



NetApp Verified Architecture

FlexPod Datacenter with SQL Server AlwaysON Availability Groups and All Flash FAS

NVA Deployment

NetApp Converged Infrastructure Engineering
December 2016 | NVA-0025-DEPLOY | Version 1.0

Reviewed by



TABLE OF CONTENTS

1	Executive Summary.....	7
1.1	FlexPod Program Benefits	7
2	Solution Overview	8
2.1	Target Audience.....	8
2.2	Solution Technology	9
2.3	Use Case Summary	10
3	Technology Overview.....	11
3.1	FlexPod.....	11
3.2	FlexPod: FCoE Direct-Connect Design.....	11
4	Technology Requirements	12
4.1	Hardware Requirements	12
4.2	Software Requirements	13
4.3	Configuration Guidelines.....	14
4.4	Physical Infrastructure	14
5	Network Switch Configuration	19
5.1	Physical Connectivity	19
5.2	Network Configuration Variables.....	19
5.3	FlexPod Cisco Nexus Base.....	20
6	FlexPod Cisco Nexus Switch Configuration.....	21
6.1	Enable Features.....	21
6.2	Set Global Configurations	22
6.3	Create VLANs	22
6.4	Add NTP Distribution Interface.....	22
6.5	Add Individual Port Descriptions for Troubleshooting	22
6.6	Create Port Channels	23
6.7	Configure Port Channel Parameters	24
6.8	Configure Virtual Port Channels.....	25
6.9	Uplink into Existing Network Infrastructure.....	26
7	Storage Configuration.....	26
7.1	NetApp Hardware Universe	26
7.2	Controllers.....	27
7.3	Disk Shelves	27
7.4	Clustered Data ONTAP 8.3.1	27

7.5	Complete the Configuration Worksheet	27
7.6	Log in to the Cluster	35
7.7	Zero All Spare Disks	35
7.8	Set Onboard Unified Target Adapter 2 Port Personality	35
7.9	Set Auto-Revert on Cluster Management	36
7.10	Set Up Management Broadcast Domain	36
7.11	Set Up Service Processor Network Interface	36
7.12	Create Aggregates	36
7.13	Verify Storage Failover	37
7.14	Disable Flow Control on UTA2 Ports	37
7.15	Configure NTP	38
7.16	Configure SNMP	38
7.17	Configure AutoSupport	39
7.18	Enable Cisco Discovery Protocol	39
7.19	Create Interface Groups	39
7.20	Create VLANs	39
7.21	Create Infrastructure Storage Virtual Machine	39
7.22	Create Load-Sharing Mirrors of SVM Root Volume	40
7.23	Configure HTTPS Access	40
7.24	Create FlexVol Volumes	41
7.25	Create Boot LUNs	41
7.26	Create Infrastructure LUNs	41
7.27	Schedule Deduplication	42
7.28	Create FCP LIFs	42
7.29	Add Infrastructure SVM Administrator	42
7.30	Create Production Storage Virtual Machine	43
7.31	Create Load-Sharing Mirrors of SVM Root Volume	43
7.32	Configure HTTPS Access	43
7.33	Create FCP LIFs	44
7.34	Add Production SVM Administrator	45
8	Server Configuration	45
8.1	Cisco UCS Base Configuration	45
8.2	Cisco UCS Configuration Variables	45
8.3	Perform Initial Setup of Cisco UCS 6248UP Fabric Interconnects for FlexPod Environments	46
8.4	Cisco UCS Setup	47
8.5	Add More Servers to FlexPod Unit	90

9	Storage Configuration—Infrastructure LUNs and Igroups	90
9.1	Clustered Data ONTAP Storage Setup.....	90
10	VMware vSphere 6.0 U1 Setup	92
10.1	VMware ESXi 6.0.....	92
10.2	VMware vCenter 6.0	100
11	FlexPod Management Tools Setup.....	104
11.1	Management Tools Configuration Variables	104
12	NetApp Virtual Storage Console 6.2P1 Deployment Procedure.....	105
12.1	Virtual Storage Console 6.2P1 Preinstallation Considerations.....	105
12.2	Install Virtual Storage Console 6.2P1	105
12.3	Register VSC with vCenter Server.....	106
12.4	Discover and Add Storage Resources	106
12.5	Optimal Storage Settings for ESXi Hosts	106
12.6	VSC 6.2P1 Backup and Recovery	107
13	NetApp OnCommand Unified Manager 6.3P4 Deployment Procedure	108
13.1	OnCommand Unified Manager 6.3P4	108
13.2	OnCommand Performance Manager 2.0	111
14	SQL Server Virtual Machine Deployment and Configuration	115
14.1	Deploy Production Datastores	116
14.2	Create Virtual Machines.....	116
14.3	Install Windows Server 2012 R2	116
14.4	Install Windows Failover Clustering	117
14.5	Create Windows Failover Cluster.....	117
14.6	Configure File Share Witness	118
14.7	Add Data Disks to the Virtual Machine.....	119
14.8	Install the SQL Server 2014 Software.....	119
14.9	Create AlwaysOn Availability Group	121
15	NetApp SnapCenter.....	122
15.1	Install and Configure SnapCenter 1.1	122
16	Solution Verification.....	124
16.1	Performance Validation.....	125
16.2	SnapCenter Validation	127
16.3	Resiliency Validation.....	128

17 Conclusion	129
Acknowledgements	129
References.....	129
Cisco Unified Computing System	129
Cisco Nexus Networking	130
NetApp FAS Storage	130
NetApp SnapCenter	130
VMware vSphere	130
Interoperability Matrices	131
Version History	131

LIST OF TABLES

Table 1) Solution hardware requirements.....	12
Table 2) Solution software requirements	13
Table 3) Required VLANs.....	14
Table 4) Cisco Nexus 9396-A cabling information.....	16
Table 5) Cisco Nexus 9396-B cabling information.....	16
Table 6) NetApp AFF8080 EX controller 1 cabling information.	17
Table 7) NetApp AFF8080 EX controller 2 cabling information.	17
Table 8) Cisco UCS6248 fabric interconnect A cabling information.	18
Table 9) Cisco UCS6248 fabric interconnect B cabling information.	18
Table 10) Network configuration variables.	19
Table 11) Storage configuration variables.....	27
Table 12) Cisco UCS configuration variables	45
Table 13) Boot LIF configuration.	81
Table 14) Storage target WWPNs.	90
Table 15) Host initiator WWPNs.....	90
Table 16) VMware configuration variables.	92
Table 17) FlexPod management tools configuration variables.....	105
Table 18) SQL Server disks and mount points.....	119
Table 19) Database LUN and volume configuration.....	125
Table 20) Validated database performance.....	127

LIST OF FIGURES

Figure 1) FlexPod component families.....	7
Figure 2) High-level solution architecture.....	9
Figure 3) FCoE direct-connect topology.....	12
Figure 4) FlexPod cabling diagram.....	15

1 Executive Summary

NetApp® Validated Architectures (NVAs) describe systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp developed to address the business needs of customers.

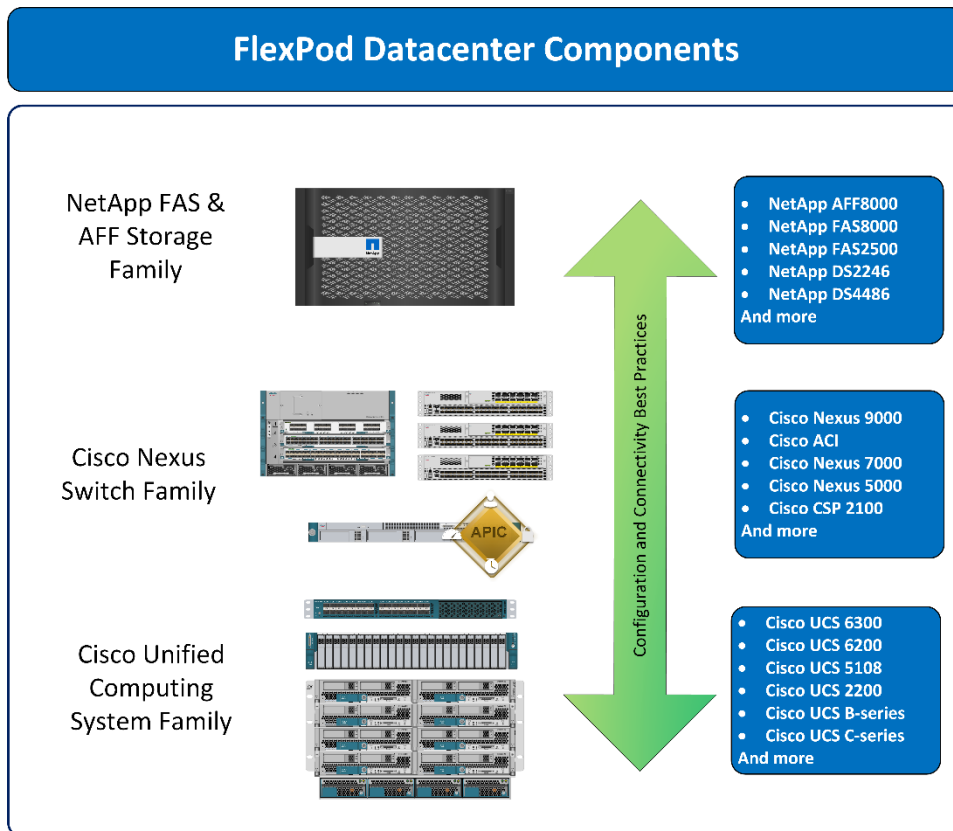
This document describes a reference architecture for Microsoft SQL Server 2014 with AlwaysOn Availability Groups built on the All Flash FAS (AFF) FlexPod® model. This solution includes Cisco UCS B200 M4 blade servers and NetApp AFF8080 EX storage arrays. The other major highlights of this solution are the use of NetApp SnapCenter® enterprise software for application-integrated database backup and recovery, VMware vSphere 6.0 virtualization, and Cisco Nexus 9000 Series switches for client traffic.

This document also discusses design choices and best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific components described in this document and are also applicable to other versions.

1.1 FlexPod Program Benefits

FlexPod is a predesigned, best practice data center architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp FAS and AFF storage systems. FlexPod is a suitable platform for running various virtualization hypervisors as well as bare-metal operating systems (OSs) and enterprise workloads. FlexPod delivers a baseline configuration and can also be sized and optimized to accommodate many different use cases and requirements. Figure 1 depicts the component families of the FlexPod solution.

Figure 1) FlexPod component families.



FlexPod provides a uniform approach to IT architecture, offering a well-characterized and documented shared pool of resources for application workloads. FlexPod delivers operational efficiency and consistency with the versatility to meet a variety of SLAs and IT initiatives, including the following:

- Application rollouts or migrations
- Business continuity and disaster recovery (DR)
- Desktop virtualization
- Cloud delivery models (public, private, and hybrid) and service models (IaaS, PaaS, and SaaS)
- Asset consolidation and virtualization
- Data center consolidation and footprint reduction

Cisco and NetApp thoroughly validated and verified the FlexPod solution architecture and its many use cases. In addition, they created a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model. This portfolio includes, but is not limited to, the following items:

- Best practice architectural design
- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for what is and what is not a FlexPod configuration)
- Frequently asked questions (FAQs)
- NVAs and Cisco Validated Designs (CVDs) that focus on a variety of use cases

Cisco and NetApp have also built a robust and experienced support team focused on FlexPod solutions, including customer account and technical sales representatives, professional services personnel, and technical support engineers. This support alliance provides customers and channel services partners with direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment. As a key FlexPod Cooperative Support partner, VMware provides the virtualization hypervisor and management tools for this verified design with VMware vSphere and VMware vCenter.

2 Solution Overview

The FlexPod architecture is designed to help manage infrastructure complexity with proven guidance and measurable value. By introducing standardization, FlexPod helps customers mitigate the risk and uncertainty intrinsic to planning, designing, and implementing a new data center infrastructure. The result is a more predictable and adaptable architecture capable of meeting and exceeding IT demands.

This FlexPod design describes the deployment of Microsoft SQL Server 2014 with AlwaysOn Availability Groups in a VMware vSphere virtualized environment. This AFF design demonstrates that, with the addition of flash storage, FlexPod architectures can meet the most demanding performance requirements and still deliver the values of a standardized shared infrastructure.

2.1 Target Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, and partner engineers. It also includes customers who want to take advantage of an infrastructure built to deliver high-performance and high-availability database services.


2.2 Solution Technology

Traditionally, Microsoft SQL Server databases relied on Microsoft Windows clustering with shared storage to achieve a high-availability (HA) configuration. Microsoft database technology has come a long way since then and is now composed of highly efficient technology based on Windows Server 2012 R2 and AlwaysOn Availability Groups.

This solution is based on SQL Server 2014 AlwaysOn Availability Groups implemented on Windows Server 2012 R2 virtual machines (VMs). SQL Server instances are deployed as VMware vSphere VMs on NetApp AFF storage arrays.

Figure 2 shows the high-level solution architecture, with the association between the primary and secondary database replicas on separate storage.

Figure 2) High-level solution architecture.



This design enables OLTP workloads to be serviced from the primary database replica while database writes are synchronously committed to the secondary database server. The secondary replica provides near-site disaster recovery and enhances the overall availability of the solution by providing database-level failover. Because AlwaysOn replication enables readable database secondaries, the secondary database replica can be used to offload backup and reporting jobs.

The secondary database replica is hosted on equivalent storage and compute resources to provide consistent performance in the event of a failover. As part of NetApp integrated data management capabilities, SnapCenter Server creates NetApp Snapshot® copies of secondary databases and creates clones of databases for secondary processing or test and development activities.

This solution uses NetApp AFF8080 EX storage arrays for both primary and secondary database copies. Doing so provides the highest levels of performance for both copies during normal operations and failover situations. When FCoE is used, LUNs are provisioned to ESXi servers to make vSphere Virtual Machine File System (VMFS) datastores for SQL Server OS disks. In addition, separate dedicated LUNs are provisioned to make datastores for each database. By using this technology, SQL Server 2014 databases deliver excellent performance, and native application integration enables seamless backup, recovery, and cloning.

Cisco UCS Fabric Interconnects (FI) and B-Series servers enhance this configuration by providing flexibility and availability in the compute layer. Redundant hardware and software components eliminate single points of failure and make sure that data traffic is uninterrupted in the event of a component failure. Cisco UCS Service profiles ease deployment and maintenance activities by creating consistent configurations that are easily updated and can be moved between physical blades as needed for maintenance or upgrades.

Cisco Nexus 9000 switches act as the access layer for the primary and secondary database environments. Cisco Nexus 9396PX switches are deployed in pairs, and the Cisco UCS Fabric Interconnects are connected to virtual port channels for maximum availability. Cisco Nexus 9000 switches provide 40Gb of switching capability and can participate in the Cisco Application Centric Infrastructure (ACI). However, these switches do not support the FC or FCoE storage protocols. To enable FCoE in this solution, storage controllers are connected directly to the Cisco UCS FI with redundant connections for each controller.

When combined into a complete infrastructure, this solution delivers the following benefits:

- Tier 1 SQL Server 2014 database performance on a standardized, shared infrastructure
- Database-level availability by using a SQL Server AlwaysOn synchronous replica (near-site disaster recovery)
- Integrated backup and recovery of SQL Server 2014 databases with NetApp SnapCenter
- Integrated cloning of databases for secondary processing or testing and development
- Hardware-level redundancy for all the major components by using Cisco UCS, Cisco Nexus, and NetApp availability features

2.3 Use Case Summary

The FlexPod Datacenter, Microsoft SQL Server 2014, and NetApp AFF solution is designed to provide enterprises with the performance, manageability, and reliability necessary for tier 1 application databases. To this end, the following use cases were configured and tested in the lab to demonstrate the performance and functionality of this design:

- SQL Server 2014 AlwaysOn Availability Groups with one primary database and one replica. This configuration generates 150,000 to 200,000 IOPS with an average read latency below 1ms for a typical OLTP workload while in a synchronized state. Secondary database copies run on separate storage and compute resources that have capabilities equivalent to the primary.
- Failover of the SQL Server 2014 AlwaysOn Availability Groups to the secondary copy. This use case demonstrates database resiliency and the comparable performance of the secondary copy in a synchronized AG configuration under a typical OLTP workload.
- Backup and recovery of SQL Server 2014 databases with NetApp SnapCenter 1.1 and the SnapCenter plug-ins for SQL Server. Backups are taken from the secondary database copy to offload backup and validation processing from the primary infrastructure.
- Cloning of SQL Server 2014 databases with NetApp SnapCenter for use in application testing and development environments.
- Failure testing of various infrastructure components while the environment is operating under a heavy OLTP workload to verify the resiliency and reliability of the overall architecture.

3 Technology Overview

3.1 FlexPod

FlexPod is a best practice data center architecture that includes three core components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus switches
- NetApp FAS and AFF storage systems

These components are connected and configured according to best practices of both Cisco and NetApp and provide the ideal platform for running various enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed). FlexPod also can scale out for environments that need multiple consistent deployments (rolling out additional FlexPod stacks). FlexPod delivers a baseline configuration and has the flexibility to be sized and optimized to accommodate many different use cases.

Typically, the more scalable and flexible a solution is, the more difficult it becomes to maintain a single unified architecture capable of offering the same features and functionality across each implementation. This capability is one of the key benefits of FlexPod. Each of the component families shown in Figure 1 offers platform and resource options to scale the infrastructure up or down. These families do so supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

FlexPod addresses four primary design principles: availability, scalability, flexibility, and manageability. These architecture goals are as follows:

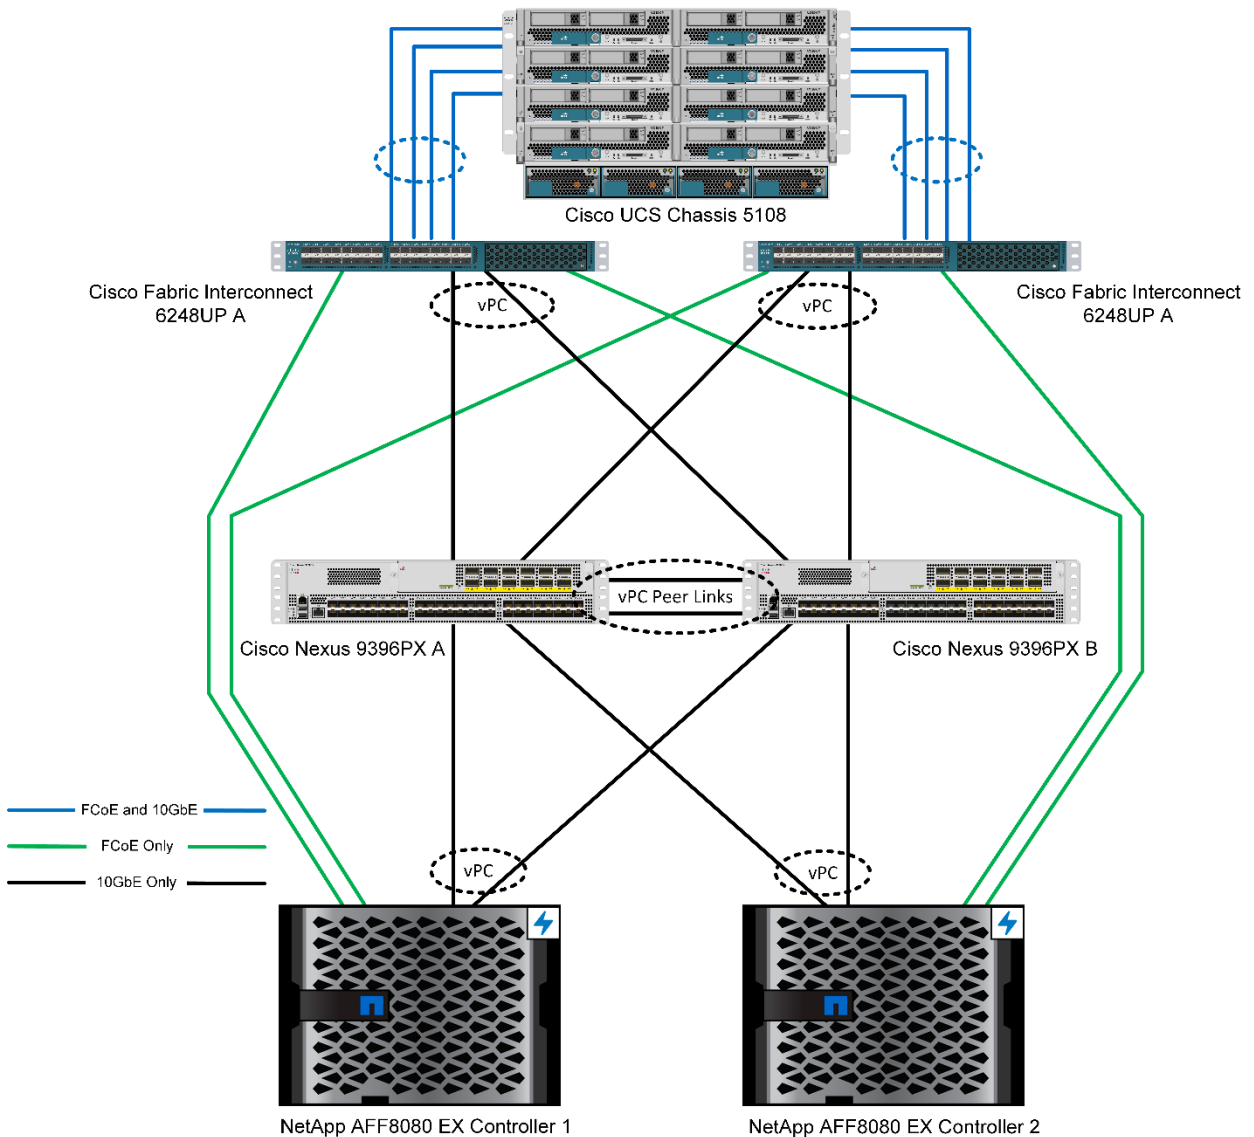
- **Application availability.** Makes sure that services are accessible and ready to use.
- **Scalability.** Addresses increasing demands with appropriate resources.
- **Flexibility.** Provides new services or recovers resources without infrastructure modification requirements.
- **Manageability.** Facilitates efficient infrastructure operations through open standards and APIs.

For details about the specific products used in this validation and the design factors considered when creating this reference architecture, refer to [NVA-0025-DESIGN FlexPod Datacenter with SQL Server AlwaysOn Availability Groups and AFF](#).

3.2 FlexPod: FCoE Direct-Connect Design

As noted previously, flexibility is a key design principle of FlexPod. The Cisco Nexus 9396PX switches used in this design support 10GbE and 40GbE networking and Cisco ACI, but the switches do not support FC or FCoE protocols. FlexPod can still support these protocols by connecting the Unified Target Adapter ports on NetApp storage controllers directly to the Cisco UCS Fabric Interconnects. Doing so applies the FC features of the FIs to provide name services and zoning capabilities to enable Cisco UCS servers to boot from and access FC or FCoE storage without additional FC switches. Figure 3 shows the basic topology of the direct-connect design.

Figure 3) FCoE direct-connect topology.



4 Technology Requirements

4.1 Hardware Requirements

Table 1 lists the hardware components required to implement the solution as tested. The hardware components used in any implementation of this solution might vary based on customer requirements.

Table 1) Solution hardware requirements.

Hardware	Configuration
Cisco UCS 6200 Series Fabric Interconnects	2x Cisco UCS 6248UP Fabric Interconnects Includes Cisco UCS Manager
Cisco UCS B200 M4	2x Xeon E5-2690 CPU (12 cores/each 2.6Ghz)

Hardware	Configuration
	128GB RAM/blade, 1 VIC 1340/blade
Cisco UCS 5108 chassis	Includes 2x Cisco UCS-IOM 2204XP
Cisco Nexus 9396PX	No expansion modules
NetApp AFF8080 EX	No additional PCI cards
NetApp DS2246 disk shelves	2x disk shelves with 800GB SSDs

4.2 Software Requirements

Table 2 lists the software components required to implement the solution as tested. The software components used in any implementation of the solution might vary based on customer requirements.

Table 2) Solution software requirements.

Software/Firmware	Version
Compute	
Cisco UCS Manager	2.2(5d)
Networking	
Cisco Nexus 9396PX	NX-OS software release 7.0(3)I1(3)
Storage	
NetApp clustered Data ONTAP®	8.3.1
NetApp Windows PowerShell toolkit	4.1.0
NetApp System Manager	8.3.1
NetApp Virtual Storage Console for VMware vSphere (VSC)	6.2
NetApp SnapCenter Server	1.1
SnapCenter Plug-In for Microsoft Windows	1.1
SnapCenter Plug-In for Microsoft SQL Server	1.1
VMware vSphere	
VMware ESXi	6.0.0 3380124
VMware vCenter Server	6.0.0 2656761
VMware vSphere PowerCLI	6.0 release 3 build 3205540
Database Server	
Microsoft Windows Server	2012 R2 Standard
Microsoft SQL Server	SQL Server 2014 Service Pack 1 Enterprise Edition

4.3 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 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document. Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are similarly configured. Also, this document details the steps for provisioning multiple Cisco UCS hosts. These examples are identified as VM-Host-Infra-01 and VM-Host-Prod-02 to represent infrastructure and production hosts deployed to each of the fabric interconnects in this document. 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:

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 3 describes the VLANs used during the validation of this solution. Customers should use values appropriate to their environment for all configuration variables.

Table 3) Required VLANs.

VLAN Name	VLAN Purpose	ID Used for Validation in This Document
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
vMotion	VLAN for VMware vMotion	3173
VM-Traffic	VLAN for production VM interfaces	3174

4.4 Physical Infrastructure

FlexPod Cabling

The information in this section is 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 the details for the prescribed and supported configuration of the NetApp AFF8080 running clustered Data ONTAP 8.3.1.

Note: For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool](#) (IMT).

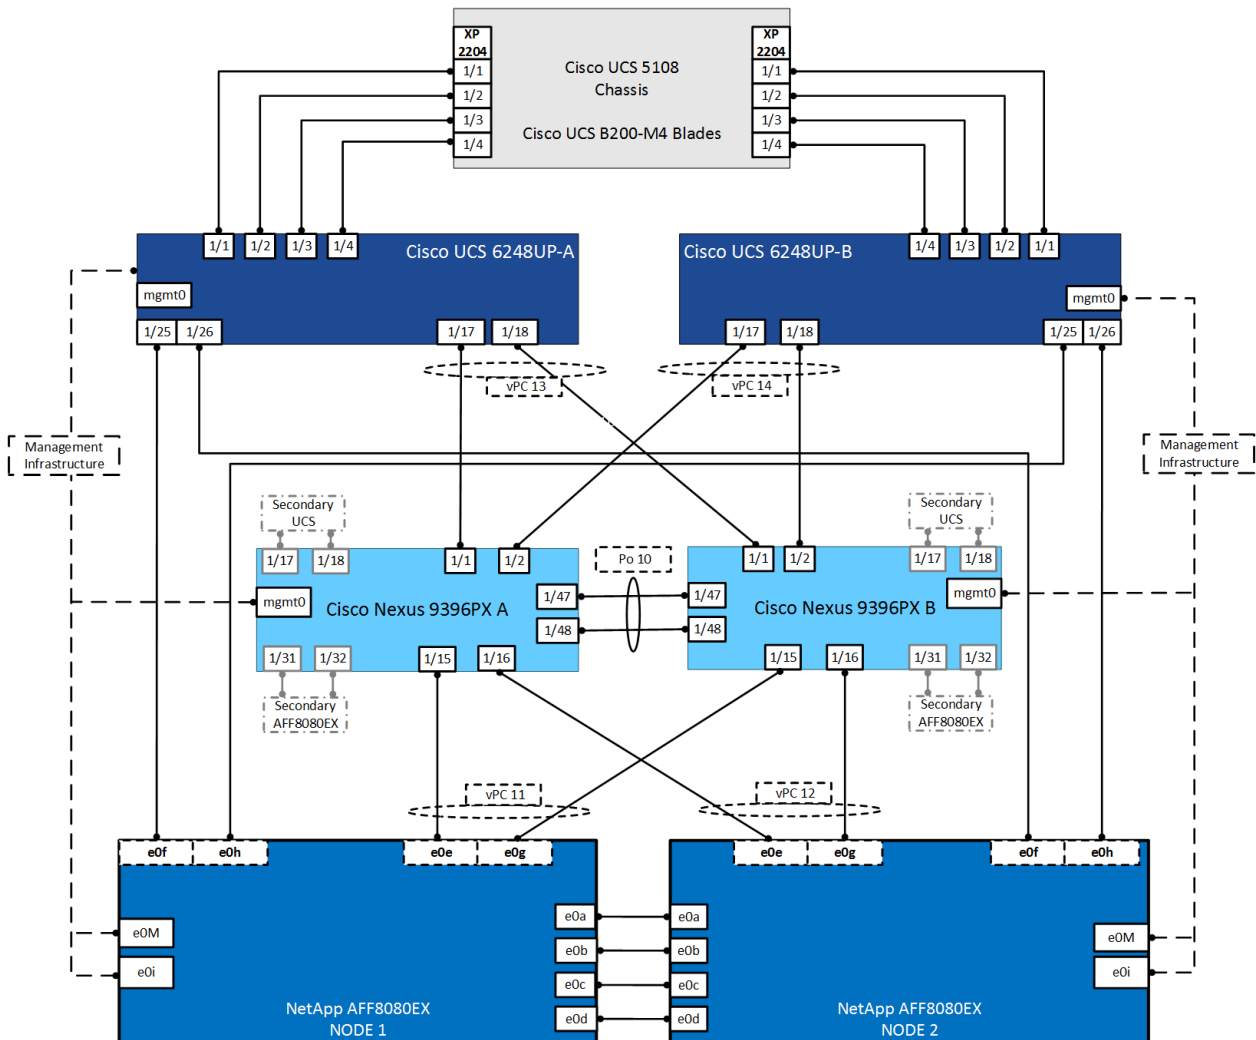
This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps

The procedures in this document are based on the cabling as depicted below. Changes to the cabling described here require changes to the deployment procedures that follow because specific port locations are referenced.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the [Universal SAS and ACP Cabling Guide](#).

In this document, the AFF8080 and Cisco UCS servers used to host one of the SQL replicas is referred to as **primary** and the equipment used to host the other replica is referred to as **secondary**. Figure 4 shows the connectivity between the primary NetApp AFF8080 and Cisco UCS 6248UP Fabric Interconnects and the Cisco Nexus 9396PX network switches. The secondary AFF8080 and Cisco UCS are connected in an identical configuration except for the Cisco Nexus switch ports, as noted in the illustrations and tables.

Figure 4) FlexPod cabling diagram.



Tables 4 through 9 provide the details of all the connections in use. The secondary storage and Cisco UCS systems use the same cabling as described in Tables 4 through 9 and Figure 4.

Table 4) Cisco Nexus 9396-A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396 A	Eth1/1	10GbE	Cisco UCS 6248UP FI A (primary replica)	eth1/17
	Eth1/2	10GbE	Cisco UCS 6248UP FI B (primary replica)	eth1/17
	Eth1/15	10GbE	NetApp AFF8080 EX controller 1 (primary replica)	e0e
	Eth1/16	10GbE	NetApp AFF8080 EX Controller 2 (primary replica)	e0e
	Eth1/17	10GbE	Cisco UCS 6248UP FI A (secondary replica)	eth1/17
	Eth1/18	10GbE	Cisco UCS 6248UP FI B (secondary replica)	eth1/17
	Eth1/31	10GbE	NetApp AFF8080 EX Controller 1 (secondary replica)	e0e
	Eth1/32	10GbE	NetApp AFF8080 EX Controller 2 (secondary replica)	e0e
	Eth1/47	10GbE	Cisco Nexus 9396PX B	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 9396PX B	Eth1/48
	MGMT0	GbE	GbE management switch	Any

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

Table 5) Cisco Nexus 9396-B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396 B	Eth1/1	10GbE	Cisco UCS 6248UP FI A (primary replica)	18
	Eth1/2	10GbE	Cisco UCS 6248UP FI B (primary replica)	18
	Eth1/15	10GbE	NetApp AFF8080 EX controller 1 (primary replica)	e0g
	Eth1/16	10GbE	NetApp AFF8080 EX controller 2 (primary replica)	e0g
	Eth1/17	10GbE	Cisco UCS 6248UP FI A (secondary replica)	18
	Eth1/18	10GbE	Cisco UCS 6248UP FI B (secondary replica)	18
	Eth1/31	10GbE	NetApp AFF8080 EX controller 1 (secondary replica)	e0g

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/48	10GbE	NetApp AFF8080 EX controller 2 (secondary replica)	e0g
	Any	10GbE	Cisco Nexus 9396PX A	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 9396PX A	
	MGMT0	GbE	GbE management switch	

Table 6) NetApp AFF8080 EX controller 1 cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1 Connections for secondary replica noted in parentheses	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp AFF8080 EX controller 2	e0a
	e0b	10GbE	NetApp AFF8080 EX controller 2	e0b
	e0c	10GbE	NetApp AFF8080 EX controller 2	e0c
	e0d	10GbE	NetApp AFF8080 EX controller 2	e0d
	0e	10GbE	Cisco Nexus 9396PX A	Eth1/15 (eth1/31)
	0f	10GbE / FCoE	Cisco UCS 6248UP FI A	eth1/25
	0g	10GbE	Cisco Nexus 9396PX B	Eth1/15 (eth1/31)
	0h	10GbE / FCoE	Cisco UCS 6248UP FI B	eth1/25

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

Table 7) NetApp AFF8080 EX controller 2 cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2 Connections for secondary replica noted in parentheses	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp AFF8080 EX controller 1	e0a
	e0b	10GbE	NetApp AFF8080 EX controller 1	e0b

Local Device	Local Port	Connection	Remote Device	Remote Port
	e0c	10GbE	NetApp AFF8080 EX controller 1	e0c
	e0d	10GbE	NetApp AFF8080 EX controller 1	e0d
	0e	10GbE	Cisco Nexus 9396PX A	Eth1/16 (eth1/32)
	0f	10GbE / FCoE	Cisco UCS 6248UP FI A	eth1/26
	0g	10GbE	Cisco Nexus 9396PX B	Eth1/16 (eth1/32)
	0h	10GbE / FCoE	Cisco UCS 6248UP FI B	eth1/26

Table 8) Cisco UCS6248 fabric interconnect A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A Connections for secondary replica noted in parentheses	Eth1/1	10GbE	Cisco UCS Chassis 1 2204 FEX A	IOM 1/1
	Eth1/2	10GbE	Cisco UCS Chassis 1 2204 FEX A	IOM 1/2
	Eth1/3	10GbE	Cisco UCS Chassis 1 2204 FEX A	IOM 1/3
	Eth1/4	10GbE	Cisco UCS Chassis 1 2204 FEX A	IOM 1/4
	Eth1/17	10GbE	Cisco Nexus 9396PX A	Eth1/1 (eth1/17)
	Eth1/18	10GbE	Cisco Nexus 9396PX B	Eth1/1 (eth1/17)
	Eth1/25	10GbE / FCoE	NetApp AFF8080 EX controller 1	e0f
	Eth1/26	10GbE / FCoE	NetApp AFF8080 EX controller 2	e0f
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI B	L1
	L2	GbE	Cisco UCS 6248UP FI B	L2

Table 9) Cisco UCS6248 fabric interconnect B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B Connections for secondary replica noted in parentheses	Eth1/1	10GbE	Cisco UCS Chassis 1 2204 FEX B	IOM 1/1
	Eth1/2	10GbE	Cisco UCS Chassis 1 2204 FEX B	IOM 1/2
	Eth1/3	10GbE	Cisco UCS Chassis 1 2204 FEX B	IOM 1/3
	Eth1/4	10GbE	Cisco UCS Chassis 1 2204 FEX B	IOM 1/4

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/17	10GbE	Cisco Nexus 9396PX A	Eth1/2 (eth1/18)
	Eth1/18	10GbE	Cisco Nexus 9396PX B	Eth1/2 (eth1/18)
	Eth1/25	10GbE / FCoE	NetApp AFF8080 EX controller 1	e0h
	Eth1/26	10GbE / FCoE	NetApp AFF8080 EX controller 2	e0h
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI A	L1
	L2	GbE	Cisco UCS 6248UP FI A	L2

5 Network Switch Configuration

This section details the procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely; failure to do so could result in an improper configuration.

5.1 Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in the section “FlexPod Cabling.”

5.2 Network Configuration Variables

Table 10 lists the variables used in the configuration of the Cisco Nexus 9396PX switches.

Table 10) Network configuration variables.

Variable	Value
<<var_password>>	
<<var_nexus_A_hostname>>	
<<var_nexus_A_mgmt0_ip>>	
<<var_nexus_A_mgmt0_netmask>>	
<<var_nexus_A_mgmt0_gw>>	
<<var_global_ntp_server_ip>>	
<<var_nexus_B_hostname>>	
<<var_nexus_B_mgmt0_ip>>	
<<var_nexus_B_mgmt0_netmask>>	
<<var_nexus_B_mgmt0_gw>>	
<<var_ib-mgmt-vlan_gateway>>	
<<var_ib-mgmt_vlan_id>>	
<<var_native_vlan_id>>	
<<var_vmotion_vlan_id>>	
<<var_vm-traffic_vlan_id>>	

Variable	Value
<<var_switch_a_ntp_ip>>	
<<var_ib-mgmt_vlan_netmask_length>>	
<<var_switch_b_ntp_ip>>	
<<var_pri_ucs_6248_clustername>>	
<<var_sec_ucs_6248_clustername>>	
<<var_pri_node01>>	
<<var_pri_node02>>	
<<var_sec_node01>>	
<<var_sec_node02>>	
<<var_nexus_vpc_domain_id>>	

5.3 FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I1(3).

Note: The following procedure includes the setup of NTP distribution on the in-band management VLAN. The `interface-vlan` feature and `ntp` commands are used for this setup. This procedure also assumes that the default VRF will be used to route the in-band management VLAN.

Set Up Initial Configuration

Cisco Nexus 9396PX A

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

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the Cisco 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]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_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: <<var_nexus_A_hostname>>
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]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus 9396PX B

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

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the Cisco 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]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_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: <<var_nexus_B_hostname>>
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]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

6 FlexPod Cisco Nexus Switch Configuration

6.1 Enable Features

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To license the Cisco Nexus switches, complete the following steps:

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

```
config t
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

6.2 Set Global Configurations

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To set global configurations, complete the following step on both switches:

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <<var_global_ntp_server_ip>> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <<var_ib-mgmt-vlan_gateway>>
copy run start
```

6.3 Create VLANs

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To create the necessary virtual LANs, complete the following step on both switches:

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

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
vlan <<var_native_vlan_id>>
name Native-VLAN
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
exit
```

6.4 Add NTP Distribution Interface

Cisco Nexus 9396PX A

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

```
ntp source <<var_switch_a_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <<var_switch_a_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

Cisco Nexus 9396PX B

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

```
ntp source <<var_switch_b_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <<var_switch_b_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

6.5 Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 9396PX 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 Eth1/1
description <<var_pri_ucs_6248_clustername>>-a:1/17
interface Eth1/2
description <<var_pri_ucs_6248_clustername>>-b:1/17
interface Eth1/15
description <<var_pri_node01>>:e0e
interface Eth1/16
description <<var_pri_node02>>:e0e
interface Eth1/17
description <<var_sec_ucs_6248_clustername>>-a:1/17
interface Eth1/18
description <<var_sec_ucs_6248_clustername>>-b:1/17
interface Eth1/31
description <<var_sec_node01>>:e0e
interface Eth1/32
description <<var_sec_node02>>:e0e
interface Eth1/47
description <<var_nexus_B_hostname>>:1/47
interface Eth1/48
description <<var_nexus_B_hostname>>:1/48
exit
```

Cisco Nexus 9396PX B

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

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

```
interface Eth1/1
description <<var_pri_ucs_6248_clustername>>-a:1/18
interface Eth1/2
description <<var_pri_ucs_6248_clustername>>-b:1/18
interface Eth1/15
description <<var_pri_node01>>:e0g
interface Eth1/16
description <<var_pri_node02>>:e0g
interface Eth1/17
description <<var_sec_ucs_6248_clustername>>-a:1/18
interface Eth1/18
description <<var_sec_ucs_6248_clustername>>-b:1/18
interface Eth1/31
description <<var_sec_node01>>:e0g
interface Eth1/32
description <<var_sec_node02>>:e0g
interface Eth1/47
description <<var_nexus_A_hostname>>:1/47
interface Eth1/48
description <<var_nexus_A_hostname>>:1/48
exit
```

6.6 Create Port Channels

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

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
interface Eth1/47-48
channel-group 10 mode active
no shutdown
interface Po11
description <<var_pri_ucs_6248_clustername>>-a
```

```

interface Eth1/1
channel-group 11 mode active
no shutdown
interface Po12
description <<var_pri_ucs_6248_clustername>>-b
interface Eth1/2
channel-group 12 mode active
no shutdown
interface Po15
description <<var_pri_node01>>
interface Eth1/15
channel-group 15 mode active
no shutdown
interface Po16
description <<var_pri_node02>>
interface Eth1/16
channel-group 16 mode active
no shutdown
interface Po21
description <<var_sec_ucs_6248_clustername>>-a
interface Eth1/17
channel-group 21 mode active
no shutdown
interface Po22
description <<var_sec_ucs_6248_clustername>>-b
interface Eth1/18
channel-group 22 mode active
no shutdown
interface Po25
description <<var_sec_node01>>
interface Eth1/31
channel-group 25 mode active
no shutdown
interface Po26
description <<var_sec_node02>>
interface Eth1/32
channel-group 26 mode active
no shutdown
exit
copy run start

```

6.7 Configure Port Channel Parameters

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To configure port channel parameters, 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 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
spanning-tree port type network
interface Po11
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po12
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
spanning-tree port type edge trunk
mtu 9216

```



```

interface Po15
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po16
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po21
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po22
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po25
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
interface Po26
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
copy run start

```

6.8 Configure Virtual Port Channels

Cisco Nexus 9396PX 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>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po11
vpc 11
interface Po12
vpc 12
interface Po15
vpc 15
interface Po16
vpc 16
interface Po21
vpc 21
interface Po22

```

```
vpc 22
interface Po25
vpc 25
interface Po26
vpc 26
exit
copy run start
```

Cisco Nexus 9396PX 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>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po11
vpc 11
interface Po12
vpc 12
interface Po15
vpc 15
interface Po16
vpc 16
interface Po21
vpc 21
interface Po22
vpc 22
interface Po25
vpc 25
interface Po26
vpc 26
exit
copy run start
```

6.9 Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. For an existing Cisco Nexus environment, NetApp recommends using vPCs to uplink the Cisco Nexus 9396PX switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is completed.

7 Storage Configuration

7.1 NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. The application provides configuration information for all of the NetApp storage appliances supported by ONTAP software. It also provides a table of component compatibilities.

Note: Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the [HWU](#) application.

1. Access the [HWU](#) application to view the System Configuration guides. Click the Platforms tab to view the compatibility between different versions of ONTAP and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.
3. HWU also provides site requirement information, such as about:
 - Site preparation
 - System connectivity requirements
 - Circuit breaker, power outlet balancing, system cabinet power cord plugs, and console pinout requirements

7.2 Controllers

Follow the physical installation procedures for the controllers found in the [AFF8000 series product documentation](#) at the [NetApp Support](#) site.

7.3 Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by AFF80xx is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to the [SAS Disk Shelves Universal SAS and ACP Cabling Guide](#) for cabling guidelines.

7.4 Clustered Data ONTAP 8.3.1

The following procedure documents the deployment of a single Data ONTAP cluster. For this reference architecture, two clusters were deployed to host primary and secondary SQL replicas. After the primary replica cluster is deployed, repeat this procedure to deploy the secondary cluster.

7.5 Complete the Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the [Clustered Data ONTAP 8.3 Software Setup Guide](#). You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

Storage Configuration Variables

Table 11 lists the values that are required for configuration of the NetApp AFF8080 EX storage systems.

Table 11) Storage configuration variables.

Variable	Value
<<var_node01_mgmt_ip>>	
<<var_node01_mgmt_mask>>	
<<var_node01_mgmt_gateway>>	
<<var_node02_mgmt_ip>>	
<<var_node02_mgmt_mask>>	
<<var_mgmt_gateway>>	
<<var_url_boot_software>>	
<<var_password>>	
<<var_clustername>>	
<<var_cluster_base_license_key>>	

Variable	Value
<<var_clustermgmt_ip>>	
<<var_clustermgmt_mask>>	
<<var_clustermgmt_port>>	
<<var_clustermgmt_gateway>>	
<<var_fcp_license>>	
<<var_dns_domain_name>>	
<<var_nameserver_ip>>	
<<var_node_location>>	
<<var_node01>>	
<<var_node02>>	
<<var_node01_sp_ip>>	
<<var_node01_sp_mask>>	
<<var_node01_sp_gateway>>	
<<var_node02_sp_ip>>	
<<var_node02_sp_mask>>	
<<var_node02_sp_gateway>>	
<<var_timezone>>	
<<var_snmp_contact>>	
<<var_snmp_location>>	
<<var_oncommand_server_fqdn>>	
<<var_snmp_community>>	
<<var_mailhost>>	
<<var_storage_admin_email>>	
<<var_cert_common_name>>	
<<var_cert_country>>	
<<var_cert_state>>	
<<var_cert_locality>>	
<<var_cert_org>>	
<<var_cert_unit>>	
<<var_cert_email>>	
<<var_cert_days>>	
<<var_infra_svm_mgmt_ip>>	
<<var_infra_svm_mgmt_mask>>	
<<var_infra_svm_mgmt_gateway>>	
<<var_prod_svm_mgmt_ip>>	
<<var_prod_svm_mgmt_mask>>	
<<var_prod_svm_mgmt_gateway>>	

Install ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Clustered Data ONTAP 8.3 Software Setup Guide](#) to learn about configuring ONTAP. Table 11 lists the information that you will need

to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Configure Node 01

To configure node 01, complete the following steps:

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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

Note: If Data ONTAP 8.3.1 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.1 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.
9. Enter the URL where the software can be found.

Note: This web server should be pingable.

```
<<var_url_boot_software>>
```

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

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-A prompt. If these actions occur, the system might deviate from this procedure.

11. Press Ctrl-C when you see the boot menu message:

```
Press Ctrl-C for Boot Menu
```

12. Select option 4 for Clean Configuration and Initialize All Disks.
13. Enter `y` to zero disks, reset config, and install a new file system.
14. Enter `y` to erase all the data on the disks.

Note: The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, complete the following steps:

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. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

Note: If Data ONTAP 8.3.1 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.1 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Enter e0M for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.
9. Enter the URL where the software can be found.

Note: This web server must be pingable.

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

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-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C when you see the boot message:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

Note: The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3.1 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```
Welcome to node setup.
```

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

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

```
This system will send event messages and weekly reports to NetApp Technical Support.  
To disable this feature, enter "autosupport modify -support disable" within 24 hours.  
Enabling AutoSupport can significantly speed problem determination and resolution should a  
problem occur on your system.  
For further information on AutoSupport, see:  
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
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>>
A node management interface on port e0M with IP address <<var_node01_mgmt_ip>> has been created

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility
from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

2. Press Enter and log in to the node with the admin user ID and no password.
3. At the node command prompt, enter the following commands to set HA mode for storage failover.
Note: If the node responds that the HA mode was already set, proceed to step 4.

```
::> storage failover modify -mode ha
Mode set to HA. Reboot node to activate HA.

::> system node reboot

Warning: Are you sure you want to reboot node "localhost"? {y|n}: y
```

4. After reboot, set up the node with the preassigned values.

```
Welcome to node setup.

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 setup wizard.
  Any changes you made before quitting will be saved.

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

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
Enter the node management interface default gateway [<<var_node01_mgmt_gateway>>]: Enter

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility
from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

5. Log in to the node as the admin user with no password.
6. Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

In ONTAP, 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 node 01.

Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup
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 with the factory default settings and then enter the `cluster setup` command.

To create a cluster, complete the following steps:

1. Enter `Create` to create a new cluster:
2. Enter `no` for the single-node cluster option.
3. Enter `no` to configure a switchless cluster.
4. The system defaults are displayed. Enter `yes` to use the system defaults. Use the following prompts to configure the cluster ports:

```
Existing cluster interface configuration found:

Port      MTU      IP              Netmask
e0a       9000     169.254.123.185 255.255.0.0
e0b       9000     169.254.116.155 255.255.0.0
e0c       9000     169.254.142.211 255.255.0.0
e0d       9000     169.254.34.22   255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: no

System Defaults:
Private cluster network ports [e0a,e0c].
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]: yes
```

5. Complete the following steps to create a cluster:

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>

It can take several minutes to create cluster interfaces...

Step 1 of 5: Create a Cluster
You can type "back", "exit", or "help" at any question.

Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_fcp_license>>
```

Note: The cluster is created. Creation can take a few minutes.

Note: For this validated architecture, NetApp recommends installing license keys for NetApp SnapRestore® data recovery software, NetApp FlexClone® data replication technology, and the NetApp SnapManager® suite. In addition, install all required storage protocol licenses and all licenses that came with the AFF bundle. After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0e]: e0i
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 the IP addresses with a comma.

7. Set up the node.

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

The node management interface has been modified to use port e0M with IP address
<<var_node01_mgmt_ip>>.
This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a
problem occur on your system.
For further information on AutoSupport, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter
Cluster "<<var_clustername>>" has been created.
To complete cluster setup, you must join each additional node to the cluster by running "cluster
setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide
for information about additional system configuration tasks. You can find the Software Setup
Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP
command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address
(<<var_clustermgmt_ip>>).
To access the command-line interface, connect to the cluster management IP address (for example,
ssh admin@<<var_clustermgmt_ip>>).

<<var_clustername>>:::
```

Note: The node management interface can be on the same subnet as the cluster management interface or it can be on a different subnet. In this document, we assumed that the node management interface is on the same subnet.

Join Node 02 to Cluster

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 node 01 and the node joining the cluster in this example is node 02.

To join node 02 to the existing cluster, complete the following steps:

1. If prompted, enter `admin` in the login prompt.
2. Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup

This node's storage failover partner is already a member of a cluster.
Storage failover partners must be members of the same cluster.
The cluster setup wizard will default to the cluster join dialog.

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?
{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.

3. Enter `join` to join a cluster:
4. Data ONTAP detects the existing cluster and agrees to join the same cluster. Follow the prompts to join the cluster.

```
Existing cluster interface configuration found:

Port      MTU      IP          Netmask
e0a       9000     169.254.82.126 255.255.0.0
e0b       9000     169.254.70.118 255.255.0.0
e0c       9000     169.254.112.77 255.255.0.0
e0d       9000     169.254.249.220 255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: no

System Defaults:
Private cluster network ports [e0a,e0c].
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]:Enter
It can take several minutes to create cluster interfaces...
```

5. Complete the following steps to join a cluster:

```
Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.

Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
Joining cluster <<var_clustername>>
Starting cluster support services ..

This node has joined the cluster <<var_clustername>>.

Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.

SFO is enabled.

Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.

Notice: HA is configured in management.
```

Note: The node should find the cluster name. Joining the cluster can take a few minutes.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node02_mgmt_mask>>]: Enter
Enter the node management interface default gateway [<<var_node02_mgmt_gateway>>]: Enter
The node management interface has been modified to use port e0M with IP address
<<var_node02_mgmt_ip>>.
This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a
problem occur on your system.
For further information on AutoSupport, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter

This node has been joined to cluster "<<var_clustername>>".
To complete cluster setup, you must join each additional node to the cluster by running "cluster
setup" on each node.
```

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for information about additional system configuration tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address (<<var_clustermgmt_ip>>).

To access the command-line interface, connect to the cluster management IP address (for example, `ssh admin@<<var_clustermgmt_ip>>`).

Note: The node management interface can be on the same subnet as the cluster management interface or it can be on a different subnet. In this document, we assume that the node management interface is on the same subnet.

7.6 Log in to the Cluster

To log in to the cluster, complete the following steps:

1. Open an SSH connection to either the cluster IP or host name.
2. Log in to the admin user with the password provided previously.

7.7 Zero All Spare Disks

1. To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Note: Disk autoassign should have assigned half of the connected disks to each node in the HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare disks can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

7.8 Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard Unified Target Adapter 2 (UTA2), complete the following steps:

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin show` command.

```
ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
<<var_node01>>						
	0e	cna	target	-	-	online
<<var_node01>>						
	0f	cna	target	-	-	online
<<var_node01>>						
	0g	cna	target	-	-	online
<<var_node01>>						
	0h	cna	target	-	-	online
<<var_node02>>						
	0e	cna	target	-	-	online
<<var_node02>>						
	0f	cna	target	-	-	online
<<var_node02>>						
	0g	cna	target	-	-	online
<<var_node02>>						
	0h	cna	target	-	-	online

8 entries were displayed.

2. Verify that the Current Mode and Current Type properties for all ports are set properly. All ports should be set to mode `cna`. The port type for all protocols should be set to `target`. Change the port personality by running the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode {fc|cna} -type target
```

Note: The ports must be offline to run this command. To take an adapter offline, run the `fcport adapter modify -node <home node of the port> -adapter <port name> -state down` command. Ports must be converted in pairs (for example, 0e and 0f). After conversion, a reboot is required and the ports must be brought back to the up state.

7.9 Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, complete the following step:

Note: A storage virtual machine (SVM) is referred to as a Vserver (or `vserver`) in the GUI and CLI.

1. Run the following command:

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

7.10 Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Ports used for data services (for example, e0b, e0d, e0e, e0f, e0g, e0h, e0j, e0k, and e0l) should be removed from the default broadcast domain, leaving just the management network ports (e0i and e0M).

1. To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports <<var_node01>>:  
e0e,<<var_node01>>:e0f,<<var_node01>>:e0g,<<var_node01>>:e0h,<<var_node01>>:e0j,<<var_node01>>:e0  
k,<<var_node01>>:e0l,<<var_node02>>:e0e,<<var_node02>>:e0f,<<var_node02>>:e0g,<<var_node02>>:e0h,  
<<var_node02>>:e0j,<<var_node02>>:e0k,<<var_node02>>:e0l  
broadcast-domain show
```

7.11 Set Up Service Processor Network Interface

1. To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_node01>> -address-family IPv4 -enable true -  
dhcp none -ip-address <<var_node01_sp_ip>> -netmask <<var_node01_sp_mask>> -gateway  
<<var_node01_sp_gateway>>  
  
system service-processor network modify -node <<var_node02>> -address-family IPv4 -enable true -  
dhcp none -ip-address <<var_node02_sp_ip>> -netmask <<var_node02_sp_mask>> -gateway  
<<var_node02_sp_gateway>>
```

Note: The service processor IP addresses should be in the same subnet as the node management IP addresses.

7.12 Create Aggregates

An aggregate containing the root volume is created during the ONTAP setup process.

To create data aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <<var_node01>> -diskcount 22  
aggr create -aggregate aggr1_node02 -node <<var_node02>> -diskcount 22
```

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: The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

2. Disable NetApp Snapshot copies for the two data aggregates recently created.

```
node run <<var_node01>> aggr options aggr1_node01 nosnap on
node run <<var_node02>> aggr options aggr1_node02 nosnap on
```

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

```
node run <<var_node01>> snap delete -A -a -f aggr1_node01
node run <<var_node02>> snap delete -A -a -f aggr1_node02
```

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>>
```

7.13 Verify Storage Failover

To confirm that storage failover is enabled, run the following commands for a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Note: Both the nodes `<<var_node01>>` and `<<var_node02>>` must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. 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.

3. Verify the HA status for a two-node cluster.

Note: This step does not apply to clusters with more than two nodes.

```
cluster ha show
```

4. Skip step 5 and continue with step 6 if high availability is already configured.
5. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because doing so causes problems with failover.

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

6. 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>>
```

7.14 Disable Flow Control on UTA2 Ports

NetApp recommends disabling flow control on all of the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

1. Run the following commands to configure node 01:

```
network port modify -node <<var_node01>> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <<var_node02>> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show -fields flowcontrol-admin
```

7.15 Configure NTP

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 <ccyymmddhhmm.ss>
```

Note: The format for the date is
`<[Century][Year][Month][Day][Hour][Minute].[Second]>` (for example,
`201309081735.17`).

3. Configure the Network Time Protocol (NTP) servers for the cluster.

```
cluster time-service ntp server create -server <<var_switch_a_ntp_ip>>
cluster time-service ntp server create -server <<var_switch_b_ntp_ip>>
```

7.16 Configure SNMP

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

1. Configure basic SNMP 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>>
```

Configure SNMPv1 Access

To configure SNMPv1 access, set the shared secret plain-text password (called a community).

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

Create SNMPv3 User

SNMPv3 requires that a user be defined and configured for authentication. To create and configure a user for SNMPv3, complete the following steps:

1. Create a user called `snmpv3user`.

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

2. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
3. Run the `security snmpusers` command to view the engine ID.
4. When prompted, enter an eight-character minimum-length password for the authentication protocol.

5. Select `des` as the privacy protocol.
6. When prompted, enter an eight-character minimum-length password for the privacy protocol.

7.17 Configure AutoSupport

The NetApp AutoSupport® tool sends support summary information to NetApp through HTTPS.

1. To configure AutoSupport, run the following command:

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

7.18 Enable Cisco Discovery Protocol

1. To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command to enable CDP on ONTAP:

```
node run -node * options cdpd.enable on
```

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

7.19 Create Interface Groups

1. To create the LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0g

ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0g

ifgrp show
```

7.20 Create VLANs

1. To create VLANs, create NFS VLAN ports and add them to the NFS broadcast domain.

```
network port modify -node <<var_node01>> -port a0a -mtu 9000
network port modify -node <<var_node02>> -port a0a -mtu 9000
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_ib-mgmt_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_ib-mgmt_vlan_id>>

broadcast-domain add-ports -broadcast-domain Default -ports <<var_node01>>:a0a-<<var_ib-mgmt_vlan_id>>, <<var_node02>>:a0a-<<var_ib-mgmt_vlan_id>>
```

7.21 Create Infrastructure Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping only `fc`.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,nfs,ndmp
```

3. Add the two data aggregates to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the FCP protocol in the infra-SVM.

```
fcpx create -vserver Infra-SVM
```

7.22 Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

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

```
volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path //Infra-SVM/rootvol -destination-path //Infra-SVM/rootvol_m01 -
type LS -schedule 15min
snapmirror create -source-path //Infra-SVM/rootvol -destination-path //Infra-SVM/rootvol_m02 -
type LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //Infra-SVM/rootvol
snapmirror show
```

7.23 Configure HTTPS Access

To configure secure access to the storage controller, complete the following steps:

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

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

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <<serial_number>>) by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the four default certificates and replace them with either self-signed certificates or certificates from a Certificate Authority (CA). To delete the default certificates, run the following command.

Use Tab Completion to select and delete each default certificate.

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server
-serial <<serial_number>>
```

Note: Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates.

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Use Tab Completion to aid in the completion of these commands.

```
security certificate create -common-name <<var_cert_common_name>> -type
server -size 2048 -country <<var_cert_country>> -state <<var_cert_state>> -
locality <<var_cert_locality>> -organization <<var_cert_org>> -unit
<<var_cert_unit>> -email-addr <<var_cert_email>> -expire-days
<<var_cert_days>> -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

5. To obtain the values for the parameters required in step 6 (<<var_cert_ca>> and <<var_cert_serial>>), run the `security certificate show` command.

6. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use Tab Completion to aid in the completion of these commands.

```
security ssl modify -vserver <<var_clustername>> -server-enabled true -client-enabled false -ca <<var_cert_ca>> -serial <<var_cert_serial>> -common-name <<var_cert_common_name>>
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -ssl3-enabled true
Warning: Modifying the cluster configuration will cause pending web service requests to be
interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -vserver <<var_clustername>>
```

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

8. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name compat -vserver * -enabled true
```

7.24 Create FlexVol Volumes

To create a NetApp FlexVol® volume, complete the following steps:

1. Collect the following information:
 - The volume name
 - The volume size
 - The aggregate on which the volume exists
2. To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node02 -size 500GB -
state online -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_node01 -size 100GB -state
online -space-guarantee none -percent-snapshot-space 0 -snapshot-policy none

volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size 100GB -state
online -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path //Infra-SVM/rootvol
```

7.25 Create Boot LUNs

1. To create boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-01 -size 15GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-02 -size 15GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-03 -size 15GB -ostype vmware -
space-reserve disabled
```

7.26 Create Infrastructure LUNs

1. To create LUNs for infrastructure VMs, run the following commands:

```
lun create -vserver Infra-SVM -volume infra_datastore_1 -lun infra_datastore_1 -size 500GB -
ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume infra_swap -lun infra_swap -size 100GB -ostype vmware -
space-reserve disabled
```

Note: Volumes and LUNs created here are used for booting the ESXi servers and hosting infrastructure virtual machines such as Virtual Center, Virtual Storage Console, and so on. Additional volumes and LUNs for production VMs and databases are created later in the deployment by using the Virtual Storage Console.

7.27 Schedule Deduplication

Deduplication is enabled by default on NetApp All Flash FAS systems. To schedule deduplication, complete the following steps:

1. After the volumes are created, assign a once-a-day dedupe schedule to `esxi_boot`:

```
efficiency modify -vserver Infra-SVM -volume esxi_boot -schedule sun-sat@0
```

2. Create the AlwaysOn Deduplication efficiency policy:

```
cron create -name lmin -minute  
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,35  
,36,37,38,39,40,41,42,43,44,45,46,47,48,49,50,51,52,53,54,55,56,57,58,59  
efficiency policy create -vserver Infra-SVM -policy Always_On_Deduplication -type scheduled -  
schedule lmin -qos-policy background -enabled true
```

3. Optionally, assign the AlwaysOn Deduplication policy to `infra_datastore_1`:

```
efficiency modify -vserver Infra-SVM -volume infr_datastore_1 -policy Always_On_Deduplication
```

Note: If you do not want to assign an AlwaysOn Deduplication policy to `infra_datastore_1`, assign the once-a-day deduplication schedule:

7.28 Create FCP LIFs

1. To create four FCP LIFs (two on each node) for each attached fabric interconnect, run the following commands:

```
network interface create -vserver Infra-SVM -lif infra_fcp01a -role data -data-protocol fcp -  
home-node <<var_node01>> -home-port 0f -status-admin up  
  
network interface create -vserver Infra-SVM -lif infra_fcp01b -role data -data-protocol fcp -  
home-node <<var_node01>> -home-port 0h -status-admin up  
  
network interface create -vserver Infra-SVM -lif infra_fcp02a -role data -data-protocol fcp -  
home-node <<var_node02>> -home-port 0e -status-admin up  
  
network interface create -vserver Infra-SVM -lif infra_fcp02b -role data -data-protocol fcp -  
home-node <<var_node02>> -home-port 0h -status-admin up  
network interface show
```

7.29 Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, complete the following steps:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data -data-protocol none -home-node  
<<var_node02>> -home-port e0i -address <<var_infra_svm_mgmt_ip>> -netmask  
<<var_infra_svm_mgmt_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-  
policy mgmt -auto-revert true
```

Note: The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway  
<<var_infra_svm_mgmt_gateway>>
```

```
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>

security login unlock -username vsadmin -vserver Infra-SVM
```

7.30 Create Production Storage Virtual Machine

To create an SVM for production workloads, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver prod-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-
security-style unix
```

2. Select the SVM data protocols to configure, keeping `fc` and `nfs`.

```
vserver remove-protocols -vserver Prod-SVM -protocols cifs,nfs,ndmp
```

3. Add the two data aggregates to the Prod-SVM aggregate list for NetApp VSC.

```
vserver modify -vserver Prod-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the FCP protocol in the Prod-SVM.

```
fc create -vserver Prod-SVM
```

7.31 Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

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

```
volume create -vserver Prod-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Prod-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path //Prod-SVM/rootvol -destination-path //Prod-SVM/rootvol_m01 -type
LS -schedule 15min
snapmirror create -source-path //Prod-SVM/rootvol -destination-path //Prod-SVM/rootvol_m02 -type
LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //Prod-SVM/rootvol
snapmirror show
```

7.32 Configure HTTPS Access

To configure secure access to the storage controller, complete the following steps:

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

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

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <<serial_number>>) by running the following command:

```
security certificate show
```

For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the four default certificates and replace them with either self-signed certificates or certificates from a Certificate Authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Prod-SVM -common-name Prod-SVM -ca Prod-SVM -type server -
serial <<serial_number>>
```

Note: Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use Tab Completion to select and delete each default certificate.

3. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Prod-SVM and the cluster SVM. Use Tab Completion to aid in the completion of these commands.

```
security certificate create -common-name <<var_cert_common_name>> -type server -size 2048 -
country <<var_cert_country>> -state <<var_cert_state>> -locality <<var_cert_locality>> -
organization <<var_cert_org>> -unit <<var_cert_unit>> -email-addr <<var_cert_email>> -expire-days
<<var_cert_days>> -protocol SSL -hash-function SHA256 -vserver Prod-SVM
```

4. To obtain the values for the parameters required in step 5 (<<var_cert_ca>> and <<var_cert_serial>>), run the `security certificate show` command.
5. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use Tab Completion to aid in the completion of these commands.

```
security ssl modify -vserver <<var_clustername>> -server-enabled true -client-enabled false -ca
<<var_cert_ca>> -serial <<var_cert_serial>> -common-name <<var_cert_common_name>>
```

6. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web service requests to be
interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -vserver <<var_clustername>>
```

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

7. Revert to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name compat -vserver * -enabled true
```

7.33 Create FCP LIFs

1. Run the following commands to create four FCP LIFs (two on each node) for each attached fabric interconnect:

```
network interface create -vserver Prod-SVM -lif prod_fcp01a -role data -data-protocol fcp -home-
node <<var_node01>> -home-port 0f -status-admin up

network interface create -vserver Prod-SVM -lif prod_fcp01b -role data -data-protocol fcp -home-
node <<var_node01>> -home-port 0h -status-admin up

network interface create -vserver Prod-SVM -lif prod_fcp02a -role data -data-protocol fcp -home-
node <<var_node02>> -home-port 0e -status-admin up

network interface create -vserver Prod-SVM -lif prod_fcp02b -role data -data-protocol fcp -home-
node <<var_node02>> -home-port 0h -status-admin up
network interface show
```

7.34 Add Production SVM Administrator

To add the production SVM administrator and SVM administration LIF in the out-of-band management network, complete the following steps:

1. Run the following commands:

```
network interface create -vserver Prod-SVM -lif vsmgmt -role data -data-protocol none -home-node
<<var_node02>> -home-port e0i -address <<var_prod_svm_mgmt_ip>> -netmask
<<var_prod_svm_mgmt_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-
policy mgmt -auto-revert true
```

Note: The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Prod-SVM -destination 0.0.0.0/0 -gateway
<<var_prod_svm_mgmt_gateway>>
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Prod-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>

security login unlock -username vsadmin -vserver Prod-SVM
```

8 Server Configuration

8.1 Cisco UCS Base Configuration

This FlexPod deployment shows configuration steps for the Cisco UCS 6248UP Fabric Interconnects in a design that supports iSCSI and Fibre Channel direct-attached connectivity to the NetApp AFF. Implementation of these protocols simultaneously should not be considered mandatory and the selection of one or the other should be acceptable depending on your environment and preferences.

Configuration steps are referenced for both fabric interconnects and are called out by the specific model where steps differ.

Note: The following procedure documents the deployment of a single Cisco UCS compute cluster. For this reference architecture, two clusters were deployed to host primary and secondary SQL replicas. After the primary replica cluster is deployed, repeat this procedure to deploy the secondary cluster.

8.2 Cisco UCS Configuration Variables

Table 12 lists the values that are required for configuration of the Cisco UCS servers.

Table 12) Cisco UCS configuration variables

Variable	Value
<<var_password>>	
<<var_ucs_clustername>>	
<<var_ucsa_mgmt_ip>>	
<<var_ucsa_mgmt_mask>>	
<<var_ucsa_mgmt_gateway>>	
<<var_ucs_cluster_ip>>	

Variable	Value
<<var_nameserver_ip>>	
<<var_dns_domain_name>>	
<<var_ucsb_mgmt_ip>>	
<<var_vm_host_infra_01_wwpn1>>	
<<var_vm_host_infra_01_wwpn2>>	
<<var_vm_host_prod_01_wwpn1>>	
<<var_vm_host_prod_01_wwpn2>>	
<<var_vm_host_prod_02_wwpn1>>	
<<var_vm_host_prod_02_wwpn2>>	
<<var_vm_host_prod_03_wwpn1>>	
<<var_vm_host_prod_03_wwpn2>>	

8.3 Perform Initial Setup of Cisco UCS 6248UP Fabric Interconnects 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 B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS 6248UP 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 6248UP fabric interconnect.

```

Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have chosen to setup 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
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

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

Cisco UCS 6248UP 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 6248UP 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. 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.
3. Repeat the previous steps for the 6248 fabric interconnects.

8.4 Cisco UCS Setup

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 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 2.2(5d)

This document assumes the use of Cisco UCS 2.2(5d). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 2.2(5d), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Set Anonymous Reporting

To set up anonymous reporting, complete the following step:

1. In the Anonymous Reporting page, select whether to send anonymous data to Cisco for improving its future products.

Add Block of IP Addresses for KVM Access

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.


The screenshot shows a web-based dialog box titled "Create Block of IPv4 Addresses". It contains several input fields for configuring an IP address block. The "From" field is set to "192.168.156.101" and the "Size" field is set to "12". The "Subnet Mask" field is set to "255.255.255.0" and the "Default Gateway" field is set to "192.168.156.1". The "Primary DNS" field is set to "0.0.0.0" and the "Secondary DNS" field is set to "0.0.0.0". At the bottom right of the dialog, there are "OK" and "Cancel" buttons.

5. Click OK to create the 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_switch_a_ntp_ip>>` and click OK.



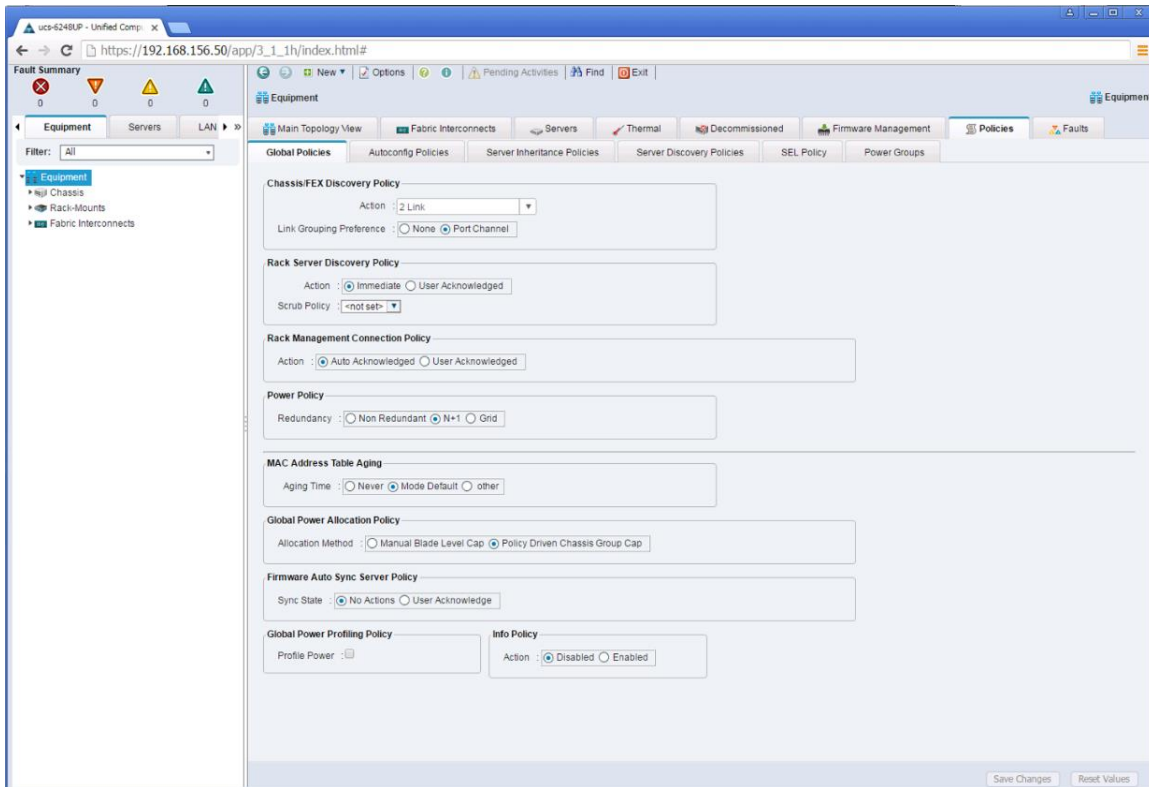
7. Click Add NTP Server.
8. Enter `<<var_switch_b_ntp_ip>>` and click OK.
9. Click OK.

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 from 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 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 FC Switching

To use direct-attached Fibre Channel or FCoE connectivity, the fabric interconnects need to be placed in Fibre Channel switching mode by completing the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the pane.
2. Expand Fabric Interconnects and select either fabric interconnect.

Note: Step 3 reboots both Cisco UCS Fabric Interconnects. If any servers run on this system, shut them down before you execute this step.

3. In the Actions pane, select Set FC Switching Mode. Click Yes and then OK.



4. After the fabric interconnects have rebooted, log back into Cisco UCS Manager.
5. Expand Fabric Interconnects on the Equipment tab.
6. Select each fabric interconnect and verify under Status that the FC mode is now Switch.

Enable FCoE Storage Ports

To enable FCoE 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 FCoE ports on the storage controllers (ports 25 and 26), right-click them, and select Configure as FCoE Storage Port. Click Yes to confirm.
5. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
6. Expand Ethernet Ports.
7. Select ports that are connected to the FCoE ports on the storage controllers (ports 25 and 26), right-click them, and select Configure as FCoE Storage Port. Click Yes to confirm.

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.

- Expand Ethernet Ports.
- Select the ports that are connected to the chassis (ports 1–4), right-click them, and select Configure as Server Port.
- Click Yes to confirm the server ports and click OK.
- Verify that the ports connected to the chassis are now configured as server ports.
- Select ports 17 and 18 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

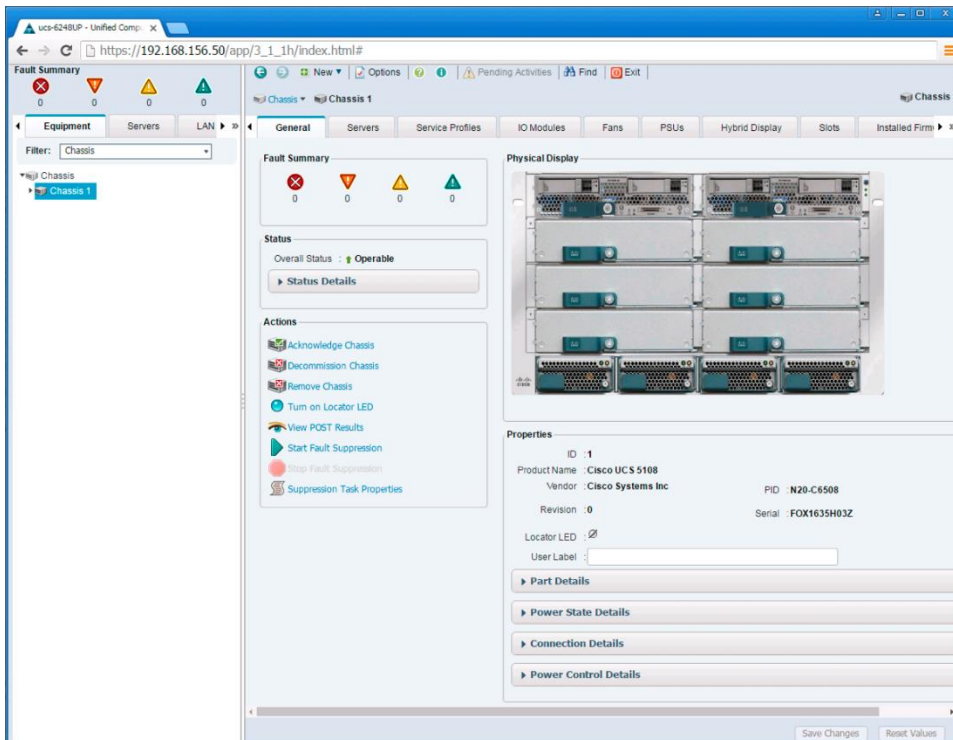


- Click Yes to confirm uplink ports and click OK.
- Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
- Expand Ethernet Ports.
- Select the ports that are connected to the chassis (ports 1–4), right-click them, and select Configure as Server Port.
- Click Yes to confirm server ports and click OK.
- Select ports 17 and 18 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
- Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis, complete the following steps:

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



4. Click Yes and then click OK to complete acknowledging the chassis.

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

15. Click Next.
16. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 17
 - Slot ID 1 and port 18
17. Click >> to add the ports to the port channel.
18. Click Finish to create the port channel.
19. Click OK.

Create a WWNN Pool for FCoE Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps in Cisco UCS Manager.

1. Select the SAN tab on the left.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization and select Create WWNN Pool to create the WWNN pool.
4. Enter `WWNN_Pool` for the name of the WWNN pool.
5. Optional: Enter a description for the WWNN pool.
6. Select Sequential for Assignment Order.
7. Click Next.
8. Click Add.
9. Modify the From field as necessary for the UCS environment.

Note: Modifications of the WWN block and the WWPN and MAC addresses can convey identifying information for the Cisco UCS domain. Within the From field in our example, the sixth octet was changed from 00 to 91 to represent identifying information for this being in building 1 on the first floor. The seventh octet was changed from 00 to 10 to represent our first UCS domain.

Note: Also, when having multiple adjacent Cisco UCS domains, it is important that these blocks, the WWN, WWPN, and MAC, hold differing values between each set.

10. Specify a size of the WWNN block sufficient to support the available server resources.

11. Click OK.

Create WWPN Pools

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.

2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.
4. Right-click WWPN Pools under the root organization and select Create WWPN Pool to create the WWPN pool.
5. Enter `WWPN_Pool_A` as the name of the WWPN pool.
6. Optional: Enter a description for the WWPN pool.
7. Select Sequential for Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting WWPN.

Note: For the FlexPod solution, NetApp recommends placing `0A` in the next-to-last octet of the starting WWPN to identify the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:11:1A:00`.

11. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources.
12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click WWPN Pools under the root organization and select Create WWPN Pool to create the WWPN pool.
16. Enter `WWPN_Pool_B` as the name of the WWPN pool.
17. Optional: Enter a description for the WWPN pool.
18. Select Sequential for Assignment Order.
19. Click Next.
20. Click Add.
21. Specify a starting WWPN.

Note: For the FlexPod solution, NetApp recommends placing `0B` in the next-to-last octet of the starting WWPN to identify the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:11:1AB:00`.

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

Create VSANs

To configure the necessary virtual SANs (VSANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.

Note: In this procedure, two VSANs are created.

2. Select SAN > SAN Cloud.
3. Right-click VSANs and select Create VSAN.
4. Enter `VSAN_A` as the name of the VSAN to be used for fabric A.
5. Select Enabled for FC Zoning.

6. Select Fabric A.
7. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. NetApp recommends using the same ID for both parameters and using something other than 1.

The screenshot shows a 'Create VSAN' dialog box with the following fields and options:

- Name :** VSAN_A
- FC Zoning Settings**
 - FC Zoning : ☐ Disabled ☒ Enabled
 - Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.
- ☐ Common/Global ☒ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently
- Instructions:**
 - You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.
 - Enter the VSAN ID that maps to this VSAN.
 - A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
 - Enter the VLAN ID that maps to this VSAN.
- VSAN ID :** 101
- FCoE VLAN :** 101
- Buttons:** OK, Cancel

8. Click OK and then click OK again.
9. Under SAN Cloud, right-click VSANs and select Create VSAN.
10. Enter VSAN_B as the name of the VSAN to be used for fabric B.
11. Select Enabled for FC Zoning.
12. Select Fabric B.
13. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. NetApp recommends using the same ID for both parameters and using something other than 1.

Create VSAN

Name : VSAN_B

FC Zoning Settings

FC Zoning : ☐ Disabled ☒ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

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

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

VSAN ID : 102

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

FCoE VLAN : 102

OK Cancel

14. Click OK and then click OK again.
15. Under Storage Cloud, right-click VSANs and select Create Storage VSAN.
16. Enter `VSAN_A` as the name of the VSAN to be used for fabric A.
17. Select Enabled for FC Zoning.
18. Select Fabric A.
19. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for fabric A previously.

Create Storage VSAN

Name : VSAN_A

FC Zoning Settings

FC Zoning : ☐ Disabled ☒ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

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

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

VSAN ID : 101

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

FCoE VLAN : 101

OK Cancel

20. Click OK and then click OK again.
21. Under Storage Cloud, right-click VSANs and select Create Storage VSAN.
22. Enter `VSAN_B` as the name of the VSAN to be used for fabric B.
23. Select Enabled for FC Zoning.
24. Select Fabric B.
25. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for fabric B previously.

Create Storage VSAN

Name : VSAN_B

FC Zoning Settings

FC Zoning : ☐ Disabled ☒ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

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

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

VSAN ID : 102

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

FCoE VLAN : 102

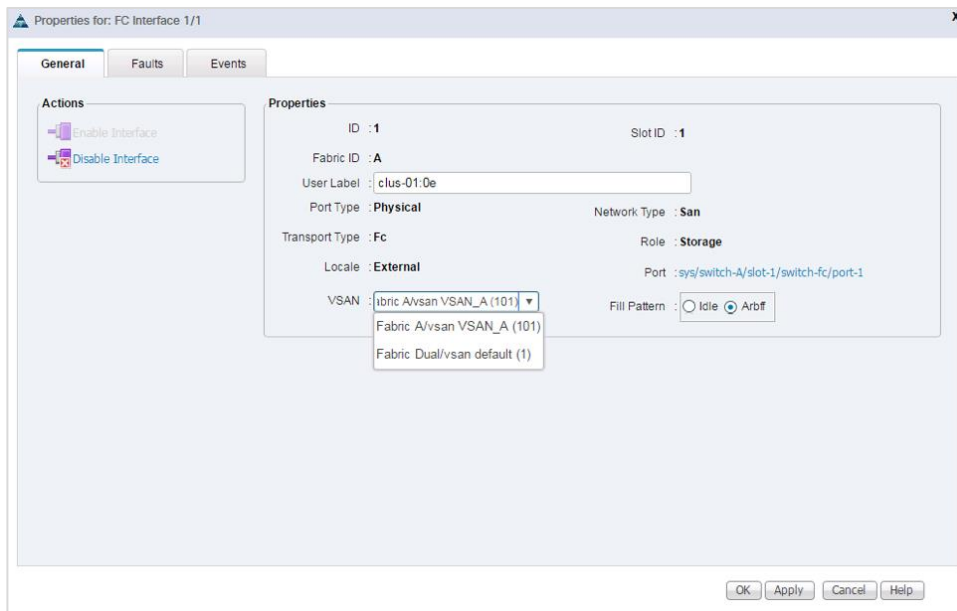
OK Cancel

26. Click OK and then click OK again.


Assign VSANs to FCoE Storage Ports

To assign the necessary VSANs to the FC storage ports for the Cisco UCS environment, complete the following steps:


1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.
2. Select SAN > Storage Cloud.
3. Expand Fabric A and Storage FCoE Interfaces.
4. Right-click FC Interface 1/25 and select Storage FCoE Interface.
5. Set the User Label to the storage controller name and port to which this interface is connected.
6. Select `VSAN_A (101)` as the VSAN.



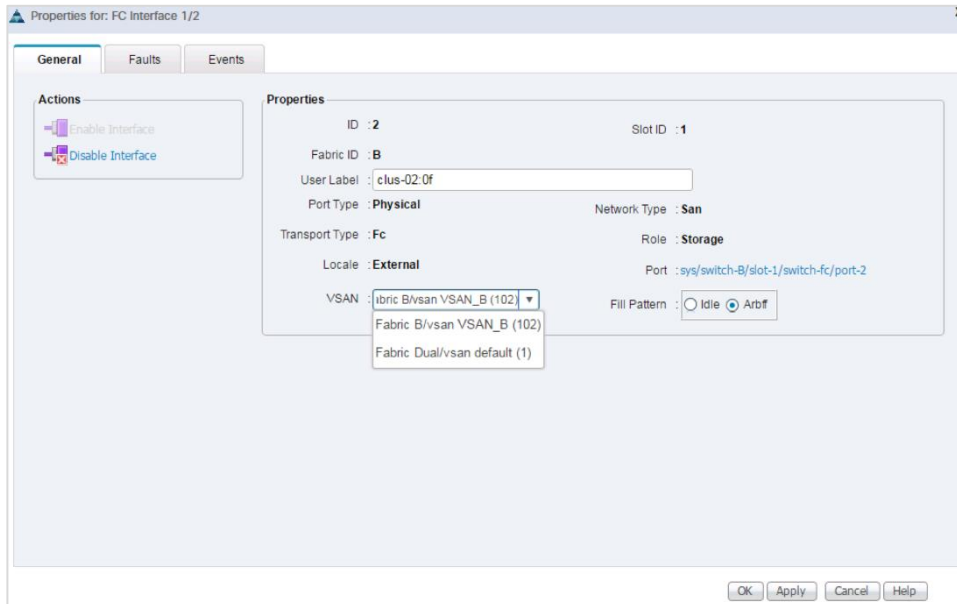
7. Click OK.
8. Expand Fabric A and Storage FCoE Interfaces.
9. Right-click FC Interface 1/26 and select Storage FCoE Interface.
10. Set the User Label to the storage controller name and port to which this interface is connected.
11. Select `VSAN_A (101)` as the VSAN.



12. Click OK.
13. Expand Fabric B and Storage FCoE Interfaces.
14. Right-click FCoE Interface 1/25 and select Storage FCoE Interface.
15. Set the User Label to the storage controller name and port to which this interface is connected.
16. Select `VSAN_B (102)` as the VSAN.



17. Click OK.
18. Expand Fabric B and Storage FCoE Interfaces.
19. Right-click FCoE Interface 1/26 and select Storage FCoE Interface.
20. Set the User Label to the storage controller name and port to which this interface is connected.
21. Select `VSAN_B (102)` as the VSAN.



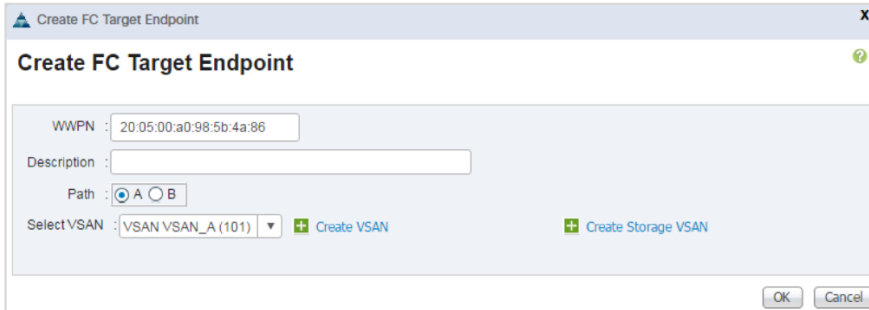
22. Click OK.

Create Storage Connection Policies for FC Zoning

To create storage connection policies for the fc zoning, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.

2. Select Policies > root.
3. Right-click Storage Connection Policies and select Create Storage Connection Policy.
4. Enter `Fabric-A` as the name of the policy.
5. Select the Single Initiator Multiple Targets Zoning Type.
6. Click the plus sign on the right to add a zoning target.
7. Enter the WWPN for `infra_fcp01a` from the storage cluster. You can obtain this WWPN by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.
8. Select Path A and VSAN_A.



Create FC Target Endpoint

WWPN : 20:05:00:a0:98:5b:4a:86

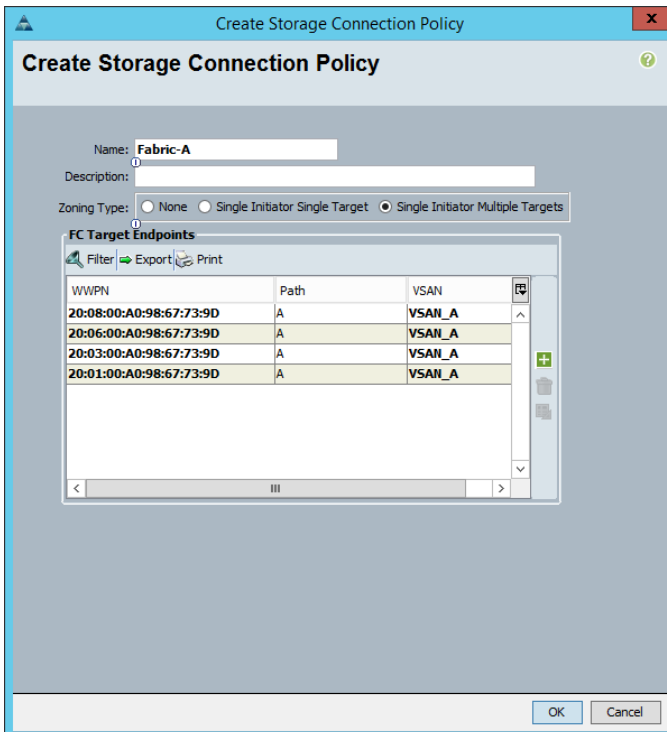
Description :

Path : ☒ A ☐ B

Select VSAN : VSAN_VSAN_A (101) + Create VSAN + Create Storage VSAN

OK Cancel

9. Click OK.
10. Repeat steps 7 through 9 for all fabric A target endpoints:
 - `Infra_fcp01a`
 - `Infra_fcp02a`
 - `Prod_fcp01a`
 - `Prod_fcp02a`



Create Storage Connection Policy

Name: `Fabric-A`

Description:

Zoning Type: ☐ None ☐ Single Initiator Single Target ☒ Single Initiator Multiple Targets

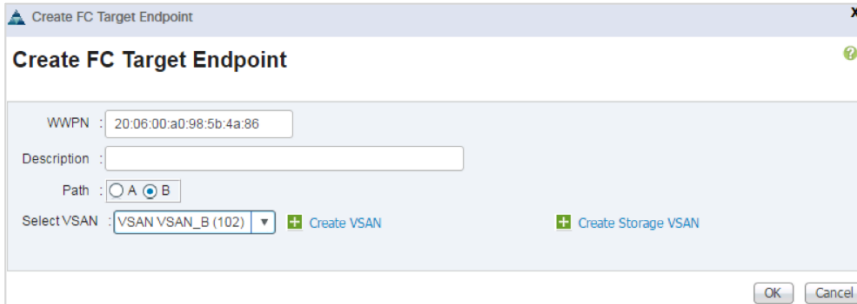
FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:08:00:A0:98:67:73:9D	A	VSAN_A
20:06:00:A0:98:67:73:9D	A	VSAN_A
20:03:00:A0:98:67:73:9D	A	VSAN_A
20:01:00:A0:98:67:73:9D	A	VSAN_A

OK Cancel

11. Click OK and then click OK again.
12. Right-click Storage Connection Policies and select Create Storage Connection Policy.
13. Enter `Fabric-B` as the name of the policy.
14. Select the Single Initiator Multiple Targets Zoning Type.
15. Click the plus sign on the right to add a zoning target.
16. Enter the WWPN `infra_fcp01b` from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.
17. Select Path B and VSAN_B.



Create FC Target Endpoint

WWPN : 20:06:00:a0:98:5b:4a:86

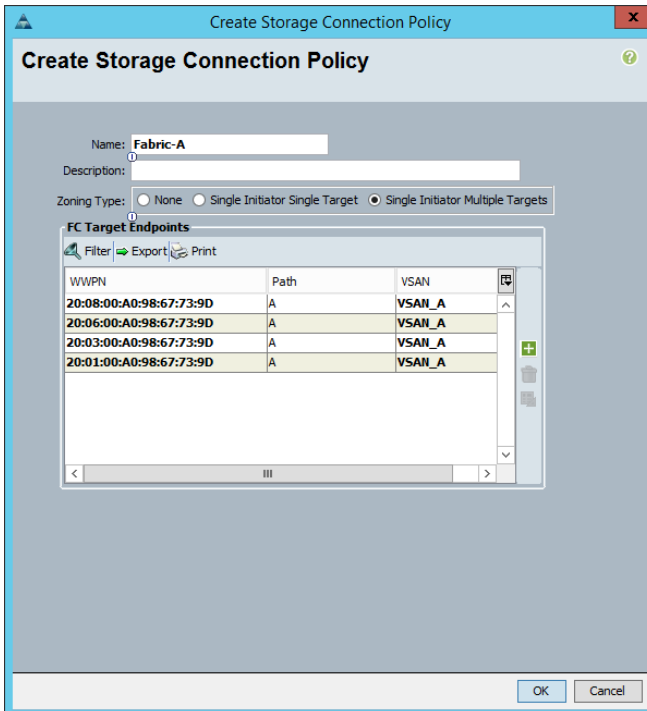
Description :

Path : ☐ A ☒ B

Select VSAN : VSAN_VSAN_B (102) + Create VSAN + Create Storage VSAN

OK Cancel

18. Click OK.
19. Repeat this process for all fabric A WWPN targets:
 - `infra_fcp01a`
 - `infra_fcp02a`
 - `prod_fcp01a`
 - `prod_fcp02a`



Create Storage Connection Policy

Name: `Fabric-A`

Description :

Zoning Type: ☐ None ☐ Single Initiator Single Target ☒ Single Initiator Multiple Targets

FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:08:00:A0:98:67:73:9D	A	VSAN_A
20:06:00:A0:98:67:73:9D	A	VSAN_A
20:03:00:A0:98:67:73:9D	A	VSAN_A
20:01:00:A0:98:67:73:9D	A	VSAN_A

OK Cancel

20. Click OK and then click OK again.

Create vHBA Templates

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates and select Create vHBA Template.
4. Enter `vHBA_Template_A` as the vHBA template name.
5. Keep Fabric A selected.
6. Select `VSAN_A`.
7. Leave Initial Template as the template type.
8. Select `WWPN_Pool_A` as the WWPN pool.
9. Click OK to create the vHBA template.
10. Click OK.

The screenshot shows the 'Create vHBA Template' dialog box. The fields are filled as follows:

- Name: `vHBA_Template_A`
- Description: (empty)
- Fabric ID: ☒ A ☐ B
- Select VSAN: `VSAN_A` (with a '+ Create VSAN' link)
- Template Type: ☒ Initial Template ☐ Updating Template
- Max Data Field Size: `2048`
- WWPN Pool: `WWPN_Pool_A(32/32)`
- QoS Policy: `<not set>`
- Pin Group: `<not set>`
- Stats Threshold Policy: `default`

Buttons: OK, Cancel

11. Right-click vHBA Templates and select Create vHBA Template.
12. Enter `vHBA_Template_B` as the vHBA template name.
13. Select Fabric B as the fabric ID.
14. Select `VSAN_B`.
15. Leave Initial Template as the template type.
16. Select `WWPN_Pool_B` as the WWPN pool.
17. Click OK to create the vHBA template.
18. Click OK.

Create SAN Connectivity Policy

To configure the necessary Infrastructure SAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the Navigation pane.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies and select Create SAN Connectivity Policy.
4. Enter Infra-SAN-Policy as the name of the policy.
5. Select the previously created WWNN_Pool for the WWNN assignment.
6. Click the Add button at the bottom to add a vHBA.
7. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
8. Select the Use vHBA Template checkbox.
9. From the vHBA Template list, select vHBA_Template_A.
10. From the Adapter Policy list, select VMWare.
11. Click OK.
12. Click the Add button at the bottom to add a second vHBA.
13. In the Create vHBA dialog box, enter Fabric-B as the name of the vHBA.
14. Select the Use vHBA Template checkbox.
15. From the vHBA Template list, select vHBA_Template_B.
16. From the Adapter Policy list, select VMWare.
17. Click OK.

Create SAN Connectivity Policy

Name :

Description :

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

World Wide Node Name

WWNN Assignment:

[+ Create WWNN Pool](#)

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

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

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

18. Click OK to create the SAN Connectivity Policy.
19. Click OK to confirm creation.
20. Right-click the Infra-SAN-Policy and select Create vHBA Initiator Group.
21. Enter Fabric-A for the name and select the Fabric-A vHBA.
22. Select the Fabric-A Storage Connection Policy and click OK.
23. Right-click the Infra-SAN-Policy and select Create vHBA Initiator Group.
24. Enter Fabric-B for the name and select the Fabric-B vHBA.
25. Select the Fabric-B Storage Connection Policy and click OK.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete 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 and select Create MAC Pool to create the MAC address pool.
4. Enter `MAC_Pool_A` as the name of the MAC pool.
5. Optional: Enter a description for the MAC pool.
6. Select Sequential as the option for Assignment Order.
7. Click Next.
8. Click Add.

9. Specify a starting MAC address.

Note: For the FlexPod solution, NetApp recommends placing 0A in the next-to-last octet of the starting MAC address to identify the MAC addresses as fabric A addresses. In our example, we have carried forward in our example of also embedding the extra building, floor, and Cisco UCS domain number information, giving us 00:25:B5:91:1A:00 as our first MAC address.

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 and select Create MAC Pool to create the MAC address pool.
15. Enter MAC_Pool_B as the name of the MAC pool.
16. Optional: Enter a description for the MAC pool.
17. Click Next.
18. Click Add.
19. Specify a starting MAC address.

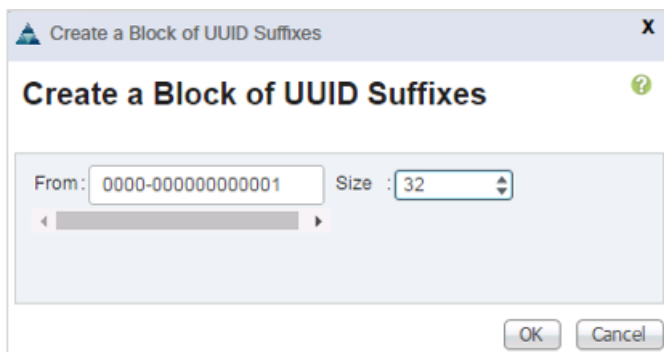
Note: For the FlexPod solution, NetApp recommends placing 0B in the next-to-last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the extra building, floor, and Cisco UCS domain number information, giving us 00:25:B5:91:1B:00 as our first MAC address.

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

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 and select Create UUID Suffix Pool.
4. Enter UUID_Pool as the name of the UUID suffix pool.
5. Optional: Enter a description for the UUID suffix pool.
6. Keep the prefix at the derived option.
7. Select Sequential for Assignment Order.
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.

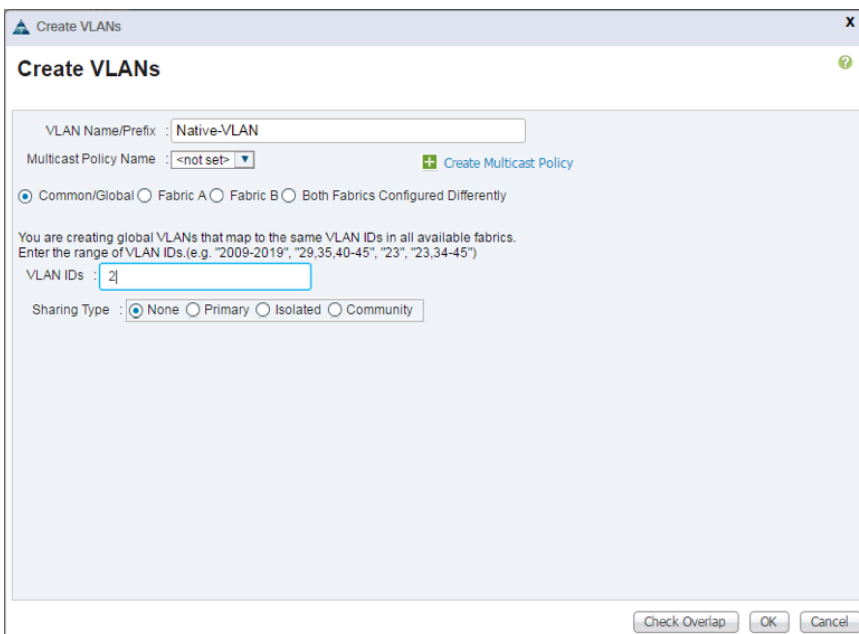
Create VLANs

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, four unique VLANs are created. See Table 3 for a list of VLANs required for this solution.

2. Select LAN > LAN Cloud.
3. Right-click VLANs and select Create VLANs.
4. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.
5. Keep the Common/Global option selected for the scope of the VLAN.
6. Enter the native VLAN ID.
7. Keep Sharing Type as None.
8. Click OK and then click OK again.



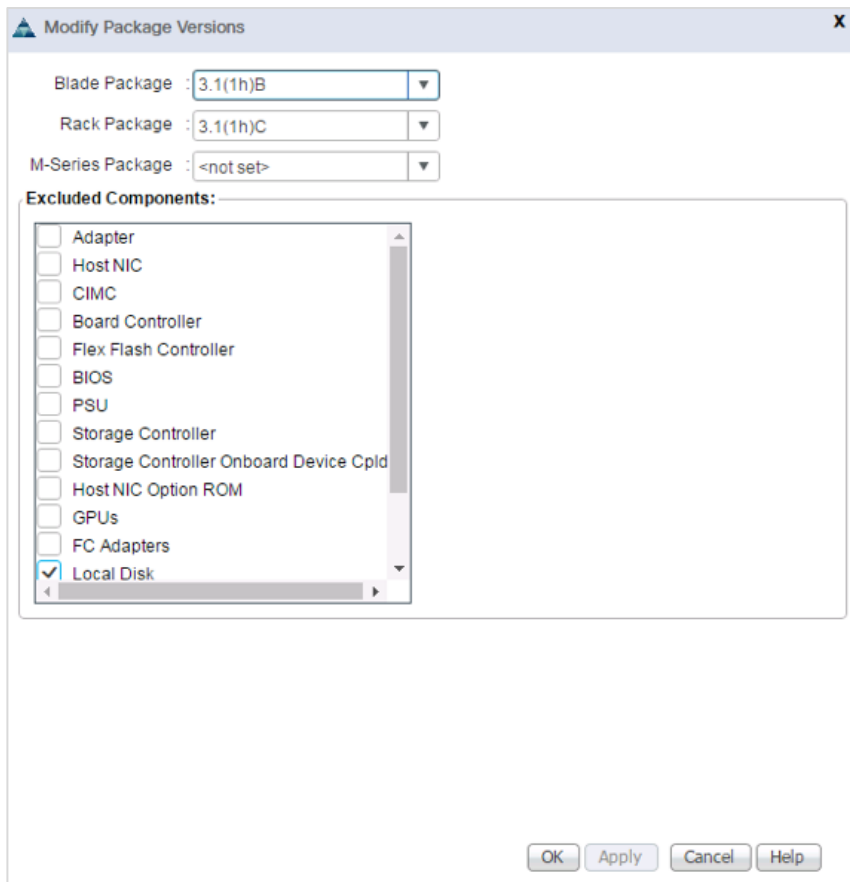
9. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN`, and select `Set as Native VLAN`.
10. Click `Yes` and then click `OK`.
11. Right-click VLANs and select `Create VLANs`.
12. Enter `IB-Mgmt` as the name of the VLAN to use for management traffic.
13. Keep the `Common/Global` option selected for the scope of the VLAN.
14. Enter the in-band management VLAN ID.
15. Keep `Sharing Type` as `None`.
16. Click `OK` and then click `OK` again.
17. Right-click VLANs and select `Create VLANs`.
18. Enter `Infra-NFS` as the name of the VLAN to use for NFS.
19. Keep the `Common/Global` option selected for the scope of the VLAN.
20. Enter the NFS VLAN ID.
21. Keep `Sharing Type` as `None`.
22. Click `OK` and then click `OK` again.
23. Right-click VLANs and select `Create VLANs`.
24. Enter `vMotion` as the name of the VLAN to be used for vMotion.
25. Keep the `Common/Global` option selected for the scope of the VLAN.
26. Enter the vMotion VLAN ID.
27. Keep `Sharing Type` as `None`.
28. Click `OK` and then click `OK` again.
29. Right-click VLANs and select `Create VLANs`.
30. Enter `VM-Traffic` as the name of the VLAN to be used for VM Traffic.
31. Keep the `Common/Global` option selected for the scope of the VLAN.
32. Enter the VM-Traffic VLAN ID.
33. Keep `Sharing Type` as `None`.
34. Click `OK` and then click `OK` again.

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. Expand `Host Firmware Packages`.
4. Select `Default`.
5. In the `Actions` pane, select `Modify Package Versions`.
6. Select the version `3.1(1h)` for `Blade Package` and for `Rack Package`.
7. Leave `M-Series Package` as `<not set>` and leave `Excluded Components` with only `Local Disk` selected.

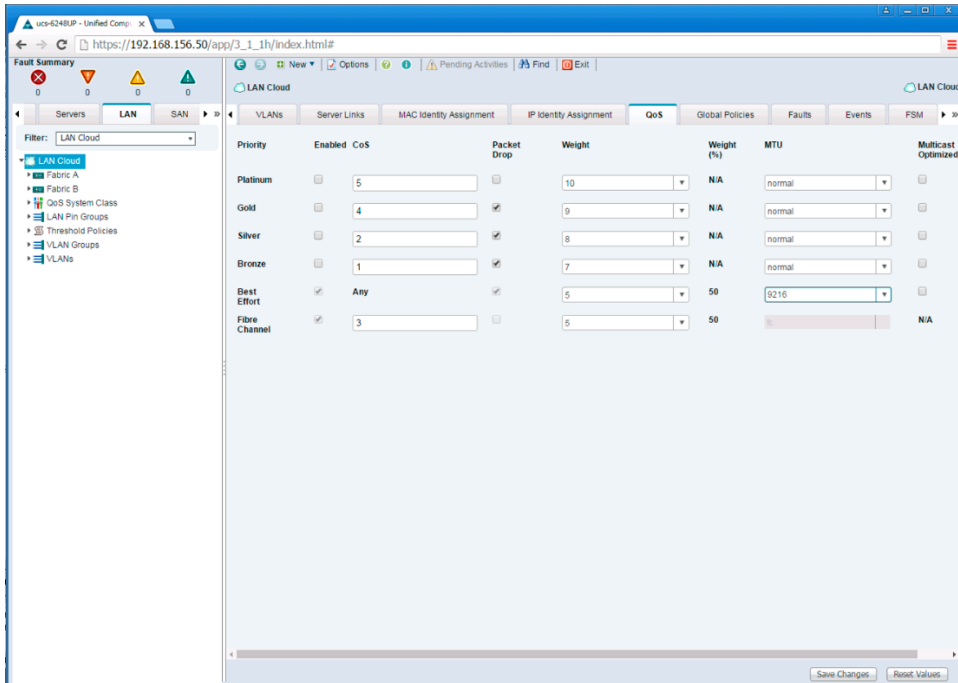


8. Click OK to modify the host firmware package.

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 and select Create Local Disk Configuration Policy.
4. Enter `SAN-Boot` as the local disk configuration policy name.
5. Change the mode to No Local Storage.
6. Click OK to create the local disk configuration policy.

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : ☒ Disable ☐ Enable

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

FlexFlash RAID Reporting State : ☒ Disable ☐ Enable

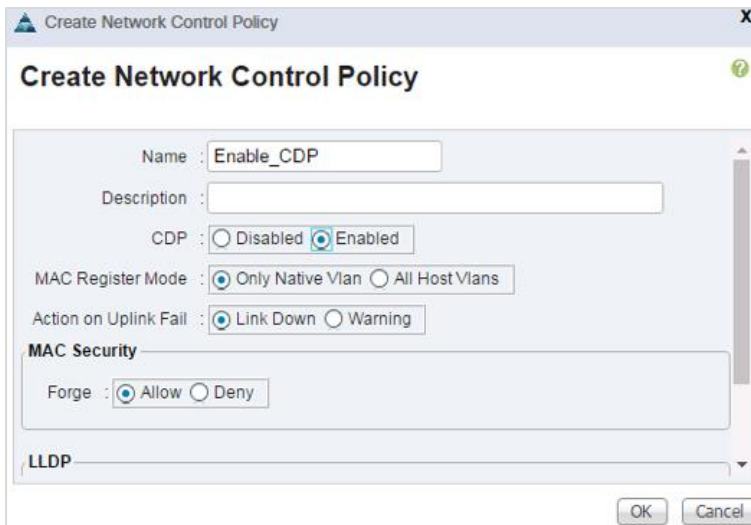
OK Cancel

7. Click OK.

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 and select Create Network Control Policy.
4. Enter `Enable_CDP` as the policy name.
5. For CDP, select the Enabled option.
6. Click OK to create the network control policy.



Create Network Control Policy

Name :

Description :

CDP : ☐ Disabled ☒ Enabled

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

Action on Uplink Fail : ☒ Link Down ☐ Warning

MAC Security

Forge : ☒ Allow ☐ Deny

LLDP

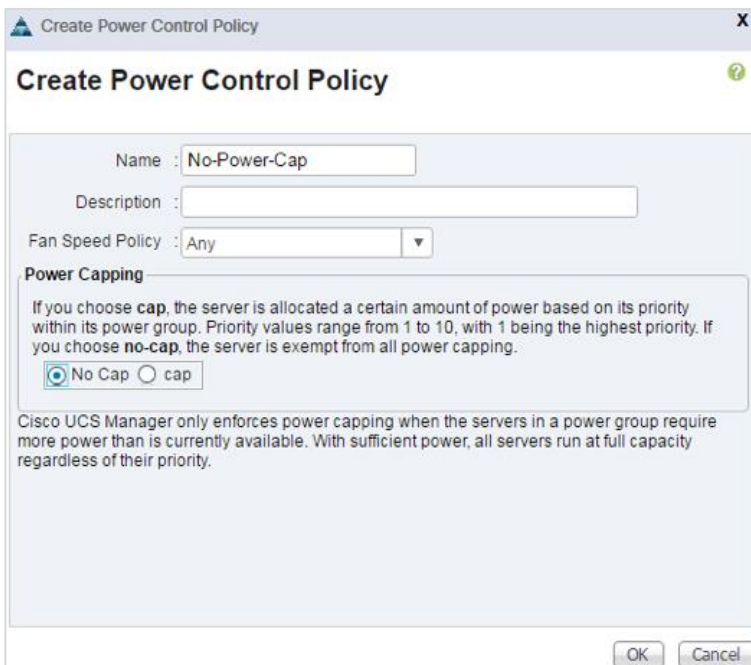
OK Cancel

7. Click OK.

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 and select Create Power Control Policy.
4. Enter `No-Power-Cap` as the power control policy name.
5. Change the power capping setting to No Cap.
6. Click OK to create the power control policy.
7. Click OK.



Create Power Control Policy

Name :

Description :

Fan Speed Policy :

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

☒ No Cap ☐ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

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 and select Create BIOS Policy.
4. Enter `VM-Host-Infra` as the BIOS policy name.
5. Change the Quiet Boot setting to Disabled.
6. Change Consistent Device Naming to Enabled.
7. Click Finish to create the BIOS policy.

The screenshot shows the 'Create BIOS Policy' window in the Unified Computing System Manager. The 'Main' tab is active, displaying the following configuration:

- Name: VM-Host-Infra
- Description: (empty field)
- Reboot on BIOS Settings Change: ☐
- Quiet Boot: ☒ disabled ☐ enabled ☐ Platform Default
- Post Error Pause: ☐ disabled ☐ enabled ☒ Platform Default
- Resume Ac On Power Loss: ☐ stay-off ☐ last-state ☐ reset ☒ Platform Default
- Front Panel Lockout: ☐ disabled ☐ enabled ☒ Platform Default
- Consistent Device Naming: ☐ disabled ☒ enabled ☐ Platform Default

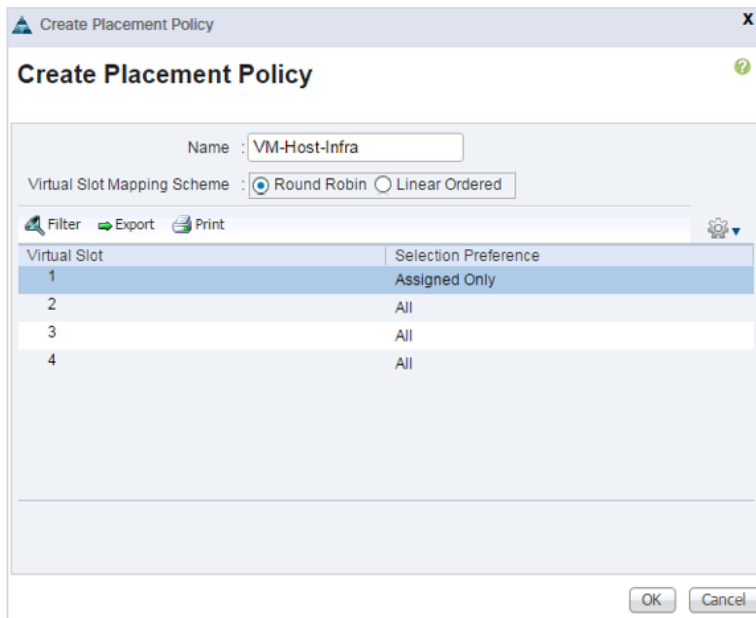
At the bottom, there are navigation buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

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 and select Create Placement Policy.
4. Enter `VM-Host-Infra` as the name of the placement policy.
5. Click 1 and select Assigned Only.
6. Click OK and then click OK again.



Create Placement Policy

Name :

Virtual Slot Mapping Scheme : ☒ Round Robin ☐ Linear Ordered

Filter Export Print

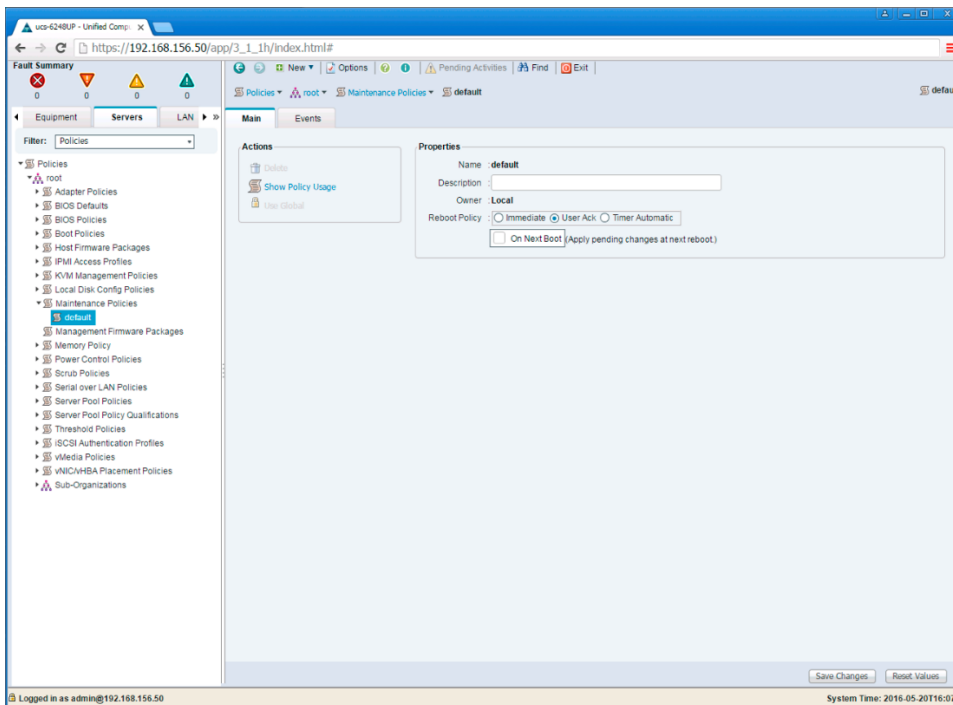
Virtual Slot	Selection Preference
1	Assigned Only
2	All
3	All
4	All

OK Cancel

Update the 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. Optional: Click Next Boot to delegate maintenance windows to server owners.



ucsm4248UP - Unified Comp: x

https://192.168.156.50/app/3_1_1h/index.html#

Fault Summary

0 0 0 0

Equipment Servers LAN

Filter: Policies

- Policies
 - root
 - Adapter Policies
 - BIOS Defaults
 - BIOS Policies
 - Boot Policies
 - Host Firmware Packages
 - IPMI Access Profiles
 - KVM Management Policies
 - Local Disk Config Policies
 - Maintenance Policies
 - default
 - Management Firmware Packages
 - Memory Policy
 - Power Control Policies
 - Scrub Policies
 - Serial over LAN Policies
 - Server Pool Policies
 - Server Pool Policy Qualifications
 - Threshold Policies
 - SCSI Authentication Profiles
 - Video Policies
 - vNIC/HBA Placement Policies
 - Sub-Organizations

Main Events

Actions

- Update
- Show Policy Usage
- Use Global

Properties

Name: default

Description:

Owner: Local

Reboot Policy: ☐ Immediate ☒ User Ack ☐ Timer Automatic

☐ On Next Boot (Apply pending changes at next reboot)

Save Changes Reset Values

Logged in as admin@192.168.156.50 System Time: 2016-05-20T16:07

6. Click Save Changes.
7. 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. A total of four vNIC templates are created.

Note: The same infra-VLANs were used on both the infrastructure (Infra) and production (Prod) hosts deployed in this environment. If the production networks differ in the VLANs used for infrastructure networks, different vNIC templates should be created for each.

Create Data vNICs

1. In Cisco UCS Manager, click the LAN tab in the Navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates and select Create vNIC Template.
4. Enter `vNIC_Template_A` as the vNIC template name.
5. Keep Fabric A selected.
6. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template as the template type.
9. Under VLANs, select the checkboxes for IB-MGMT, Infra-NFS, Native-VLAN, VM-Traffic, and vMotion VLANs.

Create vNIC Template

Name : vNIC_Template_A

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

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-NFS	<input type="radio"/>

OK Cancel

10. Set Native-VLAN as the native VLAN.
11. For MTU, enter 9000.
12. In the MAC Pool list, select MAC_Pool_A.
13. In the Network Control Policy list, select Enable_CDP.

Template Type : ☐ Initial Template ☒ Updating Template

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-NFS	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

Create VLAN

CDN Source : ☒ vNIC Name ☐ User Defined

MTU : 9000

MAC Pool : MAC_Pool_A(32/32)

QoS Policy : <not set>

Network Control Policy : Enable_CDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMO

Dynamic vNIC Connection Policy : <not set>

OK Cancel

14. Click OK to create the vNIC template.

15. Click OK.

Repeat the following equivalent steps for vNIC_Template_B:

1. In the Navigation pane, select the LAN tab.
2. Select Policies > root.
3. Right-click vNIC Templates and select Create vNIC Template.
4. Enter vNIC_Template_B as the vNIC template name.
5. Select Fabric B.
6. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template as the template type.
9. Under VLANs, select the checkboxes for IB-MGMT, INFRA-NFS, Native-VLAN, and vMotion VLANs.

Create vNIC Template

Name : vNIC_Template_B

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

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input type="checkbox"/>	Infra-NFS	<input type="radio"/>

OK Cancel

10. Set default as the native VLAN.
11. Select vNIC Name for the CDN source.
12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_B.
14. In the Network Control Policy list, select Enable_CDP.

Create vNIC Template

Template Type : ☐ Initial Template ☒ Updating Template

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-NFS	<input type="radio"/>
<input type="checkbox"/>	iSC SI-A-VLAN	<input type="radio"/>
<input type="checkbox"/>	iSC SI-B-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

Create VLAN

CDN Source : ☒ vNIC Name ☐ User Defined

MTU : 9000

MAC Pool : MAC_Pool_B(32/32)

QoS Policy : <not set>

Network Control Policy : Enable_CDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel


15. Click OK to create the vNIC template.

16. Click OK.

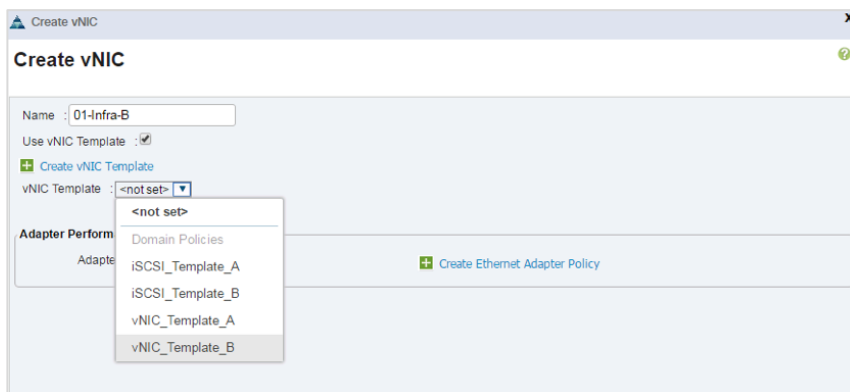
Create LAN Connectivity Policy

To configure the necessary infrastructure LAN connectivity policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the Navigation pane.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies and select Create LAN Connectivity Policy.
4. Enter Infra-LAN-Policy as the name of the policy.
5. Click the upper Add button to add a vNIC.
6. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.
7. Select the Use vNIC Template checkbox.
8. From the vNIC Template list, select vNIC_Template_A.
9. From the Adapter Policy list, select VMWare.
10. Click OK to add this vNIC to the policy.



11. Click the upper Add button to add another vNIC to the policy.
12. In the Create vNIC box, enter vNIC-01-Infra-B as the name of the vNIC.
13. Select the Use vNIC Template checkbox.
14. From the vNIC Template list, select vNIC_Template_B.
15. From the Adapter Policy list, select VMWare.
16. Click OK to add the vNIC to the policy.



Create vMedia Policy for VMware ESXi 6.0 U1b Install Boot

When setting up Data ONTAP, an HTTP web server is required for hosting NetApp Data ONTAP as well as the VMware software. The vMedia policy created here will map VMware ESXi 6.0u1b ISO to the Cisco UCS server to boot the ESXi installation. To create this policy, complete the following steps:

1. In Cisco UCS Manager, select the Servers tab.
2. Select Policies > root.
3. Right-click vMedia Policies and select Create vMedia Policy.
4. Name the policy ESXi-6.0U1b-HTTP.
5. Enter Mounts Cisco Custom ISO for ESXi 6.0U1b in the Description field.
6. Click Add.
7. Name the mount ESXi-6.0U1b-HTTP.
8. Select the CDD device type.
9. Select the HTTP protocol.
10. Enter the IP address of the web server.

Note: Because DNS server IPs were not entered into the KVM IP previously, you must enter the IP of the web server instead of the host name pool.

11. Enter Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso as the remote file name.

12. Enter the web server path to the ISO file in the Remote Path field.

The screenshot shows a 'Create vMedia Mount' window. The 'Name' field is filled with 'ESXi-6.0U1b-HTTP'. The 'Device Type' has 'CDD' selected. The 'Protocol' has 'HTTP' selected. The 'Hostname/IP Address' is '192.168.156.150'. The 'Image Name Variable' has 'None' selected. The 'Remote File' is 'Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso'. The 'Remote Path' is '/software/vSphere/'. There are 'OK' and 'Cancel' buttons at the bottom right.

13. Click OK to create the vMedia Mount.

14. Click OK and then click OK again to complete creating the vMedia policy.

Note: You can use the vMedia service profile template to install the ESXi host for any new servers added to the Cisco UCS environment. On first boot, the host will boot into the ESXi installer. After ESXi is installed, the vMedia will not be referenced as long as the boot disk is accessible.

Create Boot Policies (FCoE Boot)

This procedure applies to a Cisco UCS environment in which each controller has two FCoE connections, one to each fabric interconnect. This setup provides a total of four paths and redundant connections between servers and storage.

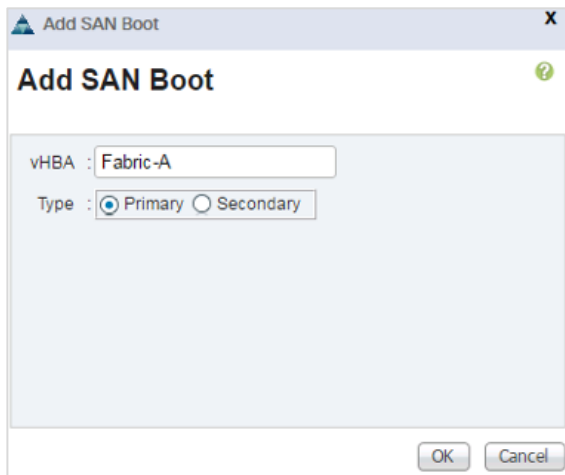
Table 13) Boot LIF configuration.

	6248UP Fabric A	6248UP Fabric B
AFF Cluster Node 1 LIF	infra_fcp01a	infra_fcp01b
AFF Cluster Node 2 LIF	infra_fcp02a	infra_fcp02b

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 and select Create Boot Policy.
4. Enter `Boot-FC-A` as the name of the boot policy.
5. Optional: Enter a description for the boot policy.
Note: Do not select the Reboot on Boot Order Change checkbox.
6. Keep the Reboot on Boot Order Change option cleared.
7. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.

8. Expand the vHBAs drop-down menu and select Add SAN Boot.
9. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA field.
10. Confirm that Primary is selected for the Type option.

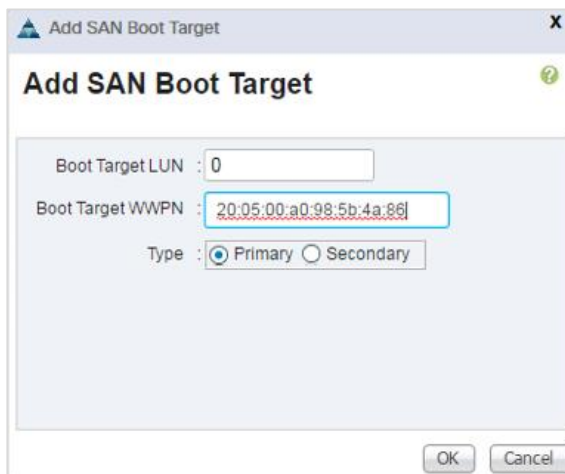


The image shows a dialog box titled "Add SAN Boot". It has a close button (X) in the top right corner and a help icon (?) in the top right corner of the main content area. The dialog contains two input fields: "vHBA" with the value "Fabric-A" and "Type" with radio buttons for "Primary" (selected) and "Secondary". At the bottom right, there are "OK" and "Cancel" buttons.

11. Click OK to add the SAN boot initiator.
12. From the vHBA drop-down menu, select Add SAN Boot Target.
13. Keep 0 as the value for the boot target LUN.
14. Enter the WWPN for `fcplif01a`.

Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.

15. Select Primary for the SAN boot target type.



The image shows a dialog box titled "Add SAN Boot Target". It has a close button (X) in the top right corner and a help icon (?) in the top right corner of the main content area. The dialog contains three input fields: "Boot Target LUN" with the value "0", "Boot Target WWPN" with the value "20:05:00:a0:98:5b:4a:86" (highlighted with a red dashed border), and "Type" with radio buttons for "Primary" (selected) and "Secondary". At the bottom right, there are "OK" and "Cancel" buttons.

16. Click OK to add the SAN boot target.
17. From the vHBA drop-down menu, select Add SAN Boot Target.
18. Enter 0 as the value for Boot Target LUN.
19. Enter the WWPN for `fcplif02a`.
20. Click OK to add the SAN boot target.
21. From the vHBA drop-down menu, select Add SAN Boot.

22. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.
The SAN boot type should automatically be set to Secondary and the Type option should be unavailable.
23. Click OK to add the SAN boot initiator.
24. From the vHBA drop-down menu, select Add SAN Boot Target.
25. Keep 0 as the value for Boot Target LUN.
26. Enter the WWPN for `fc_p_lif01b`.
27. Select Primary for the SAN boot target type.

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:06:00:a0:98:5b:4a:86

Type : ☒ Primary ☐ Secondary

OK Cancel

28. Click OK to add the SAN boot target.
29. From the vHBA drop-down menu, select Add SAN Boot Target.
30. Keep 0 as the value for Boot Target LUN.
31. Enter the WWPN for `fc_p_lif02b`.
32. Click OK to add the SAN boot target.

Create Boot Policy

Name :

Description :

Reboot on Boot Order Change : ☐

Enforce vNIC/vHBA/iSCSI Name : ☒

Boot Mode : ☒ Legacy ☐ Uefi

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

Local Devices

vNICs

vHBAs

iSCSI vNICs

CIMC Mounted vMedia

Boot Order

Name	Order	vNIC/vHBA...	Type	WWN	LUN Name
Remote CD/DVD	1				
San	2				
SAN Primary		Fabric-A	Primary		
SAN Target P...			Primary	20:05:00:A0:98:5B:4A:86	0
SAN Target S...			Secondary	20:07:00:A0:98:5B:4A:86	0
SAN Secondary		Fabric-B	Secondary		
SAN Target P...			Primary	20:06:00:A0:98:5B:4A:86	0
SAN Target S...			Secondary	20:08:00:A0:98:5B:4A:86	0

33. Click OK, then click OK again to create the boot policy.

Create Service Profile Template (FCoE Boot)

In this procedure, one service profile template for infrastructure ESXi hosts 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 and select Create Service Profile Template to open the Create Service Profile Template wizard.
4. Enter VM-Host-Prod-FC-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
5. Select the Updating Template option.
6. Under UUID, select UUID_Pool as the UUID pool.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

Identify Service Profile Template

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

Name : VM-Host-Prod-FC-A

The template will be created in the following organization. Its name must be unique within this organization.
Where : org-root

The template will be created in the following organization. Its name must be unique within this organization.
Type : ☐ Initial Template ☒ Updating Template

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

UUID

UUID Assignment:

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

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

< Prev Next > Finish Cancel

7. Click Next.

Configure Storage Provisioning

1. If you have servers with no physical disks, click the Local Disk Configuration Policy and select SAN-Boot Local Storage Policy. Otherwise, select the default local storage policy.
2. Click Next.

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Use Connectivity Policy option to configure the LAN connectivity.
3. Select Infra-LAN-Policy from the LAN Connectivity Policy drop-down list.
4. Click Next.

Configure Storage Options

1. Select the Use Connectivity Policy option for the How Would You Like to Configure SAN Connectivity? field.
2. Select the Infra-SAN-Policy option from the SAN Connectivity Policy drop-down list.



3. Click Next.

Configure Zoning Options

1. Set no zoning options and click Next.

Configure vNIC/HBA Placement

1. In the Select Placement list, leave the placement policy as Let System Perform Placement.
2. Click Next.

Configure vMedia Policy

1. From the vMedia Policy drop-down list, select ESXi-6.0U1b-HTTP.
2. Click Next.

Configure Server Boot Order

1. Select Boot-FC-A for the boot policy.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

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

Name : **Boot-FC-A**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Legacy**

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

Boot Order

Name	Order	v...	Type	WWN	L...	SI...	B...	B...	D...
CD/DVD	3								
Remote CD/DVD	1								
San	2								
SAN Primary		v...	Primary						
SAN Target Primary			Primary	20:00:00:A0:98:5B:4A:86					
SAN Target Secondary			Secondary	20:02:00:A0:98:5B:4A:86					
SAN Secondary		v...	Secondary						
SAN Target Primary			Primary	20:01:00:A0:98:5B:4A:86					

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

< Prev Next > Finish Cancel

2. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the maintenance policy to Default.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Maintenance Policy

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: **default** [+ Create Maintenance Policy](#)

Name : **default**
 Description :
 Reboot Policy : **User Ack**
 Config. Trigger State : **On Next Boot**

< Prev Next > Finish Cancel

2. Click Next.

Configure Server Assignment

To configure the server assignment, complete the following steps:

1. From the Pool Assignment list, select Infra_Pool.
2. Optional: Select a server pool qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Select UCS-Broadwell for the server pool qualification.
5. Leave Firmware Management as is because it will use the default from the Host Firmware list.

The screenshot shows the 'Create Service Profile Template' window in the Unified Computing System Manager. The left sidebar lists the steps: 1. Identify Service Profile Template, 2. Storage Provisioning, 3. Networking, 4. SAN Connectivity, 5. Zoning, 6. vNIC/vHBA Placement, 7. Media Policy, 8. Server Boot Order, 9. Maintenance Policy, 10. Server Assignment (highlighted), and 11. Operational Policies. The main area is titled 'Server Assignment' and contains the following options:

- Pool Assignment:** A dropdown menu set to 'Infra_Pool' with a '+ Create Server Pool' link.
- Power State:** Radio buttons for 'Up' and 'Down', with 'Down' selected.
- Server Pool Qualification:** A dropdown menu set to 'UCS-Broadwell'.
- Restrict Migration:** A checkbox that is unchecked.
- Firmware Management (BIOS, Disk Controller, Adapter):** A section with a note: 'If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.' Below this is a dropdown menu for 'Host Firmware Package' set to 'Select (Default policy used by default)' and a '+ Create Host Firmware Package' link.

At the bottom right, there are navigation buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

6. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. From the BIOS Policy list, select VM-Host-Infra.
2. Expand Power Control Policy Configuration and select No-Power-Cap from the Power Control Policy list.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Operational Policies
Optionally specify information that affects how the system operates.

▼ BIOS Configuration
If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile
BIOS Policy : VM-Host-Infra

► External IPMI Management Configuration

► Management IP Address

► Monitoring Configuration (Thresholds)

▼ Power Control Policy Configuration
Power control policy determines power allocation for a server in a given power group.
Power Control Policy : No-Power-Cap [Create Power Control Policy](#)

► Scrub Policy

► KVM Management Policy

< Prev Next > Finish Cancel

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profiles

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

1. Select Service Profile Templates > root > Service Template VM-Host-Infra-A.
2. Right-click VM-Host-Infra-A and select Create Service Profiles from Template.
3. Enter VM-Host-Infra-0 as the service profile prefix.
4. Enter 1 for Name Suffix Starting Number.
5. Enter 1 for Number of Instances.
6. Click OK to create the service profile.

Create Service Profiles From Template

Naming Prefix : VM-Host-Infra-0

Name Suffix Starting Number : 1

Number of Instances : 1

OK Cancel

7. Click OK in the confirmation message.
8. Select Service Profile Templates > root > Service Template VM-Host-Infra-A.
9. Right-click VM-Host-Infra-A and select Create Service Profiles from Template.
10. Enter VM-Host-Prod-0 as the service profile prefix.

11. Enter 1 for Name Suffix Starting Number.
12. Enter 1 for Number of Instances.
13. Click OK to create the service profile.
14. Click OK in the confirmation message.

8.5 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 are created, each infrastructure blade in the environment has 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 14 and Table 15.

Table 14) Storage target WWPNS.

Vserver	Target: WWPNS (FC)
Infra-SVM	
Prod-SVM	

Note: To obtain the FC WWPNS, run the `fcv show` command on the storage cluster management interface.

Table 15) Host initiator WWPNS.

Cisco UCS Service Profile Name	Initiator: WWPNS (FC)	Variables
VM-Host-Infra-01		<<var_vm_host_infra_01_wwpn1>> and <<var_vm_host_infra_01_wwpn2>>
VM-Host-Prod-01		<<var_vm_host_prod_01_wwpn1>> and <<var_vm_host_prod_01_wwpn2>>
VM-Host-Prod-02		<<var_vm_host_prod_02_wwpn1>> and <<var_vm_host_prod_02_wwpn2>>
VM-Host-Prod-03		<<var_vm_host_prod_03_wwpn1>> and <<var_vm_host_prod_03_wwpn2>>

Note: To obtain the FC vHBA WWPNS information in the Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the Storage tab and the vHBAs tab on the right. The WWPNS are displayed in the table at the bottom of the page.

9 Storage Configuration—Infrastructure LUNs and Igroups

9.1 Clustered Data ONTAP Storage Setup

Create Infrastructure Igroups

1. To create infrastructure igroups, run the following commands from the cluster management node SSH connection:

```

igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol fc -ostype vmware -initiator
<<var_vm_host_infra_01_wwpn1>>, <<var_vm_host_infra_01_wwpn2>>
igroup create -vserver Infra-SVM -igroup VM-Host-Prod-01 -protocol fc -ostype vmware -initiator
<<var_vm_host_prod_01_wwpn1>>, <<var_vm_host_prod_01_wwpn2>>
igroup create -vserver Infra-SVM -igroup VM-Host-Prod-02 -protocol fc -ostype vmware -initiator
<<var_vm_host_prod_02_wwpn1>>, <<var_vm_host_prod_02_wwpn2>>
igroup create -vserver Infra-SVM -igroup VM-Host-Prod-03 -protocol fc -ostype vmware -initiator
<<var_vm_host_prod_03_wwpn1>>, <<var_vm_host_prod_03_wwpn2>>

igroup create -vserver Infra-SVM -igroup Infra-Hosts -protocol fc -ostype vmware -initiator
<<var_vm_host_infra_01_wwpn1>>, <<var_vm_host_infra_01_wwpn2>>, <<var_vm_host_prod_01_wwpn1>>,
<<var_vm_host_prod_01_wwpn2>>, <<var_vm_host_prod_02_wwpn1>>, <<var_vm_host_prod_02_wwpn2>>,
<<var_vm_host_prod_03_wwpn1>>, <<var_vm_host_prod_03_wwpn2>>

```

Note: Use the values listed in Table 16 for the WWPN.

To view the igroups just created, run the `igroup show` command.

Map Boot LUNs to Hosts

1. To map boot LUNS to hosts, run the following commands from the storage cluster management SSH connection:

```

lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-01 -igroup VM-Host-Prod-01 -lun-id 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-02 -igroup VM-Host-Prod-02 -lun-id 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-03 -igroup VM-Host-Prod-03 -lun-id 0

```

Map Infrastructure LUNs to Hosts

1. To map infrastructure LUNs to hosts, run the following commands from the storage cluster management SSH connection:

```

lun map -vserver Infra-SVM -volume infra_datastore_1 -lun infra_datastore_1 -igroup infra-hosts -lun-id 1
lun map -vserver Infra-SVM -volume infra_swap -lun infra_swap -igroup infra-hosts -lun-id 2

```

Create Production Igroups

1. To create production igroups, run the following commands from the cluster management node SSH connection:

```

igroup create -vserver prod-SVM -igroup prod-hosts -protocol fc -ostype vmware -initiator
<<var_vm_host_prod_01_wwpn1>>, <<var_vm_host_prod_01_wwpn2>>, <<var_vm_host_prod_02_wwpn1>>,
<<var_vm_host_prod_02_wwpn2>>, <<var_vm_host_prod_03_wwpn1>>, <<var_vm_host_prod_03_wwpn2>>

```

Map Production LUNs to Hosts

1. To map production LUNs to hosts, run the following commands from the storage cluster management SSH connection:

```

lun map -vserver prod-SVM -volume prod_datastore_1 -lun prod_datastore_1 -igroup prod-hosts -lun-id 3
lun map -vserver prod-SVM -volume prod_swap -lun prod_swap -igroup prod-hosts -lun-id 4

```

10 VMware vSphere 6.0 U1 Setup

10.1 VMware ESXi 6.0

This section provides detailed instructions for installing VMware ESXi 6.0 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

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 boot LUNs.

VMware vSphere Configuration Variables

Table 16 lists the values that are required for the installation and configuration of VMware vSphere.

Table 16) VMware configuration variables.

Variable	Value
<<var_vm_host_infra_01_ip>>	
<<var_vm_host_prod_01_ip>>	
<<var_vm_host_prod_02_ip>>	
<<var_vm_host_prod_03_ip>>	
<<var_vmotion_vlan_ip_host_01>>	
<<var_vmotion_vlan_ip_host_02>>	
<<var_vmotion_vlan_ip_host_03>>	
<<var_vmotion_vlan_ip_host_04>>	
<<var_vcenter_ip>>	
<<var_vcenter_fqdn>>	
<<var_vcenter_subnet_mask>>	
<<var_vcenter_gateway>>	

Download Cisco Custom Image for ESXi 6.0 U1

1. Go to the [VMware login page](#).
2. Enter your e-mail or customer number and the password and click Log In.
3. Download [CiscoCustomImage6.0 U1b](#).
4. Save the image to your destination folder.

Note: This ESXi 6.0 U1b Cisco custom image includes updates for the fNIC and eNIC drivers. The versions that are part of this image are eNIC: 2.3.0.6 and fNIC: 1.6.0.24.

Log in to Cisco UCS Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the OS through remote media. You need 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. To download the Cisco UCS Manager software, click the Launch UCS Manager link.

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. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.
7. Select Servers > Service Profiles > root > VM-Host-Infra-01.
8. Right-click VM-Host-Infra-01 and select KVM Console.
9. If prompted to accept an unencrypted KVM session, accept as necessary.
10. Select Servers > Service Profiles > root > VM-Host-Prod-02.
11. Right-click VM-Host-Prod-02 and select KVM Console.
12. If prompted to accept an unencrypted KVM session, accept as necessary.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-01, -02, and -03

Note: Skip this procedure if using vMedia policies. The ISO file will already be connected to the KVM.

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

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Boot the server by selecting Boot Server and clicking OK. Click OK again.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-01, -02, and -03

To install VMware ESXi to the FC-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 boot menu displayed.
2. After the installer is loaded, 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 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 the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, click the Virtual Media tab and clear the ✓ mark next to the ESXi installation media. Click Yes.

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

9. 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:

1. After the server reboots, press F2 to customize the system.
 2. Log in as `root`, enter the corresponding password, and press Enter to log in.
 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. Select the Network Adapters option, select `vmnic04`, and press Enter.
 7. From the Configure Management Network menu, select IP Configuration and press Enter.
 8. Select the Set Static IP Address and Network Configuration option by using the space bar.
 9. Enter the appropriate IP address for managing the ESXi host: `<<var_vm_host_infra_01_ip>>`, `<<var_vm_host_prod_01_ip>>`, `<<var_vm_host_prod_02_ip>>`, `<<var_vm_host_prod_03_ip>>`.
 10. Enter the subnet mask for the ESXi host.
 11. Enter the default gateway for the ESXi host.
 12. Press Enter to accept the changes to the IP configuration.
 13. Select the IPv6 Configuration option and press Enter.
 14. Using the space bar, select Disable IPv6 (restart required) and press Enter.
 15. Select the DNS Configuration option and press Enter.
- Note:** Because the IP address is assigned manually, the DNS information must also be entered manually.
16. Enter the IP address of the primary DNS server.
 17. Optional: Enter the IP address of the secondary DNS server.
 18. Enter the fully qualified domain name (FQDN) for the ESXi host.
 19. Press Enter to accept the changes to the DNS configuration.
 20. Press Esc to exit the Configure Management Network submenu.
 21. Press Y to confirm the changes and return to the main menu.
 22. Select Test Management Network to verify that the management network is set up correctly. 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

To download VMware vSphere Client, complete the following steps:

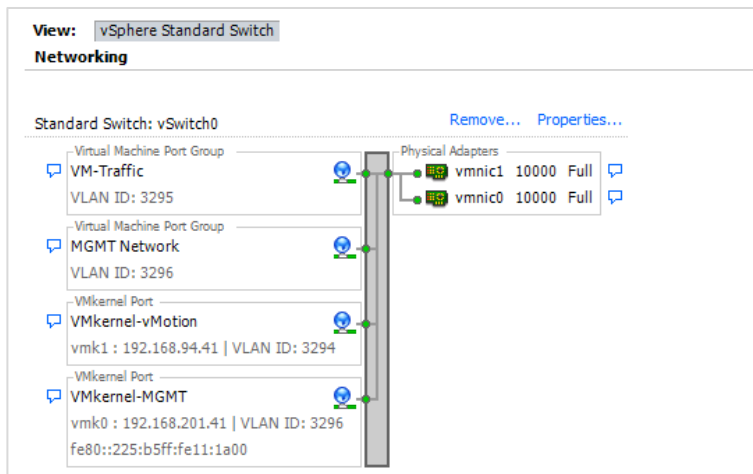
1. Open a web browser on the management workstation and navigate to the `VM-Host-Infra-01` management IP address.
2. Download and install vSphere Client.

Note: This application is downloaded from the VMware website. Internet access is required on the management workstation.

Configure ESXi Host Networking

To set up the VMkernel ports and the virtual switches on the ESXi hosts, complete the following steps on all hosts:

1. Log in to vSphere Client.
2. From vSphere Client, select the host in the inventory.
3. Click the Configuration tab.
4. In the Hardware pane, click Networking.
5. On the right side of vSwitch0, click Properties.
6. On the Network Adapters tab, click Add.
7. Select the vmnic1 checkbox and click Next.
8. Click Next to accept the default failover order configuration, then click Finish.
9. In the Ports tab, select the vSwitch configuration and click Edit.
10. From the General tab, change the MTU to 9000.
11. Click OK.
12. Select the Management Network configuration and click Edit.
13. Change the network label to `VMkernel-MGMT` and select the Management Traffic checkbox.
14. Click OK to finalize the edits for the management network.
15. Select the VM Network configuration and click Edit.
16. Change the network label to `MGMT Network` and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.
17. Click OK to finalize the edits for VM Network.
18. Click Add to create a vMotion VMkernel port.
19. Select the VMkernel option and click Next.
20. Change the network label to `VMkernel-vMotion` and enter `<<var_vmotion_vlan_id>>` in the VLAN ID (Optional) field.
21. Select the checkbox to enable vMotion on this port.
22. Click Next.
23. Enter the IP address `<<var_vmotion_vlan_ip_host_01>>` and the subnet mask `<<var_vmotion_vlan_ip_mask>>` for the vMotion VLAN interface for VM-Host-Prod-01.
24. To continue with the vMotion VMkernel creation, click Next.
25. To finalize the creation of the vMotion VMkernel interface, click Finish.
26. Select the `VMkernel-vMotion` configuration and click Edit.
27. Change the MTU to 9000.
28. Click OK to finalize the edits for the VMkernel-vMotion network.
29. Click Add to create the VM traffic network.
30. Select the Virtual Machine option and click Next.
31. Change the network label to `VM-Traffic` and enter `<<var_vm-traffic_vlan_id>>` in the VLAN ID (Optional) field.
32. Click Next and then click Finish to create the virtual network.
33. To finalize the ESXi host networking setup, close the dialog box.
34. Repeat this process for each ESXi host.
35. Networking for the ESXi host should be similar to the following example:



Mount Infrastructure Datastores

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

1. From vSphere Client, select the host in the inventory.
2. Click the Configuration tab and select Storage in the Hardware pane.
3. From the Datastores area, click Add Storage to open the Add Storage wizard.
4. Select Disk/LUN and click Next.
5. Select LUN 1 (500GB) and click Next.
6. Click Next to accept the default partitioning.
7. Enter `infra_datastore_1` as the name for the datastore and click Next.
8. Click Next to accept the partition sizing.
9. Click Finish to create the datastore.
10. From the Datastores area, click Add Storage to open the Add Storage wizard.
11. Select Disk/LUN and click Next.
12. Select LUN 2 (100GB) and click Next.
13. Click Next to accept the default partitioning.
14. Enter `infra_swap` as the name for the datastore and click Next.
15. Click Next to accept the partition sizing.
16. Click Finish to create the datastore.
17. After the datastores are created on one host, use the Refresh option on the Storage page to discover the datastores on the remaining hosts.

View: Datastores Devices							
Datastores							
Identification	Device	Drive Type	Capacity	Free	Type	Last Update	Hardware Acceleration
datastore1	NETAPP Fibre Ch...	Non-SSD	7.50 GB	6.66 GB	VMFSS	10/12/2016 11:02:51 ...	Supported
infra_datastore_1	NETAPP Fibre Ch...	Non-SSD	500.00 GB	499.05 GB	VMFSS	10/12/2016 11:02:51 ...	Supported
infra_swap	NETAPP Fibre Ch...	Non-SSD	99.75 GB	98.80 GB	VMFSS	10/12/2016 11:02:51 ...	Supported

Install VMware Drivers for the Cisco Virtual Interface Card

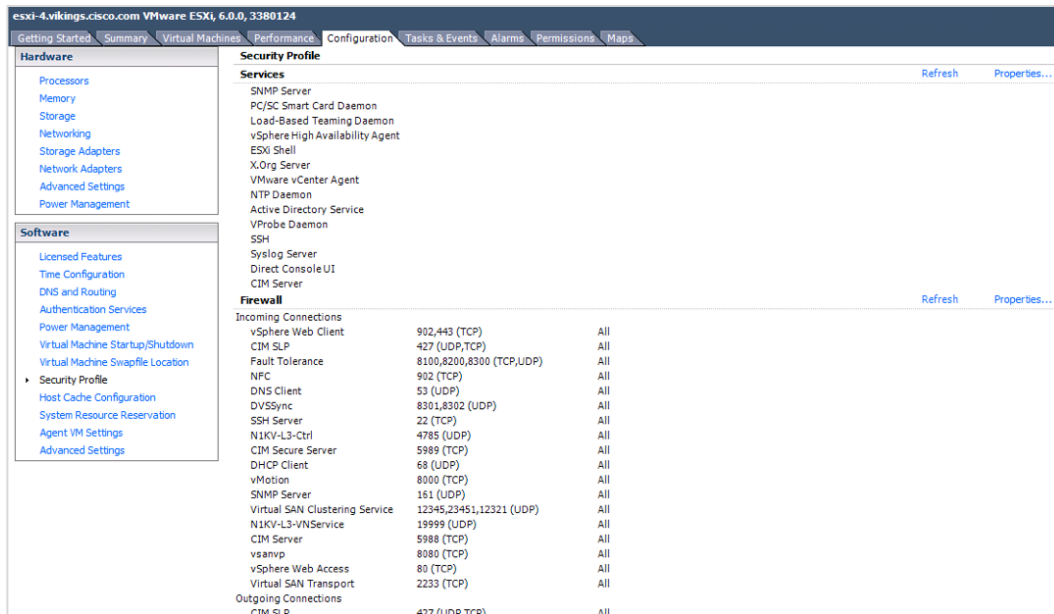
Download and extract the following VMware Cisco Virtual Interface Card (VIC) drivers to the management workstation:

- [fNIC driver version 1.6.0.25](#)
- [eNIC driver version 2.3.0.7](#)


All ESXi Hosts

To install VMware VIC drivers on the ESXi hosts, complete the following steps:

1. From vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. Click Enter Maintenance Mode within the Commands section of the Summary tab.
4. Click Yes for any dialog box presented.
5. From Resources > Storage, right-click infra_datastore_1 and select Browse Datastore.
6. Click the fourth button and select Upload File.
7. Navigate to the saved location for the downloaded VIC drivers and select fnic_driver_1.6.0.25-3741467.zip.
8. Click Open and then Yes to upload the file to infra_datastore_1.
9. Click the fourth button and select Upload File.
10. Navigate to the saved location for the downloaded VIC drivers and select ESXi60-enic-2.3.0.7-3642661.zip.
11. Click Open and then Yes to upload the file to infra_datastore_1.
12. After the files are uploaded to the shared datastore, all hosts can access the files.
13. Within vSphere Client, select the Configuration tab and click Security Profile within the Software section.



14. Click Properties under the Services section at the top.
15. Select ESXi Shell and click Options.



16. Select Start Automatically If Any Ports Are Open and Stop When All Ports Are Closed within the Startup Policy section.
17. Click Start and then OK.
18. Select SSH and click Options.
19. Select Start Automatically If Any Ports Are Open and Stop When All Ports Are Closed within the Startup Policy section.
20. Click Start and then OK.
21. Click OK to exit from the Service Properties configuration window.

Note: Enabling SSH can be considered optional if the VMware vSphere Remote CLI is installed and used.

22. Connect to each ESXi host through SSH from a shell connection or Putty terminal.
23. Log in as root with the password specified for <<var_password>>.
24. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore/fnic_driver_1.6.0.25-offline_bundle-3642682.zip
esxcli software vib install -d /vmfs/volumes/datastore/ESXi60-enic-2.3.0.7-offline_bundle-3642661.zip
```

25. Reboot each host after both commands have been run.

Configure NTP on ESXi Hosts

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

1. From vSphere Client, select the host in the inventory.
2. Click the Configuration tab.


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 (ntpd) 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_switch_a_ntp_ip>> as the IP address of the NTP server and click OK.
8. Click Add.
9. In the Add NTP Server dialog box, enter <<var_switch_b_ntp_ip>> as the IP address of the NTP server and click OK.
10. In the NTP Daemon Options dialog box, select the Restart NTP Service to Apply Changes checkbox and click OK.
11. 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 might vary slightly from the host time.

Move VM Swap File Location

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

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the Configuration tab.
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 the `infra_swap` datastore in which to house the swap files.



- Click OK to finalize moving the swap file location.

10.2 VMware vCenter 6.0

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.0 U1b server appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Install the Client Integration Plug-In

To install the client integration plug-in, complete the following steps:

- Download the .iso installer for the [vCenter Server Appliance and Client Integration Plug-in](#).
- Mount the ISO image to the Windows VM or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter server appliance.
- In the software installer directory, navigate to the VCSA directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-in installation wizard appears.
- On the Welcome page, click Next.
- Read and accept the terms in the EULA and click Next.
- Click Next.
- Click Install.

Building the VMware vCenter Server Appliance

To build the VMware vCenter virtual machine, complete the following steps:

- In the software installer directory, double-click vcsa-setup.html.
- Allow the plug-in to run on the browser when prompted.
- On the Home page, click Install to start the vCenter Server Appliance Deployment wizard.

4. Read and accept the license agreement and click Next.
5. In the Connect to Target Server page, enter the ESXi host name, user name, and password.
6. Click Yes to accept the certificate.
7. In the Set Up Virtual Machine page, enter the appliance name and password details.
8. In the Select Deployment Type page, select Install vCenter Server with an Embedded Platform Services Controller. Click Next
9. In the Set Up Single Sign-On page, select Create a New SSO Domain. Enter the SSO password, domain, and site name. Click Next.
10. In the Select appliance size page, select the appliance size; for example, Tiny (up to 10 hosts, 100 VMs). Click Next.
11. In the Select Datastore page, select infra_datastore_1 and click Next.
12. In the Configure Database page, select Use an Embedded Database. Click Next.
13. In the Network Settings page, configure the settings as follows:
 - a. Choose a network: MGMT-Network
 - b. IP address family: IPV4
 - c. Network type: static
 - d. Network address: <<var_vcenter_ip>>
 - e. System name: <<var_vcenter_fqdn>>
 - f. Subnet mask: <<var_vcenter_subnet_mask>>
 - g. Network gateway: <<var_vcenter_gateway>>
 - h. Network DNS servers: <<var_nameserver_ip>>
 - i. Configure time sync: Use NTP servers
 - j. (Optional): Enable SSH

The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard at the 'Network Settings' step. The left sidebar shows a progress list with steps 1 through 10, where step 9 'Network Settings' is currently selected. The main area is titled 'Network Settings' and contains the following fields:

- Choose a network:** A dropdown menu set to 'Management Network'.
- IP address family:** A dropdown menu set to 'IPv4'.
- Network type:** A dropdown menu set to 'static'.
- Network address:** A text box containing '10.1.156.100'.
- System name [FQDN or IP address]:** A text box containing 'vc.vikings.cisco.com'.
- Subnet mask:** A text box containing '255.255.255.0'.
- Network gateway:** A text box containing '10.1.156.1'.
- Network DNS Servers (separated by commas):** A text box containing '192.168.156.9'.
- Configure time sync:** Two radio buttons. The first is 'Synchronize appliance time with ESXi host' (unselected). The second is 'Use NTP servers (Separated by commas)' (selected).

At the bottom of the window are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

14. Review the configuration and click Finish.
15. The vCenter appliance installation takes a few minutes to complete.

Setting Up VMware vCenter Server


1. Using a web browser, navigate to https://<var_vcenter_ip>.
2. Click Log in to vSphere Web Client.
3. Click OK if the Launch Application window appears.
4. Log in using the Single Sign-On user name and password created during vCenter installation.
5. Navigate to vCenter Inventory Lists on the left pane.
6. Under Resources, click Datacenters in the left pane.
7. To create a data center, click the left-most icon in the center pane that has a green plus symbol above it.
8. Type `FlexPod_DC` in the Datacenter Name field. Select the vCenter Name/IP option and click OK.
9. Right-click the data center `FlexPod_DC` in the list in the center pane. Click New Cluster.
10. Name the cluster `FlexPod_Management`.
11. Select the DRS checkbox and retain the default values.
12. Select the vSphere HA checkbox and retain the default values.

New Cluster

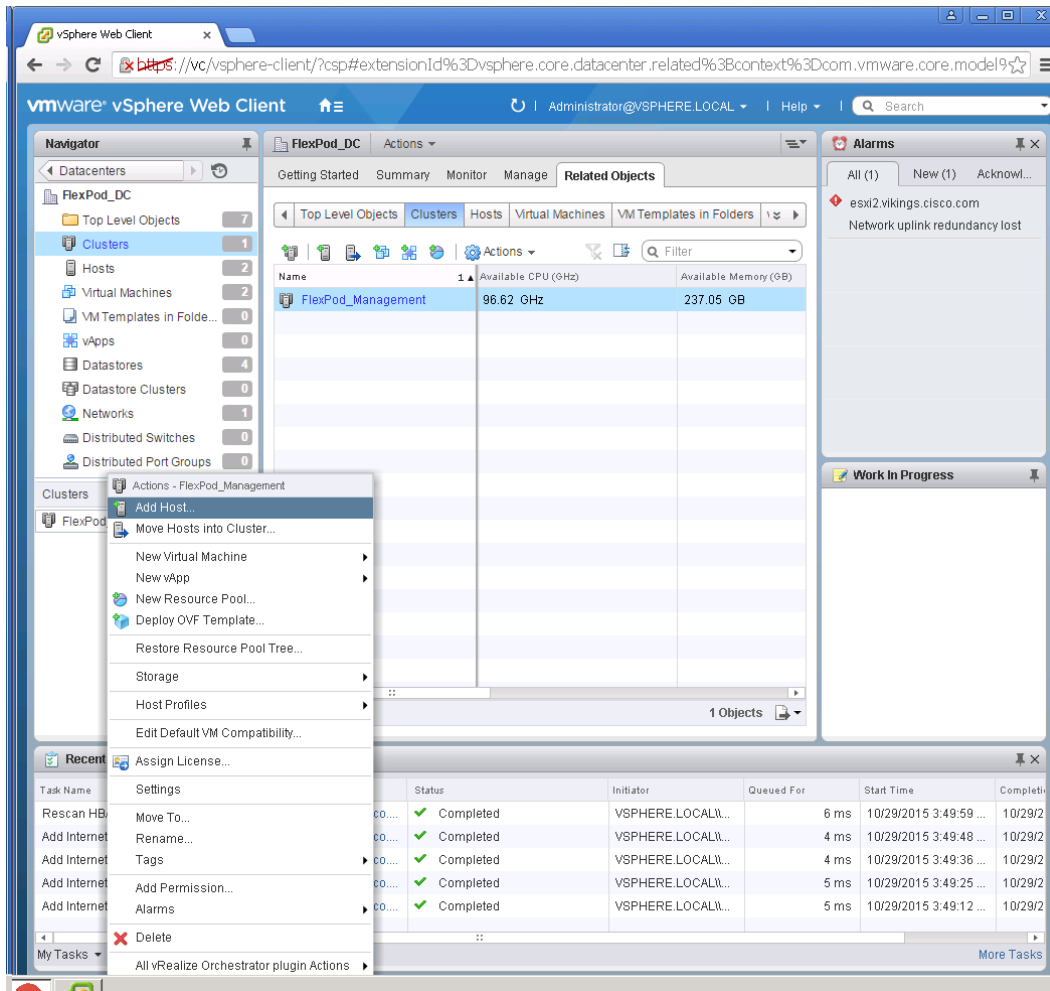
Name	FlexPod_Management
Location	FlexPod_DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative — Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control
Admission Control Status	Admission control will prevent powering on VMs that violate availability constraints
Policy	Specify the type of the policy that admission control should enforce. <input checked="" type="radio"/> Host failures cluster tolerates: 1 <input type="radio"/> Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: 25 % CPU Reserved failover Memory capacity: 25 % Memory
VM Monitoring	<input type="checkbox"/> Turn ON
VM Monitoring Status	Disabled
Monitoring Sensitivity	Low — High
EVC	Disable
Virtual SAN	<input type="checkbox"/> Turn ON

OK Cancel

13. Click OK to create the cluster.
14. On the left pane, double-click `FlexPod_DC`.
15. Click Clusters.



16. Under the Clusters pane, right-click FlexPod_Management and select Settings.
17. Select Configuration > General from the list on the left and select Edit to the right of General.
18. Select Datastore Specified by Host and click OK.
19. Under the Clusters pane, right-click FlexPod_Management and click Add Host.



20. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.
21. Type root as the user name and the root password. Click Next to continue.
22. Click Yes to accept the certificate.
23. Review the host details and click Next to continue.
24. Assign a license and click Next to continue.
25. Click Next to continue.
26. Click Next to continue.
27. Review the configuration parameters. Then click Finish to add the host.
28. Repeat steps 21 through 29 to add the remaining VMware ESXi hosts to the cluster.

11 FlexPod Management Tools Setup

11.1 Management Tools Configuration Variables

Table 17 lists the values that are required for installation of the FlexPod management tools.

Table 17) FlexPod management tools configuration variables.

Variable	Value
<<var_oncommand_server_ip>>	
<<var_oncommand_server_mask>>	
<<var_oncommand_server_gateway>>	
<<var_oncommand_pm_ip>>	
<<var_oncommand_pm_mask>>	
<<var_oncommand_pm_gateway>>	

12 NetApp Virtual Storage Console 6.2P1 Deployment Procedure

This section describes the deployment procedures for the NetApp VSC.

12.1 Virtual Storage Console 6.2P1 Preinstallation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3.1:

- Protocol licenses (NFS and FCP)
- NetApp FlexClone (for provisioning and cloning only)
- NetApp SnapRestore (for backup and recovery)
- NetApp SnapManager suite

12.2 Install Virtual Storage Console 6.2P1

To install the VSC 6.2P1 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 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.
2. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain.
3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure feature on the VM.
4. Install all Windows updates on the VM.
5. Log in to the VSC VM as FlexPod admin user.
6. Download the x64 version of the [Virtual Storage Console 6.2P1](#) from the [NetApp Support](#) site.
7. From the VMware Console, right-click the VSC .exe file downloaded in step 6 and select Run As Administrator.
8. Select the appropriate language and click OK.
9. On the Installation Wizard Welcome page, click Next.

Note: Review and accept the terms and click Next.
10. The Backup and Recovery capability requires an extra license.
11. Click Next to accept the default installation location.
12. Click Install.
13. After the installation completes, click Finish to exit the installer.

12.3 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. If the URL does not open automatically, open <https://localhost:8143/Register.html> in Internet Explorer.
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 or root), and user password for the vCenter Server. Click Register to complete the registration.
5. After successful registration, the storage controllers are discovered automatically.

Note: The storage discovery process might take some time.

12.4 Discover and Add Storage Resources

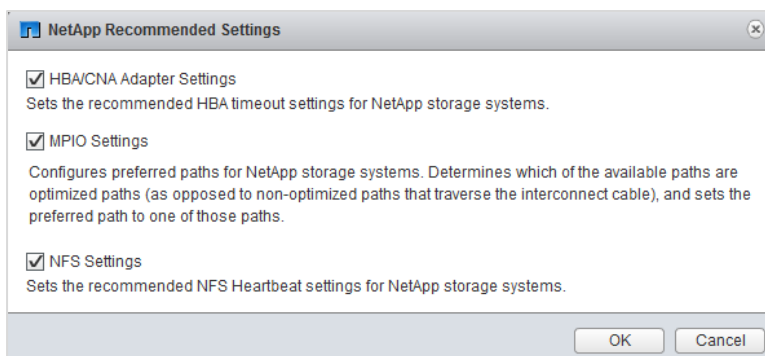
To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, complete the following steps:

1. Using vSphere Web Client, log in to the vCenter Server as FlexPod admin user or root. If vSphere Web Client was previously opened, close it and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.
3. Select Storage Systems. Under the Objects tab, click Actions > Modify.
4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for the password. Confirm that Use SSL to Connect to This Storage System is selected. Click OK.
5. Click OK to accept the controller privileges.

12.5 Optimal Storage Settings for ESXi Hosts

VSC allows the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for These Hosts.



2. Check the settings that are to be applied to the 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).
3. Click OK.

4. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

12.6 VSC 6.2P1 Backup and Recovery

Prerequisites

Before you use the backup and recovery capability to schedule backups and restore your datastores, VMs, or virtual disk files, you must confirm that the storage systems that contain the datastores and VMs for which you are creating backups have valid storage credentials.

Note: If you plan to use the `snapmirror update` option, add all the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

To set up backup and recovery using VSC 6.2P1, complete the following steps:

1. From the Home screen, select the Home tab and click Storage.
2. On the left, expand the Datacenter and select Datastores.
3. Right-click the datastore that you want to back up. Select NetApp VSC > Backup > Schedule Backup Job.

Note: If you prefer a one-time backup, then select Backup Now instead of Schedule Backup.

4. Type a backup job name and description and click Next.

Note: If you want to create a VMware snapshot for each backup, select Perform VMware Consistency Snapshot in the Options pane.

5. In the Options page, select the options to include in the backup and click Next.
6. In the Spanned Entities page, click Next.
7. In the Scripts page, select one or more backup scripts if available and click Next.
8. In the Schedule and Retention page, Select the hourly, daily, weekly, or monthly schedule for this backup job and click Next.
9. In the Credentials and Alerts page, use the default vCenter credentials or enter the user name and password for vCenter Server and click Next.
10. Specify backup retention details as per the requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate the addresses. Click Next.
11. 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.
12. Click OK when notified about the successful backup.
13. On the storage cluster interface, you can disable automatic Snapshot copies of the volume by entering this command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

14. To delete any existing automatic Snapshot copies that were created on the volume, enter this command:

```
volume snapshot show -volume infra_datastore_1  
volume snapshot delete -volume infra_datastore_1 -vserver Infra-SVM -snapshot <snapshot name>
```

Note: You can use the wildcard character (*) in snapshot names in the previous command.

13 NetApp OnCommand Unified Manager 6.3P4 Deployment Procedure

This section describes the deployment procedures for NetApp OnCommand® Unified Manager.

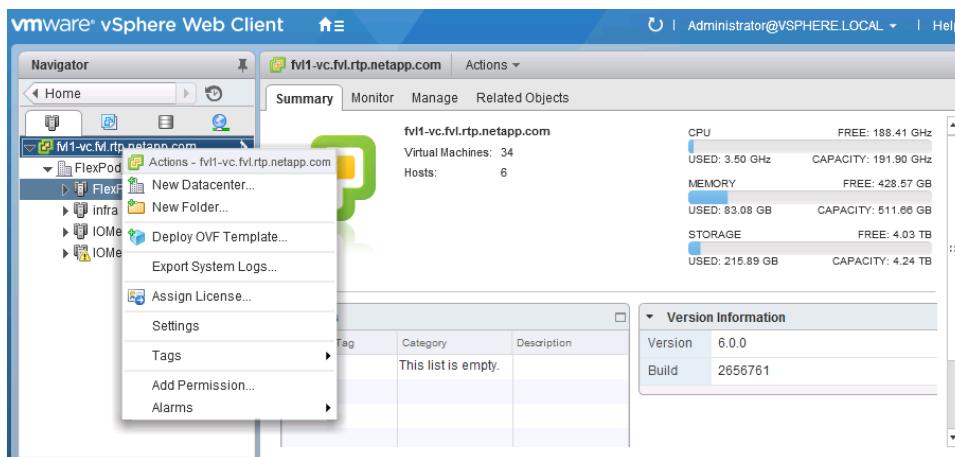
13.1 OnCommand Unified Manager 6.3P4

OnCommand Unified Manager OVF Deployment

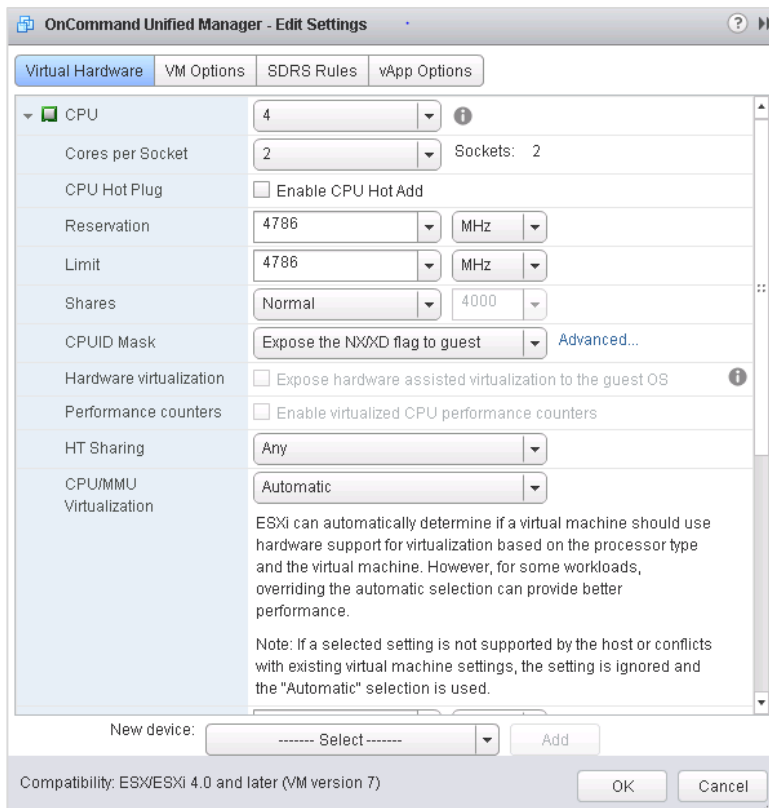
To install OnCommand Unified Manager, complete the following steps.

Note: Download and review the [OnCommand Unified Manager Installation and Setup Guide](#).

1. Download [OnCommand Unified Manager version 6.3P2](#) (OnCommandUnifiedManager-6.3P2.ova).
2. Log in to vSphere Web Client. Go to vCenter > VMs and Templates.
3. At the top of the center pane, select Actions > Deploy OVF Template.



4. Browse the .ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.
5. Select Accept Extra Configuration Options and click Next.
6. Review the deployment details and click Next.
7. Read and accept the EULA and then click Next.
8. Enter the name of the VM and select the FlexPod_DC folder to contain it. Click Next.
9. Select FlexPod_Management within the FlexPod_DC data center as the destination compute resource pool to host the VM. Click Next.
10. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the virtual disk format. Click Next.
11. Select IB-MGMT-VLAN as the destination network for the nat source network. Click Next.
12. In the Customize template page, enter the host name, IP address, network mask, gateway, and primary and secondary DNS fields. Click Next.
13. Clear the Power On After Deployment checkbox.
14. Review the configuration details and click Finish to deploy the VM with the provided configuration information.
15. In the left pane, select vCenter > Virtual Machines. After the OVF deployment is complete, right-click the newly created VM and select Edit Settings.



16. Size the VM's CPU and memory parameters according to the [OnCommand Unified Manager 6.3 Installation and Setup Guide](#).
17. Click OK to accept the changes.
18. Right-click the VM in the left-hand pane and click Power On.

OnCommand Unified Manager Basic Setup

To set up OnCommand Unified Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.
2. In the VMRC window, select VMRC > Manage > Install VMware Tools. VMware Tools will install in the VM.
3. Set up OnCommand Unified Manager by answering the following questions in the console window:

Geographic area: <<Enter the number corresponding to your time zone>>

```

Performing setup of VMware Tools. This may take a few minutes.
Please wait.....
VMware Tools installation/upgrade complete

Configuring timezone...

Configuring tzdata
-----

Please select the geographic area in which you live. Subsequent configuration
questions will narrow this down by presenting a list of cities, representing the
time zones in which they are located.

1. Africa          6. Asia           11. System V timezones
2. America         7. Atlantic Ocean 12. US
3. Antarctica      8. Europe         13. None of the above
4. Australia       9. Indian Ocean
5. Arctic Ocean    10. Pacific Ocean

Geographic area: _

```

These commands complete the network configuration checks, generate SSL certificates for HTTPS, and start the OnCommand Unified Manager services.

4. To create a Maintenance User account, run the following commands:

Note: The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

```

Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>

```

```

OnCommand Unified Manager

System IP addresses:
-----
10.29.128.170

Log in to OnCommand Unified Manager in a web browser by using
https://10.29.128.170/
or
https://ocum.ciscorobo.com/

The maintenance console should be used when the web interface is not available.
For normal usage of OnCommand Unified Manager, use the web interface.

ocum login: _


```

5. Using a web browser, navigate to OnCommand Unified Manager with the following link: https://<<var_oncommand_server_ip>>.
6. Log in using the maintenance user account credentials.
7. Select Yes to enable AutoSupport capabilities.
8. Click Continue.
9. Enter the NTP server IP address <<var_switch_a_ntp_ip>>.
10. Enter the storage admin e-mail address <<var_storage_admin_email>>.
11. Enter the SMTP server host name and click Save.
12. Click Add Cluster.

13. Provide the cluster management IP address, user name, password, protocol, and port. Click Add.
14. Click Yes to trust the certificate from the controller.

Note: The Cluster Add operation might take a couple of minutes.

15. After the cluster is added, you can access it by clicking the Storage tab and selecting Clusters.

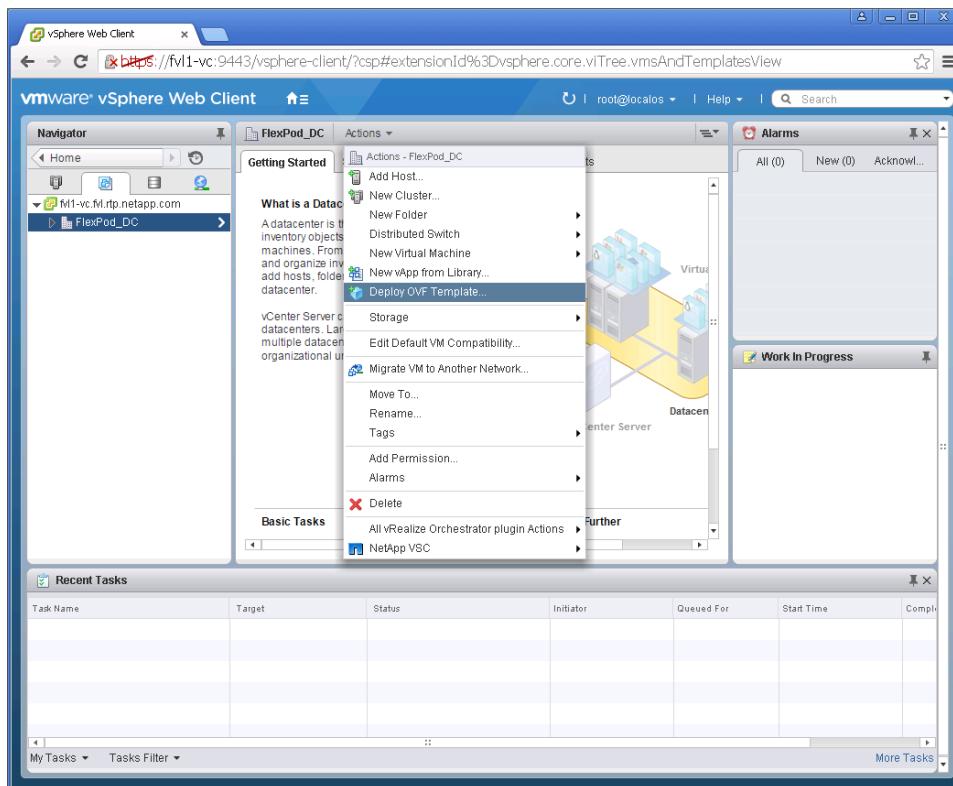


13.2 OnCommand Performance Manager 2.0

OnCommand Performance Manager OVF Deployment

To install OnCommand Performance Manager, complete the following steps:

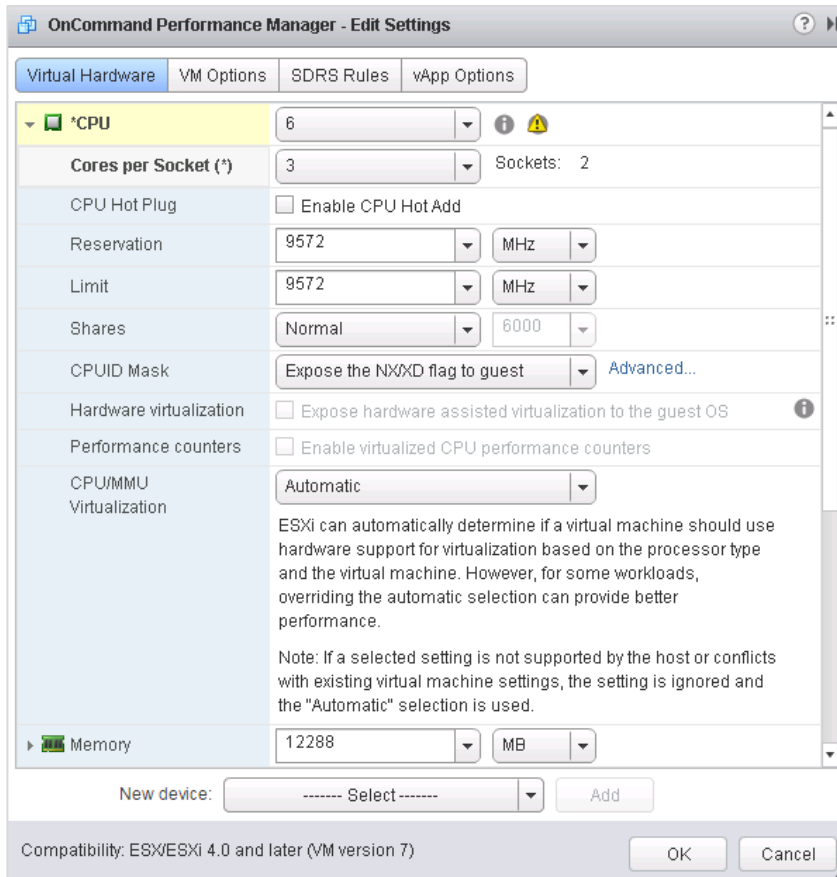
1. Download and review the [OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances](#).
2. Download [OnCommand Performance Manager version 2.0](#) (OnCommandPerformanceManager-netapp-2.0.0.ova).
3. Log in to vSphere Web Client. Select Home > VMs and Templates.
4. At the top of the center pane, click Actions > Deploy OVF Template.



5. Browse to the `OnCommandPerformanceManager-netapp-2.0.0.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.
6. Review the details and click Next.
7. Read and accept the EULA and then click Next.
8. Enter the name of the VM and select the `FlexPod_DC` folder to hold the VM. Click Next.
9. Select `FlexPod_Management` within the `FlexPod_DC` data center as the destination compute resource pool to host the VM. Click Next.
10. In the Select storage space, select `infra_datastore_1` as the storage target for the VM. Select Thin Provision as the virtual disk format and click Next.
11. Select `IB—MGMT-VLAN` as the destination network for the nat source network. Click Next.
12. Enter the host name, IP address, network mask, gateway, primary DNS, and secondary DNS. Click Next.
13. Deselect Power On After Deployment.
14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.
15. In the left pane, navigate to Home > Hosts and Clusters. Expand the `FlexPod_Management` cluster and select the newly created OnCommand Performance Manager VM. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.
16. Expand the CPU options.
 - a. The minimum required CPU reservation is 9572MHz. Determine the CPU frequency of the host.
 - b. Set the required number of CPUs ($9572 / \text{CPU frequency of the host}$).
 - c. Set the number of cores per socket where the socket number on the right matches the number of CPU sockets in the host. For example, if a host has two CPUs operating at a speed of 1999MHz,

then the VM requires six virtual CPUs ($9572 / 1999 = 4.79$ – rounded to 6 virtual CPUs). If the host has two physical CPU sockets, then set three cores per socket.

Note: For detailed information, see [OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances](#).



17. Click OK to accept the changes.
18. Right-click the VM in the left-hand pane and click Power On.

OnCommand Performance Manager Basic Setup

To set up OnCommand Performance Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.
2. In the VMware Remote Console window, select VMRC > Manage > Install VMware Tools. VMware Tools installs in the VM.
3. Set up OnCommand Performance Manager by answering the following questions in the console window:

```
Geographic area: <<Enter your geographic location>>
Time zone: <<Select the city or region corresponding to your time zone>>
```

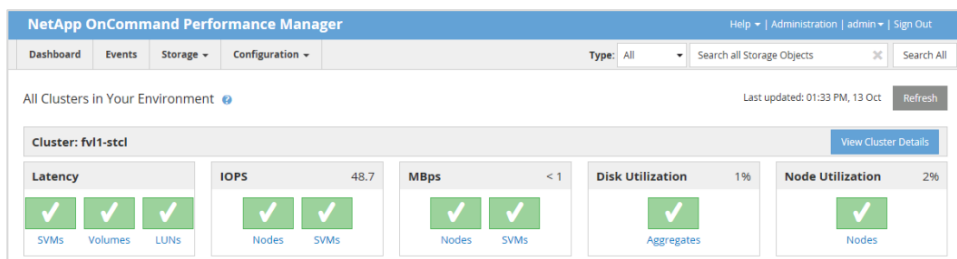
These commands complete the network configuration checks, generate SSL certificates, and start OnCommand Performance Manager services.

4. To create a maintenance user account, run the following commands:

Note: The maintenance user manages and maintains the settings on the OnCommand Performance Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

5. Using a web browser, navigate to OnCommand Performance Manager using the URL `https://<<var_oncommand_pm_ip>>`.
6. Log in using the maintenance user account (admin) credentials.
7. Enter a maintenance user e-mail address, SMTP mail server information, and the NTP server IP address. Click Save.
8. Select the Yes option to enable AutoSupport capabilities. Click Save.
9. Click Save and go to the next step to not change the admin password.
10. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster and then click Save and Complete Configuration. It can take up to 15 minutes for the cluster to be visible in OnCommand Performance Manager.



11. After the cluster is added, it can be accessed by clicking Administration > Manage Data Sources.

The screenshot shows the 'Manage Data Sources' page in the NetApp OnCommand Performance Manager. The left sidebar contains a 'Management' section with links to 'Users', 'Data Sources', and 'Setup'. The 'Data Sources' link is selected. The main content area shows a table of data sources. The table has columns for Name, Host Name or IP A, Protocol, Port, User Name, Status, and Status Message. A single data source is listed with the name 'fvl1-stdl', host name '172.20.70.10', protocol 'HTTPS', port '443', user name 'admin', and status 'Normal'.

Name	Host Name or IP A	Protocol	Port	User Name	Status	Status Message
fvl1-stdl	172.20.70.10	HTTPS	443	admin	Normal	

Link OnCommand Performance Manager to OnCommand Unified Manager

To link OnCommand Performance Manager to OnCommand Unified Manager, complete the following steps:

1. Using a web browser, navigate to OnCommand Unified Manager using the URL `https://<<var_oncommand_server_ip>>`. Log in with the maintenance user ID and password set up previously.
2. In the OnCommand Unified Manager web interface, select Administration > Manage Users to set up an Event Publication user.
3. Click Add to add a user.
4. Leave Type set to Local User. Use eventpub as the name and enter and confirm a password. Enter an e-mail address for this user and set Role to Event Publisher. Click Add.

5. In the OnCommand Performance Manager console window, log into the CLI with the maintenance user (admin) defined previously.
6. Enter 5 to select Unified Manager Connection.

```

OnCommand Performance Manager - VMware Remote Console
VMRC
For regular system operation and usage, use the UI.
ocpm login: admin
Password:
Linux OnCommand 3.2.0-4-amd64 #1 SMP Debian 3.2.68-1+deb7u1 x86_64

OnCommand Performance Manager Maintenance Console

Version      : 2.0.0
System ID    : 8cde1672-8af5-446b-8e08-45deef876778
Status       : Running

Main Menu
-----
1 ) Upgrade
2 ) Network Configuration
3 ) System Configuration
4 ) Support/Diagnostics
5 ) Unified Manager Connection
6 ) External Data Provider
7 ) Backup/Restore

x ) Exit

Enter your choice: 5_

```

7. Enter 2 to Add / Modify Unified Manager Server Connection.
8. Enter *y* to continue.
9. Enter the OnCommand Unified Manager FQDN or IP address.
10. Click Enter to accept the default port 443.
11. Enter *y* to accept the Unified Manager security certificate.
12. Enter *eventpub* for Event Publisher User Name.
13. Enter the *eventpub* password.
14. Enter *y* to accept the entered settings.
15. Press any key to continue.
16. Exit the OnCommand Performance Manager console. OnCommand Performance Manager events now appear in the OnCommand Unified Manager Dashboard.

14 SQL Server Virtual Machine Deployment and Configuration

This section documents the deployment procedure for the SQL Server virtual machines to be used in a SQL Server failover cluster instance. This procedure creates a single AlwaysOn Availability Group. Repeat these procedures as necessary to provision the required number of SQL Servers and Availability Groups.

The configuration created in this procedure is the bare minimum necessary to operate SQL Server 2014 AlwaysOn Availability Groups. Customer implementations should be sized and configured based on actual database and application requirements.

This solution is based on separate compute and storage resources for primary and secondary SQL database replicas. Each pair of SQL Servers should be configured using one SQL Server VM from the primary infrastructure and one SQL Server VM from the secondary infrastructure.

14.1 Deploy Production Datastores

Dedicated storage should be created to host production virtual machine boot disks as well as SQL data. To create the necessary datastores, complete the following steps:

1. Log in to vSphere Web Client.
2. Select Hosts and Clusters.
3. Right-click the `SQL-PRI` cluster and select NetApp VSC > Provision Datastore.
4. Enter `prod_datastore_1` as the datastore name. Select VMFS for the datastore type and FC/FCoE for the VMFS protocol. Click Next.
5. In the Storage system page, make sure that the `sqlpri` storage system is selected. Select prod-SVM from the SVM drop-down list and click Next.
6. In the Details page, select Thin Provision and enter 500GB for the size. Leave the Create New Volume box selected and choose an appropriate aggregate. Click Next.
7. Review the details and click Finish to provision the datastore.
8. Repeat steps 1 through 7 to provision a datastore on the secondary storage system for secondary SQL replicas.
9. Repeat steps 1 through 7 to provision the following datastores for each SQL database. These datastores should be sized based on the database requirements and will be hosted on the primary and secondary storage systems as appropriate. Database instances should also be balanced across controllers within each storage system, with all LUNs for each database on the same controller.
 - `<db01-primary>_data`
 - `<db01-primary>_log`
 - `<db01-primary>_tempdb`
 - `<db01-primary>_snapinfo`
 - `<db01-secondary>_data`
 - `<db01-secondary>_log`
 - `<db01-secondary>_tempdb`
 - `<db01-secondary>_snapinfo`

14.2 Create Virtual Machines

Create two virtual machines for use with primary and secondary SQL replicas. For optimal performance and availability, each pair of virtual SQL Servers should consist of one VM from the primary infrastructure and one from the secondary.

Solution validation testing was performed using VMs with 12 vCPUs and 8GB RAM to maximize use of both storage and compute resources. Production virtual machines should be sized based on database requirements and include CPU and memory headroom for SnapCenter operations.

14.3 Install Windows Server 2012 R2


1. Use vSphere Web Client to open a remote console session to the VM.
2. Map the Windows Server 2012 R2 ISO to the virtual CD-ROM device of the VM.
3. Power on the VM and install the OS using the best practices as outlined by Microsoft at [https://technet.microsoft.com/en-us/library/jj134246\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/jj134246(v=ws.11).aspx).
4. Update the latest Windows patches, making sure that [Microsoft KB2887595](#) is applied.
5. Add the server to your Active Directory domain.

14.4 Install Windows Failover Clustering

This section describes the steps to install the Windows Failover Cluster on the two virtual machines that will host the SQL Server failover cluster instance later.

To install Windows Failover Clustering on each of the VMs, complete the following steps:

1. Start Windows Server Manager and select Add Roles and Features from the Dashboard page. Use the Add Roles and Features wizard to add the Failover Clustering feature and accept the additional tools to be installed. No other roles or features are required at this time.
2. After the feature installation completes, start the Failover Cluster Manager from the Tools menu of the Windows Server Manager.
3. Select Validate Configuration from the Actions pane on the right.
4. Enter the host names of all SQL Servers to be included in the cluster and add them to the select servers list. Click Next.
5. Leave the Run All Tests option selected and click Next.
6. Review the report created by the Validate a Configuration wizard.



Storage and network warnings are expected at this point in the validation. Database storage disks and the cluster file share witness are created in later steps. Network connectivity is sufficient for basic setup, with redundancy provided by the infrastructure. Additional network interfaces can be configured as needed for customer-specific SQL application requirements.

7. Close the Configuration Validation wizard.

14.5 Create Windows Failover Cluster

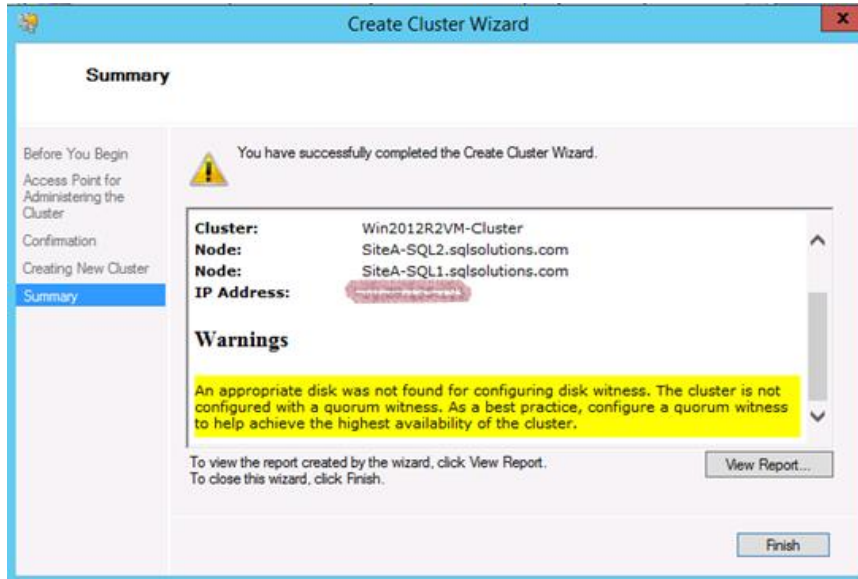
To create a Windows failover cluster, complete the following steps:

1. In the Create Cluster wizard, click Next on the Before You Begin page.
2. In the Access Point for Administering the Cluster page, enter the cluster name, select the appropriate network, provide the cluster IP address, and click Next.
3. In the Confirmation page, verify the cluster details and click Next. Because storage will be added later, the Add All Eligible Storage to the Cluster option should be deselected.

The Creating New Cluster page displays the progress of the cluster installation.


4. After the cluster is successfully created, the Summary page displays the summary of the cluster configuration details. The Summary page displays the warning that An appropriate disk was

not found for configuring disk witness. This warning can be ignored because a file share witness will be configured later.



5. Click Finish.

You can view the cluster configuration details by connecting to the newly created cluster from the Failover Cluster Manager window on the VMs.



14.6 Configure File Share Witness

The Windows failover cluster will be configured using the File Share majority quorum model. The cluster is configured so that only production VMs participate in cluster voting. This approach makes sure that the failure of a DR virtual machine does not affect the cluster status.

To configure the File Share Witness, complete the following steps:

1. In the Failover Cluster Manager page, navigate to Failover Cluster Manager > <Cluster Name>.
2. Right-click the cluster and select More Actions > Configure Custom Quorum Settings. The Configure Cluster Quorum wizard opens.
3. Click Next on the Before You Begin page.
4. In the Select Quorum Configuration Option page, choose the Select the Quorum Witness option and click Next.
5. In the Select Quorum Witness page, select the Configure a File Share Witness option and click OK.
6. In the Configure File Share Witness page, specify the file share that will be used by the file share witness resource and click Next.
7. In the Confirmation page, verify the configuration details and click Next.

The Summary displays the summary of quorum settings that are configured for the cluster.

8. Click Finish.

14.7 Add Data Disks to the Virtual Machine

Datstores for database storage were created in the section “Deploy Production Datstores.” To allocate these datstores to the SQL Server VMs, complete the following steps:

1. Log in to vSphere Web Client.
2. Right-click the SQL Server VM and select Edit Settings.
3. Click the New Device drop-down list at the bottom of the page, select New Hard Disk, and click Add.
4. Expand the New Hard Disk object and enter the following values:
 - a. Enter the hard disk size.
 - b. Select the appropriate datastore.
 - c. Select Thin Provisioned.
5. Repeat steps 3 and 4 for each disk to be added. Each SQL VM should be provisioned with four dedicated disks for database storage in addition to the operating system disk (C:).
6. Mount and format the disks in the SQL Server. Table 18 shows the provisioned disks and mount points for each.

Table 18) SQL Server disks and mount points.

VMDK Name	Description	Disk Type	Location
SQL_Server_x_System.vmdk	SQL Server system database disk		System Databases (C:\MSSQL\System)
SQL_Server_x_TempDB.vmdk	SQL Server temp database disk	Fixed	Temp Database (C:\MSSQL\TempDB)
SQL_Server_x_Data.vmdk	SQL Server database data disk	Fixed	Data files (C:\MSSQL\Data)
SQL_Server_x_Log.vmdk	SQL Server database log disk	Fixed	Logs (C:\MSSQL\Log)
SQL_Server_x_SnapInfo.vmdk	SnapCenter 1.1 SnapInfo log directory	Fixed	SnapInfo for SnapCenter (C:\MSSQL\SnapInfo)

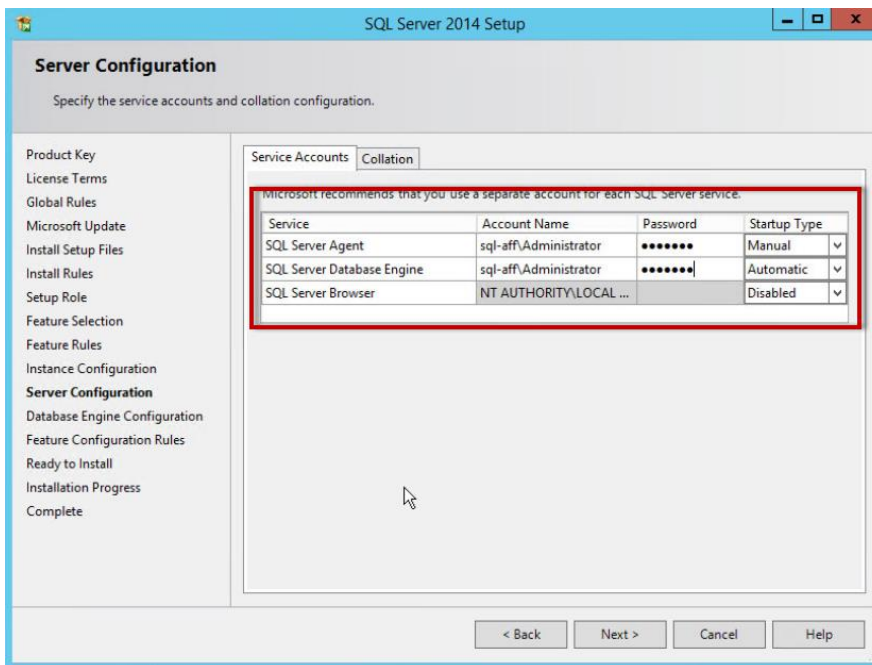
14.8 Install the SQL Server 2014 Software

This section describes the procedure to create a SQL Server cluster instance. This procedure needs to be completed first on one SQL Server VM; the second SQL Server node should then be added to the cluster.

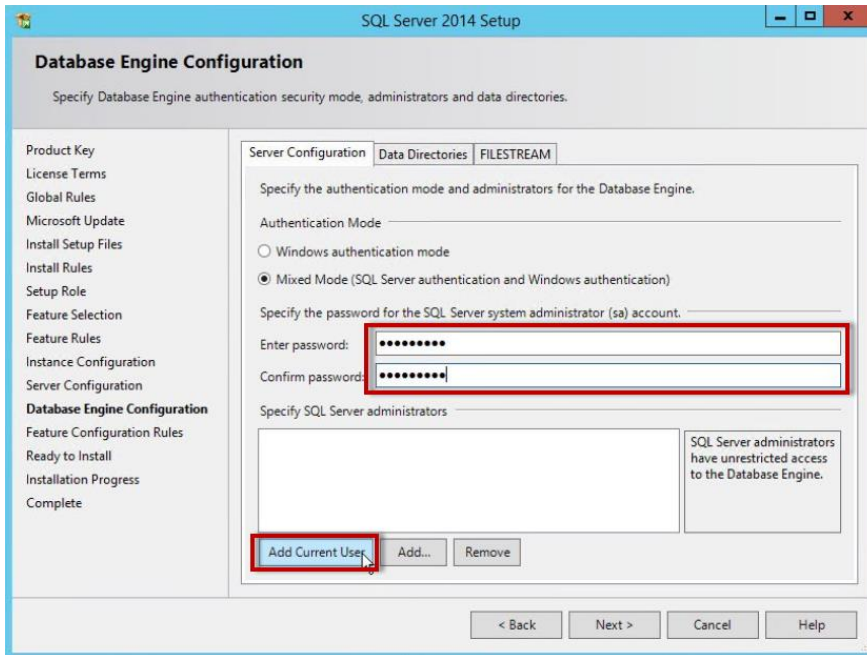
To complete the installation of SQL Server 2014 on each VM and create a SQL Server cluster, complete the following steps:

1. Mount the SQL Server ISO image to the VM.
 - a. Log in to vSphere Web Client.
 - b. Right-click the VM and select Edit Settings.
 - c. Expand the CD/DVD object, select Datastore ISO file, and select the Connected checkbox.
 - d. Browse to the datastore containing the SQL Server 2014 ISO file and select the ISO image to mount.

- e. Click OK to complete the CD/DVD configuration.
2. Inside the virtual machine, double-click the CDROM drive to open the SQL Server Installation Center.
3. Select the Installation link on the left, then select New SQL Server Stand-alone Installation or Add Features to an Existing Installation to start the SQL Server 2014 Setup wizard.
4. In the Product Key page, enter the product key and click Next.
5. Read and accept the EULA and click Next.
6. Optional: The Microsoft Update page appears next if the Microsoft Update checkbox in Control Panel\All Control Panel Items\Windows Update\Change settings is not selected. Selecting the Microsoft Update option will change the computer settings to include the latest updates for Microsoft Windows and SQL Server 2014.
7. On the Product Updates page, the latest available SQL Server product updates are displayed. If no product updates are discovered, SQL Server Setup does not display this page and auto advances to the Install Setup Files page.
8. Review the Install Rules test results and click Next to continue.
9. In the Setup Role page, select SQL Server Feature Installation.
10. Select the Database Engine Services and Management tools for installation. You can install additional features as required.
11. The Feature Rule page displays the rule executions and will automatically advance if all rules pass. Remediate any failures before continuing.
12. In the Instance Configuration page, specify the SQL Server network name and the instance ID and click Next.
13. In the Server Configuration page, specify the service accounts and collation configuration details and click Next.



14. In the Database Engine Configuration page, specify the database engine authentication security mode, administrators, and data directory details. In the Data Directory tab, make sure that the root directory and the temp database directory are set appropriately.



15. The Feature Configuration Rules automatically run the feature configuration rules. Verify the output and click Next.
16. In the Ready to Install page, verify the installation options and click Install to begin the SQL Server Failover Cluster installation.
17. After the installation is complete, verify the installation summary and click Close to close the wizard. Repeat the process on the other nodes in the Availability Group.

14.9 Create AlwaysOn Availability Group

This section describes the procedure to create an AlwaysOn Availability Group (AG). The steps to follow are outlined at <https://blogs.technet.microsoft.com/arnabm/2016/05/14/step-by-step-configuration-manager-site-database-hosted-on-sql-server-alwayson-availability-group/>.

Enable AlwaysOn Availability Groups on Each SQL Server

1. Open the SQL Server Configuration Manager.
2. Double-click SQL Server (MSSQLSERVER) Service to open Properties.
3. Open the AlwaysOn Availability Groups tab.
4. Select the Enable AlwaysOn Availability Groups option.
5. Click OK.
6. Restart the service.

Prepare All Servers for AlwaysOn Availability Group

1. Add the computer account of the primary site server to the administrator's group on each SQL Server.
2. Add the Installation account to the administrators group on each SQL Server.
3. Add the Installation account to the sysadmin role on each SQL Server.

Note: Make sure that the database to be used for the Availability Group has Recovery Mode set to Full.

Create and Configure the AlwaysOn Availability Group

1. Open SQL Management Studio.
2. Right-click AlwaysOn High Availability and select New Availability Group Wizard.
3. Enter the Availability Group name.
4. Select the database to use for the AG after making sure that the database meets prerequisites. Click Next.
5. Add Replicas and make sure that they are all set to Synchronous Commit and Readable Secondary. Set the Automatic Failover based on the requirements of your environment.
6. Click the Endpoints tab and verify that the servers are participating in the AG.
7. Click the Listener tab and create an Availability Group Listener. Click Next.
8. On the Select Your Data Synchronization preference, click Next to complete validation.
9. Confirm selections and click Finish.
The wizard performs a backup and then a restore on your replica servers.
10. Validate the Availability Group in SQL Management Studio.

15 NetApp SnapCenter

SnapCenter includes the SnapCenter Server and the SnapCenter Plug-ins Package for Windows (which includes the SnapCenter Plug-in for Microsoft SQL Server and the SnapCenter Plug-in for Microsoft Windows). SnapCenter also includes the SnapCenter Plug-ins Package for Linux (which includes the SnapCenter Plug-in for Oracle Database and the SnapCenter Plug-in for UNIX). In addition, SnapCenter interacts with VSC to support database backup, restore, recovery, and cloning on Raw Device Mappings (RDMs) and Virtual Machine Disks (VMDKs).

15.1 Install and Configure SnapCenter 1.1

This section describes the installation and configuration of SnapCenter 1.1. You can find detailed information at <http://mysupport.netapp.com/NOW/download/software/snapcenter/1.1/>.

Download SnapCenter 1.1 from

<http://mysupport.netapp.com/NOW/download/software/snapcenter/1.1/download.shtml>.

Create a Windows Server 2012R2 VM, making sure that the VM has available a minimum of 4 CPU cores, 16GB of RAM, and 20GB of space. Although these requirements are more than the minimum requirements detailed in the installation guide, these elements provide a more responsive SnapCenter environment.

The installation in this environment was performed by following the instructions on page 27 of the [SnapCenter Software 1.1 Installation and Setup Guide](#).

Install SnapCenter 1.1


1. Make sure that all prerequisites detailed in the [SnapCenter Software 1.1 Installation and Setup Guide](#) are met.
2. Double-click SnapCenter1.1.exe.
3. Click Next on the Welcome to SnapCenter 1.1 wizard page.
4. On the Network Load Balancing page, leave the settings at the default unless you want to enable a high-availability installation of SC1.1.
5. Set the SnapCenter service credentials.
6. Choose to install SQLExpress or to use an existing version of Microsoft SQL Server for the SnapCenter 1.1 database repository.

- Click Install to begin installation.

Configure SnapCenter 1.1

This section describes the configuration of SnapCenter 1.1 as outlined on page 28 of the [SnapCenter Software 1.1 Installation and Setup Guide](#).

- Log into SnapCenter using the domain admin credentials that were used while setting up SnapCenter 1.1.
- Click Hosts and then click Add.
- In the Host section, enter the host name of one of the servers in the AlwaysOn Availability Group and select Add All the Hosts in the Cluster. Click Next.
- In the Discover Plug-ins section, no plug-ins will be found. Click Next.
- In the Plug-ins section, select SnapCenter Plug-ins 1.1 for Windows Package. Select Microsoft Windows and Microsoft SQL Server. Make sure that Run As Name is selected and the licensing model is set to per-storage systems. Click Next.
- Validate the installation parameters and click Finish to install the SnapCenter 1.1 agents on the servers in the Availability Group.
- To monitor the installation, click Monitor, select the installation job, and click Details.

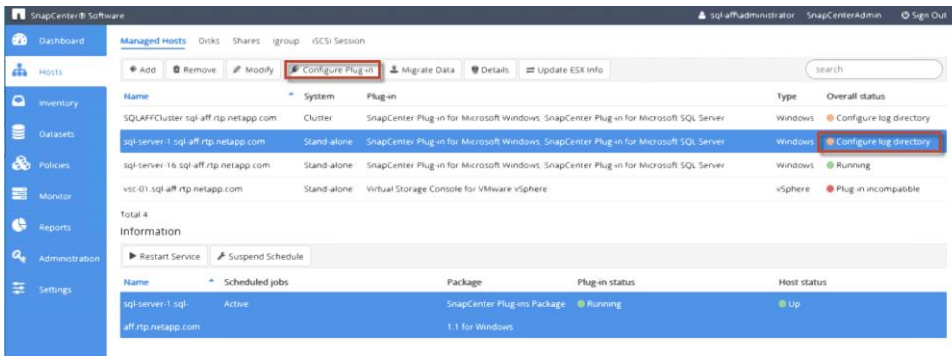


Configure SVM Connection Settings

- In SnapCenter, click Settings > New.
- Enter the SVM name and login credentials. Select the communications protocol and the preferred IP address if it is different from the DNS name of your SVM. Click OK.

Configure Host Plug-Ins

- In SnapCenter, click Hosts in the left-hand pane. Select the cluster and click Configure Plug-In.




- In the Log Directory section, click Browse and locate the SnapInfo directory. Repeat for all members of the AlwaysOn Availability Group and then click Next.
- In the Verification Server section, select the verification server and click Next.

4. In the Summary section, validate the configuration and click Finish.

Inventory Databases

1. In SnapCenter, click Inventory in the left-hand pane. Select the host server to be configured and select the type of resource to be added to the SnapCenter inventory. Repeat for all servers.



Create a Dataset for Backup

1. In SnapCenter, click Datasets from the left-hand pane. Click New and then select Backup Dataset.
2. Assign a name, policy, and backup credentials to the dataset and then click Next.
3. Set Resource Type to Availability Groups. Select Availability Group Resources in the left window and click the arrow in the middle to move it to the Selected Resources window on the right. Click Next.
4. Select the verification server and then click Next.
5. Set the notification setting to meet the requirements of your environment. Then click Next.
6. Review your settings and select Finish to create the backup dataset.

16 Solution Verification

This reference architecture is based on a standard FlexPod infrastructure hosting Microsoft SQL Server 2014 databases. Cisco Nexus 9396PX switches in NX-OS mode connect the Cisco UCS compute nodes and NetApp storage arrays to the Ethernet network for client and NAS storage access. The storage arrays are connected directly to the Cisco UCS Fabric Interconnects for FCoE storage access. Cisco UCS B200-M4 blades running VMware vSphere ESXi 6.0 host Microsoft Windows Server 2012 VMs and use dedicated FCoE LUNs for individual database instances. Specific details of the configuration can be found in the section “Technology Requirements.”

To provide enterprise-class performance, management, and reliability, this solution was verified with the following test cases:

- The performance of the primary databases was validated with the Microsoft TPC-E toolkit and the TPC-E workload configured for Microsoft SQL Server 2014 AlwaysOn Availability Groups. Successful validation required that the primary Microsoft SQL Server 2014 databases deliver 150,000 to 200,000 I/O operations per second with an average read latency below 1ms.
- Performance of the secondary database configuration was validated after failover from the primary database instance. Successful validation required that the primary Microsoft SQL Server 2014 databases deliver 150,000 to 200,000 I/O operations per second with an average read latency below 1ms.
- SnapCenter was used to create application-integrated Snapshot copies of secondary database instances at the same time that primary instances were in operation.
- SnapCenter was used to recover databases from Snapshot copies after intentionally introduced database corruption.

- SnapCenter was used to clone databases and to mount to alternate servers for validation or test/dev purposes.
- The infrastructure was subjected to several hardware resiliency tests while under the same workload used for performance validation. This step was taken to make sure that the workload would continue to run during the following failure and maintenance scenarios:
 - Disconnection of an FCoE link between the storage array and the fabric interconnect
 - Disconnection of a 10GbE link between the storage array and the fabric interconnect
 - Failover of each storage controller and subsequent takeover by the partner controller
 - Reboot of the primary fabric interconnect

During solution testing, Cisco UCS blade servers hosted the infrastructure and SQL Server 2014 VMs. The database and infrastructure servers were hosted on discrete compute resources so that the workload to the NetApp AFF system could be precisely measured. Separating production VMs from the infrastructure VMs is a NetApp and industry best practice because noisy neighbors or bully VMs can affect the infrastructure. This result can have a negative effect on all users, applications, and performance results.

Detailed results for the validation testing can be found in [NVA-0025-DESIGN: FlexPod Datacenter with SQL Server AlwaysOn Availability Groups and AFF](#).

16.1 Performance Validation

The performance requirement for this solution was the delivery of 150,000 to 200,000 IOPS with submillisecond read latency using SQL Server 2014 in an Availability Group configuration. The following sections describe the methodology and design considerations used to test the AFF8080 EX running a standard SQL Server workload.

Database Configuration

For this validation, a total of six 2.5TB databases hosted the simulated OLTP environment. Each storage system controller had a single data aggregate of 22 800GB SSDs, as shown in Table 19. Databases were balanced across the controllers by placing odd-numbered SQL Servers on controller 1 and even-numbered servers on controller 2.

The database layout shown in Table 19 was repeated for each of the six SQL Server databases. For each SQL Server database, the data files, log files, and tempdb were each contained in a separate LUN within a separate volume. LUNs were allocated to the vSphere ESXi hosts and formatted as vSphere datastores by using VMFS. Virtual disks were then allocated to each SQL Server VM and formatted with default settings.

Table 19) Database LUN and volume configuration.

LUN Name	Datastore Name	VMDK Size	File Name	File Size
/vol/sqlX_data/sqlX_data	sqlX_data	2.5TB	MSSQL_tpce_root.mdf	8MB
			Fixed_1.ndf	5MB
			Growing_1.ndf	580GB
			Growing_2.ndf	580GB
			Growing_3.ndf	580GB
			Growing_3.ndf	580GB

LUN Name	Datastore Name	VMDK Size	File Name	File Size
			Scaling_1.ndf	20GB
			Scaling_2.ndf	20GB
			Scaling_3.ndf	20GB
			Scaling_4.ndf	20GB
/vol/sqlX_log/sqlX_log	sqlX_log	500GB	TPCE_Log.ldf	200GB
/vol/sqlX_tempdb/sqlX_tempdb	sqlX_tempdb	300GB	tempdev01.ndf	8MB
			tempdev02.ndf	10GB
			tempdev03.ndf	10GB
			tempdev04.ndf	10GB
			tempdev05.ndf	10GB
			tempdev06.ndf	10GB
			tempdev07.ndf	10GB
			tempdev08.ndf	10GB

Test Methodology

For this validation, the Microsoft TPC-E toolkit was used to generate an industry-standard OLTP warehouse transaction workload against the SQL Server 2014 test configuration. This process generated a workload of approximately 90% reads and 10% writes against the SQL Server databases in the test configuration. The goal of these tests was not to measure the maximum performance of the configuration. Instead, the goal was to validate that the performance available at an acceptable read latency of approximately 1ms was within the limits specified for the solution of 150,000 to 200,000 IOPS.

For these tests, a total of six SQL Server 2014 VMs were hosted on three ESXi servers. Each ESXi server was connected through FCoE to the AFF8080 EX storage array, as shown in Figure 3. Four LUNs were provisioned for each database server and formatted as VMFS datastores. Then a VMDK was provisioned from each datastore to the appropriate SQL Server VM. Each server created a database of 2.6TB total capacity, for a total size requirement of 16TB across all six primary SQL Server databases.

With these six SQL Server databases and the Microsoft TPC-E load generator, performance of the standalone databases was measured under successively heavier workloads until the read latency as observed by the database server exceeded 1ms. Then the workload was reduced until the read latency was under 1ms, and this configuration was used for all use-case testing. Each performance test was run for a total of 2.5 hours so that a steady state was achieved. The IOPS and latency results reported are the aggregate across all six primary or secondary database servers.

Microsoft performance-monitoring tools captured IOPS and latency data from each of the six SQL Server database servers. This step made sure that the observed performance met the validation requirements described previously for each of the tested configurations.

Test Results

The results of the testing are listed in Table 20. In all cases, the performance observed at the database level met or exceeded the requirements set for validation. That performance demonstrates that a SQL

Server environment based on FlexPod can deliver enterprise-class throughput and latencies for the most demanding applications when configured as an Availability Group.

Table 20 shows the performance of the six primary databases configured in AlwaysOn Availability Groups and the secondary databases running after failover from the primary replica while still in a synchronized state.

Table 20) Validated database performance.

	AG Nominal	After AG Failover
IOPS	211,789	205,776
Read latency	0.97ms	0.94ms

In addition to the performance data described previously, we observed VM, ESXi host, and storage controller CPU utilization during the testing. In all test configurations, we observed CPU use between 60% and 65% on the database server instances and storage controllers and 70% on ESXi hosts. This use indicates that this specific configuration was well balanced between compute and storage resources. Both resources had adequate headroom to handle additional workloads.

16.2 SnapCenter Validation

To demonstrate the data management capabilities of this solution, NetApp SnapCenter created Snapshot copies of running databases, recovered databases after simulated corruption, and cloned databases for data verification or test/dev purposes.

SnapCenter 1.1 Prerequisites

SnapCenter operations on each SQL Server require memory beyond the requirements of the SQL database instances. For SnapCenter processes to operate efficiently, allocate 4GB of memory to each SQL Server and to the memory requirements of the databases.

SnapCenter Server requires at least two vCPUs and 8GB RAM. Additional CPU and memory resources can improve performance depending on workload and job schedules.

All Microsoft Windows servers should have current patch updates applied. Specifically, Microsoft KB2887595 must be applied for successful operation.

Database Backup

For this solution validation, backups were performed against both primary and secondary replicas of the availability group to verify the functionality of both sites. SnapCenter 1.1 does not currently support one-button backup of individual availability group databases. To back up a database in an Availability Group, create a backup dataset that includes one or more databases in the Availability Group. The backup operations can then be performed using the backup dataset.

Database Cloning

Cloning of databases provides a means to rapidly provision a copy of production data for further processing in analytics, test, or development environments. SnapCenter 1.1 can create clones from either the active database or an existing backup set. For validation purposes, clones were created using both methods.

Database Restore

SnapCenter 1.1 requires that a corrupted database be removed from the Availability Group before proceeding with the restore operation. After the database is removed from the AG configuration,

SnapCenter restores the database using the NORECOVERY option, which restores the data files but does not replay any existing log files. When the storage recovery is complete, the database logs from the surviving database instance are replayed to synchronize the restored copy and rejoin the restored database replica to the Availability Group.

16.3 Resiliency Validation

We then demonstrated the enterprise-class resiliency of this FlexPod solution. To do so, we induced various failure scenarios into the system while the database servers and storage were subjected to the identical workloads used for performance validation. The goal was to subject the FlexPod infrastructure to a heavy workload and measure the effect of the specific failure scenario on the overall performance and stability of the system. To pass each of these tests, the database and storage had to continue to serve I/O during the event. Some of these tests were extremely disruptive. Therefore, minimal drops in overall performance were considered acceptable as long as the overall system functioned nominally and performance returned to prefailure levels after the failure was corrected.

Link Failure Tests

These tests intentionally disconnected and then reconnected a single FCoE and 10GbE link in the environment while under a heavy OLTP workload of approximately 200,000 IOPS. The workload was allowed to run for 5 minutes to make sure of a steady state before introducing the first failure. The workload was then allowed to run for 10 minutes between each failure before the failure was corrected.

Because primary storage access in this solution uses FCoE, there is a brief drop in IOPS when the FCoE link is disconnected. After the host multipath driver activates alternate paths, the observed throughput returns to nominal values in less than 30 seconds.

No noticeable performance impact was observed during removal or replacement of the 10GbE link. In this solution, 10GbE is used for synchronous replication of the SQL Server databases in the AlwaysOn Availability Group. Although the link failure might have affected latency of the replication, the effect was not enough to lower the overall throughput of the primary database.

Storage Controller Failover Tests

These tests intentionally induced a catastrophic failure in the one controller of the AFF8080 EX storage system while under the 200,000 IOPS OLTP workload. The result of this failure is that the surviving storage controller is required to service the entire workload previously handled by both.

For these tests, the following methodology was used:

1. The OLTP workload was started and allowed to run on both controllers for 10 minutes.
2. With the forced failover of controller 2, controller 1 services the entire workload. During this time, controller 2 reboots to a preoperational state and then enters a wait period. After the wait period, the system starts an automatic giveback and controller 2 resumes normal operations approximately 15 minutes later.
3. The workload was allowed to run for 10 minutes on both controllers again to ensure nominal operation.
4. Controller 1 was failed over; controller 2 then serviced the entire workload. Controller 1 resumed operation approximately 15 minutes later.
5. The workload was allowed to run for another 10 minutes again to ensure nominal operation.

The following results were observed:

6. There was a noticeable reduction in IOPS at the moment each controller went offline. This reduction in IOPS is the result of host failover driver timeout values and represents the time necessary for the host to recognize the path failure and redirect I/O to the surviving controller. I/O did not stop

completely because half of the database instances were not affected by each controller failure and so did not need to wait for path failover timeouts.

7. The performance recovered when the controller was offline but was on average roughly 10% lower than when both controllers were online. Storage controller CPU utilization was approximately 90%, with latencies less than 10% higher than those during normal operations.
8. After restoring each controller, performance quickly returned to prefailure levels.

Cisco UCS Fabric Interconnect Failover Test

This test intentionally rebooted the primary Cisco UCS fabric interconnect while it ran the 200,000 IOPS OLTP workload. During the failure, all I/O was forced to traverse the secondary fabric interconnect. For this test, the workload was allowed to run for five minutes at a steady state before the fabric interconnect was rebooted. During this test, a brief drop in IOPS occurred before a quick recovery to prefailure levels was observed and there was no apparent effect when the fabric interconnect resumed operation.

17 Conclusion

FlexPod Datacenter is the optimal infrastructure foundation on which to deploy Microsoft SQL Server 2014. Cisco and NetApp created a platform that is both flexible and scalable for multiple use cases and designs. The flexibility and scalability of FlexPod enable you to start out with a right-sized infrastructure that can grow with and adapt to your evolving business requirements. The system can grow from single SQL Server deployments to consolidated database farms.

In the verification tests of this reference architecture, the total IOPS and latency as measured at the database met performance expectations. The NetApp FlexPod Datacenter configuration with AFF reached over 200,000 combined IOPS and read latencies of less than 1ms while averaging 60% CPU utilization during most operations. In addition, this verification demonstrated that the solution as documented could maintain these performance levels even in the event of a storage, host, or network failure.

Acknowledgements

This document is a result of the work, prior documentation, and assistance of Lindsey Street, David Arnette, Jens Dickmeis, Pat Sinthusan, Atul Bhalodia, Chris Lemmons, Troy Mangum, and the Converged Infrastructure team at NetApp.

References

This section provides links to additional information and reference material for the subjects contained in this document.

Cisco Unified Computing System

The following links provide additional information about Cisco Unified Computing System:

- Cisco Design Zone for FlexPod
http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/landing_flexpod.html
- Cisco Unified Computing System
<http://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>
- Cisco UCS 6200 Series Fabric Interconnects
<http://www.cisco.com/en/US/products/ps11544/index.html>

- Cisco UCS 5100 Series Blade Server Chassis
<http://www.cisco.com/en/US/products/ps10279/index.html>
- Cisco UCS B-Series Blade Servers
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>
- Cisco UCS Adapters
http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html
- Cisco UCS Manager
<http://www.cisco.com/en/US/products/ps10281/index.html>

Cisco Nexus Networking

The following links provide additional information about Cisco Nexus 9000 Series switches:

- Cisco Nexus 9000 Series Switches
<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>
- Cisco Nexus 9396PX Switch Information
<http://www.cisco.com/c/en/us/support/switches/nexus-9396px-switch/model.html>

NetApp FAS Storage

The following links provide additional information about NetApp FAS storage:

- Clustered Data ONTAP 8.3.1 Documentation
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=62175&language=en-US>
- TR-3982: NetApp Clustered Data ONTAP 8.3
<http://www.netapp.com/us/media/tr-3982.pdf>
- Guest OS Tunings for a VMware vSphere Environment (KB)
<https://kb.netapp.com/support/index?page=content&id=3013622>
- TR-4403: NetApp AFF8080 EX Performance and Server Consolidation with Microsoft SQL Server 2014
<http://www.netapp.com/us/media/tr-4403.pdf>

NetApp SnapCenter

The following links provide additional information about NetApp SnapCenter software:

SnapCenter 1.1 Getting Started Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2440145

SnapCenter 1.1 Installation and Setup Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2440193

SnapCenter 1.1 Administration Guide
https://library.netapp.com/ecm/ecm_download_file/ECMLP2439718

VMware vSphere

The following links provide additional information about VMware vSphere:

- VMware vSphere Documentation Center
<http://pubs.vmware.com/vsphere-60/index.jsp>
- SQL Server on VMware Best Practices Guide
https://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf

Interoperability Matrices

The following links provide information about interoperability tools:

- Cisco UCS Hardware and Software Interoperability Tool
<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>
- NetApp Interoperability Matrix Tool
<http://support.netapp.com/matrix>
- VMware Compatibility Guide
<http://www.vmware.com/resources/compatibility>

Version History

Version	Date	Document Version History
Version 1.0	December 2016	Engineering content creation

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE 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 NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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 AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.
NVA-0025-DEPLOY