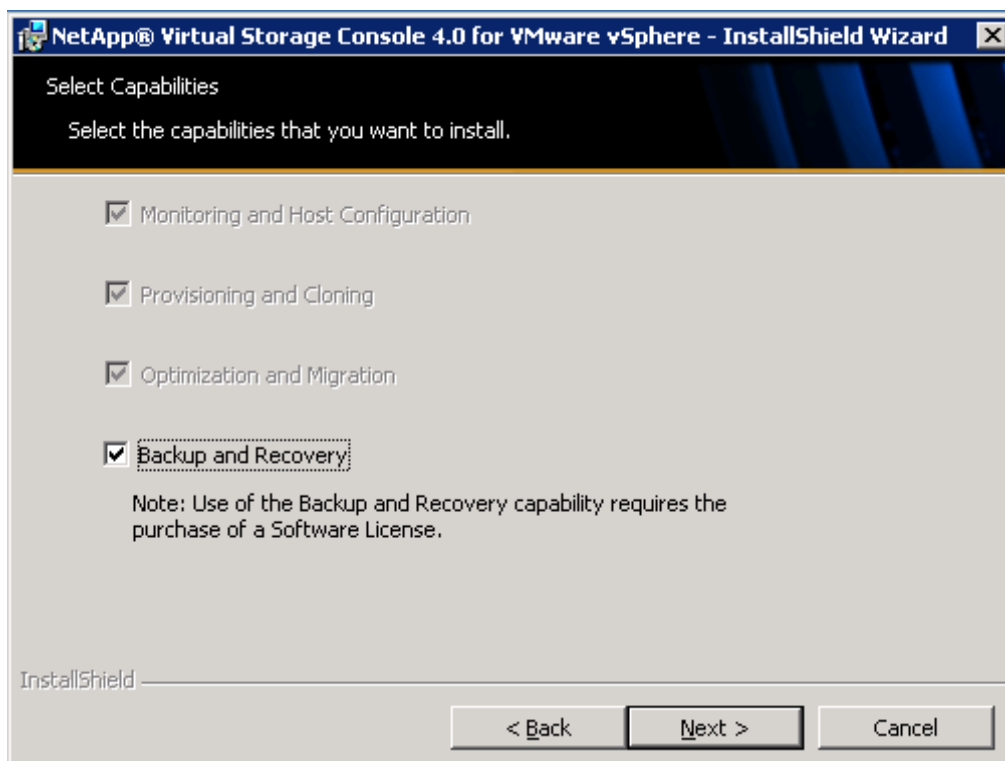
NetApp Virtual Storage Console Deployment Procedure

The following subsections (through “Provisioning and Cloning Setup”) provide detailed procedures for installing the NetApp Virtual Storage Console. The deployment procedures that follow are customized to include the environment variables discussed in previous sections. By the end of this section, a VSC will be a configured and operational plug-in with VMware vCenter.

Installing NetApp Virtual Storage Console 4.0

1. Using the previous instructions for virtual machine creation, build a VSC and OnCommand virtual machine with 4GB RAM, two CPUs, and two virtual network interfaces, one in the **<MGMT VLAN ID> VLAN** and the other in the **<NFS VLAN ID> VLAN**. The second virtual network interface should be a VMXNET 3 adapter. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory® domain. Install the current version of Adobe Flash Player on the virtual machine. Install all Windows updates on the VM, but do not install Internet Explorer 9. Keep Internet Explorer 8 on the virtual machine.
2. Configure jumbo frames on the network adapter in the **<NFS VLAN ID> VLAN**. Open Server Manager and click **View Network Connections**. Right-click the **network connection** in the **<NFS VLAN ID> VLAN** and select **Properties**. Click **Configure Select the Advanced** tab. Select the **Jumbo Packet** property and use the pull-down menu to select **Jumbo 9000**. Click **OK**. Close the **Network Connections** window.
3. Download the Virtual Storage Console 4.0 from the NetApp Support site.

4. To install the VSC plug-in, double-click the file that you downloaded in step 2 (for example, VSC-4.0-win64.exe).
5. On the installation wizard landing page, select **Next** at the bottom of the screen to proceed with the software installation.
6. Select the **Backup and Recovery** checkbox and click **Next** if you have the necessary licenses.



7. Select the location where VSC will be installed and click **Next**.
8. Make a note of the registration URL. This URL is needed to register the VSC plug-in with vCenter after the installation. Click **Install**.

**Note**

A browser window with the URL opens automatically when the installation phase is complete. However, some browser settings might interfere with this function. If the browser window does not open automatically, open one manually and enter the URL.

9. Open a Web browser.
10. Enter the URL provided by the installation wizard or replace **localhost** with the host name or IP address of the VSC server:
<https://localhost:8143/Register.html>
11. In the Plug-in service information section, select the IP address that the vCenter server uses to access the VSC server.
12. In the vCenter Server information section, enter the host name or IP address, port, user name, and user password and click **Register** to complete the registration.

vSphere Plugin Registration - Windows Internet Explorer

https://localhost:8143/f

vSphere Plugin Registration

vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information

Host name or IP Address: 192.168.175.215

vCenter Server information

Host name or IP Address: 192.168.175.213

Port: 443

User name: icefxp7-vc\Administrator

User password:

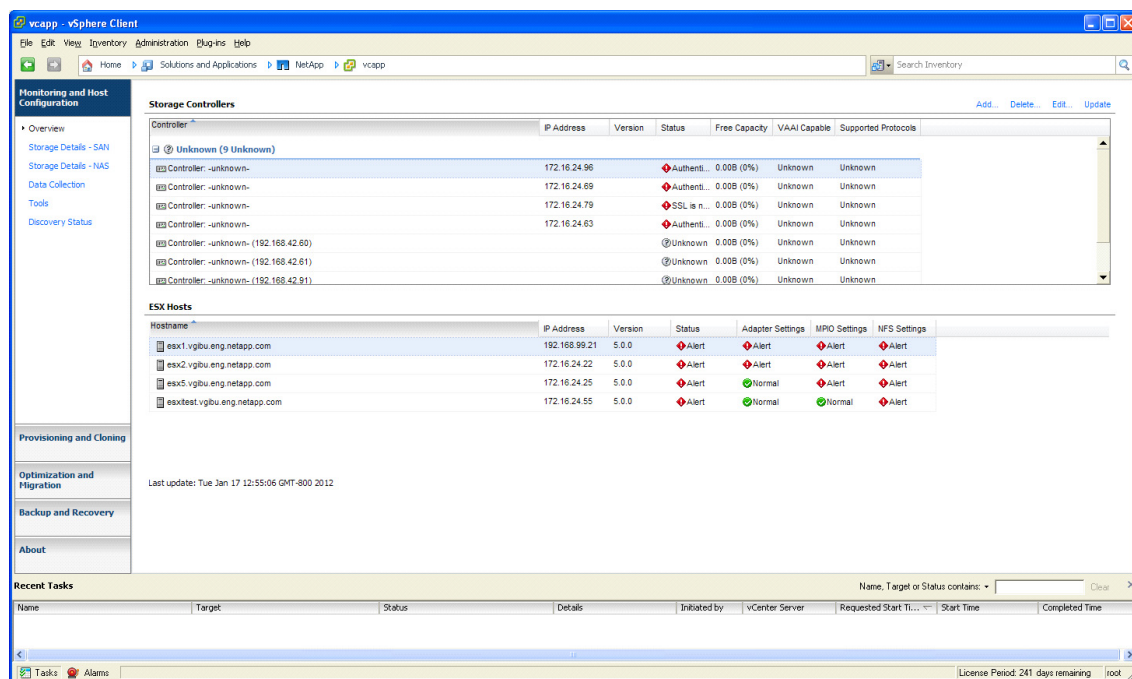
Register

13. Click **Finish** in the install wizard window.

Discover and Add Storage Resources

Do the following steps to add storage resources for monitoring, host configuration, and for provisioning and cloning.

1. Log in to **vCenter** using the vSphere client.
2. Click the **Home** tab in the upper left of the window.
3. In the Solutions and Applications section, click **NetApp**.
4. If the discovery process does not start automatically, click **Update** in the Overview panel of the Monitoring and Host Configuration screen. All NetApp storage elements associated to hosts that are managed by vCenter are discovered and displayed.



5. Add credentials for discovered storage controllers:
 - a. Right-click a discovered storage controller.
 - b. Select **Modify Credentials**.
 - c. Make sure that the storage controller IP address is in the NFS VLAN.
 - d. Enter the user name and password.
 - e. Select the **Use SSL** checkbox.
 - f. Click **OK**.

**Note**

This discovery process applies to monitoring and host configuration and to provisioning and cloning. Storage controllers must be manually added for backup and recovery.

Optimal Storage Settings for ESXi Hosts

VSC allows the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. These steps describe setting the recommended values for NetApp attached ESXi hosts.

1. Select individual or multiple ESXi hosts and right-click them to open the drop-down menu.
2. Select **Set Recommended Values** for these hosts.

**Note**


This functionality sets values for HBAs and CNAs, sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

Storage Controllers
















| Controller | IP Address | Version | Status | Free Capacity | VAAI Capable | Supported Protocols |
|------------|-------------|-------------------------|--------|----------------|--------------|---------------------|
| netapp | 10.10.10.11 | NetApp Release 8.0.1... | Normal | 20.18GB (74... | Enabled | iSCSI, NFS |
| netapp2 | 10.10.10.12 | NetApp Release 8.0.1... | Normal | 24.76GB (91... | Enabled | iSCSI, NFS |

ESX Hosts

| Hostname |
|------------|
| 10.10.10.1 |
| 10.10.10.2 |
| 10.10.10.3 |



| NFS Settings |
|--------------|
| Normal |
| Normal |
| Normal |

| ESX Hosts | | | | | | |
|--|------------|---------|--|--|--|--|
| Hostname | IP Address | Version | Status | Adapter Settings | MPIO Settings | NFS Settings |
|  10.10.10.1 | 10.10.10.1 | 4.1.0 |  Normal |  Normal |  Normal |  Normal |
|  10.10.10.2 | 10.10.10.2 | 4.1.0 |  Normal |  Normal |  Normal |  Normal |
|  10.10.10.3 | 10.10.10.3 | 4.1.0 |  Normal |  Normal |  Normal |  Normal |




Note

Depending on what changes have been made, servers might require a restart for network-related parameter changes to take effect. If no reboot is required, the Status value is set to Normal. If required, the Status is set to Pending Reboot. If required, the ESXi hosts should be evacuated, placed into maintenance mode, and restarted before proceeding.

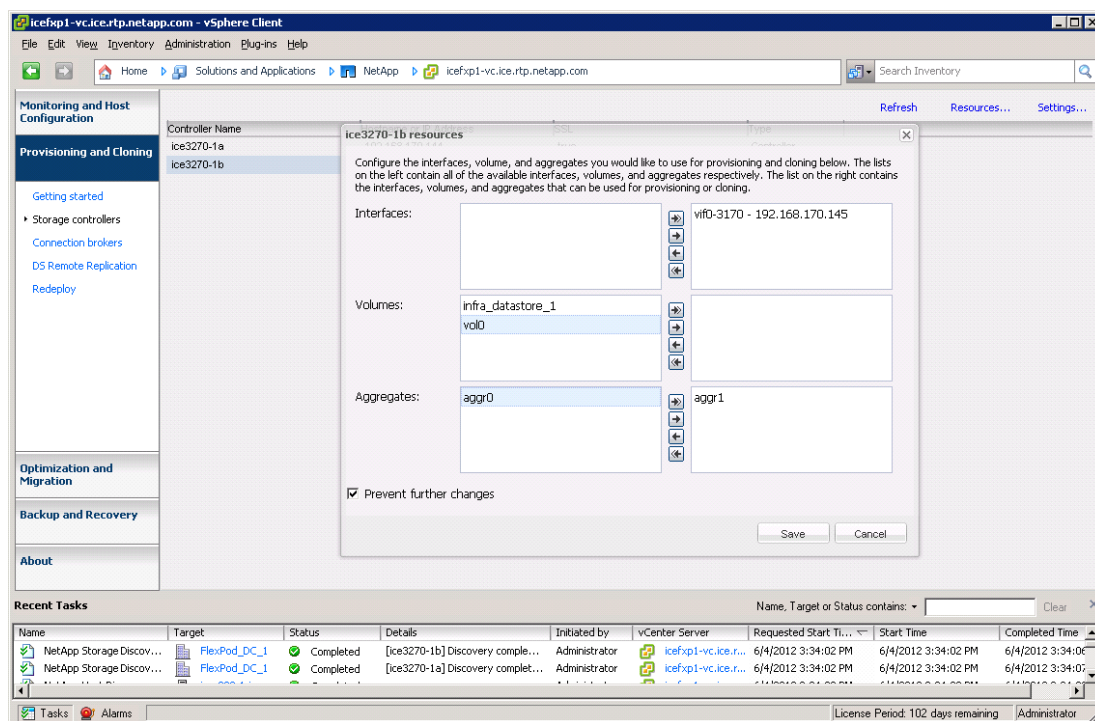
Provisioning and Cloning Setup

Provisioning and Cloning in VSC 4.0 help administrators to provision both VMFS and NFS datastores at the data center, datastore cluster, or host level in VMware environments. Furthermore, the target storage can be storage controllers running in either Cluster-Mode or 7-Mode.

1. In a vSphere client connected to vCenter, select **Home > Solutions and Applications > NetApp** and select the **Provisioning and Cloning** tab on the left. Select **Storage controllers**.
2. In the main part of the window, right-click **<Storage Controller A>** and select **Resources**.
3. In the **<Storage Controller A>** resources window, use the arrows to move volumes **vol0** and **infra_swap** and **aggregate aggr0** to the left. Also select the **Prevent further changes** checkbox as shown in the following screenshot.



4. Click **Save**.
5. In the main part of the window, right-click **<Storage Controller B>** and select **Resources**.
6. In the **<Storage Controller B>** resources window, use the arrows to move volumes **vol0** and **infra_datastore_1** and aggregate **aggr0** to the left. Select the **Prevent Further changes** checkbox as shown in the following screenshot.



7. Click **Save**.

NetApp OnCommand Deployment Procedure

The following subsections provide detailed procedures for installing the NetApp OnCommand software. The deployment procedures that follow are customized to include the environment variables described in previous sections. By the end of this section (through “Configure Host Services”), an OnCommand System Manager will be configured and operational.

Install NetApp OnCommand

1. Open the console on the VSC OnCommand (DataFabric Manager) virtual machine in vCenter.
2. Open a Web browser, navigate to the NetApp Support site, and download OnCommand Core 5.0: http://support.netapp.com/NOW/download/software/occore_win/5.0/.
3. Verify the solution components and the system requirements.
4. Click **Continue** at the bottom of the screen.
5. Click **Accept** to accept the license agreement.
6. Click the **Windows 64-bit (occore-setup-5-0-win-x64.exe)** OnCommand Core package file and save it.
7. Double-click the **64-bit installer (occore-setup-5-0-win-x64.exe)** to start the installation of the DataFabric Manager software.

**Note**

The 32 bit OnCommand Core package is not supported on 64-bit systems.

8. Click **Next** to continue installation.
9. Select **I accept for the AutoSupport notice** and click **Next**.
10. Select **Standard Edition** and click **Next**.

**Note**

The Express Edition is for a small environment (four storage systems and one virtualized host) and also has limited monitoring capabilities. Refer to the Install Guide for more details.

11. Enter the license key and click **Next**.
12. Click **Next** to accept the default installation location, or modify the location as necessary.
13. Click **Install** to install DataFabric Manager.
14. Verify that the installation is completed, then click **Next**.
15. Make sure that **Launch OnCommand Console** is selected, then click **Finish** to complete the installation.

**Note**

To install DFM as an HA service in a Microsoft MSCS environment, follow the instructions given in the Install Guide.

16. Log in to the DataFabric Manager server as the local administrator.
17. Enter the following command from an MS-DOS command prompt to generate the SSL keys on the host: **dfm ssl server setup -z 1024 -d 365 -c <global ssl country> -s <global ssl state> -l <global ssl locality> -o <global ssl org> -u <global ssl org unit> -n <global ntap dfm hostname> -e <ntap admin email address>.**

**Note**

The command in step 17 is one complete command. Take care when pasting this command into the command line. Also, any two-word parameters, such as “North Carolina” for the state, should be enclosed in double quotation marks.

**Note**

On a clustered system, the SSL keys must be copied to the other cluster nodes. The files are located in the directory **<install-dir>/conf**. For more information, refer to the product installation and upgrade guide.

18. Log in to the DataFabric Manager server and configure it to communicate by using SSL. Disable HTTP access by enabling or disabling the following options:

dfm option set httpsEnabled=yes

dfm option set httpEnabled=no

dfm option set httpsPort=8443

dfm option set hostLoginProtocol=ssh

dfm option set hostAdminTransport=https

dfm option set discoverEnabled=no

**Note**

Storage controllers that are being monitored by Operations Manager must have the HTTPS and SSH protocols enabled.

19. Restart the HTTP service on the DataFabric Manager server to make sure that the changes take effect.

```
dfm service stop http
```

```
dfm service start http
```

20. Issue the following command to configure Operations Manager to use SNMPv3 to poll configuration information from the storage devices:

```
dfm snmp modify -v 3 -c <ntap SNMP community> -U <ntap SNMP users> -P <ntap SNMP password> -A md5 default
```

**Note**

Verify that the storage controllers have the same SNMP community defined as specified in this command. NetApp recommends deleting the public SNMP domain and creating a custom SNMP community. NetApp also recommends using SNMPv3.

- g. If SNMPv3 is not in use, enter the following command: **dfm snmp modify -v 1 -c <ntap SNMP community> default.**
21. Enter the following commands to configure the DataFabric Manager AutoSupport feature:


```
dfm option set SMTPServerName=<ntap autosupport mailhost>
```

```
dfm option set autosupportAdminContact=<ntap admin email address>
```

```
dfm option set autosupportContent=complete
```

```
dfm option set autosupportProtocol=https
```

Manually Add Storage Controllers into DataFabric Manager

The following steps provide details for the manual addition of storage controllers into the DataFabric Manager server.

1. Log in to the DataFabric Manager server.
2. Use the following commands to add a storage system manually:


```
dfm host add <ntap A hostname>
```

```
dfm host add <ntap B hostname>
```
3. Use the following commands to set the array login and password credentials:


```
dfm host set <ntap a hostname> hostlogin=root
```

```
dfm host set <ntap A hostname> hostPassword=<global default password>
```

```
dfm host set <ntap B hostname> hostlogin=root
```

```
dfm host set <ntap B hostname> hostPassword= <global default password>
```
4. List the storage systems discovered by DataFabric Manager and list the storage system properties:


```
dfm host list
```

```
dfm host get <ntap A hostname>
```

```
dfm host get <ntap B hostname>
```

**Note**

If the arrays being added or discovered use a custom SNMP community, then the correct SNMP community string must be defined before the arrays can be discovered because the default discovery method uses the “ro public” SNMP community. This does not work if a custom SNMP string is defined.

Run Diagnostics for Verifying DataFabric Manager Communication

The following steps provide details for verifying DataFabric Manager communication by running diagnostics. This helps identify misconfigurations that can prevent the DataFabric Manager server from monitoring or managing a particular appliance, and it should be the first command you use when troubleshooting.

1. Log in to the DataFabric Manager Server.
2. Use the following commands to run diagnostics:


```
dfm host diag <ntap A hostname>
dfm host diag <ntap B hostname>
```
3. You can also refresh host properties after a change or do a force discovery using the following commands:


```
dfm host discover <ntap A hostname>
dfm host discover <ntap B hostname>
```

Configure Additional Operations Manager Alerts

The following steps provide details for configuring an SNMP trap host as well as configuring daily e-mails from Operations Manager for critical alerts.

1. Log in to the DFM server.
2. Use the following command to configure an SNMP traphost.


```
dfm alarm create -T <ntap SNMP traphosts>
```
3. Use the following command to configure daily e-mails from Operations Manager:


```
dfm alarm create -E <ntap admin email address> -v Critical
```

Deploy the NetApp OnCommand Host Package

The following steps provide details for deploying the OnCommand Host Package.

1. Log in to the DataFabric Manager server where you intend to install the Host Package.
2. Install the required software for the Host Package.
 - a. Open **Server Manager**.
 - b. Click **Features**.
 - c. Click **Add Feature**.
 - d. Expand the **.NET Framework** and select **.NET Framework 3.5.1**. Do not select **WCF Activation**.
 - e. Click **Next**.

- f. Click **Install**.
 - g. Click **Close**.
3. Open a Web browser and navigate to
<http://support.netapp.com/NOW/download/software/ochost/1.1/>
4. Verify the solution components and the system requirements.
5. Click **Continue** at the bottom of the screen.
6. Click **Accept** to accept the license agreement.
7. Verify that the prerequisites for the Host Package installation are satisfied.
8. Click the appropriate **OnCommand Host Package file: Windows 64-bit (ochost-setup-1-1-x64.exe)** to download the Host Package installer.
9. Double-click the **64-bit installer (ochost-setup-1-1-x64.exe)** to start the installation of the OnCommand Host Package software.
10. Click **Next** to continue installation.
11. Click **Next** to accept the default installation location or modify the location as necessary.
12. Specify the Service Credentials by entering the username and password in the domainname\username format, then click **Next**.
13. On the Configure Communication Ports page, enter the port numbers to use, or accept the default port numbers; then click **Next**.
14. On the Configure DataFabric Manager server page, enter the IP address and the username and password to access the DataFabric Manager server and click **Next**. You can skip the validation of the DataFabric Manager server if you do not have the server credentials available during the installation.

**Note**

In this configuration, the OnCommand Host Server is the same as the DataFabric Manager server.

15. Enter the vCenter server information on this page:
 - a. Enter the IP address of the system on which you are installing the OnCommand Host Package (use the management VLAN address).
 - b. Enter the host name or IP address of the system on which the vCenter server is installed and the username and password that allow the vSphere client to communicate with the vCenter server, then click **Next**.
16. Click **Install** on the summary page, then click **Finish**.

**Note**

After you finish, the host service must be configured to perform backups. You must associate storage systems with a host service when you finish installing the OnCommand Host Package. For the purposes of this reference architecture, NetApp recommends using VSC.

Set a Shared Lock Directory to Coordinate Mutually Exclusive Activities on Shared Resources



Note

If you plan to install the OnCommand Host Package on the same system as Virtual Storage Console 4.0, a best practice is to set up a shared lock directory. A shared lock directory is used for products that share resources through the vSphere client. This makes certain that mutually exclusive functions do not happen in parallel.

1. Stop the OnCommand Host Service VMware plug-in by using “Services” in Administrative Tools on the system:
 - a. Select **Start >Administrative Tools**.
 - b. Select **Services**.
 - c. Locate the OnCommand Host Service VMware Plug-in.
 - d. Right-click the **OnCommand Host Service VMware Plug-in** and select **Stop**.
2. Delete the **locks subdirectory** in the OnCommand Host Package installation directory:
`<OC_install_directory>\locks`.
3. Locate and open the **smvi.override** file. This file is installed in the following location by default:
`<OC_install_directory>\VMware Plugin\etc\smvi.override`.
4. Add the following line in the smvi.override file:
`shared.subplugin.lock.directory=<VSC installation directory>\locks`
5. Save and close the **smvi.override** file.
6. Restart the **OnCommand Host Service VMware Plug-in**.

Install NetApp OnCommand Windows PowerShell Cmdlets

The following steps describe installing the NetApp OnCommand Windows PowerShell cmdlets.

1. Navigate to the installation folder for OnCommand Core Package, then navigate to the folder that has the Windows PowerShell installation package: `<DFM_Install_dir>\DFM\web\clients` folder.
2. Copy the Windows PowerShell executable installation file to the system that has the host package software installed.
3. Execute the installation package and follow the installation wizard prompts and finish the install.

Configure Host Services

The following steps provide details for configuring host services.

1. To open Windows Firewall with Advanced Security, click **Start > Administration Tools > Windows Firewall with Advanced Security**.
2. Select **Inbound Rules**.
3. For each OnCommand rule and the SnapManager® for Virtual Infrastructure rule, right-click and select **Properties**. Select the **Advanced** tab. Select the **Public** checkbox. Click **OK**. When all changes have been made, all of these rules should show All under Profile.
4. Close **Windows Firewall with Advanced Security**.
5. Verify that the host service is registered with the DataFabric Manager server to correctly discover objects.

- a. On the system that has the OnCommand core package installed, open a Web browser and enter **https://localhost:8443** to open the OnCommand console.
 - b. Log in to the **console**.
 - c. Select **Administration > Host Services**.
 - d. In the Host Services list, verify that the name of the host service is listed. The status of the host service should be Pending Authorization.
 - e. If the host service is not displayed, add and configure the new host service by clicking **Add** on the Host Services tab and entering the correct properties.
6. Authorize the new host service to access the storage system credentials. You must authorize the host service to access the storage system to create jobs and to see the reporting statistics.
 - a. Select **Administration > Host Services**.
 - b. In the Host Services list, select a host service and click **Edit**.
 - c. In the Edit Host Service dialog box, click **Authorize**, review the certificate, and click **Yes**. If the Authorize area is not available, then the host service is already authorized.
7. Associate a host service with vCenter to provide part of the communication needed for discovery, monitoring, and backup and recovery of virtual server objects such as virtual machines and datastores:
 - a. Select **Administration > Host Services**.
 - b. In the Host Services list, select a host service and click **Edit**.
 - c. Enter the vCenter server properties by specifying the host name or FQDN for the host service registration, then click **OK**. If the details were specified during the installation, then the property fields are already populated.
8. Verify communication between the host service and the OnCommand Plug-in.
 - a. Select **View Menu > Server**.
 - b. Scroll through the list of virtual machines to verify that the VMs related to the host services are listed.
9. Associate storage systems with a host service. You must associate one or more storage systems that host virtual machines for the host service. This enables communication between the service and storage to provide that storage objects such as virtual disks are discovered and that the host service features work properly.
 - a. Click **Administration > Host Services**.
 - b. In the Host Services list, select the host service with which you want to associate storage and click **Edit**.
 - c. In the Storage Systems area, click **Associate**. To associate storage systems shown in the available storage systems list, select the system names and click the right-arrow button, then click **OK**. To associate a storage system that is not listed, click **Add**, enter the required information, and click **OK**.
 - d. The newly associated storage systems are displayed in the storage systems area.
 - e. In the list of storage systems, verify that the status is Good for the login and NDMP credentials. If the status is other than Good for any storage system, you must edit the storage system properties to provide the correct credentials before you can use that storage system.
 - f. Click **OK**.
10. Click **Refresh** to see the associated storage systems.

NetApp VASA Provider Deployment Procedure

The following subsections provide detailed procedures for installing the NetApp VASA Provider. The VASA Provider provides information to VMware vCenter about the capabilities of both VMFS and NFS datastores. The deployment procedures that follow are customized to include the specific environment variables that have been noted in previous sections. By the end of this section, a VASA Provider virtual machine will be configured and connected with VMware vCenter as a storage provider.

Install NetApp VASA Provider

Using the previous instructions for virtual machine creation, build a VASA Provider virtual machine with 2GB RAM, two CPUs, and one virtual network interface in the <MGMT VLAN ID> VLAN. Because of TCP port conflicts, the VASA Provider should be installed on a separate virtual machine and not installed on either the vCenter or VSC/OC VMs. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain. Install all Windows updates on the VM, but do not install Internet Explorer 9. Keep Internet Explorer 8 on the VM.

The following steps describe installing the VASA Provider virtual machine:

1. Download the NetApp FAS/V-Series VASA Provider 1.0 from NetApp Support.
2. Double-click the file downloaded in step 1 (for example, **netappvp-1-0-winx64.exe**).
3. On the installation wizard landing page, select **Next** at the bottom of the screen to proceed with the software installation.
4. Choose the location where the VASA Provider will be installed, then click **Next**.
5. Click **Install**.
6. Make sure the **Launch VASA Configuration** checkbox is selected and click **Finish**.
7. In the NetApp FAS/V-Series VASA Provider 1.0 window, fill in the vCenter administrator user name and password, then click **Save**.

<https://localhost:8143/Register.html>


8. Click **Add**. Fill in the management network IP address or host name for <Controller A>. Fill in the root user name and password and click **OK**.



Note

Repeat for Controller B.

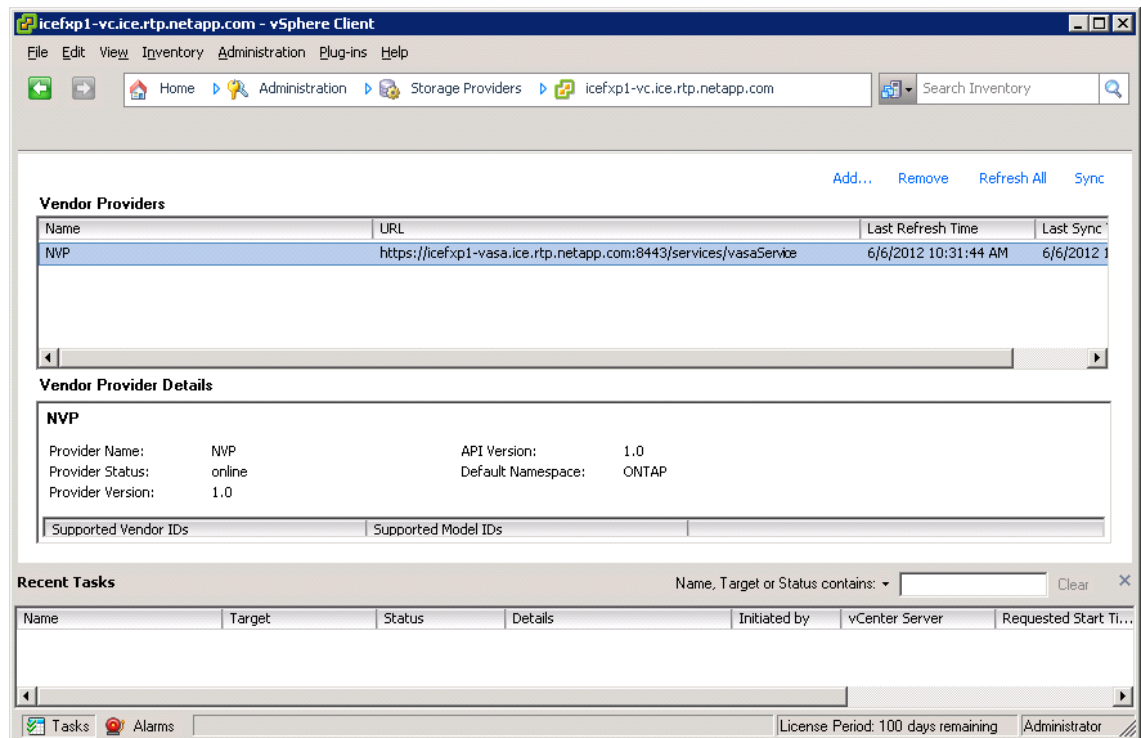
9. In the VMware vCenter Server information section, enter the host name or IP address, port, administrator user name, and password and click **Register Provider** to complete the registration. Click **OK** to acknowledge VASA Provider registration.




10. Click **OK** to close the VASA Provider window.

Verify VASA Provider in vCenter

1. Log in to vCenter using the vSphere client.
2. Click the **Home** tab at the upper left of the window.
3. In the Administration section, click **Storage Providers**.
4. Click **Refresh All**. The NetApp VASA Provider (NVP) should now appear as a vendor provider.



5. Click the **Home** tab in the upper-left portion of the window.
6. In the **Administration** section, click **Datastores and Datastore Clusters**.
7. Expand the **vCenter** and the data center. Click a **datastore**.
8. Click the **Summary** tab. Verify that a **System Storage Capability** appears under **Storage Capabilities**.



9. Repeat for all datastores.

Appendix

Sample Nexus A Config

```
!Command: show running-config
!Time: Tue Jul 10 12:35:53 2012
```

```
version 5.1(3)N2(1)
feature fcoe
hostname ice5548-1
feature npiv
feature fport-channel-trunk
no feature telnet
no feature http-server
cfs eth distribute
feature lacp
feature vpc
feature lldp
```

```

username admin password 5 $1$2BqlI4bJ$ne27Q1AVBEglur9H.BG171 role
network-admin
ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
policy-map type network-qos jumbo
    class type network-qos class-fcoe
        pause no-drop
        mtu 2158
    class type network-qos class-default
        mtu 9000
        multicast-optimize
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
    service-policy type network-qos jumbo
slot 1
    port 31-32 type fc
snmp-server user admin network-admin auth md5
0x2e8af112d36e9af1466f4e4db0ce36a3 priv 0x2e8af112d36e9af1466f4e4db0ce36a3
localizedkey
ntp server 192.168.175.4 use-vrf management
vrf context management
    ip route 0.0.0.0/0 192.168.175.1
vlan 1
vlan 2
    name Native-VLAN
vlan 101
    fcoe vsan 101
    name FCoE_Fabric_A
vlan 3170

```



```

    name NFS-VLAN
vlan 3173
    name vMotion-VLAN
vlan 3174
    name VM-Traffic-VLAN
vlan 3175
    name MGMT-VLAN
vlan 3176
    name Pkt-Ctrl-VLAN
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vpc domain 23
    role priority 10
    peer-keepalive destination 192.168.175.70 source 192.168.175.69
vsan database
    vsan 101 name "Fabric_A"
device-alias database
    device-alias name ice3270-1a_2a pwnn 50:0a:09:81:8d:7d:92:bc
    device-alias name ice3270-1b_2a pwnn 50:0a:09:81:9d:7d:92:bc
    device-alias name VM-Host-Infra-01_A pwnn 20:00:00:25:b5:01:0a:1f
    device-alias name VM-Host-Infra-02_A pwnn 20:00:00:25:b5:01:0a:1e

device-alias commit

fcdomain fcid database
    vsan 101 wwn 50:0a:09:81:8d:7d:92:bc fcid 0x550000 dynamic
!
    [ice3270-1a_2a]
    vsan 101 wwn 50:0a:09:81:9d:7d:92:bc fcid 0x550001 dynamic
!
    [ice3270-1b_2a]
    vsan 101 wwn 24:01:54:7f:ee:23:52:40 fcid 0x550002 dynamic
    vsan 101 wwn 20:00:00:25:b5:01:0a:1f fcid 0x550003 dynamic
!
    [VM-Host-Infra-01_A]
    vsan 101 wwn 20:00:00:25:b5:01:0a:1e fcid 0x550004 dynamic
!
    [VM-Host-Infra-02_A]

interface san-port-channel 1
    channel mode active

interface port-channel10
    description vPC peer-link
    switchport mode trunk

```

```
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
spanning-tree port type network
vpc peer-link

interface port-channel11
description ice3270-1a
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 101,3170
spanning-tree port type edge trunk
vpc 11

interface port-channel12
description ice3270-2b
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 101,3170
spanning-tree port type edge trunk
vpc 12

interface port-channel13
description iceucsm-2a
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
spanning-tree port type edge trunk
vpc 13

interface port-channel14
description iceucsm-2b
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
spanning-tree port type edge trunk
vpc 14

interface port-channel20
description icecore uplink
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3175
spanning-tree port type network
```

```
vpc 20

interface vfc11
  bind interface port-channel11
  switchport description ice3270-1a:2a
  no shutdown

interface vfc12
  bind interface port-channel12
  switchport description ice3270-1b:2a
  no shutdown
vsan database
  vsan 101 interface vfc11
  vsan 101 interface vfc12
  vsan 101 interface san-port-channel 1

interface fc1/31
  switchport description iceucsm-2a:fc1/31
  channel-group 1 force
  no shutdown

interface fc1/32
  switchport description iceucsm-2a:fc1/32
  channel-group 1 force
  no shutdown

interface Ethernet1/1
  description ice3270-1a:e2a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3170
  channel-group 11 mode active

interface Ethernet1/2
  description ice3270-1b:e2a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3170
  channel-group 12 mode active

interface Ethernet1/3
  description iceucsm-2a:Eth1/19
  switchport mode trunk
```

```
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 13 mode active

interface Ethernet1/4
description iceucsm-2b:Eth1/19
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 14 mode active

interface Ethernet1/5
description ice5548-2:Eth1/5
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 10 mode active

interface Ethernet1/6
description ice5548-2:Eth1/6
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 10 mode active

interface Ethernet1/7
description ice1010-1:Eth1
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000

interface Ethernet1/8
description ice1010-2:Eth1
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000

interface Ethernet1/9

interface Ethernet1/10
```

```
interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20
    description icecore:Eth1/21
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3175
    channel-group 20 mode active

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29
```

```

interface Ethernet1/30

interface mgmt0
    ip address 192.168.175.69/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
interface fc1/31
interface fc1/32
interface fc1/31
interface fc1/32
!Full Zone Database Section for vsan 101
zone name VM-Host-Infra-01_A vsan 101
    member pwnn 20:00:00:25:b5:01:0a:1f
!
    [VM-Host-Infra-01_A]
    member pwnn 50:0a:09:81:8d:7d:92:bc
!
    [ice3270-1a_2a]
    member pwnn 50:0a:09:81:9d:7d:92:bc
!
    [ice3270-1b_2a]

zone name VM-Host-Infra-02_A vsan 101
    member pwnn 20:00:00:25:b5:01:0a:1e
!
    [VM-Host-Infra-02_A]
    member pwnn 50:0a:09:81:8d:7d:92:bc
!
    [ice3270-1a_2a]
    member pwnn 50:0a:09:81:9d:7d:92:bc
!
    [ice3270-1b_2a]

zoneset name flexpod vsan 101
    member VM-Host-Infra-01_A
    member VM-Host-Infra-02_A

zoneset activate name flexpod vsan 101

```

Sample Nexus B Config

```

!Command: show running-config
!Time: Tue Jul 10 12:40:36 2012

version 5.1(3)N2(1)

```

```

feature fcoe
hostname ice5548-2
feature npiv
feature fport-channel-trunk
no feature telnet
no feature http-server
cfs eth distribute
feature lacp
feature vpc
feature lldp
username admin password 5 $1$6TphYWHr$u7wmo7lDYKdeyhT5PS7Y01 role
network-admin
ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
policy-map type network-qos jumbo
    class type network-qos class-fcoe
        pause no-drop
        mtu 2158
    class type network-qos class-default
        mtu 9000
        multicast-optimize
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
    service-policy type network-qos jumbo
slot 1
    port 31-32 type fc
snmp-server user admin network-admin auth md5
0xe481d1d2fee4aaa498237df1852270e8 priv 0xe481d1d2fee4aaa498237df1852270e8
localizedkey

```

```

ntp server 192.168.175.4 use-vrf management
vrf context management
    ip route 0.0.0.0/0 192.168.175.1
vlan 1
vlan 2
    name Native-VLAN
vlan 102
    fcoe vsan 102
    name FCoE_Fabric_B
vlan 3170
    name NFS-VLAN
vlan 3173
    name vMotion-VLAN
vlan 3174
    name VM-Traffic-VLAN
vlan 3175
    name MGMT-VLAN
vlan 3176
    name Pkt-Ctrl-VLAN
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vpc domain 23
    role priority 20
    peer-keepalive destination 192.168.175.69 source 192.168.175.70
vsan database
    vsan 102 name "Fabric_B"
device-alias database
    device-alias name ice3270-1a_2b pwwn 50:0a:09:82:8d:7d:92:bc
    device-alias name ice3270-1b_2b pwwn 50:0a:09:82:9d:7d:92:bc
    device-alias name VM-Host-Infra-01_B pwwn 20:00:00:25:b5:01:0b:0f
    device-alias name VM-Host-Infra-02_B pwwn 20:00:00:25:b5:01:0b:1f

device-alias commit

fcdomain fcid database
    vsan 102 wwn 50:0a:09:82:8d:7d:92:bc fcid 0x3f0000 dynamic
!
    [ice3270-1a_2b]
    vsan 102 wwn 50:0a:09:82:9d:7d:92:bc fcid 0x3f0001 dynamic
!
    [ice3270-1b_2b]
    vsan 102 wwn 24:02:54:7f:ee:23:8b:00 fcid 0x3f0002 dynamic
    vsan 102 wwn 20:00:00:25:b5:01:0b:0f fcid 0x3f0003 dynamic
!
    [VM-Host-Infra-01_B]

```



```

vsan 102 wwn 20:00:00:25:b5:01:0b:1f fcid 0x3f0004 dynamic
!
[VM-Host-Infra-02_B]

interface san-port-channel 2
    channel mode active

interface port-channel10
    description vPC peer-link
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3170,3173-3176
    spanning-tree port type network
    vpc peer-link

interface port-channel11
    description ice3270-1a
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 102,3170
    spanning-tree port type edge trunk
    vpc 11

interface port-channel12
    description ice3270-1b
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 102,3170
    spanning-tree port type edge trunk
    vpc 12

interface port-channel13
    description iceucsm-2a
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3170,3173-3176
    spanning-tree port type edge trunk
    vpc 13

interface port-channel14
    description iceucsm-2b
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3170,3173-3176

```

```
spanning-tree port type edge trunk
vpc 14

interface port-channel20
  description icecore uplink
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3175
  spanning-tree port type network
  vpc 20

interface vfc11
  bind interface port-channel11
  switchport description ice3270-1a:2b
  no shutdown

interface vfc12
  bind interface port-channel12
  switchport description ice3270-1b:2b
  no shutdown
vsan database
  vsan 102 interface vfc11
  vsan 102 interface vfc12
  vsan 102 interface san-port-channel 2

interface fc1/31
  switchport description iceucsm-2b:fc1/31
  channel-group 2 force
  no shutdown

interface fc1/32
  switchport description iceucsm-2b:fc1/32
  channel-group 2 force
  no shutdown

interface Ethernet1/1
  description ice3270-1a:e2b
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102,3170
  channel-group 11 mode active

interface Ethernet1/2
```

```

description ice3270-1b:e2b
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 102,3170
channel-group 12 mode active

interface Ethernet1/3
description iceucsm-2a:Eth1/20
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 13 mode active

interface Ethernet1/4
description iceucsm-2b:Eth1/20
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 14 mode active

interface Ethernet1/5
description ice5548-1:Eth1/5
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 10 mode active

interface Ethernet1/6
description ice5548-1:Eth1/6
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
channel-group 10 mode active

interface Ethernet1/7
description ice1010-1:Eth2
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000

interface Ethernet1/8
description ice1010-2:Eth2

```

```
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20
description icecore:Eth1/22
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3175
channel-group 20 mode active

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
```

```

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface mgmt0
    ip address 192.168.175.70/24
    line console
    line vty
    boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
    boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
    interface fc1/31
    interface fc1/32
    interface fc1/31
    interface fc1/32
    !Full Zone Database Section for vsan 102
    zone name VM-Host-Infra-01_B vsan 102
        member pwnn 20:00:00:25:b5:01:0b:0f
    !
        [VM-Host-Infra-01_B]
        member pwnn 50:0a:09:82:8d:7d:92:bc
    !
        [ice3270-1a_2b]
        member pwnn 50:0a:09:82:9d:7d:92:bc
    !
        [ice3270-1b_2b]

    zone name VM-Host-Infra-02_B vsan 102
        member pwnn 20:00:00:25:b5:01:0b:1f
    !
        [VM-Host-Infra-02_B]
        member pwnn 50:0a:09:82:8d:7d:92:bc
    !
        [ice3270-1a_2b]
        member pwnn 50:0a:09:82:9d:7d:92:bc
    !
        [ice3270-1b_2b]

    zoneset name flexpod vsan 102
        member VM-Host-Infra-01_B
        member VM-Host-Infra-02_B

    zoneset activate name flexpod vsan 102

```