



FlexPod Data Center with VMware vSphere 5.1 and Cisco Nexus 7000

Deployment Guide for FlexPod with VMware vSphere 5.1, Cisco Nexus 7000, and NetApp Clustered Data ONTAP 8.1


Last Updated: November 22, 2013



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors

Francis Guillier, Technical Marketing Engineer, Unified Fabric Switching Services Product Group, Cisco Systems

Francis Guillier is a Technical Marketing Engineer for the Nexus 7000 Unified Fabric Switching products focusing on IP based storage solution and FabricPath designs. Francis is currently focused on multiple DC LAN and SAN switching technology like vPC, FEX, FabricPath and FCOE. Francis is deeply involved in design architecture with customers from all segments ranging from commercial to enterprise including service providers. He supports complex network validation at CISCO CPOC and enjoys interacting with customers.

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

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 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.

Derek Huckaby, Technical Marketing Engineer, Unified Fabric Switching Services Product Group, Cisco Systems

Derek Huckaby is a Technical Marketing Engineer for the Nexus 7000 Unified Fabric Switching products focusing on Nexus 7000 integration into FlexPod designs and Nexus 7000 services. Prior to joining the Nexus 7000 Product Marketing team, Derek led the team of Technical Marketing Engineers for the Data Center Application Services BU within Cisco. He began his work in network services at Cisco over 13 years ago specializing in application delivery and SSL termination solutions.

John Kennedy, Technical Leader, Server Access Virtualization Business Unit, Cisco Systems

John Kennedy is focusing 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 Inc. for eight and a half years as a Senior Systems Engineer supporting channel partners outside the US and serving on the HP Alliance team. He is a VMware Certified Professional on every version of VMware's ESX / ESXi, vCenter, and Virtual Infrastructure including vSphere 5. He has presented at various industry conferences in over 20 countries.

Chris O'Brien, Technical Marketing Manager, Server Access Virtualization Business Unit, Cisco Systems

Chris O'Brien is currently focused on developing infrastructure best practices and solutions that are designed, tested, and documented to facilitate and improve customer deployments. Previously, O'Brien was an application developer and has worked in the IT industry for more than 15 years.

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

Chris Reno is a Reference Architect in the NetApp Infrastructure and Cloud Enablement group and is focused on creating, validating, supporting, and evangelizing solutions based on NetApp products. Before being employed in his current role, he worked with NetApp product engineers designing and developing innovative ways to perform QA for NetApp products, including enablement of a large grid infrastructure using physical and virtualized compute resources. In these roles, Chris gained expertise in stateless computing, netboot architectures, and virtualization.

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

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 Bachelors of Science degree in Computer Networking and her Master's of Science in Information Security from East Carolina University.

About Cisco Validated Design (CVD) Program

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

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

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

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

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

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

© 2013 Cisco Systems, Inc. All rights reserved



VMware vSphere 5.1 on FlexPod with the Nexus 7000 Deployment Guide

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®, NetApp®, and VMware® virtualization that uses IP-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 with NetApp clustered Data ONTAP®.

Architecture

The FlexPod architecture is highly modular or "podlike." Although each customer's FlexPod unit varies in its exact configuration, after a FlexPod unit is deployed, 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 both virtualized and nonvirtualized solutions. VMware vSphere® built on FlexPod includes NetApp storage, NetApp Data ONTAP, Cisco networking, the Cisco Unified Computing System™ (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.



Corporate Headquarters:

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

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

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 IP-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 built on FlexPod components and the network connections for a configuration with IP-based storage. This design uses Cisco Nexus® 7000, 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 PortChannels (vPCs). This infrastructure is deployed to provide iSCSI-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.

With respect to the NetApp storage controller, it is a fundamental design decision to leverage clustered Data ONTAP or 7-Mode as these cannot be run simultaneously on the same HA pair and the choice will influence hardware requirements, the logical construction of the FlexPod stack and ultimately the operational practices of the enterprise. Organizations with the following requirements should consider adopting clustered Data ONTAP.

- Large to midsize enterprises that are seeking scalable, shared IT solutions for nondisruptive operations
- New installations
- Existing clustered Data ONTAP 8.x and Data ONTAP GX organizations that are looking to upgrade
- Organizations deploying an enterprise content repository

Organizations that have the following characteristics or needs might want to use the 7-Mode design:

- Existing Data ONTAP 7G and Data ONTAP 8.x 7-Mode customers who are looking to upgrade
- Midsize enterprises: customers who are primarily interested in the FAS2000 series
- Customers who absolutely require SnapVault®, synchronous SnapMirror®, MetroCluster™, SnapLock® software, IPv6, or Data ONTAP Edge

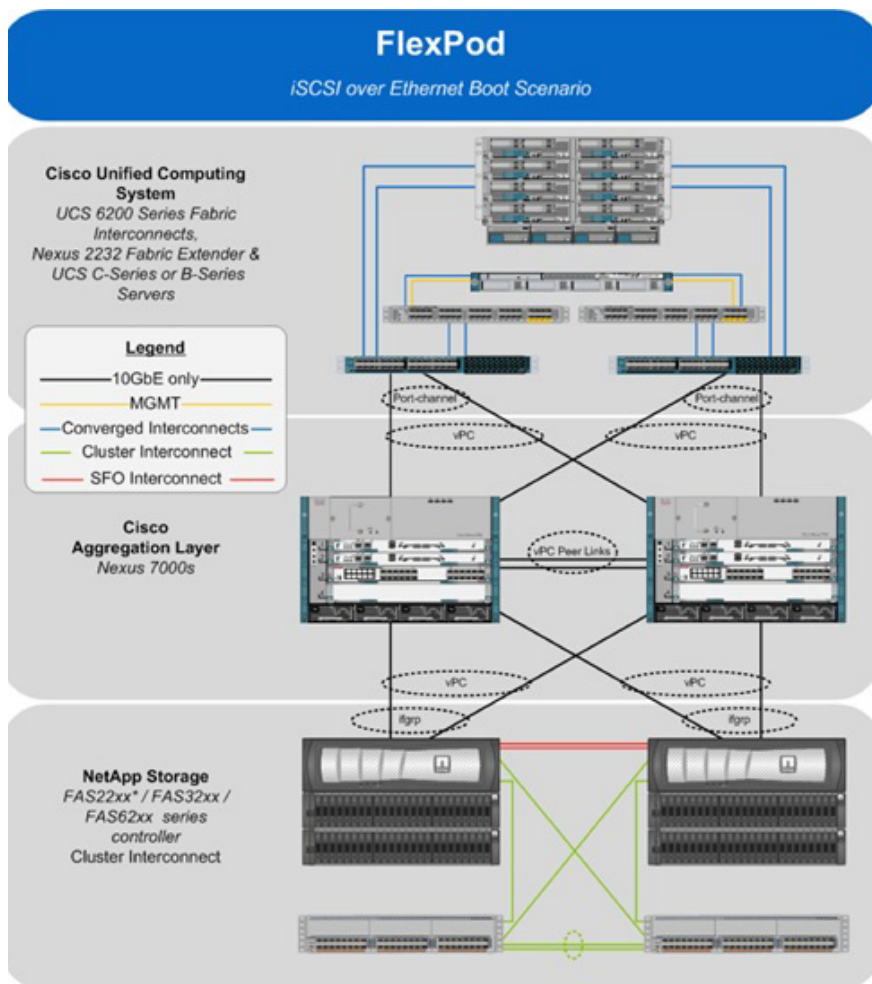

Note

This document provides detailed deployment steps for both clustered Data ONTAP solutions and those operating in 7-Mode. Steps for clustered Data ONTAP are found in the main sections of the document while those specific to controllers operating in 7-Mode are found in the appendix.


Note

Deployment of a FAS22xx would require modification to the data and cluster interconnect.

Figure 1 VMware vSphere built on FlexPod components



The reference configuration includes:

- Two Cisco Nexus 7000 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 FAS3250-AE (HA pair) running clustered Data ONTAP

Storage is provided by a NetApp FAS3250-AE (HA configuration in two chassis) operating in either clustered Data ONTAP or 7-Mode. 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.


Note

Although this document leverages the NetApp Unified Target Adapter (UTA) for storage connectivity, this adapter is not required. The UTA provides the greatest flexibility when migrating to an end-to-end FCoE design, which is why it is used in this design. However, a standard 10GbE can be used for IP-based storage designs.

FlexPod Benefits

One of the founding design principles of the FlexPod architecture is flexibility. Previous FlexPod architectures have highlighted FCoE-, FC-, or 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 7000 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. This VMware vSphere 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 shared storage using storage protocols such as iSCSI, CIFS, and NFS. For a detailed study of several practical solutions deployed on FlexPod, refer to [NetApp Technical Report 3884: FlexPod Solutions Guide](#).

Benefits of 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 TCO. Building on Cisco's strength in enterprise networking, Cisco 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, interprocess communication, and virtual machine traffic with security isolation, visibility, and control equivalent to those provided by 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 are 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 and 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 results are increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender (FEX) 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 lead to simplicity rather than complexity. The results are 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 C220-M3, B200-M3, B230-M2, and 6248UP 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, which enable iSCSI 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 7000

The modular Cisco Nexus 7000 Series offers a comprehensive one-platform solution for the data center core network. It also offers aggregation, high density, and end-of-row and top-of-rack server connectivity. For campus core deployments, it provides a scalable, highly resilient, high-performance solution.

The Cisco Nexus 7000 Series platform runs on Cisco NX-OS Software. It was specifically designed for the most mission-critical deployments in the data center and on campus.

The Cisco Nexus 7000 Series was designed around four principles:

- Infrastructure scalability
 - Design that provides scalability to more than 15Tbps for ongoing investment protection
 - Support for consolidated networks with virtual port channel innovations to scale beyond 1500 ports
 - Multicore, multithreaded OS to optimize CPU resources and offload tasks to processors distributed across the modules
 - Cisco Trusted Security (Cisco TrustSec®) for scalable security with link-layer encryption, security group access control lists, and role-based access control

- Flexible NetFlow to optimize the network infrastructure, reducing operating costs and improving capacity planning capabilities
- Operational continuity
 - Lossless nondisruptive upgrades for zero-service downtime through no single point of failure in the system hardware and a modular operating system
 - Connectivity management processor (CMP) for integrated out-of-band management access
 - Innovative stateful process restart for nondisruptive operations in event of process termination
 - Comprehensive Extensible Markup Language (XML) API for total platform control
- Transport flexibility
 - Foundation for unified fabrics with Cisco DCE unified I/O and FCoE
 - Virtualized control plane and data plane forwarding for optimized performance
 - Virtual device contexts (VDCs) to maximize software and hardware resource utilization while providing strong security and software fault isolation
 - Built to currently support high-density GbE and 10GbE and the emerging 40Gbps and 100Gbps Ethernet standards
- Data center switching features
 - In-Service Software Upgrade (ISSU) enables hitless upgrades with zero packet loss
 - NetFlow provides visibility and flexible monitoring and control over the network
 - Multihop FCoE provides director-class FCoE on a modular platform to offer rich LAN and SAN services
 - OTV and LISP enable seamless workload mobility across geographically separated data centers
 - MPLS L3 VPN service supports multi-tenant segmentation within and between data centers
 - Virtual Device Contexts consolidate data center switching hardware through virtualization
 - Fabric Path/TRILL allows scalable data center networks to be built without the tree protocol

Benefits of In-Service Software Upgrade

The modular Cisco Nexus 7000 Series is designed for highly scalable networks and utilizes Cisco NX-OS, which is a data center-class operating system to make sure of continuous availability for mission-critical data center environments. The Cisco Nexus 7000 provides key high-availability features, which include:

- Hitless ISSU
- Layer 2 ISSU
- Layer 3 ISSU
- Stateful supervisor switchover
- Stateful process restart
- NSF awareness

In-Service Software Upgrade (ISSU) is a key component of Cisco NX-OS that enables nondisruptive software upgrades and downgrades. By leveraging ISSU, users can migrate from one Cisco NX-OS version to another without removing the Cisco Nexus 7000 from the production environment. In the majority of cases, zero packet loss is observed.

One of the most significant challenges to meeting and exceeding 99.999% uptime is the ability to minimize planned outages. In many highly redundant environments, complex and tedious preparation must be performed by network operators to minimize downtime during software upgrades. By utilizing ISSU for nondisruptive upgrades, revenues can be recaptured while simultaneously reducing the complexity required for maintaining mission-critical data center environments. ISSU enables faster adoption of new data center features, while allowing users to meet stringent network uptime SLAs.

Recommended support documents include:

- Cisco Nexus 7000 Family of switches: <http://www.cisco.com/en/US/products/ps9402/index.html>

Benefits of NetApp FAS Family of Storage Controllers

NetApp solutions are user friendly, easy to manage, and quick to deploy and offer increased availability while consuming fewer IT resources. This means that they dramatically lower the lifetime total cost of ownership. Where others manage complexity, NetApp eliminates it. A NetApp solution includes hardware in the form of controllers and disk storage and the NetApp Data ONTAP operating system, the #1 branded storage OS¹.

NetApp offers the NetApp Unified Storage Architecture. The term "unified" refers to a family of storage systems that simultaneously support storage area network (SAN), network-attached storage (NAS), and iSCSI across many operating environments such as VMware, Windows®, and UNIX®. This single architecture provides access to data by using industry-standard protocols, including NFS, CIFS, iSCSI, FCP, SCSI, FTP, and HTTP. Connectivity options include standard Ethernet (10/100/1000 or 10GbE) and Fibre Channel (1, 2, 4, or 8Gb/sec) as well as Fibre Channel over Ethernet (FCoE). In addition, all systems can be configured with high-performance solid state drives (SSDs) or serial ATA (SAS) disks for primary storage applications, low-cost SATA disks for secondary applications (backup, archive, and so on), or a mix of the different disk types.

A storage system running Data ONTAP has a main unit, also known as the controller or storage engine, which is the hardware device that receives and sends data. This unit detects and gathers information about the hardware configuration, the storage system components, the operational status, hardware failures, and other error conditions.

A storage system uses storage on disk shelves. The disk shelves are the containers or device carriers that hold disks and associated hardware such as power supplies, connectivity interfaces, and cabling.

If storage requirements change over time, NetApp storage offers the flexibility to change quickly as needed without expensive and disruptive "forklift" upgrades. For example, a LUN can be changed from FC access to iSCSI access without moving or copying the data. Only a simple dismount of the FC LUN and a mount of the same LUN using iSCSI would be required. In addition, a single copy of data can be shared between Windows and UNIX systems while allowing each environment to access the data through native protocols and applications. If a system was originally purchased with all SATA disks for backup applications, high-performance SAS disks could be added to support primary storage applications such as Oracle®, Microsoft Exchange Server, or ClearCase.

NetApp storage solutions provide redundancy and fault tolerance through clustered storage controllers, hot-swappable redundant components (such as cooling fans, power supplies, disk drives, and shelves), and multiple network interfaces. This highly available and flexible architecture enables customers to manage all data under one common infrastructure while achieving mission requirements. The NetApp Unified Storage Architecture allows data storage with higher availability and performance, easier dynamic expansion, and easier management than any other solution.

1. Source: IDC Worldwide Quarterly Disk Storage Systems Tracker Q4 2012, March 2013 (Open Networked Disk Storage Systems revenue)

The storage efficiency built into Data ONTAP provides substantial space savings, allowing more data to be stored at lower cost. Data protection provides replication services, making sure that valuable data is backed up and recoverable. The following features provide storage efficiency and data protection:

- Thin provisioning. Volumes are created using "virtual" sizing. They appear to be provisioned at their full capacity, but are actually created much smaller and use additional space only when it is actually needed. Extra unused storage is shared across all volumes, and the volumes can grow and shrink on demand.
- NetApp Snapshot™. Automatically scheduled point-in-time copies that write only changed blocks, with no performance penalty. The Snapshot copies consume minimal storage space, since only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- FlexClone® volumes. Near-zero space, instant "virtual" copies of datasets. The clones are writable, but only changes to the original are stored, so they provide rapid, space-efficient creation of additional data copies ideally suited for dev/test environments.
- Deduplication. Removes redundant data blocks in primary and secondary storage with flexible policies to determine when the deduplication process is run.
- Compression. Compresses data blocks. Compression can be run whether or not deduplication is enabled and can provide additional space savings, whether run alone or together with deduplication.
- SnapMirror. Allows volumes to be asynchronously replicated either within the cluster or to another cluster.

For more information, see:

- 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 Subsystem Resiliency Guide: www.netapp.com/us/system/pdf-reader.aspx?m=tr-3437.pdf&cc=us
- NetApp TR-3749: NetApp and VMware vSphere Storage Best Practices: www.netapp.com/us/system/pdf-reader.aspx?m=tr-3749.pdf&cc=us
- NetApp TR-3884: FlexPod Solutions Guide: www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61208-16&m=tr-3884.pdf
- NetApp TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010: www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61277-16&m=tr-3824.pdf
- NetApp Data ONTAP 8: <http://www.netapp.com/us/products/platform-os/data-ontap-8/index.aspx>

Clustered Data ONTAP

With the release of clustered Data ONTAP 8.1, NetApp introduces enterprise-ready, unified scale-out storage. Developed from a solid foundation of proven Data ONTAP technology and innovation, clustered Data ONTAP is the basis for large virtualized shared storage infrastructures that are architected for nondisruptive operations over the system lifetime. Controller nodes are deployed in HA pairs with these HA pairs participating in a single storage domain or cluster.

Scale-out is a way to respond to growth in a storage environment. All storage controllers have physical limits to their expandability: number of CPUs, memory slots, and space for disk shelves that dictate the maximum capacity and controller performance. If more storage or performance capacity is needed, it might be possible to add CPUs and memory or install additional disk shelves, but ultimately the controller becomes completely populated, with no further expansion possible. At this stage, the only option is to acquire another controller. One way to do this is to "scale up": that is, to add additional

controllers in such a way that each is a completely independent management entity that does not provide any shared storage resources. If the original controller is to be completely replaced by the newer and larger controller, data migration is required to transfer the data from the old to the new. This is time-consuming and potentially disruptive and likely requires configuration changes on all of the attached host systems.

If the newer controller can coexist with the original controller, there are now two storage controllers to be individually managed, and there are no native tools to balance or reassign workloads across them. The situation becomes worse as the number of controllers increases. If the scale-up approach is used, the operational burden increases consistently as the environment grows, and the end result is a very unbalanced and difficult-to-manage environment. Technology refresh cycles require substantial planning in advance, lengthy outages, and configuration changes, which introduce risk into the system.

By contrast, using scale-out means that as the storage environment grows, additional controllers are added seamlessly to the resource pool residing on a shared storage infrastructure. Host and client connections as well as datastores can move seamlessly and nondisruptively anywhere in the resource pool, so that existing workloads can be easily balanced over the available resources, and new workloads can be easily deployed. Technology refreshes (replacing disk shelves, adding or completely replacing storage controllers) are accomplished while the environment remains online and serving data.

Although scale-out products have been available for some time, these were typically subject to one or more of the following shortcomings:

- Limited protocol support. NAS only
- Limited hardware support. Supported only a particular type of storage controller or a very limited set
- Little or no storage efficiency. Thin provisioning, deduplication, compression
- Little or no data replication capability

Therefore, while these products are well positioned for certain specialized workloads, they are less flexible, capable, and robust for broad deployment throughout the enterprise.

Data ONTAP is the first product to offer a complete scale-out solution, and it offers an adaptable, always-available storage infrastructure for today's highly virtualized environments.


Note

The use of clustered Data ONTAP is addressed in the body of this document.

Data ONTAP Operating in 7-Mode

As previously mentioned, customers have a choice to deploy their NetApp storage environment operating in 7-Mode or leverage clustered Data ONTAP. Data ONTAP operating in 7-Mode provides customers a broad suite of application integrations, storage efficiencies, and legacy of customer satisfaction.

As well-known and trusted as Data ONTAP operating in 7-Mode is, technology companies must always have an eye toward new innovations. For this reason NetApp has continually invested in clustered Data ONTAP, which truly changes the conversation of storage from a cost-center discussion to one in which storage can add value to the company.

It is acknowledged that clustered Data ONTAP is the future for NetApp. However, customers can choose to join NetApp on this journey at their own pace. Data ONTAP operating in 7-Mode is deployed on a given HA pair of controllers that is a discrete pair from any other storage systems in the environment and managed as such. For this reason, the scalability with clustered Data ONTAP is superior to that of 7-Mode.


Note

The use of Data ONTAP operating in 7-Mode is addressed in the appendix of this document.

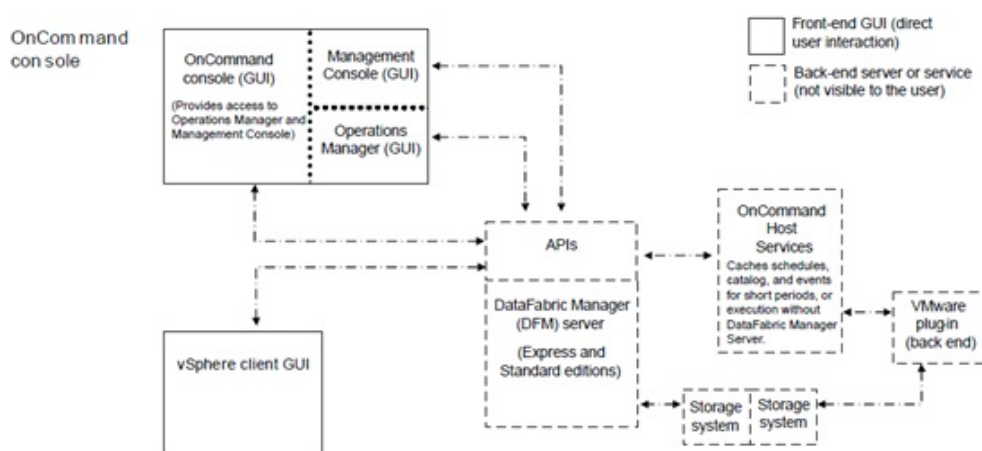
Benefits of NetApp OnCommand Unified Manager Software

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 create 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 2 *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 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 names 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>

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 5.1 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	2.1(1b)	Embedded management
	Cisco UCS C 200 M2	2.1(1b)	Software bundle release
	Cisco UCS C 220 M3	2.1(1b)	Software bundle release
	Cisco UCS B 200 M2	2.1(1b)	Software bundle release
	Cisco UCS B 200 M3	2.1(1b)	Software bundle release
	Cisco enic	2.1.2.38	Ethernet driver for Cisco VIC
Network	Cisco Nexus 7000	6.1(2)	Operating system version
Storage	NetApp FAS3250-A	Clustered Data ONTAP 8.1.2	Operating system version
Software	Cisco UCS hosts	VMware vSphere ESXi 5.1	Operating system version
	Microsoft .NET Framework	3.5.1	Feature enabled within Windows operating system
	Microsoft SQL Server	Microsoft SQL Server 2008 R2 SP1	VM (1 each): SQL Server DB
	VMware vCenter	5.1	VM (1 each): VMware vCenter
	NetApp OnCommand	5.1	VM (1 each): OnCommand
	NetApp Virtual Storage Console (VSC)	4.1	Plug-in within VMware vCenter
	Cisco Nexus 1110-x	4.2.1.SP1.5.1a	Virtual services appliance
	Cisco Nexus 1000v	4.2.1.SV2.1.1a	Virtual services blade within the 1110-x
	NetApp NFS Plug-in for VMware vStorage APIs for Array Integration (VAAI)	1.0-018	Plug-in within VMware vCenter

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the 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 network port vlan create command within clustered Data ONTAP:

Usage:

```
network port vlan create ?
  [-node] <nodename>           Node
  { [-vlan-name] {<netport>|<ifgrp>} VLAN Name
  | -port {<netport>|<ifgrp>}    Associated Network Port
  [-vlan-id] <integer> }        Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <node01> -vlan-name i0a-<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 2 lists the virtual storage area networks (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.

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	3171
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for NFS traffic	3170
iSCSI-A	VLAN for iSCSI traffic for fabric A	911
iSCSI-B	VLAN for iSCSI traffic for fabric B	912
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 (Cisco Nexus 1000v)	3176

Table 3 **VMware virtual machines (VMs) created**

Virtual Machine Description	Host Name
vCenter SQL Server database	
vCenter Server	
NetApp Virtual Storage Console (VSC) and NetApp OnCommand core	

Table 4 **Configuration variables**

Variable	Description	Customer Implementation Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01	
<<var_node01_mgmt_mask>>	Out-of-band management network netmask	
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_url_boot_software>>	Data ONTAP 8.1.2 URL; format: http://	
<<var_#_of_disks>>	Number of disks to assign to each storage controller	
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02	
<<var_node02_mgmt_mask>>	Out-of-band management network netmask	
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_clustername>>	Storage cluster host name	

<<var_cluster_base_license_key>>	Cluster base license key	
<<var_password>>	Global default administrative password	
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster	
<<var_clustermgmt_mask>>	In-band management network netmask	
<<var_clustermgmt_gateway>>	In-band management network default gateway	
<<var_dns_domain_name>>	DNS domain name	
<<var_nameserver_ip>>	DNS server IP(s)	
<<var_node_location>>	Node location string for each node	
<<var_node01>>	Cluster node 01 host name	
<<var_node02>>	Cluster node 02 host name	
<<var_raidsize>>	RAID group size for each node	
<<var_num_disks>>	Number of disks to assign to each storage data aggregate	
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP	
<<var_node01_sp_mask>>	Out-of-band management network netmask	
<<var_node01_sp_gateway>>	Out-of-band management network default gateway	
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP	
<<var_node02_sp_mask>>	Out-of-band management network netmask	
<<var_node02_sp_gateway>>	Out-of-band management network default gateway	
<<var_timezone>>	FlexPod time zone (for example, America/New York)	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_snmp_contact>>	Administrator e-mail address	
<<var_snmp_location>>	Cluster location string	
<<var_oncommand_server_fqdn>>	VSC or OnCommand virtual machine fully qualified domain name (FQDN)	
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name	
<<var_mailhost>>	Mail server host name	
<<var_storage_admin_email>>	Administrator e-mail address	
<<var_security_cert_vserver_common_name>>	Infrastructure Vserver FQDN	
<<var_country_code>>	Two-letter country code	
<<var_state>>	State or province name	
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_security_cert_cluster_common_name>>	Storage cluster FQDN	

<<var_security_cert_node01_common_name>>	Cluster node 01 FQDN	
<<var_security_cert_node02_common_name>>	Cluster node 02 FQDN	
<<var_esxi_host1_nfs_ip>>	NFS VLAN IP address for each VMware ESXi host	
<<var_node01_nfs_lif_ip>>	Cluster node 01 NFS VLAN IP address	
<<var_node01_nfs_lif_mask>>	NFS VLAN netmask	
<<var_node02_nfs_lif_ip>>	Cluster node 02 NFS VLAN IP address	
<<var_node02_nfs_lif_mask>>	NFS VLAN netmask	
<<var_node01_iscsi_A_IP>>	Cluster node 01 iSCSI IP address for fabric A	
<<var_node01_iscsi_A_mask>>	iSCSI netmask for fabric A	
<<var_node01_iscsi_B_IP>>	Cluster node 01 iSCSI IP address for fabric B	
<<var_node01_iscsi_B_mask>>	iSCSI netmask for fabric B	
<<var_node02_iscsi_A_IP>>	Cluster node 02 iSCSI IP address for fabric A	
<<var_node02_iscsi_A_mask>>	iSCSI netmask for fabric A	
<<var_node02_iscsi_B_IP>>	Cluster node 02 iSCSI IP address for fabric B	
<<var_node02_iscsi_B_mask>>	iSCSI netmask for fabric B	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_native_vlan_id>>	Native VLAN ID	
<<var_nfs_vlan_id>>	NFS VLAN ID	
<<var_iscsi_vlan_A_id>>	iSCSI VLAN for fabric A	
<<var_iscsi_vlan_B_id>>	iSCSI VLAN for fabric B	
<<var_pkt-ctrl_vlan_id>>	Cisco Nexus 1000v packet control VLAN ID	
<<var_vmotion_vlan_id>>	VMware vMotion VLAN ID	
<<var_vm-traffic_vlan_id>>	VM traffic VLAN ID	
<<var_nexus_vpc_domain_id>>	Unique Cisco Nexus switch VPC domain ID	
<<var_nexus_1110x-1>>	Cisco Nexus 1110X-1 host name	

<<var_nexus_1110x-2>>	Cisco Nexus 1110X-2 host name	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucs_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	
<<var_ucs_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucs_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	
<<var_cimc_ip>>	Out-of-band management IP for each Cisco Nexus 1110-X CIMC	
<<var_cimc_mask>>	Out-of-band management network netmask	
<<var_cimc_gateway>>	Out-of-band management network default gateway	
<<var_1110x_domain_id>>	Unique Cisco Nexus 1110-X domain ID	
<<var_1110x_vsa>>	Virtual storage appliance (VSA) host name	
<<var_1110x_vsa_ip>>	In-band VSA management IP address	
<<var_1110x_vsa_mask>>	In-band management network netmask	
<<var_1110x_vsa_gateway>>	In-band management network default gateway	
<<var_vsm_domain_id>>	Unique Cisco Nexus 1000v virtual supervisor module (VSM) domain ID	
<<var_vsm_mgmt_ip>>	Cisco Nexus 1000v VSM management IP address	
<<var_vsm_mgmt_mask>>	In-band management network netmask	
<<var_vsm_mgmt_gateway>>	In-band management network default gateway	
<<var_vsm_hostname>>	Cisco Nexus 1000v VSM host name	
<<var_vcenter_server_ip>>	vCenter Server IP	
<<var_nodename>>	Name of node	
<<var_node01_rootaggrname>>	Root aggregate name of Node 01	
<<var_clustermgmt_port>>	Port for cluster management	
<<var_global_domain_name>>	Domain name	
<<var_dns_ip>>	IP address of the DNS server	
<<var_vsadmin_password>>	Password for VS admin account	
<<var_vserver_mgmt_ip>>	Management IP address for Vserver	
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver	
<<var_rule_index>>	Rule index number	
<<var_ftp_server>>	IP address for FTP server	
<<var_vm_host_infra_01_iqn_A>>	iSCSI Qualified Name (IQN) of host infra 01 A	
<<var_vm_host_infra_01_iqn_B>>	IQN of host infra 01 B	

<<var_vm_host_infra_02_iqn_A>>	IQN of host infra 02 A	
<<var_vm_host_infra_02_iqn_B>>	IQN of host infra 02 B	
<<var_vmhost_infra01_ip>>	VMware ESXi host 01 in-band management IP	
<<var_vmhost_infra02_ip>>	VMware ESXi host 02 in-band management IP	
<<var_nfs_vlan_id_ip_host-01>>	NFS VLAN IP address for ESXi host 01	
<<var_nfs_vlan_id_mask_host-01>>	NFS VLAN netmask for ESXi host 01	
<<var_vmotion_vlan_id_ip_host-01>>	vMotion VLAN IP address for ESXi host 01	
<<var_vmotion_vlan_id_mask_host-01>>	vMotion VLAN netmask for ESXi host 01	
<<var_nfs_vlan_id_ip_host-02>>	NFS VLAN IP address for ESXi host 02	
<<var_nfs_vlan_id_mask_host-02>>	NFS VLAN netmask for ESXi host 02	
<<var_vmotion_vlan_id_ip_host-02>>	vMotion VLAN IP address for ESXi host 02	
<<var_vmotion_vlan_id_mask_host-02>>	vMotion VLAN netmask for ESXi host 02	

Physical Infrastructure

This section describes the steps to deploy base infrastructure components as well as 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 procedure is outlined as follows:

1. Initial configuration of the NetApp controller.
2. Initial configuration of Cisco UCS.
3. Initial configuration of Cisco Nexus.
4. Creation of necessary VLANs for management, basic functionality, and virtualized infrastructure specific to VMware.
5. Creation of vPCs to provide high availability among devices.
6. Creation of service profile pools: MAC, UUID, server, and so forth.
7. Creation of service profile policies: adapter, boot, and so forth.
8. Creation of two service profile templates from the created pools and policies: one each for fabric A and B.
9. Provisioning of two servers from the created service profiles in preparation for OS installation.
10. Initial configuration of the infrastructure components residing on the NetApp controller.
11. Installation of VMware vSphere 5.1.
12. Installation and configuration of VMware vCenter.
13. Enablement of NetApp Virtual Storage Console (VSC).
14. Configuration of NetApp OnCommand.
15. 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 might vary, 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.

FlexPod Cabling on Clustered Data ONTAP

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 FAS3250-AE running clustered Data ONTAP 8.1. This configuration uses a dual-port 10 GbE adapter, and external SAS disk shelves. The built-in FC ports are not required for this design but are available to support FC protocols in the future. For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).



Note

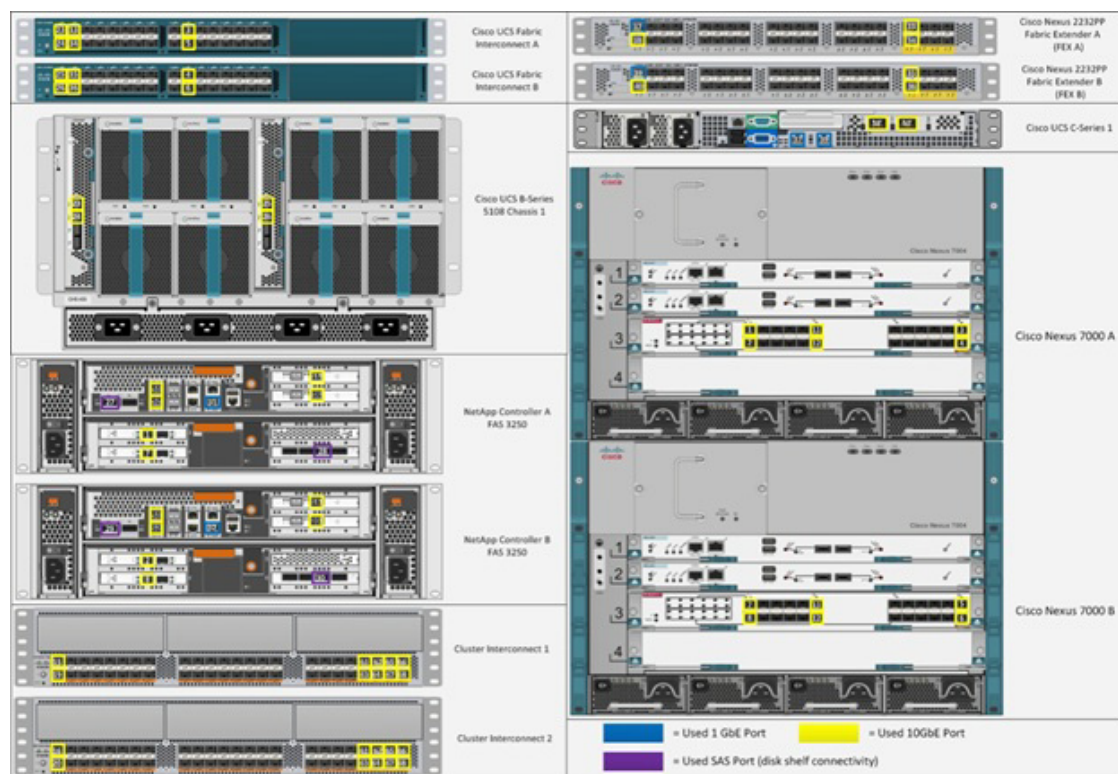
To use a storage controller's built-in FC ports, a switch capable of supporting native Fibre channel is required.

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 FAS3250 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 3 shows a FlexPod cabling diagram. The labels indicate connections to endpoints rather than port numbers on the physical device. For example, 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. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide at https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 3 FlexPod cabling diagram in clustered Data ONTAP

The information provided in Table 5 through Table 15 corresponds to each connection shown in Figure 3.

**Note**

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

Table 5 Cisco Nexus 7000 A cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 7000 A	Eth3/1	10GbE	NetApp controller A	e3a
	Eth3/2	10GbE	NetApp controller B	e3a
	Eth3/11	10GbE	VPC peer link	Eth3/11
	Eth3/12	10GbE	VPC peer link	Eth3/12
	Eth3/23	10GbE	Cisco UCS fabric interconnect A	Eth1/19
	Eth3/24	10GbE	Cisco UCS fabric interconnect B	Eth1/19
	MGMT0	100MbE	100MbE management switch	Any

Table 6 *Cisco Nexus 7000 B cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 7000 B	Eth3/1	10GbE	NetApp controller A	e4a
	Eth3/2	10GbE	NetApp controller B	e4a
	Eth3/11	10GbE	VPC peer link	Eth3/11
	Eth3/12	10GbE	VPC peer link	Eth3/12
	Eth3/23	10GbE	Cisco UCS fabric interconnect A	Eth1/20
	Eth3/24	10GbE	Cisco UCS fabric interconnect B	Eth1/20
	MGMT0	100MbE	100MbE management switch	Any

Table 7 *Cisco Nexus 5596 A cluster interconnect cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5596 A	Eth1/1	10GbE	NetApp controller 1	e1a
	Eth1/2	10GbE	NetApp controller 2	e1a
	Eth1/41	10GbE	Cisco Nexus 5596 B	Eth1/41
	Eth1/42	10GbE	Cisco Nexus 5596 B	Eth1/42
	Eth1/43	10GbE	Cisco Nexus 5596 B	Eth1/43
	Eth1/44	10GbE	Cisco Nexus 5596 B	Eth1/44
	Eth1/45	10GbE	Cisco Nexus 5596 B	Eth1/45
	Eth1/46	10GbE	Cisco Nexus 5596 B	Eth1/46
	Eth1/47	10GbE	Cisco Nexus 5596 B	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 5596 B	Eth1/48
	MGMT0	GbE	GbE management switch	Any

Table 8 *Cisco Nexus 5596 B cluster interconnect cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5596 B	Eth1/1	10GbE	NetApp controller 1	e2a
	Eth1/2	10GbE	NetApp controller 2	e2a
	Eth1/41	10GbE	Cisco Nexus 5596 A	Eth1/41
	Eth1/42	10GbE	Cisco Nexus 5596 A	Eth1/42
	Eth1/43	10GbE	Cisco Nexus 5596 A	Eth1/43
	Eth1/44	10GbE	Cisco Nexus 5596 A	Eth1/44
	Eth1/45	10GbE	Cisco Nexus 5596 A	Eth1/45
	Eth1/46	10GbE	Cisco Nexus 5596 A	Eth1/46
	Eth1/47	10GbE	Cisco Nexus 5596 A	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 5596 A	Eth1/48
	MGMT0	GbE	GbE management switch	Any

**Note**

When the term e0M is used, the physical Ethernet port to which the table is referring is the port indicated by a wrench icon on the rear of the chassis.

Table 9 *NetApp controller 1 cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	100MbE	100MbE management switch	Any
	e0a	GbE	GbE management switch	Any
	e0b	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	c0a	10GbE	NetApp controller 2	c0a
	c0b	10GbE	NetApp controller 2	c0b
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/1
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/1
	e3a	10GbE	Cisco Nexus 7000 A	Eth3/1
	e4a	10GbE	Cisco Nexus 7000 B	Eth3/1

Table 10 *NetApp controller 2 cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	100MbE	100MbE management switch	Any
	e0a	GbE	GbE management switch	Any
	e0b	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	c0a	10GbE	NetApp controller 1	c0a
	c0b	10GbE	NetApp controller 1	c0b
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/2
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/2
	e3a	10GbE	Cisco Nexus 7000 A	Eth3/2
	e4a	10GbE	Cisco Nexus 7000 B	Eth3/2

Table 11 *Cisco UCS fabric interconnect A cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/19	10GbE	Cisco Nexus 7000 A	Eth3/23
	Eth1/20	10GbE	Cisco Nexus 7000 B	Eth3/23
	Eth1/1	10GbE	Cisco UCS chassis fabric extender (FEX) A / Cisco Nexus 2232PP FEX A	Port 2/1

	Eth1/2	10GbE	Cisco UCS chassis FEX A / Cisco Nexus 2232PP FEX A	
	Eth1/3	10GbE	Cisco UCS chassis FEX A / Cisco Nexus 2232PP FEX A	
	Eth1/4	10GbE	Cisco UCS chassis FEX A / Cisco Nexus 2232PP FEX A	
	Eth1/5	10GbE	Cisco UCS chassis FEX A / Cisco Nexus 2232PP FEX A	
	Eth1/6	10GbE	Cisco UCS chassis FEX A / Cisco Nexus 2232PP FEX A	
	Eth1/2	10GbE	Cisco Nexus 2232PP FEX A	Port 2/2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 12 *Cisco UCS fabric interconnect B cabling information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/19	10GbE	Cisco Nexus 7000 A	Eth3/24
	Eth1/20	10GbE	Cisco Nexus 7000 B	Eth3/24
	Eth1/1	10GbE	Cisco UCS chassis fabric extender (FEX) B / Cisco Nexus 2232PP FEX B	
	Eth1/2	10GbE	Cisco UCS chassis FEX B / Cisco Nexus 2232PP FEX B	
	Eth1/3	10GbE	Cisco UCS chassis FEX B / Cisco Nexus 2232PP FEX B	
	Eth1/4	10GbE	Cisco UCS chassis FEX B / Cisco Nexus 2232PP FEX B	
	Eth1/5	10GbE	Cisco UCS chassis FEX B / Cisco Nexus 2232PP FEX B	
	Eth1/6	10GbE	Cisco UCS chassis FEX B / Cisco Nexus 2232PP FEX B	
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1
	L2	GbE	Cisco UCS fabric interconnect A	L2

Table 13 *Cisco Nexu 2232PP FEX A*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX A	Port 1	GbE	Cisco UCS C-Series 1	M1
	Port 2	10GbE	Cisco UCS C-Series 1	Port 0
	Port 3	GbE	Cisco UCS C-Series 2	M1

	Port 4	10GbE	Cisco UCS C-Series 2	Port 0
	Port 5	GbE	Cisco UCS C-Series 3	M1
	Port 6	10GbE	Cisco UCS C-Series 3	Port 0
	Port 7	GbE	Cisco UCS C-Series 4	M1
	Port 8	10GbE	Cisco UCS C-Series 4	Port 0
	Port 2/1	10GbE	Cisco UCS fabric interconnect A	
	Port 2/2	10GbE	Cisco UCS fabric interconnect A	

Table 14 *Cisco Nexus 2232PPFEX B*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX B	Port 1	GbE	Cisco UCS C-Series 1	M2
	Port 2	10GbE	Cisco UCS C-Series 1	Port 1
	Port 3	GbE	Cisco UCS C-Series 2	M2
	Port 4	10GbE	Cisco UCS C-Series 2	Port 1
	Port 5	GbE	Cisco UCS C-Series 3	M2
	Port 6	10GbE	Cisco UCS C-Series 3	Port 1
	Port 7	GbE	Cisco UCS C-Series 4	M2
	Port 8	10GbE	Cisco UCS C-Series 4	Port 1
	Port 2/1	10GbE	Cisco UCS fabric interconnect B	
	Port 2/2	10GbE	Cisco UCS fabric interconnect B	

Table 15 *Cisco UCS C-Series 1*

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



Note

In this iSCSI boot FlexPod configuration, the X1140A-R6 Unified Target Adapters in slots 3 and 4 are not required; however, they were used in this lab validation. If these cards are not used in the FlexPod implementation, for the rest of this document substitute port e1b for e3a and e2b for e4a. Also, move the cards in slots 5 and 6 up to slots 3 and 4.

Storage Configuration

Controller FAS32xx Series

Table 16 *Controller FAS32xx series prerequisites*

Requirement	Reference	Comments
Physical site where storage system needs to be installed must be ready	Site Requirements Guide	Refer to the “Site Preparation” section.
Storage system connectivity requirements	Site Requirements Guide	Refer to the “System Connectivity Requirements” section.
Storage system general power requirements	Site Requirements Guide	Refer to the “Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements” section.
Storage system model-specific requirements	Site Requirements Guide	Refer to the “FAS32xx/V32xx Series Systems” section.

System Configuration Guides

System configuration guides provide supported hardware and software components for the specific Data ONTAP version. These online guides provide configuration information for all NetApp storage appliances currently supported by the Data ONTAP software. They also provide a table of component compatibilities.

1. Make sure that the hardware and software components are supported with the version of Data ONTAP that you plan to install by checking the System Configuration Guides at the NetApp Support site.
2. Click the appropriate NetApp storage appliance and then click the component you want to view. Alternatively, to compare components by storage appliance, click a component and then click the NetApp storage appliance you want to view.

Controllers

Follow the physical installation procedures for the controllers in the FAS32xx documentation at the NetApp Support site.

Disk Shelves DS2246 Series

DS2246 Disk Shelves

Follow the procedures in the Disk Shelf Installation and Setup section of the DS2246 Disk Shelf Overview to install a disk shelf for a new storage system.

Follow procedures for proper cabling with the controller model as described in SAS Disk Shelves Universal SAS and ACP Cabling Guide.

The following information applies to DS2246 disk shelves:

- SAS disk drives use software-based disk ownership. Ownership of a disk drive is assigned to a specific storage system by writing software ownership information on the disk drive rather than by using the topography of the storage system's physical connections.
- Connectivity terms used: shelf-to-shelf (daisy-chain), controller-to-shelf (top connections), and shelf-to controller (bottom connections).
- Unique disk shelf IDs must be set per storage system (a number from 0 through 98).
- Disk shelf power must be turned on to change the digital display shelf ID. The digital display is on the front of the disk shelf.
- Disk shelves must be power-cycled after the shelf ID is changed for it to take effect.
- Changing the shelf ID on a disk shelf that is part of an existing storage system running Data ONTAP requires that you wait at least 30 seconds before turning the power back on so that Data ONTAP can properly delete the old disk shelf address and update the copy of the new disk shelf address.
- Changing the shelf ID on a disk shelf that is part of a new storage system installation (the disk shelf is not yet running Data ONTAP) requires no wait; you can immediately power-cycle the disk shelf.

Cisco Nexus 5596 Cluster Network Switch Configuration



Note

If your FlexPod implementation uses 7-Mode instead of clustered Data ONTAP, go to section "Alternate 7-Mode NetApp FAS3250 Deployment Procedure: Part 2" from the appendix.

Table 17 *Cisco Nexus 5596 cluster network switch configuration prerequisites*

Description
<ul style="list-style-type: none"> ▪ Rack and connect power to the new Cisco Nexus 5596 switches ▪ Provide a terminal session that connects to the switch's serial console port (9600, 8, n, 1) ▪ Connect the <code>mgmt0</code> port to the management network and be prepared to provide IP address information ▪ Obtain password for admin ▪ Determine switch name ▪ Identify SSH keytype (dsa, rsa, or rsa1) ▪ Set up an e-mail server for Cisco Smart Call Home and IP connectivity between the switch and the e-mail server ▪ Provide SNMP contact information for Cisco Smart Call Home (name, phone, street address) ▪ Identify a CCO ID associated with an appropriate Cisco SMARTnet[®] Service contract for Cisco Smart Call Home ▪ Enable Cisco SMARTnet Service for the device to be registered for Cisco Smart Call home

Initial Setup of Cisco Nexus 5596 Cluster Interconnect

The first time a Cisco Nexus 5596 cluster interconnect is accessed, it runs a setup program that prompts the user to enter an IP address and other configuration information needed for the switch to communicate over the management Ethernet interface. This information is required to configure and manage the switch. If the configuration must be changed later, the setup wizard can be accessed again by running the setup command in EXEC mode.

To set up the Cisco Nexus 5596 cluster interconnect, complete the following steps. These steps will need to be completed on both cluster interconnects.

1. Provide applicable responses to the setup prompts displayed on the Cisco Nexus 5596 cluster interconnect.

```
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <switchname>
Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <ic_mgmt0_ip>
Mgmt0 IPv4 netmask: <ic_mgmt0_netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <ic_mgmt0_gw>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa): rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <ntp_server_ip>
Enter basic FC configurations (yes/no) [n]: Enter
```

2. At the end of the setup, the configuration choices are displayed. Verify the information and save the configuration at this time.

```
Would you like to edit the configuration? (yes/no) [n]: <n>
Use this configuration and save it? (yes/no) [y]: <y>
```

Download and Install NetApp Cluster Switch Software

When the Cisco Nexus 5596 is being used as a cluster network switch with Data ONTAP 8.1.2, it should be running NX-OS version 5.2(1)N1(1). The show version command from the switch command line interface will show the switch version currently running on the switch. If the currently running version is not 5.2(1)N1(1), go to the NetApp Support site and download and install NX-OS 5.2(1)N1(1) for the Cisco Nexus 5596 switch. Make sure both cluster interconnects are running NX-OS version 5.2(1)N1(1).

Download and Merge of NetApp Cluster Switch Reference Configuration File

Cluster network and management network switches are shipped without the configuration files installed. These files must be downloaded to the switches during deployment. Configuration files must be downloaded when the cluster network and management network switches are first installed or after the Cisco switch software is updated or reinstalled.

After the initial setup is complete, the NetApp cluster network switch reference configuration must be transferred to the switch and merged with the existing configuration. Instructions for this task and the reference configuration files for the appropriate switches are available on the NetApp Support site.

To download configuration files to a host and install them on a Cisco Nexus 5596 switch, complete the following steps on both cluster interconnects:

1. Obtain a console connection to the switch. Verify the existing configuration on the switch by running the show run command.

2. Log in to the switch. Make sure that the host recognizes the switch on the network (for example, use the ping utility).

3. Enter the following command:

```
copy <transfer protocol>: bootflash: vrf management
```

4. Verify that the configuration file is downloaded.
5. Merge the configuration file into the existing running-config. Run the following command, where <config file name> is the file name for the switch type. A series of warnings regarding PortFast is displayed as each port is configured.

```
copy <config file name> running-config
```

6. Verify the success of the configuration merge by running the show run command and comparing its output to the contents of the configuration file (a .txt file) that was downloaded.
 - a. The output for both installed-base switches and new switches should be identical to the contents of the configuration file for the following items:
 - banner (should match the expected version)
 - Switch port descriptions such as description Cluster Node x
 - The new ISL algorithm port-channel load-balance Ethernet source-dest-port
 - b. The output for new switches should be identical to the contents of the configuration file for the following items:
 - Port channel
 - Policy map
 - System QoS
 - Interface
 - Boot
 - c. The output for installed-base switches should have the flow control receive and send values on for the following items:
 - Interface port-channel 1 and 2
 - Ethernet interface 1/41 through Ethernet interface 1/48.

7. Copy the running-config to the startup-config.

```
copy running-config startup-config
```

Cisco Smart Call Home Setup

To configure Smart Call Home on a Cisco Nexus 5596 switch, complete the following steps:

1. Enter the mandatory system contact using the snmp-server contact command in global configuration mode. Then run the callhome command to enter callhome configuration mode.

```
NX-5596#config t
```

```
NX-5596(config)#snmp-server contact <sys-contact>
```

```
NX-5596(config)#callhome
```

2. Configure the mandatory contact information (phone number, e-mail address, and street address).

```
NX-5596(config-callhome)#email-contact <email-address>
```

```
NX-5596(config-callhome)#phone-contact <+1-000-000-0000>
```

```
NX-5596(config-callhome)#streetaddress <a-street-address>
```

3. Configure the mandatory e-mail server information. The server address is an IPv4 address, IPv6 address, or the domain-name of a SMTP server to which Call Home will send e-mail messages. Optional port number (default=25) and VRF may be configured.

```
NX-5596(config-callhome)#transport email smtp-server <ip-address> port 25
use-vrf <vrf-name>
```

4. Set the destination profile CiscoTAC-1 e-mail address to callhome@cisco.com

```
NX-5596(config-callhome)#destination-profile CiscoTAC-1 email-addr
callhome@cisco.com vrf management
```

5. Enable periodic inventory and set the interval.

```
NX-5596(config-callhome)#periodic-inventory notification
NX-5596(config-callhome)#periodic-inventory notification interval 30
```

6. Enable callhome, exit, and save the configuration.

```
NX-5596(config-callhome)#enable
NX-5596(config-callhome)#end
NX-5596#copy running-config startup-config
```

7. Send a callhome inventory message to start the registration process.

```
NX-5596#callhome test inventory
trying to send test callhome inventory message
successfully sent test callhome inventory message
```

8. Watch for an e-mail from Cisco regarding the registration of the switch. Follow the instructions in the e-mail to complete the registration for Smart Call Home.

SNMP Monitoring Setup

1. Configure SNMP by using the following example as a guideline. This example configures a host receiver for SNMPv1 traps and enables all link up/down traps.

```
NX-5596(config)# snmp-server host <ip-address> traps { version 1 } <community>
[udp_port <number>]
NX-5596(config)# snmp-server enable traps link
```

Clustered Data ONTAP 8.1.2

Node 1

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort
```

2. From the Loader-A prompt:

```
printenv
```

3. If the last-OS-booted-ver parameter is not set to 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.

4. Allow the system to boot up.

```
boot_ontap
```

5. Press Ctrl-C when the Press Ctrl-C for Boot Menu message appears.

**Note**

If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 is the version being booted, then select option 8 and yes to reboot the node. Then proceed with step 15.

6. To install new software, first select option 7.

7

7. Answer yes to perform a nondisruptive upgrade.

y

8. Select e0M for the network port you want to use for the download.

e0M

9. Select yes to reboot now.

y

10. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_node01_mgmt_ip>> <<var_node01_mgmt_mask>> <<var_node01_mgmt_gateway>>
```

11. Enter the URL where the software can be found.

**Note**

This web server must be pingable.

```
<<var_url_boot_software>>
```

12. Press Enter for the user name, indicating no user name.

Enter

13. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

y

14. Enter yes to reboot the node.

y

**Note**

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the LOADER prompt. If these actions occur, the system might deviate from this procedure.

15. Press Ctrl-C to exit autoboot when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

16. From the LOADER-A prompt, enter:

```
printenv
```

**Note**

If bootarg.init.boot_clustered true is not listed, the system is not set to boot in clustered Data ONTAP.

17. If the system is not set to boot in clustered Data ONTAP, at the LOADER prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. When you see Press Ctrl-C for Boot Menu:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

4

21. Answer yes to Zero disks, reset config and install a new file system.

y

22. Enter yes to erase all the data on the disks.

y



Note

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue to node 02 configuration while the disks for node 01 are zeroing.

Node 2

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

2. From the Loader-A prompt, enter:

printenv

3. If the last-OS-booted-ver parameter is not set to 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.

4. Allow the system to boot up.

boot_ontap

5. Press Ctrl-C when Press Ctrl-C for Boot Menu is displayed.

Ctrl-C



Note

If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 is the version being booted, then select option 8 and yes to reboot the node. Then proceed with step 15.

6. To install new software first select option 7.

7

7. Answer yes to perform a nondisruptive upgrade.

y

8. Select e0M for the network port you want to use for the download.

e0M

9. Select yes to reboot now.

y

10. Enter the IP address, netmask, and default gateway for e0M in their respective places.

<<var_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_mgmt_gateway>>

11. Enter the URL where the software can be found.



Note

This web server must be pingable.

<<var_url_boot_software>>

12. Press Enter for the user name, indicating no user name.

Enter

13. Select yes to set the newly installed software as the default to be used for subsequent reboots.

y

14. Select yes to reboot the node.

y



Note

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the LOADER prompt. If these actions occur, the system might deviate from this procedure.

15. Press Ctrl-C to exit autoboot when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

16. From the LOADER-A prompt, enter:

printenv



Note

If bootarg.init.boot_clustered true is not listed, the system is not set to boot in clustered Data ONTAP.

17. If the system is not set to boot in clustered Data ONTAP, at the LOADER prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
```

```
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

4

21. Answer yes to Zero disks, reset config and install a new file system.

y

22. Enter yes to erase all the data on the disks.

y



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.

Cluster Create in Clustered Data ONTAP

Table 18 Cluster create in clustered Data ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Clustered Data ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered Node01.

1. During the first node boot, the Cluster Setup wizard starts running on the console.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```



Note

If a login prompt appears instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings and then enter the cluster setup command.

2. Enter the following command to create a new cluster:

```
create
```

3. The system defaults are displayed.

```
System Defaults:
Private cluster network ports [e1a,e2a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.



Note

Cluster is created; this can take a minute or two.

5. The steps to create a cluster are displayed.

```
Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
```

Enter additional license key[]:

**Note**

For this validated architecture we recommend you install license keys for SnapRestore®, NFS, iSCSI, FlexClone, and SnapManager® Suite. After you finish entering the license keys, press Enter.

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster management interface port [e0a]: e0a
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway:
<<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

```
Enter the DNS domain names:<<var_dns_domain_name>>
Enter the name server IP addresses:<<var_nameserver_ip>>
```

**Note**

If you have more than one name server IP address, separate them with a comma.

7. Set up the node.

```
Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0b
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
enter the node management interface netmask:<<var_node01_mgmt_mask>>
Enter the node management interface default gateway:<<var_node01_mgmt_gateway>>
```

**Note**

The node management interface should be in a different subnet than the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

8. Press Enter to accept the AutoSupport™ message.**9. Reboot node 01.**

```
system node reboot <<var_node01>>
y
```

10. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

11. Select 5 to boot into maintenance mode.

```
5
```

12. When prompted Continue with boot?, enter y.**13. To verify the HA status of your environment, run the following command:**

```
ha-config show
```

**Note**

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

14. To see how many disks are unowned, enter:

```
disk show -a
No disks should be owned in this list.
```

15. Assign disks.

**Note**

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

```
disk assign -n <<var_#_of_disks>>
```

16. Reboot the controller.

```
halt
```

17. At the LOADER-A prompt, enter:

```
autoboot
```

Cluster Join in Clustered Data ONTAP

Table 19 *Cluster Join in clustered Data ONTAP prerequisites*

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered Node01, and the node joining the cluster in this example is Node02.

1. During the node boot, the Cluster Setup wizard starts running on the console.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
```

```
"back" - if you want to change previously answered questions, and
```

```
"exit" or "quit" - if you want to quit the cluster setup wizard.
```

```
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".
```

```
To accept a default or omit a question, do not enter a value.
```

```
Do you want to create a new cluster or join an existing cluster?
```

```
{create, join}:
```

**Note**

If a login prompt displays instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings, and then enter the cluster setup command.

2. Enter the following command to join a cluster:

```
join
```

3. The system defaults are displayed.

```
System Defaults:
```

```
Private cluster network ports [e1a,e2a].
```

```
Cluster port MTU values will be set to 9000.
```

```
Cluster interface IP addresses will be automatically generated.
```

```
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.

**Note**

The cluster creation can take a minute or two.

5. The steps to create a cluster are displayed.

Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter

**Note**

The node should find the cluster name.

6. Set up the node.

Enter the node management interface port [e0M]: e0b

Enter the node management interface IP address: <<var_node02_mgmt_ip>>

Enter the node management interface netmask: Enter

Enter the node management interface default gateway: Enter

7. The node management interface should be in a subnet different from the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

8. Press Enter to accept the AutoSupport message.

9. Log in to the Cluster Interface with the admin user id and <<var_password>>.

10. Reboot node 02.

system node reboot <<var_node02>>

Y

11. When you see Press Ctrl-C for Boot Menu, enter:

Ctrl - C

12. Select 5 to boot into maintenance mode.

5

13. At the question, Continue with boot? enter:

Y

14. To verify the HA status of your environment, enter:

**Note**

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

ha-config show

15. To see how many disks are unowned, enter:

disk show -a

16. Assign disks.

**Note**

This reference architecture allocates half the disks to each controller. Workload design could dictate different percentages, however. Assign all remaining disks to node 02.

disk assign -n <<var_#_of_disks>>

17. Reboot the controller:

halt

18. At the LOADER-A prompt, enter:

autoboot

19. Press Ctrl-C for boot menu when prompted.

Ctrl-C

Log Into the Cluster

To log into the cluster, do the following:

- Open an SSH connection to cluster IP or host name and log in to the admin user with the password you provided earlier.

Zero All Spare Disks

To zero all spare disks, enter:

```
disk zerospares
```

Set Auto-Revert on Cluster Management

To set the auto-revert parameter on the cluster management interface, enter:

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt
-auto-revert true
```

Failover Groups Management in Clustered Data ONTAP

To create a management port failover group, enter:

```
network interface failover-groups create -failover-group fg-cluster-mgmt -node
<<var_node01>> -port e0a
network interface failover-groups create -failover-group fg-cluster-mgmt -node
<<var_node02>> -port e0a
```

Assign Management Failover Group to Cluster Management LIF

To assign the management port failover group to the cluster management LIF, enter:

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt
-failover-group fg-cluster-mgmt
```

Failover Groups Node Management in Clustered Data ONTAP

To create a management port failover group, enter:

```
network interface failover-groups create -failover-group fg-node-mgmt-01 -node
<<var_node01>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-01 -node
<<var_node01>> -port e0M
network interface failover-groups create -failover-group fg-node-mgmt-02 -node
<<var_node02>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-02 -node
<<var_node02>> -port e0M
```

Assign Node Management Failover Groups to Node Management LIFs

To assign the management port failover group to the cluster management LIF, enter:

```
network interface modify -vserver <<var_node01>> -lif mgmt1 -auto-revert true
-use-failover-group enabled -failover-group fg-node-mgmt-01
network interface modify -vserver <<var_node02>> -lif mgmt1 -auto-revert true
-use-failover-group enabled -failover-group fg-node-mgmt-02
```

Flash Cache in Clustered Data ONTAP

Complete the following steps to enable Flash Cache™ on each node.

Run the following commands from the cluster management interface:

```
system node run -node <<var_node01>> options flexscale.enable on
system node run -node <<var_node01>> options flexscale.lopri_blocks off
system node run -node <<var_node01>> options flexscale.normal_data_blocks on
system node run -node <<var_node02>> options flexscale.enable on
system node run -node <<var_node02>> options flexscale.lopri_blocks off
system node run -node <<var_node02>> options flexscale.normal_data_blocks on
```



Note

Data ONTAP 8.1 and later does not require a separate license for Flash Cache.



Note

For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, refer to TR-3832: Flash Cache Best Practices Guide. Before customizing the settings, determine whether the custom settings are required or if the default settings are sufficient.

64-Bit Aggregates in Clustered Data ONTAP

A 64-bit aggregate containing the root volume is created during the Data ONTAP setup process. To create additional 64-bit aggregates, determine the aggregate name, the node on which to create it, and the number of disks it will contain.

1. Execute the following command to create new aggregates:

```
aggr create -aggregate aggr01 -nodes <<var_node01>> -B 64 -s <<var_raidsize>>
-diskcount <<var_num_disks>>
aggr create -aggregate aggr02 -nodes <<var_node02>> -B 64 -s <<var_raidsize>>
-diskcount <<var_num_disks>>
```



Note

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.



Note

Calculate the RAID group size to allow for roughly balanced (same size) RAID groups of from 12 through 20 disks (for SAS disks) within the aggregate. For example, if 52 disks were being assigned to the aggregate, select a RAID group size of 18. A RAID group size of 18 would yield two 18-disk RAID groups and one 16-disk RAID group. Keep in mind that the default RAID group size is 16 disks, and that the larger the RAID group size, the longer the disk rebuild time in case of a failure.



Note

The aggregate cannot be created until disk zeroing completes. Use the aggr show command to display aggregate creation status. Do not proceed until both aggr01 and aggr02 are online.

2. Disable Snapshot copies for the two data aggregates just created.

```
node run <<var_node01>> aggr options aggr01 nosnap on
node run <<var_node02>> aggr options aggr02 nosnap on
```

3. Delete any existing Snapshot copies for the two data aggregates.

```
node run <<var_node01>> snap delete -A -a -f aggr01
node run <<var_node02>> snap delete -A -a -f aggr02
```

4. Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02.

```
aggr show
aggr rename -aggregate aggr0 -newname <<var_node01_rootaggrname>>
```

Service Processor

Gather information about the network and the AutoSupport settings before configuring the Service Processor (SP).

Configure the SP using DHCP or static addressing. If the SP uses a static IP address, verify that the following SP prerequisites have been met:

- An available static IP address
- The network netmask
- The network gateway IP
- AutoSupport information

A best practice is to configure the AutoSupport recipients and mail host before configuring the SP. Data ONTAP automatically sends AutoSupport configuration to the SP, allowing the SP to send alerts and notifications through an AutoSupport message to the system administrative recipients specified in AutoSupport. When configuring the SP, enter the name or the IP address of the AutoSupport mail host, when prompted.

- A service processor needs to be set up on each node.

Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.1.2, you must upgrade to the latest service processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Using a web browser, connect to <http://support.netapp.com/NOW/cgi-bin/fw>.
2. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
3. Proceed to the download page for the latest release of the SP firmware for your storage platform.
4. Using the instructions on this page, update the SPs on both nodes in your cluster. You will need to download the .zip file to a web server that is reachable from the cluster management interface. In step 1a of the instructions substitute the following command: `system image get -node * -package http://web_server_name/path/SP_FW.zip`. Also, instead of `run local`, use `system node run <<var_nodename>>`, then execute steps 2-6 on each node.

Configure the Service Processor on Node 01

1. From the cluster shell, enter the following command:

```
system node run <<var_node01>> sp setup
```

2. Enter the following to set up the SP:


```

Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node01_sp_ip>>
Please enter the netmask of the SP[]: <<var_node01_sp_mask>>
Please enter the IP address for the SP gateway[]: <<var_node01_sp_gateway>>

```

Configure the Service Processor on Node 02

1. From the cluster shell, enter the following command:

```
system node run <<var_node02>> sp setup
```

2. Enter the following to set up the SP:

```

Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node02_sp_ip>>
Please enter the netmask of the SP[]: <<var_node02_sp_mask>>
Please enter the IP address for the SP gateway[]: <<var_node02_sp_gateway>>

```

Storage Failover in Clustered Data ONTAP

Run the following commands in a failover pair to enable storage failover.

1. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_node01>> -enabled true
```



Note

Enabling failover on one node enables it for both nodes.

2. Enable HA mode for two-node clusters only.



Note

Do not run this command for clusters with more than two nodes because it will cause problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

3. Verify that hardware assist is correctly configured and if needed modify the partner IP address.

```

storage failover hwassist show
storage failover modify -hwassist-partner-ip <<var_node02_mgmt_ip>> -node
<<var_node01>>
storage failover modify -hwassist-partner-ip <<var_node01_mgmt_ip>> -node
<<var_node02>>

```

IFGRP LACP in Clustered Data ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Therefore, make sure that the switch is configured properly.

1. Run the following commands on the command line to create interface groups (ifgrps).

```

ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e4a

```

```
ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e4a
```

**Note**

All interfaces must be in the down status before being added to an interface group.

**Note**

The interface group name must follow the standard naming convention of a0x.

VLAN in Clustered Data ONTAP

1. Create NFS VLANs.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>
```

2. Create iSCSI VLANs.

```
network port vlan create -node <<var_node01>> -vlan-name
a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_node01>> -vlan-name
a0a-<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_node02>> -vlan-name
a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_node02>> -vlan-name
a0a-<<var_iscsi_vlan_B_id>>
```

Jumbo Frames in Clustered Data ONTAP

1. To configure a clustered Data ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), run the following command from the cluster shell:

```
network port modify -node <<var_node01>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.

Do you want to continue? {y|n}: y

```
network port modify -node <<var_node02>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.

Do you want to continue? {y|n}: y

NTP in Clustered Data ONTAP

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

```
timezone <<var_timezone>>
```

**Note**

For example, in the Eastern United States, the time zone is America/New_York.

2. Set the date for the cluster.

```
date <ccyyymmddhhmm>
```

**Note**

The format for the date is <[Century][Year][Month][Day][Hour][Minute]>; for example, 201208081240.

3. Configure the Network Time Protocol (NTP) for each node in the cluster.

```
system services ntp server create -node <<var_node01>> -server
<<var_global_ntp_server_ip>> system services ntp server create -node
<<var_node02>> -server <<var_global_ntp_server_ip>>
```

4. Enable the NTP for the cluster.

```
system services ntp config modify -enabled true
```

SNMP in Clustered Data ONTAP

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the sysLocation and sysContact variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <<var_oncommand_server_fqdn>>
```

SNMPv1 in Clustered Data ONTAP

1. Set the shared secret plain-text password, which is called a community.

```
snmp community delete all
snmp community add ro <<var_snmp_community>>
```

**Note**

Use the delete all command with caution. If community strings are used for other monitoring products, the delete all command will remove them.

SNMPv3 in Clustered Data ONTAP

SNMPv3 requires that a user be defined and configured for authentication.

1. Create a user called snmpv3user.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

2. Select all of the default authoritative entities and select md5 as the authentication protocol.
3. Enter an eight-character minimum-length password for the authentication protocol, when prompted.
4. Select des as the privacy protocol.

5. Enter an eight-character minimum-length password for the privacy protocol, when prompted

AutoSupport HTTPS in Clustered Data ONTAP

AutoSupport sends support summary information to NetApp through HTTPS.

1. Execute the following commands to configure AutoSupport:

```
system node autosupport modify -node * -state enable -mail-hosts
<<var_mailhost>> -transport https -support enable -noteto
<<var_storage_admin_email>>
```

Cisco Discovery Protocol in Clustered Data ONTAP

Enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers by using the following procedure.



Note

To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

To enable CDP on the NetApp storage controllers, complete the following step:

1. Enable CDP on Data ONTAP.

```
node run -node <<var_node01>> options cdpd.enable on
node run -node <<var_node02>> options cdpd.enable on
```

Vserver

To create an infrastructure Vserver, complete the following steps:

1. Run the Vserver setup wizard.

```
vserver setup
```

Welcome to the Vserver Setup Wizard, which will lead you through the steps to create a virtual storage server that serves data to clients.

You can enter the following commands at any time:

"help" or "?" if you want to have a question clarified,
 "back" if you want to change your answers to previous questions, and
 "exit" if you want to quit the Vserver Setup Wizard. Any changes
 you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To accept a default
 or omit a question, do not enter a value.

Step 1. Create a Vserver.

You can type "back", "exit", or "help" at any question.

2. Enter the Vserver name.

Enter the Vserver name:Infra_Vserver

3. Select the Vserver data protocols to configure.

Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi}:nfs, iscsi

4. Select the Vserver client services to configure.

Choose the Vserver client services to configure {ldap, nis, dns}:Enter

5. Enter the Vserver's root volume aggregate:

Enter the Vserver's root volume aggregate {aggr01, aggr02} [aggr01]:aggr01

6. Enter the Vserver language setting. English is the default [C].

Enter the Vserver language setting, or "help" to see all languages [C]:Enter

7. Enter the Vserver's security style:

Enter the Vservers root volume's security style {unix, ntfs, mixed}} [unix]:
Enter

8. Answer no to Do you want to create a data volume?

Do you want to create a data volume? {yes, no} [Yes]: no

9. Answer no to Do you want to create a logical interface?

Do you want to create a logical interface? {yes, no} [Yes]: no

10. Answer no to Do you want to Configure FCP? {yes, no} [yes]: no.

Do you want to Configure FCP? {yes, no} [yes]: no

11. Add the two data aggregates to the Infra_Vserver aggregate list for NetApp Virtual Console.

```
vserver modify -vserver Infra_Vserver -aggr-list aggr01, aggr02
```

Create Load Sharing Mirror of Vserver Root Volume in Clustered Data ONTAP

1. Create a volume to be the load sharing mirror of the infrastructure Vserver root volume on each node.

```
volume create -vserver Infra_Vserver -volume root_vol_m01 -aggregate aggr01  
-size 20MB -type DP volume create -vserver Infra_Vserver -volume root_vol_m02  
-aggregate aggr02 -size 20MB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path //Infra_Vserver/root_vol -destination-path  
//Infra_Vserver/root_vol_m01 -type LS  
snapmirror create -source-path //Infra_Vserver/root_vol -destination-path  
//Infra_Vserver/root_vol_m02 -type LS
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //Infra_Vserver/root_vol
```

4. Set an hourly (at 5 minutes past the hour) update schedule on each mirroring relationship.

```
snapmirror modify -source-path //Infra_Vserver/root_vol -destination-path *  
-schedule hourly
```

iSCSI in Clustered Data ONTAP

1. Create the iSCSI service on each Vserver. The following command starts the iSCSI service and creates the IQN of the Vserver.

```
iscsi create -vserver Infra_Vserver
```

HTTPS Access in Clustered Data ONTAP

Secure access to the storage controller must be configured.

1. Increase the privilege level to access the certificate commands.

```
set -privilege advanced
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Check it with the following command:

```
security certificate show
```

3. Run the following commands as one-time commands to generate and install self-signed certificates:



Note

You can also use the security certificate delete command to delete expired certificates

```
security certificate create -vserver Infra_Vserver -common-name
<<var_security_cert_vserver_common_name>> -size 2048 -country
<<var_country_code>> -state <<var_state>> -locality <<var_city>> -organization
<<var_org>> -unit <<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>> -size 2048 -country
<<var_country_code>> -state <<var_state>> -locality <<var_city>> -organization
<<var_org>> -unit <<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node01>> -common-name
<<var_security_cert_node01_common_name>> -size 2048 -country
<<var_country_code>> -state <<var_state>> -locality <<var_city>> -organization
<<var_org>> -unit <<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node02>> -common-name
<<var_security_cert_node02_common_name>> -size 2048 -country
<<var_country_code>> -state <<var_state>> -locality <<var_city>> -organization
<<var_org>> -unit <<var_unit>> -email <<var_storage_admin_email>>
```

4. Configure and enable SSL and HTTPS access and disable Telnet access.

```
system services web modify -external true -sslv3-enabled true
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -action allow
system services firewall policy create -policy mgmt -service http -action deny
-ip-list 0.0.0.0/0
system services firewall policy delete -policy mgmt -service telnet -action
allow
system services firewall policy create -policy mgmt -service telnet -action deny
-ip-list 0.0.0.0/0
security ssl modify -vserver Infra_Vserver -certificate
<<var_security_cert_vserver_common_name>> -enabled true
y
security ssl modify -vserver <<var_clustername>> -certificate
<<var_security_cert_cluster_common_name>> -enabled true
y
security ssl modify -vserver <<var_node01>> -certificate
<<var_security_cert_node01_common_name>> -enabled true
y
security ssl modify -vserver <<var_node02>> -certificate
<<var_security_cert_node02_common_name>> -enabled true
y
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
```

```
vserver services web access create -name spi -role admin -vserver
<<var_clustername>>
vserver services web access create -name ontapi -role admin -vserver
<<var_clustername>>\
vserver services web access create -name compat -role admin -vserver
<<var_clustername>>
```

**Note**

It is normal for some of these commands to return an error message stating that the entry does not exist

NFSv3 in Clustered Data ONTAP

Run all commands to configure NFS on the Vserver.

1. Secure the default rule for the default export policy and create the FlexPod export policy.

```
vserver export-policy rule modify -vserver Infra_Vserver -policyname default
-ruleindex 1 -rorule never -rwrule never -superuser never
vserver export-policy create -vserver Infra_Vserver FlexPod
```

2. Create a new rule for the FlexPod export policy.

**Note**

For each ESXi host being created, create a rule. Each host will have its own rule index. Your first ESXi host will have rule index 1, your second ESXi host will have rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra_Vserver -policyname FlexPod
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys
-rwrule sys -superuser sys -allow-suid false
```

3. Assign the FlexPod export policy to the infrastructure Vserver root volume.

```
volume modify -vserver Infra_Vserver -volume root_vol -policy FlexPod
```

FlexVol in Clustered Data ONTAP

The following information is required to create a FlexVol® volume: the volume's name and size, and the aggregate on which it will exist.

1. Create two VMware datastore volumes, a server boot volume, and a volume to hold the OnCommand database LUN. Also, update the Vserver root volume load sharing mirrors to make the NFS mounts accessible.

```
volume create -vserver Infra_Vserver -volume infra_datastore_1 -aggregate aggr02
-size 500g -state online -policy FlexPod -junction-path /infra_datastore_1
-space-guarantee none -percent-snapshot-space 0
```

```
volume create -vserver Infra_Vserver -volume infra_swap -aggregate aggr01 -size
100g -state online -policy FlexPod -junction-path /infra_swap -space-guarantee
none -percent-snapshot-space 0 -snapshot-policy none
```

```
volume create -vserver Infra_Vserver -volume esxi_boot -aggregate aggr01 -size
100g -state online -policy default -space-guarantee none -percent-snapshot-space
0
```

```
volume create -vserver Infra_Vserver -volume OnCommandDB -aggregate aggr02 -size
200g -state online -policy default -space-guarantee none -percent-snapshot-space
0
```

```
snapmirror update-ls-set -source-path //Infra_Vserver/root_vol
```

LUN in Clustered Data ONTAP

1. Create two boot LUNS: VM-Host-Infra-01 and VM-Host-Infra-02.

```
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -size 10g -ostype vmware -space-reserve disabled
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-02 -size 10g -ostype vmware -space-reserve disabled
```

Deduplication in Clustered Data ONTAP

1. Enable deduplication on appropriate volumes.

```
volume efficiency on -vserver Infra_Vserver -volume infra_datastore_1
volume efficiency on -vserver Infra_Vserver -volume esxi_boot
volume efficiency on -vserver Infra_Vserver -volume OnCommandDB
```

Failover Groups NAS in Clustered Data ONTAP

1. Create an NFS port failover group.

```
network interface failover-groups create -failover-group
fg-nfs-<<var_nfs_vlan_id>> -node <<var_node01>> -port a0a-<<var_nfs_vlan_id>>
network interface failover-groups create -failover-group
fg-nfs-<<var_nfs_vlan_id>> -node <<var_node02>> -port a0a-<<var_nfs_vlan_id>>
```

NFS LIF in Clustered Data ONTAP

1. Create an NFS logical interface (LIF).

```
network interface create -vserver Infra_Vserver -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>>
-address <<var_node01_nfs_lif_ip>> -netmask <<var_node01_nfs_lif_mask>>
-status-admin up -failover-policy nextavail -firewall-policy data -auto-revert
true -use-failover-group enabled -failover-group fg-nfs-<<var_nfs_vlan_id>>

network interface create -vserver Infra_Vserver -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>>
-address <<var_node02_nfs_lif_ip>> -netmask <<var_node02_nfs_lif_mask>>
-status-admin up -failover-policy nextavail -firewall-policy data -auto-revert
true -use-failover-group enabled -failover-group fg-nfs-<<var_nfs_vlan_id>>
```

Create iSCSI LIF in Clustered Data ONTAP

1. Create four iSCSI LIFs, two on each node.

```
network interface create -vserver Infra_Vserver -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_node01>> -home-port 3a -address
<<var_node01_iscsi_A_IP>> -netmask <<iscsi_A_mask>>
```



```

network interface create -vserver Infra_Vserver -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_node01>> -home-port 4a -address
<<var_node01_iscsi_B_IP>> -netmask <<iscsi_B_mask>

network interface create -vserver Infra_Vserver -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_node02>> -home-port 3a -address
<<var_node02_iscsi_A_IP>> -netmask <<iscsi_A_mask>

network interface create -vserver Infra_Vserver -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_node02>> -home-port 4a -address
<<var_node02_iscsi_B_IP>> -netmask <<iscsi_B_mask>

```

Add Infrastructure Vserver Administrator

1. Add the infrastructure Vserver administrator and Vserver administration logical interface in the out-of-band management network with the following commands:

```

network interface create -vserver Infra_Vserver -lif vsmgmt -role data
-data-protocol none -home-node <<var_node02>> -home-port e0a -address
<<var_vserver_mgmt_ip>> -netmask <<var_vserver_mgmt_mask>> -status-admin up
-failover-policy nextavail -firewall-policy mgmt -auto-revert true
-use-failover-group enabled -failover-group fg-cluster-mgmt

network routing-groups route create -vserver Infra_Vserver -routing-group
d<<var_clustermgmt_ip>> -destination 0.0.0.0/0 -gateway
<<var_clustermgmt_gateway>>
security login password -username vsadmin -vserver Infra_Vserver
Please enter a new password: <<var_vsadmin_password>>
Please enter it again: <<var_vsadmin_password>>

security login unlock -username vsadmin -vserver Infra_Vserver

```

Server Configuration

FlexPod Cisco Unified Computing System

Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlexPod environment. The steps are necessary to provision the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS 6248 A

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect.

```

Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have choosen to setup a a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y

```

```

Enter the password for "admin": <<var_password>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)?
(yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
2. Review the settings printed to the console. If they are correct, answer yes to apply and save the
   configuration.
3. Wait for the login prompt to make sure that the configuration has been saved.

```

Cisco UCS 6248 B

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

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

```

Enter the configuration method: console
Installer has detected the presence of a peer Fabric interconnect. This Fabric
interconnect will be added to the cluster. Do you want to continue {y|n}? y
Enter the admin password for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no): y

```

2. Wait for the login prompt to make sure that the configuration has been saved.

FlexPod Cisco UCS iSCSI vSphere on Clustered Data ONTAP

Log Into Cisco UCS Manager

To log in to the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.
2. Click 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 as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

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

This document assumes the use of Cisco UCS 2.1(1b). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to version 2.1(1b), refer to Cisco UCS Manager Install and Upgrade Guides.

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:



Note

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools > IP Pool ext-mgmt.
3. In the Actions pane, select Create Block of IP Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.
5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

1. In Cisco UCS Manager, click the Admin tab in the navigation pane.
2. Select All > Timezone Management.
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter <<var_global_ntp_server_ip>> and click OK.
7. Click OK.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity.

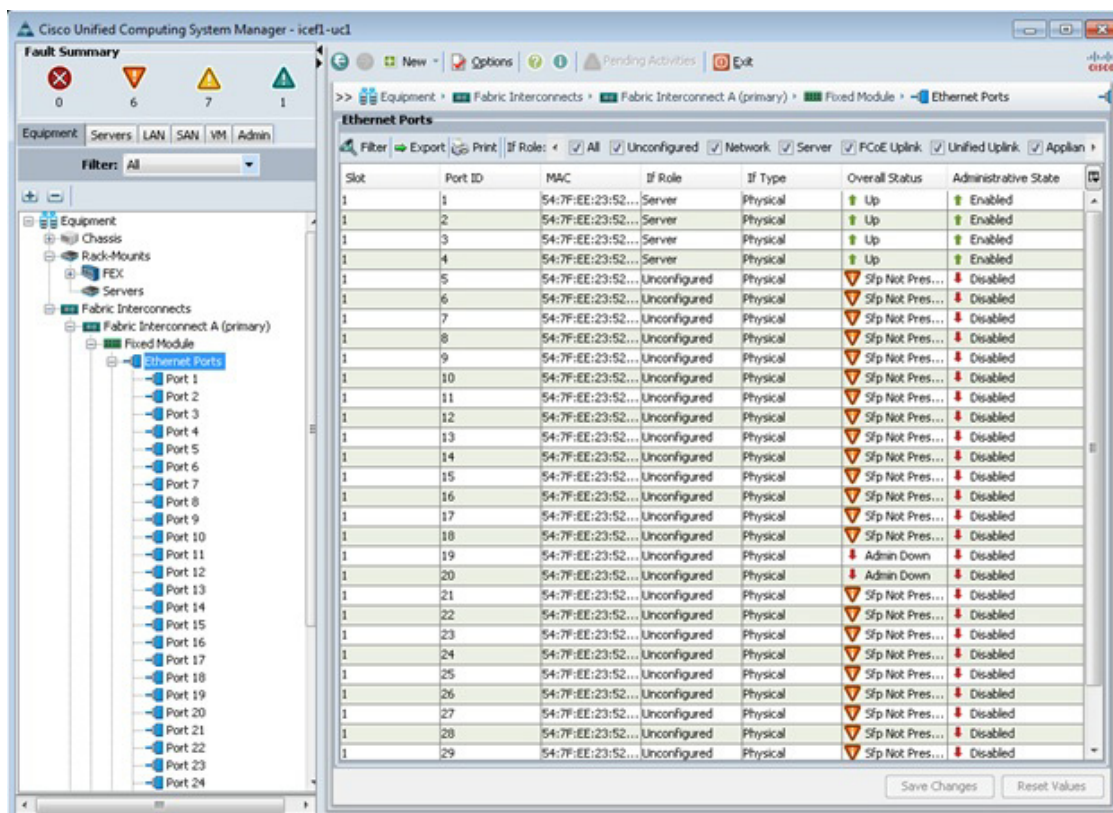
To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to 2-link or set it to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXs) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.
5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis or to the Cisco 2232 FEX are now configured as server ports.



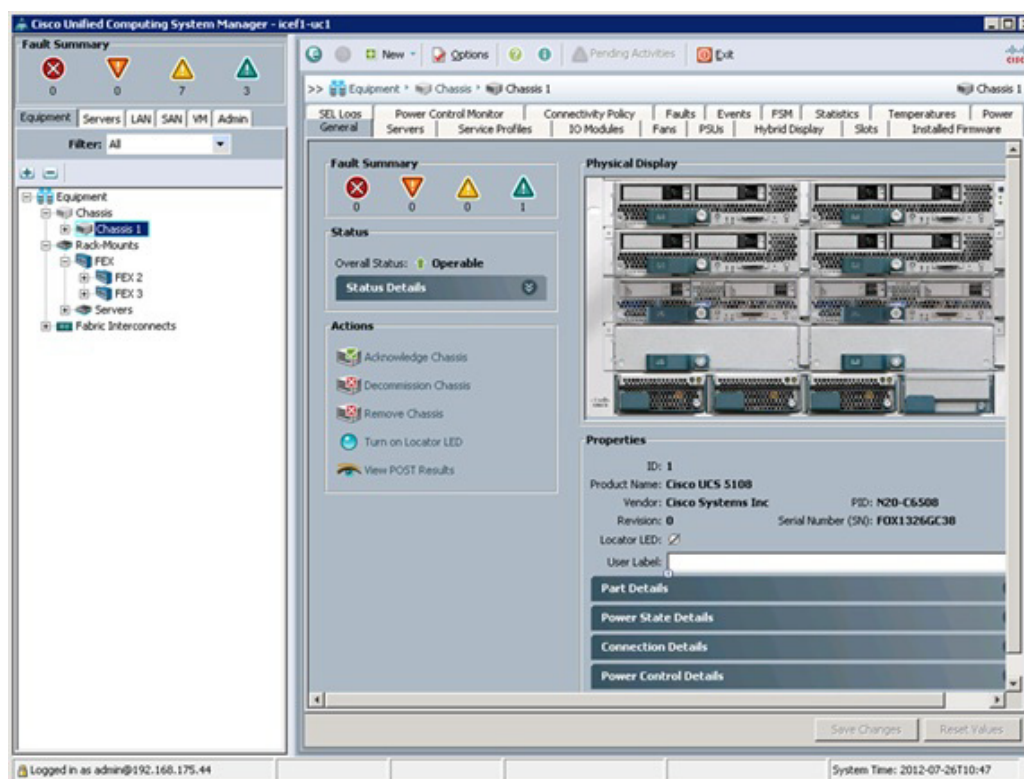
7. Select ports 19 and 20 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.

13. Select ports 19 and 20 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

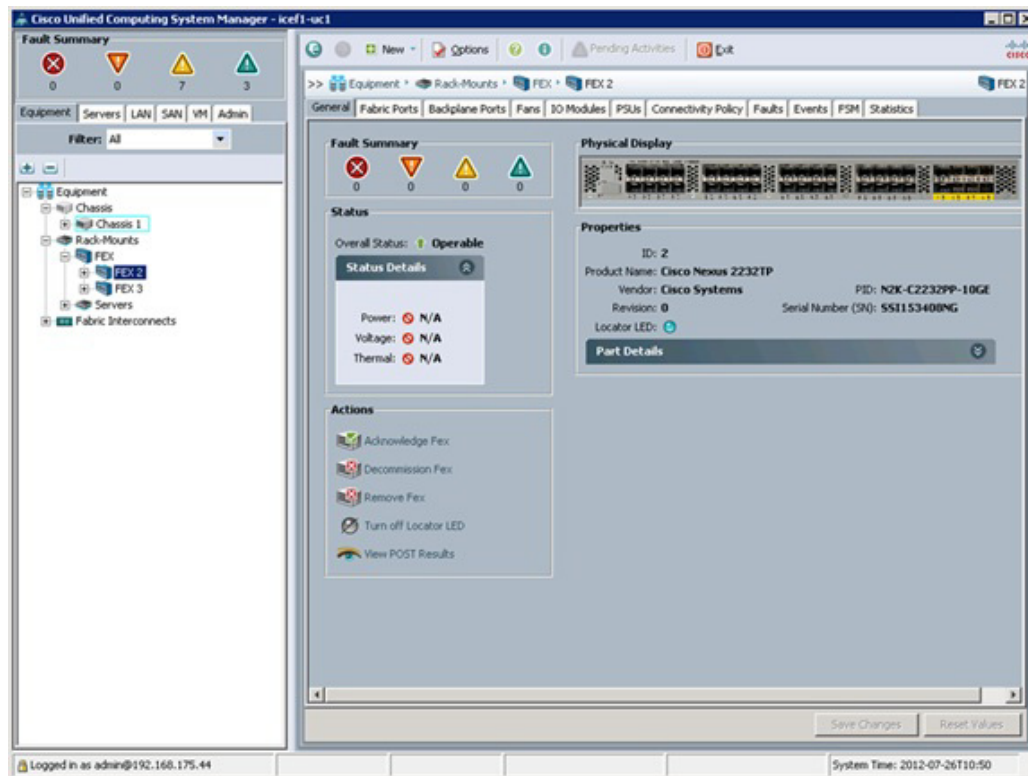
Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and external 2232 FEX modules, do the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.



4. Click Yes and then click OK to complete acknowledging the chassis.
5. If C-Series servers are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.



7. Click Yes and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.



Note

In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches.

2. Under LAN > 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 port channel.
6. Enter vPC-13-Nexus as the name of the port channel.
7. Click Next.

8. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 19
 - Slot ID 1 and port 20
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-Nexus as the name of the port channel.
17. Click Next.
18. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 19
 - Slot ID 1 and port 20
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create an Organization

Organizations are used to organize resources and restrict access to various groups within the IT organization, thereby enabling multi-tenancy of the compute resources.



Note

Although this document does not assume the use of organizations this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization.
4. Click OK.
5. Click OK in the confirmation message.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, do the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root.



Note

In this procedure, two MAC address pools are created, one for each switching fabric.

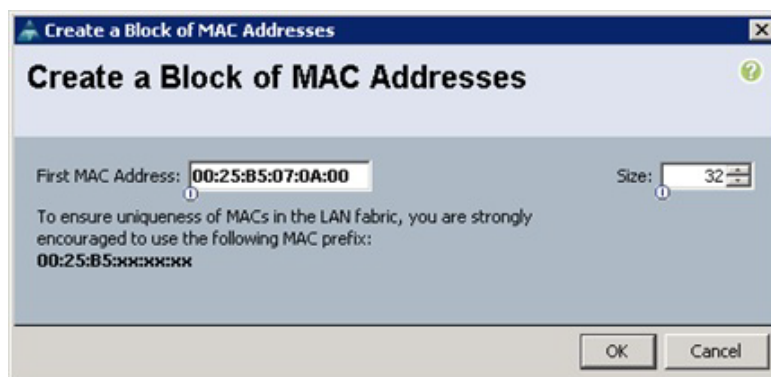
3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC_Pool_A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Click Next.
8. Click Add.
9. Specify a starting MAC address.



Note

For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.



11. Click OK.
12. Click Finish.
13. In the confirmation message, 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 as the name of the MAC pool.
17. Optional: Enter a description for the MAC pool.
18. Click Next.
19. Click Add.
20. Specify a starting MAC address.

**Note**

For the FlexPod solution, the recommendation is to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

21. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
22. Click OK.
23. Click Finish.
24. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps.

Cisco UCS Manager

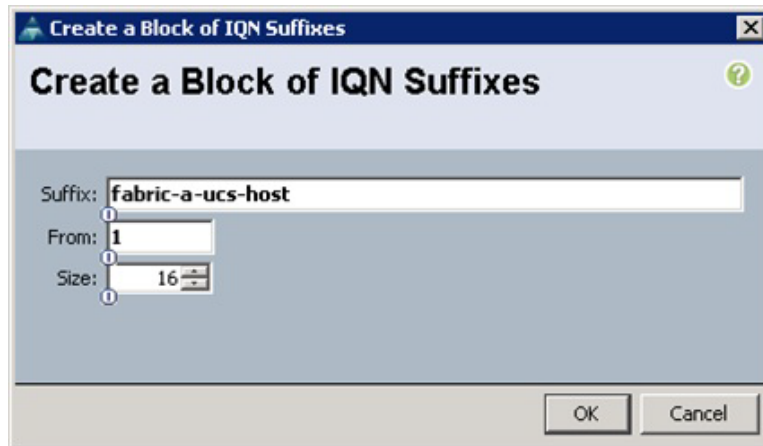
1. Select the SAN tab on the left.
2. Select Pools > root.

**Note**

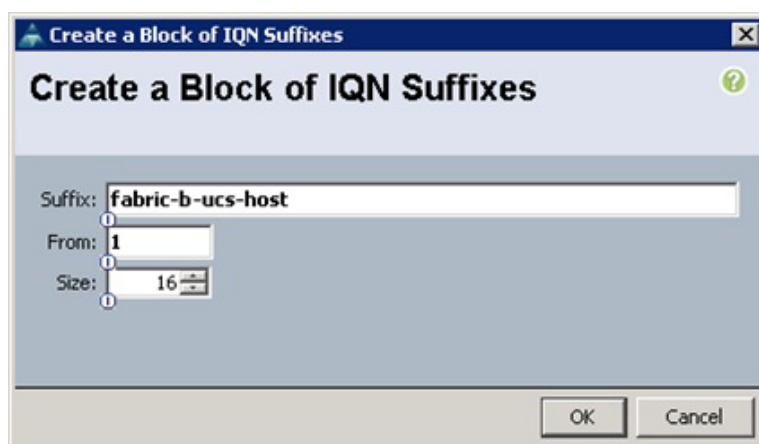
Two IQN pools are created, one for each switching fabric.

3. Right-click IQN Pools under the root organization.

4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN_Pool_A for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.1992-08.com.cisco as the prefix
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter fabric-a-ucs-host as the suffix.
12. Enter 1 in the From field.
13. Specify a size of the IQN block sufficient to support the available server resources.
14. Click OK.



15. Click Finish.
16. In the message box that displays, click OK.
17. Right-click IQN Pools under the root organization.
18. Select Create IQN Suffix Pool to create the IQN pool.
19. Enter IQN_Pool_B for the name of the IQN pool.
20. Optional: Enter a description of the IQN pool.
21. Enter iqn.1992-08.com.cisco as the prefix.
22. Select Sequential for Assignment Order.
23. Click Next.
24. Click Add.
25. Enter fabric-b-ucs-host as the suffix.
26. Enter 1 in the From field.
27. Specify a size of the IQN block sufficient to support the available server resources.
28. Click OK.



29. Click Finish.
30. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

These steps provide details for configuring the necessary IP pools iSCSI boot for the Cisco UCS environment.

Cisco UCS Manager

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



Note

Two IP pools are created, one for each switching fabric.

3. Right-click IP Pools under the root organization.
4. Select Create IP Pool to create the IP pool.
5. Enter iSCSI_IP_Pool_A for the name of the IP pool.
6. Optional: Enter a description of the IQN pool.
7. Select Sequential for Assignment Order.
8. Click Next.
9. Click Add.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.
12. Click OK.
13. Click Finish.
14. Right-click IP Pools under the root organization.
15. Select Create IP Pool to create the IP pool.
16. Enter iSCSI_IP_Pool_B for the name of the IP pool.
17. Optional: Enter a description of the IQN pool.

18. Select Sequential for Assignment Order.
19. Click Next.
20. Click Add.
21. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
22. Set the size to enough addresses to accommodate the servers.
23. Click OK.
24. Click Finish.

Create Block of IP Addresses

Create a Block of IP Addresses

From: Size:

Subnet Mask: Default Gateway:

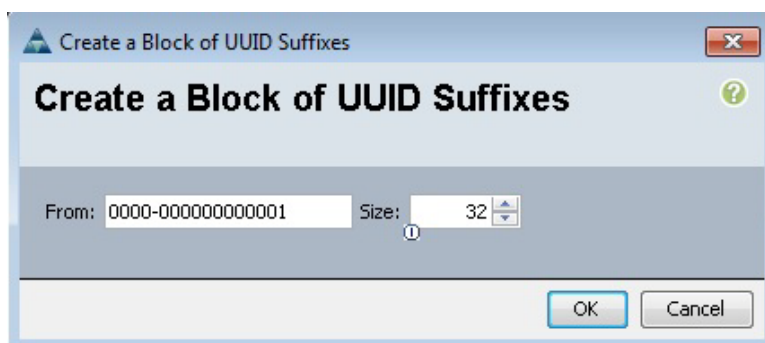
Primary DNS: Secondary DNS:

OK Cancel

Create UUID Suffix Pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Click Next.
9. Click Add to add a block of UUIDs.
10. Keep the From field at the default setting.
11. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



12. Click OK.
13. Click Finish.
14. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



Note

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Infra_Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two servers to be used for the VMware management cluster and click >> to add them to the Infra_Pool server pool.
9. Click Finish.
10. Click OK.

Create VLANs

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

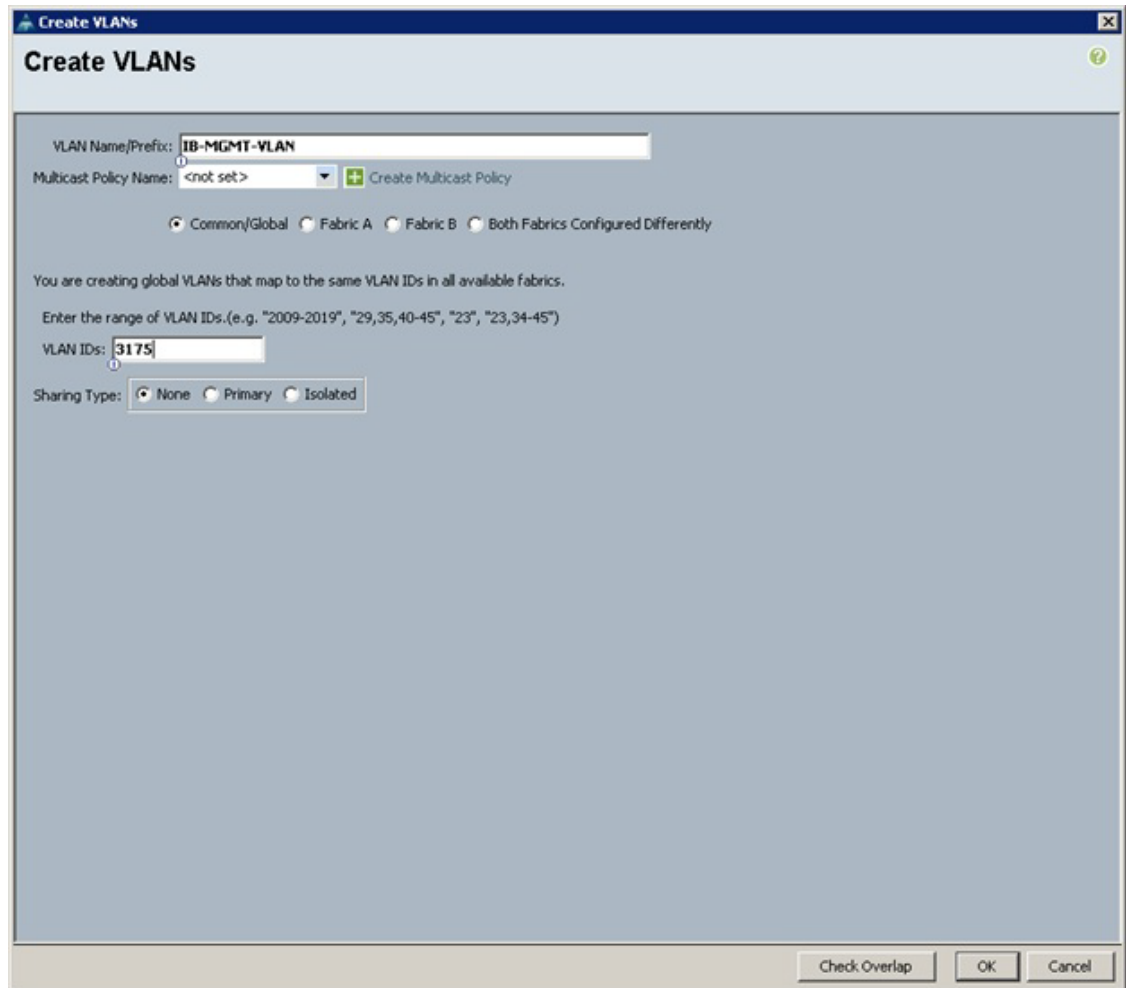


Note

In this procedure, five VLANs are created.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.

4. Select Create VLANs.
5. Enter IB-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 <<var_ib-mgmt_vlan_id>> as the ID of the management VLAN.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.



10. Right-click VLANs.
11. Select Create VLANs.
12. Enter NFS-VLAN as the name of the VLAN to be used for NFS.
13. Keep the Common/Global option selected for the scope of the VLAN.
14. Enter the <<var_nfs_vlan_id>> for the NFS VLAN.
15. Keep the Sharing Type as None.
16. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name: [+ Create Multicast Policy](#)

☒ 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.

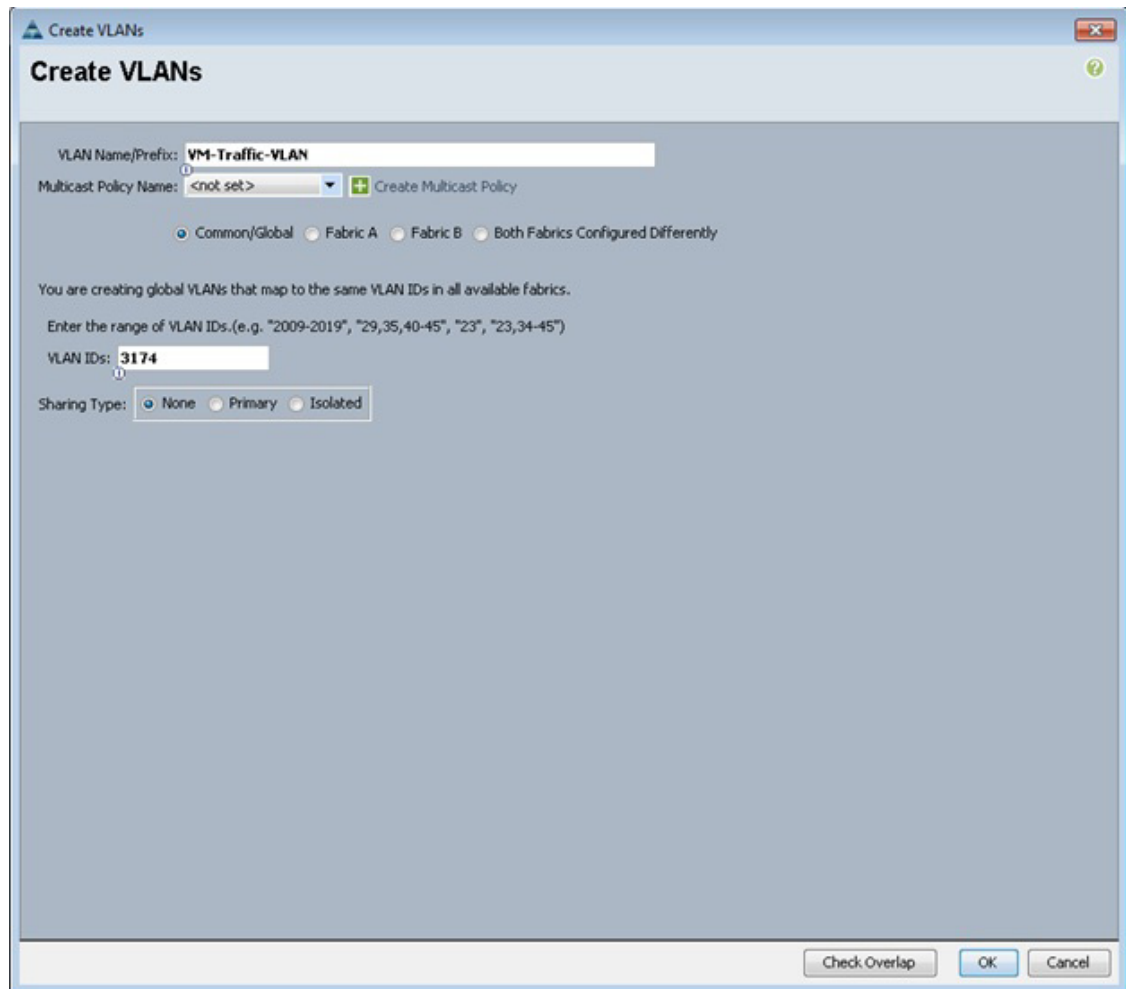
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

[Check Overlap](#)
[OK](#)
[Cancel](#)

17. Right-click VLANs.
18. Select Create VLANs.
19. Enter vMotion-VLAN as the name of the VLAN to be used for vMotion.
20. Keep the Common/Global option selected for the scope of the VLAN.
21. Enter the `<<var_vmotion_vlan_id>>` as the ID of the vMotion VLAN.
22. Keep the Sharing Type as None.
23. Click OK, and then click OK again.
24. Right-click VLANs.
25. Select Create VLANs.
26. Enter VM-Traffic-VLAN as the name of the VLAN to be used for the VM traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the `<<var_vm-traffic_vlan_id>>` for the VM Traffic VLAN.
29. Keep the Sharing Type as None.
30. Click OK, and then click OK again.



31. Right-click VLANs.
32. Select Create VLANs.
33. Enter iSCSI-A-VLAN as the name of the VLAN to be used for the first iSCSI VLAN.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the VLAN ID for the first iSCSI VLAN.
36. Click OK, then OK.

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name: [+ Create Multicast Policy](#)

☒ 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.

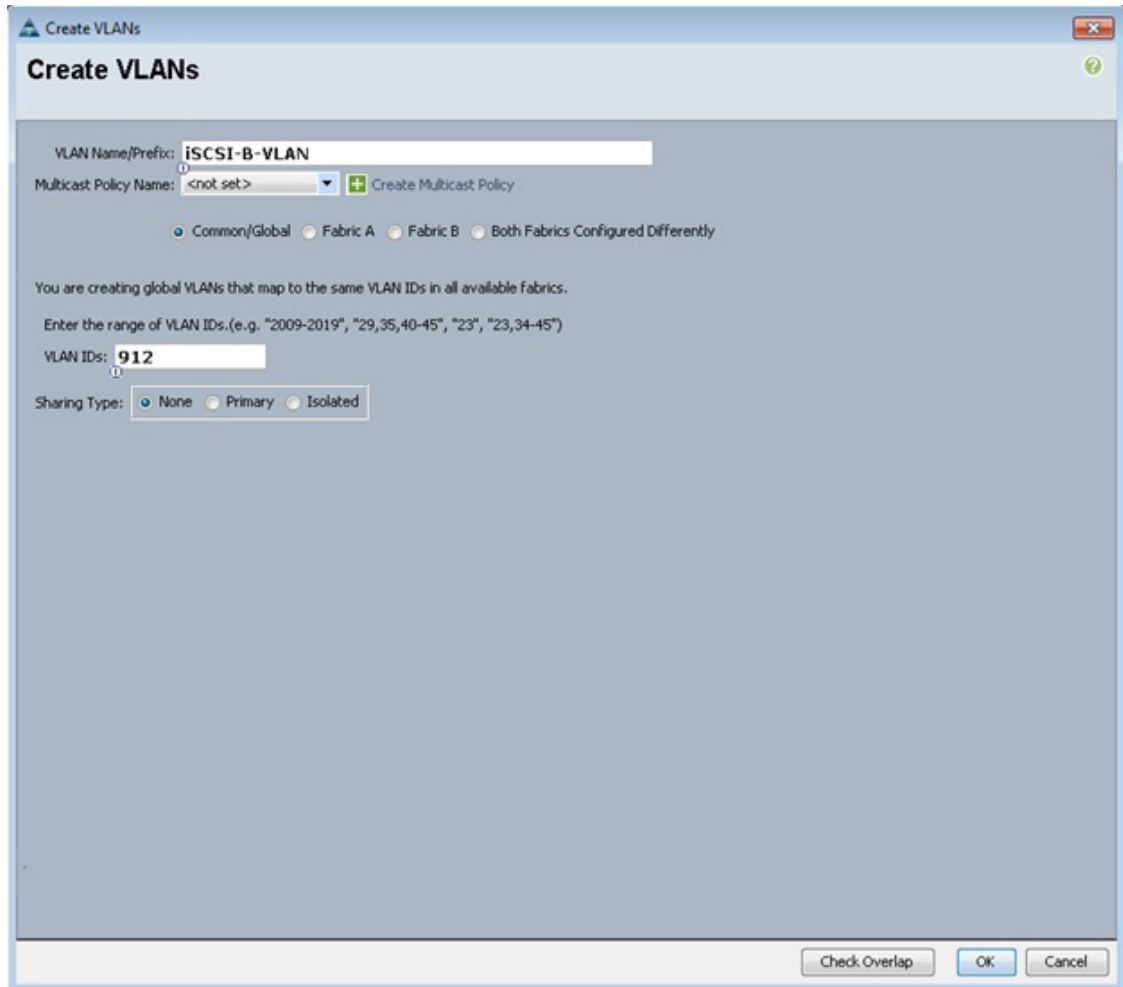
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

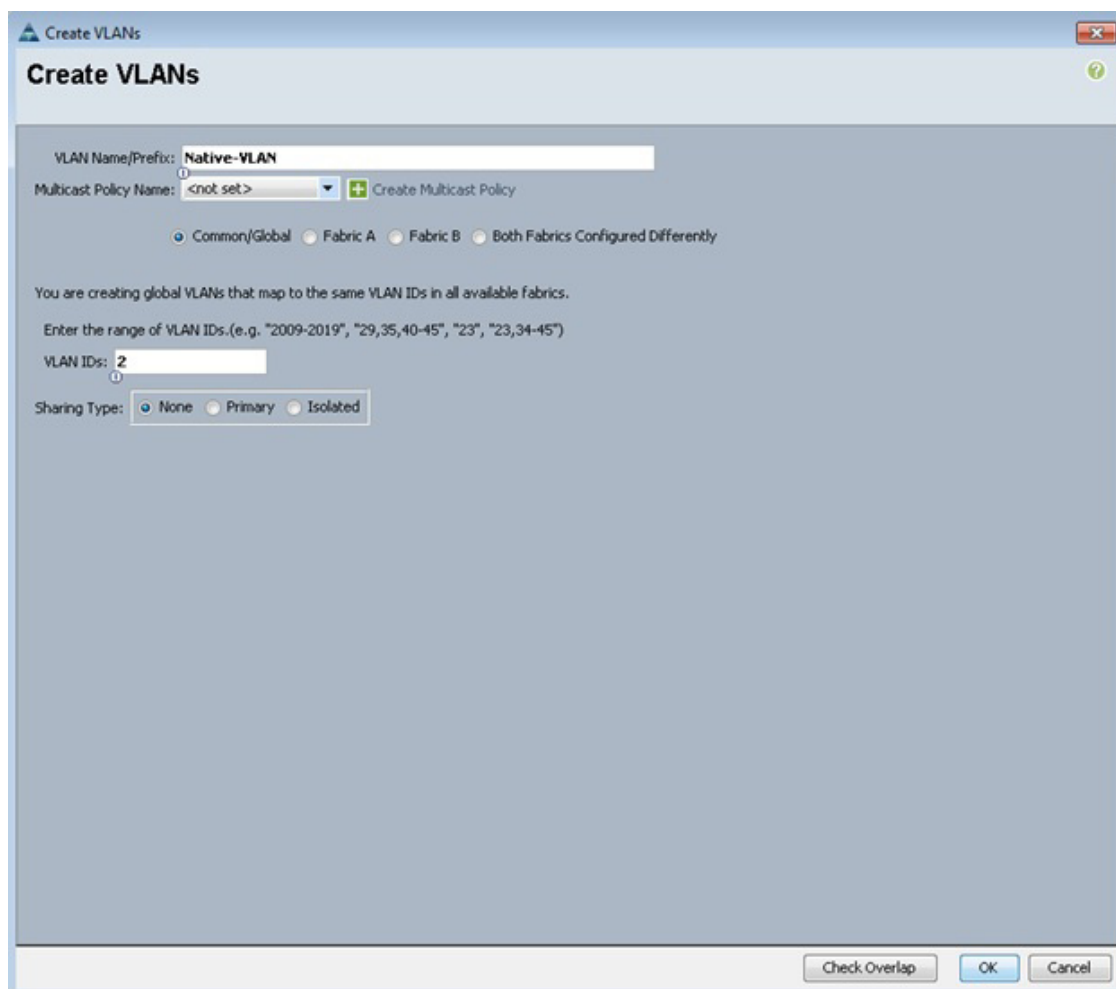
Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

[Check Overlap](#)
[OK](#)
[Cancel](#)

37. Right-click VLANs.
38. Select Create VLANs.
39. Enter iSCSI-B-VLAN as the name of the VLAN to be used for the second iSCSI VLAN.
40. Keep the Common/Global option selected for the scope of the VLAN.
41. Enter the VLAN ID for the second iSCSI VLAN.
42. Click OK, then OK.



43. Right-click VLANs.
44. Select Create VLANs.
45. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
46. Keep the Common/Global option selected for the scope of the VLAN.
47. Enter the <<var_native_vlan_id>> as the ID of the native VLAN.
48. Keep the Sharing Type as None.
49. Click OK, and then click OK again.



50. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
51. Click Yes, and then click OK.

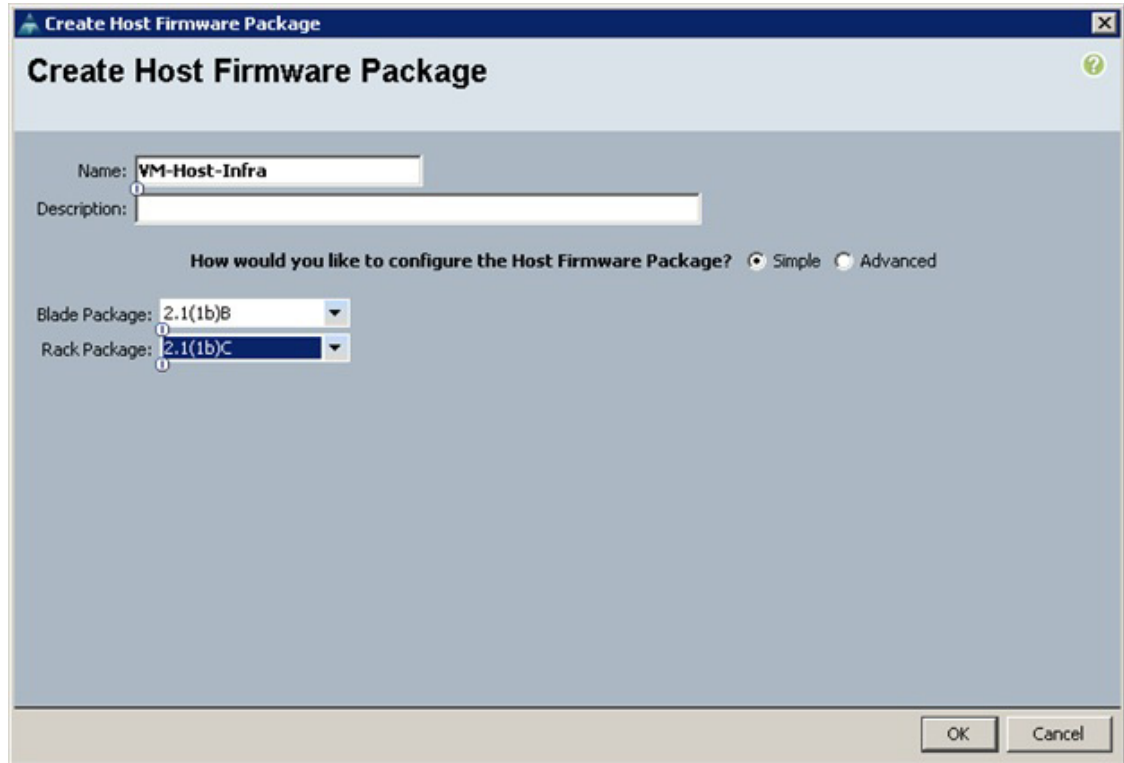
Create Host Firmware Package

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

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
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. Leave Simple selected.

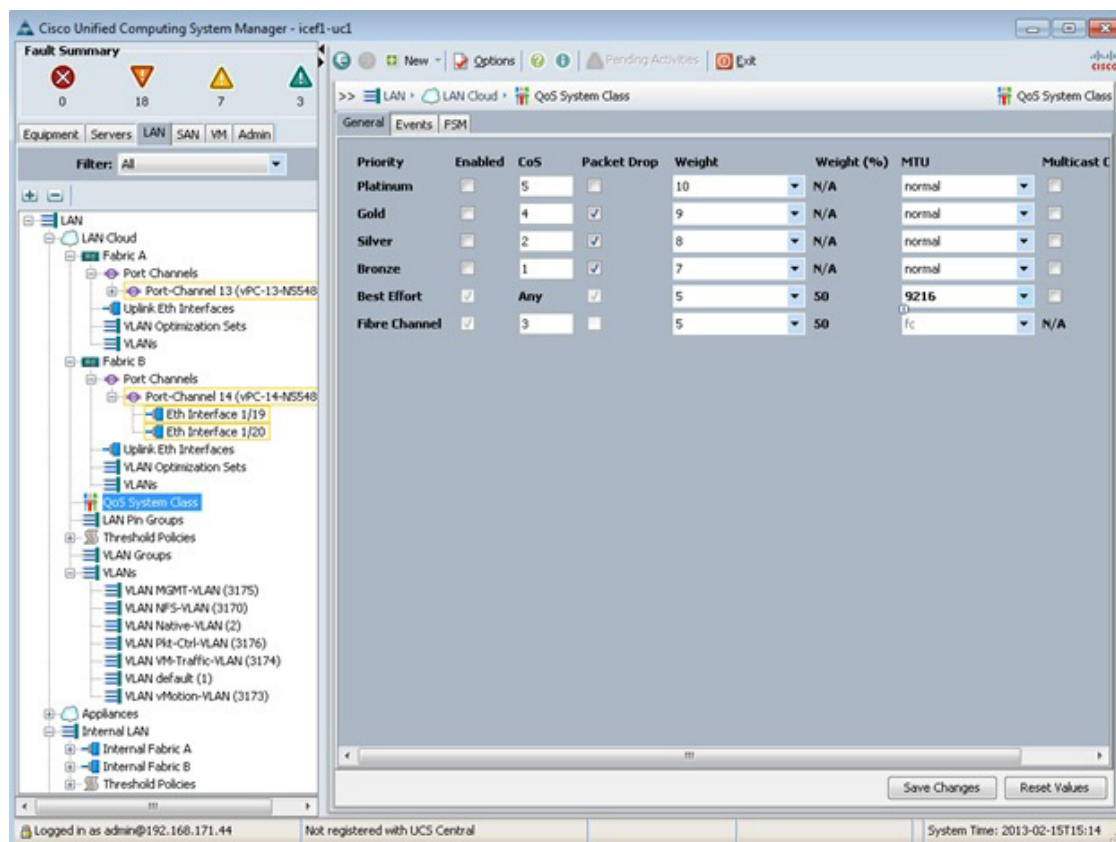
7. Select the version 2.1(1b) for both the Blade and Rack Packages.
8. Click OK to create the host firmware package.
9. Click OK.



Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.
6. Click OK.



Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.

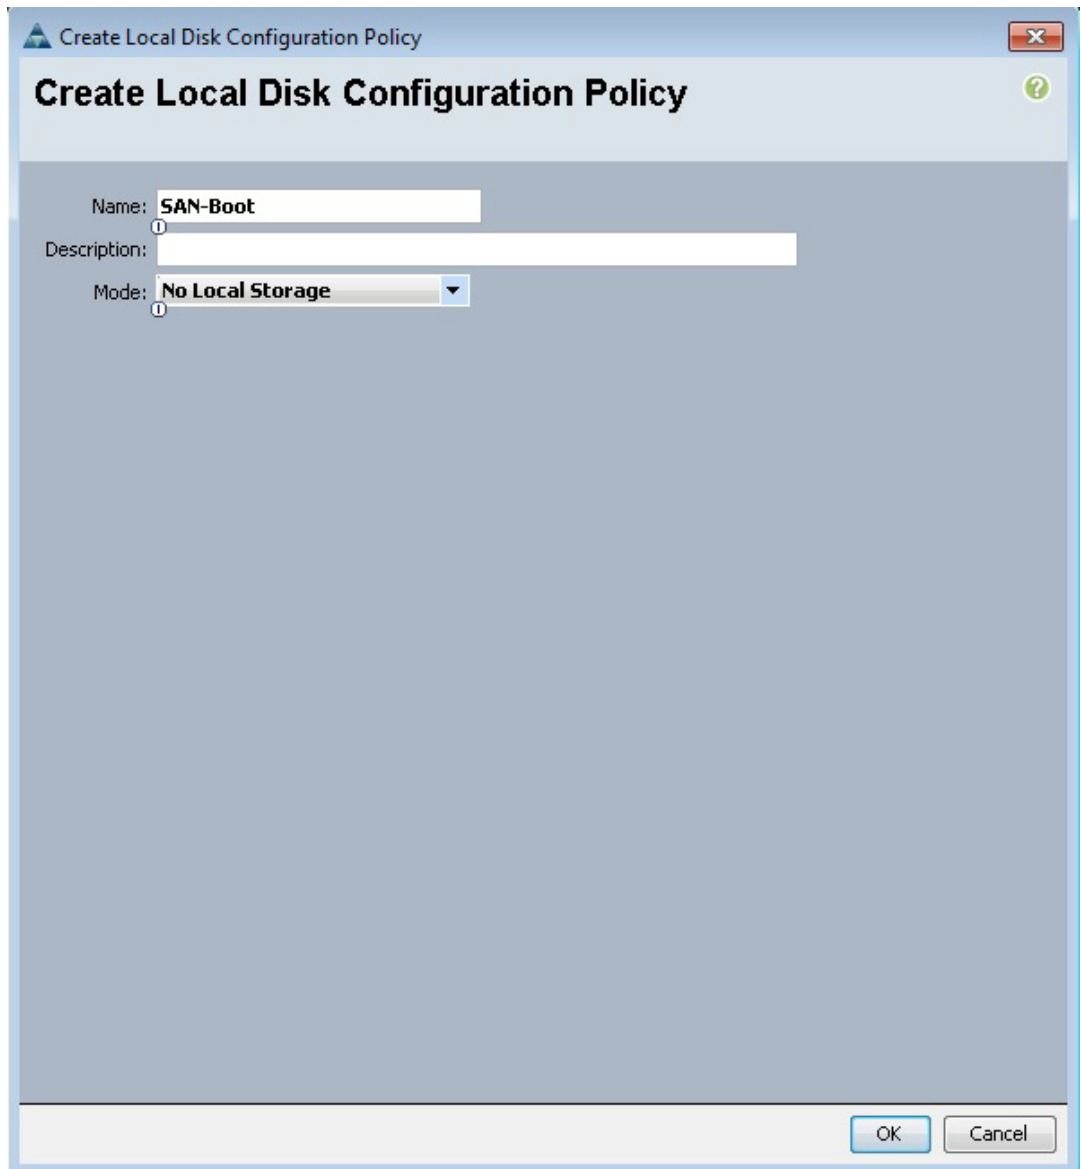


Note

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

To create a local disk configuration policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select 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.
7. Click OK to create the local disk configuration policy.
8. Click OK.



Create Local Disk Configuration Policy

Name: **SAN-Boot**

Description:

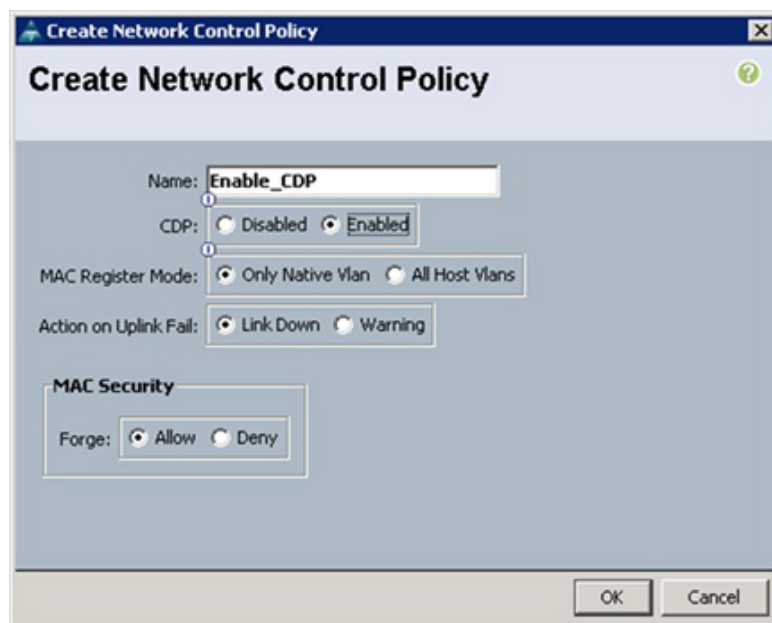
Mode: **No Local Storage**

OK Cancel

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

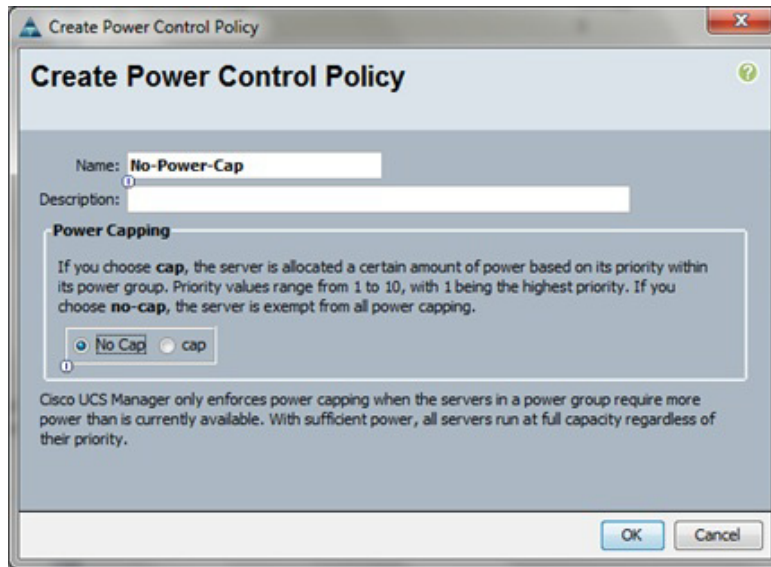
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable_CDP as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.



Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select 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 setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.



Create Server Pool Qualification Policy (Optional)

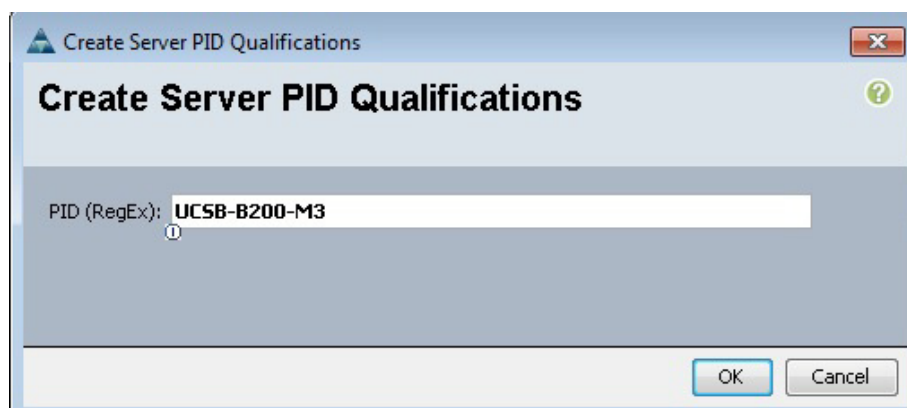
To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:



Note

This example creates a policy for a B200-M3 server.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Enter UCSB-B200-M3 as the name for the policy.
6. Select Create Server PID Qualifications.
7. Enter UCSB-B200-M3 as the PID.
8. Click OK to create the server pool qualification policy.
9. Click OK, and then click OK again.



Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select 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 setting to Disabled.
7. Click Finish to create the BIOS policy.



8. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter VM-Host-Infra as the name of the placement policy.
6. Click 1 and select Assigned Only.
7. Click OK, and then click OK again.

Create Placement Policy

Name: **VM-Host-Infra**

Virtual Slot Mapping Scheme: ☒ Round Robin ☐ Linear Ordered

Filter Export Print

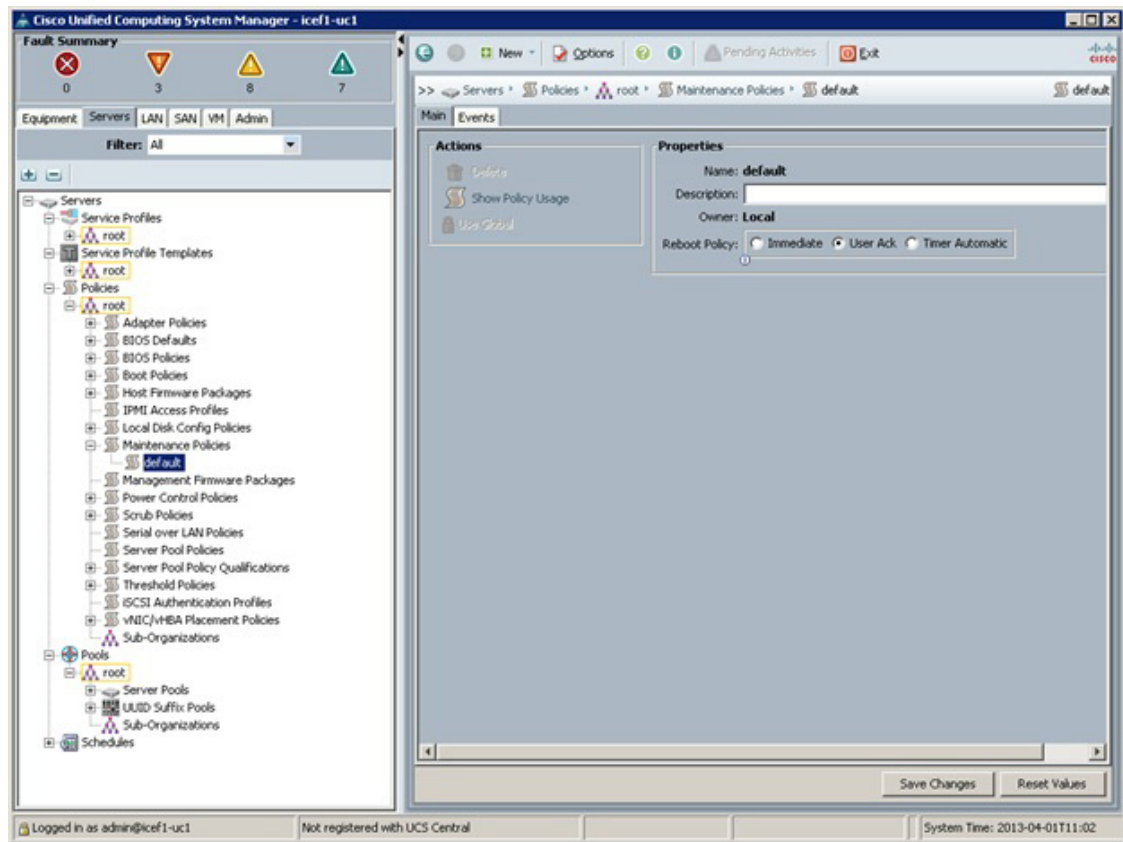
Virtual Slot	Selection Preference
1	Assigned Only
2	All
3	All
4	All

OK Cancel

Update Default Maintenance Policy

To update the default maintenance policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Click Save Changes.
6. Click OK to accept the change.



Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for IB-MGMT-VLAN, NFS-VLAN, Native-VLAN, VM-Traffic-VLAN, and vMotion-VLAN.
11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. From the MAC Pool list, select MAC_Pool_A.

14. From the Network Control Policy list, select Enable_CDP.
15. Click OK to create 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"/>	IB-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:

Dynamic vNIC Connection Policy:

OK Cancel

17. In the navigation pane, select the LAN tab.
18. Select Policies > root.
19. Right-click vNIC Templates.
20. Select Create vNIC Template.
21. Enter vNIC_Template_B as the vNIC template name.
22. Select Fabric B.

23. Do not select the Enable Failover checkbox.
24. Under Target, make sure the VM checkbox is not selected.
25. Select Updating Template as the template type.
26. Under VLANs, select the checkboxes for IB-MGMT-VLAN, NFS-VLAN, Native-VLAN, VM-Traffic-VLAN, and vMotion-VLAN.
27. Set Native-VLAN as the native VLAN.
28. For MTU, enter 9000.
29. From the MAC Pool list, select MAC_Pool_B.
30. From the Network Control Policy list, select Enable_CDP.
31. Click OK to create the vNIC template.
32. 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"/>	IB-MGMT-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

33. Select the LAN tab on the left.
34. Select Policies > root.
35. Right-click vNIC Templates.
36. Select Create vNIC Template.
37. Enter iSCSI_Template_A as the vNIC template name.
38. Leave Fabric A selected. Do not select the Enable Failover checkbox. Under Target, make sure that the VM checkbox is not selected. Select Updating Template for Template Type. Under VLANs, select iSCSI-A-VLAN. Set iSCSI-A-VLAN as the native VLAN. Under MTU, enter 1500. From the MAC Pool list, select MAC_Pool_A. From the Network Control Policy list, select Enable_CDP.
39. Click OK to complete creating the vNIC template.

40. 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"/>	VM-Traffic-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI-A-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	iSCSI-B-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

+ Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

OK Cancel

41. Select the LAN tab on the left.
42. Select Policies > root.
43. Right-click vNIC Templates.
44. Select Create vNIC Template.
45. Enter iSCSI_Template_B as the vNIC template name.
46. Select Fabric B. Do not select the Enable Failover checkbox. Under Target, make sure that the VM checkbox is not selected. Select Updating Template for Template Type. Under VLANs, select iSCSI-B-VLAN. Set iSCSI-B-VLAN as the native VLAN. Under MTU, enter 1500. From the MAC Pool list, select MAC_Pool_B. From the Network Control Policy list, select Enable_CDP.

47. Click OK to complete creating the vNIC template.
48. 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"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI-B-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

OK Cancel

Create Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi lif01a and iscsi lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi lif02a and iscsi lif02b). Also, it is assumed that the "a" LIFs are connected to fabric A (Cisco Nexus A) and the "b" LIFs are connected to fabric B (Cisco Nexus B).

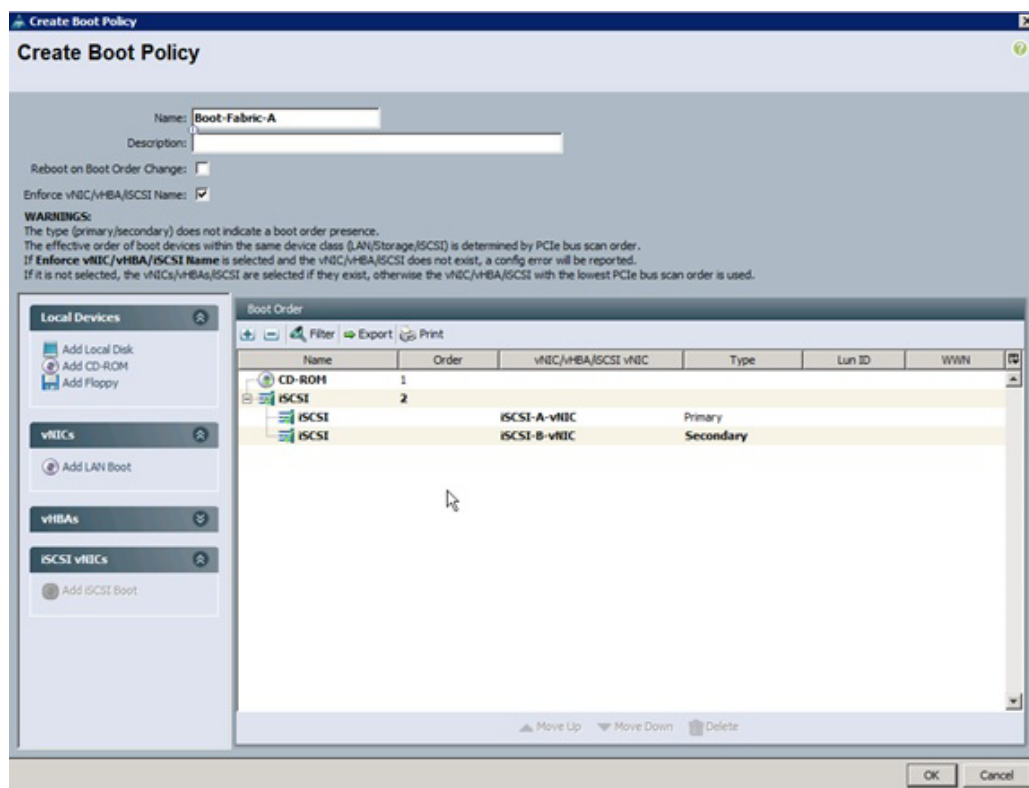
**Note**

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, it is assumed that two iSCSI network interfaces (one in each iSCSI VLAN) are configured on each storage controller.

One boot policy is configured in this procedure. This policy configures the primary target to be `iscsi_lif01a`.

To create boot policies for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter Boot-Fabric-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select Add CD-ROM.
9. Expand the iSCSI vNICs section and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
14. Click OK.
15. Click OK to save the boot policy. Click OK to close the Boot Policy window.



Create Service Profile Templates

In this procedure, one service profile template is created for fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
 - b. Select the Updating Template option.
 - c. Under UUID, select UUID_Pool as the UUID pool.
 - d. Click Next.

6. Configure the networking options:
 - a. Keep the default setting for Dynamic vNIC Connection Policy.
 - b. Select the Expert option to configure the LAN connectivity.
 - c. Click the upper Add button to add a vNIC to the template.
 - d. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
 - e. Select the Use vNIC Template checkbox.
 - f. In the vNIC Template list, select vNIC_Template_A.
 - g. In the Adapter Policy list, select VMWare.
 - h. Click OK to add this vNIC to the template.

Create vNIC

Name:

Use vNIC Template: ☒

+ Create vNIC Template

vNIC Template:

Adapter Performance Profile

Adapter Policy:

+ Create Ethernet Adapter Policy

OK Cancel

7. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.
8. In the Create vNIC box, enter vNIC-B as the name of the vNIC.
9. Select the Use vNIC Template checkbox.
10. In the vNIC Template list, select vNIC_Template_B.
11. In the Adapter Policy list, select VMWare.
12. Click OK to add the vNIC to the template.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Networking

Optionally specify LAN configuration information.

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

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

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

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

Delete Add Modify

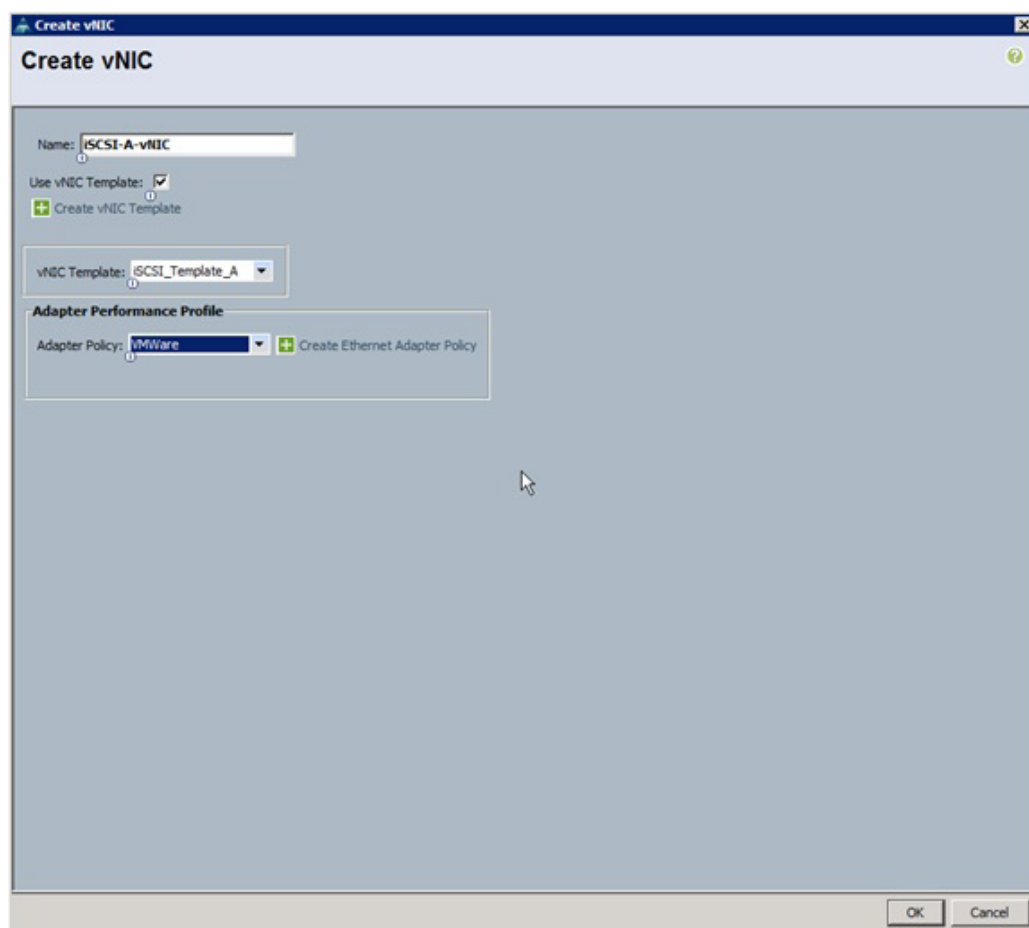
Click **Add** to specify one or more iSCSI vNICs that the server should use.

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
------	-------------------	----------------------	-------------

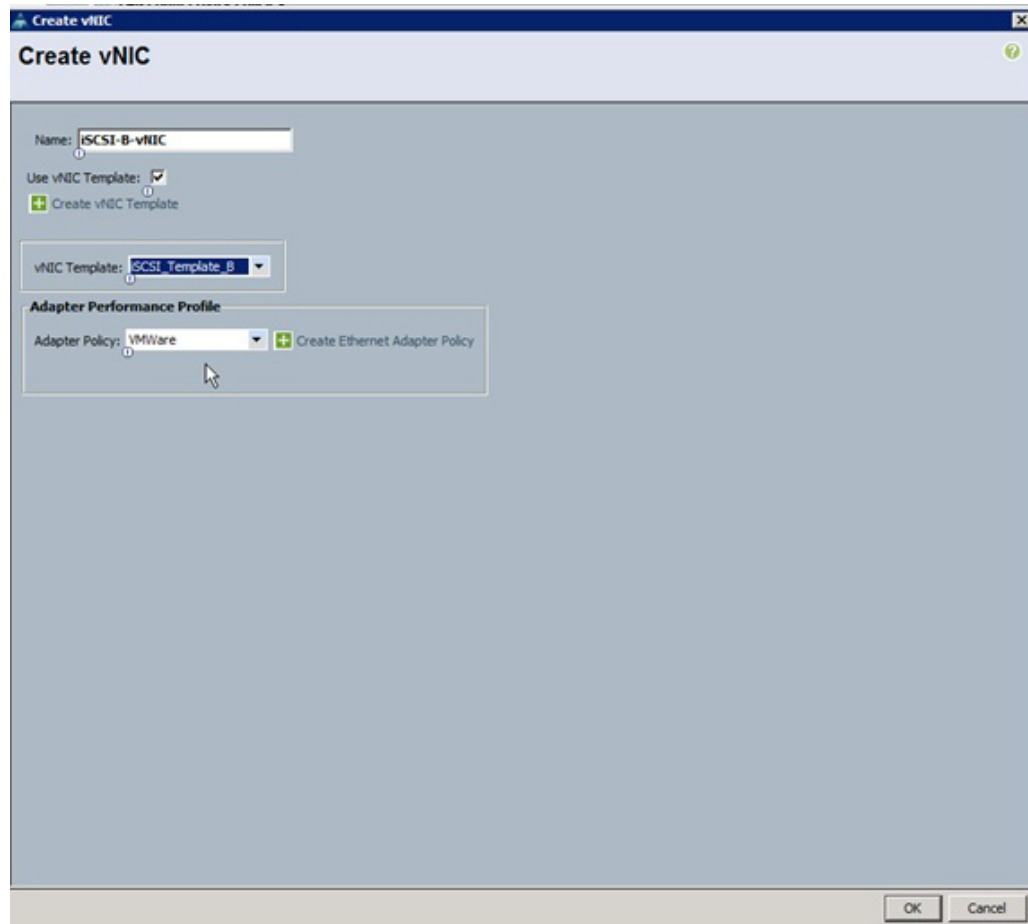
Add Delete Modify

< Prev Next > Finish Cancel

13. Click the upper Add button to add a vNIC to the template.
14. In the Create vNIC dialog box, enter iSCSI-A-vNIC as the name of the vNIC.
15. Select the Use vNIC Template checkbox.
16. In the vNIC Template list, select iSCSI_Template_A.
17. In the Adapter Policy list, select VMWare.
18. Click OK to add this vNIC to the template.



19. Click the upper Add button to add a vNIC to the template.
20. In the Create vNIC dialog box, enter iSCSI-B-vNIC as the name of the vNIC.
21. Select the Use vNIC Template checkbox.
22. In the vNIC Template list, select iSCSI_Template_B.
23. In the Adapter Policy list, select VMWare.
24. Click OK to add this vNIC to the template.



25. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
26. Enter iSCSI-A-vNIC as the name of the vNIC.
27. Select iSCSI-A-vNIC for Overlay vNIC.
28. Set the iSCSI Adapter Policy to default.
29. Set the VLAN to iSCSI-A-VLAN.
30. Leave the MAC Address set to None.
31. Click OK.

Create iSCSI vNIC

Name:

Overlay vNIC:

iSCSI Adapter Policy: [+ Create iSCSI Adapter Policy](#)

VLAN:

iSCSI MAC Address

MAC Address Assignment:

[+ Create MAC Pool](#)

OK Cancel

32. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
33. Enter iSCSI-B-vNIC as the name of the vNIC.
34. Set the Overlay vNIC to iSCSI-B-vNIC
35. Set the iSCSI Adapter Policy to default.
36. Set the VLAN to iSCSI-B-VLAN
37. Leave the MAC Address set to None.
38. Click OK.

Create iSCSI vNIC

Name:

Overlay vNIC:

iSCSI Adapter Policy: [+ Create iSCSI Adapter Policy](#)

VLAN:

iSCSI MAC Address

MAC Address Assignment:

[+ Create MAC Pool](#)

39. Click OK.
40. Review the table in the Networking page to make sure that all vNICs were created.
41. Click Next.

Create LAN Connectivity Policy

Name:

Description:

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

Name	MAC Address	Native VLAN
vNIC iSCSI-B-vNIC	Derived	
vNIC iSCSI-A-vNIC	Derived	
vNIC vNIC-B	Derived	
vNIC vNIC-A	Derived	

Delete Add Modify

Add iSCSI vNICs

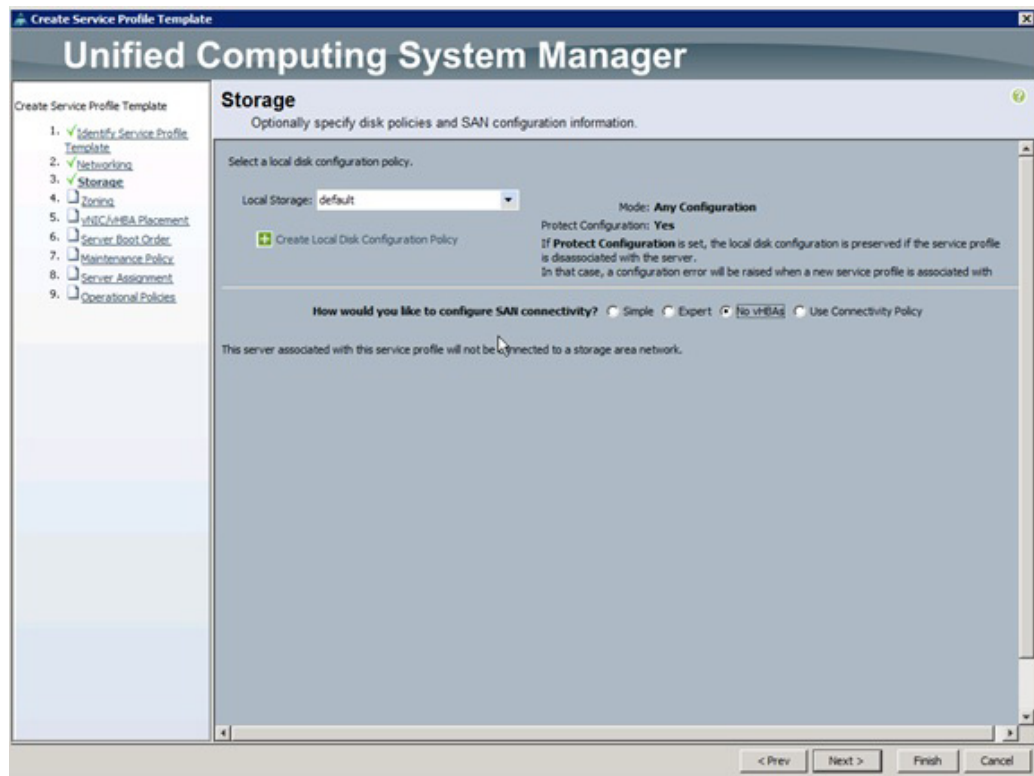
Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-B-vNIC	iSCSI-B-vNIC	default	Derived
iSCSI vNIC iSCSI-A-vNIC	iSCSI-A-vNIC	default	Derived

Add Delete Modify

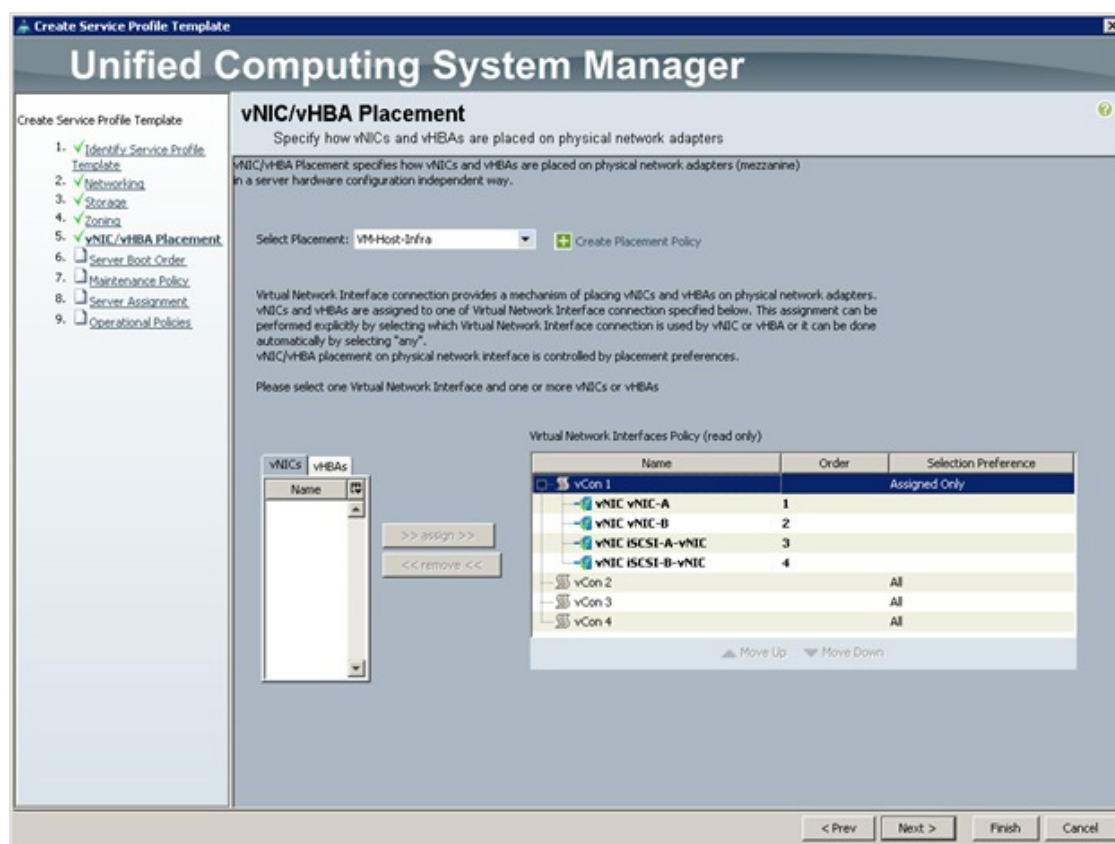
OK Cancel

42. Configure the storage options:

- a. Select a local disk configuration policy:
 - If the server in question has local disks, select default in the Local Storage list.
 - If the server in question does not have local disks, select SAN-Boot.
- b. Select the No vHBAs option for the How would you like to configure SAN connectivity? field.
- c. Click Next.



43. Set no Zoning options and click Next.
44. Set the vNIC/vHBA placement options.
 - a. In the Select Placement list, select the VM-Host-Infra placement policy.
 - b. Select vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - vNIC-A
 - vNIC-B
 - iSCSI-vNIC-A
 - iSCSI-vNIC-B
 - c. Review the table to verify that all vNICs and vHBAs were assigned to the policy in the appropriate order.
 - d. Click Next.



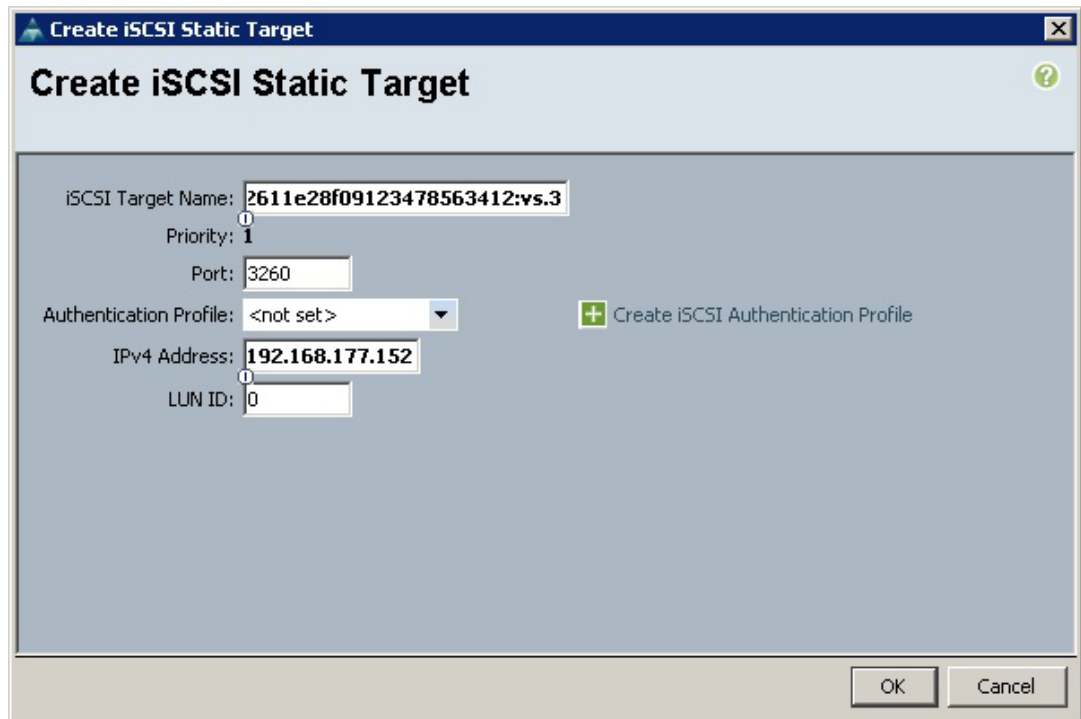
45. Set the server boot order:



Note

If your FlexPod implementation uses NetApp 7-Mode storage instead of clustered Data ONTAP, replace this section with the section "Alternate Cisco UCS Configuration with 7-Mode Storage" from the appendix.

- a. Select Boot-Fabric-A for Boot Policy.
- b. In the Boot Order pane, select iSCSI-A-vNIC.
- c. Click the Set iSCSI Boot Parameters button.
- d. In the Set iSCSI Boot Parameters dialog box, set the initiator name assignment to IQN_Pool_A.
- e. In the Set iSCSI Boot Parameters dialog box, set iSCSI_IP_Pool_A as the initiator IP address policy.
- f. Keep the iSCSI Static Target Interface button selected and click the button.
- g. Log in to the storage cluster management interface and run the following command:
iscsi nodename
- h. Note or copy the iSCSI target name for Infra_Vserver.
- i. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_Vserver into the iSCSI Target Name field.
- j. Enter the IP address of iscsi_lif02a for the IPv4 Address field.




The image shows a 'Create iSCSI Static Target' dialog box. It has a title bar with a green icon and a close button. The main area contains several input fields: 'iSCSI Target Name' with the value '2611e28f09123478563412:vs.3', 'Priority' with the value '1', 'Port' with the value '3260', 'Authentication Profile' with a dropdown menu showing '<not set>', 'IPv4 Address' with the value '192.168.177.152', and 'LUN ID' with the value '0'. There is a green plus icon and the text 'Create iSCSI Authentication Profile' next to the 'Authentication Profile' dropdown. At the bottom right, there are 'OK' and 'Cancel' buttons.

Create iSCSI Static Target

iSCSI Target Name: 2611e28f09123478563412:vs.3

Priority: 1

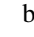
Port: 3260

Authentication Profile: <not set>  Create iSCSI Authentication Profile

IPv4 Address: 192.168.177.152

LUN ID: 0

OK Cancel

- k. Click OK to add the iSCSI static target.
- l. Keep the iSCSI Static Target Interface option selected and click the  button.
- m. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra_Vserver into the iSCSI Target Name field.
- n. Enter the IP address of iscsi_lif01a in the IPv4 Address field.
- o. Click OK.

Set iSCSI Boot Parameters

Name: **iSCSI-A-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: IQN_Pool_A(16/16)

Initiator Name:

+ Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_A(16/16)

IPv4 Address: 0.0.0.0
Subnet Mask: 255.255.255.0
Default Gateway: 0.0.0.0
Primary DNS: 0.0.0.0
Secondary DNS: 0.0.0.0

+ Create IP Pool

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

☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id
iqn.1992-08.c...	2	3260		192.168.177.151	0
iqn.1992-08.c...	1	3260		192.168.177.152	0

OK Cancel

- p. Click OK.
- q. In the Boot Order pane, select iSCSI-vNIC-B.
- r. Click the Set iSCSI Boot Parameters button.
- s. In the Set iSCSI Boot Parameters dialog box, set the set the initiator name assignment to IQN_Pool_B.
- t. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.

- u. Keep the iSCSI Static Target Interface option selected and click the + button.
- v. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra_Vserver into the iSCSI Target Name field (same target name as above).
- w. Enter the IP address of iscsi_lif02b in the IPv4 address field.

Create iSCSI Static Target

iSCSI Target Name: 2611e28f09123478563412:vs.3

Priority: 1

Port: 3260

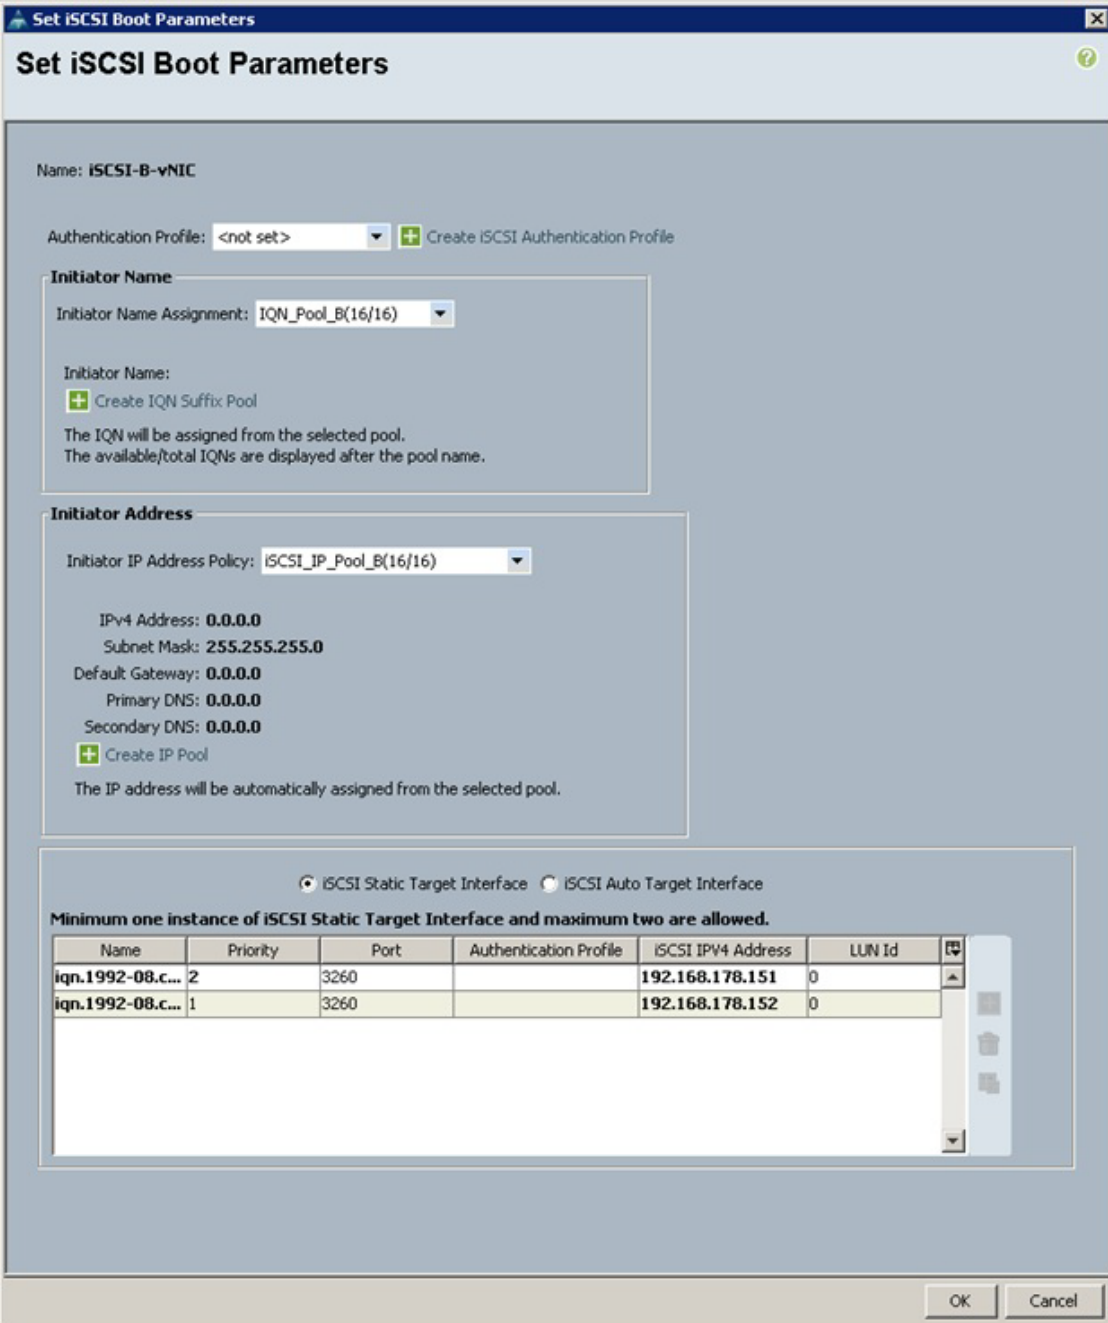
Authentication Profile: <not set> + Create iSCSI Authentication Profile

IPv4 Address: 192.168.178.152

LUN ID: 0

OK Cancel

- 46. Click OK to add the iSCSI static target.
- 47. Keep the iSCSI Static Target Interface option selected and click the + button.
- 48. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_Vserver into the iSCSI Target Name field.
- 49. Enter the IP address of iscsi_lif01b in the IPv4 Address field.
- 50. Click OK.



Set iSCSI Boot Parameters

Name: **iSCSI-B-vNIC**

Authentication Profile: **<not set>** + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: **IQN_Pool_B(16/16)**

Initiator Name:

+ Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: **iSCSI_IP_Pool_B(16/16)**

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

+ Create IP Pool

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

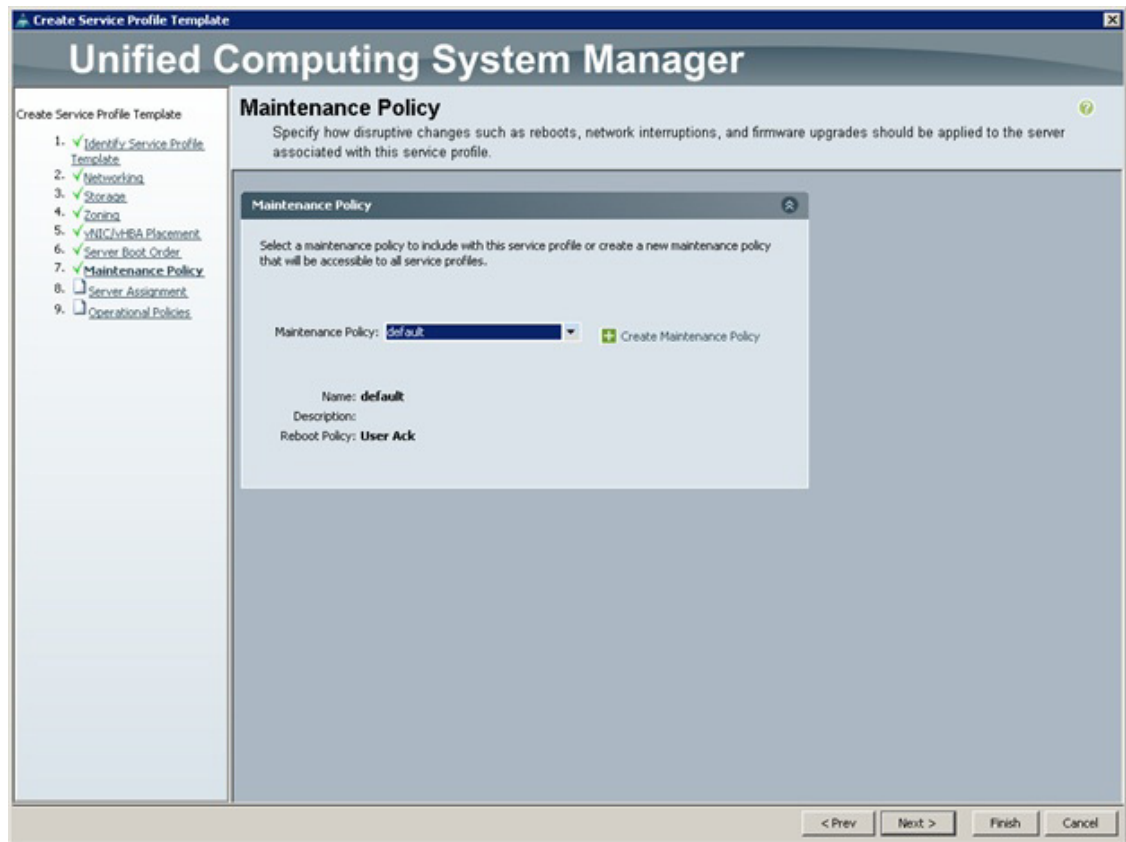
☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

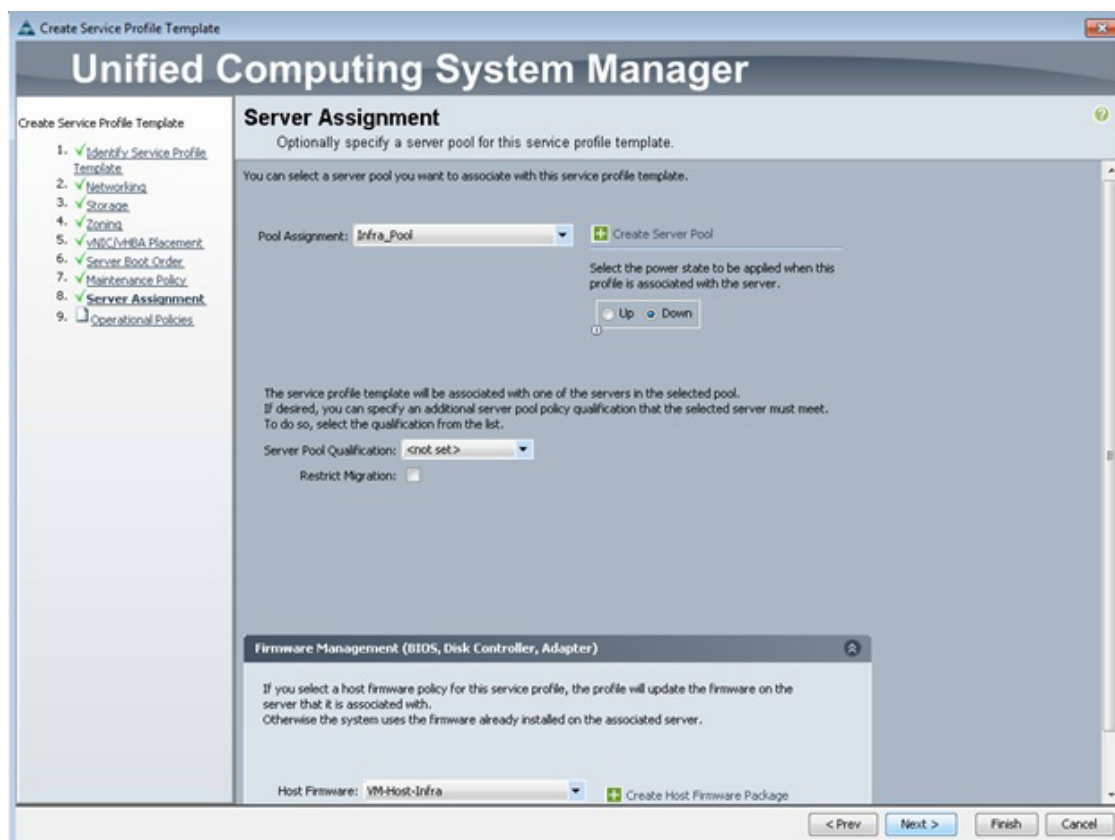
Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id
iqn.1992-08.c...	2	3260		192.168.178.151	0
iqn.1992-08.c...	1	3260		192.168.178.152	0

OK Cancel

- x. Click OK.
 - y. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
 - z. Click Next to continue to the next section.
51. Add a maintenance policy:
- a. Select the default Maintenance Policy.
 - b. Click Next.

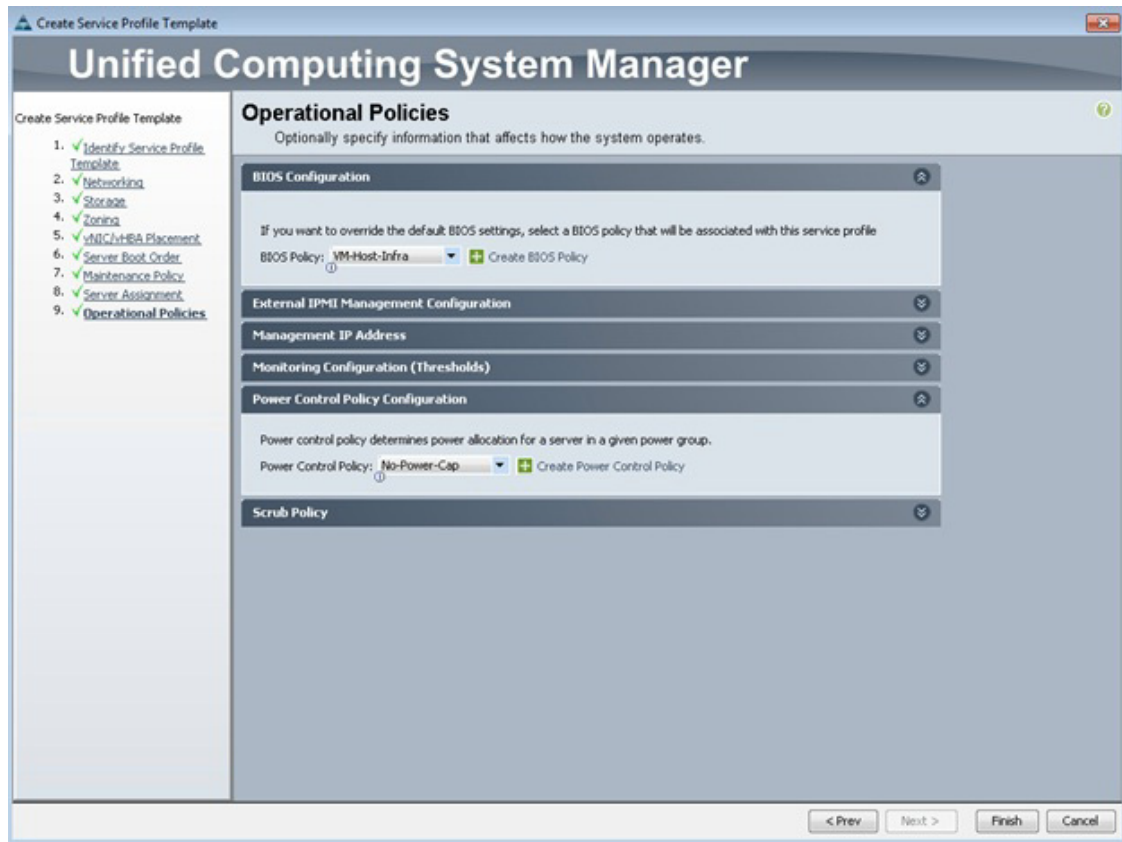


52. Specify the server assignment:
 - a. In the Pool Assignment list, select Infra_Pool.
 - b. Optional: Select a Server Pool Qualification policy.
 - c. Select Down as the power state to be applied when the profile is associated with the server.
 - d. Expand Firmware Management at the bottom of the page and select VM-Host-Infra from the Host Firmware list.
 - e. Click Next.



53. Add operational policies:

- a. In the BIOS Policy list, select VM-Host-Infra.
- b. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

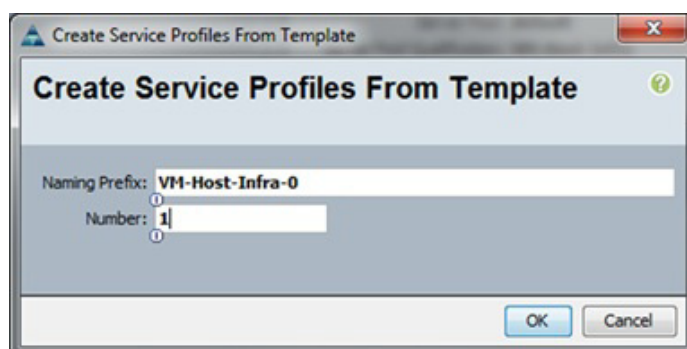


54. Click Finish to create the service profile template.
55. Click OK in the confirmation message.

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-Fabric-A.
3. Right-click VM-Host-Infra-Fabric-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 2 as the number of service profiles to create.
6. Click OK to create the service profile.



7. Click OK in the confirmation message.
8. Verify that the service profiles VM-Host-Infra-01 and VM-Host-Infra-02 have been created. The service profiles are automatically associated with the servers in their assigned server pools.
9. Optional: Select each newly created service profile and enter the server host name or the fully qualified domain name (FQDN) in the User Label field in the General tab. Click Save Changes to map the server host name to the service profile name.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created 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, each infrastructure blade in the environment will 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 into Table 20 and Table 21.

Table 20 *ISC SI LIFs for iSCSI IQN*

Vserver	iSCSI Target IQN
Infra_Vserver	

Table 21 *vNIC iSC SI IQNs for fabric A and fabric B*

Cisco UCS Service Profile Name	Fabric A iSCSI IQN	Fabric B iSCSI IQN
VM-Host-Infra-01		
VM-Host-Infra-02		

To gather the vNIC IQN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Click each service profile and then click the Boot Order tab on the right. Expand iSCSI and select each iSCSI vNIC. Click Set iSCSI Boot Parameters. In Table 21, record the IQN information that is displayed in the right pane for each iSCSI vNIC in each service profile in the table above.

Networking Configuration

The following section provides the detailed procedure for configuring the Cisco Nexus 7000 4-Slot 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 provide guidance for configuring the Cisco Nexus 7000 running release 6.1(2). 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.

Cisco Nexus 7000 Network Aggregation Configuration

Initial Setup of the Cisco Nexus 7000 Switch

These steps provide details for the initial Cisco Nexus 7000 switch setup.

Cisco Nexus A

To set up the initial configuration for the first Cisco Nexus, complete the following steps:

1. Configure the switch.



Note

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

```
Abort Power On Auto Provisioning and continue with normal setup ?(yes/no) [n]: y
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
```

```
---- Basic System Configuration Dialog VDC: 1 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

```
Would you like to enter the basic configuration dialog (yes/no): yes
```

```
Create another login account (yes/no) [n]: Enter
```

```
Configure read-only SNMP community string (yes/no) [n]: Enter
```

```
Configure read-write SNMP community string (yes/no) [n]: Enter
```

```

Enter the switch name : <<var_nexus_A_hostname>>

Enable license grace period? (yes/no) [n]:  Enter

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Enter

Mgmt0 IPv4 address : <<var_nexus_A_mgmt0_ip>>

Mgmt0 IPv4 netmask : <<var_nexus_A_mgmt0_netmask>>

Configure the default gateway? (yes/no) [y]:  Enter

IPv4 address of the default gateway : <<var_nexus_A_mgmt0_gw>>

Configure advanced IP options? (yes/no) [n]:  Enter

Enable the telnet service? (yes/no) [n]:  Enter

Enable the ssh service? (yes/no) [y]:  Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]:  Enter

Number of rsa key bits <1024-2048> [1024]:  Enter

Configure the ntp server? (yes/no) [n]:  Enter

Configure default interface layer (L3/L2) [L3]: L2

Configure default switchport interface state (shut/noshut) [shut]:  Enter

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
Enter

The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_hostname>>
no license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
no system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address <<var_nexus_A_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown

Would you like to edit the configuration? (yes/no) [n]:  Enter

Use this configuration and save it? (yes/no) [y]:  Enter

Disabling ssh: as its enabled right now:
generating rsa key(1024 bits).....
.

```

```
generated rsa key
Enabling ssh: as it has been disabled
% All 0s mask is invalid
```

2. Review the configuration summary before enabling the configuration.

```
Would you like to save the running-config to startup-config? (yes/no) [n]: y
```

```
[#####] 100%
Copy complete.
```

Cisco Nexus B

To set up the initial configuration for the second Cisco Nexus switch complete the following steps:

1. Configure the switch.



Note

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

```
Abort Power On Auto Provisioning and continue with normal setup ?(yes/no) [n]: y
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
```

```
---- Basic System Configuration Dialog VDC: 1 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

```
Would you like to enter the basic configuration dialog (yes/no): yes
```

```
Create another login account (yes/no) [n]: Enter
```

```
Configure read-only SNMP community string (yes/no) [n]: Enter
```

```
Configure read-write SNMP community string (yes/no) [n]: Enter
```

```
Enter the switch name : <<var_nexus_B_hostname>>
```

```
Enable license grace period? (yes/no) [n]: Enter
```

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Enter
```

```
Mgmt0 IPv4 address : <<var_nexus_B_mgmt0_ip>>
```

```
Mgmt0 IPv4 netmask : <<var_nexus_B_mgmt0_netmask>>
```

```
Configure the default gateway? (yes/no) [y]: Enter
```



```

IPv4 address of the default gateway : <<var_nexus_B_mgmt0_gw>>

Configure advanced IP options? (yes/no) [n]:  Enter

Enable the telnet service? (yes/no) [n]:  Enter

Enable the ssh service? (yes/no) [y]:  Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]:  Enter

Number of rsa key bits <1024-2048> [1024]:  Enter

Configure the ntp server? (yes/no) [n]:  Enter

Configure default interface layer (L3/L2) [L3]: L2

Configure default switchport interface state (shut/noshut) [shut]:  Enter

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
Enter

The following configuration will be applied:
password strength-check
switchname <<var_nexus_B_hostname>>
no license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
no system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address <<var_nexus_B_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown

Would you like to edit the configuration? (yes/no) [n]:  Enter

Use this configuration and save it? (yes/no) [y]:  Enter

Disabling ssh: as its enabled right now:
generating rsa key(1024 bits).....
.
generated rsa key
Enabling ssh: as it has been disabled
% All 0s mask is invalid
2. Review the configuration summary before enabling the configuration.

Would you like to save the running-config to startup-config? (yes/no) [n]: y

[#####] 100%
Copy complete.

```

Enable Appropriate Cisco Nexus Features

Cisco Nexus A

To license the first Cisco Nexus switch, complete the following steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature udd
feature interface-lanvpc
feature lacp
feature vpc
```

Cisco Nexus B

To license the second Cisco Nexus switch, complete the following steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature udd
feature interface-lanvpc
feature lacp
feature vpc
```

Set Global Configurations

Cisco Nexus A

To set up the global configurations for the first Cisco Nexus switch, complete the following step:

1. From the global configuration mode, run the following commands:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
system jumbomtu 9216
copy run start
```

Cisco Nexus B

To set up the global configurations for the second Cisco Nexus B, complete the following step:

1. From the global configuration mode, run the following commands:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
system jumbomtu 9216
copy run start
```

Create VLANs

Both Cisco Nexus Switches

To create the necessary virtual local area networks (VLANs), complete the following step on both switches:

From the global configuration mode, run the following commands:

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
exit
```

```

vlan <<var_native_vlan_id>>
name Native-VLAN
exit
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_pkt_ctrl_vlan_id>>
name Packet-Control-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_vm_traffic_vlan_id>>
name VM-Traffic-VLAN
exit
vlan <<var_isCSI_A_vlan_id>>
name iSCSI-A-VLAN
exit
vlan <<var_isCSI_B_vlan_id>>
name iSCSI-B-VLAN
exit

```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```

interface Eth3/1
description <<var_node01>>:e3a
exit
interface Eth3/2
description <<var_node02>>:e3a
exit
interface Eth3/11
description VPC Peer <<var_nexus_B_hostname>>:3/11
exit
interface Eth3/12
description VPC Peer <<var_nexus_B_hostname>>:3/12
exit
interface eth3/23
description <<var_ucs_clustername>>-A:1/19
exit
interface eth3/24
description <<var_ucs_clustername>>-B:1/20
exit

```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B, complete the following steps:

1. From the global configuration mode, run the following commands:

```

interface Eth3/1
description <<var_node01>>:e4a
exit
interface Eth3/2

```

```

description <<var_node02>>:e4a
exit
interface Eth3/11
description VPC Peer <<var_nexus_A_hostname>>:3/11
exit
interface Eth3/12
description VPC Peer <<var_nexus_B_hostname>>:3/12
exit
interface eth3/23
description <<var_ucs_clustername>>-A:1/19
exit
interface eth3/24
description <<var_ucs_clustername>>-B:1/20
exit

```

Create Port Channels

Both Cisco Nexus Switches

To create the necessary port channels between devices, complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```

interface Po10
description vPC peer-link
exit
interface Eth3/11-12
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth3/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth3/2
channel-group 12 mode active
no shutdown
exit
interface Po13
description <<var_ucs_clustername>>-A
exit
interface Eth3/23
channel-group 13 mode active
no shutdown
exit
interface Po14
description <<var_ucs_clustername>>-B
exit
interface Eth3/24
channel-group 14 mode active
no shutdown
exit

```

```
copy run start
```

Configure Port Channels

Both Cisco Nexus Switches

To configure the port channels, complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_pkt_ctrl_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm_traffic_vlan_id>>,
<<var_iscsi_a_vlan_id>>, <<var_iscsi_b_vlan_id>>
spanning-tree port type network
no shutdown
exit
interface Po11
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>, <<var_iscsi_a_vlan_id>>,
<<var_iscsi_b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
no shutdown
exit
interface Po12
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>, <<var_iscsi_a_vlan_id>>,
<<var_iscsi_b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
no shutdown
exit
interface Po13
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm_traffic_vlan_id>>, <<var_iscsi_a_vlan_id>>,
<<var_iscsi_b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
no shutdown
exit
interface Po14
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm_traffic_vlan_id>>, <<var_iscsi_a_vlan_id>>,
<<var_iscsi_b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
no shutdown
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source
<<var_nexus_A_mgmt0_ip>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, complete the following step:

1. From the global configuration mode, run the following commands.

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source
<<var_nexus_B_mgmt0_ip>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
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 or use the Cisco Nexus 7000 switch of the FlexPod environment as your distribution block. If an existing Cisco Nexus environment is present, Cisco recommends using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the existing infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after configuration is completed.

Storage Part 2

Clustered Data ONTAP SAN Boot Storage Setup



Note

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, replace this section with the section "Alternate 7-Mode NetApp FAS3250 Deployment Procedure: Part 2" from the appendix.

Create Igroups

1. From the cluster management node SSH connection, enter the following:

```
igroup create -vserver Infra_Vserver -igroup VM-Host-Infra-01 -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_01_iqn_A>>,
<<var_vm_host_infra_01_iqn_B>>
```

```
igroup create -vserver Infra_Vserver -igroup VM-Host-Infra-02 -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_02_iqn_A>>,
<<var_vm_host_infra_02_iqn_B>>
```

```
igroup create -vserver Infra_Vserver -igroup MGMT-Hosts -protocol iscsi -ostype
vmware -initiator <<var_vm_host_infra_01_iqn_A>>,
<<var_vm_host_infra_01_iqn_B>>, <<var_vm_host_infra_02_iqn_A>>,
<<var_vm_host_infra_02_iqn_B>>
```



Note

To view the three igroups just created, type `igroup show`.

Map Boot LUNs to Igroups

1. From the cluster management SSH connection, enter the following:

```
lun map -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -igroup
VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-02 -igroup
VM-Host-Infra-02 -lun-id 0
```

VMware vSphere 5.1 Setup

FlexPod VMware ESXi 5.1 iSCSI on Clustered Data ONTAP

This section provides detailed instructions for installing VMware ESXi 5.1 in a FlexPod environment. After the procedures are completed, two iSCSI-booted ESXi hosts will be provisioned. These deployment procedures are customized to include the environment variables.



Note

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in Keyboard, Video, Mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their iSCSI boot logical unit numbers (LUNs).

Log in to Cisco UCS 6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Log in to Cisco UCS Manager by using the admin user name and password.
3. From the main menu, click the Servers tab.
4. Select Servers > Service Profiles > root > VM-Host-Infra-01.
5. Right-click VM-Host-Infra-01 and select KVM Console.
6. Select Servers > Service Profiles > root > VM-Host-Infra-02.
7. Right-click VM-Host-Infra-02 and select KVM Console Actions > KVM Console.

Set Up VMware ESXi Installation

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

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click the Virtual Media tab.
2. Click Add Image.
3. Browse to the ESXi installer ISO image file and click Open.
4. Select the Mapped checkbox to map the newly added image.
5. Click the KVM tab to monitor the server boot.
6. Boot the server by selecting Boot Server and clicking OK. Then click OK again.

Install ESXi

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

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the NetApp LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.
8. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.



Note

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, simply click Yes to unmap the image.
10. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi Host VM-Host-Infra-01

To configure the VM-Host-Infra-01 ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the <<var_ib-mgmt_vlan_id>> and press Enter.
6. From the Configure Management Network menu, select IP Configuration and press Enter.
7. Select the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the first ESXi host: <<var_vm_host_infra_01_ip>>.
9. Enter the subnet mask for the first ESXi host.

10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Select the IPv6 Configuration option and press Enter.
13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
14. Select the DNS Configuration option and press Enter.

**Note**

Because the IP address is assigned manually, the DNS information must also be entered manually.

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

ESXi Host VM-Host-Infra-02

To configure the VM-Host-Infra-02 ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the <<var_ib-mgmt_vlan_id>> and press Enter.
6. From the Configure Management Network menu, select IP Configuration and press Enter.
7. Select the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the second ESXi host: <<var_vm_host_infra_02_ip>>.
9. Enter the subnet mask for the second ESXi host.
10. Enter the default gateway for the second ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Select the IPv6 Configuration option and press Enter.
13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
14. Select the DNS Configuration option and press Enter.

**Note**

Because the IP address is assigned manually, the DNS information must also be entered manually.

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the FQDN for the second ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

Download VMware vSphere Client and vSphere Remote CLI

To download the VMware vSphere Client and install Remote CLI, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.
2. Download and install both the vSphere Client and the Windows version of vSphere Remote Command Line.

These applications are downloaded from the VMware website and Internet access is required on the management workstation.

Log in to VMware ESXi Hosts Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

To log in to the VM-Host-Infra-01 ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Infra-01 as the host you are trying to connect to: <<var_vm_host_infra_01_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

ESXi Host VM-Host-Infra-02

To log in to the VM-Host-Infra-02 ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Infra-02 as the host you are trying to connect to: <<var_vm_host_infra_02_ip>>.
2. Enter root for the user name.

3. Enter the root password.
4. Click Login to connect.

Download Updated Cisco VIC enic Driver

To download the Cisco virtual interface card (VIC) enic driver, complete the following steps:

The enic version used in this configuration is 2.1.2.38.

1. Open a Web browser on the management workstation and navigate to [http://software.cisco.com/download/release.html?mdfid=283853163&softwareid=283853158&release=2.0\(5\)&reind=AVAILABLE&rellifecycle=&reltype=latest](http://software.cisco.com/download/release.html?mdfid=283853163&softwareid=283853158&release=2.0(5)&reind=AVAILABLE&rellifecycle=&reltype=latest).
2. Login and select the driver ISO for version 2.1(1b). Download the ISO file.
3. When the ISO file is downloaded, either burn the ISO to a CD or map the ISO to a drive letter. Extract the following files from within the VMware directory for ESXi 5.1:
 - Network - net-enic-2.1.2.38-1OEM.500.0.0.472560.x86_64.zip
4. Document the saved location.

Load Updated Cisco VIC enic Drivers

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

To load the updated versions of the enic driver for the Cisco VIC, complete the following steps for the hosts on each vSphere Client:

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.
4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded enic driver version and select net-enic-2.1.2.38-1OEM.500.0.0.472560.x86_64.zip.
6. Click Open to open the file.
7. Click Yes to upload the .zip file to datastore1.
8. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
9. At the command prompt, run the following commands to account for each host (enic):

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/net-enic-2.1.2.38-1OEM.500.0.0.472560.x86_64.zip
```

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/net-enic-2.1.2.38-1OEM.500.0.0.472560.x86_64.zip
```

```

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp!23 software vib install --no-sig-check -d /vmfs/volumes/datastore1/net-enic-2.1.2.38-10EM.500.0.0.472560.x86_64.zip
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_2.1.2.38-10EM.500.0.0.472560
  VIBs Removed: VMware_bootbank_net-enic_1.4.2.15a-1vmw.510.0.0.799733
  VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.59 -u root -p NetApp!23 software vib install --no-sig-check -d /vmfs/volumes/datastore1/net-enic-2.1.2.38-10EM.500.0.0.472560.x86_64.zip
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_2.1.2.38-10EM.500.0.0.472560
  VIBs Removed: VMware_bootbank_net-enic_1.4.2.15a-1vmw.510.0.0.799733
  VIBs Skipped:

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

```

10. From the vSphere Client, right-click each host in the inventory and select Reboot.
11. Select Yes to continue.
12. Enter a reason for the reboot and click OK.
13. After the reboot is complete, log back in to both hosts by using the vSphere Client.

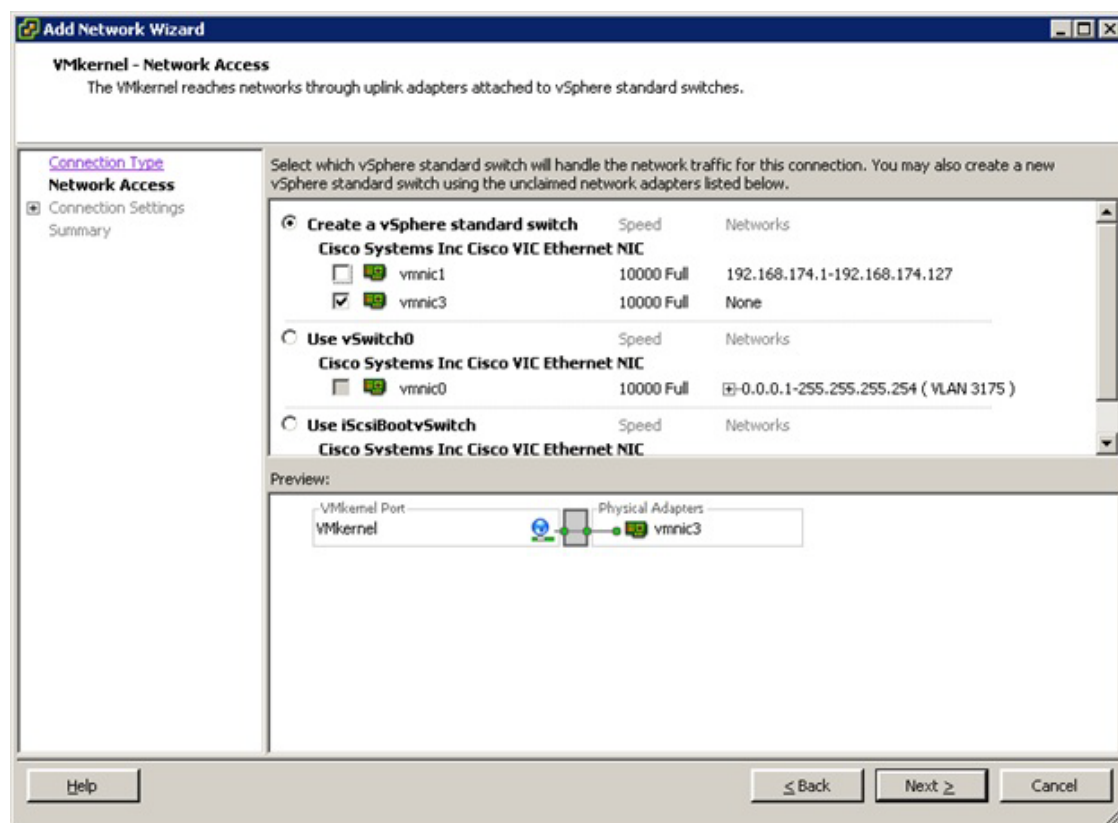
Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

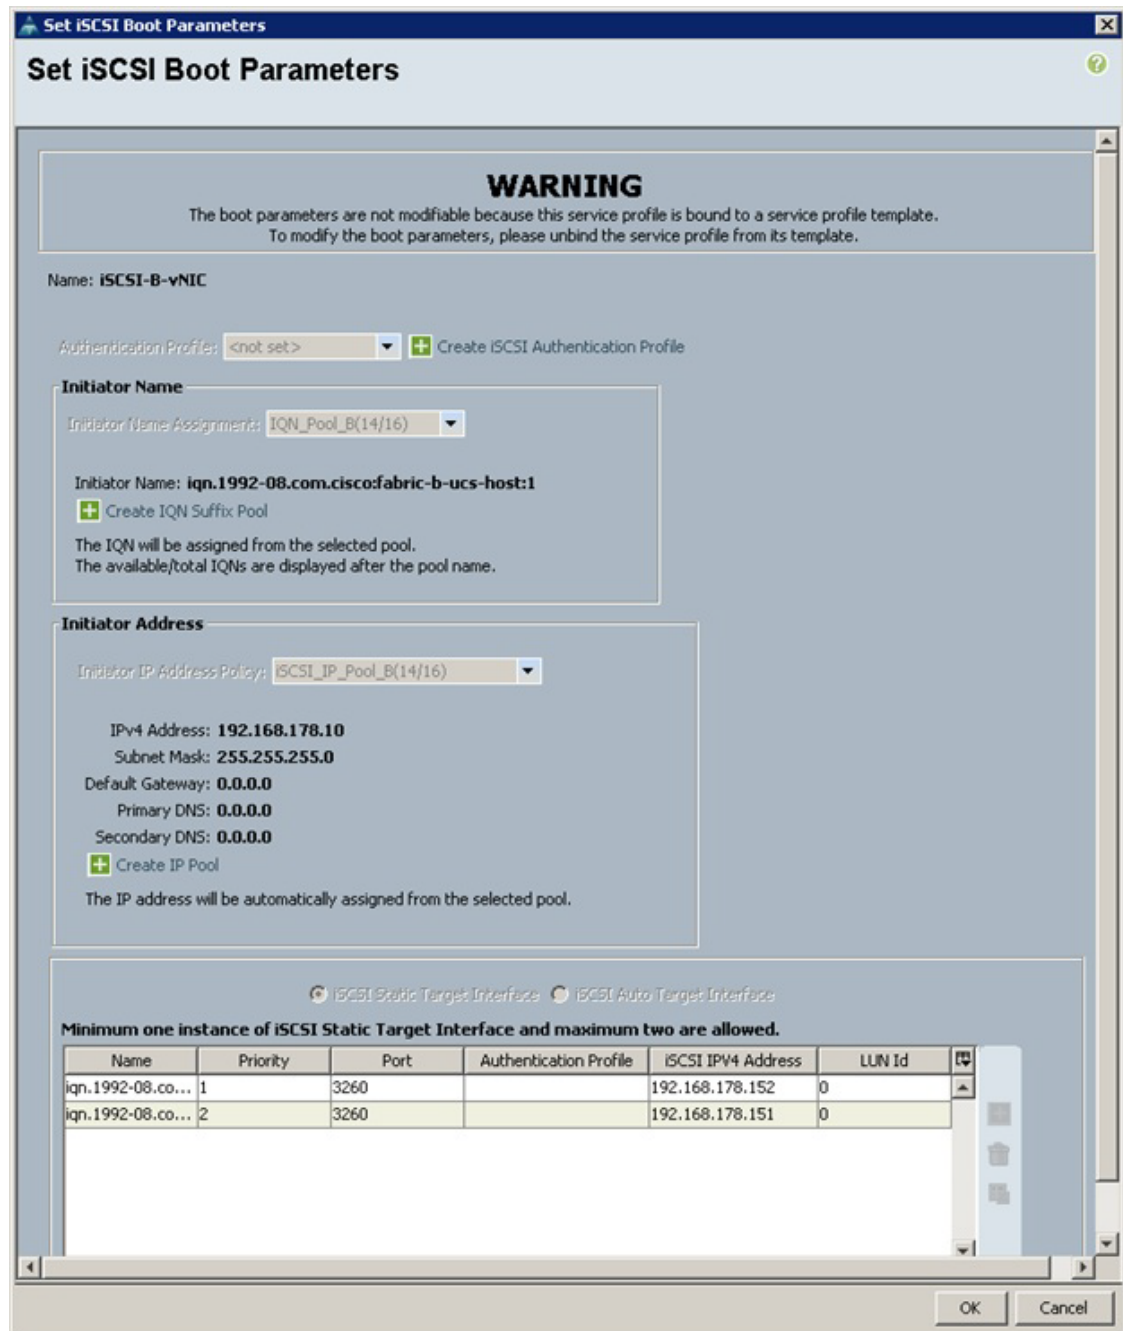
To set up the VMkernel ports and the virtual switches on the VM-Host-Infra-01 ESXi host, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Networking in the Hardware pane.
4. Click Properties on the right side of vSwitch0.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to 9000.
7. Click OK to close the properties for vSwitch0.
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 to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to IB-MGMT Network and enter <<var_ib-mgmt_vlan_id>> in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Add to add a network element.
15. Select VMkernel and click Next.

16. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
17. Click Next to continue with the NFS VMkernel creation.
18. Enter the IP address <<var_nfs_vlan_id_ip_host-01>> and the subnet mask <<var_nfs_vlan_id_mask_host01>> for the NFS VLAN interface for VM-Host-Infra-01.
19. Click Next to continue with the NFS VMkernel creation.
20. Click Finish to finalize the creation of the NFS VMkernel interface.
21. Select the VMkernel-NFS configuration and click Edit.
22. Change the MTU to 9000.
23. Click OK to finalize the edits for the VMkernel-NFS network.
24. Click Add to add a network element.
25. Select VMkernel and click Next.
26. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.
27. Select the Use This Port Group for vMotion checkbox.
28. Click Next to continue with the vMotion VMkernel creation.
29. Enter the IP address <<var_vmotion_vlan_id_ip_host-01>> and the subnet mask <<var_vmotion_vlan_id_mask_host-01>> for the vMotion VLAN interface for VM-Host-Infra-01.
30. Click Next to continue with the vMotion VMkernel creation.
31. Click Finish to finalize the creation of the vMotion VMkernel interface.
32. Select the VMkernel-vMotion configuration and click Edit.
33. Change the MTU to 9000.
34. Click OK to finalize the edits for the VMkernel-vMotion network.
35. Close the dialog box to finalize the ESXi host networking setup.
36. Click Properties on the right side of iScsiBootvSwitch.
37. Select the iScsiBootPG configuration and click Edit.
38. Change the Network Label to VMkernel-iSCSI-A. Do not set a VLAN ID.
39. Click OK to save changes to the VMkernel port.
40. Click Close to close the vSwitch Properties window.
41. On the right, click Add Networking.
42. Select VMkernel and click Next.
43. Leave Create a vSphere standard switch selected. Clear vmnic1 and select vmnic3. Click Next.



44. Change the Network Label to VMkernel-iSCSI-B. Leave the VLAN ID set to None.
45. Click Next.
46. Set the VMkernel-iSCSI-B IP address and subnet mask. To get this information, select the VM-Host-Infra-01 Service Profile in Cisco UCS Manager. Select the Boot Order tab and select iSCSI-B-vNIC. Click Set iSCSI Boot Parameters. Obtain the IPv4 address and subnet mask from this window.
47. Click Next.



Set iSCSI Boot Parameters

WARNING
The boot parameters are not modifiable because this service profile is bound to a service profile template. To modify the boot parameters, please unbind the service profile from its template.

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name
Initiator Name Assignment: IQN_Pool_B(14/16)
Initiator Name: **iqn.1992-08.com.cisco:fabric-b-ucs-host:1**
+ Create IQN Suffix Pool
The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

Initiator Address
Initiator IP Address Policy: iSCSI_IP_Pool_B(14/16)
IPv4 Address: **192.168.178.10**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**
+ Create IP Pool
The IP address will be automatically assigned from the selected pool.

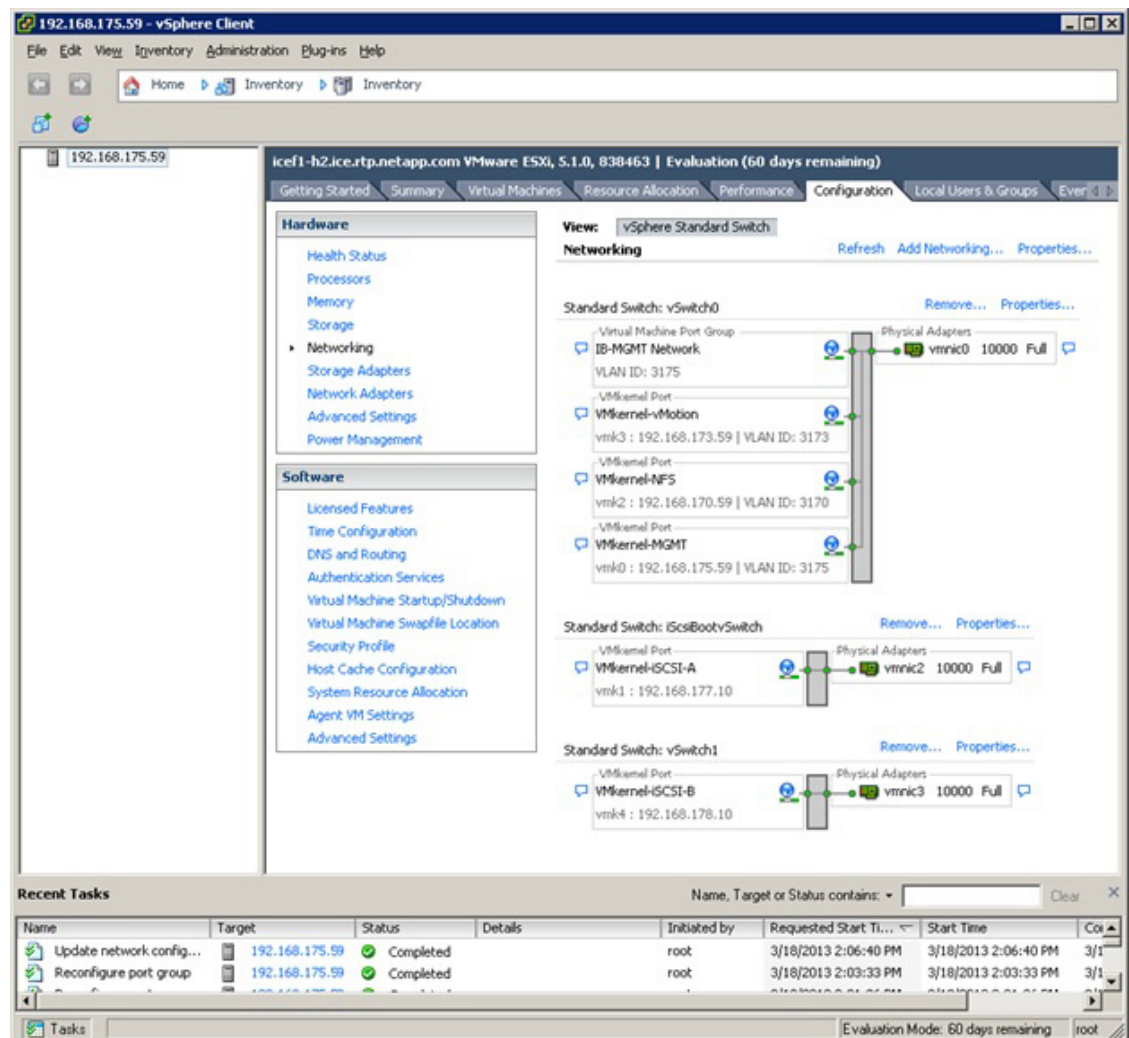
☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id
iqn.1992-08.co...	1	3260		192.168.178.152	0
iqn.1992-08.co...	2	3260		192.168.178.151	0

OK Cancel

48. Click Finish. vSwitch 1 is created.
49. The networking for the ESXi host should be similar to the following example:



50. Click Storage Adapters in the Hardware pane.
51. Select the iSCSI Software Adapter and click Properties in the Details pane.
52. Select the Network Configuration tab.
53. Click Add.
54. Select VMkernel-iSCSI-A and click OK.
55. Click Add.
56. Select VMkernel-iSCSI-B and click OK.
57. Select the Static Discovery tab.
58. Click Settings.
59. Select the entire iSCSI Target Name field, right-click, and select Copy to copy this target name to the clipboard.
60. Click Close to close the Static Target Server Settings window.
61. Click Add.
62. Right-click the iSCSI Target Name field and select Paste.

63. Type the IP address of one of the iSCSI LIFs that does not already appear in the list.

**Note**

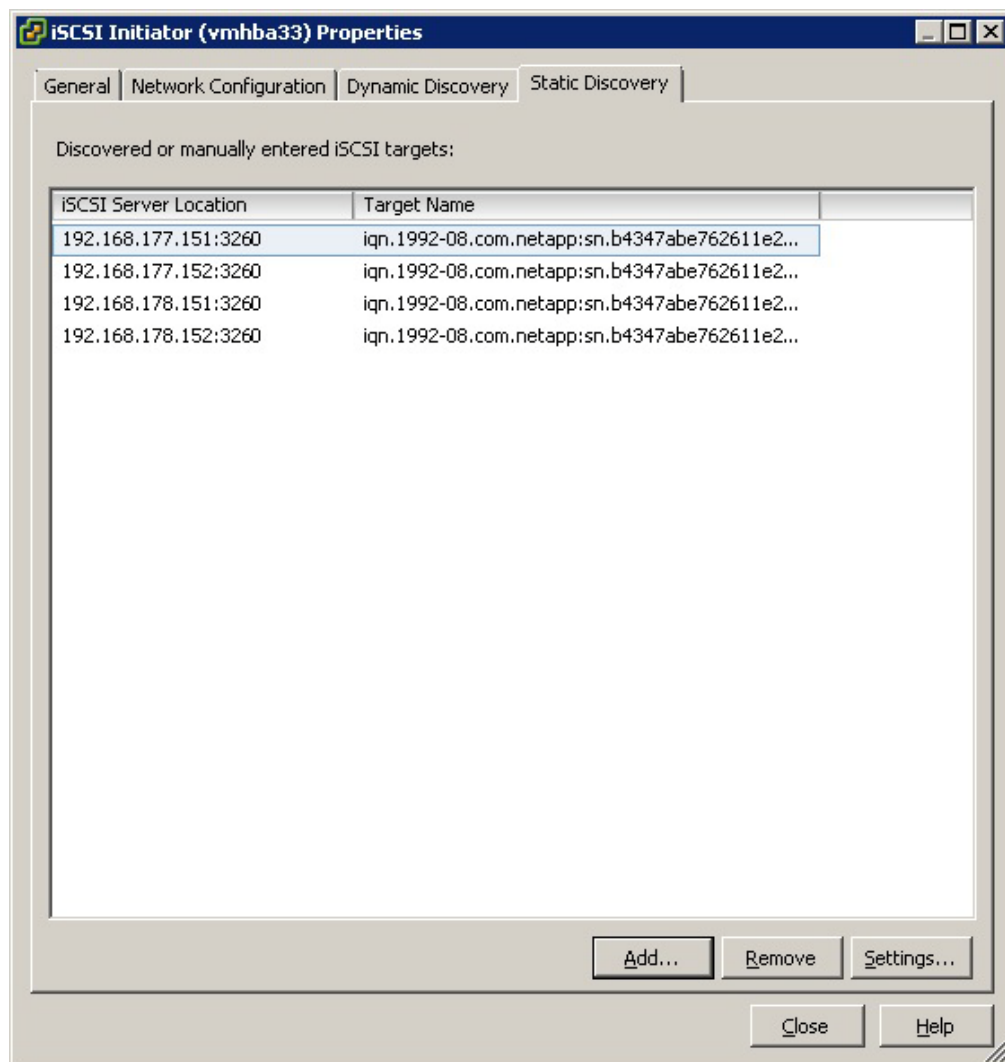
If your FlexPod configuration uses 7-Mode of clustered Data ONTAP, in this step, enter the IP address of controller 1's ifgrp0-<<var_iscsi_vlan_B_id>> interface.

64. Click OK.

65. Repeat steps 60 through 63 until the IP addresses of all four iSCSI LIFs are in the list.

**Note**

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, for the third and fourth targets, enter the iSCSI target node name from controller 2 and use the two iSCSI IPs from controller 2.



66. Click Close to close the iSCSI Initiator Properties windows.

67. Click Yes to rescan the host bus adapter.

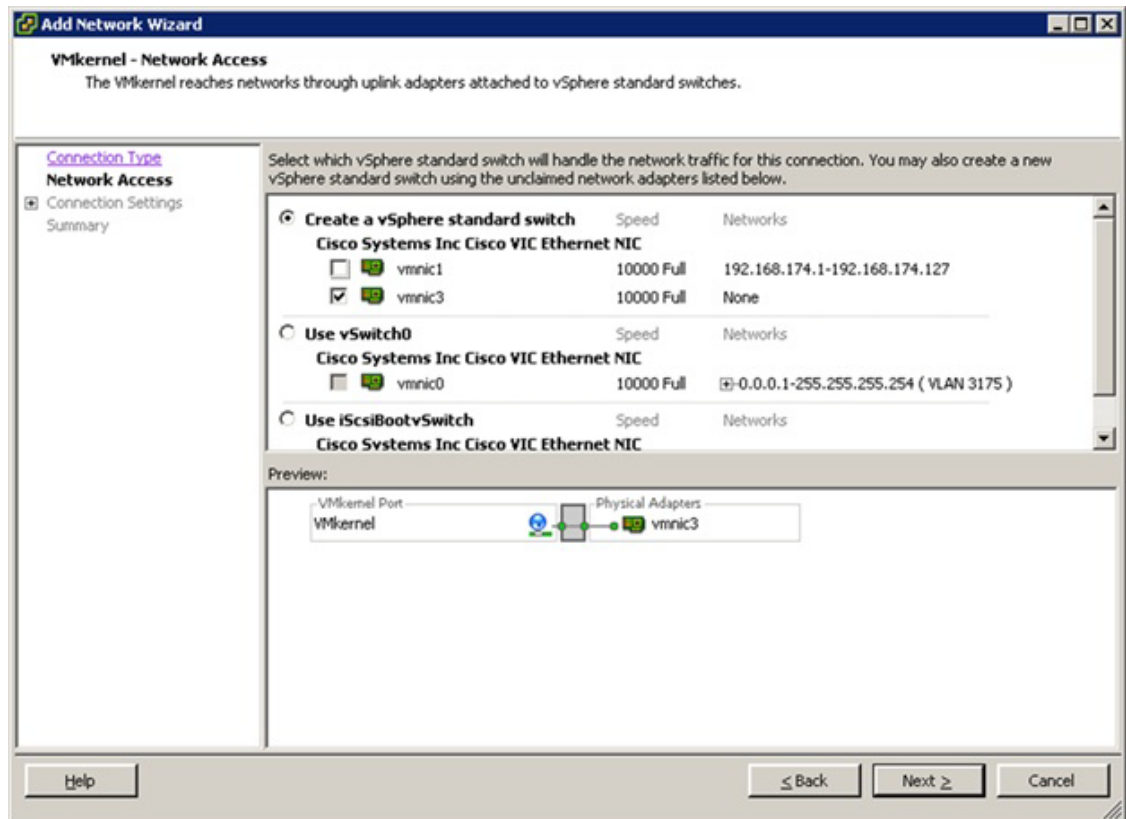
68. If your FlexPod configuration uses 7-Mode storage instead of clustered Data ONTAP, right-click NETAPP iSCSI Disk and select Manage Paths. Change the Path Selection parameter to Round Robin (VMware) and click Change. Click Close to close the Manage Paths window.
69. Right-click the host in the left pane and select Reboot.
70. Click Yes.
71. Enter a reason for the reboot and click OK.
72. After the host has rebooted, log back into the host by using vSphere Client.

ESXi Host VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the VM-Host-Infra-02 ESXi host, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Networking in the Hardware pane.
4. Click Properties on the right side of vSwitch0.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to 9000.
7. Click OK to close the properties for vSwitch0.
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 to finalize the edits for the Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to IB-MGMT Network and enter <<var_ib-mgmt_vlan_id>> in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for the VM Network.
14. Click Add to add a network element.
15. Select VMkernel and click Next.
16. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
17. Click Next to continue with the NFS VMkernel creation.
18. Enter the IP address <<var_nfs_vlan_id_ip_host-02>> and the subnet mask <<var_nfs_vlan_id_mask_host-02>> for the NFS VLAN interface for VM-Host-Infra-02.
19. Click Next to continue with the NFS VMkernel creation.
20. Click Finish to finalize the creation of the NFS VMkernel interface.
21. Select the VMkernel-NFS configuration and click Edit.
22. Change the MTU to 9000.
23. Click OK to finalize the edits for the VMkernel-NFS network.
24. Click Add to add a network element.
25. Select VMkernel and click Next.
26. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.

27. Select the Use This Port Group for vMotion checkbox.
28. Click Next to continue with the vMotion VMkernel creation.
29. Enter the IP address <<var_vmotion_vlan_id_ip_host-02>> and the subnet mask <<var_vmotion_vlan_id_mask_host-02>> for the vMotion VLAN interface for VM-Host-Infra-02.
30. Click Next to continue with the vMotion VMkernel creation.
31. Click Finish to finalize the creation of the vMotion VMkernel interface.
32. Select the VMkernel-vMotion configuration and click Edit.
33. Change the MTU to 9000.
34. Click OK to finalize the edits for the VMkernel-vMotion network.
35. Click Properties on the right side of iScsiBootvSwitch
36. Select the iScsiBootPG configuration and click Edit.
37. Change the Network Label to VMkernel-iSCSI-A. Do not set a VLAN ID.
38. Click OK to save changes to the VMkernel port.
39. Click Close to close the vSwitch Properties.
40. On the right, click Add Networking.
41. Select VMkernel and click Next.
42. Leave Create a vSphere standard switch selected. Unselect vmnic1 and select vmnic3. Click Next.



43. Change the Network Label to VMkernel-iSCSI-B. Leave the VLAN ID set to None.

44. Click Next.
45. Set the VMkernel-iSCSI-B IP Address and Subnet Mask. To get this information, select the VM-Host-Infra-02 Service Profile in UCS Manager. Select the Boot Order tab and select iSCSI-B-vNIC. Select Set iSCSI Boot Parameters. Obtain the IPv4 Address and Subnet Mask from this window.
46. Click Next.

Set iSCSI Boot Parameters

WARNING
The boot parameters are not modifiable because this service profile is bound to a service profile template. To modify the boot parameters, please unbind the service profile from its template.

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name
Initiator Name Assignment: IQN_Pool_B(14/16)
Initiator Name: **iqn.1992-08.com.cisco:fabric-b-ucs-host:2**
+ Create IQN Suffix Pool
The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

Initiator Address
Initiator IP Address Policy: iSCSI_IP_Pool_B(14/16)
IPv4 Address: **192.168.178.11**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**
+ Create IP Pool
The IP address will be automatically assigned from the selected pool.

☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

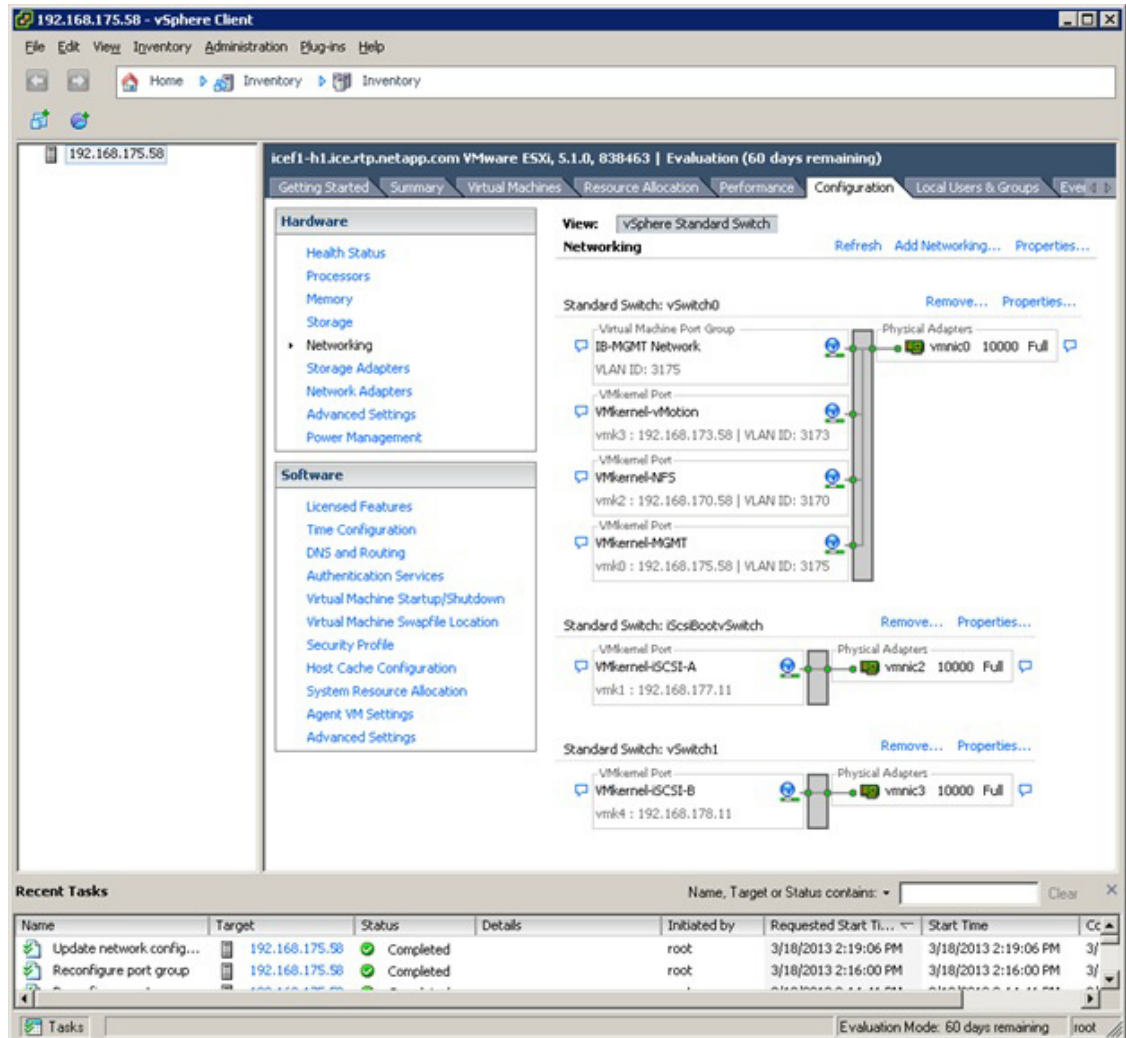
Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id
iqn.1992-08.co...	1	3260		192.168.178.152	0
iqn.1992-08.co...	2	3260		192.168.178.151	0

OK Cancel

47. Click Finish. vSwitch 1 is created.

48. Close the dialog box to finalize the ESXi host networking setup. The networking for the ESXi host should be similar to the following example:



49. Click Storage Adapters in the Hardware pane.
50. Select the iSCSI Software Adapter and click Properties in the Details pane.
51. Select the Network Configuration tab.
52. Click Add.
53. Select VMkernel-iSCSI-A and click OK.
54. Click Add
55. Select VMkernel-iSCSI-B and click OK.
56. Select the Static Discovery tab.
57. Click Settings.
58. Select the entire iSCSI Target Name field, right-click and select Copy to copy this target name to the clipboard.
59. Click Close to close the Static Target Server Settings window.

60. Click Add.
61. Right-click the iSCSI Target Name field and select Paste.
62. Type the IP address of one of the iSCSI LIFs that does not already appear in the list.

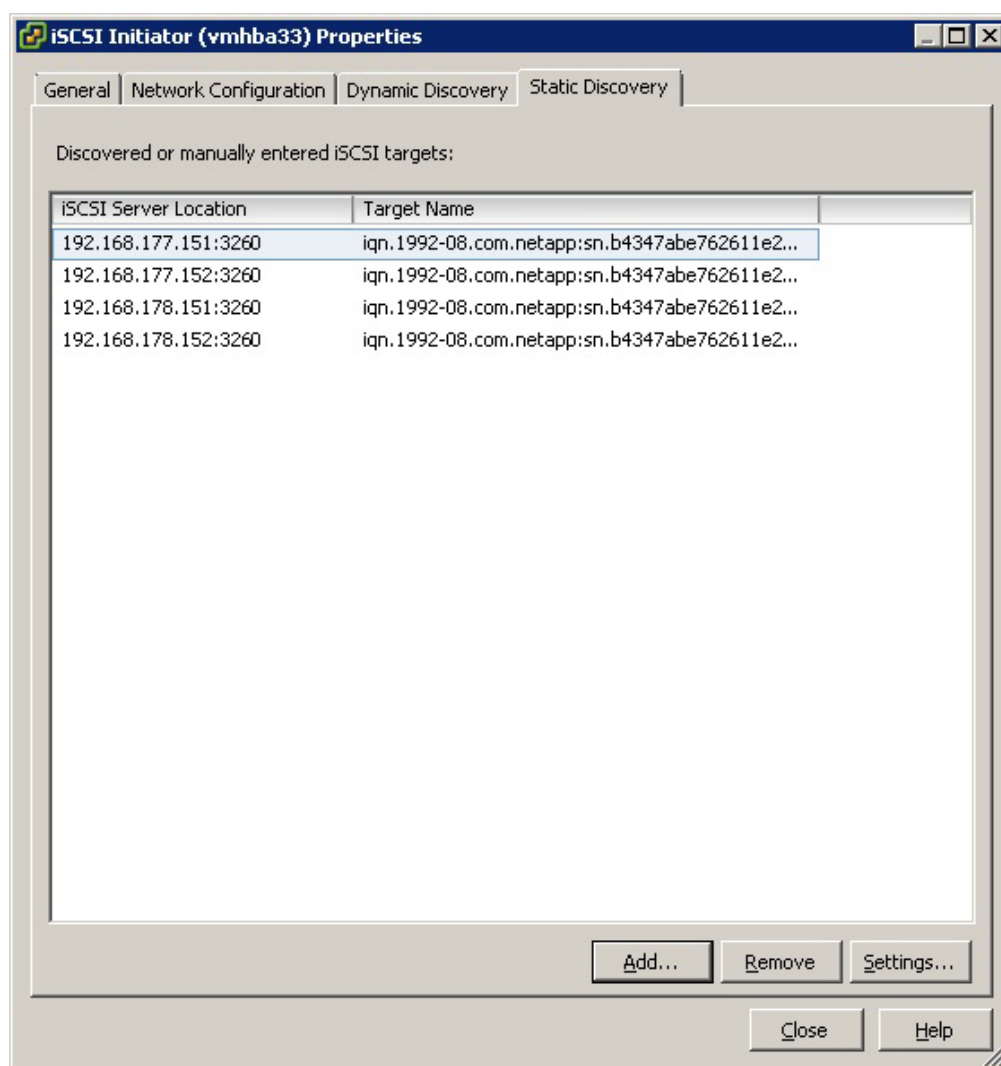
**Note**

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, in this step, enter the IP address of controller 1's ifgrp0-`<<var_iscsi_vlan_B_id>>` interface.

63. Click OK.
64. Repeat steps 60-63 until the IP addresses of all four iSCSI LIFs are in the list.

**Note**

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, for the third and fourth targets, enter the iSCSI target node name from controller 2 and use the two iSCSI IPs from controller 2.



65. Click Close to close the iSCSI Initiator Properties.

66. Click Yes to rescan the host bus adapter.

**Note**

If your FlexPod configuration uses 7-Mode instead of clustered Data ONTAP, right-click NETAPP iSCSI Disk and select Manage Paths. Change the Path Selection parameter to Round Robin (VMware) and click Change. Click Close to close the Manage Paths window.

67. Right-click the host in the left pane and select Reboot.
68. Click Yes.
69. Enter a reason for the reboot and click OK.
70. After the host has rebooted, log back into the host with the vSphere Client.

Mount Required Datastores

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

To mount the required datastores, complete the following steps on each ESXi host:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab to enable configurations.
3. Click Storage in the Hardware pane.
4. From the Datastore area, click Add Storage to open the Add Storage wizard.
5. Select Network File System and click Next.
6. The wizard prompts for the location of the NFS export. Enter <<var_nfs_lif02_ip>> as the IP address for nfs_lif02.

**Note**

For 7-Mode storage, this IP will be the NFS IP address of controller 2.

7. Enter /infra_datastore_1 as the path for the NFS export.

**Note**

For 7-Mode storage, this path will be /vol/infra_datastore_1.

8. Make sure that the Mount NFS read only checkbox is NOT selected.
9. Enter infra_datastore_1 as the datastore name.
10. Click Next to continue with the NFS datastore creation.
11. Click Finish to finalize the creation of the NFS datastore.
12. From the Datastore area, click Add Storage to open the Add Storage wizard.
13. Select Network File System and click Next.
14. The wizard prompts for the location of the NFS export. Enter <<var_nfs_lif01_ip>> as the IP address for nfs_lif01.

**Note**

For 7-Mode storage, this IP will be the NFS IP address of controller 1.

15. Enter /infra_swap as the path for the NFS export.

**Note**

For 7-Mode storage, this path will be /vol/infra_swap.

16. Make sure that the Mount NFS read only checkbox is NOT selected.
17. Enter infra_swap as the datastore name.
18. Click Next to continue with the NFS datastore creation.
19. Click Finish to finalize the creation of the NFS datastore.

Configure NTP on ESXi Hosts

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

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab to enable configurations.
3. Click Time Configuration in the Software pane.
4. Click Properties at the upper right side of the window.
5. At the bottom of the Time Configuration dialog box, click Options.
6. In the NTP Daemon Options dialog box, complete the following steps:
 - a. Click General in the left pane and select Start and stop with host.
 - b. Click NTP Settings in the left pane and click Add.
7. In the Add NTP Server dialog box, enter <<var_global_ntp_server_ip>> as the IP address of the NTP server and click OK.
8. In the NTP Daemon Options dialog box, select the Restart NTP Service to Apply Changes checkbox and click OK.
9. In the Time Configuration dialog box, complete the following steps:
 - a. Select the NTP Client Enabled checkbox and click OK.
 - b. Verify that the clock is now set to approximately the correct time.

**Note**

The NTP server time may vary slightly from the host time.

Move VM Swap File Location

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

To move the VM swap file location, complete the following steps on each ESXi host:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab to enable configurations.
3. Click Virtual Machine Swapfile Location in the Software pane.
4. Click Edit at the upper right side of the window.
5. Select Store the swapfile in a swapfile datastore selected below.

6. Select infra_swap as the datastore in which to house the swap files.
7. Click OK to finalize moving the swap file location.

FlexPod VMware vCenter 5.1

The procedures in the following subsections provide detailed instructions for installing VMware vCenter 5.1 in a FlexPod environment. After the procedures are completed, a VMware vCenter Server will be configured along with a Microsoft SQL Server database to provide database support to vCenter. These deployment procedures are customized to include the environment variables.



Note

This procedure focuses on the installation and configuration of an external Microsoft SQL Server 2008 R2 database, but other types of external databases are also supported by vCenter. To use an alternative database, refer to the VMware vSphere 5.1 documentation for information about how to configure the database and integrate it into vCenter.

To install VMware vCenter 5.1, an accessible Windows Active Directory® (AD) Domain is necessary. If an existing AD Domain is not available, an AD virtual machine, or AD pair, can be set up in this FlexPod environment. Refer to the section "Build Windows Active Directory Server VM(s)" in the appendix.

Build Microsoft SQL Server VM

ESXi Host VM-Host-Infra-01

To build a SQL Server virtual machine (VM) for the VM-Host-Infra-01 ESXi host, complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the vSphere Client, select the host in the inventory pane.
3. Right-click the host and select New Virtual Machine.
4. Select Custom and click Next.
5. Enter a name for the VM. Click Next.
6. Select infra_datastore_1. Click Next.
7. Select Virtual Machine Version: 8. Click Next.
8. Verify that the Windows option and the Microsoft Windows Server® 2008 R2 (64-bit) version are selected. Click Next.
9. Select two virtual sockets and one core per virtual socket. Click Next.
10. Select 4GB of memory. Click Next.
11. Select one network interface card (NIC).
12. For NIC 1, select the IB-MGMT Network option and the VMXNET 3 adapter. Click Next.
13. Keep the LSI Logic SAS option for the SCSI controller selected. Click Next.
14. Keep the Create a New Virtual Disk option selected. Click Next.
15. Make the disk size at least 60GB. Click Next.
16. Click Next.

17. Select the checkbox for Edit the Virtual Machine Settings Before Completion. Click Continue.
18. Click the Options tab.
19. Select Boot Options.
20. Select the Force BIOS Setup checkbox.
21. Click Finish.
22. From the left pane, expand the host field by clicking the plus sign (+).
23. Right-click the newly created SQL Server VM 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 then select Connect to ISO Image on Local Disk.
26. Navigate to the Windows Server 2008 R2 SP1 ISO, select it, and click Open.
27. Click in the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
28. The Windows Installer boots. Select the appropriate language, time and currency format, and keyboard. Click Next.
29. Click Install Now.
30. Make sure that the Windows Server 2008 R2 Standard (Full Installation) option is selected. Click Next.
31. Read and accept the license terms and click Next.
32. Select Custom (Advanced). Make sure that Disk 0 Unallocated Space is selected. Click Next to allow the Windows installation to complete.
33. After the Windows installation is complete and the VM has rebooted, click OK to set the Administrator password.
34. Enter and confirm the Administrator password and click the blue arrow to log in. Click OK to confirm the password change.
35. After logging in to the VM desktop, from the VM console window, select the VM menu. Under Guest, select Install/Upgrade VMware Tools. Click OK.
36. If prompted to eject the Windows installation media before running the setup for the VMware tools, click OK, then click OK.
37. In the dialog box, select Run setup64.exe.
38. In the VMware Tools installer window, click Next.
39. Make sure that Typical is selected and click Next.
40. Click Install.
41. Click Finish.
42. Click Yes to restart the VM.
43. After the reboot is complete, select the VM menu. Under Guest, select Send Ctrl+Alt+Del and then enter the password to log in to the VM.
44. Set the time zone for the VM, IP address, gateway, and host name. Add the VM to the Windows AD domain.

**Note**

A reboot is required.

45. If necessary, activate Windows.

46. Log back in to the VM and download and install all required Windows updates.

**Note**

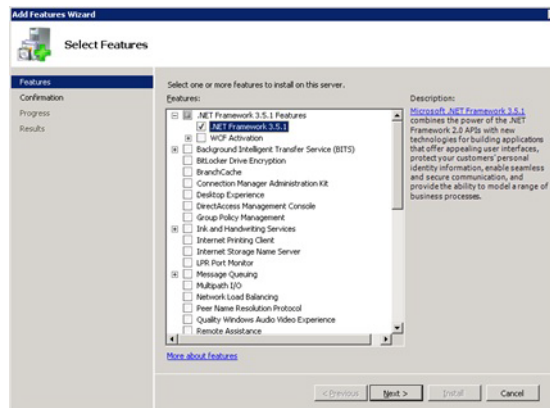
This process requires several reboots.

Install Microsoft SQL Server 2008 R2

vCenter SQL Server VM

To install SQL Server on the vCenter SQL Server VM, complete the following steps:

1. Connect to an AD Domain Controller in the FlexPod Windows Domain and add an admin user for the FlexPod using the Active Directory Users and Computers tool. This user should be a member of the Domain Administrators security group.
2. Log in to the vCenter SQL Server VM as the FlexPod admin user and open Server Manager.
3. Expand Features and click Add Features.
4. Expand .NET Framework 3.5.1 Features and select only .NET Framework 3.5.1.



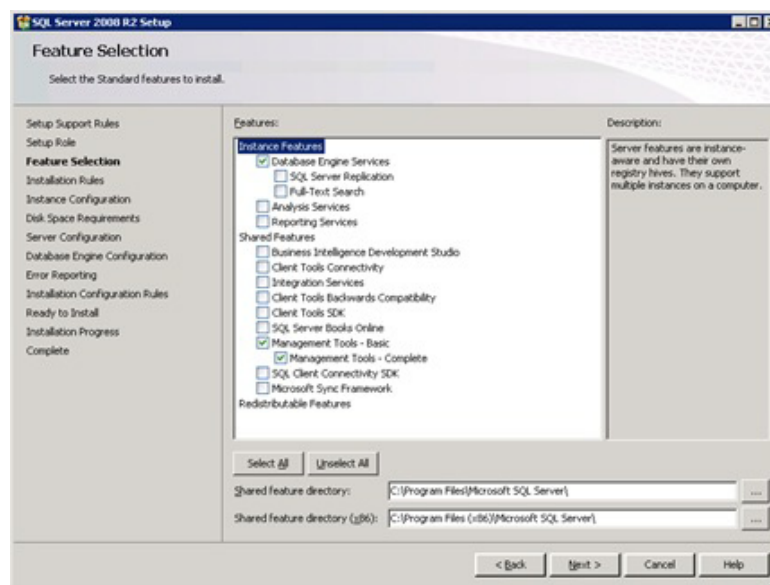
5. Click Next.
6. Click Install.
7. Click Close.
8. Open Windows Firewall with Advanced Security by navigating to Start > Administrative Tools > Windows Firewall with Advanced Security.
9. Select Inbound Rules and click New Rule.
10. Select Port and click Next.
11. Select TCP and enter the specific local port 1433. Click Next.
12. Select Allow the Connection. Click Next, and then click Next again.
13. Name the rule SQL Server and click Finish.
14. Close Windows Firewall with Advanced Security.

15. In the vCenter SQL Server VMware console, click the ninth button (CD with a wrench) to map the Microsoft SQL Server 2008 R2 ISO. Select Connect to ISO Image on Local Disk.
16. Navigate to the SQL Server 2008 R2 ISO, select it, and click Open.
17. In the dialog box, click Run setup.exe.
18. In the SQL Server Installation Center window, click Installation on the left.
19. Select New Installation or Add Features to an Existing Installation.
20. Click OK.
21. Select Enter the Product Key. Enter a product key and click Next.
22. Read and accept the license terms and choose whether to select the second checkbox. Click Next.
23. Click Install to install the setup support files.
24. Address any warnings except for the Windows firewall warning. Click Next.

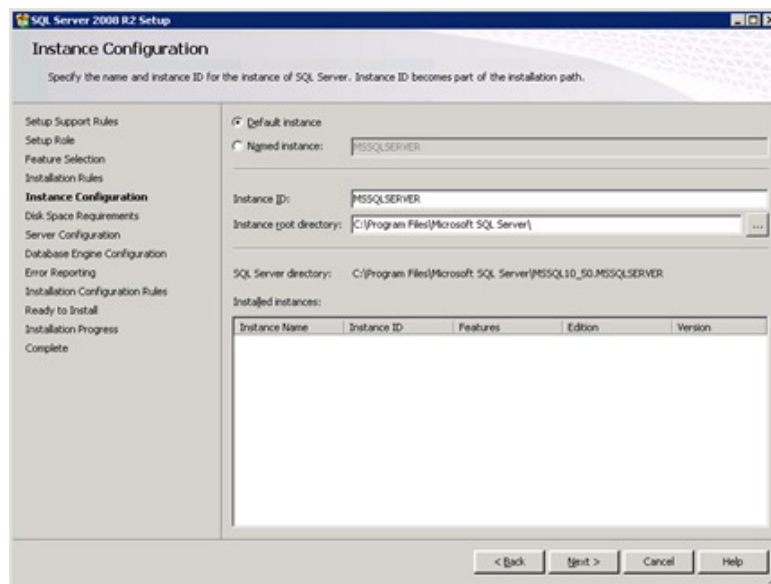
**Note**

The Windows firewall issue was addressed in step 13.

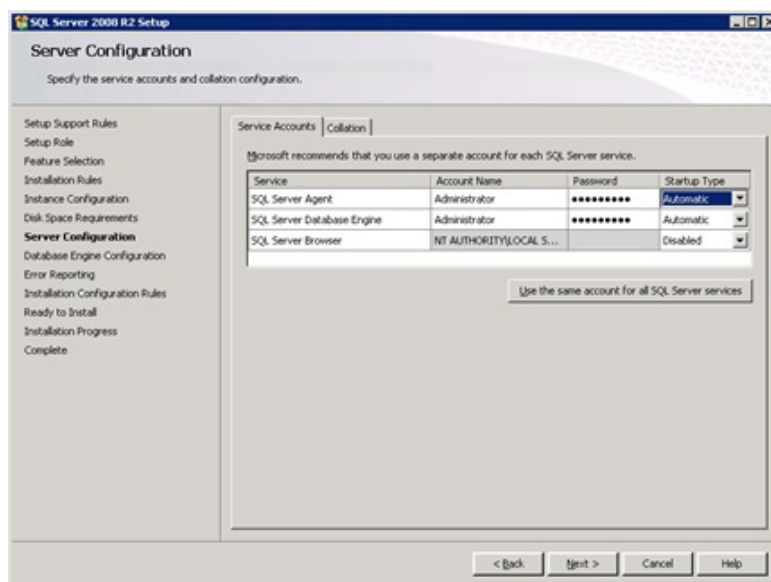
25. Select SQL Server Feature Installation and click Next.
26. Under Instance Features, select only Database Engine Services.
27. Under Shared Features, select Management Tools - Basic and Management Tools - Complete.
28. Click Next.



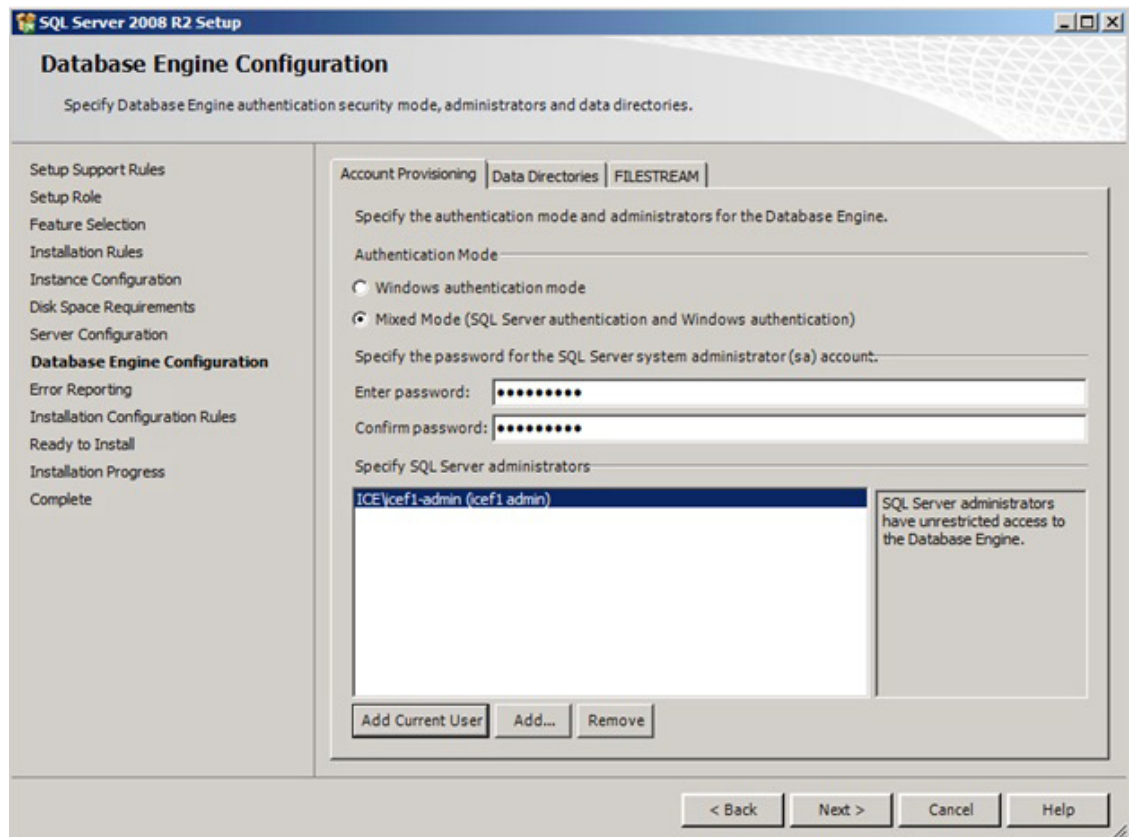
29. Click Next.
30. Keep Default Instance selected. Click Next.



31. Click Next for Disk Space Requirements.
32. For the SQL Server Agent service, click in the first cell in the Account Name column and then click <<Browse...>>.
33. Enter the local machine administrator name (for example, systemname\Administrator), click Check Names, and click OK.
34. Enter the administrator password in the first cell under Password.
35. Change the startup type for SQL Server Agent to Automatic.
36. For the SQL Server Database Engine service, select Administrator in the Account Name column and enter the administrator password again. Click Next.



37. Select Mixed Mode (SQL Server Authentication and Windows Authentication). Enter and confirm the password for the SQL Server system administrator (sa) account, click Add Current User, and Click Next.



38. Choose whether to send error reports to Microsoft. Click Next.
39. Click Next.
40. Click Install.
41. After the installation is complete, click Close to close the SQL Server installer.
42. Close the SQL Server Installation Center.
43. Install all available Microsoft Windows updates by navigating to Start > All Programs > Windows Update.
44. Open the SQL Server Management Studio by selecting Start > All Programs > Microsoft SQL Server 2008 R2 > SQL Server Management Studio.
45. Under Server Name, select the local machine name. Under Authentication, select SQL Server Authentication. Enter sa in the Login field and enter the sa password. Click Connect.
46. Click New Query.
47. Run the following script, substituting the vpxuser password for <Password>:

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

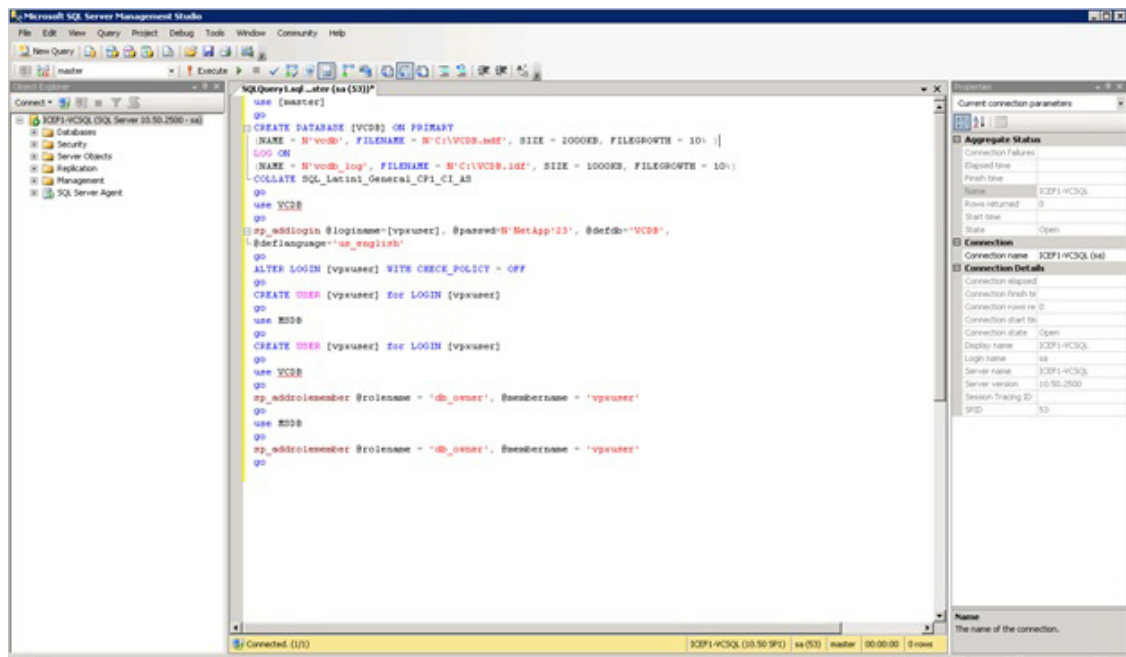
```

(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
use MSDB
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use VCDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go

```

**Note**

This example illustrates the script.



48. Click Execute and verify that the query executes successfully.
49. Close Microsoft SQL Server Management Studio.
50. Disconnect the Microsoft SQL Server 2008 R2 ISO from the SQL Server VM.

Build and Set Up VMware vCenter VM

Build VMware vCenter VM

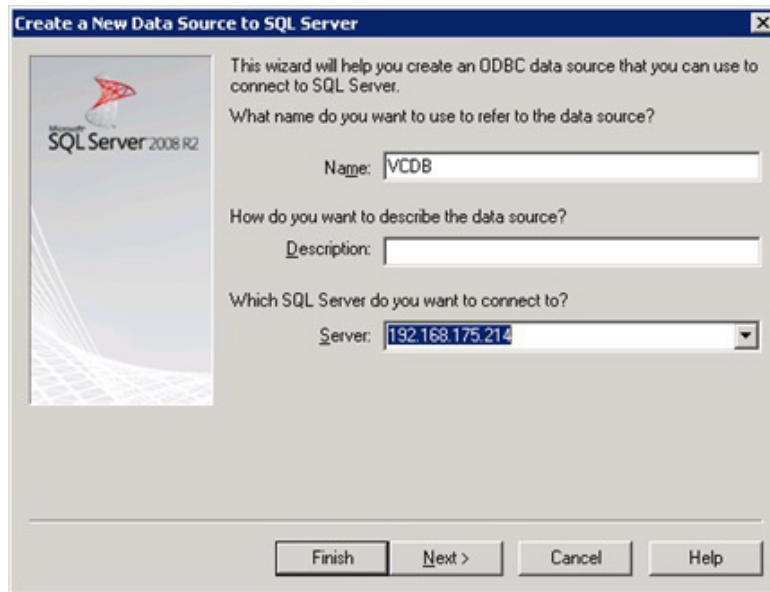
To build the VMware vCenter VM, complete the following steps:

1. Using the instructions for building a SQL Server VM provided in the section "Build Microsoft SQL Server VM," build a VMware vCenter VM with the following configuration in the <<var_ib-mgmt_vlan_id>> VLAN:
 - 4GB RAM
 - Two CPUs
 - One virtual network interface
2. Start the VM, install VMware Tools, and assign an IP address and host name to it in the Active Directory domain.

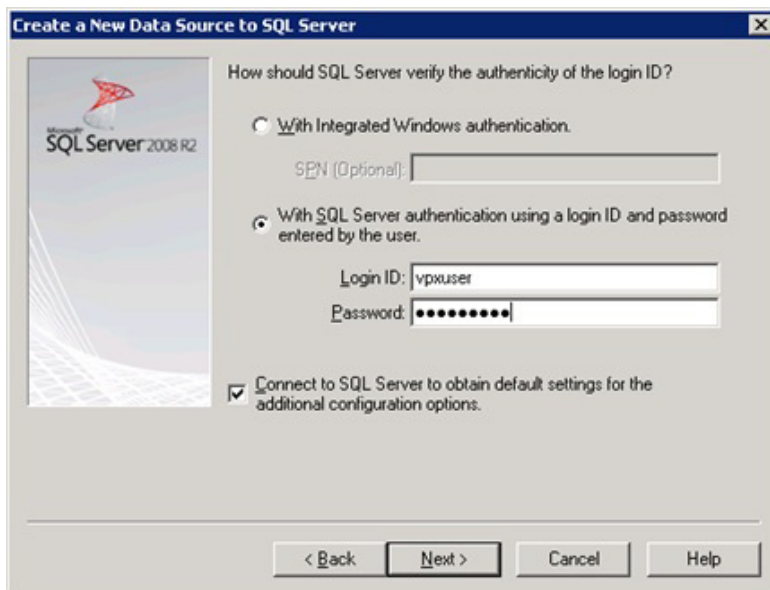
Set Up VMware vCenter VM

To set up the newly built VMware vCenter VM, complete the following steps:

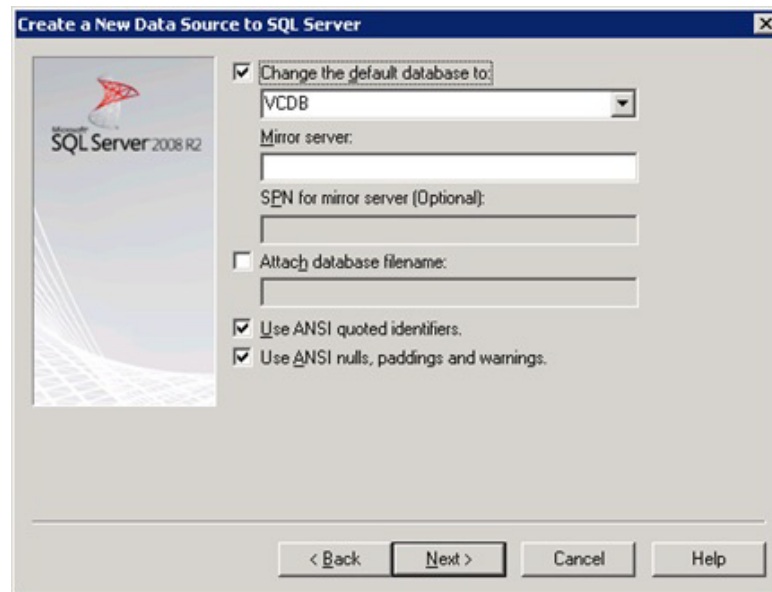
1. Log in to the vCenter VM as the FlexPod admin user and open Server Manager.
2. Expand Features and click Add Features.
3. Expand .NET Framework 3.5.1 Features and select only .NET Framework 3.5.1.
4. Click Next.
5. Click Install.
6. Click Close to close the Add Features wizard.
7. Close Server Manager.
8. Download and install the client components of the Microsoft SQL Server 2008 R2 Native Client from the Microsoft Download Center.
9. Create the vCenter database data source name (DSN). Open Data Sources (ODBC) by selecting Start > Administrative Tools > Data Sources (ODBC).
10. Click 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 using a login ID and password entered by the user. Enter vpxuser as the login ID and the vpxuser password. Click Next.



16. Select Change the Default Database To and select VCDB from the list. Click Next.



17. Click Finish.
18. Click Test Data Source. Verify that the test completes successfully.



19. Click OK and then click OK again.
20. Click OK to close the ODBC Data Source Administrator window.
21. Install all available Microsoft Windows updates by navigating to Start > All Programs > Windows Update.

**Note**

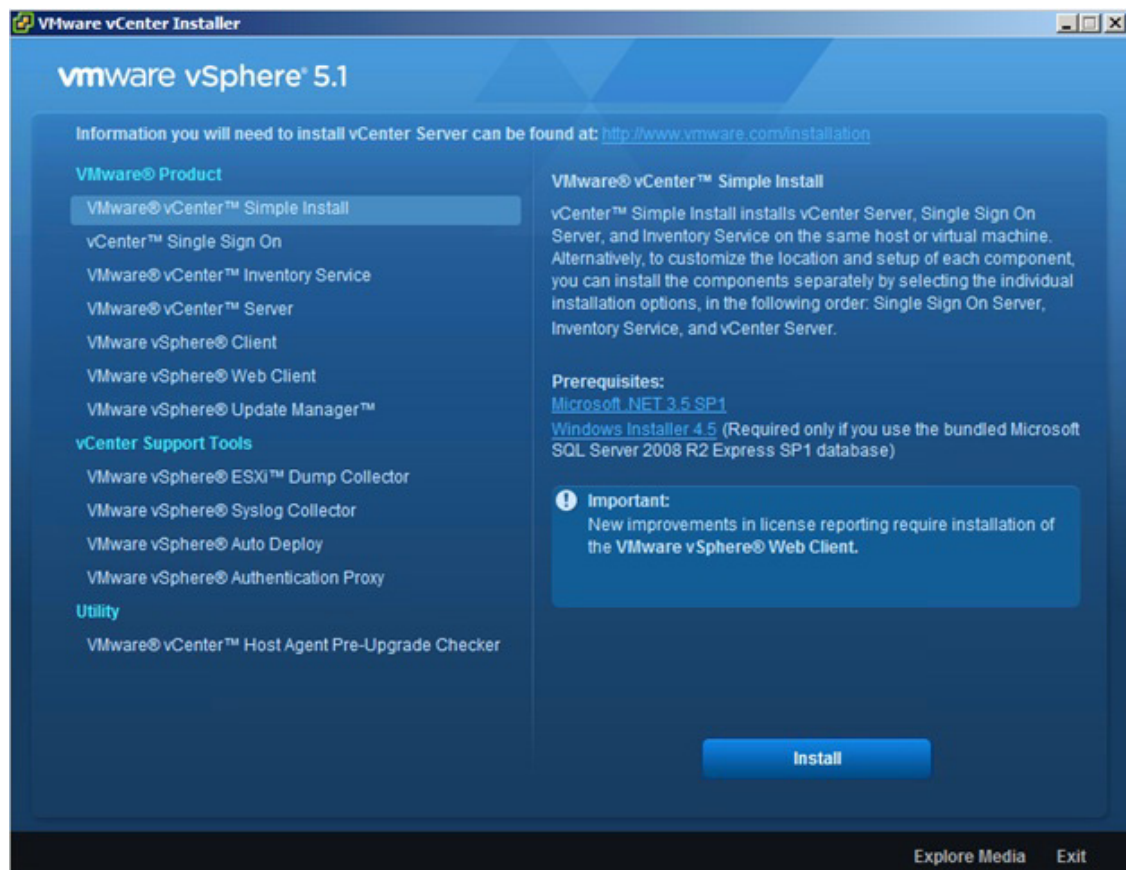
A restart might be required.

Install VMware vCenter Server

vCenter Server VM

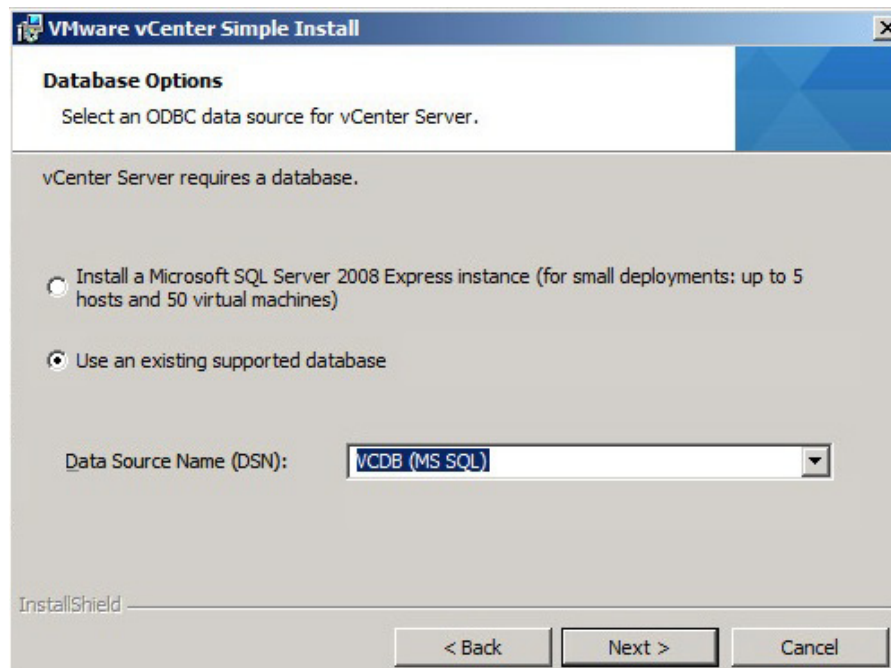
To install vCenter Server on the vCenter Server VM, complete the following steps:

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 5.1 (VIMSetup) ISO, select it, and click Open.
3. In the dialog box, click Run autorun.exe.
4. In the VMware vCenter Installer window, make sure that VMware vCenter Simple Install is selected and click Install.

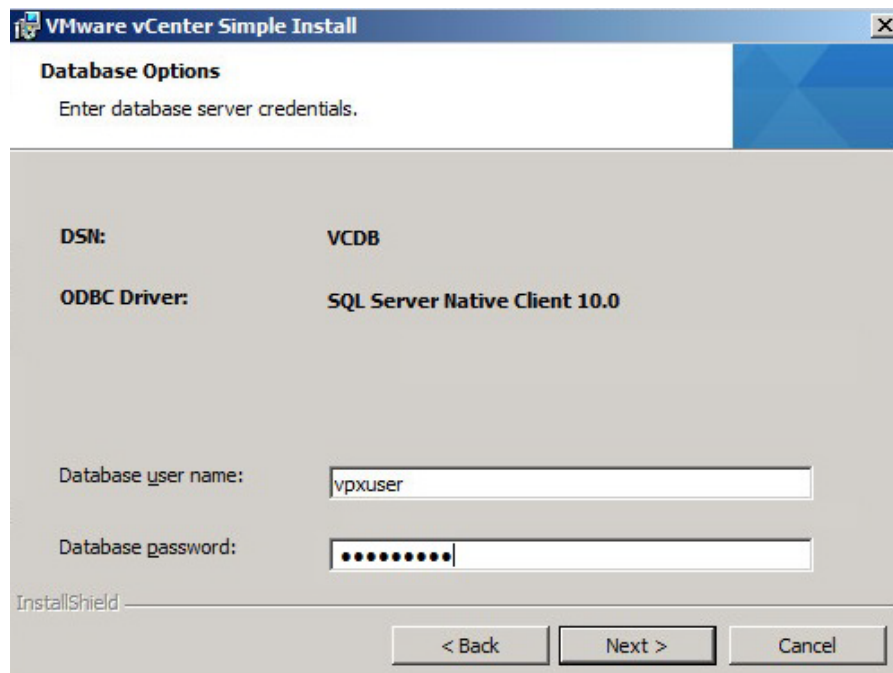


5. Click Yes at the User Account Control warning.
6. Click Next to install vCenter Single Sign On.
7. Click Next.
8. Accept the terms of the license agreement and click Next.
9. Enter and confirm <<var_password>> for admin@System-Domain. Click Next.

10. Keep the radio button selected to install a local Microsoft SQL Server 2008 R2 Express instance and click Next.
11. Enter and confirm <<var_password>> for both user names. Click Next.
12. Verify the vCenter VM FQDN and click Next.
13. Leave Use network service account selected and click Next.
14. Click Next to select the default destination folder.
15. Click Next to select the default HTTPS port.
16. Click Install to install vCenter Single Sign On.
17. Click Yes at the User Account Control warning.
18. Click Yes at the User Account Control warning.
19. Enter the vCenter 5.1 license key and click Next.
20. Select Use an Existing Supported Database. Select VCDB from the Data Source Name list and click Next.



21. Enter the vpxuser password and click Next.



22. Review the warning and click OK.
23. Click Next to use the SYSTEM Account.
24. Click Next to accept the default ports.
25. Select the appropriate inventory size. Click Next.
26. Click Install.
27. Click Finish.
28. In the VMware vCenter Installer window, under vCenter Support Tools, select VMware VSphere ESXi Dump Collector.
29. On the right, click Install.
30. Click Yes.
31. Select the appropriate language and click OK.
32. In the vSphere ESXi Dump Collector Installation Wizard, click Next.
33. Click Next.
34. Accept the terms in the License Agreement and click Next.
35. Click Next to accept the default Destination Folders.
36. Click Next to accept a Standalone installation.
37. Click Next to accept the default ESXi Dump Collector Server Port (6500).
38. Select the VMware vCenter Server IP address from the drop-down menu. Click Next.
39. Click Install to complete the installation.
40. Click Finish.
41. Click Exit in the VMware vCenter Installer window.
42. Disconnect the VMware vCenter ISO from the vCenter VM.

43. Install all available Microsoft Windows updates by navigating to Start > All Programs > Windows Updates.

**Note**

A restart might be required.

44. Back on the Management Workstation, open the VMware vSphere CLI command prompt.
45. Set each ESXi Host to coredump to the ESXi Dump Collector by running the following commands:

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system
coredump network set --interface-name vmk0 --server-ipv4
<<var_vcenter_server_ip> --server-port 6500
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system
coredump network set --interface-name vmk0 --server-ipv4
<<var_vcenter_server_ip> --server-port 6500
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system
coredump network set --enable true
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system
coredump network set --enable true
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system
coredump network check
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system
coredump network check
```

```

C:\Program Files (x86)\VMware\VMware vSphere CLI>
C:\Program Files (x86)\VMware\VMware vSphere CLI>
C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.59 -u root -p NetApp123 system coredump network set --interface-name vmk0 --server-ipv4 192.168.175.188 --server-port 6500

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp123 system coredump network set --interface-name vmk0 --server-ipv4 192.168.175.188 --server-port 6500

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.59 -u root -p NetApp123 system coredump network set --enable true

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp123 system coredump network set --enable true

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.59 -u root -p NetApp123 system coredump network check
Verified the configured netdump server is running

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.58 -u root -p NetApp123 system coredump network check
Verified the configured netdump server is running

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

```

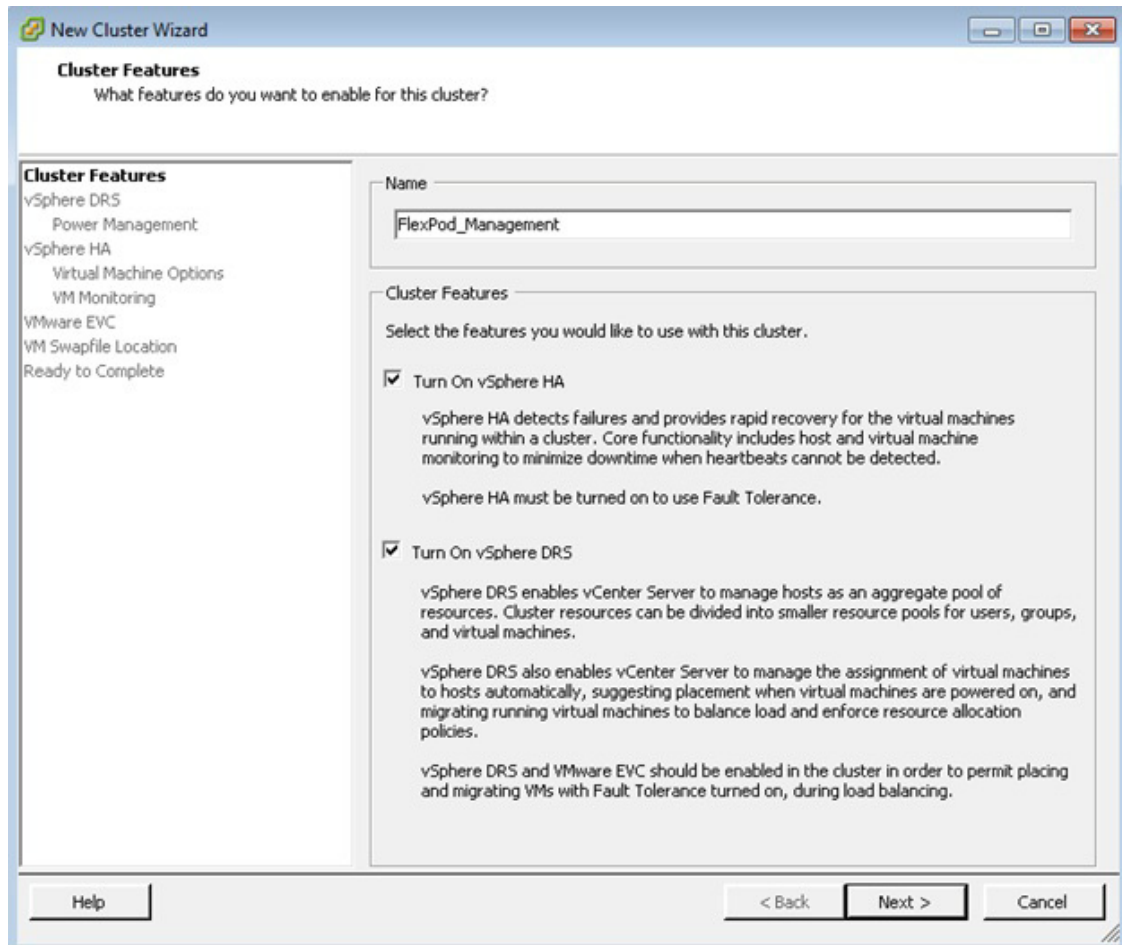
Set Up vCenter Server

vCenter Server VM

To set up vCenter Server on the vCenter Server VM, complete the following steps:

1. Using the vSphere Client, log in to the newly created vCenter Server as the FlexPod admin user.
2. Click Create a data center.
3. Enter FlexPod_DC_1 as the data center name.
4. Right-click the newly created FlexPod_DC_1 data center and select New Cluster.

5. Name the cluster FlexPod_Management and select the checkboxes for Turn On vSphere HA and Turn on vSphere DRS. Click Next.



6. Accept the defaults for vSphere DRS. Click Next.
7. Accept the defaults for Power Management. Click Next.
8. Accept the defaults for vSphere HA. Click Next.
9. Accept the defaults for Virtual Machine Options. Click Next.
10. Accept the defaults for VM Monitoring. Click Next.
11. Accept the defaults for VMware EVC. Click Next.

**Note**

If mixing UCS B or C-Series M2 and M3 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

12. Select Store the swapfile in the datastore specified by the host. Click Next.
13. Click Finish.
14. Right-click the newly created FlexPod_Management cluster and select Add Host.

15. In the Host field, enter either the IP address or the host name of the VM-Host-Infra_01 host. Enter root as the user name and the root password for this host. Click Next.
16. Click Yes.
17. Click Next.
18. Select Assign a New License Key to the Host. Click Enter Key and enter a vSphere license key. Click OK, and then click Next.
19. Click Next.
20. Click Next.
21. Click Finish. VM-Host-Infra-01 is added to the cluster.
22. Repeat this procedure to add VM-Host-Infra-02 to the cluster.

FlexPod Cisco Nexus 1110-X and 1000V vSphere

This section provides detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1110-X Virtual Services Appliances (VSAs) in a FlexPod configuration. The Cisco Nexus 1110-X appliances can be directly connected to the Cisco Nexus 7000 FlexPod switches by using the F2-Series or M1-Series XL modules or by attaching to an existing management infrastructure. This validation effort used a preexisting management infrastructure to support the VSA devices and therefore does not document the cabling configuration.



Note

If connecting the Cisco Nexus 1110-X appliances to a dedicated management network, make sure that all relevant VLANs are available to the VSAs.



Note

If attaching the Cisco Nexus 1110-X to the FlexPod switches, the ports in these line cards will be operating at 1Gbps.

Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the 1110-Xs. By the end of this section, a Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned. This procedure assumes that the Cisco Nexus 1000V software version 4.2(1)SV2(1.1a) has been downloaded from www.cisco.com and expanded. This procedure also assumes that VMware vSphere 5.1 Enterprise Plus licensing is installed.

Configure CIMC Interface on Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure the Cisco Integrated Management Controller (CIMC) interface on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Using the supplied dongle, connect a monitor and USB keyboard to the KVM console port on the front of the Cisco Nexus 1110-X virtual appliance.
2. Reboot the virtual appliance.
3. Press F8 when prompted to configure the CIMC interface.
4. Using the spacebar, set the NIC mode to Dedicated.
5. Clear the checkbox for DHCP enabled.

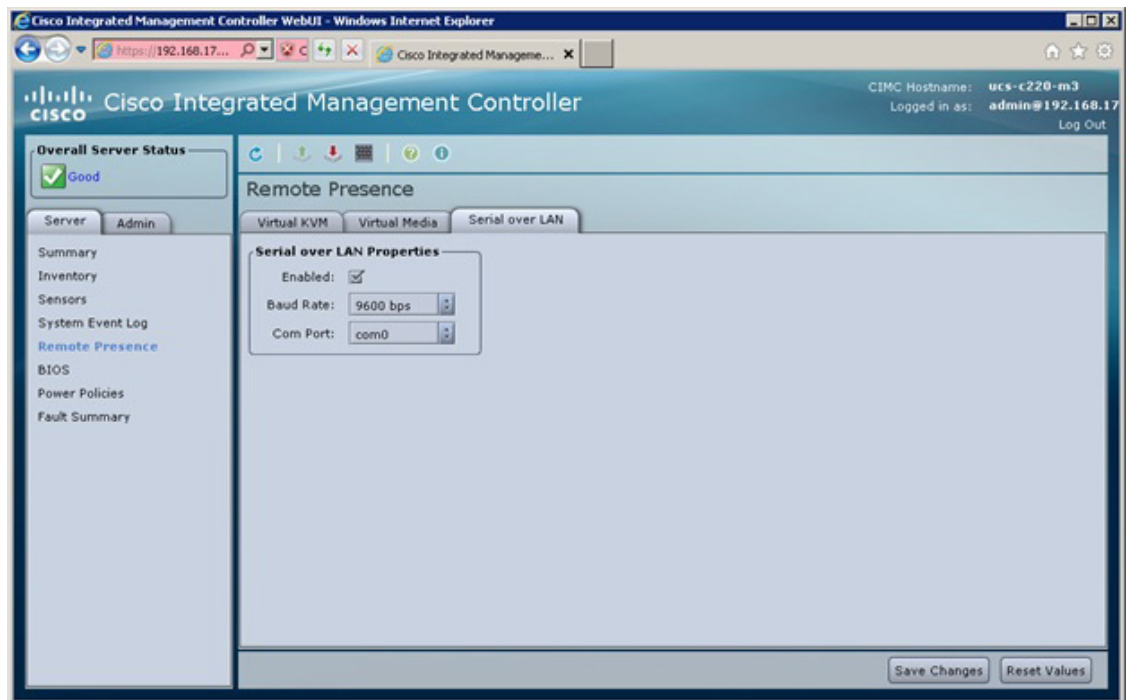
6. Set the CIMC IP address (<<var_cimc_ip>>) in the out-of-band management VLAN.
7. Set the CIMC subnet mask (<<var_cimc_mask>>).
8. Set the CIMC gateway (<<var_cimc_gateway>>).
9. Set the NIC redundancy to None.
10. Set and reenter the CIMC default password (<<var_password>>).
11. Press F10 to save the configuration.
12. Continue pressing F5 until Network settings configured is shown.
13. Press Esc to reboot the virtual appliance.

Configure Serial over LAN for Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

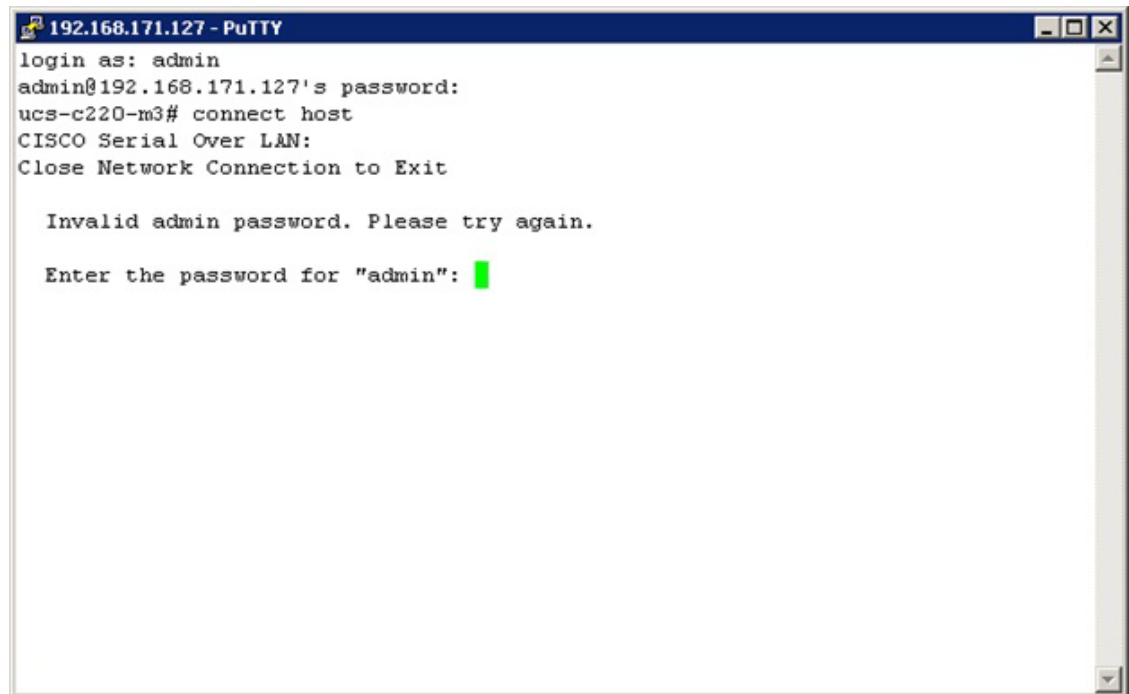
To configure serial over LAN on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Use a Web browser to open the URL at http://<<var_cimc_ip>>.
2. Log in to the CIMC with the admin user id and the CIMC default password (<<var_password>>).
3. In the left column, click Remote Presence.
4. Click the Serial over LAN tab.
5. Select the Enabled checkbox for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, select 9600 bps.
7. Click Save Changes.



8. Log out of the CIMC Web interface.
9. Use an SSH client to connect to <<var_cimc_ip>> with the default CIMC user name and password.

10. Run connect host.



Configure Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.

```

Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Enter HA role[primary/secondary]: primary
Enter network-uplink type <1-5>: 1
Enter control VLAN <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
Enter the domain<1-4095>: <<var_1110x_domain_id>>
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
Would you like to enter the basic system configuration dialogue (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the VSA name : <<var_1110x_vsa>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IP address type V4/V6? (V4): Enter
Mgmt0 IPv4 address : <<var_1110x_vsa_ip>>
Mgmt0 IPv4 netmask : <<var_1110x_vsa_mask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <<var_1110x_vsa_gateway>>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (das/rsa) [rsa]: Enter
Number of rsa key bits <768-2048> [1024]: Enter

```

```
Enable the http server? (yes/no) [y]: Enter
Configure the ntp server? (yes/no) [n]: y
2. NTP server IPv4 address: <<var_global_ntp_server_ip>>
```

3. Review the configuration summary. If everything is correct, enter no to skip editing the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter
```

4. The Cisco Nexus 1110-X saves the configuration and reboots. After reboot, log back in as admin.

Cisco Nexus 1110-X B

To configure the Cisco Nexus 1110-X B, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.

```
Enter the password for "admin": <<var_password>>
```



Note

This is the same password that you entered on the primary Cisco Nexus 1110-X.

2. Enter the admin password again to confirm: <<var_password>>.

```
Enter HA role[primary/secondary]: secondary
Enter network-uplink type <1-5>: 1
Enter control vlan <1-3967, 4048-4093>: <<var_pkt_ctrl_vlan_id>>
Enter the domain id<1-4095>: <<var_1110x_domain_id>>
```



Note

This is the same unique Cisco Nexus 1110 domain ID entered on Cisco Nexus 1110-X A.

```
Enter management vlan <1-3967, 4048-4093>: <<var_ib_mgmt_vlan_id>>
```

3. The Cisco Nexus 1110-X saves the configuration and reboots.

Set Up the Primary Cisco Nexus 1000V VSM

Cisco Nexus 1110-X A

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:

1. Continue periodically running the following command until module 2 (Cisco Nexus 1110-X B) has a status of ha-standby.

```
show module
```

2. Enter the global configuration mode and create a virtual service blade.

```
config t
virtual-service-blade VSM-1
dir /repository
```

3. If the desired Cisco Nexus 1000V ISO file (nexus-1000v.4.2.1.SV2.1.1a.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must place the file either on an FTP server or on a UNIX or Linux® machine (using scp) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy ftp://<<var_ftp_server>>/nexus-1000v.4.2.1.SV2.1.1a.iso /repository/.

```
virtual-service-blade-type new nexus-1000v.4.2.1.SV2.1.1a.iso
interface control vlan <<var_pkt_ctrl_vlan_id>>
interface packet vlan <<var_pkt_ctrl_vlan_id>>
enable primary
Enter vsb image:[nexus-1000v.4.2.1.SV2.1.1a.iso] Enter
```

```
Enter domain id[1-4095]: <<var_vsm_domain_id>>
```

**Note**

This domain ID should be different than the VSA domain ID.

```
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
Enter Management subnet mask: <<var_vsm_mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
Enter the password for 'admin': <<var_password>>
copy run start
```

4. Run show virtual-service-blade summary. Continue periodically entering this command until the primary VSM-1 has a state of VSB POWERED ON.

Set Up the Secondary Cisco Nexus 1000V VSM

To set up the secondary Cisco Nexus 1000V VSM on Cisco Nexus 1110-X B, complete the steps in the following two subsections:

Cisco Nexus 1110-X A

1. Run system switchover to activate Cisco Nexus 1110-X B.

Cisco Nexus 1110-X B

1. Log in to Cisco Nexus 1110-X B as the admin user.

```
config t
virtual-service-blade VSM-1
dir /repository
```

2. If the desired Cisco Nexus 1000V ISO file (nexus-1000v.4.2.1.SV2.1.1a.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must place the file either on an FTP server or on a UNIX or Linux machine (using the scp command) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy ftp:// <<var_ftp_server>>/nexus-1000v.4.2.1.SV2.1.1a.iso /repository/.

```
enable secondary
Enter vsb image: [nexus-1000v.4.2.1.SV2.1.1a.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
Enter Management subnet mask: <<var_vsm_mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
```

3. Enter the admin password <<var_password>>.
4. Type show virtual-service-blade summary. Continue periodically entering this command until both the primary and secondary VSM-1s have a state of VSB POWERED ON.

```
copy run start
```

5. Run system switchover on Cisco Nexus 1110-X B to activate Cisco Nexus 1110-X A. This causes Cisco Nexus 1110-X B to reboot.

Install Virtual Ethernet Module on Each ESXi Host

vCenter Server VM

To install the Virtual Ethernet Module (VEM) on the ESXi hosts, complete the following steps:

1. Launch a Web browser to `http://<var_vsm_mgmt_ip>`.
2. Right-click the `cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib` hyperlink and select Save target as.
3. Save the file as `cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib`, type All Files, on the Desktop of the management workstation.
4. From the main window in the vSphere Client connected to vCenter, select the first server in the list under the FlexPod Management cluster.
5. Click the Summary tab.
6. Under Storage on the right, right-click `infra_datastore_1` and select Browse Datastore.
7. Select the root folder (/) and click the third button at the top to add a folder.
8. Name the folder VEM and click OK.
9. On the left, select the VEM folder.
10. Click the fourth button at the top and select Upload File.
11. Navigate to the `cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib` file and click Open.
12. Click Yes. The VEM file should now appear in the VEM folder in the datastore.
13. Open the VMware vSphere CLI command prompt.
14. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <Host Server IP> -u root -p <Root Password> software vib install -v
/vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vi
b
```

```

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.62 -u root -p NetApp!23 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  UIBs Installed: Cisco_bootbank_cisco-vem-v152-esx_4.2.1.2.1.1a.0-3.1.1
  UIBs Removed:
  UIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.101 -u root -p NetApp!23 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  UIBs Installed: Cisco_bootbank_cisco-vem-v152-esx_4.2.1.2.1.1a.0-3.1.1
  UIBs Removed:
  UIBs Skipped:

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

Register Cisco Nexus 1000V as a vCenter Plug-in

To register the Cisco Nexus 1000V as a vCenter plug-in, complete the following steps:

1. Using a web browser, navigate to the `<<var_vsm_mgmt_ip>` using `http://<var_vsm_mgmt_ip>`.
2. Right-click the `cisco_nexus_1000v_extension.xml` hyperlink and select Save target as.
3. Save the XML file to the local desktop.

4. In the vSphere Client connected to vCenter, select Plug-ins > Manage Plug-ins.
5. Right-click 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 that was previously saved. 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.

Perform Base Configuration of the Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t
svs connection vCenter
protocol vmware-vim
remote ip address <<var_vcenter_server_ip>> port 80
vmware dvs datacenter-name FlexPod_DC_1
connect
exit
ntp server <<var_global_ntp_server_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
system mtu 9000
state enabled
port-profile type ethernet iscsi-a-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_a_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_a_vlan_id>>
```

```

no shutdown
system vlan <<var_iscsi_a_vlan_id>>
state enabled
port-profile type ethernet iscsi-b-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
state enabled
port-profile type vethernet IB-MGMT-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet NFS-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_id>>
no shutdown
system vlan <<var_nfs_vlan_id>>
state enabled
port-profile type vethernet vMotion-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>>
no shutdown
system vlan <<var_vmotion_vlan_id>>
state enabled
port-profile type vethernet VM-Traffic-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_id>>
no shutdown
system vlan <<var_vm-traffic_vlan_id>>
state enabled
port-profile type vethernet nlkv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet iSCSI-A-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
state enabled
port-profile type vethernet iSCSI-B-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_id>>

```



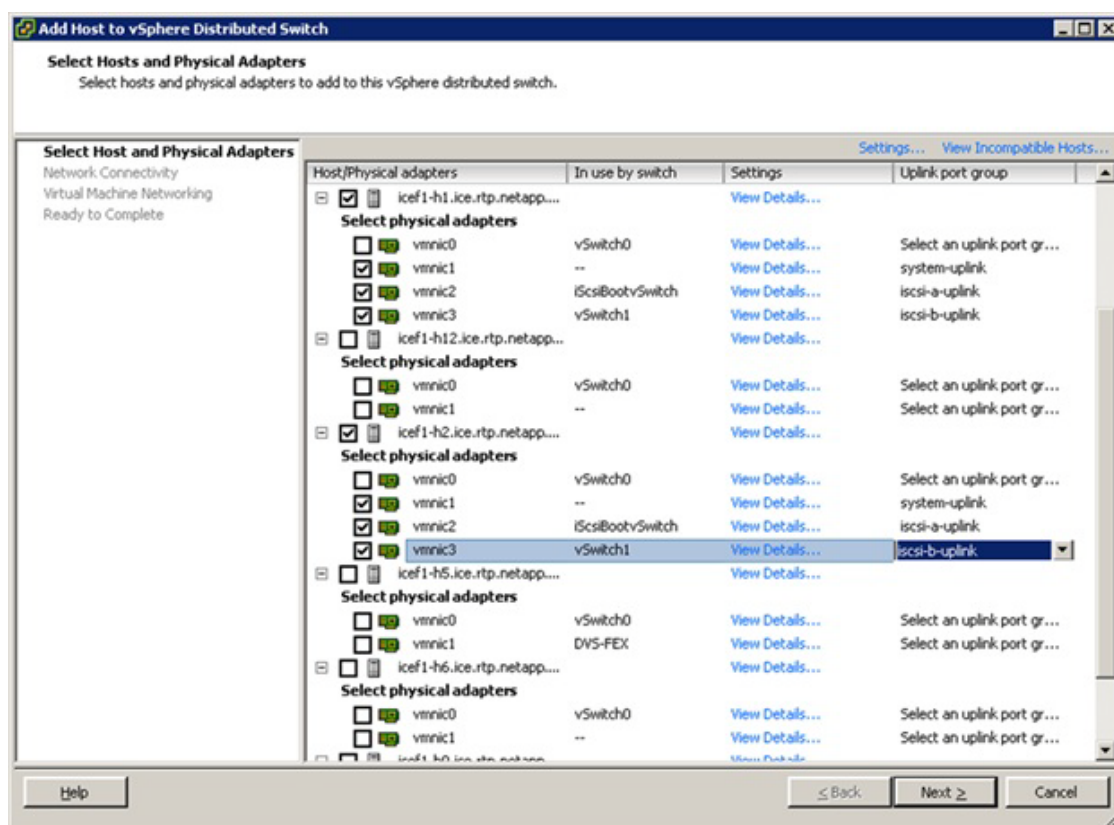
```
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
state enabled
exit
copy run start
```

Migrate Networking Components for ESXi Hosts to Cisco Nexus 1000V

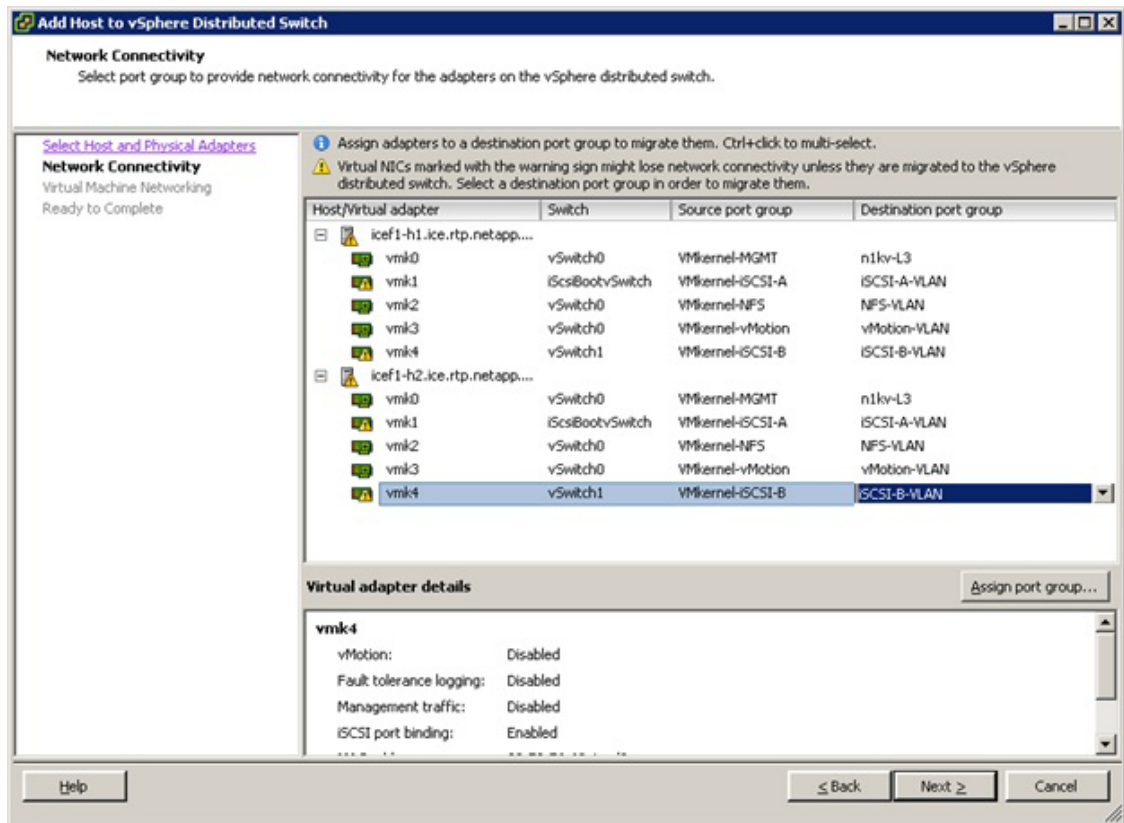
vSphere Client Connect to vCenter

To migrate the networking components for the ESXi hosts to the Cisco Nexus 1000V, complete the following steps:

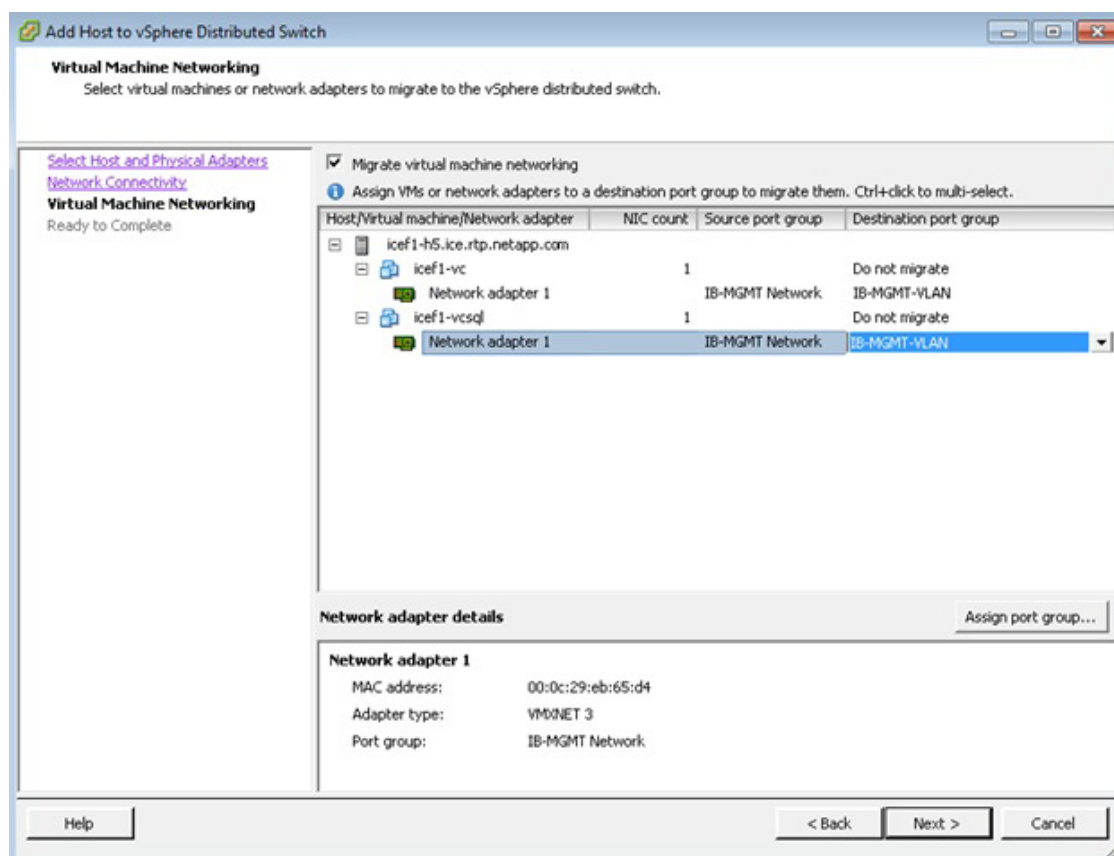
1. In the VMware vSphere Client connected to vCenter, select Home > Networking.
2. Expand the vCenter, DataCenter, and Cisco Nexus 1000V folders. Select the Cisco Nexus 1000V switch.
3. Under Basic Tasks for the vSphere distributed switch, select Add a Host.
4. For both hosts, select vmnic1 and select the system-uplink Uplink port group. Also, for vmnic2 select iscsi-a-uplink and for vmnic 3 select iscsi-b-uplink. Click Next.



5. For all VMkernel ports, select the appropriate Destination Port Group from the Cisco Nexus 1000V, making sure to select the "n1kv-L3" destination port group for the MGMT VMkernel ports. Click Next.



6. On the popup, select Continue without resolving the errors and click Close.
7. Select the Migrate Virtual Machine Networking checkbox. Expand each VM and select the port groups for migration individually. Click Next.



8. Click Finish. Wait for the migration process to complete.
9. In the vSphere Client window, select Home > Hosts and Clusters.
10. Select the first ESXi host and select the Configuration tab. In the Hardware box, select Networking.
11. Make sure that vSphere Standard Switch is selected at the top next to View. None of the three vSwitches should have any active VMkernel or VM network ports on them. On the upper right of each vSwitch, click Remove.
12. Click Yes.
13. Remove all three vSwitches.
14. After all vSwitches have disappeared from the screen, click vSphere Distributed Switch at the top next to View.
15. Click Manage Physical Adapters.
16. Scroll down to the system-uplink box and click <Click to Add NIC>.
17. Select vmnic0 and click OK.
18. Click OK to close the Manage Physical Adapters window. Two system uplinks should now be present.
19. Select the second ESXi host and click the Configuration tab. In the Hardware box, select Networking.
20. Make sure vSphere Standard Switch is selected at the top next to View. None of the vSwitches should have any active VMkernel or VM network ports on them. At the upper right of each vSwitch, click Remove.

21. Click Yes.
22. Remove all three vSwitches.
23. After all three vSwitches have disappeared from the screen, click vSphere Distributed Switch at the top next to View.
24. Click Manage Physical Adapters.
25. Scroll down to the system-uplink box and click <Click to Add NIC>.
26. Select vmnic0 and click OK.
27. Click OK to close the Manage Physical Adapters dialog box. Two system-uplinks should now be present.
28. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	up	routed	full	1000	--
Eth3/1	--	up	trunk	full	10G	--
Eth3/2	--	up	trunk	full	10G	--
Eth3/3	--	up	trunk	full	10G	--
Eth3/4	--	up	trunk	full	10G	--
Eth4/1	--	up	trunk	full	10G	--
Eth4/2	--	up	trunk	full	10G	--
Eth4/3	--	up	trunk	full	10G	--
Eth4/4	--	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	3178	auto	auto	--
Veth3	VMware VMkernel, v	up	3170	auto	auto	--
Veth4	VMware VMkernel, v	up	3173	auto	auto	--
Veth5	VMware VMkernel, v	up	3177	auto	auto	--
Veth6	VMware VMkernel, v	up	3175	auto	auto	--
Veth7	VMware VMkernel, v	up	3177	auto	auto	--

--More--

29. Run show module and verify that the two ESXi hosts are present as modules.

```

icefl-vsm
icefl-vsm(config)# show module
Mod  Ports  Module-Type                Model                Status
---  -
1    0      Virtual Supervisor Module  Nexus1000V           ha-standby
2    0      Virtual Supervisor Module  Nexus1000V           active *
3    248    Virtual Ethernet Module    NA                    ok
4    248    Virtual Ethernet Module    NA                    ok

Mod  Sw                Hw
---  -
1    4.2 (1) SV2 (1.1a) 0.0
2    4.2 (1) SV2 (1.1a) 0.0
3    4.2 (1) SV2 (1.1a) VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
4    4.2 (1) SV2 (1.1a) VMware ESXi 5.1.0 Releasebuild-838463 (3.1)

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--

```

30. Run copy run start.
31. Type exit two times to log out of the Cisco Nexus 1000v.

FlexPod Management Tool Setup

NetApp Virtual Storage Console (VSC) 4.1 Deployment Procedure

VSC 4.1 Preinstallation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.1.2:

- Protocol licenses (NFS and iSCSI)
- FlexClone (for provisioning and cloning only)
- SnapRestore (for backup and recovery)
- SnapManager suite

Install VSC 4.1

To install the VSC 4.1 software, complete the following steps:

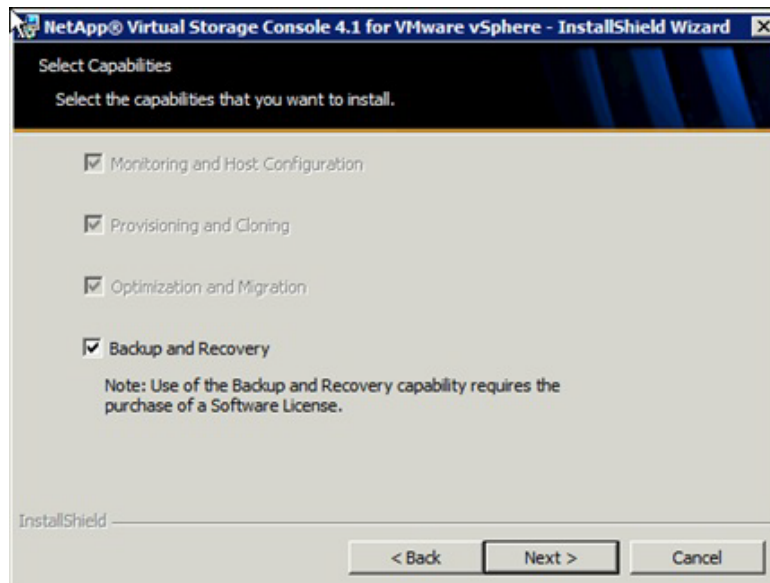
1. Using the instructions in section "Build Microsoft SQL Server VM," build a VSC and an OnCommand virtual machine with 4GB RAM, two CPUs, and one virtual network interface in the <<var_ib-mgmt_vlan_id>> VLAN. The virtual network interface should be a VMXNET 3 adapter.

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 VM. Install all Windows updates on the VM.


2. Log in to the VSC and OnCommand VM as the FlexPod admin user.
3. Download the x64 version of the Virtual Storage Console 4.1 from the NetApp Support site.
4. Right-click the file downloaded in step 3 and select Run As Administrator.
5. Click Yes at the User Access Control warning.
6. On the Installation wizard Welcome page, click Next.
7. Select the backup and recovery capability. Click Next.

**Note**

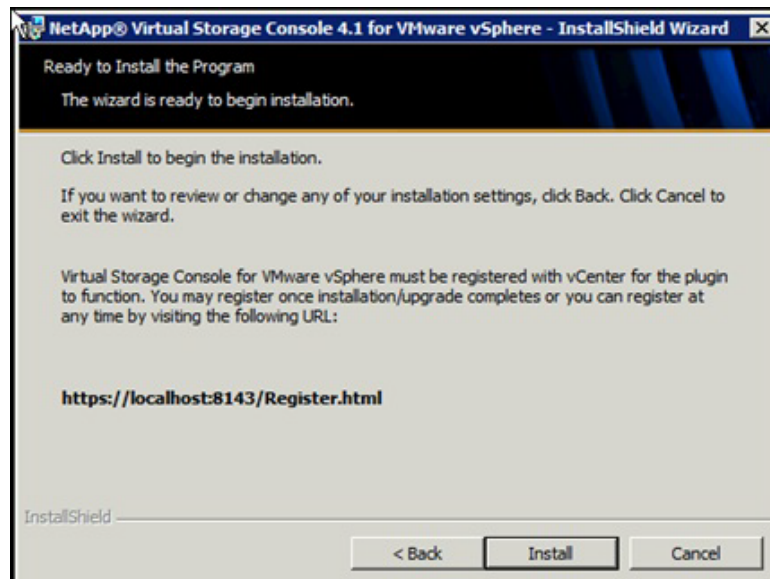
The backup and recovery capability requires an additional license.



8. Click Next to accept the default installation location.



9. Click Install.
10. Click Finish.



Register VSC with vCenter Server

To register the VSC with the vCenter Server, complete the following steps.

1. A browser window with the registration URL opens automatically when the installation phase is complete.
2. Click Continue to this website (not recommended).
3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.

4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user), and user password for the vCenter Server. Click Register to complete the 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: 192.168.175.191

vCenter Server information

Host name or IP Address: 192.168.175.188

Port: 443

User name: ice\icef1-admin

User password: ••••••••

Register

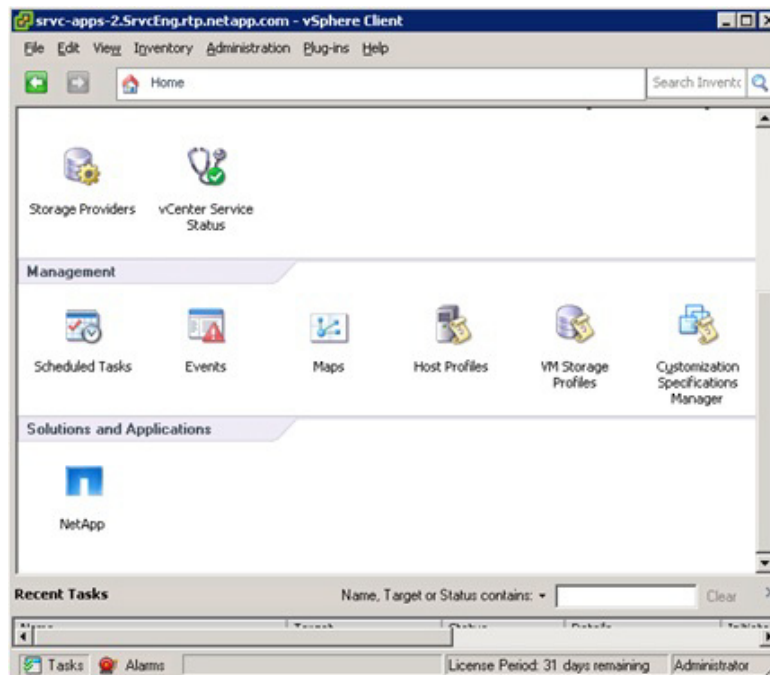
**Note**

If using 7-Mode storage, iSCSI and NFS network addresses need to be configured in VSC. Go to the section "Set Up 7-Mode iSCSI and NFS Networks in VSC" in the appendix and execute that procedure.

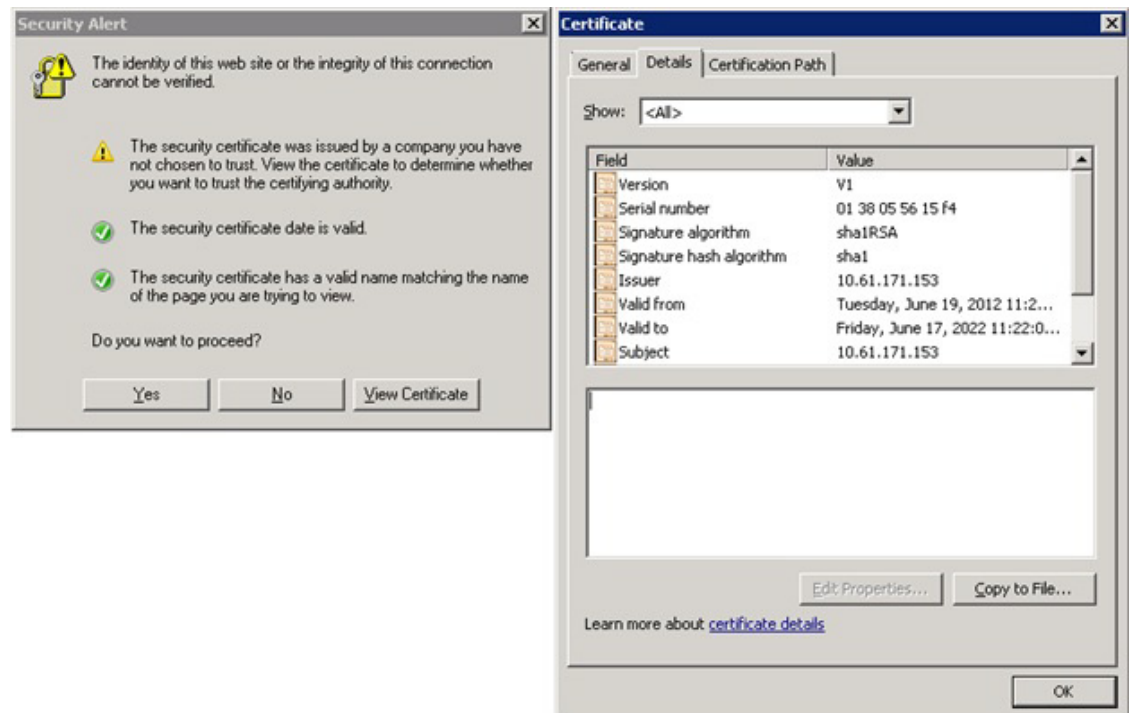
Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, complete the following steps:

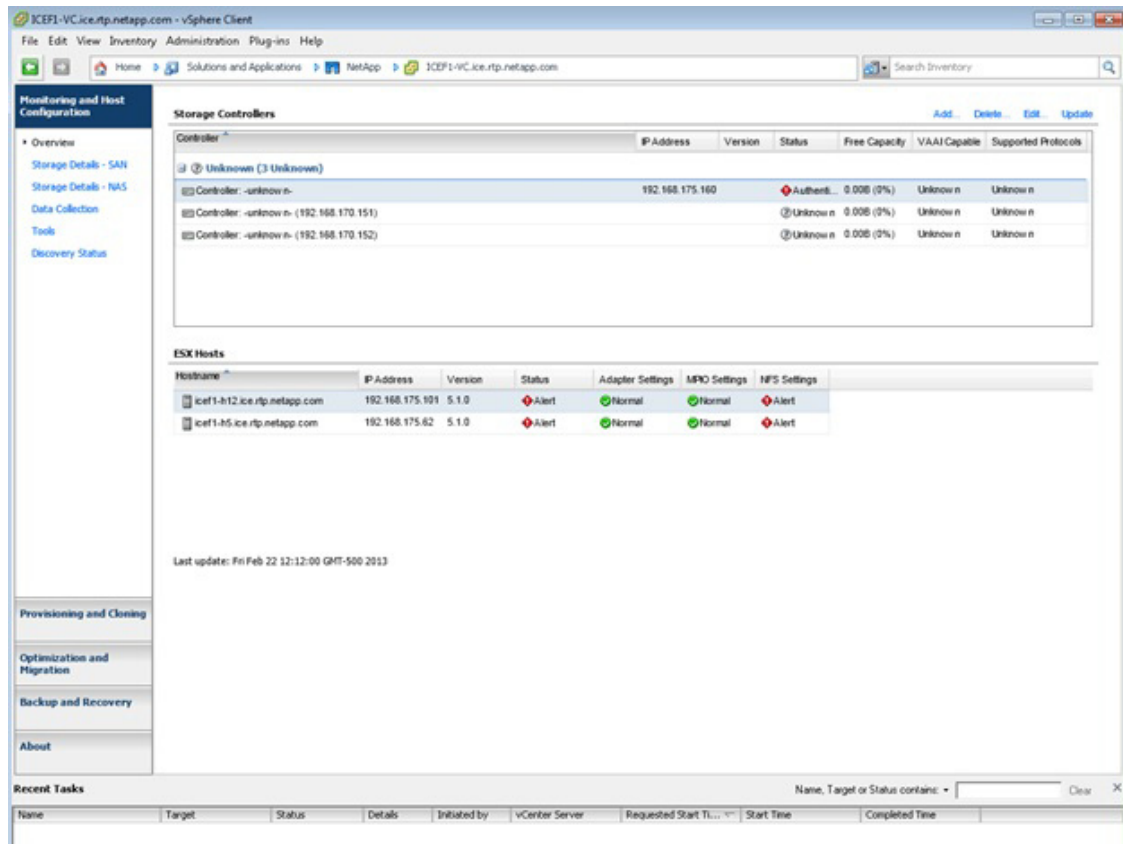
1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user. If the vSphere Client was previously opened, close it and then reopen it.
2. Click the Home tab in the left side of the vSphere Client window.
3. Under Solutions and Applications, click the NetApp icon.



4. Click Yes when the security certificate warning appears. To view the certificate, click View Certificate.



5. In the navigation pane, select Monitoring and Host Configuration if it is not selected by default.



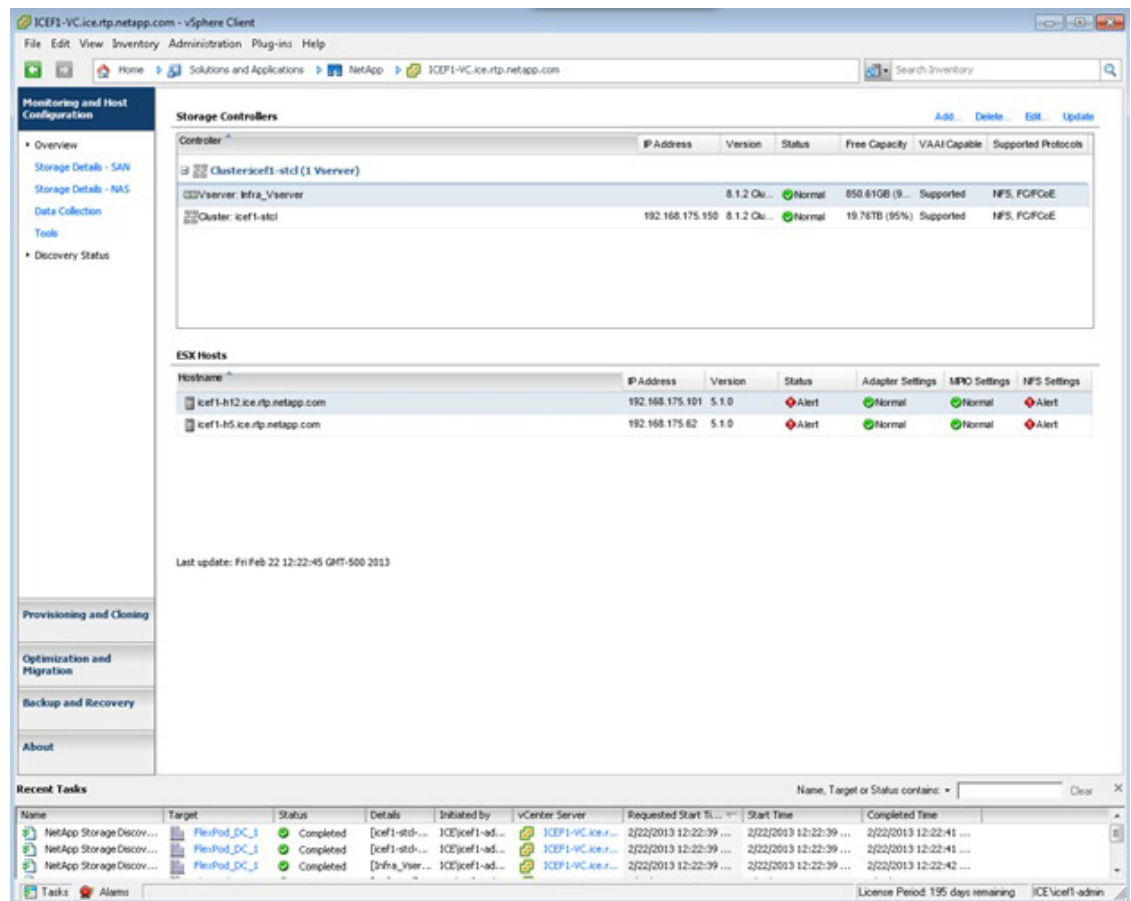
6. In the list of storage controllers, right-click the first controller listed and select Modify Credentials.
7. Enter the storage cluster management IP address in the Management IP address field. Enter admin for the User name, and the admin password for the Password. Make sure that Use SSL is selected. Click OK.



Note

If you are using 7-Mode storage, you will need to modify credentials for two storage systems. Make sure the out-of-band management IP address is being used and use the root user ID.

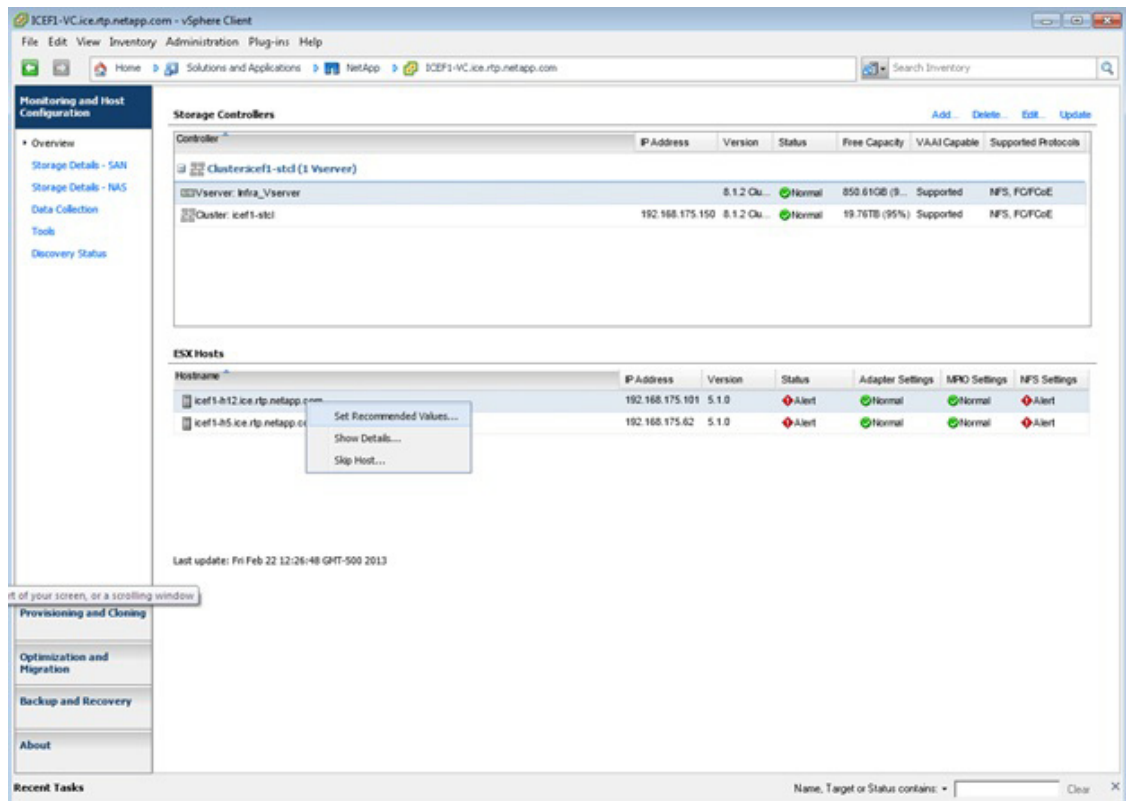
8. Click OK to accept the controller privileges.



Optimal Storage Settings for ESXi Hosts

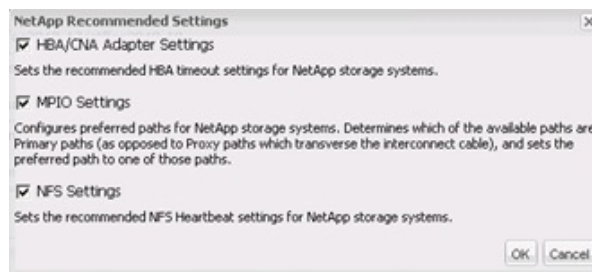
VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. Select individual or multiple ESXi hosts.
2. Right-click and select Set Recommended Values for these hosts.



3. Check the settings to apply to selected vSphere hosts. Click OK to apply the settings.

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).



Note

Depending on what changes have been made, the 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 a reboot is required, the Status value is set to Pending Reboot. If a reboot is required, the ESX or ESXi servers should be placed into Maintenance Mode, evacuated (if necessary), and restarted before proceeding.



Note

If using 7-Mode storage, it is necessary to set up storage system resources under VSC Provisioning and Clone. Go to the section "VSC 4.1 Provisioning and Cloning Setup for 7-Mode Storage" in the appendix and execute that procedure.

VSC 4.1 Backup and Recovery

Adding Storage Systems to the Backup and Recovery Capability

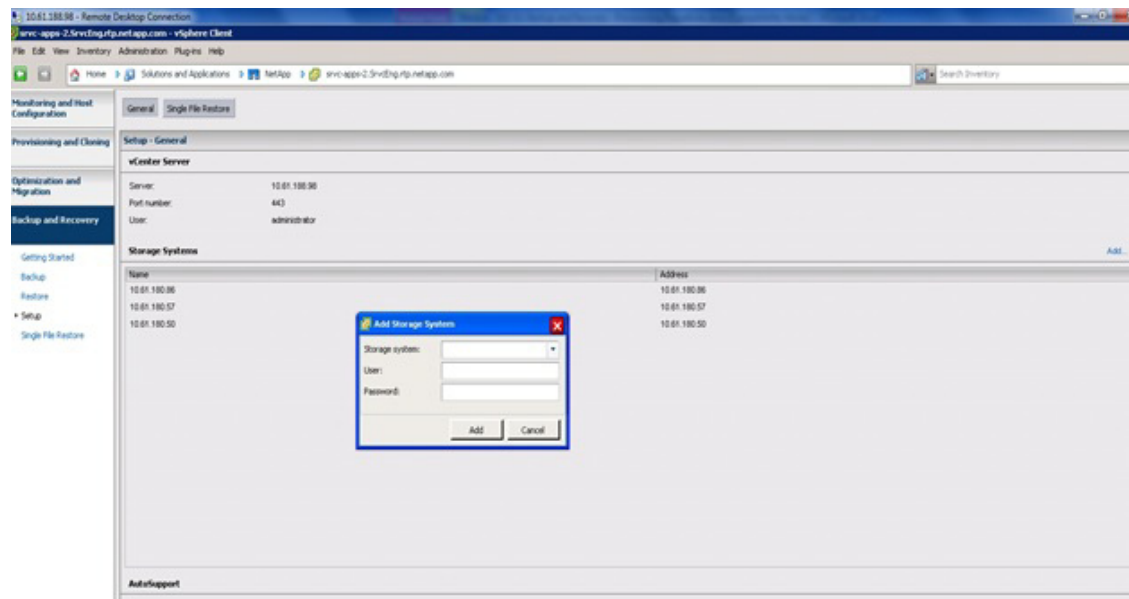
Before you begin using the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must add the storage systems that contain the datastores and virtual machines for which you are creating backups.



Note

The Backup and Recovery capability does not use the user credentials from the Monitoring and Host Configuration capability.

Follow these steps to add the storage systems to the Backup and Recovery capability.



4. Click Backup and Recovery and then Select Setup.
5. Click Add. The Add Storage System dialog box appears.
6. Type the DNS name or IP address and the user credentials of the storage cluster.
7. Click Add to add the storage cluster.



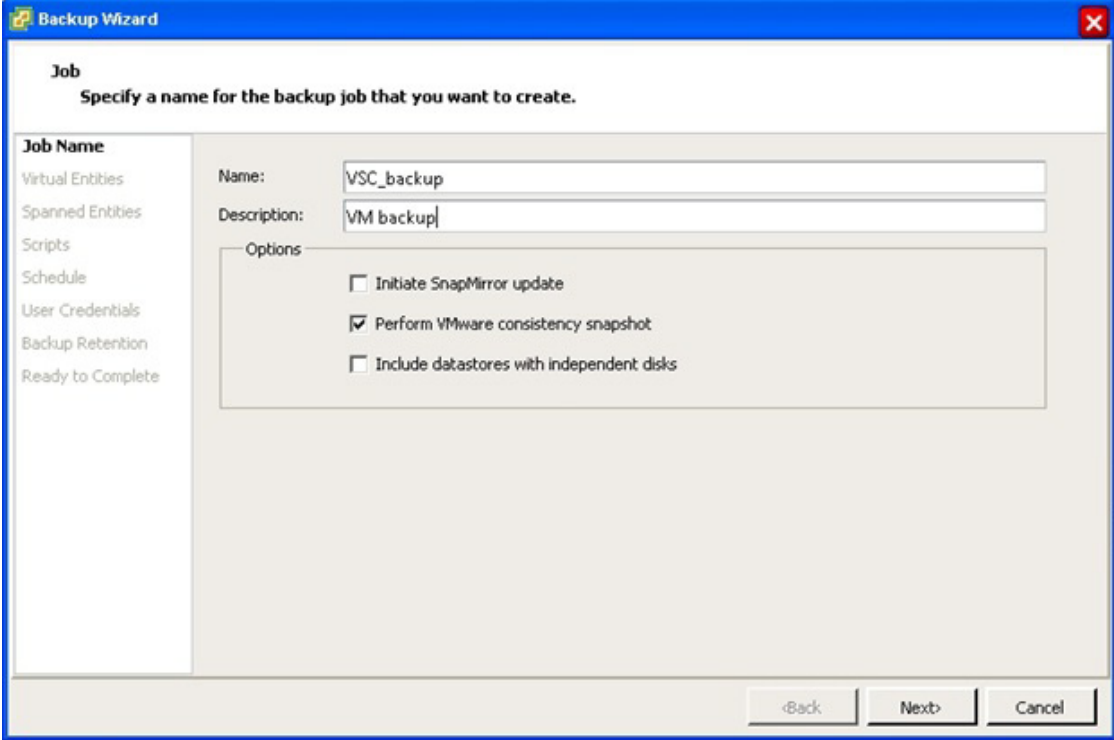
Note

If using 7-Mode storage, enter both storage controllers.

Backup and Recovery Configuration

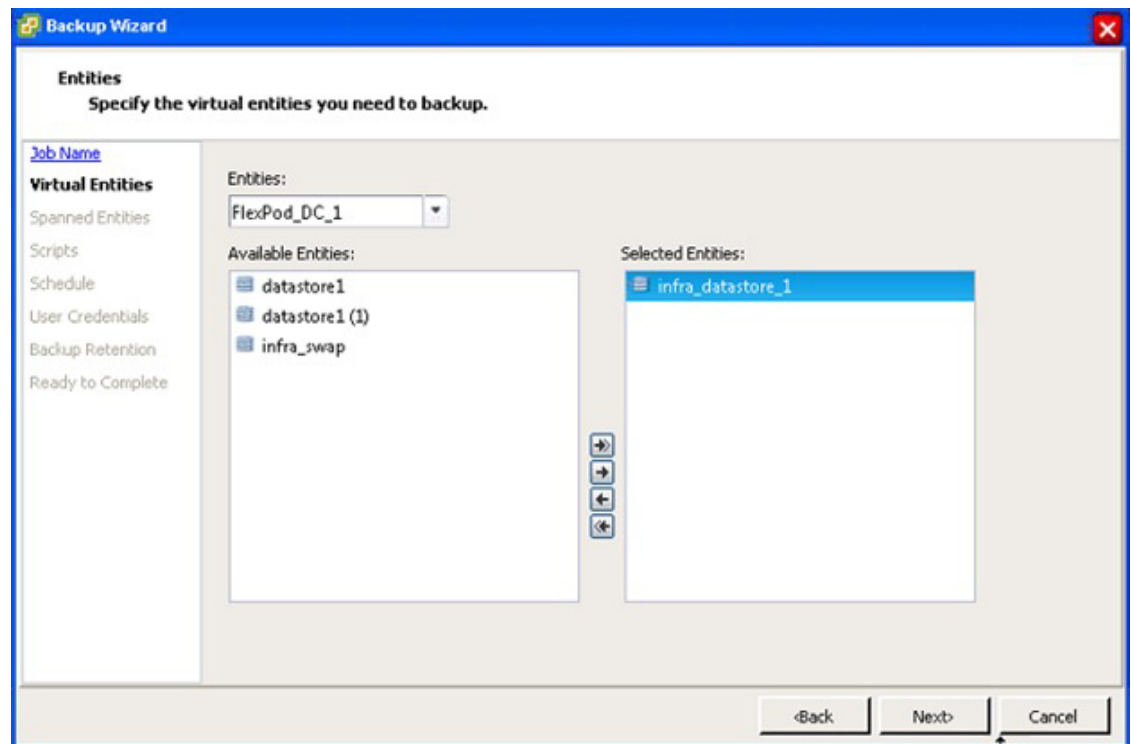
The following steps detail the procedure to configure a backup job for a datastore.

1. Click Backup and Recovery, then select Backup.
2. Click Add. The Backup wizard appears.

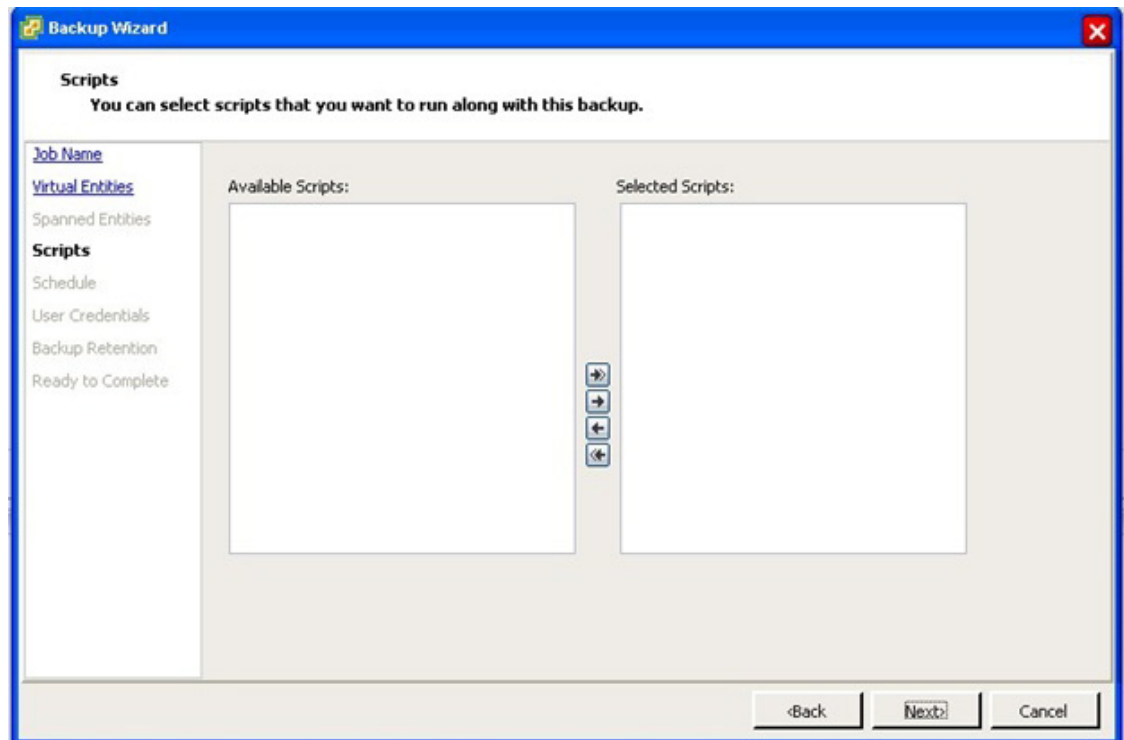


The screenshot shows the 'Backup Wizard' window with the 'Job' tab selected. The window title is 'Backup Wizard'. The main heading is 'Job' with the instruction 'Specify a name for the backup job that you want to create.' On the left is a sidebar with the following options: 'Job Name' (selected), 'Virtual Entities', 'Spanned Entities', 'Scripts', 'Schedule', 'User Credentials', 'Backup Retention', and 'Ready to Complete'. The main area contains three input fields: 'Name' with the value 'VSC_backup', 'Description' with the value 'VM backup', and an 'Options' section with three checkboxes: 'Initiate SnapMirror update' (unchecked), 'Perform VMware consistency snapshot' (checked), and 'Include datastores with independent disks' (unchecked). At the bottom right are three buttons: '<Back', 'Next>', and 'Cancel'.

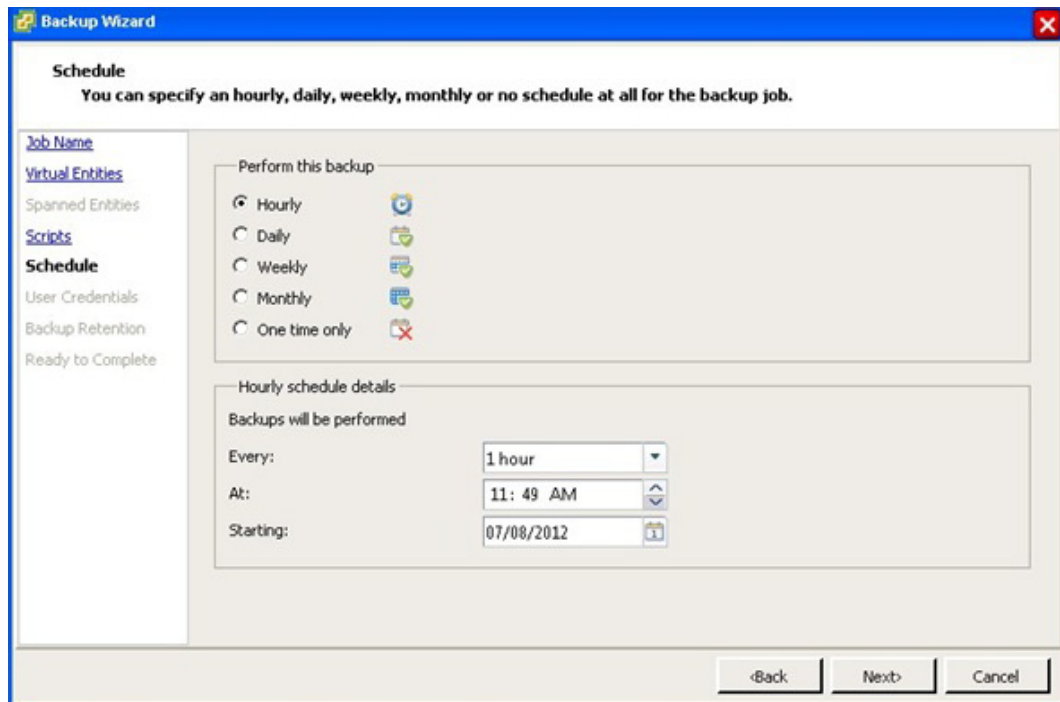
3. Type a backup job name and description.
4. If you want to create a VMware snapshot for each backup, select Perform VMware consistency snapshot in the options pane.
5. Click Next.
6. Select infra_datastore_1 and then click the ' button to move it to the selected entities. Click Next.



7. Select one or more backup scripts if available and click Next.



8. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.



9. Use the default vCenter credentials or type the user name and password for the vCenter Server and click Next.
10. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate e-mail addresses. Click Next.

Backup Wizard

Retention and Alerts
You can specify backup retention based on maximum days, maximum no of backups or backup indefinitely.

Job Name
[Virtual Entities](#)
[Spanned Entities](#)
[Scripts](#)
[Schedule](#)
[User Credentials](#)
Backup Retention
Ready to Complete

Retention

☒ A maximum of days: 1

☐ A maximum of backups: 1

☐ Never expires

Email alerts

Source email address: test1@example.com

Destination email address (s): test2@example.com

SMTP host: smtp.example.com

Notify on: Always

<Back Next> Cancel

- Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.

Backup Wizard

Summary
Review this summary before completing this wizard.

Job Name
[Virtual Entities](#)
[Spanned Entities](#)
[Scripts](#)
[Schedule](#)
[User Credentials](#)
[Backup Retention](#)
Ready to Complete

The Backup Job will be created with the following options:

Name: vsc_backup1

Description: VM backup

Perform VMware consistency snapshot: Yes

Virtual entities to be backed up: ab_esx_test
vSphere51_1

Perform this backup: Every 1 hour at 11:49 starting 7/8/2012

Backup retention: Maximum of 1 day

Email notification will be sent on: Always

Email notification will be sent from: test1@example.com

Email notification will be sent to: test2@example.com

☒ Run Job Now

<Back Finish Cancel

12. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by typing the command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

**Note**

The 7-Mode equivalent of this command is `snap sched infra_datastore_1 0 0 0`.

Also, to delete any existing automatic Snapshot copies that have been created on the volume, type the following command:

```
volume snapshot show -volume infra_datastore_1
volume snapshot delete -volume infra_datastore_1 <snapshot name>
```

**Note**

The 7-Mode equivalents of these commands are `snap list infra_datastore_1` and `snap delete infra_datastore_1 <snapshot name>`.

OnCommand Unified Manager 5.1

Create Raw Device Mapping (RDM) Datastore

From the VMware vCenter Client, do as follows:

1. In the VMware vCenter Client, from Home > Inventory > Hosts and Clusters, right-click the FlexPod_Management cluster.
2. Select NetApp > Provisioning and Cloning > Provision Datastore.
3. Make sure the Infra_Vserver is selected in Vserver drop-down menu and click Next.

**Note**

For 7-Mode storage, select controller 2 as the target storage controller.

4. Select VMFS as the Datastore type and click Next.
5. Select iSCSI as the Protocol type, set the Size to 100, enter the datastore name as RDM_Map, select the checkbox to create new volume container, select aggr02 for Aggregate, select the Thin Provision checkbox, and click Next.

**Note**

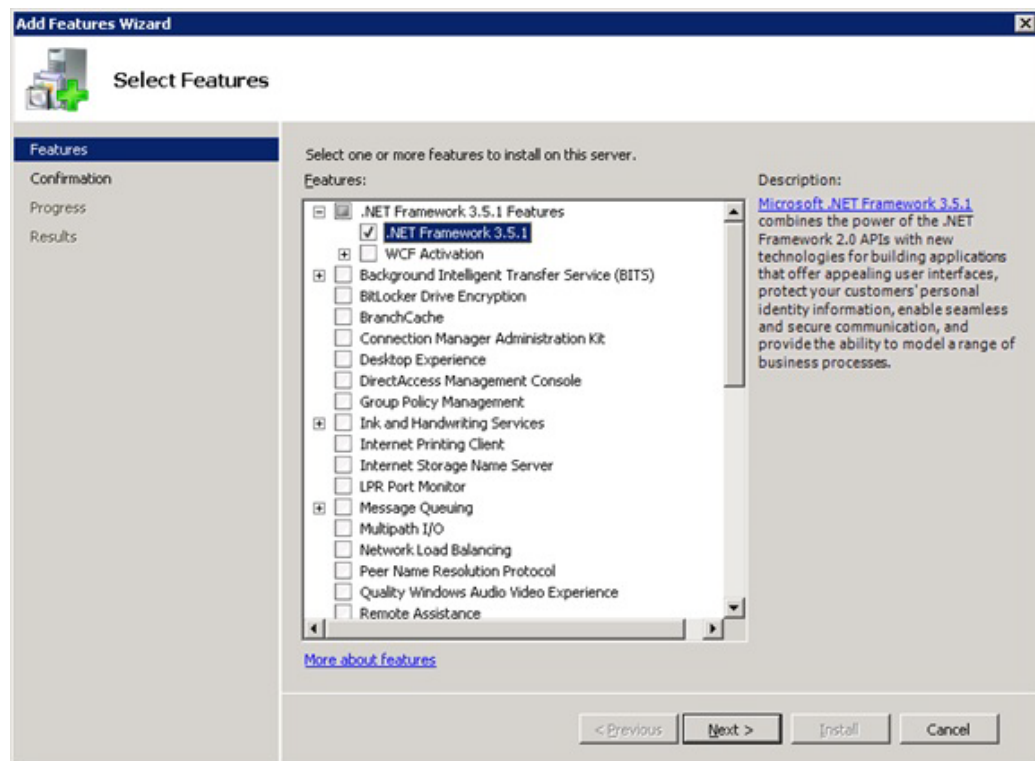
For 7-Mode storage, select aggr1 as aggregate.

6. Verify settings and click Apply.

Install .NET Framework 3.5.1 Feature

From the Virtual Storage Console (VSC) and OnCommand VM:

1. Log in to the VSC and OnCommand VM as the FlexPod admin and open Server Manager.
2. Click 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. Close Server Manager.

Install SnapDrive 6.4.2

Complete the following steps to install SnapDrive® 6.4.2.

1. Download SnapDrive 6.4.2 from the NetApp Support site.
2. Browse to the location of the SnapDrive installation package and double-click the executable file. This launches the SnapDrive installation wizard and opens the Welcome page.
3. Click Next in the Welcome page of the SnapDrive installation wizard.
4. If this is a new SnapDrive installation, read and accept the license agreement. Click Next.
5. If this is a SnapDrive upgrade, select Modify/Upgrade in the Program Maintenance page. Click Next.
6. Select "Per Storage System" as the license type. Click Next.



Note

In the case of upgrading SnapDrive, the license information will already be populated.

**Note**

In the case of selecting storage system licensing, SnapDrive can be installed without entering a license key. SnapDrive operations can be executed only on storage systems that have a SnapDrive or SnapManager license installed.

**Note**

In the case of clustered Data ONTAP 8.1-based systems, the storage system licensing for SnapDrive is bundled with the other SnapManager product licenses. They are now a single license called the SnapManager_suite license.

**Note**

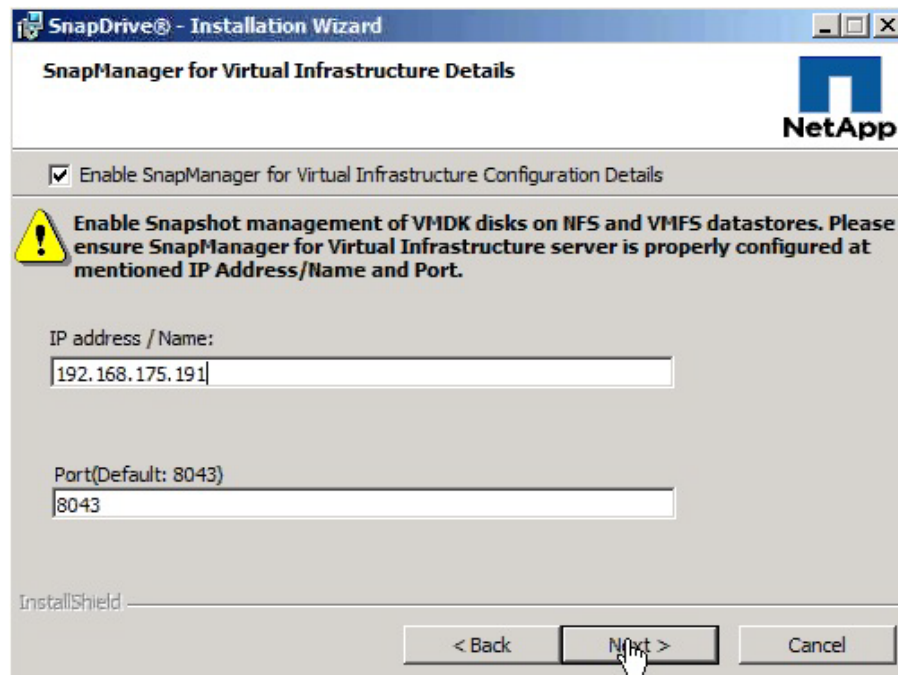
For 7-Mode storage, select Per Server, enter the SnapDrive License Key, and click Next.

7. In the Customer Information page, type the user name and organization name. Click Next.
8. The Destination Folder page prompts for a directory in which to install SnapDrive on the host. For new installations, by default this directory is C:\Program Files\NetApp\SnapDrive\. To accept the default, click Next.
9. Select the Enable VirtualCenter or ESX Server Settings checkbox. Enter the IP address, user name, and password for the vCenter Server and click Next.

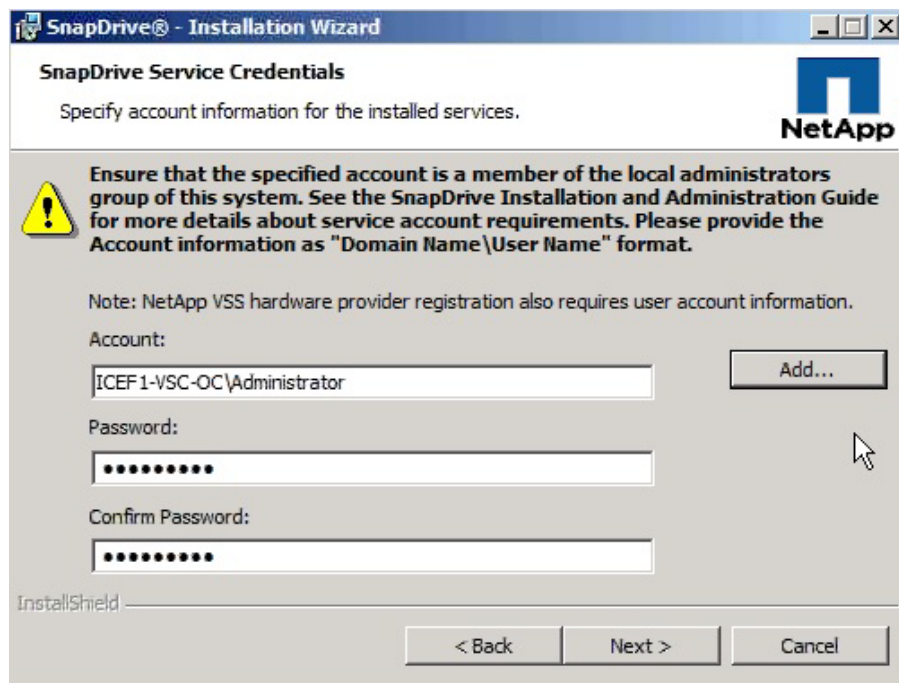
**Note**

Selecting Enable VirtualCenter or ESX Server Settings enables SnapDrive to use RDM pass-through LUNs. Select this option to use RDM pass-through disks. By default, this option is not selected.

10. Select the Enable SnapManager for Virtual Infrastructure Configuration Details checkbox. Enter the IP address of the VSC and OnCommand Server, and accept the default port. Click Next.

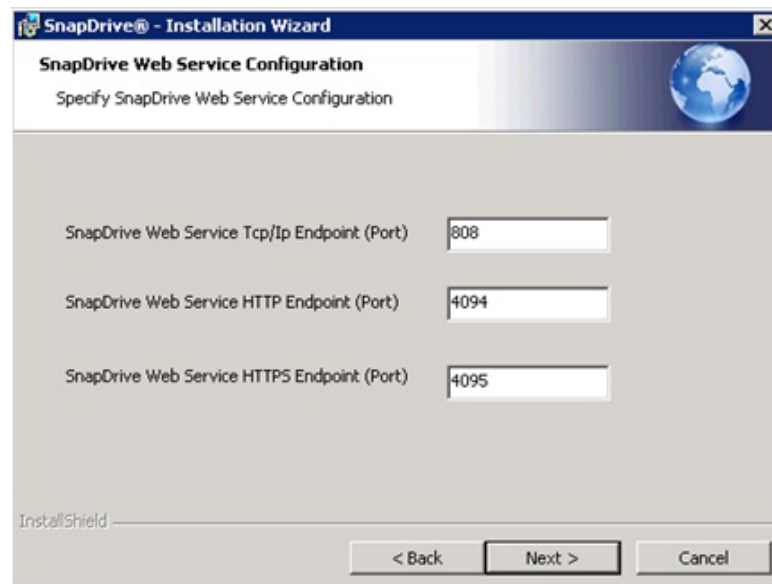


11. Type the account or complete the following steps to select a user account:
 - a. In the Enter object name to select box, enter the local machine administrator in Domain name\user name format. Click Add.
 - b. Click Check Names.
 - c. Click OK.
 - d. Enter the Administrator password.
 - e. Click Next.
 - f. Click OK.

**Note**

The specified account must be a member of the local administrators' group of this system.

12. In the SnapDrive Web Service Configuration page, keep the default ports unless any of them are already being used exclusively by another service. Click Next.



13. In the Transport Protocol Default Setting window:
 - a. Select Enable Transport Protocol Settings.
 - b. Select HTTPS as the transport protocol.

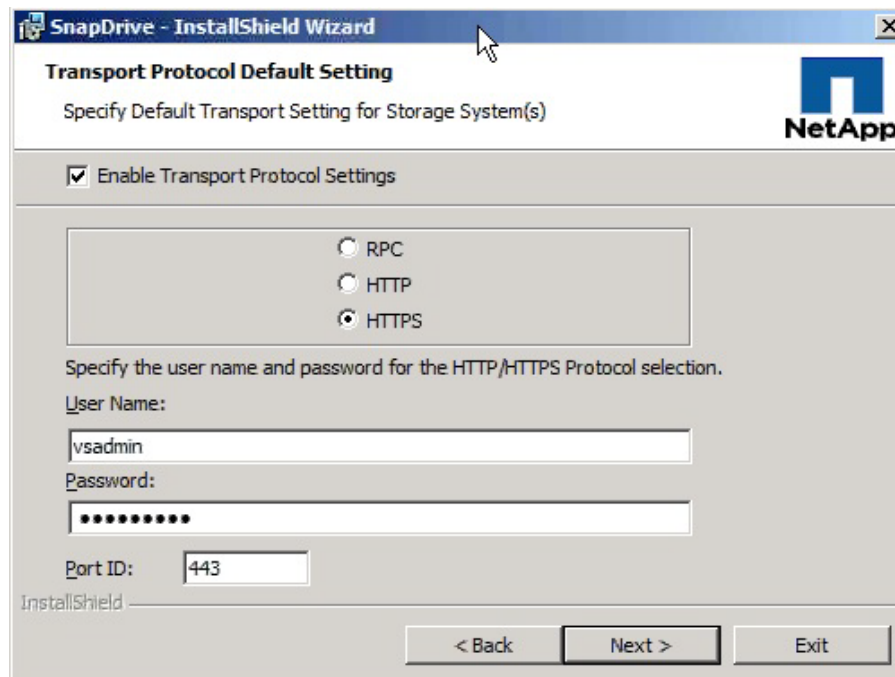
- c. Enter the user name (vsadmin) and password for the Infra_Vserver vserver.



Note

If 7-Mode storage is being used, enter root for the user name and the storage systems' root password.

- d. Verify that port ID is set to 443 and click Next.



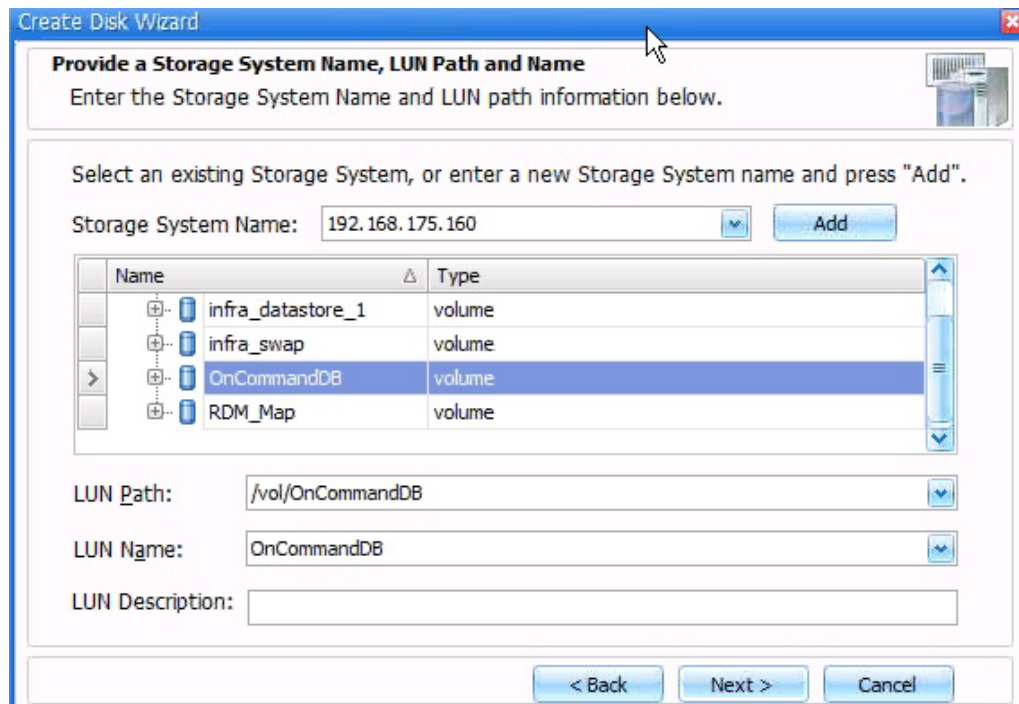
14. Click Next > Next > Install > Finish.
15. From the Start menu, open SnapDrive.
16. In the left pane, expand the local machine and select Disks.
17. In the right pane, select Create Disk.
18. In the create disk Wizard Window, click Next.
19. In the storage system name field, enter the Infra_Vserver management IP address, and click Add.



Note

If using 7-Mode storage, enter the IP address of controller 2.

20. In the list that appears, select OnCommandDB.
21. Enter OnCommandDB for the LUN Name and click Next.



Create Disk Wizard

Provide a Storage System Name, LUN Path and Name
Enter the Storage System Name and LUN path information below.

Select an existing Storage System, or enter a new Storage System name and press "Add".

Storage System Name: 192.168.175.160 Add

Name	Type
infra_datastore_1	volume
infra_swap	volume
OnCommandDB	volume
RDM_Map	volume

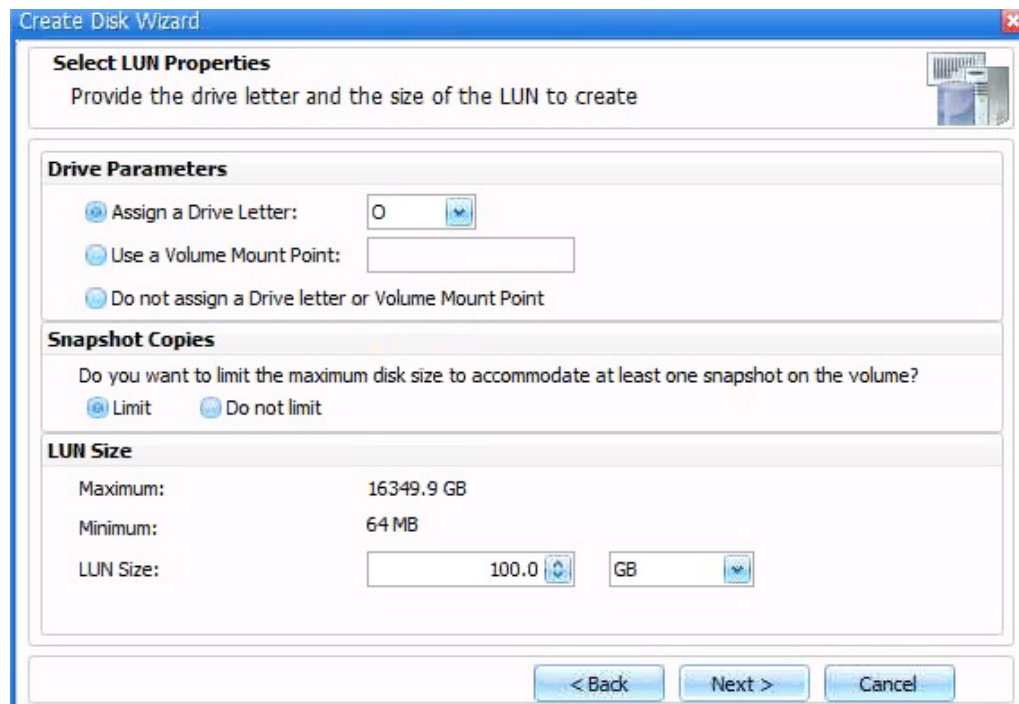
LUN Path: /vol/OnCommandDB

LUN Name: OnCommandDB

LUN Description:

< Back Next > Cancel

22. Make sure the LUN type is set to Dedicated and click Next.
23. Assign drive letter O and set LUN size to 100GB. Click Next.



Create Disk Wizard

Select LUN Properties
Provide the drive letter and the size of the LUN to create

Drive Parameters

☒ Assign a Drive Letter: O

☐ Use a Volume Mount Point:

☐ Do not assign a Drive letter or Volume Mount Point

Snapshot Copies

Do you want to limit the maximum disk size to accommodate at least one snapshot on the volume?

☒ Limit ☐ Do not limit

LUN Size

Maximum: 16349.9 GB

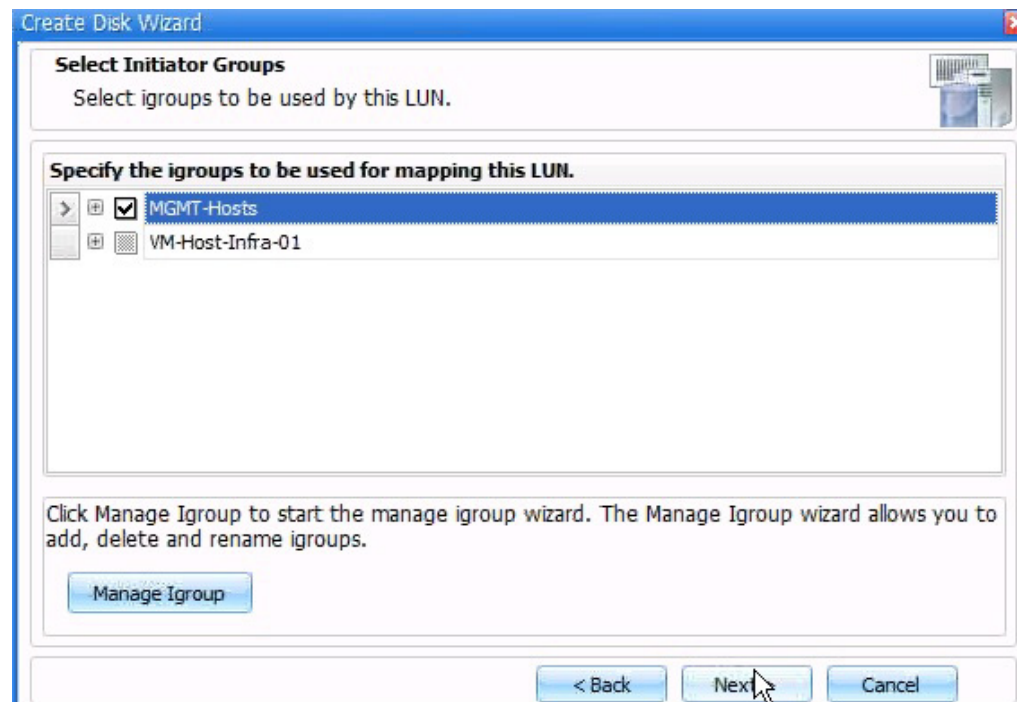
Minimum: 64 MB

LUN Size: 100.0 GB

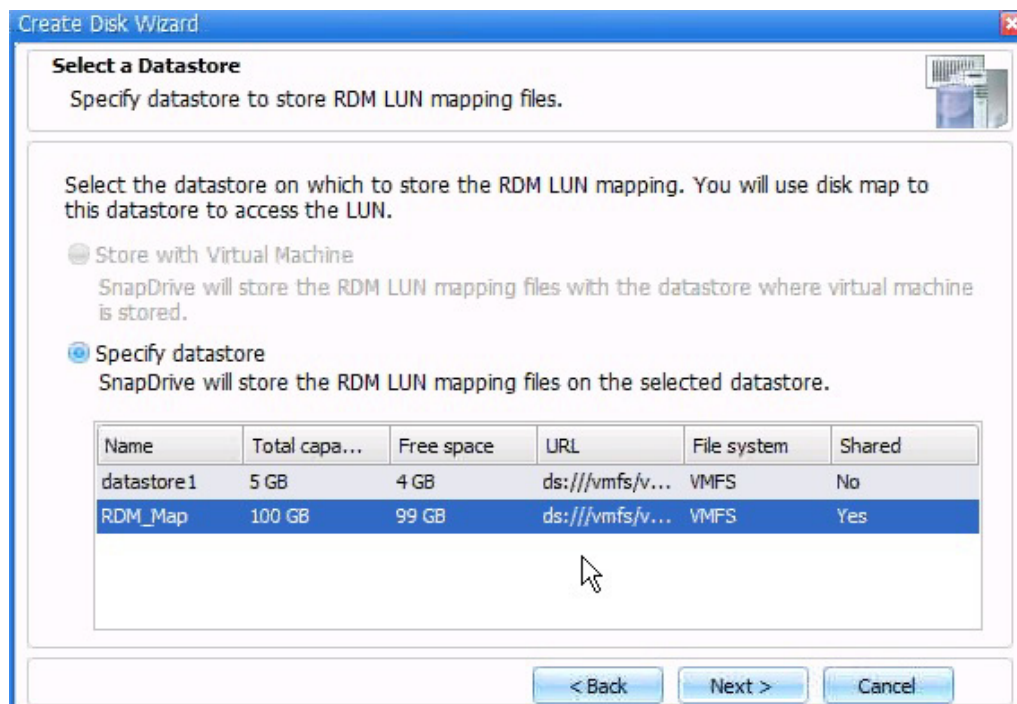
< Back Next > Cancel

24. Select all initiators on the Initiator List, and click Next.
25. Select manual as the Initiator group management, and click Next.

26. Select the MGMT-Hosts igroup, and click Next.



27. Select the RDM_Map Datastore in the Select Datastore section. Click Next.



28. Click Finish to create the disk.

29. Close SnapDrive.

**Note**

For 7-Mode storage, in vSphere Client connected to vCenter, go to the Storage Adapters window for each ESXi host. Select the iSCSI Software Adapter. Right-click the disk that was just created (the last one in the list), and select Manage Paths. Change the Path Selection to Round Robin (VMware) and click Change. Close the Manage Paths window. Change this setting on both ESXi management hosts.

Install NetApp OnCommand Core Package

To install the OnCommand Unified Manager Core Package, complete the following steps:

1. To download the OnCommand Unified Manager Core Package for Windows, click [here](#).
2. Using the FlexPod admin credentials, log in to the VSC and OnCommand VM.
3. Identify the DataFabric Manager Server license key before starting the installation.
4. Navigate to the path or directory containing the downloaded file and launch the file.
5. In the Security Warning message, click Yes to start the installation.
6. In the Welcome page, click Next.
7. Accept the AutoSupport notice and click Next.
8. Identify whether the OnCommand Unified Manager instance should manage systems with clustered Data ONTAP or 7-Mode and click Next.

**Note**

For a 7-Mode environment, either the Express edition or the Standard edition of the software is available.

**Note**

For a clustered Data ONTAP environment, only the Standard edition of the software is available.

**Note**

If the infrastructure has both 7-Mode and clustered Data ONTAP systems, two OnCommand instances are needed to manage the respective 7-Mode or clustered Data ONTAP systems.

9. Enter the 14-character license key when prompted and click Next.
10. Select the installation location, if different from the default.

**Note**

Do not change the default location of the local Temp Folder directory, or the installation will fail. The installer automatically extracts the installation files to the %TEMP% location.

11. Follow the remaining setup prompts to complete the installation.
12. From the Start Menu, right-click Command Prompt, and select Run as administrator. Run the following commands from this command prompt.
13. In preparation for the database movement to the previously created LUN from local storage, stop all OnCommand Unified Manager services and verify that the services have stopped.

```
dfm service stop
dfm service list
```

14. Move the data to the previously created LUN.

**Note**

The dfm datastore setup help command provides switch options available with the command.

```
dfm datastore setup 0:\
```

15. Start OnCommand Unified Manager and then verify that all services have started.

```
dfm service start
```

```
dfm service list
```

16. Generate an SSL key.

```
dfm ssl server setup
```

```
Key Size (minimum = 512..1024..2048..) [default=512]: 1024
```

```
Certificate Duration (days) [default=365]: Enter
```

```
Country Name (e.g., 2 letter code): <<var_country_code>>
```

```
State or Province Name (full name): <<var_state>>
```

```
Locality Name (city): <<var_city>>
```

```
Organization Name (e.g., company): <<var_org>>
```

```
Organizational Unit Name (e.g., section): <<var_unit>>
```

```
Common Name (fully-qualified hostname): <<var_oncommand_server_fqdn>>
```

```
Email Address: <<var_admin_email>>
```

**Note**

The SSL key command fails if certain command line option inputs do not follow specified character lengths (for example, a two-letter country code), and any multiword entries must be encased in double quotation marks, for example, "North Carolina."

17. Turn off automatic discovery.

```
dfm option set discoverEnabled=no
```

18. Set the protocol security options for communication with various devices.

```
dfm service stop http
```

```
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
```

**Note**

The HTTPS and SSH protocols must be enabled on the storage controllers that are monitored by OnCommand Unified Manager.

19. Restart the DataFabric Manager HTTP services to make sure that the security options take effect.

```
dfm service start http
```

20. Configure OnCommand Unified Manager to use SNMPv3 to poll configuration information from the storage devices. Use the user name and password generated for SNMPv3.

```
dfm snmp modify -v 3 -c <<var_snmp_community>> -U snmpv3user -P <<var_password>>
```

```
-A MD5 -X <<var_password>> default
```

**Note**

For 7-Mode storage, leave -X <<var_password>> out of this command. Use dfm snmp modify -v 3 -c <<var_snmp_community>> -U snmpv3user -P <<var_snmp_password>> -A MD5 default.

21. Set up OnCommand Unified Manager to send AutoSupport through HTTPS to NetApp.

```
dfm option set SMTPServerName=<<var_mailhost>>
```

```
dfm option set autosupportAdminContact=<<var_storage_admin_email>>
```

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

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

**Note**

For 7-Mode storage, add both storage controllers.

22. Manually add the storage cluster to the OnCommand server.

```
dfm host add <<var_clustername>>
```

23. Set the array login and password credentials in OnCommand Unified Manager. This is the root or administrator account.

```
dfm host set <<var_clustername>> hostlogin=admin
```

```
dfm host set <<var_clustername>> hostPassword=<<var_password>>
```

**Note**

For 7-Mode storage, set credentials for both storage controllers.

24. List the storage systems discovered by OnCommand Unified Manager and their properties.

```
dfm host list
```

```
dfm host get <<var_clustername>>
```

25. Test the network configuration and connectivity between the OnCommand server and the named host. This test helps identify misconfigurations that prevent the OnCommand server from monitoring or managing a particular appliance. The test should be the first command used if a problem using the OnCommand server occurs with only some of the appliances.

```
dfm host diag <<var_clustername>>
```

26. Configure an SNMP trap host (optional).

```
dfm alarm create -T <<var_oncommand_server_fqdn>>
```

27. Configure OnCommand Unified Manager to generate and send e-mails for every event whose importance ranks as critical or higher.

```
dfm alarm create -E <<var_storage_admin_email>> -v Critical
```

28. Create a manual backup.

```
dfm backup create -t snapshot
```

29. Schedule backups to a virtual backup directory on the 100GB iSCSI LUN.

```
dfm option set backupRetentionCount=20
```

```
dfm backup schedule set -t snapshot -D 21:00
```

30. To open Windows Firewall with Advanced Security, click Start > Administrative Tools > Windows Firewall with Advanced Security.

31. Select Inbound Rules.

32. Click New Rule.

33. Select Port and click Next.

34. Leave TCP selected and enter 8443 in the Specific local ports textbox. Click Next.

35. Click Next.

36. Click Next.

37. Name the rule OnCommand Console External Access and click Finish.

38. Click New Rule.

39. Select Port and click Next.

40. Select UDP and enter 162 in the Specific local ports textbox. Click Next.

41. Click Next.

42. Click Next.
43. Name the rule OnCommand SNMP Trap and click Finish.
44. Close Windows Firewall with Advanced Security.

NetApp NFS Plug-In 1.0 for VMware VAAI

Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in clustered Data ONTAP, complete the following steps:

1. From an SSH session to the storage cluster management address, log in with the admin user name and password.
2. Enable vStorage on the Vserver.

```
vserver nfs modify -vserver Infra_Vserver -vstorage enabled
```



Note

For 7-Mode storage, enter options `nfs.vstorage.enable` on both storage controllers. Steps 3 and 4 are not necessary when using 7-Mode storage.

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show -vserver Infra_Vserver
```

4. The access protocol for the FlexPod policy name should be NFS. If the access protocol is not NFS for a given rule index, run the following command to set NFS as the access protocol:

```
vserver export-policy rule modify -vserver Infra_Vserver -policyname FlexPod  
-ruleindex <<var_rule_index>> -protocol nfs
```

Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), complete the following steps:

1. From the vSphere console of the VSC and OnCommand virtual machine (VM), go to the Software Downloads page in the NetApp Support site.
2. Scroll down to locate the NetApp NFS Plug-in for VMware VAAI, select the ESXi platform, and click Go.
3. Download the .vib file of the most recent plug-in version.
4. Verify that the file name of the .vib file matches the predefined name that VSC 4.1 for VMware vSphere uses: `NetAppNasPlugin.vib`.



Note

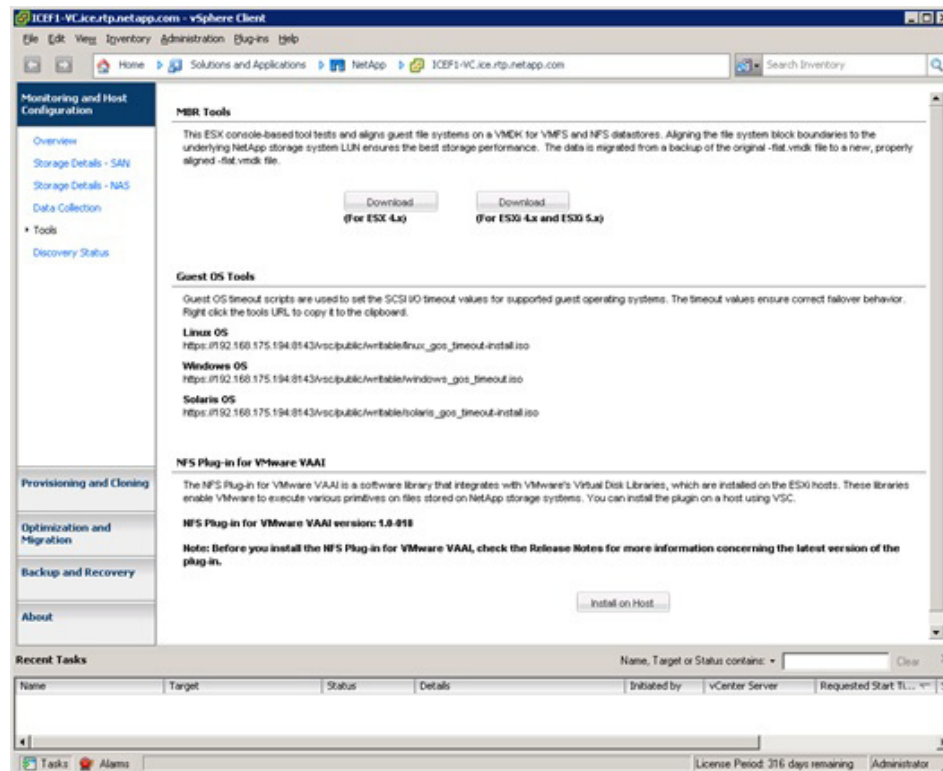
If the .vib file name does not match the predefined name, rename the .vib file. Neither the VSC client nor the NetApp vSphere Plug-in Framework (NVPF) service needs to be restarted after the .vib file is renamed.

5. Copy the plug-in .vib file (`NetAppNasPlugin.vib`) to `C:\Program Files\Virtual Storage Console\etc\vsc\web`.

**Note**

The default directory path is C:\Program Files\NetApp\Virtual Storage Console\. However, VSC 4.1 for VMware vSphere lets you change this directory. For example, if you are using the default installation directory, the path to the NetAppNasPlugin.vib file is the following: C:\Program Files\Virtual Storage Console\etc\vsc\web\NetAppNasPlugin.vib.

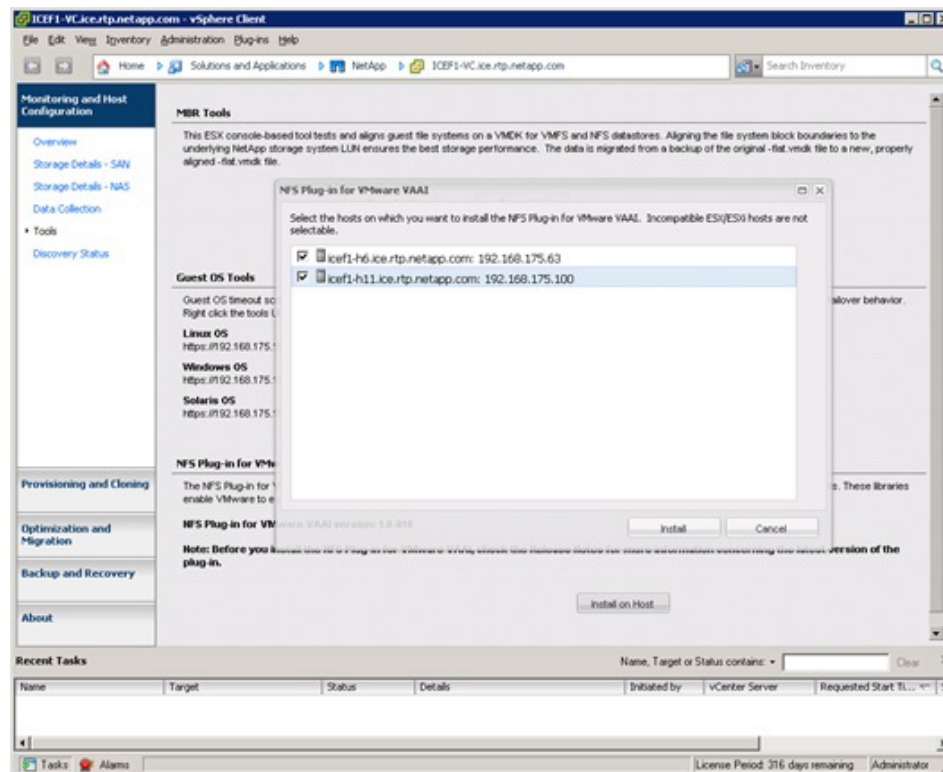
6. In the VMware vSphere Client connected to the vCenter Server, select Home > Solutions and Applications > NetApp.
7. In the Monitoring and Host Configuration capability navigation pane, select Tools.
8. Under NFS Plug-in for VMware VAAI, click Install on Host.



9. Select all ESXi hosts and click Install, and then click Yes.

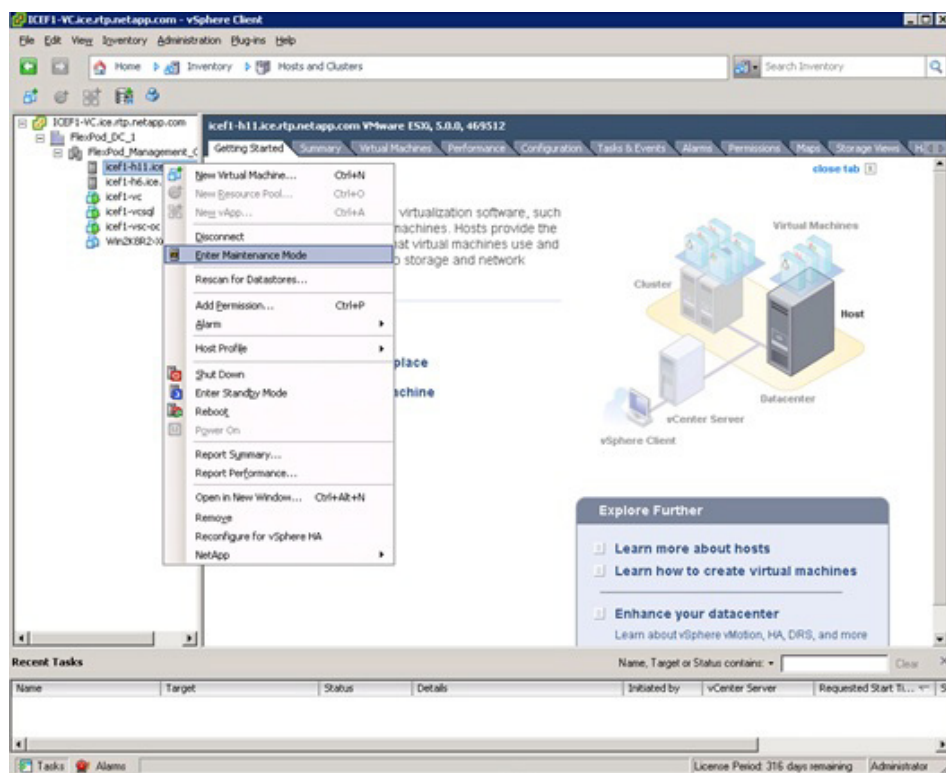
**Note**

The Monitoring and Host Configuration capability automatically installs the plug-in on the hosts selected.



10. Select Home > Inventory > Host and Clusters.

11. For each host (one at a time), right-click the host and select Enter Maintenance Mode.



12. Click Yes, click Yes again, and then click OK.

**Note**

It might be necessary to migrate all VMs away from the host.

13. After the host is in maintenance mode, right-click the host and select Reboot.
14. Enter a reason for the reboot and click OK.
15. After the host reconnects to the vCenter Server, right-click the host and select Exit Maintenance Mode.
16. Make sure that all ESXi hosts get rebooted.

Appendix

Build Windows Active Directory Server Virtual Mahines

ESXi Host VM-Host-Infra-01

To build an Active Directory Server virtual machine (VM) for the VM-Host-Infra-01 ESXi host, complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the vSphere Client, select the host in the inventory pane.
3. Right-click the host and select New Virtual Machine.

4. Select Custom and click Next.
5. Enter a name for the VM. Click Next.
6. Select infra_datastore_1. Click Next.
7. Select Virtual Machine Version: 8. Click Next.
8. Verify that the Windows option and the Microsoft Windows Server 2008 R2 (64-bit) version are selected. Click Next.
9. Select two virtual sockets and one core per virtual socket. Click Next.
10. Select 4GB of memory. Click Next.
11. Select one network interface card (NIC).
12. For NIC 1, select the IB-MGMT Network option and the VMXNET 3 adapter. Click Next.
13. Keep the LSI Logic SAS option for the SCSI controller selected. Click Next.
14. Keep the Create a New Virtual Disk option selected. Click Next.
15. Make the disk size at least 60GB. Click Next.
16. Click Next.
17. Select the checkbox for Edit the Virtual Machine Settings Before Completion. Click Continue.
18. Click the Options tab.
19. Select Boot Options.
20. Select the Force BIOS Setup checkbox.
21. Click Finish.
22. From the left pane, expand the host field by clicking the plus sign (+).
23. Right-click the newly created AD Server VM 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 then select Connect to ISO Image on Local Disk.
26. Navigate to the Windows Server 2008 R2 SP1 ISO, select it, and click Open.
27. Click in the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
28. The Windows Installer boots. Select the appropriate language, time and currency format, and keyboard. Click Next.
29. Click Install now.
30. Make sure that the Windows Server 2008 R2 Standard (Full Installation) option is selected. Click Next.
31. Read and accept the license terms and click Next.
32. Select Custom (Advanced). Make sure that Disk 0 Unallocated Space is selected. Click Next to allow the Windows installation to complete.
33. After the Windows installation is complete and the VM has rebooted, click OK to set the Administrator password.
34. Enter and confirm the Administrator password and click the blue arrow to log in. Click OK to confirm the password change.

35. After logging in to the VM desktop, from the VM console window, select the VM menu. Under Guest, select Install/Upgrade VMware Tools. Click OK.
36. If prompted to eject the Windows installation media before running the setup for the VMware tools, click OK, then click OK.
37. In the dialog box, select Run setup64.exe.
38. In the VMware Tools installer window, click Next.
39. Make sure that Typical is selected and click Next.
40. Click Install.
41. Click Finish.
42. Click Yes to restart the VM.
43. After the reboot is complete, select the VM menu. Under Guest, select Send Ctrl+Alt+Del. Then enter the password to log in to the VM.
44. Set the time zone for the VM, IP address, gateway, and host name.

**Note**

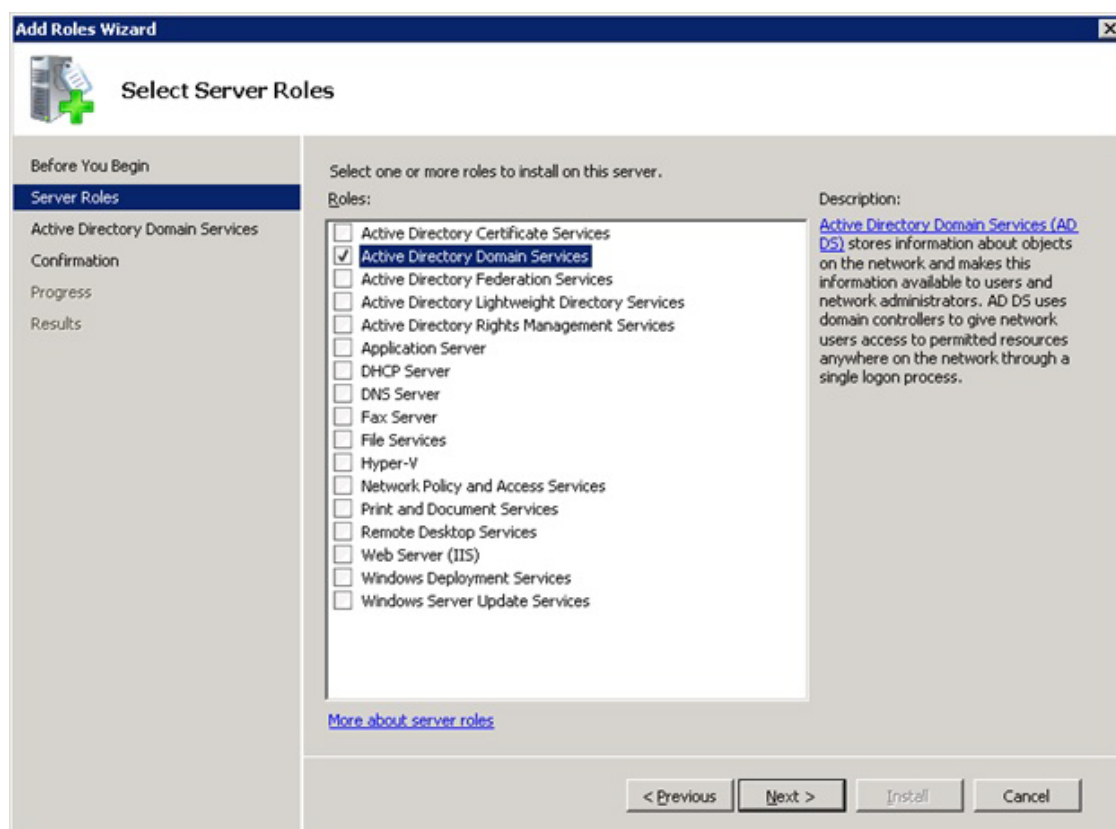
A reboot is required.

45. If necessary, activate Windows.
46. Download and install all required Windows updates.

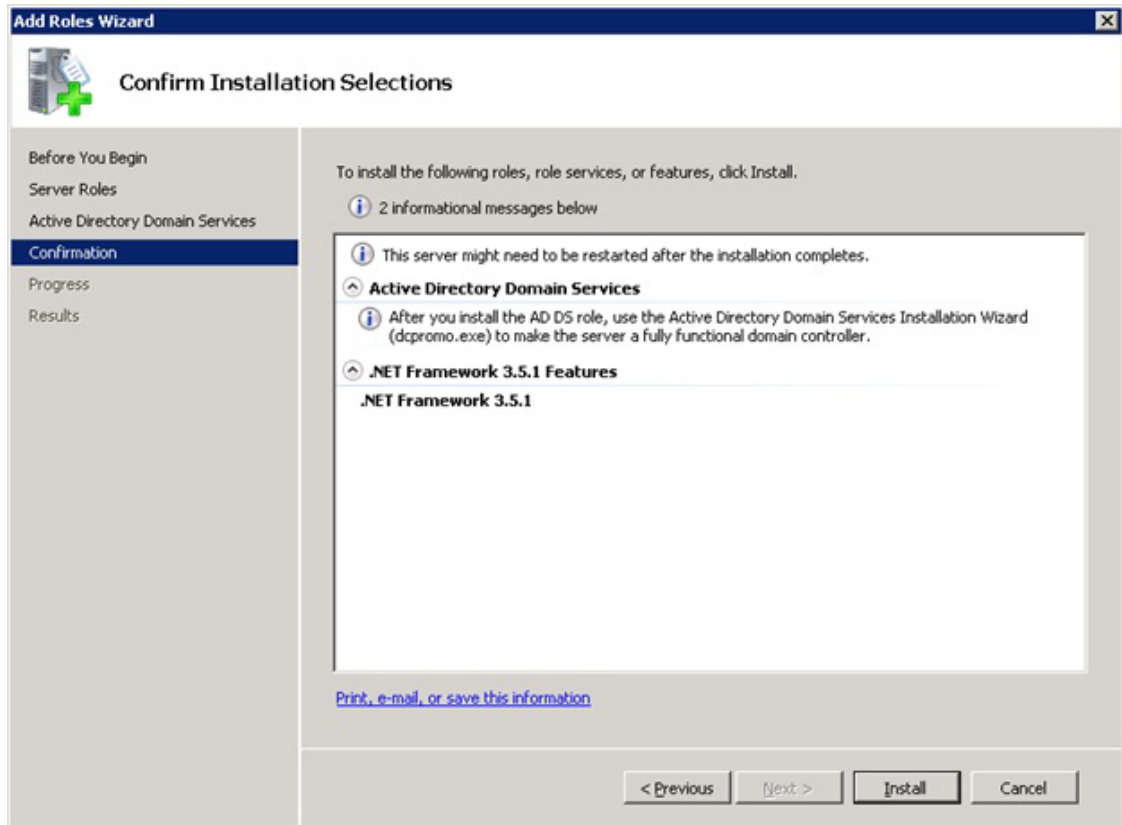
**Note**

This process requires several reboots.

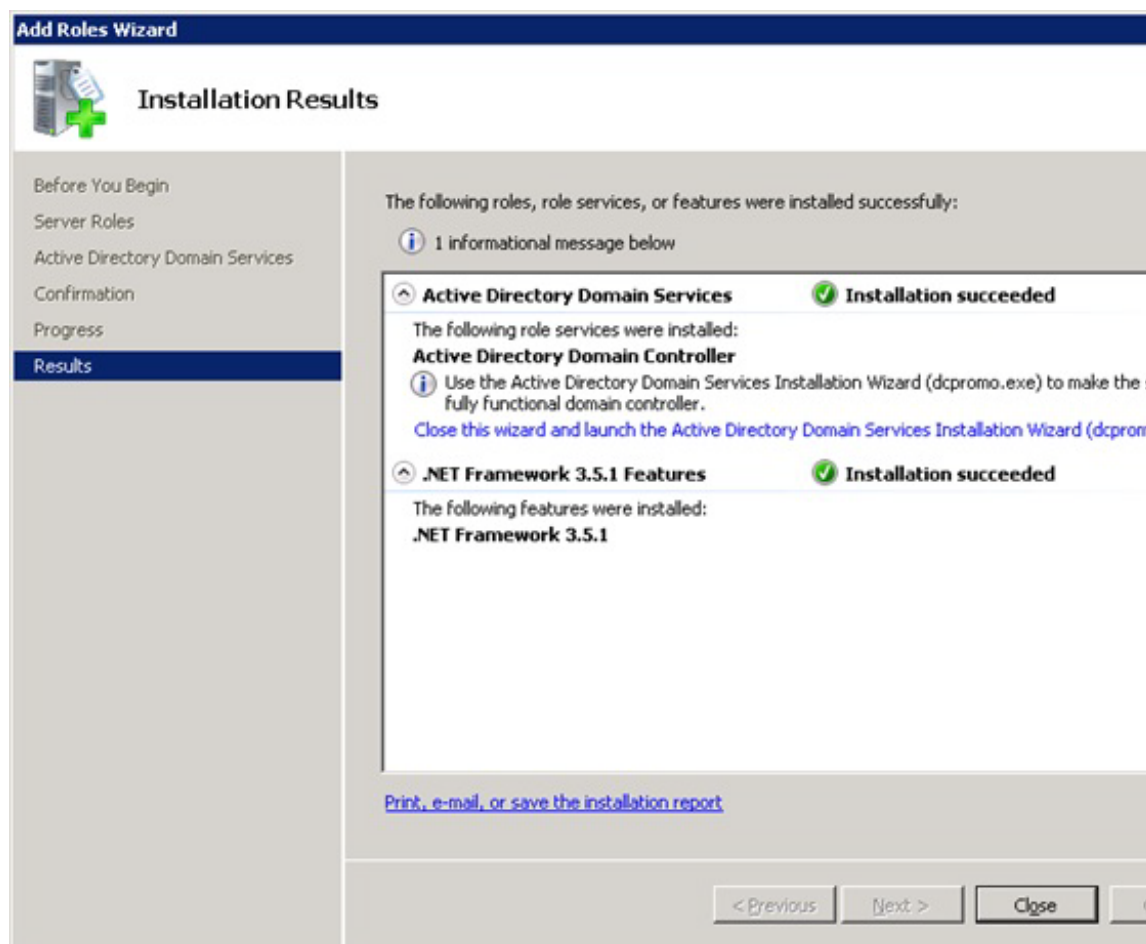
47. Open Server Manager.
48. On the left, click Roles, then select Add Roles on the right.
49. Click Next.
50. In the list, select the checkbox next to Active Directory Domain Services.
51. In the popup, click Add Required Features to add .NET Framework 3.5.1.



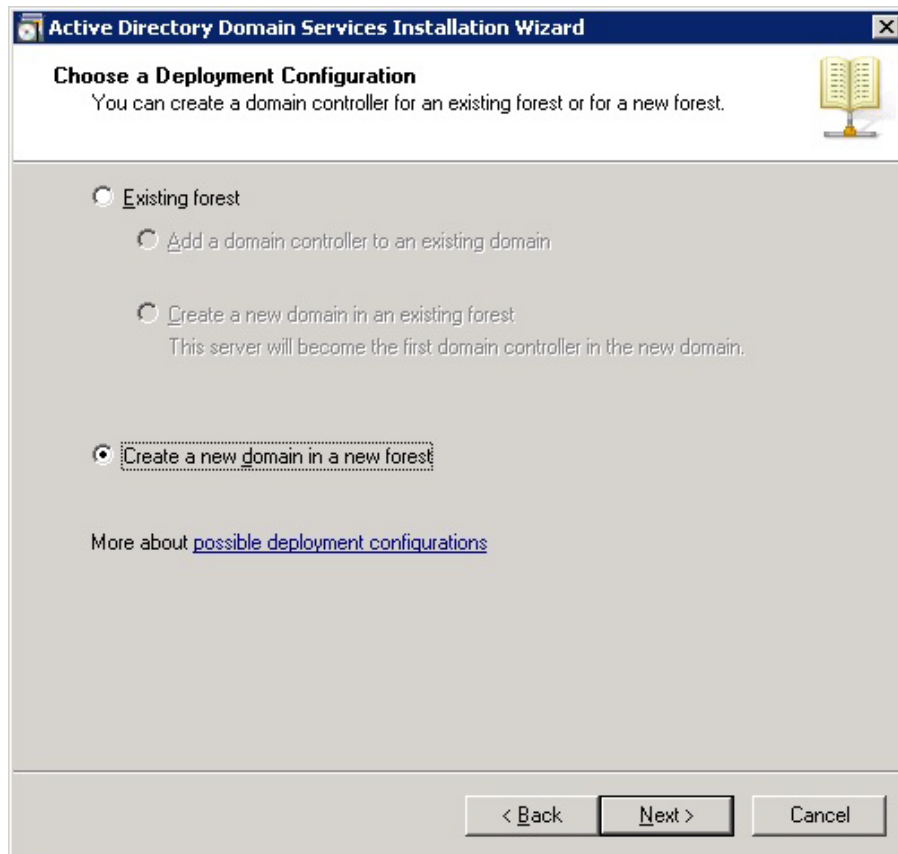
52. Click Next.



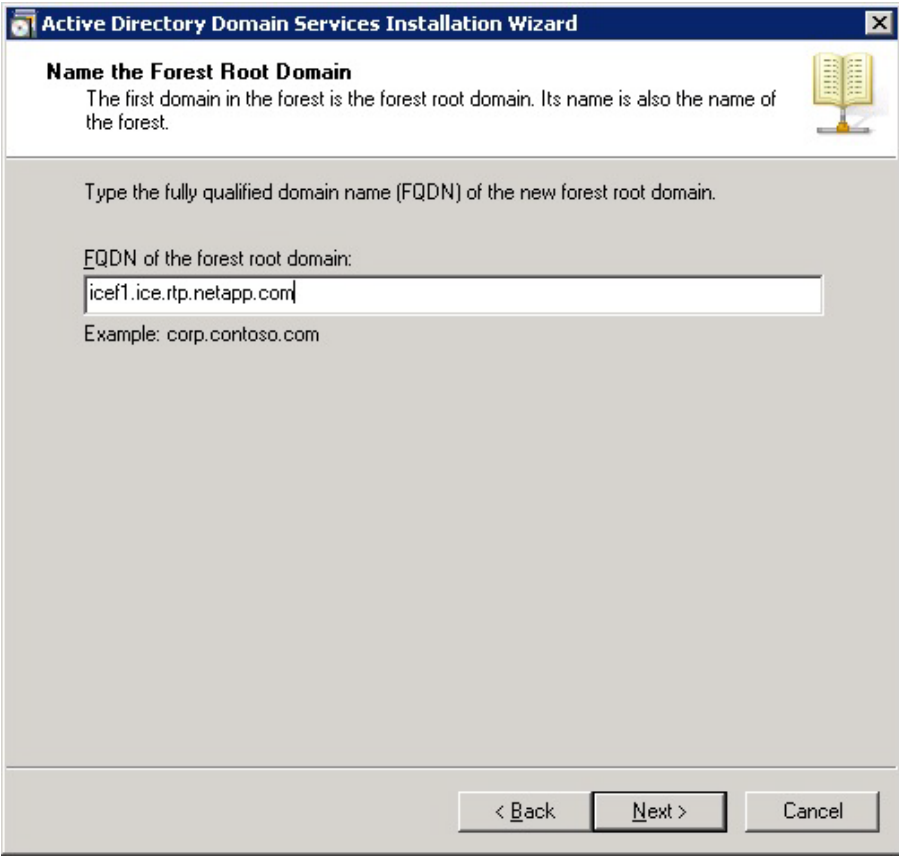
53. Click Install.



54. In the middle of the window, click Close this wizard and launch the Active Directory Domain Services Installation Wizard (dcpromo.exe).
55. In the Active Directory Domain Services Installation Wizard, click Next.
56. Click Next.
57. Select Create a new domain in a new forest and click Next.



58. Type the FQDN of the Windows domain for this FlexPod and click Next.



The screenshot shows the 'Active Directory Domain Services Installation Wizard' window. The title bar is blue with the text 'Active Directory Domain Services Installation Wizard' and a close button. The main area has a light gray background. At the top, the section 'Name the Forest Root Domain' is highlighted in a darker gray. Below this, a text box contains the instruction: 'The first domain in the forest is the forest root domain. Its name is also the name of the forest.' To the right of this text is a small icon of an open book. Below the instruction, a larger text box prompts the user: 'Type the fully qualified domain name (FQDN) of the new forest root domain.' Underneath this, a label reads 'FQDN of the forest root domain:' followed by a text input field containing the value 'icef1.ice.rtp.netapp.com'. Below the input field, an example is provided: 'Example: corp.contoso.com'. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Name the Forest Root Domain

The first domain in the forest is the forest root domain. Its name is also the name of the forest.

Type the fully qualified domain name (FQDN) of the new forest root domain.

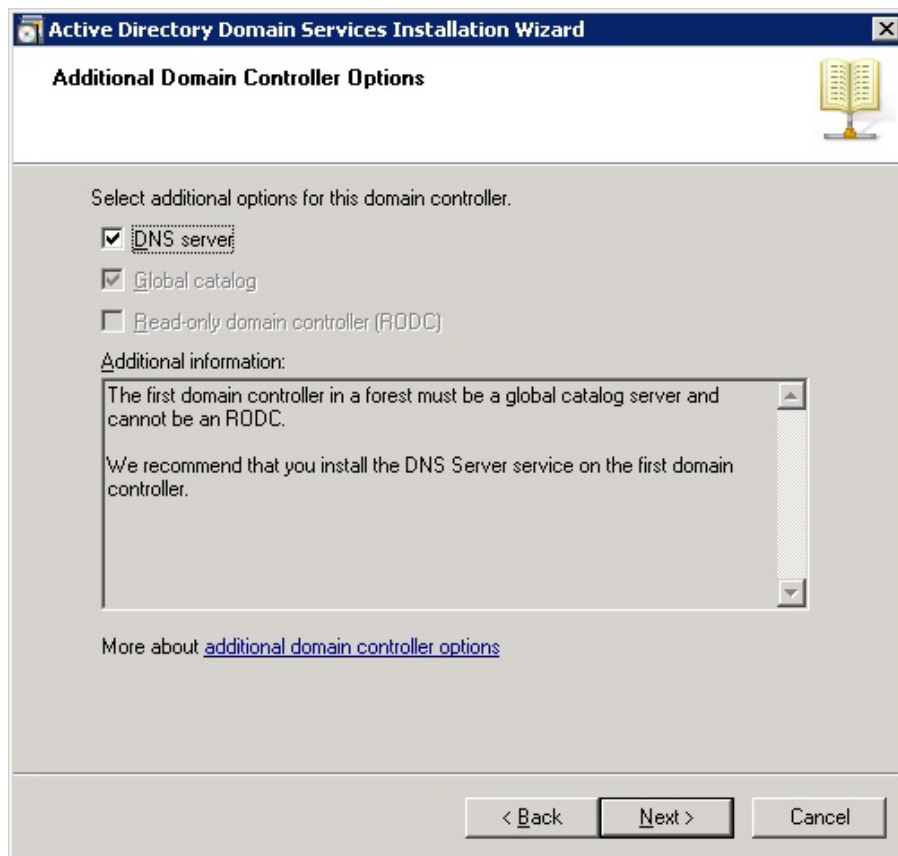
FQDN of the forest root domain:

icef1.ice.rtp.netapp.com

Example: corp.contoso.com

< Back Next > Cancel

59. Select the appropriate forest functional level and click Next.
60. Keep DNS server selected and click Next.

**Note**

If one or more DNS servers exist that this domain can resolve from, select Yes to create a DNS delegation. If this is AD server is being created on an isolated network, select No, to not create a DNS delegation. The remaining steps in this procedure assume a DNS delegation is not created.

61. Click Next.
62. Click Next to accept the default locations for database and log files.
63. Enter and confirm <<var_password>> for the Directory Services Restore Mode Administrator Password. Click Next.
64. Review the Summary information and click Next. Active Directory Domain Services will install.
65. Click Finish.
66. Click Restart Now to restart the AD Server.
67. After the machine has rebooted, log in as the domain Administrator.
68. Open the DNS Manager by clicking Start > Administrative Tools > DNS.
69. Optional: Add Reverse Lookup Zones for your IP address ranges.
70. Expand the Server and Forward Lookup Zones. Select the zone for the domain. Right-click and select New Host (A or AAAA). Populate the DNS Server with Host Records for all components in the FlexPod.

71. Optional: Build a second AD server VM. Add this server to the newly created Windows Domain and activate Windows. Install Active Directory Domain Services on this machine. Launch dcpromo.exe at the end of this installation. Choose to add a domain controller to a domain in an existing forest. Add this domain controller to the domain created earlier. Complete the installation of this second domain controller. After vCenter Server is installed, affinity rules can be created to keep the two AD servers running on different hosts.

Configuring Cisco VM-FEX with the UCS Manager

Background

FlexPod for VMware utilizes distributed virtual switching to manage the virtual access layer from a central point. While previous versions of FlexPod have only described the use of the Cisco Nexus 1000V, there exists an option to use the built-in virtual switching functionality delivered through hardware on the Cisco Unified Computing System known as VM-FEX. This has several advantages:

- There is no need for extra HW such as Cisco Nexus 1110-X.
- Cisco UCS provides a central configuration environment with which the administrator is already familiar.
- Compared to using the Cisco Nexus 1000v as virtual appliances within vCenter itself, this setup avoids an SPOF and common restart issues when running the distributed switches in an environment in which they are required for the network functionality of the ESX servers on which they are running. This is a common problem that needs to be addressed in the solution design.

In other words, it dramatically simplifies the hardware setup and operation by optimally utilizing the new hardware features.

Process Overview

This section provides a detailed overview of VM-FEX setup, configuration, and operation using Cisco UCS Manager.

This section describes:

- Initial setup and configuration
- Operation, that is, adding networks for additional tenants

For configuration guide, go to

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/2.1/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide_2_1.html.

Initial Setup

The initial setup is a five-step procedure:

1. Create a vNIC connection policy in Cisco UCS Manager.
2. Create a server BIOS policy.
3. Clone an existing service profile.
4. Install the VEM software on the ESX server.
5. Install the plug-in into vCenter.

Create a Dynamic vNIC Connection Policy

To define the dynamic vNIC connection policy that vNICs created from a vNIC template should use, complete the following steps in Cisco UCS Manager:

1. Log in to Cisco UCS Manager.
2. Select the LAN tab in the left navigation pane and click LAN > Policies > root > Sub-organizations > (name of the suborganization if applicable) > Dynamic vNIC Connection Profile.
3. Right-click and select Create Dynamic vNIC Connection Policy to start the wizard.
4. Type a name and description for the vNIC connection policy. Select VMWare from the Adapter Policy drop-down menu. Select the Protected option. Click OK.



Note

The Protected option allows the vNIC to use both fabric A and fabric B.



Note

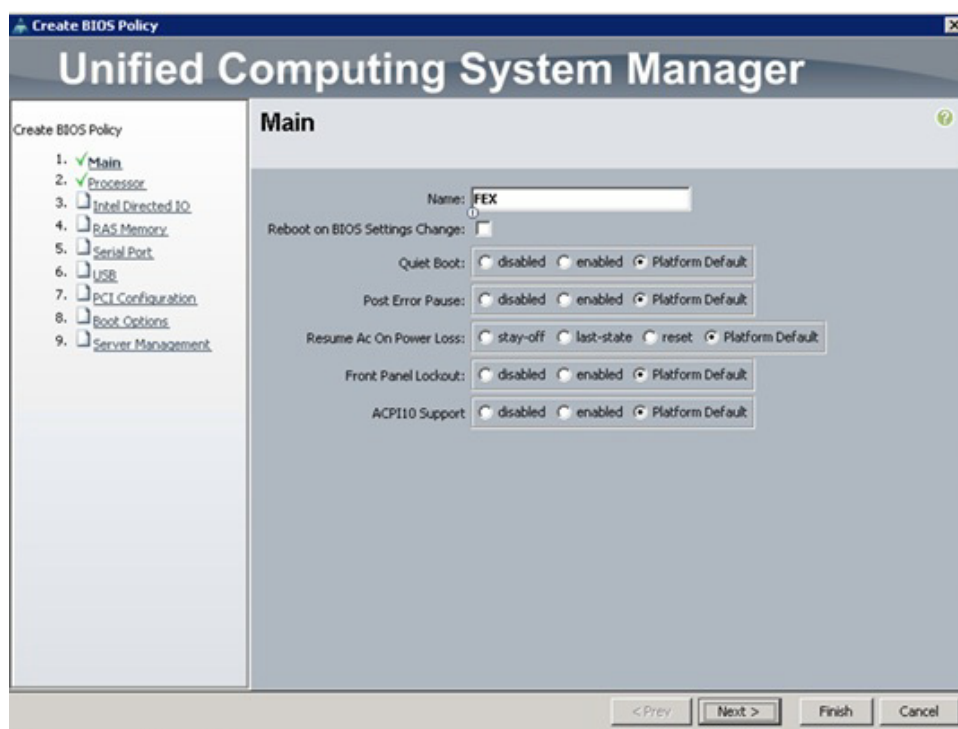
With Cisco UCS C-Series servers, the number of dynamic vNICs that can be used depends on the hardware in use. Refer to section 14.3, "VM-FEX Virtual Interfaces" in the appendix.



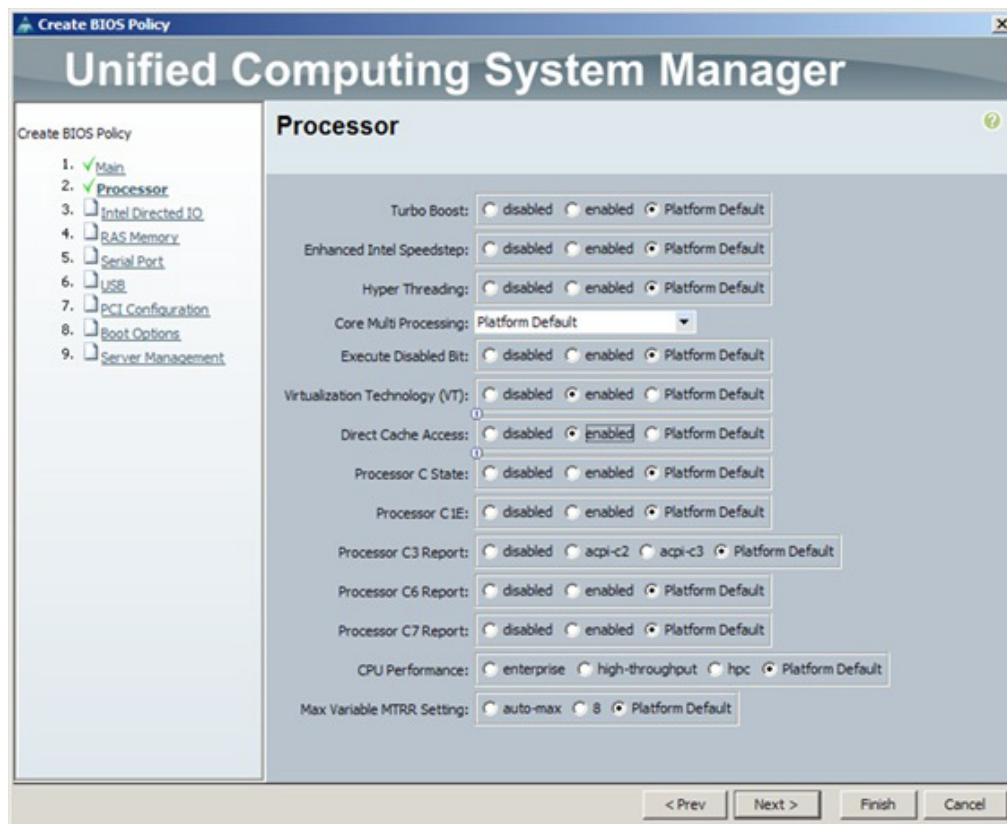
Create a Server BIOS Policy

To define the BIOS policy for a service profile that supports VM-FEX on ESXi, complete the following steps in Cisco UCS Manager:

1. Select the Server tab in the left navigation pane, and click Server > Policies > root > Sub-organizations (name of the suborganization if applicable) > BIOS Policies.
2. Right-click and select Create BIOS Policy to start the wizard.
3. Type a name for the policy and retain the platform defaults.



4. For Virtualization Technology (VT) and Direct Cache Access, select enabled.



5. Click Next.
6. For VT For Directed IO, select enabled.




7. Click Next.
8. The remaining sections of the Create BIOS Policy wizard (RAS Memory, Serial Port, USB, PCI Configuration, Boot Options, and Server Management) can retain the Platform Default option. Click Next on each of these windows and then click Finish to complete the wizard.

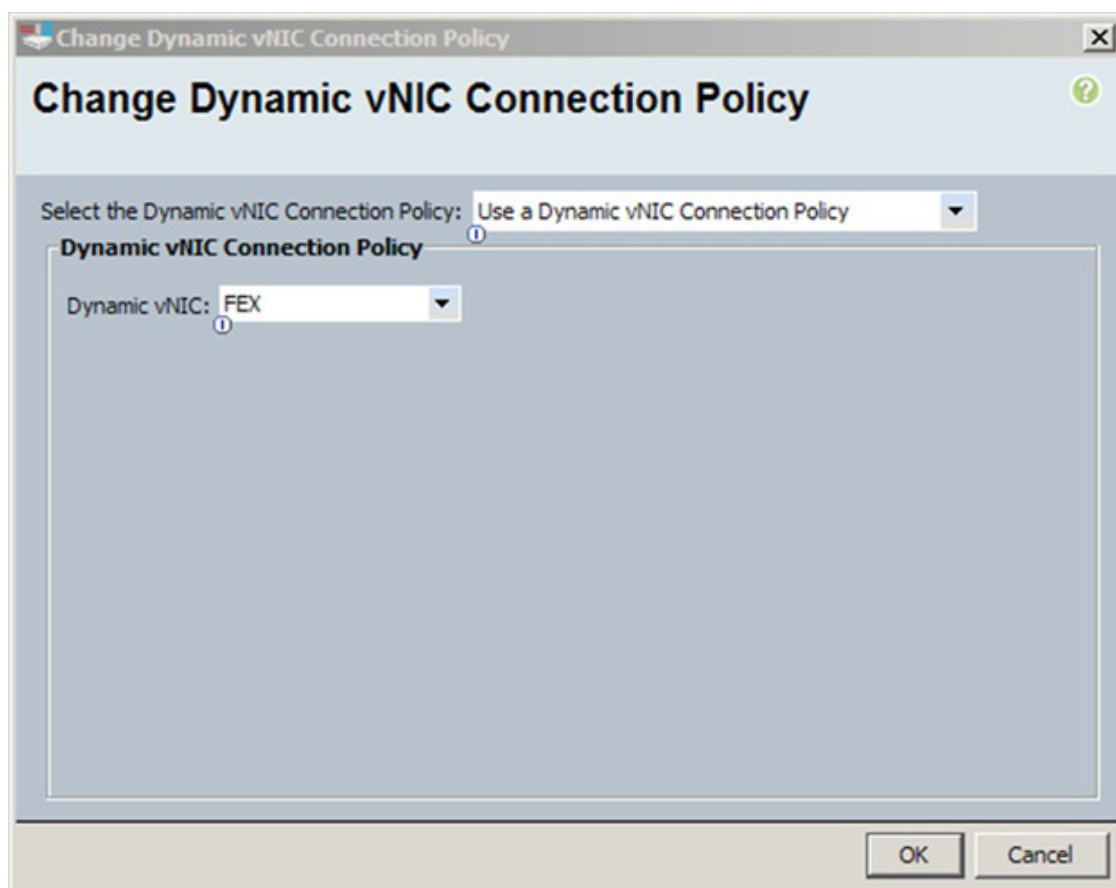
Create a VM-FEX Enabled Service Profile Template

To create a Cisco UCS service profile using VM-FEX, clone a previously defined Cisco UCS service profile and apply the dynamic vNIC and BIOS policies by completing the following steps in the Cisco UCS Manager:

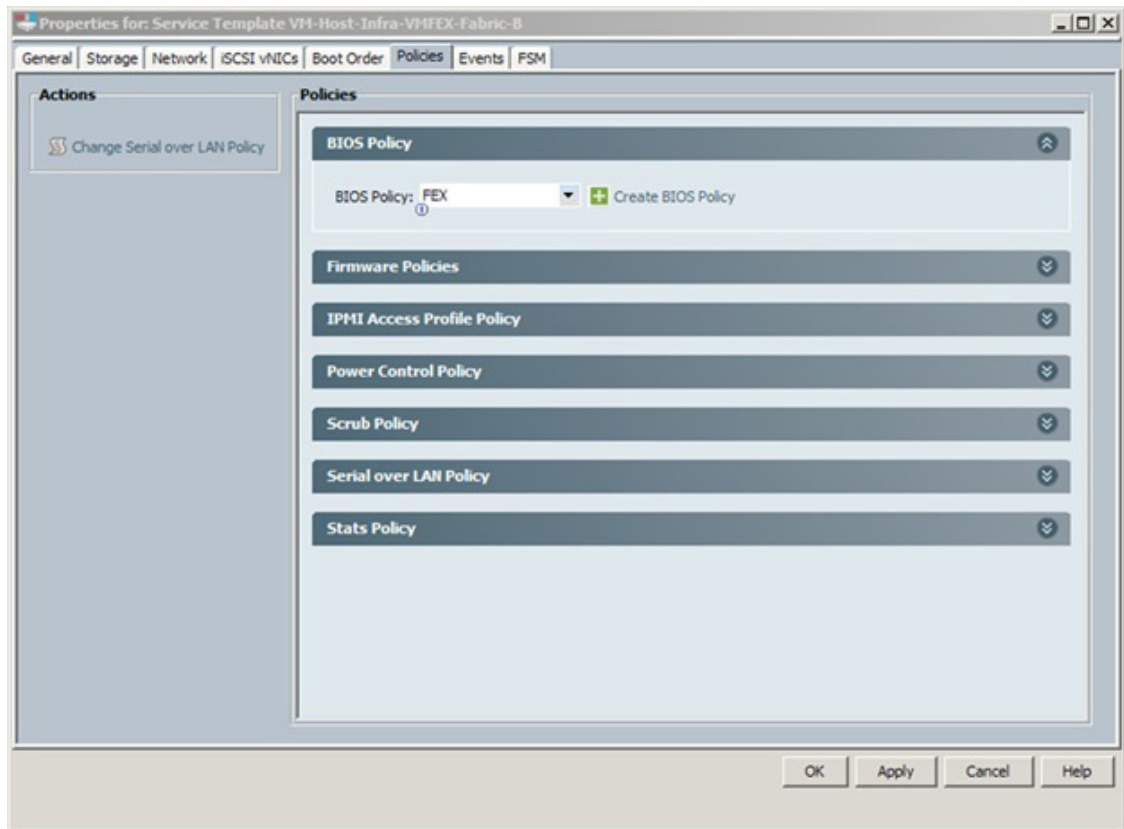
1. Select the Server tab in the left navigation pane and expand the Service Profile Templates.
2. Right-click VM-Host-Infra-Fabric-A and select Create a Clone.
3. Type a clone name and select an organizational owner for the new service profile template.



4. Click OK when notified that the service profile clone was successfully created. The Service Template navigation window appears.
5. Select the Network tab and click the Change Dynamic vNIC Connection Policy under the Actions section of the working pane. The Change Dynamic vNIC Connection Policy form appears.
6. Select Use a Dynamic vNIC Connection Policy from the drop-down menu and the previously created Dynamic vNIC policy previously defined. Click OK.



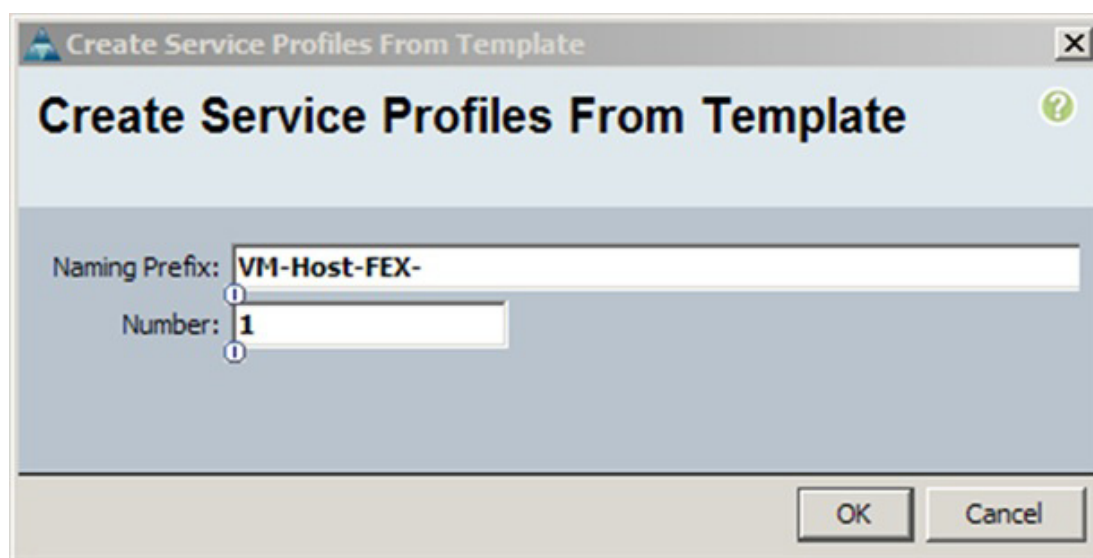
7. Click OK when notified that the vNIC connection policy was successfully modified.
8. From the Service Template properties window, select the Policies tab.
9. Expand the BIOS Policies in the Policies section of the working pane.
10. Select the previously defined FEX BIOS policy and click OK.



Create VM-FEX Service Profile

To create service profiles from the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > Service Template VM-Host-Infra-VMFEX-Fabric-A.
3. Right-click VM-Host-Infra-FEX-Fabric-A and select Create Service Profiles from Template.
4. Enter VM-Host-FEX-0 as the service profile prefix.
5. Enter 1 as the number of service profiles to create.
6. Click OK to create the service profile.



7. Click OK in the confirmation message.
8. Verify that the service profile VM-Host-FEX-1 has been created. The service profile is automatically associated with the servers in their assigned server pools.

Install and Set Up VMware ESXi

Refer to section “FlexPod VMware ESXi 5.1 iSCSI on Clustered Data ONTAP” to install and completely set up VMware ESXi version 5.1 on the two ESXi hosts. After ESXi setup is complete, add the two new hosts to VMware vCenter.

Download Cisco VEM Software Bundle

To download the Cisco UCS B-Series or C-Series server drivers, complete the following steps:

The following bundle was used during validation `cisco-vem-v151-5.1-1.1.1.1.vib`.

1. Open a Web browser on the management workstation and navigate to the following Cisco Download Software pages:
 - a. Downloads Home > Products > Servers - Unified Computing > Cisco UCS B-Series Blade Server Software > Unified Computing System (UCS) Drivers-2.1(1d)
 - b. Downloads Home > Products > Servers - Unified Computing > Cisco UCS C-Series Rack-Mount UCS-Managed Server Software > Unified Computing System (UCS) Drivers-1.4(5b)
2. Follow the steps necessary to download the software bundles located on the ISO image.
3. Mount the ISO image and copy the appropriate vib file from the VMware > VM-FEX > Cisco directory to the local machine.
4. From the vCenter vSphere Client, select the `infra_datastore_1` in the Inventory > Datastores and Datastore Clusters navigation menu.
5. Under the Basic Tasks choose Browse this Datastore
6. Select the root folder (`/`) and click the third button at the top to add a folder.

7. Name the folder VM-FEX and click OK.
8. On the left, select the VM-FEX folder.
9. Click the fourth button at the top and select Upload File.
10. Navigate to the cisco-vem-v151-5.1-1.1.1.1.vib file and click Open.
11. Click Yes to upload the .vib file to infra_datastore_1.

The VM-FEX file should now appear in the VM-FEX folder in the datastore.

Install the FEX Virtual Ethernet Module on Each ESXi Host

To install the Virtual Ethernet Module (VEM) on the ESXi hosts, complete the following steps:

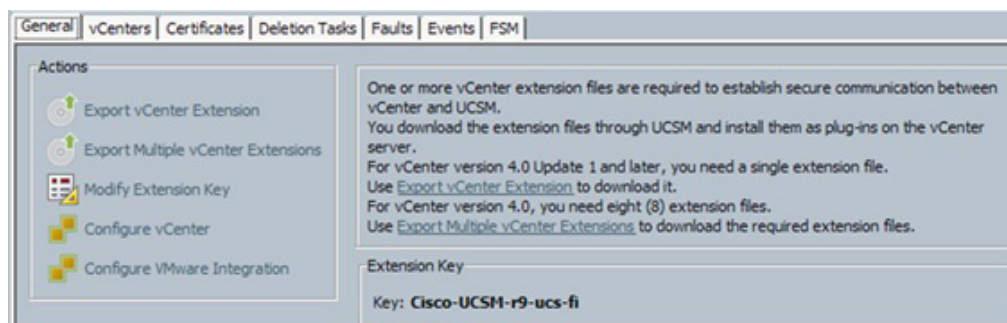
1. Open the VMware vSphere CLI command prompt.
2. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <host_ip> -u root -p <host_password> software vib install -v
/vmfs/volumes/infra_datastore_1/VM-FEX/cisco-vem-v151-5.1-1.1.1.1.vib
```

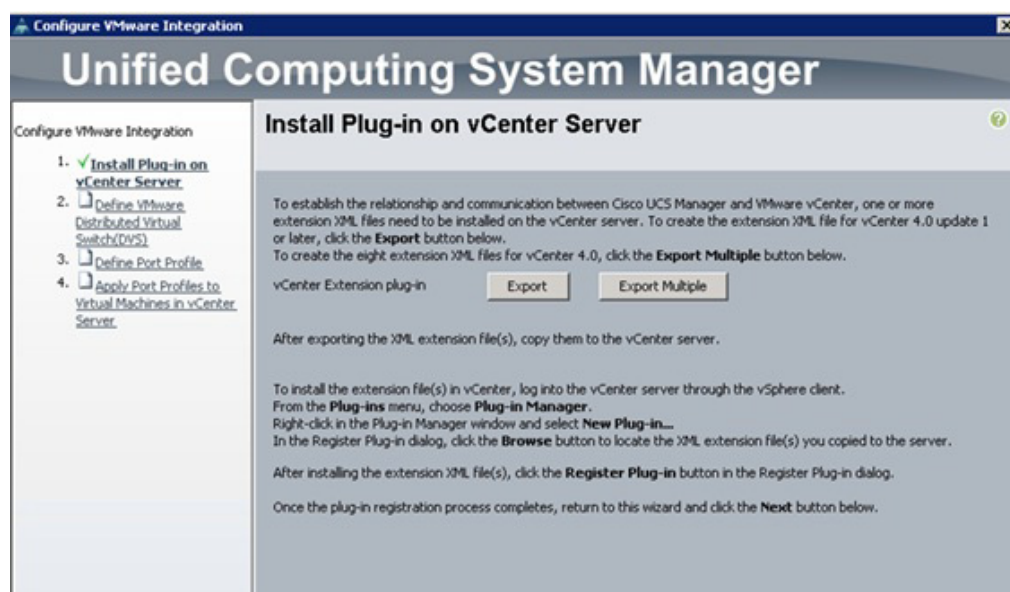
Integrate Cisco Unified Computing System with vCenter

To integrate Cisco UCS Manager and vCenter, complete the following steps:

1. Log in to the Cisco UCS Manager.
2. In the navigation pane, click the VM tab, and in the VM tab, expand the All folder. Select the VMware node, and in the Working Area, click the General tab.



3. Select the Configure VMware Integration menu in the Actions area to start the Configuration wizard.
4. Follow the instructions and click the Export button and complete the steps to install the UCSM extension file in vCenter.



5. Click Next.
6. Enter the VMware vCenter Server name, vCenter Server host name or IP address, vCenter data center name, DVS folder, and DVS name.
7. Click Next.

The screenshot shows the 'Define VMware Distributed Virtual Switch (DVS)' configuration window in the Unified Computing System Manager. The window is titled 'Configure VMware Integration' and 'Unified Computing System Manager'. On the left, a sidebar lists the steps: 1. Install Plug-in on vCenter Server (checked), 2. Define VMware Distributed Virtual Switch (DVS) (selected), 3. Define Port Profile, and 4. Apply Port Profiles to Virtual Machines in vCenter Server. The main area is divided into sections for vCenter Server, Datacenter, DVS Folder, and DVS. The vCenter Server section has fields for vCenter Server Name (set to <var vcenter_server_name>), Description, and vCenter Server Hostname or IP Address (set to <var vcenter_server_ip>). The Datacenter section has fields for vCenter Datacenter Name (set to FlexPod_DC_1) and Description. The DVS Folder section has fields for Folder Name (set to DVS-FEX) and Description. The DVS section has fields for DVS Name (set to DVS-FEX), Description, and a radio button for DVS (set to Enable). At the bottom, there are buttons for '< Prev', 'Next >', 'Finish', and 'Cancel'.

Configure VMware Integration

Unified Computing System Manager

Define VMware Distributed Virtual Switch(DVS)

Configure VMware Integration

1. ✓ Install Plug-in on vCenter Server
2. ✗ Define VMware Distributed Virtual Switch(DVS)
3. Define Port Profile
4. Apply Port Profiles to Virtual Machines in vCenter Server

vCenter Server

vCenter Server Name: <var vcenter_server_name>

Description:

vCenter Server Hostname or IP Address: <var vcenter_server_ip>

Datacenter

vCenter Datacenter Name: FlexPod_DC_1

Description:

DVS Folder

Folder Name: DVS-FEX

Description:

DVS

DVS Name: DVS-FEX

Description:

DVS ☐ Disable ☒ Enable

< Prev Next > Finish Cancel

8. Create the FEX-MGMT port profile, select the MGMT-VLAN, and indicate it is the native VLAN.

Configure VMware Integration

Unified Computing System Manager

Configure VMware Integration

1. ☒ Install Plug-in on vCenter Server
2. ☒ Define VMware Distributed Virtual Switch (DVS)
3. ☒ Define Port Profile
4. ☐ Apply Port Profiles to Virtual Machines in vCenter Server

Define Port Profile

Port Profile

Name:

QoS Policy:

Network Control Policy:

Max Ports:

Pin Group:

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	FooBar1_public	<input type="radio"/>
<input checked="" type="checkbox"/>	MGMT-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	Packet-Control-VLAN	<input type="radio"/>
<input type="checkbox"/>	Service-MA	<input type="radio"/>
<input type="checkbox"/>	ServiceNodeServices	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	yMotion-VLAN	<input type="radio"/>

Profile Client

Name:

Description:

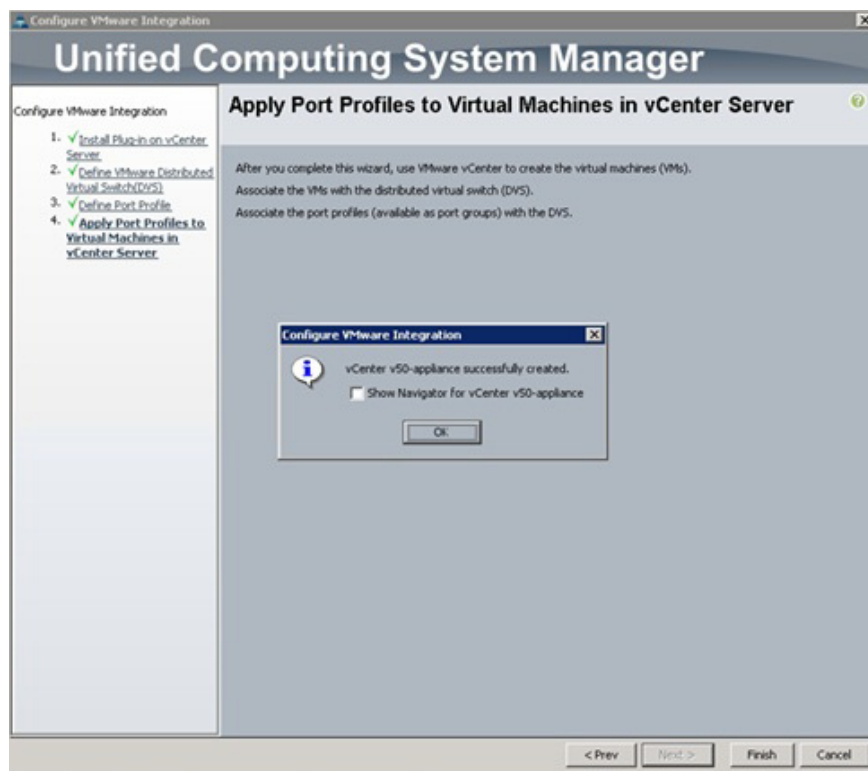
Datacenter:

Folder:

Distributed Virtual Switch:

< Prev Next > Finish Cancel

9. Click Next.
10. When finishing the wizard, the Cisco UCS Manager connects to vCenter and adds the plug-in.

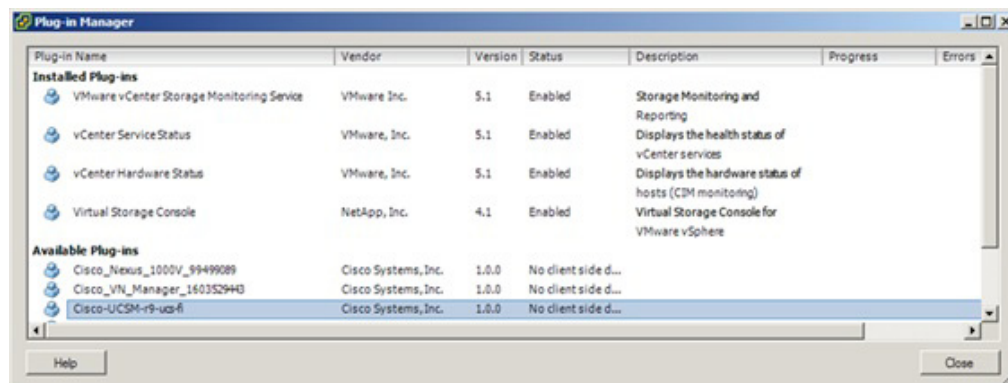
**Note**

The ESXi host will require additional hypervisor vNICs to support VMware vMotion, and NFS traffic uses the generic port-profile creation steps documented in section "Standard Operations" to establish a FEX-vMotion and FEX-NFS Port Profile.

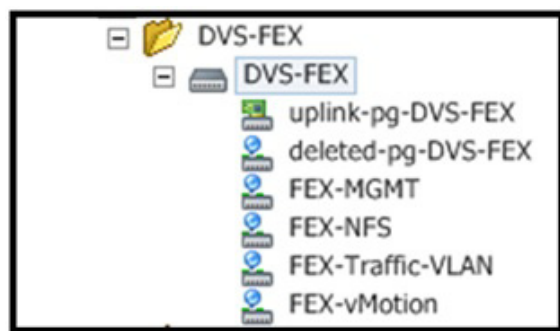
Validate Setting in VMware vCenter

To validate the successful installation of the Cisco UCS Manager plug-in, complete the following steps:

1. Log in to the vCenter Server.
2. In the Main menu, select Plug-ins > Manage Plug-ins.
3. The popup will show that Cisco UCS Manager is already integrated in vCenter.



- Click Inventory > Networking to see FEX added to distributed switch from Cisco UCS.



Standard Operations

The VM-FEX environment supports the addition of port profiles to the distributed switch. The following section describes how to add these distributed port groups.

Add Distributed Port Group to the VDS (vSphere Distributed Switch)

Port Profiles

Port profiles contain the properties and settings that you can use to configure virtual interfaces in Cisco UCS for VM-FEX. The port profiles are created and administered in Cisco UCS Manager. After a port profile is created, assigned to, and actively used by one or more distributed virtual switches (DVSs), any changes made to the networking properties of the port profile in Cisco UCS Manager are immediately applied to those DVSs.

In VMware vCenter, a port profile is represented as a port group. Cisco UCS Manager pushes the port profile names to VMware vCenter, which displays the names as port groups. None of the specific networking properties or settings in the port profile is visible in VMware vCenter. You must configure at least one port profile client for a port profile if you want Cisco UCS Manager to push the port profile to VMware vCenter.

Port Profile Client

The port profile client determines the DVSs to which a port profile is applied. By default, the port profile client specifies that the associated port profile applies to all DVSs in VMware vCenter. However, you can configure the client to apply the port profile to all DVSs in a specific data center or data center folder or to only one DVS.

Create a VM-FEX Port Profile

Complete the following steps to create VM-FEX port profiles for use on the Cisco UCS distributed virtual switch.

- Log in to Cisco UCS Manager.
- Click the VM tab.
- Right-click Port Profile > Create Port Profile.
- Enter the name of the Port Profile.
- Optional: Enter a description.
- Optional: Select a QoS policy.

7. Optional: Select a network control policy.
8. Enter the maximum number of ports that can be associated with this port profile. The default is 64 ports.

**Note**

The maximum number of ports that can be associated with a single DVS is 4096. If the DVS has only one associated port profile, that port profile can be configured with up to 4096 ports. However, if the DVS has more than one associated port profile, the total number of ports associated with all of those port profiles combined cannot exceed 4096.

9. Optional: Select High Performance.

**Note**

Select None-Traffic to and from a virtual machine passes through the DVS.

**Note**

Select High Performance-Traffic to and from a virtual machine bypasses the DVS and hypervisor and travels directly between the virtual machines and a virtual interface card (VIC) adapter.

10. Select the VLAN.
11. Select Native-VLAN.
12. Click OK.

Create Port Profile

Name:

Description:

QoS Policy:

Network Control Policy:

Max Ports:

Host Network IO Performance: ☒ None ☐ High Performance

Pin Group:

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	FooBar1_public	<input type="radio"/>
<input type="checkbox"/>	MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	Packet-Control-VLAN	<input type="radio"/>
<input type="checkbox"/>	Service-HA	<input type="radio"/>
<input type="checkbox"/>	ServiceNodeServices	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

OK Cancel

OR

Create Port Profile

Name:

Description:

QoS Policy:

Network Control Policy:

Max Ports:

Host Network IO Performance: ☐ None ☒ High Performance

Pin Group:

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	FooBar1_public	<input type="radio"/>
<input type="checkbox"/>	MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	Packet-Control-VLAN	<input type="radio"/>
<input type="checkbox"/>	Service-HA	<input type="radio"/>
<input type="checkbox"/>	ServiceNodeServices	<input type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

OK Cancel

The port profile created will appear in the working pane.

Create the Port Profile Client

To create the client profile for use in the Cisco UCS virtual distributed switch, complete the following steps:

1. In the navigation pane under the VM tab, expand All > Port Profiles. Right-click the Port Profile and click Create Profile Client.
2. Choose the data center created in your vCenter Server, folder, and distributed virtual switch created in section "Integrate Cisco UCS with vCenter."
3. Click OK.

Create Profile Client

Name:

Description:

Datacenter: All

Folder: All

Distributed Virtual Switch: All

OK Cancel

OR

Create Profile Client

Name: FEX-Traffic-VLAN

Description:

Datacenter: r9-dc-1

Folder: DVS-FEX

Distributed Virtual Switch: DVS-FEX

OK Cancel

The client profile created will appear in your distributed virtual switch DVS-FEX in vCenter as a port group.

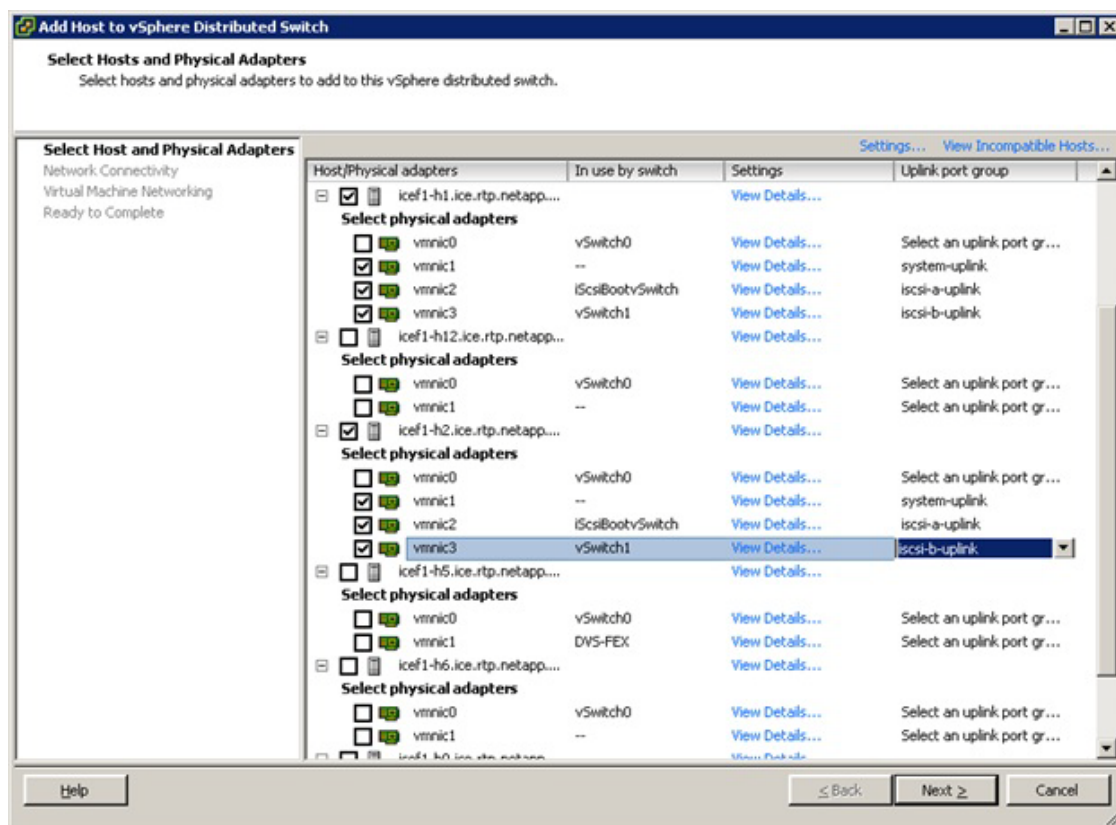
Repeat these steps as necessary for the workloads in the environment.

Migrate Networking Components for ESXi Hosts to Cisco DVS-FEX

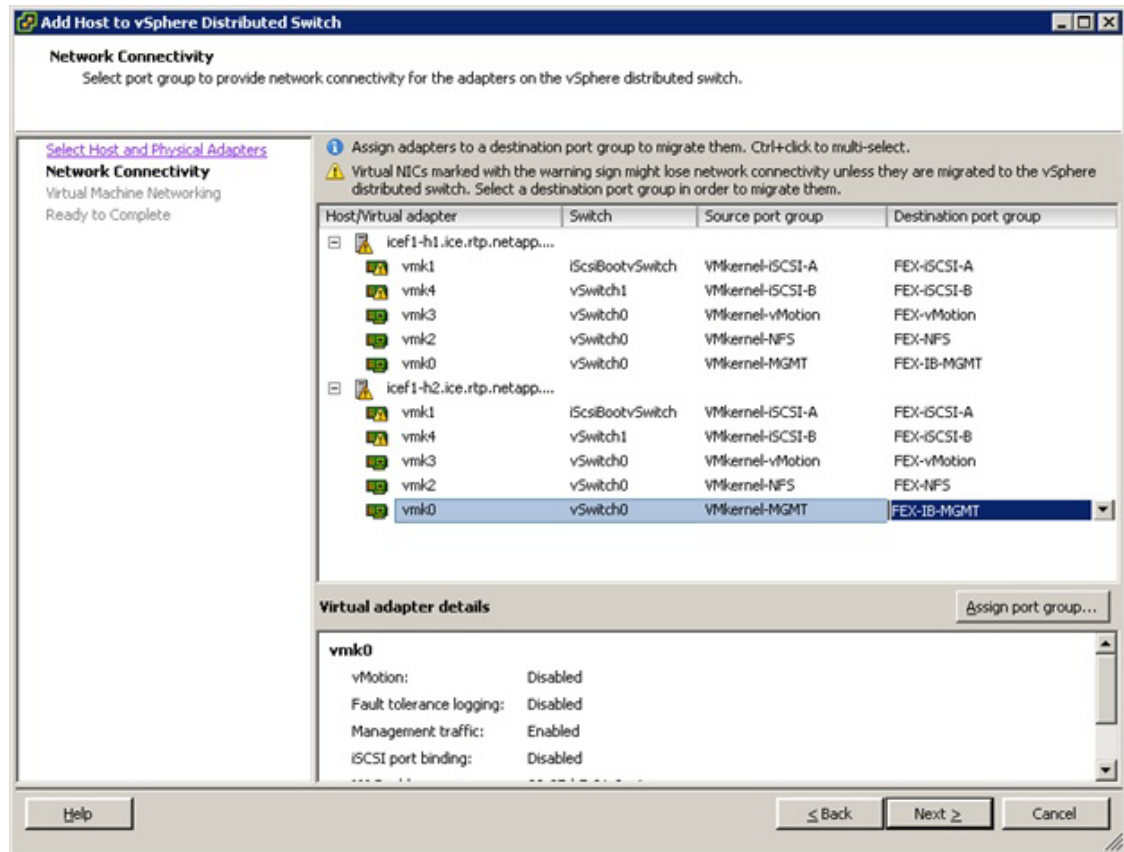
vCenter Server VM

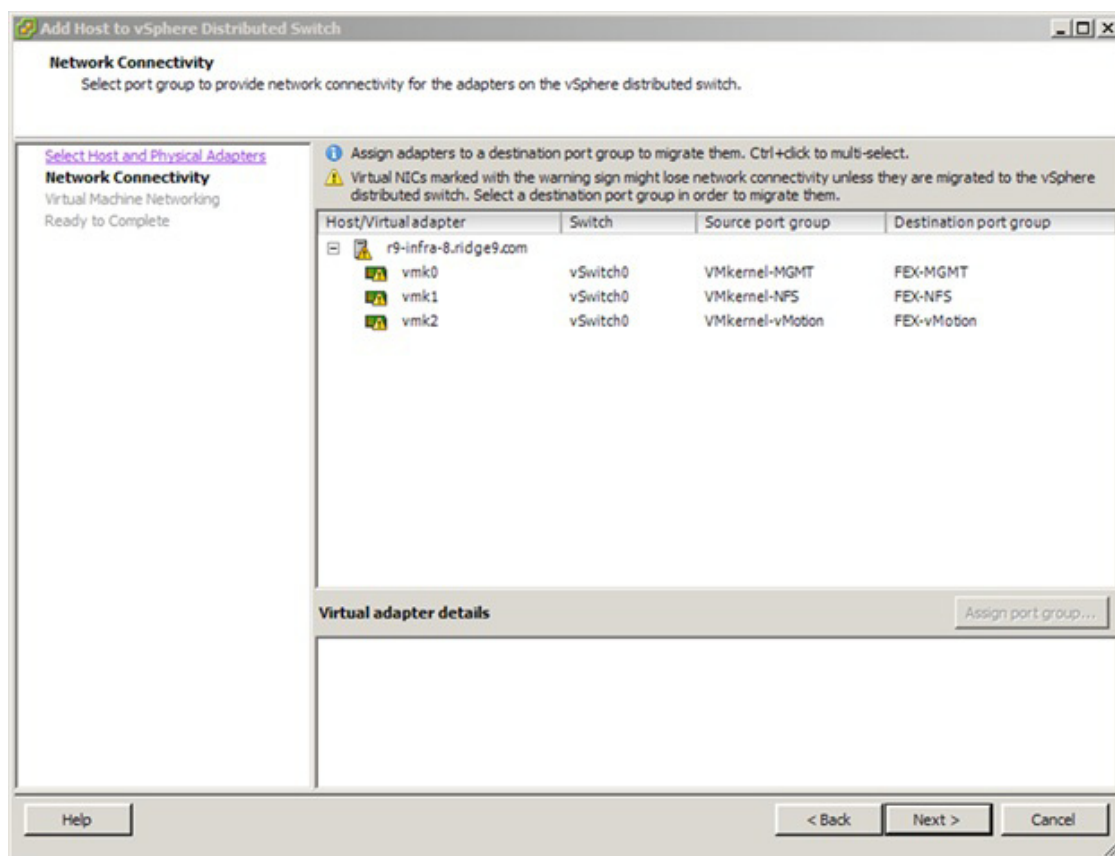
To migrate the networking components for the ESXi hosts to the Cisco FEX-DVS, complete the following steps:

1. In the VMware vSphere client connected to vCenter, select Home > Networking.
2. Expand the vCenter, DataCenter, and DVS-FEX folders. Select the DVS-FEX switch.
3. Under Basic Tasks for the vSphere distributed switch, select Add a Host.
4. For both hosts, select vmnic1, vmnic2, and vmnic 3. Select the uplink-pg-DVS-FEX Uplink port group. Click Next.

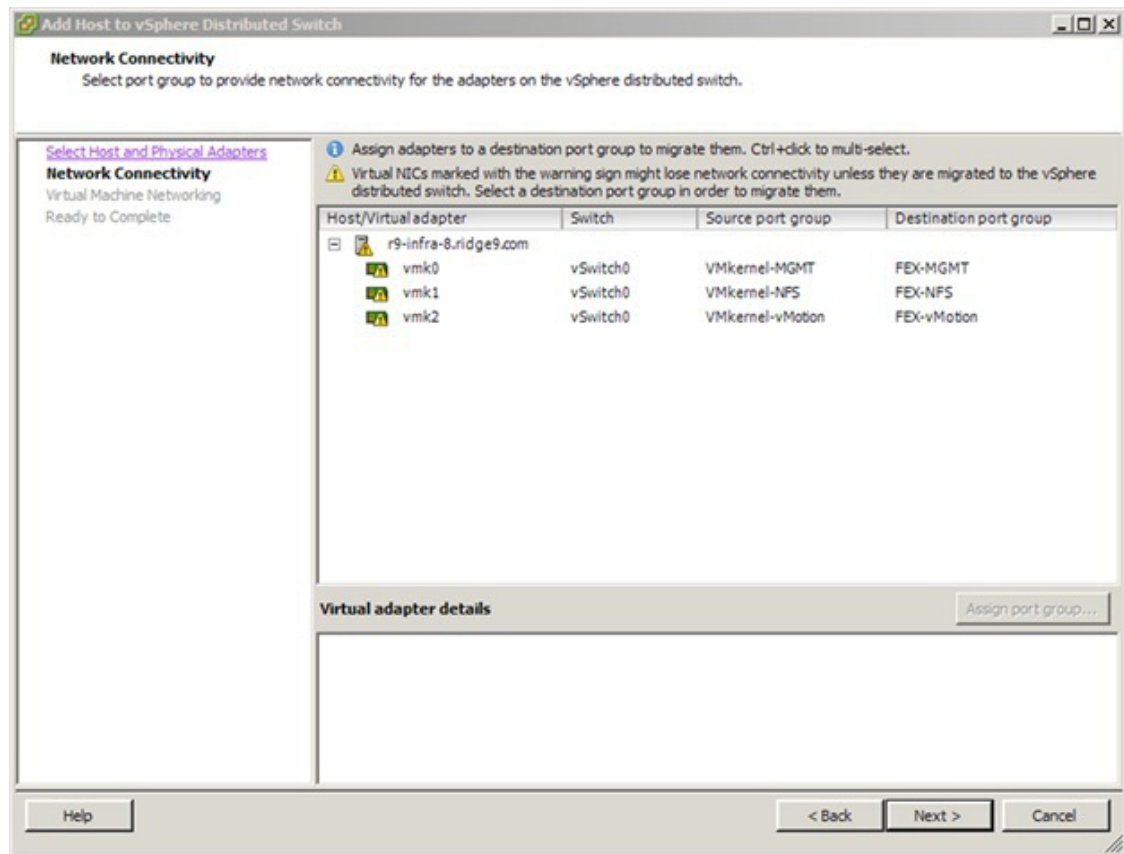


5. For all VMkernel ports, select the appropriate destination Port Group from the Cisco DVS-FEX. Click Next.





6. Select Migrate Virtual Machine Networking. 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 window, select Home > Hosts and Clusters.
9. Select the first ESXi host and select the Configuration tab. In the Hardware field, select Networking.
10. Make sure that vSphere Standard Switch is selected at the top next to View. None of the vSwitches should have any active VMkernel or VM Network ports on them. On the upper right of each vSwitch, click Remove.
11. Click Yes.
12. Remove all three vSwitches.
13. After all the vSwitches have disappeared from the screen, click vSphere Distributed Switch at the top next to View.
14. Click Manage Physical Adapters.
15. In the uplink-pg-DVS-FEX field, click <Click to Add NIC>.
16. Select vmnic0 and click OK.
17. Click OK to close the Manage Physical Adapters window. Four uplinks should now be present.
18. Select the second ESXi host and click the Configuration tab. In the Hardware field, select Networking.
19. Make sure vSphere Standard Switch is selected at the top next to View. None of the vSwitches should have any active VMkernel or VM Network ports on them. On the upper right of each vSwitch, click Remove.

20. Click Yes.
21. Remove all three vSwitches.
22. After all of the vSwitches have disappeared from the screen, click vSphere Distributed Switch at the top next to View.
23. Click Manage Physical Adapters.
24. In the uplink-pg-ADVS-FEX field, click <Click to Add NIC>.
25. Select vmnic0 and click OK.
26. Click OK to close the Manage Physical Adapters window. Four uplinks should now be present.

VM-FEX Virtual Interfaces

In a blade server environment, the number of vNICs and vHBAs configurable for a service profile is determined by adapter capability and the amount of virtual interface (VIF) namespace available in the adapter. In Cisco UCS, portions of VIF namespace are allotted in chunks called VIFs. Depending on your hardware, the maximum number of VIFs is allocated on a predefined, per-port basis.

The maximum number of VIFs varies based on hardware capability and port connectivity. For each configured vNIC or vHBA, one or two VIFs are allocated. Standalone vNICs and vHBAs use one VIF, and failover vNICs and vHBAs use two.

The following variables affect the number of VIFs available to a blade server, and therefore, the number of vNICs and vHBAs you can configure for a service profile.

- The maximum number of VIFs supported on your fabric interconnect
- How the fabric interconnects are cabled
- If the fabric interconnect and IOM are configured in fabric port channel mode

For more information about the maximum number of VIFs supported by your hardware configuration, refer to the Cisco UCS 6100 and 6200 Series Configuration Limits for Cisco UCS Manager for your software release. Tables 22 and 23 reference these limits.

Table 22 VM-FEX environment configuration limits

Feature	Cisco UCS 6200 Series
Hosts per DVS	52
DVSs per Cisco UCS Domain	1
vCenter Server units per Cisco UCS Domain	4
Port profiles per Cisco UCS Domain	512
Dynamic ports per port profile	4096
Dynamic ports per DVS	4096

Table 23 Cisco UCS fabric interconnect and C-Series server VIF support

Acknowledged Link Between FEX and Fabric Interconnect	Maximum VIFs (vNICs + vHBAs) Per VIC Adapter in Single-Wire Management	Maximum VIFs (vNICs + vHBAs) Per VIC Adapter in Dual-Wire Management
1	12	13
2	27	28
4	57	58
8	117	118

**Note**

For a non-VIC adapter the maximum number of vNICs is two and the maximum number of vHBAs is two.

**Note**

If the server in single-wire mode has two VIC adapters, the maximum number of VIFs (vNICs + vHBAs) available for the second adapter would be same as for an adapter in a dual-wire mode server.

**Note**

For more information on C-Series integration into UCSM, go to http://www.cisco.com/en/US/docs/unified_computing/ucs/c-series_integration/ucsm2.1/b_UCSM2-1_C-Integration.pdf.

Cisco Nexus 7000 Example Configurations

Cisco Nexus 7000 A

```
!Command: show running-config
!Time: Mon Mar 7 22:21:35 2013

version 6.1(2)
hostname 7K1-VPC1

cfs eth distribute
feature udld
feature interface-vlan
feature lacp
feature vpc

username admin password 5 $1$polJITl9$rnQzCMXRfgPqBRiWQRDZZ1 role vdc-admin
no password strength-check
ip domain-lookup
service unsupported-transceiver
snmp-server user admin vdc-admin auth md5 0xe98f5e9df8db7f3c7721915210dde612 pri
v 0xe98f5e9df8db7f3c7721915210dde612 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vrf context management
 ip route 0.0.0.0/0 172.26.164.1
vlan 1,2,911-912,3170-3171,3173-3174,3176
vlan 2
 name Native-VLAN
vlan 911
 name iSCSI-A-VLAN
vlan 912
 name iSCSI-B-VLAN
vlan 3170
 name NFS-VLAN
```

```

vlan 3175
    name IB-MGMT-VLAN
vlan 3173
    name vMotion-VLAN
vlan 3174
    name VM-Traffic-VLAN
vlan 3176
    name Packet-Control-VLAN

vpc domain 1
    role priority 10
    peer-keepalive destination 172.26.164.79 source 172.26.164.78
    auto-recovery

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default

interface Vlan1

interface port-channel10
    description vPC peer-link
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 911-912,3170,3173-3174,3176
    spanning-tree port type network
    mtu 9216
    vpc peer-link

interface port-channel11
    description FAS3250-A
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 911-912,3170
    spanning-tree port type edge trunk
    mtu 9216
    vpc 11

interface port-channel12
    description FAS3250-B
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 911-912,3170
    spanning-tree port type edge trunk
    mtu 9216
    vpc 12

interface port-channel13
    description UCSM-B
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 911-912,3170,3173-3175
    spanning-tree port type edge trunk

```

```
mtu 9216
vpc 13

interface port-channel14
  description UCSM-B
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3175
  spanning-tree port type edge trunk
  mtu 9216
  vpc 14

interface Ethernet3/1
  description FAS3250-A:e3a
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170
  mtu 9216
  channel-group 11 mode active
  no shutdown

interface Ethernet3/2
  description FAS3250-B:e3a
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170
  mtu 9216
  channel-group 12 mode active
  no shutdown

interface Ethernet3/3

interface Ethernet3/4

interface Ethernet3/5

interface Ethernet3/6

interface Ethernet3/7

interface Ethernet3/8

interface Ethernet3/9

interface Ethernet3/10

interface Ethernet3/11
  description VPC Peer Nexus 7000-B:3/11
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3176
  mtu 9216
  channel-group 10 mode active
  no shutdown
```



```

interface Ethernet3/12
  description VPC Peer Nexus 7000-B:3/11
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3176
  mtu 9216
  channel-group 10 mode active
  no shutdown

interface Ethernet3/13

interface Ethernet3/14

interface Ethernet3/15

interface Ethernet3/16

interface Ethernet3/17

interface Ethernet3/18

interface Ethernet3/19

interface Ethernet3/20

interface Ethernet3/21

interface Ethernet3/22

interface Ethernet3/23
  description UCSM-A:1/19
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3175
  mtu 9216
  channel-group 13 mode active
  no shutdown

interface Ethernet3/24
  description UCSM-B:1/19
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3175
  mtu 9216
  channel-group 14 mode active
  no shutdown

cli alias name wr copy r s
cli alias name shpcs sh port-channel sum
line vty
  exec-timeout 0

```

Cisco Nexus 7000 B

```

!Command: show running-config
!Time: Tue Feb 12 01:26:09 2013

version 6.1(2)
hostname 7K2-VPC1

cfs eth distribute
feature udld
feature interface-vlan
feature lacp
feature vpc

username admin password 5 $1$/jld0R/i$ABSkUbKKwbFRMBL30.udZ0 role vdc-admin
no password strength-check
ip domain-lookup
service unsupported-transceiver
snmp-server user admin vdc-admin auth md5 0xf7bd86250e9767f264324f674500409c
priv 0xf7bd86
250e9767f264324f674500409c localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vrf context management
  ip route 0.0.0.0/0 172.26.164.1
vlan 1,2,911-912,3170-3171,3173-3174,3176
vlan 2
  name Native-VLAN
vlan 911
  name iSCSI-A-VLAN
vlan 912
  name iSCSI-B-VLAN
vlan 3170
  name NFS-VLAN
vlan 3175
  name IB-MGMT-VLAN
vlan 3173
  name vMotion-VLAN
vlan 3174
  name VM-Traffic-VLAN
vlan 3176
  name Packet-Control-VLAN

vpc domain 1
  role priority 20
  peer-keepalive destination 172.26.164.78 source 172.26.164.79
  auto-recovery

interface Vlan1

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

```

```

switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170,3173-3176
spanning-tree port type network
mtu 9216
vpc peer-link

interface port-channel11
description FAS3250-A
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170
spanning-tree port type edge trunk
mtu 9216
vpc 11

interface port-channel12
description FAS3250-B
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170
spanning-tree port type edge trunk
mtu 9216
vpc 12

interface port-channel13
description UCSM-B
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170,3173-3175
spanning-tree port type edge trunk
mtu 9216
vpc 13

interface port-channel14
description UCSM-B
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170,3173-3175
spanning-tree port type edge trunk
mtu 9216
vpc 14

interface Ethernet3/1
description FAS3250-A:e4a
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170
mtu 9216
channel-group 11 mode active
no shutdown

interface Ethernet3/2
description FAS3250-B:e4a

```

```
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170
mtu 9216
channel-group 12 mode active
no shutdown

interface Ethernet3/3

interface Ethernet3/4

interface Ethernet3/5

interface Ethernet3/6

interface Ethernet3/7

interface Ethernet3/8

interface Ethernet3/9

interface Ethernet3/10

interface Ethernet3/11
description VPC Peer Nexus 7000-A:3/11
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170,3173-3176
mtu 9216
channel-group 10 mode active
no shutdown

interface Ethernet3/12
description VPC Peer Nexus 7000-A:3/11
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 911-912,3170,3173-3176
mtu 9216
channel-group 10 mode active
no shutdown

interface Ethernet3/13

interface Ethernet3/14

interface Ethernet3/15

interface Ethernet3/16

interface Ethernet3/17

interface Ethernet3/18

interface Ethernet3/19
```

```

interface Ethernet3/20

interface Ethernet3/21

interface Ethernet3/22

interface Ethernet3/23
  description UCSM-A:1/20
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3175
  mtu 9216
  channel-group 13 mode active
  no shutdown

interface Ethernet3/24
  description UCSM-B:1/20
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 911-912,3170,3173-3175
  mtu 9216
  channel-group 14 mode active
  no shutdown

cli alias name wr copy r s
cli alias name shpcs sh port-channel sum
line vty
  exec-timeout 0

```

Alternate NetApp FAS 7-Mode Storage Configuration

Assign Controller Disk Ownership and Initialize Storage

The following steps provide details for assigning disk ownership and disk initialization and verification.

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.

Table 24 **Detail and Detail Value**

Detail	Detail Value
Controller 1 mgmt IP	<<var_controller1_e0m_ip>>
Controller 1 netmask	<<var_controller1_mask>>
Controller 1 gateway	<<var_controller1_mgmt_gateway>>
URL of the Data ONTAP boot software	<<var_url_boot_software>>
Controller 2 mgmt IP	<<var_controller2_e0m_ip>>

Controller 2 netmask	<<var_controller2_mask>>
Controller 2 gateway	<<var_controller2_mgmt_gateway>>

Controller 1

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, Press Ctrl-C to exit the Autoboot loop when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

2. If the system is at the LOADER prompt, enter the following command to boot Data ONTAP:

autoboot

3. During system boot, press Ctrl-C when prompted for the Boot Menu:
4. Press Ctrl-C for Boot Menu...

**Note**

If 8.1.2 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.2 is the version being booted, then proceed with step 14, maintenance mode boot.

5. To install new software first select option 7.

7

6. Answer yes for performing a nondisruptive upgrade.

y

7. Select e0M for the network port you want to use for the download.

e0M

8. Select yes to reboot now.

9. y

10. Enter the IP address, netmask, and default gateway for e0M in their respective places.

<<var_controller1_e0m_ip>>

<<var_controller1_mask>>

<<var_controller1_mgmt_gateway>>

11. Enter the URL where the software can be found.

**Note**

This Web server must be pingable.

<<var_url_boot_software>>

12. Press Enter for the username, indicating no user name.

Enter

13. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

y

14. Enter yes to reboot the node.

y

15. When you see "Press Ctrl-C for Boot Menu":

Ctrl-C

16. To enter Maintenance mode boot, select option 5.

5

17. When you see the question "Continue to Boot?" type yes.

y

18. To verify the HA status of your environment, enter:

```
ha-config show
```

**Note**

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

19. To see how many disks are unowned, enter:

```
disk show -a
```

**Note**

No disks should be owned in this list.

20. Assign disks.

```
disk assign -n <<var_#_of_disks>>
```

**Note**

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

21. Reboot the controller.

```
halt
```

22. At the LOADER-A prompt, enter:

```
autoboot
```

23. Press Ctrl-C for Boot Menu when prompted.

```
Ctrl-C
```

24. Select option 4 for Clean configuration and initialize all disks.

```
4
```

25. Answer yes to zero disks, reset config and install a new file system.

```
Y
```

26. Enter yes to erase all the data on the disks.

```
Y
```

**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. You can continue to controller 2 configuration while the disks for controller 1 are zeroing.

Controller 2

1. Connect to the storage system console port. You should see a Loader-A prompt. However if the storage system is in a reboot loop, Press Ctrl-C to exit the Autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. If the system is at the LOADER prompt, enter the following command to boot Data ONTAP:

```
autoboot
```

3. During system boot, press Ctrl-C when prompted for the Boot Menu:

```
Press Ctrl-C for Boot Menu...
```

**Note**

If 8.1.2 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.2 is the version being booted, then proceed with step 14, maintenance mode boot

4. To install new software first select option 7.

7

5. Enter yes for performing a nondisruptive upgrade.

Y

6. Select e0M for the network port you want to use for the download.

e0M

7. Enter yes to reboot now.

Y

8. Enter the IP address, netmask and default gateway for e0M in their respective places.

<<var_controller2_e0m_ip>>

<<var_controller2_mask>>

<<var_controller2_mgmt_gateway>>

9. Enter the URL where the software can be found.

**Note**

This Web server must be pingable.

<<var_url_boot_software>>

10. Press Enter for the username, indicating no user name.

Enter

11. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

Y

12. Enter yes to reboot the node.

Y

13. When you see "Press Ctrl-C for Boot Menu":

Ctrl-C

14. To enter Maintenance mode boot, select option 5:

5

15. If you see the question "Continue to Boot?" type yes.

Y

16. To verify the HA status of your environment, enter:

ha-config show

**Note**

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

17. To see how many disks are unowned, enter:

disk show -a

**Note**

The remaining disks should be shown.

18. Assign disks by entering:

disk assign -n <<var_#_of_disks>>

**Note**

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

19. Reboot the controller.

halt

20. At the LOADER prompt, enter:

autoboot

21. Press Ctrl-C for Boot Menu when prompted.

Ctrl-C

22. Select option 4 for a Clean configuration and initialize all disks.

4

23. Answer yes to zero disks, reset config and install a new file system.

y

24. Enter yes to erase all the data on the disks.

y



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.

Run Setup Process

When Data ONTAP is installed on a new storage system, the following files are not populated:

- /etc/rc
- /etc/exports
- /etc/hosts
- /etc/hosts.equiv

Controller 1

- 1.** Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
- 2.** Enter the following information:

Please enter the new hostname []:<<var_controller1>>

Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter

Please enter the IP address for Network Interface e0a []: Enter



Note

Press Enter to accept the blank IP address.

Should interface e0a take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e0b []:Enter

Should interface e0b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e1a []:Enter

Should interface e1a take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e1b []:Enter

Should interface e1b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for Network Interface e0M []:

<<var_controller1_e0m_ip>>

Please enter the netmaskfor the Network Interface e0M [255.255.255.0]:

<<var_controller1_mask>>

```
Should interface e0M take over a partner IP address during failover? [n]: y
Please enter the IPv4 address or interface name to be taken over by e0M []: e0M
Please enter flow control for e0M {none, receive, send, full} [full]: Enter
```

**Note**

If additional interface cards are installed in your storage controller, you will have additional questions about the interfaces on those cards.

3. Enter the following information:

```
Please enter the name or IP address of the IPv4 default gateway:
<<var_controller1_mgmt_gateway>>
```

The administration host is given root access to the storage system's / etc files for system administration. To allow /etc root access to all NFS clients enter RETURN below.

```
Please enter the name or IP address for administrative host:
<<var_adminhost_ip>>
```

```
Please enter timezone [GTM]: <<var_timezone>>
```

**Note**

Example time zone: America/New_York.

```
Where is the filer located? <<var_location>>
Enter the root directory for HTTP files [home/http]: Enter
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:
```

**Note**

Optionally enter up to three name server IP addresses.

```
Do you want to run NIS client? [n]: Enter
Press the Return key to continue through AutoSupport message
Would you like to configure SP LAN interface [y]: Enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_sp_ip>>
Please enter the netmask for the SP []: <<var_sp_mask>>
Please enter the IP address for the SP gateway: <<var_sp_gateway>>
Please enter the name or IP address of the mail host [mailhost]:
<<var_mailhost>>
Please enter the IP address for <<var_mailhost>> []: <<var_mailhost_ip>>
New password: <<var_password>>
Retype new password <<var_password>>
```

4. Enter the root password to log in to controller 1.

Controller 2

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.

2. Enter the following information:

```
Please enter the new hostname []: <<var_controller2>>
```

Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter

Please enter the IP address for Network Interface e0a []: Enter



Note

Press Enter to accept the blank IP address.

Should interface e0a take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e0b []: Enter

Should interface e0b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e1a []: Enter

Should interface e1a take over a partner IP address during failover? [n]: Enter

Please enter the IP address for the Network Interface e1b []: Enter

Should interface e1b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for Network Interface e0M []:

<<var_controller2_e0m_ip>>

Please enter the netmask for the Network Interface e0M [255.255.255.0]:

<<var_controller2_mask>>

Should interface e0M take over a partner IP address during failover? [n]: y

Please enter the IPv4 address or interface name to be taken over by e0M []: e0M

Please enter flow control for e0M {none, receive, send, full} [full]: Enter



Note

If additional interface cards are installed in your storage controller, you will have additional questions about the interfaces on those cards.

3. Enter the following information:

Please enter the name or IP address of the IPv4 default gateway:

<<var_controller2_mgmt_gateway>>

The administration host is given root access to the storage system's / etc files for system administration. To allow /etc root access to all NFS clients enter RETURN below.

Please enter the name or IP address for administrative host:

<<var_adminhost_ip>>

Please enter timezone [GTM]: <<var_timezone>>



Note

Example time zone: America/New_York.

Where is the filer located? <<var_location>>

Enter the root directory for HTTP files [home/http]: Enter

Do you want to run DNS resolver? [n]: y

Please enter DNS domain name []: <<var_dns_domain_name>>

Please enter the IP address for first nameserver []: <<var_nameserver_ip>>

Do you want another nameserver? [n]:



Note

Optionally enter up to three name server IP addresses.

Do you want to run NIS client? [n]: Enter

Press the Return key to continue through AutoSupport message

Would you like to configure SP LAN interface [y]: Enter

```

Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_sp_ip>>
Please enter the netmask for the SP []: <<var_sp_mask>>
Please enter the IP address for the SP gateway: <<var_sp_gateway>>
Please enter the name or IP address of the mail host [mailhost]:
<<var_mailhost>>
Please enter the IP address for <<var_mailhost>> []: <<var_mailhost_ip>>
New password: <<var_password>>
Retype new password <<var_password>>
4. Enter the root password to log in to controller 2.

```

Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.1.2, you must upgrade to the latest Service Processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Using a web browser, connect to <http://support.netapp.com/NOW/cgi-bin/fw>.
2. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
3. Proceed to the Download page for the latest release of the SP Firmware for your storage platform.
4. Using the instructions on this page, update the SPs on both controllers. You will need to download the .zip file to a web server that is reachable from the management interfaces of the controllers.

64-Bit Aggregates in Data ONTAP 7-Mode

A 64-bit aggregate containing the root volume is created during the Data ONTAP setup process. To create additional 64-bit aggregates, determine the aggregate name, the node on which to create it, and how many disks it will contain. Calculate the RAID group size to allow for roughly balanced (same size) RAID groups of between 12 and 20 disks (for SAS disks) within the aggregate. For example, if 52 disks were being assigned to the aggregate, select a RAID group size of 18. A RAID group size of 18 would yield two 18-disk RAID groups and one 16-disk RAID group. Keep in mind that the default RAID group size is 16 disks, and that the larger the RAID group size, the longer the disk rebuild time in case of a failure.

Controller 1

1. Execute the following command to create a new aggregate:

```
aggr create aggr1 -B 64 -r <<var_raidsize>> <<var_#_of_disks>>
```

2. Leave at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Controller 2

1. Execute the following command to create a new aggregate:

```
aggr create aggr1 -B 64 -r <<var_raidsize>> <<var_#_of_disks>>
```

2. Leave at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Flash Cache

Controller 1 and Controller 2

1. Execute the following commands to enable Flash Cache:

```
options flexscale.enable on
options flexscale.lopri_blocks off
options flexscale.normal_data_blocks on
```

**Note**

For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, refer to TR-3832: Flash Cache and PAM Best Practices Guide. Before customizing the settings, determine whether the custom settings are required or whether the default settings are sufficient.

IFGRP LACP

Since this type of interface group requires two or more Ethernet interfaces and a switch that supports LACP, make sure that the switch is configured properly.

Controller 1 and Controller 2

Run the following command on the command line and also add it to the /etc/rc file, so it is activated upon boot:

```
ifgrp create lacp ifgrp0 -b port e1a e1b
wrfile -a /etc/rc "ifgrp create lacp ifgrp0 -b ip e1a e1b"
```

All interfaces must be in down status before being added to an interface group.

VLAN

Controller 1 and Controller 2

Run the following commands to create VLAN interfaces for NFS and iSCSI data traffic.

```
vlan create ifgrp0 <<var_nfs_vlan_id>>
wrfile -a /etc/rc "vlan create ifgrp0 <<var_nfs_vlan_id>>"
vlan create ifgrp0 <<var_iscsi_vlan_A_id>>
wrfile -a /etc/rc "vlan create ifgrp0 <<var_iscsi_vlan_A_id>>"
vlan create ifgrp0 <<var_iscsi_vlan_B_id>>
wrfile -a /etc/rc "vlan create ifgrp0 <<var_iscsi_vlan_B_id>>"
```

IP Config

Controller 1 and Controller 2

Run the following commands on the command line.

```
ifconfig ifgrp0-<<var_nfs_vlan_id>> <<var_nfs_ip>> netmask <<var_nfs_mask>>
mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_nfs_vlan_id>> <<var_nfs_ip>> netmask
<<var_nfs_mask>> mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>"

ifconfig ifgrp0-<<var_iscsi_vlan_A_id>> <<var_iscsi_A_ip>> netmask
<<var_iscsi_A_mask>> mtusize 1500 partner ifgrp0-<<var_iscsi_vlan_A_id>>
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_iscsi_vlan_A_id>> <<var_iscsi_A_ip>>
netmask <<var_iscsi_A_mask>> mtusize 1500 partner
ifgrp0-<<var_iscsi_vlan_A_id>>"

ifconfig ifgrp0-<<var_iscsi_vlan_B_id>> <<var_iscsi_B_ip>> netmask
<<var_iscsi_B_mask>> mtusize 1500 partner ifgrp0-<<var_iscsi_vlan_B_id>>
```

```
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_iscsi_vlan_B_id>> <<var_iscsi_B_ip>>
netmask <<var_iscsi_B_mask>> mtusize 1500 partner
ifgrp0-<<var_iscsi_vlan_B_id>>"
```

Cisco Discovery Protocol

Use the following steps to enable Cisco Discovery Protocol (CDP) on controller 1 and controller 2.

Controller 1 and Controller 2

1. Enable CDP

```
options cdpd.enable on
```

Active-Active Controller Configuration

Controller 1 and Controller 2

To enable two storage controllers to an active-active configuration, complete the following steps.

1. Enter the cluster license on both nodes.

```
license add <<var_cf_license>>
```

2. Reboot each storage controller.

```
reboot
```

3. Log back in to both controllers.

Controller 1

Enable failover on Controller 1, if it is not enabled already.

```
cf enable
```

NFSv3

Controller 1 and Controller 2

1. Add a license for NFS.

```
license add <<var_nfs_license>>
```

2. Set the following recommended options that enable NFS version 3.

```
options nfs.tcp.enable on
```

```
options nfs.udp.enable off
```

```
options nfs.v3.enable on
```

3. Enable NFS.

```
nfs on
```

iSCSI

Controller 1 and Controller 2

1. License iSCSI.

```
license add <<var_iscsi_license>>
```

2. Start the iSCSI service.

```
iscsi start
```

3. Record the iSCSI node name for later use.

```
iscsi nodename
```

NTP

The following commands configure and enable time synchronization on the storage controller. You must have either a publically available IP address or your company's standard NTP server name or IP address.

Controller 1 and Controller 2

1. Run the following commands to configure and enable the NTP server:

```
date <<var_date>>
```

2. Enter the current date in the format of [[[CC]yy]mm]dd]hhmm[.ss]].

3. For example: date 201208311436 would set the date to August 31, 2012, at 14:36.

```
options timed.servers <<var_global_ntp_server_ip>>
```

```
options timed.enable on
```

Data ONTAP SecureAdmin

Secure API access to the storage controller must be configured.

Controller 1

1. Issue the following as a one-time command to generate the certificates used by the web services for the API.

```
secureadmin setup ssl
```

```
SSL Setup has already been done before. Do you want to proceed? [no] y
```

```
Country Name (2 letter code) [US]: <<var_country_code>>
```

```
State or Province Name (full name) [California]: <<var_state>>
```

```
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
```

```
Organization Name (company) [Your Company]: <<var_org>>
```

```
Organization Unit Name (division): <<var_unit>>
```

```
Common Name (fully qualified domain name) [<<var_controller1_fqdn>>]: Enter
```

```
Administrator email: <<var_admin_email>>
```

```
Days until expires [5475] : Enter
```

```
Key length (bits) [512] : <<var_key_length>>
```



Note

NetApp recommends that your key length be 1024.

After the initialization, the CSR is available in the file /etc/keymgr/csr/secureadmin_tmp.pem.

2. Configure and enable SSL and HTTPS for API access using the following options.

```
options httpd.access none
```

```
options httpd.admin.enable off
```

```
options httpd.admin.ssl.enable on
```

```
options ssl.enable on
```

Controller 2

1. Issue the following as a one-time command to generate the certificates used by the web services for the API.

```
secureadmin setup ssl
```

```
SSL Setup has already been done before. Do you want to proceed? [no] y
```

```
Country Name (2 letter code) [US]: <<var_country_code>>
```

```

State or Province Name (full name) [California]: <<var_state>>
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
Organization Name (company) [Your Company]: <<var_org>>
Organization Unit Name (division): <<var_unit>>
Common Name (fully qualified domain name) [<<var_controller2_fqdn>>]: Enter
Administrator email: <<var_admin_email>>
Days until expires [5475] : Enter
Key length (bits) [512] : <<var_key_length>>

```

**Note**

NetApp recommends that your key length be 1024.

After the initialization, the CSR is available in the file `/etc/keymgr/csr/secureadmin_tmp.pem`.

2. Configure and enable SSL and HTTPS for API access using the following options.

```

options httpd.access none
options httpd.admin.enable off
options httpd.admin.ssl.enable on
options ssl.enable on

```

Secure Shell

SSH must be configured and enabled.

Controller 1 and Controller 2

1. Use the following one-time command to generate host keys.

```

secureadmin disable ssh
secureadmin setup -f -q ssh 768 512 1024

```

2. Use the following options to configure and enable SSH.

```

options ssh.idle.timeout 60
options autologout.telnet.timeout 5

```

SNMP

Controller 1 and Controller 2

1. Run the following commands to configure SNMP basics, such as the local and contact information. When polled, this information displays as the `sysLocation` and `sysContact` variables in SNMP.

```

snmp contact "<<var_admin_email>>"
snmp location "<<var_location>>"
snmp init 1
options snmp.enable on

```

2. Configure SNMP traps to send them to remote hosts, such as a DFM server or another fault management system.

```

snmp traphost add <<var_oncommand_server_fqdn>>

```

SNMPv1

Controller 1 and Controller 2

1. Set the shared secret plain-text password, which is called a community.

```

snmp community delete all

```



```
snmp community add ro <<var_snmp_community>>
```

**Note**

Use the delete all command with caution. If community strings are used for other monitoring products, the delete all command will remove them.

SNMPv3

SNMPv3 requires a user to be defined and configured for authentication.

Controller 1 and Controller 2

1. Create a user called snmpv3user.

```
useradmin role add snmp_requests -a login-snmp
useradmin group add snmp_managers -r snmp_requests
useradmin user add snmpv3user -g snmp_managers
New Password: <<var_password>>
Retype new password: <<var_password>>
```

AutoSupport HTTPS

AutoSupport sends support summary information to NetApp through HTTPS.

Controller 1 and Controller 2

1. Execute the following commands to configure AutoSupport:

```
options autosupport.noteto <<var_admin_email>>
```

Security Best Practices

Apply the following commands according to local security policies.

Controller 1 and Controller 2

1. Run the following commands to enhance security on the storage controller:

```
options rsh.access none
options webdav.enable off
options security.passwd.rules.maximum 14
options security.passwd.rules.minimum.symbol 1
options security.passwd.lockout.numtries 6
options autologout.console.timeout 5
```

Install Remaining Required Licenses and Enable MultiStore

Controller 1 and Controller 2

1. Install the following licenses to enable SnapRestore and FlexClone.

```
license add <<var_snaprestore_license>>
license add <<var_flex_clone_license>>
options licensed_feature.multistore.enable on
```

Enable NDMP

Run the following commands to enable NDMP.

Controller 1 and Controller 2

```
options ndmpd.enable on
```

Create FlexVol Volumes

Controller 1

1. Create two volumes on controller 1 by using the following steps:

```
vol create esxi_boot -s none aggr1 100g
snap reserve esxi_boot 0
sis on /vol/esxi_boot
vol create infra_swap -s none aggr1 100g
snap reserve infra_swap 0
snap sched infra_swap 0 0 0
```

Controller 2

1. Create two volumes on controller 2 using the following steps:

```
vol create infra_datastore_1 -s none aggr1 500g
snap reserve infra_datastore_1 0
sis on /vol/infra_datastore_1
vol create OnCommandDB -s none aggr1 200g
snap reserve OnCommandDB 0
sis on /vol/OnCommandDB
```

NFS Exports

Use the following steps to create NFS exports on each controller.

Controller 1

```
exportfs -p
sec=sys,rw=<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,root=<<
var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,nosuid
/vol/infra_swap
```

```
exportfs -p
sec=sys,ro,rw=<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra
02_nfs_host_ip>>,root==<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var
_vm_infra02_nfs_host_ip>>,nosuid /vol/vol0
```

Controller 2

```
exportfs -p
sec=sys,rw=<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,root=<<
var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,nosuid
/vol/infra_datastore_1
```

```
exportfs -p
sec=sys,ro,rw=<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra
02_nfs_host_ip>>,root==<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var
_vm_infra02_nfs_host_ip>>,nosuid /vol/vol0
```

LUN Creation

Use the following steps to create two LUNs on controller 1.

Controller 1

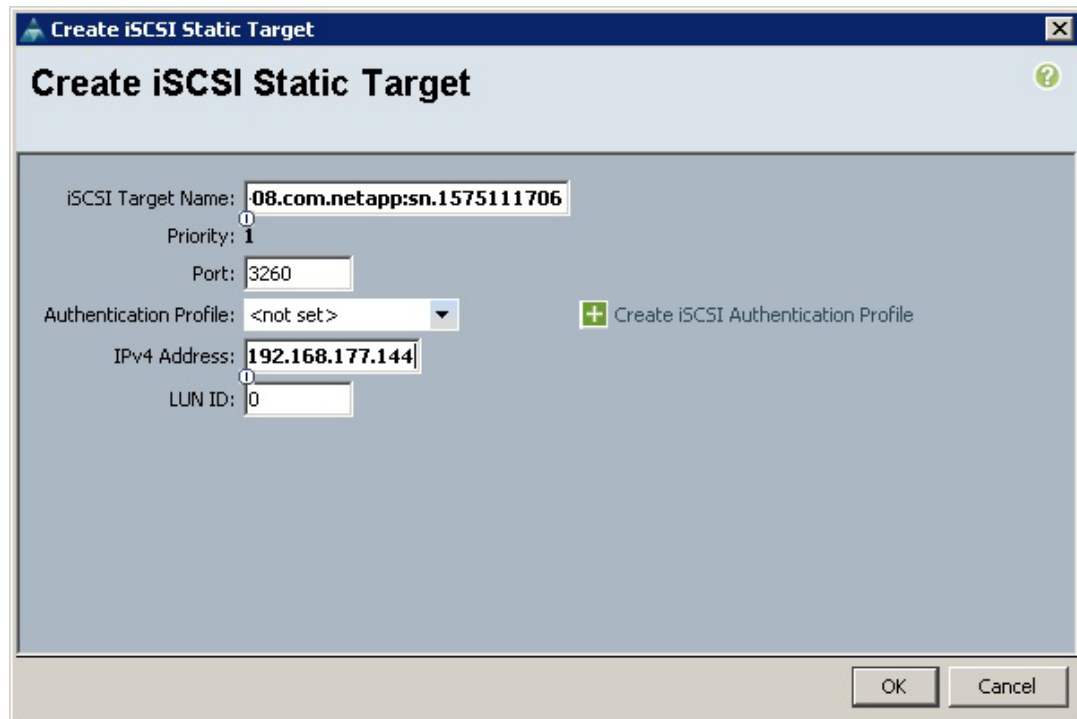
```
lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-01
lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-02
```

Alternate Cisco UCS Configuration with 7-Mode Storage

When using 7-Mode storage, the only deviation required in the Cisco UCS setup is the setup of the server boot order when creating the service profile template.

1. Set the server boot order.
 - a. Select Boot-Fabric-A for Boot Policy.
 - b. In the Boot Order pane, select iSCSI-A-vNIC.
 - c. Click the Set iSCSI Boot Parameters button.
 - d. In the Set iSCSI Boot Parameters dialog box, enter IQN_Pool_A in the Initiator Name Assignment field.
 - e. In the Set iSCSI Boot Parameters dialog box, enter iSCSI_IP_Pool_A in the Initiator IP field.
 - f. Keep the iSCSI Static Target Interface button selected and click the + button.
 - g. Log in to the controller 1's management interface and run the following command:


```
iscsi nodename
```
 - h. Note or copy the iSCSI target node name.
 - i. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from controller 1 into the iSCSI Target Name field.
 - j. Enter the IP address for controller 1's ifgrp0-`<<var_iscsi_vlan_A_id>>` in the IPv4 Address field.

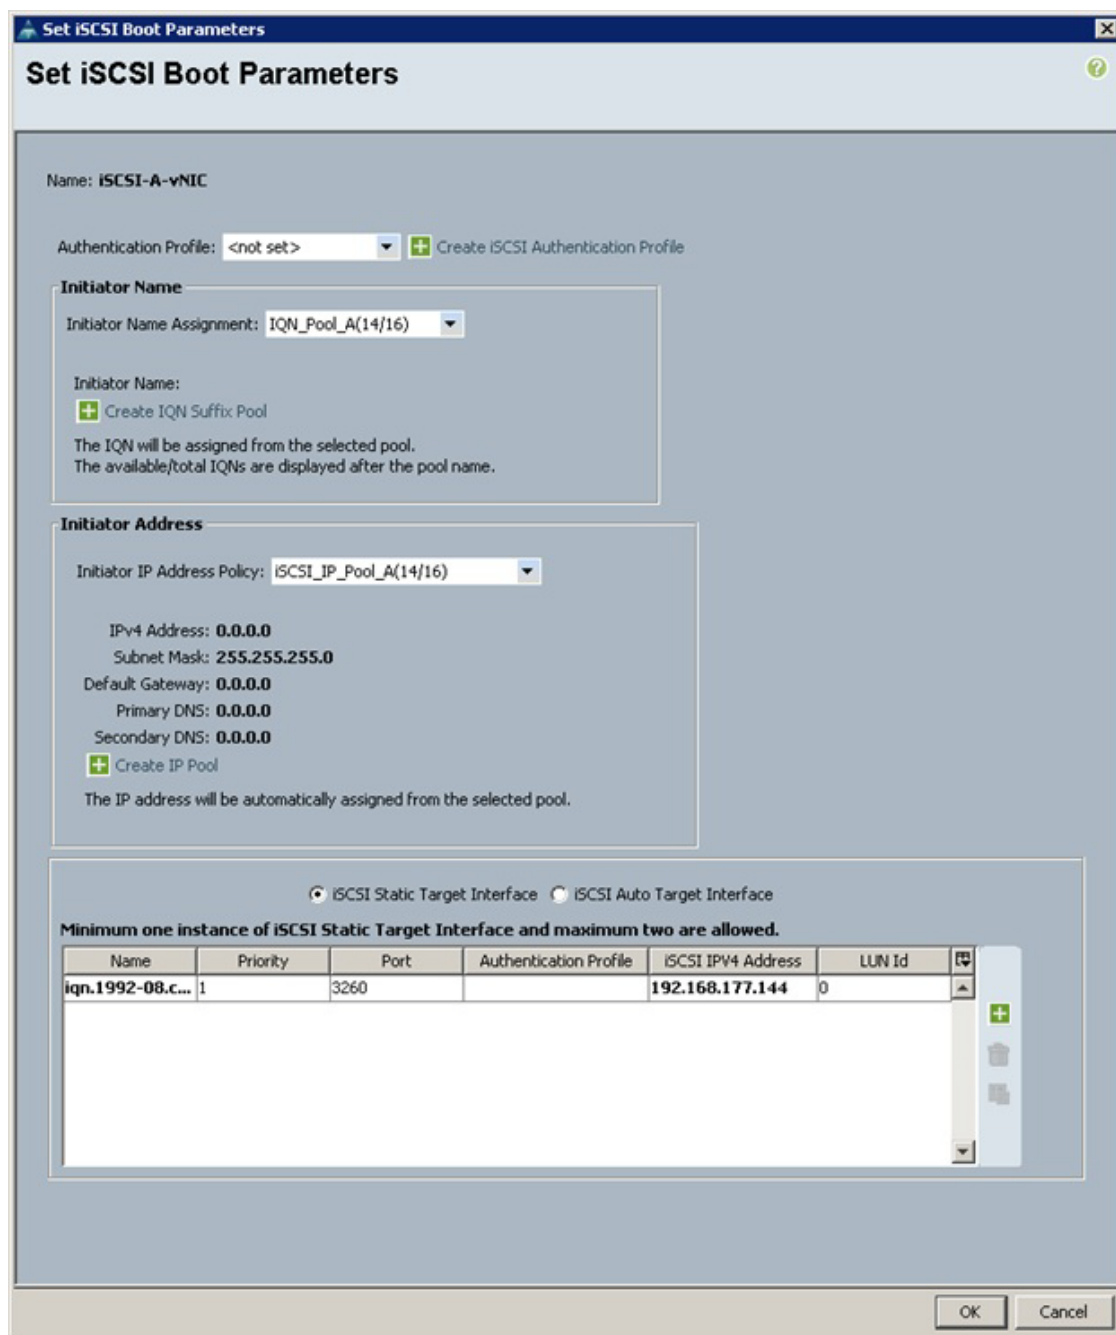


The image shows a Windows-style dialog box titled "Create iSCSI Static Target". The dialog has a blue header bar with the title and a close button (X). Below the header, the title "Create iSCSI Static Target" is repeated in a larger font. The main area of the dialog contains several input fields and a button:

- iSCSI Target Name:** A text box containing "08.com.netapp:sn.1575111706".
- Priority:** A text box containing "1".
- Port:** A text box containing "3260".
- Authentication Profile:** A dropdown menu showing "<not set>".
- IPv4 Address:** A text box containing "192.168.177.144".
- LUN ID:** A text box containing "0".
- Create iSCSI Authentication Profile:** A green button with a plus sign icon.

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

2. Click OK to add the iSCSI static target.



Set iSCSI Boot Parameters

Name: **iSCSI-A-vNIC**

Authentication Profile: **<not set>** + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: **IQN_Pool_A(14/16)**

Initiator Name:

+ Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: **iSCSI_IP_Pool_A(14/16)**

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

+ Create IP Pool

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

☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id	
iqn.1992-08.c...	1	3260		192.168.177.144	0	

+ Add Remove Refresh

OK Cancel

3. In the Boot Order pane, select iSCSI-vNIC-B.
4. Click the Set iSCSI Boot Parameters button.
5. In the Set iSCSI Boot Parameters dialog box, enter IQN_Pool_B in the Initiator Name Assignment field.
6. In the Set iSCSI Boot Parameters dialog box, enter iSCSI_IP_Pool_B in the Initiator IP Address Policy field.
7. Keep the iSCSI Static Target Interface button selected and click the + button.

8. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from controller 1 into the iSCSI Target Name field (same target name as above).
9. Enter the IP address for controller 1's ifgrp0-`<<var_iscsi_vlan_B_id>>` in the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name: 08.com.netapp:sn.1575111706

Priority: 1

Port: 3260

Authentication Profile: <not set> [+ Create iSCSI Authentication Profile](#)

IPv4 Address: 192.168.178.144

LUN ID: 0

OK Cancel

10. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters

Name: **iSCSI-B-vNIC**

Authentication Profile: **<not set>** + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: **IQN_Pool_B(14/16)**

Initiator Name:

+ Create iSCSI Suffix Pool

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

Initiator Address

Initiator IP Address Policy: **iSCSI_IP_Pool_B(14/16)**

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

+ Create IP Pool

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

☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPv4 Address	LUN Id	
iqn.1992-08.c...	1	3260		192.168.178.144	0	

+ Trash Refresh

11. Click OK.
12. 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.
13. Click Next to continue to the next section.
14. Return to step 11 (Add a maintenance policy) in section "Create Service Profile Templates."

Alternate 7-Mode NetApp FAS3250 Deployment Procedure: Part 2

Create iSCSI Igrouops

To create the Fibre Channel igroups on the storage controller for SAN boot of the Cisco UCS hosts, complete the following steps.

Controller A

1. Run the following command:

```
igroup create -i -t vmware VM-Host-Infra-01 <VM-Host-Infra-01 iSCSI-A-vNIC IQN>
<VM-Host-Infra-01 iSCSI-B-vNIC IQN>.
```

2. Run the following command:

```
igroup create -i -t vmware VM-Host-Infra-02 <VM-Host-Infra-02 iSCSI-A-vNIC IQN>
<VM-Host-Infra-02 iSCSI-B-vNIC IQN>.
```

Controller B

1. Run the following command:

```
igroup create -i -t vmware MGMT-Hosts <VM-Host-Infra-01 iSCSI-A-vNIC IQN>
<VM-Host-Infra-01 iSCSI-B-vNIC IQN> <VM-Host-Infra-02 iSCSI-A-vNIC IQN>
<VM-Host-Infra-02 iSCSI-B-vNIC IQN>.
```



Note

The ESXi Host iSCSI IQNs can be obtained by selecting the iSCSI vNICs in the Boot Order tab of the Service Profile, and clicking Set iSCSI Boot Parameters.

Map LUNs to Igrouops

To map the boot LUNs to the Fibre Channel igroups on the storage controller for SAN boot of the Cisco UCS hosts, complete the following steps:

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.
```

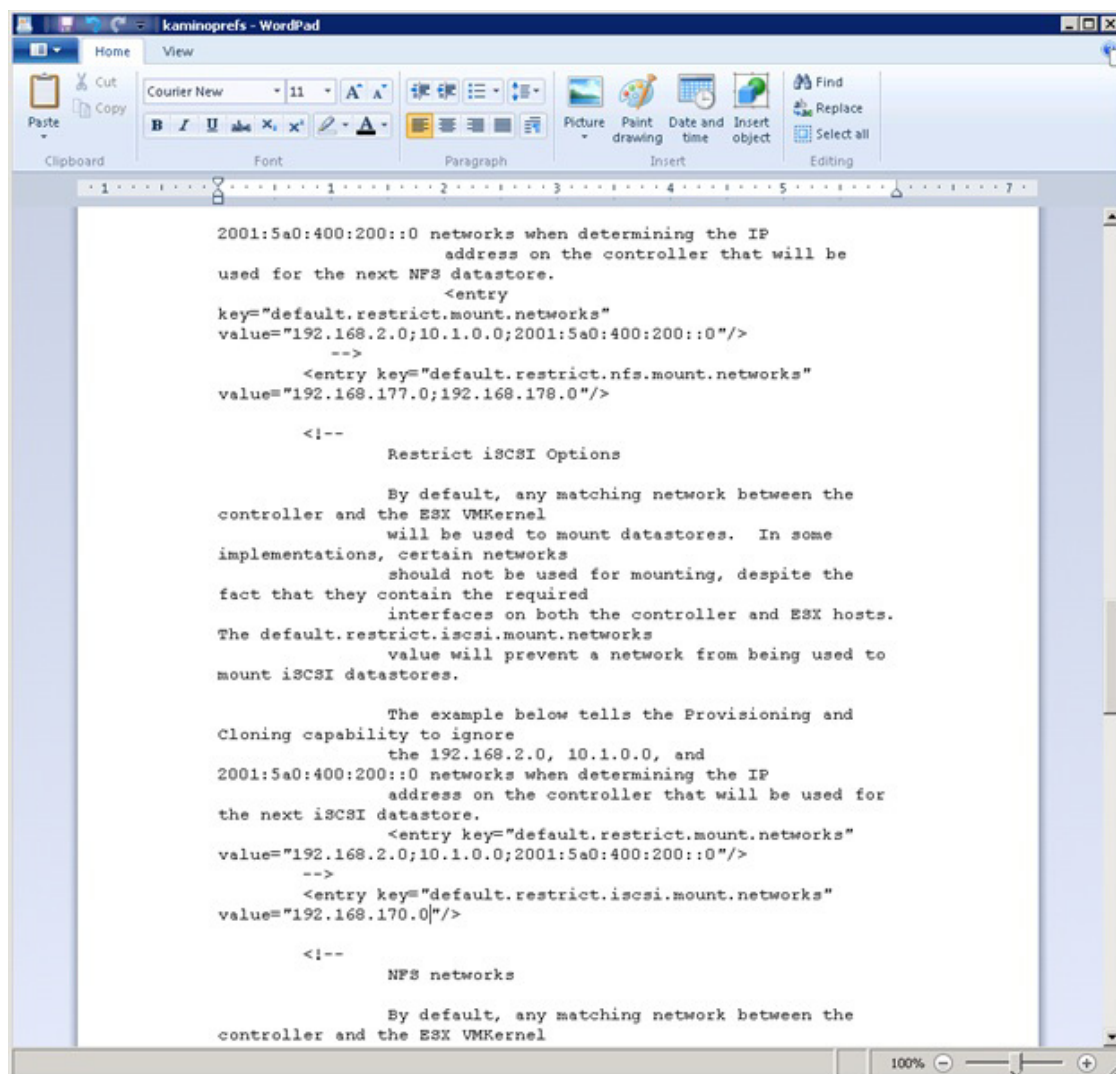
4. Verify that the created LUNs are mapped correctly.

Set Up 7-Mode iSCSI and NFS Networks in VSC

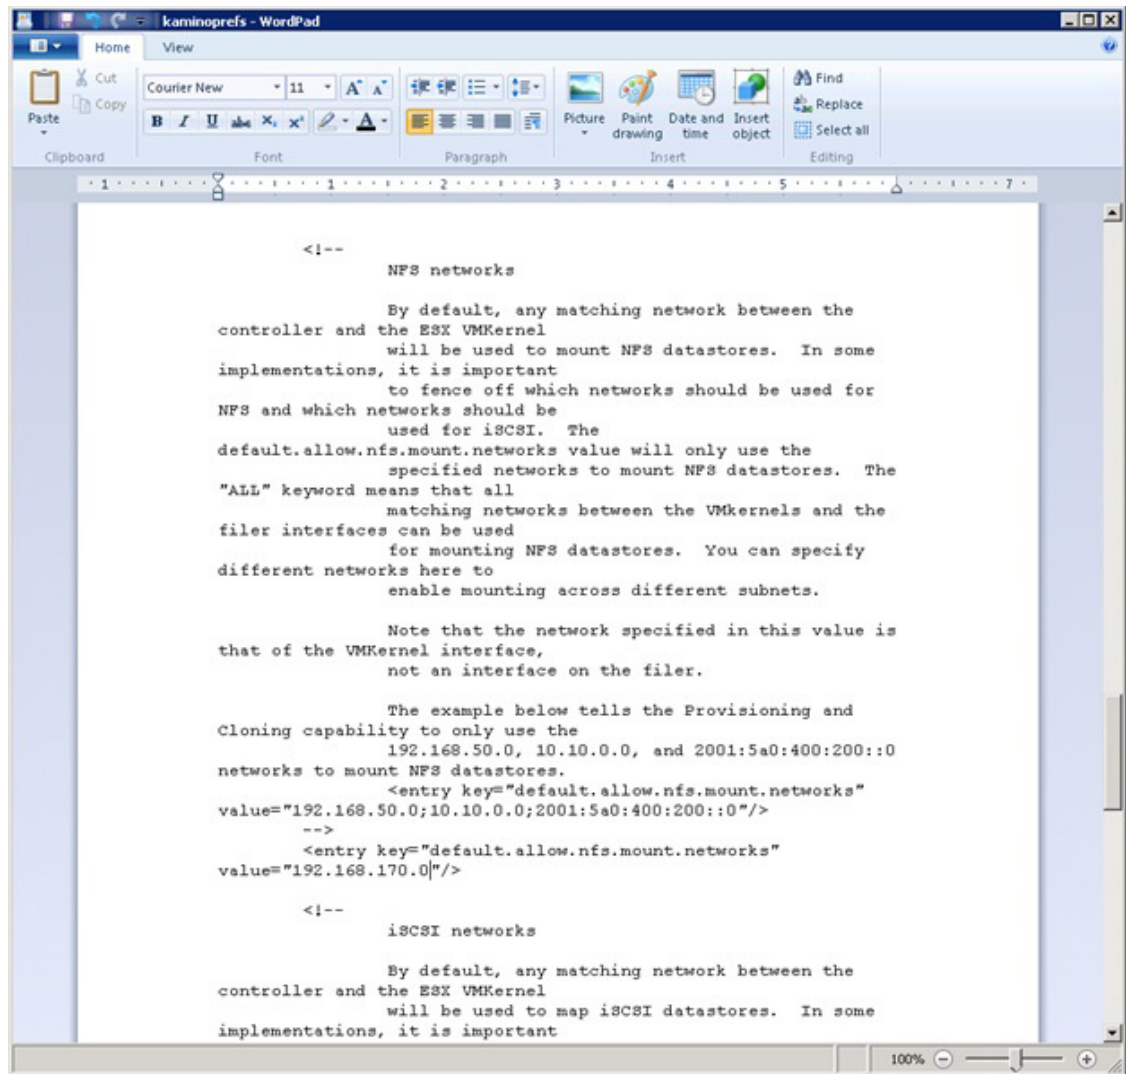
OnCommand and VSC VM Console

1. Open a console window on the VM containing VSC. Log into the VM as the FlexPod admin and open Windows Explorer.
2. Go to Start > All Programs > Accessories and right-click WordPad. Select Run as administrator. Click Yes to answer the User Access Control question.

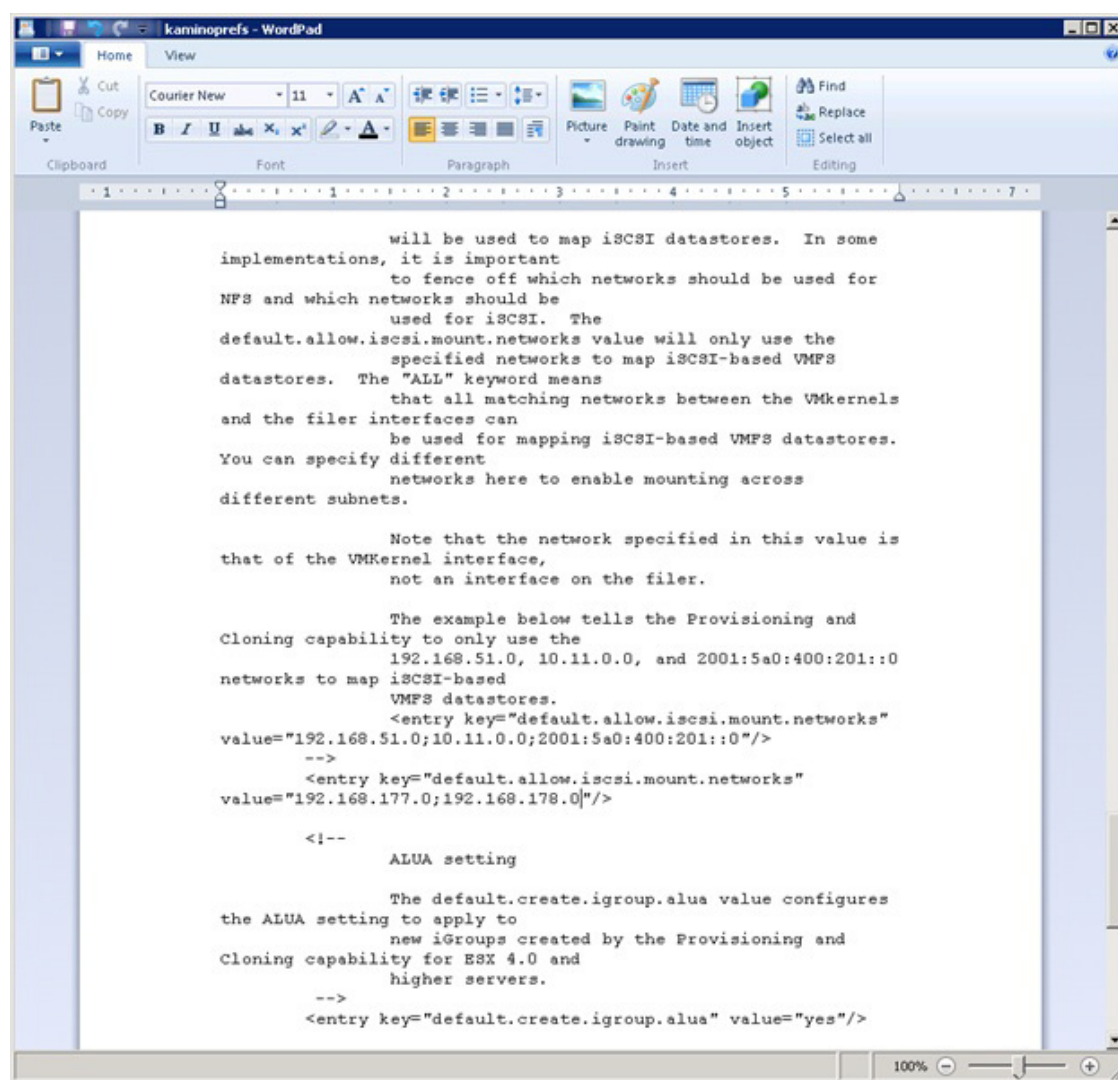
3. From the drop-down menu on the left, select Open. From the lower-right drop-down menu, select All Documents (*.*). Navigate to C:\Program Files\NetApp\Virtual Storage Console\etc\kamino.
4. Select the kaminoprefs file and click Open.
5. Scroll down to the Restrict NFS options section. In the entry key input the two iSCSI VLAN network addresses separated by a semicolon as shown.



6. Move to the Restrict iSCSI options section. In the entry key input the NFS VLAN network address as shown above.
7. Scroll down to the NFS networks section. In the entry key input the NFS VLAN network addresses as shown.



8. Move to the iSCSI networks section. In the entry key input the two iSCSI VLAN network addresses separated by a semicolon shown below.



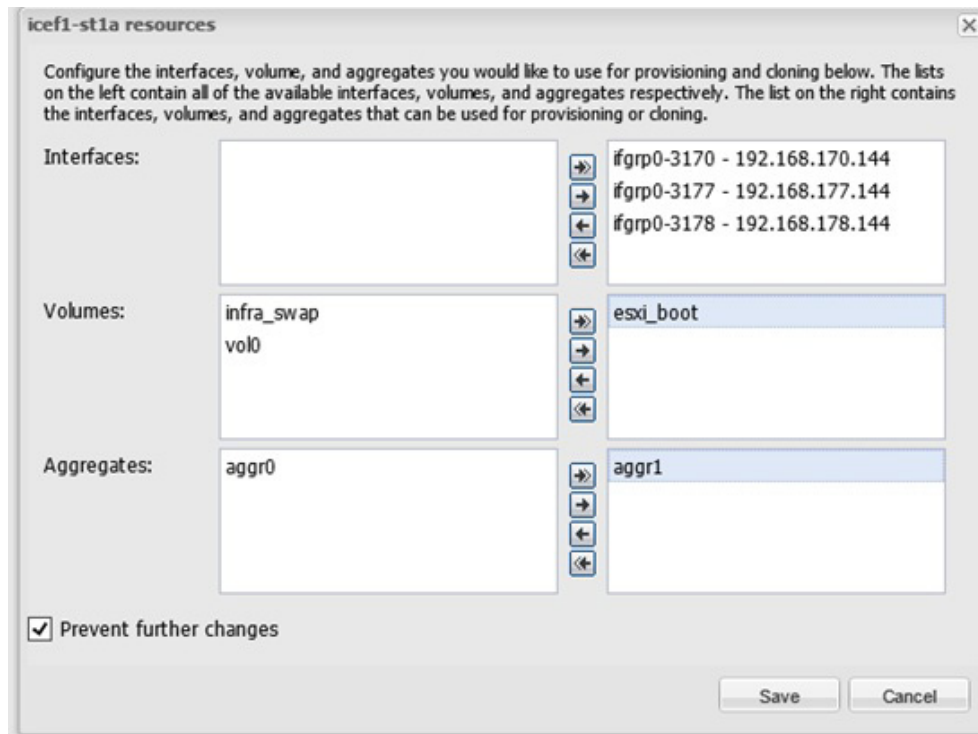
9. Close WordPad, selecting to save the changes to the file.
10. Reboot the VSC VM. After reboot is complete, close and log back into the vSphere client connected to vCenter.
11. Return to Discover and Add Storage Resources in the VSC Configuration section in the main part of this document.

VSC 4.1 Provisioning and Cloning Setup for 7-Mode Storage

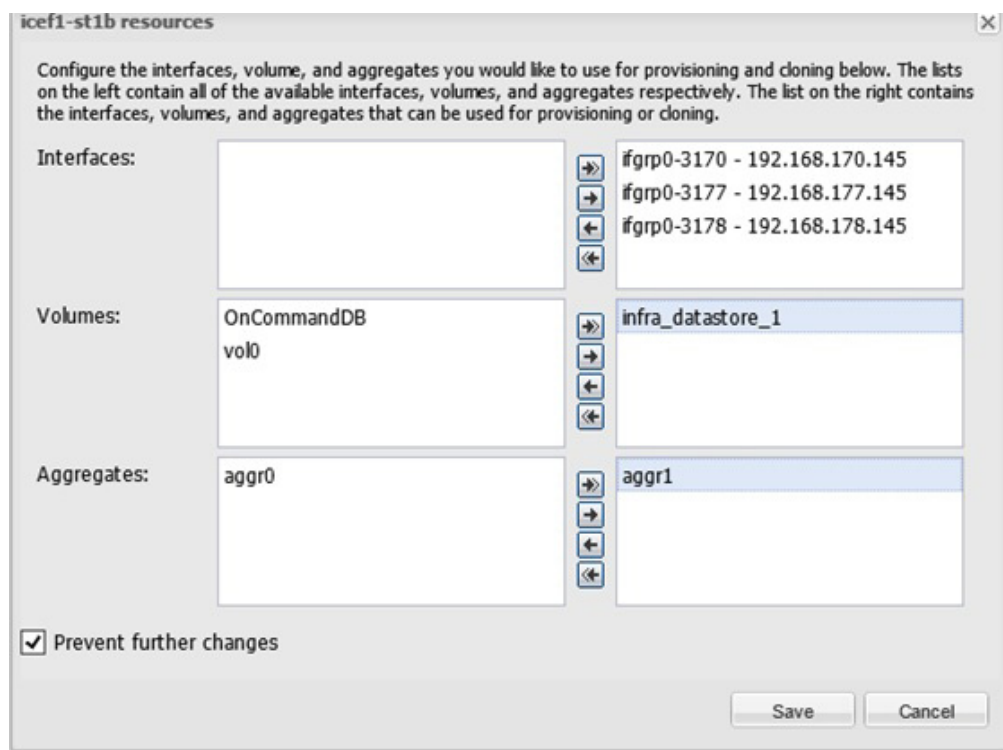
Provisioning and cloning in VSC 4.1 helps administrators to provision both VMFS and NFS datastores at the data center, datastore cluster, or host level in VMware environments.

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 <<var_controller1>> and select Resources.

3. In the <<var_controller1>> resources window, use the arrows to move all three Interfaces, Volume esxi_boot and aggr1 to the right. Select the Prevent further changes checkbox.



4. Click Save.
5. In the main part of the window, right-click <<var_controller2>> and select Resources.
6. In the <<var_controller2>> resources window, use the arrows to move all three Interfaces, Volume infra_datastore_1 and aggr1 to the right. Select the Prevent Further changes checkbox.



7. Click Save.
8. Return to section "VSC 4.1 Backup and Recovery."

NetApp VASA Provider (7-Mode Storage Only)

Install NetApp VASA Provider

To install NetApp VASA Provider, complete the following steps:

1. 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 <<var_ib-mgmt_vlan_id>> VLAN. The 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.
2. Log into the VASA Provider VM as the FlexPod admin user.
3. Download NetApp VASA Provider from the NetApp Support site.
4. Run the executable file netappvp-1-0-winx64.exe to start the installation.



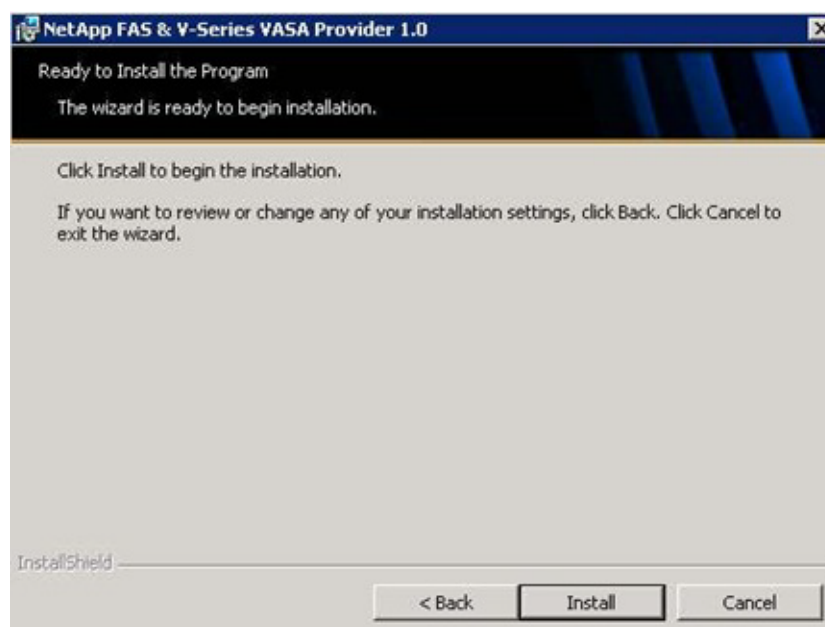
5. On the Welcome page of the installation wizard, click Next.



6. Select the installation location and click Next.



7. On the Ready to Install page, click Install.



8. Click Finish to complete the installation.



Configure NetApp VASA Provider

After NetApp VASA Provider is installed, it must be configured to communicate with the vCenter Server and retrieve storage system data. During configuration, specify a user name and password to register NetApp VASA Provider with the vCenter Server, and then add the storage systems before completing the process.

Add Storage Systems

The NetApp VASA Provider dialog box can be used to add the storage systems from which NetApp VASA Provider collects storage information. Storage systems can be added at any time.

To add a storage system, complete the following steps:

1. Double-click the VASA Configuration icon on your Windows desktop or right-click the icon and select Open to open the NetApp FAS/V-Series VASA Provider dialog box.
2. Click Add to open the Add Storage System dialog box.

NetApp FAS/V-Series VASA Provider 1.0

VASA Provider

Enter a user name and password for initial communication with vCenter Server

User Name: ice|icef1-admin Save

Password: Edit

Status: NetApp VASA Provider service is running

Alarm Thresholds

Threshold values are saved when you click the OK button

	Volume	Aggregate
Nearly Full Threshold (%):	85	90
Full Threshold (%):	90	95

VMware vCenter

Server Address: Port: 443 Register Provider

User Name: Unregister Provider

Password: OK Cancel

Or copy the URL below to register VASA Provider from VMware vSphere Client

VASA URL: https://ICEF1-VASA.ice.rtp.netapp.com:8443/services/vasaService

Storage Systems

Registered Storage Systems

Add Remove Edit

3. Enter the host name or IP address, port number, and user name and password for the storage system.

Add Storage System

Enter Storage System Credential Information

Storage System: 192.168.171.144

Protocol: ☒ HTTPS ☐ HTTP

Port: 443

User: root

Password: OK Cancel

4. Click OK to add the storage system.
5. Add both storage systems to the VASA Provider.

Register NetApp VASA Provider with vCenter Server

To establish a connection between the vCenter Server and NetApp VASA Provider, NetApp VASA Provider must be registered with the vCenter Server. The vCenter Server communicates with NetApp VASA Provider to obtain the information that NetApp VASA Provider collects from registered storage systems.

To register NetApp VASA Provider with the vCenter Server, complete the following steps:

1. Under Alarm Thresholds, accept or change the default threshold values for volume and aggregate. These values specify the percentages at which a volume or aggregate is full or nearly full.

The default threshold values are the following:

- 85% for a nearly full volume
- 90% for a full volume
- 90% for a nearly full aggregate
- 95% for a full aggregate



Note

After you finish registering NetApp VASA Provider with the vCenter Server, any changes made to the default threshold values are saved only when you click OK.

2. Under VMware vCenter, enter the host name or IP address of the vCenter Server machine and the user name and password for the vCenter Server.
3. Specify the port number to use, or accept the default port number for the vCenter Server.
4. Click Register Provider.
5. Click OK to commit all the details and register NetApp VASA Provider with the vCenter Server.



Note

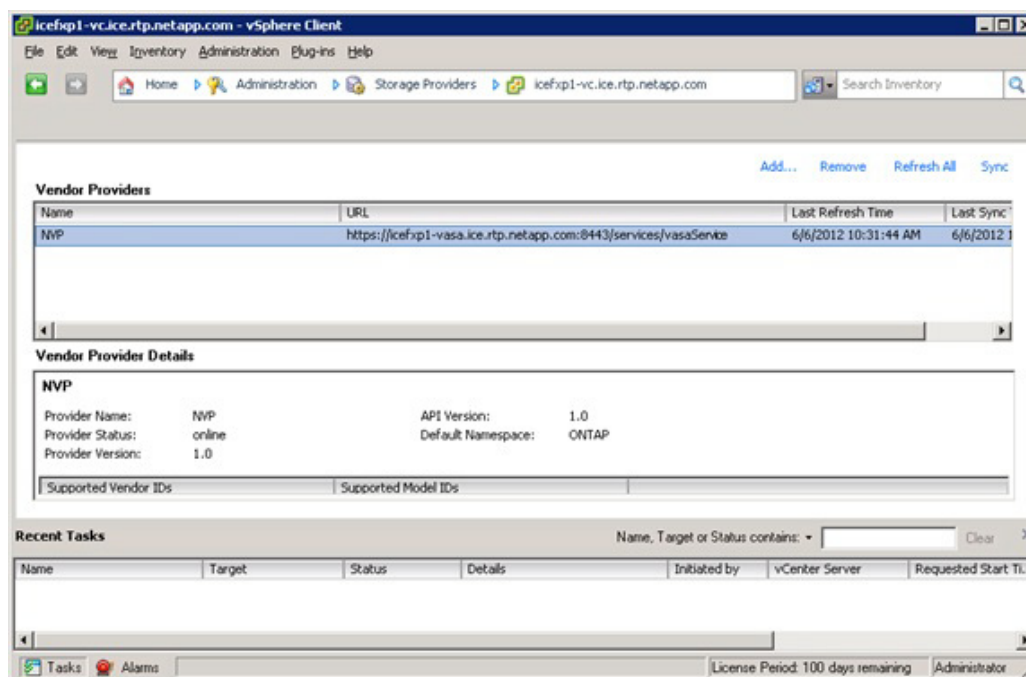
To use the vSphere Client to register NetApp VASA Provider with the vCenter Server, copy the URL from the VASA URL field and paste it into the vCenter Server.

6. Click OK to close the VASA Configuration window.

Verify VASA Provider in vCenter

1. Log in to vCenter using vSphere Client.
2. Click the Home tab at the upper-left portion 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 Inventory 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.

