



Technical Report

FlexPod Solution Delivery Guide

Jimmie Cox, NetApp
July 2019 | TR-4790

Abstract

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information toward shared infrastructures and virtualized environments, and will eventually move toward cloud computing to increase agility and reduce costs. This transformation appears daunting and complex because companies must address both organizational and technical resistance to this new IT model. NetApp® partnered with Cisco to develop the FlexPod® platform to address these virtualization needs and to simplify the evolution toward shared, virtualized, and cloud infrastructures.

TABLE OF CONTENTS

1	Solution Overview	4
1.1	Solution Technology	4
1.2	Top Ten Benefits of FlexPod.....	4
1.3	Use Case Summary	5
1.4	Target Audience.....	5
2	Technology Requirements	6
2.1	Hardware Requirements	6
2.2	Software Requirements	7
2.3	Cisco Software Licensing Options	7
2.4	NetApp Software Licensing.....	8
3	Deployment Procedures	8
3.1	NetApp Storage Configuration	8
3.2	Cisco UCS Server Configuration.....	25
3.3	Cisco Nexus Storage Networking Configuration	49
4	Extensibility.....	55
4.1	Multitenant	55
4.2	Architecture Overview	56
4.3	FlexPod Scale Up and Scale Out.....	57
4.4	NetApp Cloud Volumes ONTAP	58
4.5	FlexPod Managed Private Cloud	59
5	Conclusion	61
	Where to Find Additional Information	62
	Version History	62
	LIST OF TABLES	
	Table 1) Hardware requirement.....	7
	Table 2) Cisco software licenses option.	8
	Table 3) NetApp software licensing options.	8
	Table 4) Cluster details for the ONTAP software configuration.	10
	Table 5) Cluster details for the cluster-join operation.	13
	Table 6) iSCSI LIFs for iSCSI IQN.	49
	Table 7) vNIC iSCSI IQNs for fabric A and fabric B.....	49
	LIST OF FIGURES	
	Figure 1) FlexPod layout.	5

Figure 2) Cabling diagram for the FlexPod use case on ONTAP.	9
Figure 3) FlexPod multitenant capability.	56
Figure 4) Architecture overview.	57
Figure 5) Example of scaling a FlexPod configuration.	58
Figure 6) Place cloud at the source of demand with a FlexPod managed private cloud.	61

1 Solution Overview

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information toward shared infrastructures and virtualized environments and eventually will move toward cloud computing to increase agility and reduce costs. This transformation appears daunting and complex because companies must address both organizational and technical resistance to this new IT model. Cisco and NetApp developed FlexPod to address these virtualization needs and to simplify the evolution toward shared, virtualized, and cloud infrastructures.

FlexPod is a validated infrastructure solution built on the Cisco Unified Computing System (Cisco UCS), Cisco Nexus data center switches, NetApp FAS storage components, and software from a range of partners. FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. A FlexPod design provides a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different business solutions. This document describes how to build several different solutions on top of FlexPod. Cisco and NetApp developed FlexPod as a platform that can address current data center needs and simplify the evolution toward an IT-as-a-service (ITaaS) infrastructure. FlexPod serves as a base infrastructure layer for a variety of IT solutions.

1.1 Solution Technology

The FlexPod platform is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp® storage controllers (FAS or AFF). FlexPod is a suitable platform for running a variety of virtualization hypervisors as well as bare-metal operating systems and enterprise workloads. FlexPod delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements.


1.2 Top Ten Benefits of FlexPod

The FlexPod architecture provides the following benefits:

- Innovative, reliable technology
- Integrated, single SKU solution
- Robust partner ecosystem
- Ability to cut across storage silos
- Support for multiple protocols on a single platform
- Validated designs
- Storage efficiency
- High availability
- Performance
- Flexibility

Figure 1 shows the technical components of the solution.

Figure 1) FlexPod layout.



1.3 Use Case Summary

The base hardware, which is detailed in the FlexPod technical specifications, applies to the following use cases:

- Two Cisco Nexus switches
- Two Cisco UCS 6248UP fabric interconnects
- Cisco UCS B-Series blades with two fabric extenders per chassis and/or Cisco UCS C-Series rack servers
- NetApp FAS80xx (highly available configuration with dual controllers in a single enclosure)

In Figure 1, storage is provided by a NetApp FAS3250 with accompanying disk shelves. All systems and fabric links feature redundancy, providing end-to-end high availability (HA). Although this is the default base design, each of the components can be scaled flexibly to support the customer's specific business requirements. For example, more (or different) UCS blades and chassis can be deployed to increase compute capacity. Additional disk shelves can be deployed to increase capacity and improve input/output (I/O) throughput. And special hardware or software features can be added to introduce new features, such as NetApp Flash Cache™ intelligent caching for deduplication-aware caching.

1.4 Target Audience

This document describes the basic architecture of FlexPod as well as several solutions that can be built on FlexPod. The target audience for this document includes, but is not limited to, sales engineers, field

consultants, professional services personnel, IT managers, partner engineering personnel, and customers who want to deploy solutions on top of a FlexPod configuration.

2 Technology Requirements

This section provides the FlexPod technology requirements (hardware, software, and licensing).

2.1 Hardware Requirements

A FlexPod Datacenter configuration has minimum hardware requirements, including, but not limited to, switches, Fabric Interconnects, servers, and NetApp storage controllers. You must use Cisco UCS servers. Both C-Series and B-Series servers have been used in the validated designs. Cisco Nexus Fabric Extenders (FEXs) are optional with C-Series servers. A FlexPod configuration has the following minimum hardware requirements:

- Two Cisco Nexus switches in a redundant configuration. This configuration can consist of two redundant switches from the Cisco Nexus 5000, 7000, or 9000 Series. The two switches should be of the same model and should be configured in the same mode of operation.
Note: If you are deploying an ACI architecture, you must observe the following additional requirements:
 - Deploy the Cisco Nexus 9000 Series Switches in a leaf-spine topology.
 - Use three Cisco Application Policy Infrastructure Controllers (APICs).
- Two Cisco UCS 6200 or 6300 Series Fabric Interconnects in a redundant configuration.
- Cisco UCS servers:
 - If the solution uses B-Series servers, one Cisco UCS 5108 B-Series Blade Server Chassis plus two Cisco UCS B-Series Blade Servers plus two 2104, 2204, 2208, or 2304 I/O modules (IOMs).
 - If the solution uses C-Series servers, two Cisco UCS C-Series rack-mount servers.**Note:** For larger deployments of Cisco UCS C-Series rack-mount servers, you might want a pair of 2232PP FEX modules. However, the 2232PP is not a hardware requirement.
- Two NetApp storage controllers in an HA pair configuration:
Note: This configuration can consist of products from the NetApp FAS2000, FAS2500, FAS2650, FAS2700, FAS3000, FAS3000s, FAS3100, FAS3200, FAS6000, FAS6000s, FAS6200, FAS8000, FAS8200, or FAS9000 series of controllers. Or it can consist of products from the NetApp AFF A200, AFF A300, AFF A700, AFF A700s, AFF A800, or AFF8000 series of controllers.
 - The HA configuration requires two redundant interfaces per controller for data access; the interfaces can be FCoE, FC, or 10Gb Ethernet (10GbE).
 - If the solution uses NetApp ONTAP® data management software, a cluster interconnect topology that is approved by NetApp is required.
 - If the solution uses ONTAP, at least two additional 10GbE ports per controller are required for data access.
 - For ONTAP clusters with two nodes, a two-node switchless cluster can be configured.
 - For ONTAP clusters with more than two nodes, a pair of cluster interconnect switches are required.
- One NetApp disk shelf with any supported disk type.

Table 1 lists the hardware components that are required to implement the use case.

Table 1) Hardware requirement.

Storage Controller Model	Minimum Data ONTAP Version	Maximum Data ONTAP Version
FAS2600 Series	9.1RC1	9.4.x
FAS2700 Series	9.4RC1	9.4.x
FAS8200	9.1RC1	9.4.x
FAS9000	9.1RC2	9.4.x
AFF-A200	9.1RC2	9.4.x
AFF-A220	9.4RC1	9.4.x
AFF-A300	9.1RC1	9.4.x
AFF-A700	9.1RC2	9.4.x
AFF-A700s	9.1	9.4.x
AFF-A800	9.4RC1	9.4.x

2.2 Software Requirements

A FlexPod configuration has the following minimum software requirements:

- NetApp ONTAP:
 - ONTAP software version requires ONTAP 9.1 or later
 - Cisco UCS Manager releases:
 - Cisco UCS 6200 Series Fabric Interconnect—2.2(8a)
 - Cisco UCS 6332, 6332-16UP Fabric Interconnect—3.1(1e)
 - For Cisco Nexus 5000 Series Switches, Cisco NX-OS software release 5.0(3)N1(1c) or later, including NX-OS 5.1.x
 - For Cisco Nexus 7000 Series Switches:
 - The 4-slot chassis requires Cisco NX-OS software release 6.1(2) or later
 - The 9-slot chassis requires Cisco NX-OS software release 5.2 or later
 - The 10-slot chassis requires Cisco NX-OS software release 4.0 or later
 - The 18-slot chassis requires Cisco NX-OS software release 4.1 or later
 - For Cisco Nexus 9000 Series Switches, Cisco NX-OS software release 6.1(2) or later
- Note:** The software that is used in a FlexPod configuration must be listed and supported in the NetApp IMT. Some features might require more recent releases of the software than the ones that are listed.

2.3 Cisco Software Licensing Options

To enable storage protocols on the Cisco Nexus switches, licenses are required. The Cisco Nexus 5000 and 7000 Series of switches all require a storage services license to enable the FC or FCoE protocol for SAN boot implementations. The Cisco Nexus 9000 Series Switches currently do not support FC or FCoE.

Note: The required licenses and the part numbers for those licenses vary depending on the options that you select for each component of the FlexPod solution. For example, software license part numbers vary depending on the number of ports and which Cisco Nexus 5000 or 7000 Series Switches you choose. Consult your sales representative for the exact part numbers. Table 2 lists the Cisco software licensing options.

Table 2 lists the licenses components that are required to implement the use case.

Table 2) Cisco software licenses option.

Cisco Software Licenses	Part Number
Cisco Nexus 5500 Storage License, 8-, 48-, and 96-port	N55-8P-SSK9/N55-48PSSK9/N55-96P-SSK9
Cisco Nexus 5010/5020 Storage Protocols License	N5010-SSK9/N5020SSK9
Cisco Nexus 5600 Storage Protocols License	N56-16p-SSK9/N567272P-SSK9/N56128128P-SSK9
Cisco Nexus 7000 Storage Enterprise License	N7K-SAN1K9
Cisco Nexus 9000 Enterprise Services License	N95-LAN1K9/N93LAN1K9

2.4 NetApp Software Licensing

Table 3 lists the NetApp software licensing options that are available for the FlexPod Datacenter architecture. NetApp software is licensed at the FAS and AFF controller level.

Table 3) NetApp software licensing options.

NetApp Software Licenses	Part Number
SW, Complete BNDL (Controller), -C	SW-8XXX-COMP-BNDL-C
SW, ONTAP Essentials (Controller), -C	SW-8XXX-ONTAP9-C

3 Deployment Procedures

The procedures in the following sections must be completed in a sequential manner to deploy this FlexPod Datacenter solution.

Deploying the solution involves the following tasks:

- NetApp storage configuration
- Cisco UCS server configuration
- Cisco Nexus storage networking configuration

3.1 NetApp Storage Configuration

AFF Series Controller

For instructions on the physical installation of AFF controllers, follow the procedures in the AFF Series documentation on the [NetApp Support site](#). When planning the physical location of the storage systems, refer to the [NetApp Hardware Universe](#).

NetApp Hardware Universe

The NetApp Hardware Universe is an online application that provides information about supported hardware and software components for specific ONTAP versions. This tool provides configuration information for all NetApp storage appliances that are currently supported by the ONTAP software. It can also compare component compatibilities.

To verify configuration information in the Hardware Universe, complete the following steps:


1. Access the [Hardware Universe site](#) to verify that the hardware and software components are supported with the version of ONTAP that you plan to install.
2. Click the Platforms tab to view the compatibility between ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, click the Compare Storage Systems tab to compare components by storage appliance.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. Visit the [NetApp Support site](#) to view a complete list of supported disk shelves. This solution is built on DS224C disk shelves with SSDs. These disks provide the highest level of performance available. When using SAS disk shelves with NetApp storage controllers, refer to the SAS Disk Shelves Universal SAS and ACP Cabling Guide for information about cabling guidelines.

Figure 2 illustrates the cabling diagram for this FlexPod use case on ONTAP.

Figure 2) Cabling diagram for the FlexPod use case on ONTAP.



ONTAP

This procedure assumes that the storage system has been installed and cabled and is ready for setup. For detailed information about storage system installation, see the preceding resources.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the ONTAP 9.1 Software Setup Guide to learn about the information required to configure ONTAP. Table 4 lists the information that you need to configure two ONTAP nodes. You should customize the cluster detail values with the information that is applicable to your deployment (Table 4).

Table 4) Cluster details for the ONTAP software configuration.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node 01 IP address	<<var_node01_mgmt_ip>>
Cluster node 01 netmask	<<var_node01_mgmt_mask>>
Cluster node 01 gateway	<<var_node01_mgmt_gateway>>
Cluster node 01 service processor IP address	<<var_node01_sp_ip>>
Cluster node 01 service processor netmask	<<var_node01_sp_netmask>>
Cluster node 01 service processor gateway	<<var_node01_sp_gateway>>
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>
Cluster node 02 service processor IP address	<<var_node02_sp_ip>>
Cluster node 02 service processor netmask	<<var_node02_sp_netmask>>
Cluster node 02 service processor gateway	<<var_node02_sp_gateway>>
Cluster password	<<var_password>>
Cluster DNS domain name	<<var_dns_domain_name>>
Nameserver IP	<<var_nameserver_ip>>
Controller location	<<var_node_location>>
Cluster node 01 name	<<var_node01>>
Cluster node 02 name	<<var_node02>>
Cluster node 01 aggregate name	<<var_node01_rootaggrname>>

Set Up Node

Before you start the process of creating a storage cluster, you need to set up the individual nodes, which involves enabling the NetApp AutoSupport® remote support diagnostics system, assigning the node management IP addresses, and so on.

1. To perform the node setup, connect to the storage cluster node 01 console port. Console settings are:

- Baud rate: 115200
- Data bits: 8
- Parity: none
- Stop bit: 1
- Flow control: none

2. Enable AutoSupport on the node.

```
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
```

3. Assign the node management IP address, netmask, and gateway.

```
Enter the node management interface port [e0M]: e0M 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>>
```

4. After the node management IP is assigned, press Ctrl+C to get out of the cluster setup, log in to the shell, and set the storage failover mode to HA.

```
login: admin ***** * This is a serial console
session. Output from this * * session is mirrored on the SP console session. *
*****
::> storage failover modify -mode ha Mode is already set to HA.
::> system node reboot Warning: Are you sure you want to reboot node "localhost"? {y|n}: y
```

5. After the node reboots, 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 [10.61.184.218]: Enter
Enter the node management interface netmask [255.255.255.0]: Enter Enter the node management
interface default gateway [10.61.184.1]: 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: 10.61.184.218.
Alternatively, you can use the "cluster setup" command to configure the cluster.
Wed Jul 13 13:10:49 UTC 2016 login:
```

6. Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

In the ONTAP data management software, the first node in a cluster performs the cluster-create operation. All other nodes perform a cluster-join operation. The first node in the cluster is considered node 01. Use the values from Table 4 to complete the configuration of the cluster and each node.

To create a cluster on node 01, complete the following steps:

1. Keep using the console connection that you have connected to the storage cluster node 01.
2. At the login prompt, enter admin and run the following command:

```
login: admin *****
* This is a serial console session. Output from this *
* session is mirrored on the SP console session.      *
*****
:~> cluster setup
```

3. The Cluster Setup wizard starts on the console.

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

4. Run the following command to create a new cluster:

```
create
```

5. Enter no for the single-node cluster option.

```
Do you intend for this node to be used as a single node cluster?
{yes, no} [no]: no
```

6. Enter no for the option to use network switches for the cluster network.

```
Will the cluster network be configured to use network switches? [yes]:no
```

```
Existing cluster interface configuration found:
Port      MTU      IP              Netmask
e0a       9000     169.254.122.114 255.255.0.0
e0b       9000     169.254.124.58  255.255.0.0
Do you want to use this configuration? {yes, no} [yes]:no
System Defaults:
Private cluster network ports [e0a,e0b].
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]: no
Enter the cluster administrator's (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
List the private cluster network ports [e0a,e0b]: Enter Enter the cluster ports' MTU size [9000]:
Enter Enter the cluster network netmask [255.255.0.0]: Enter
```

7. The system defaults are displayed. Enter no for the option to use the system defaults. Follow these prompts to configure the cluster ports.

```
Existing cluster interface configuration found:
Port      MTU      IP              Netmask
e0a       9000     169.254.122.114 255.255.0.0
e0b       9000     169.254.124.58  255.255.0.0
Do you want to use this configuration? {yes, no} [yes]:no
System Defaults:
Private cluster network ports [e0a,e0b].
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]: no
Enter the cluster administrator's (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
List the private cluster network ports [e0a,e0b]: Enter
Enter the cluster ports' MTU size [9000]: Enter
Enter the cluster network netmask [255.255.0.0]: Enter
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0a [169.254.28.157]: Enter
```

```
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0b [169.254.32.222]: Enter
```

8. Use the information in Table 4 to create a cluster.

```
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_nfs_license>>
Enter an additional license key []:<<var_iscsi_license>>
```

Note: The cluster-create process can take a minute or two.

Note: For this design, we need to provide additional licenses for NetApp SnapRestore® data recovery software, NetApp FlexClone® data replication software, and the NetApp SnapManager suite.

```
Enter the cluster management interface port [e0c]: e0c
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>>
```

9. 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 server IP address, separate them with commas.

The cluster-create operation is done. Follow the steps in the next subsection to join node 02 to the cluster that we just created.

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. Table 5 lists the cluster network information required for joining node 02 to the existing cluster. You should customize the cluster detail values with the information that is applicable to your deployment.

Table 5) Cluster details for the cluster-join operation.

Cluster Details	Cluster Detail Value
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>
Cluster node 02 service processor IP address	<<var_node02_sp_ip>>
Cluster node 02 service processor netmask	<<var_node02_sp_netmask>>
Cluster node 02 service processor gateway	<<var_node02_sp_gateway>>

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

1. At the login prompt, enter admin and run the following command:

```
login: admin
*****
* This is a serial console session. Output from this *
* session is mirrored on the SP console session.    *
*****
::> cluster setup
```

2. The Cluster Setup wizard starts on the console.

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

3. Run the following command to join a cluster:

```
join
```

4. The ONTAP software detects the existing cluster and agrees to join the same cluster. Follow these prompts to join the cluster:

```
Existing cluster interface configuration found:
Port      MTU      IP              Netmask
e0a       9000     169.254.195.20  255.255.0.0
e0b       9000     169.254.209.147 255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: no
System Defaults:
Private cluster network ports [e0a,e0b].
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]: no
Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.
List the private cluster network ports [e0a,e0b]: Enter
Enter the cluster ports' MTU size [9000]: Enter
Enter the cluster network netmask [255.255.0.0]: Enter
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0a [169.254.73.101]: Enter
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0b [169.254.191.73]: Enter
```

5. Use the information in Table 5 to join node 02 to the cluster.

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

Note: The node should find the cluster name.

Note: The cluster-join process can take a minute or two.

6. Node 02 has successfully joined the cluster.

Configure Initial Cluster Settings

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

1. Open a Secure Shell (SSH) connection to the cluster IP address or to the host name.
2. Log in as the admin user with the password that you entered earlier.

Note: Use Table 4 and Table 5 for the input parameters for this section.

Assign Disks for Optimal Performance

SSDs are capable of significant I/O throughput, and proper disk assignment is required for optimal performance. To achieve optimal performance with SSDs, the disks in each chassis should be split between the controllers. Do not use the default allocation method of assigning all disks in a shelf to a single controller. In this solution, disks 0–11 on each chassis should be assigned to a controller, and disks 12–23 should be assigned to the other controller.

To assign the disks as required for this solution, complete the following steps:

1. Verify the current disk allocation.

```
disk show
```

2. Assign disks to the appropriate controller.

```
disk assign -disk <disk path name> -owner <<var_node01/02>> [-force]
```

Note: The `-force` option might be required if the disks are already assigned to another node. Verify that the disk is not a member of an existing aggregate before changing ownership.

Zero All Spare Disks

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

```
disk zerospares
```

Create Aggregates

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks that the aggregate contains.

This solution uses one aggregate on each controller, with 46 drives per aggregate. To create the aggregates required for this solution, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate n01_ssd01 -nodes <<var_node01>> -diskcount 23  
aggr create -aggregate n02_ssd01 -nodes <<var_node02>> -diskcount 23
```

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 per controller.

Note: The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until both `n01_ssd01` and `n02_ssd01` are online.

2. Disable NetApp Snapshot™ copies for the two data aggregates that you created in step 1

```
system node run -node <<var_node01>> aggr options n01_ssd01 nosnap on  
system node run -node <<var_node02>> aggr options n02_ssd01 nosnap on
```

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

```
system node run -node <<var_node01>> snap delete -A -a -f n01_ssd01  
system node run -node <<var_node02>> snap delete -A -a -f n02_ssd01
```

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>>
```

Verify Storage Failover

To confirm that storage failover is enabled, complete the following steps for a failover pair:

1. Verify the status of storage failover.

```
Storage failover show
```

2. Both nodes, `<<var_node01>>` and `<<var_node02>>`, must be capable of performing a takeover. If the nodes are capable of performing a takeover, go to step 4.
3. 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.

4. Verify the HA status for the two-node cluster.

Note: This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

5. If HA is configured, go to step 7.
6. Enable the HA mode only for the two-node cluster.

Note: 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
```

7. Verify that the hardware-assisted failover feature 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>>
```

Set Onboard UTA2 Ports Personality

To set the personality of the onboard unified target adapter 2 (UTA2) ports, complete the following steps:

1. Run the `ucadmin show` command to verify the current mode and current type of the ports.

```
Barsoom::> ucadmin show
```

Node	Adapter	Mode	Type	Current Mode	Current Type	Pending Status	Admin
Barsoom-01	0e	cna	target	-	-		online
Barsoom-01	0f	cna	target	-	-		online
Barsoom-01	0g	cna	target	-	-		online
Barsoom-01	0h	cna	target	-	-		online
Barsoom-02	0e	cna	target	-	-		online
Barsoom-02	0f	cna	target	-	-		online
Barsoom-02	0g	cna	target	-	-		online
Barsoom-02	0h	cna	target	-	-		online

8 entries were displayed.

2. Verify that the current mode of the ports in use is `cna` and the current type is `target`. If this is not the case, change the port personality by running the following command:

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

Note: The ports must be offline in order to run this command:

```
network fcp adapter modify -node Colts-stcl-1-02 -adapter 0h -status-admin down
```

Disable Flow Control on 10GbE and UTA2 Ports

A NetApp best practice is to disable flow control on all the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```
network port modify -node * -port e0e..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
```

Note: The `-node` and `-port` parameters in this example take advantage of the range operator available in the ONTAP shell.

Set Auto-Revert on Cluster Management Interface

To set the `auto-revert` parameter on the cluster management interface, run the following command:

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

Set Up Management Broadcast Domain

To set up the default broadcast domain for the management network interfaces, 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_node02>>:e0e,  
<<var_node02>>:e0f, <<var_node02>>:e0g, <<var_node02>>:e0h  
broadcast-domain show
```

Set Up Service Processor Network Interface

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_netmask>> -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_netmask>> -gateway  
<<var_node02_sp_gateway>>
```

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

Create Jumbo Frame MTU Broadcast Domains in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000  
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000  
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Create LACP Interface Groups

The LACP interface group (ifgrp) requires two or more Ethernet interfaces and a switch that supports the Link Aggregation Control Protocol (LACP). Therefore, confirm that the switch is configured properly.

To create interface groups, 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 e0f  
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0g  
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0h  
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 e0f  
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0g  
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0h  
ifgrp show
```

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

Note: The interface group name must follow the standard naming convention of `<number><letter>`, where:

`<number>` is an integer in the range of 0 to 999 without leading zeros.

<letter> is a lowercase letter.

Configure Jumbo Frames

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

```
network port modify -node * -port a0a -mtu 9000
WARNING: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

Note: Modifications to an interface group cause the underlying physical ports to inherit the same configuration. If the ports are later removed from the interface group, they retain these same settings. However, the inverse is not true; modifying the individual ports does not modify the interface group of which the ports are a member.

Note: After the MTU for the interface group is set to 9,000, all new VLAN interfaces created on that interface group also have an MTU of 9,000 bytes. Existing VLAN interfaces retain their original MTU after the ifgroup is changed.

Create VLANs

To create NFS and iSCSI VLANs and add them to their respective broadcast domains, run the following commands:

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_NFS_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_NFS_vlan_id>>
broadcast-domain add-ports -broadcast-domain <<var_NFS_broadcast_domain>> -ports
<<var_node01>>:a0a-<<var_NFS_vlan_id>>, <<var_node02>>:a0a-<<var_NFS_vlan_id>>
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-A_vlan_id>>
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-B_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-A_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-B_vlan_id>>
broadcast-domain add-ports -broadcast-domain <<var_iSCSI-A_broadcast_domain>> -ports
<<var_node01>>:a0a-<<var_iSCSI-A_vlan_id>>, <<var_node02>>:a0a-<<var_iSCSI-A_vlan_id>>
broadcast-domain add-ports -broadcast-domain <<var_iSCSI-B_broadcast_domain>> -ports
<<var_node01>>:a0a-<<var_iSCSI-B_vlan_id>>, <<var_node02>>:a0a-<<var_iSCSI-B_vlan_id>>
```

Enable Cisco Discovery Protocol

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

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

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

Configure NTP

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

1. Set the time zone for the cluster.

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

2. Set the date for the cluster.

```
date <ccyyymmddhhmm.ss>
```

Note: The format for the date is
<[century] [year] [month] [day] [hour] [minute] . [second]> (for example,
201707241320.00).

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

```
cluster time-service ntp server create -server <<var_global_ntp_server_ip>>
```

Configure SNMP

To configure SNMP, complete the following steps:

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

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

2. Configure SNMP traps to send to remote hosts.

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

Configure SNMPv1 Access

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

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

Use the `delete all` command with caution. If community strings are used for other monitoring products, then the `delete all` command removes them.

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. Run the security `snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -user-or-group-name snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
4. When prompted, enter a password for the authentication protocol. The password must have a minimum of eight characters.
5. Select des as the privacy protocol.
6. When prompted, enter a password for the privacy protocol. The password must have a minimum of eight characters.

Configure AutoSupport HTTPS

AutoSupport sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

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

Note: To enable AutoSupport to send messages using SMTP, change the `-transport` value in the preceding command to `smtp`. When configuring AutoSupport to use SMTP, be sure to enable mail relay on the mail server for the cluster management and node management IP addresses.

Configure HTTPS Access

Secure access to the storage controller is configured by default with self-signed certificates.

Set Up Storage VM

Create Storage VM

To create an infrastructure storage virtual machine (SVM), complete the following steps:

Note: The SVM is referred to as a Vserver in the ONTAP command-line interface (CLI).

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate n01_ssd01 -rootvolume-  
securitystyle unix
```

2. Select the SVM data protocols to configure.

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

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

```
vserver modify -vserver Infra-SVM -aggr-list n01_ssd01, n02_ssd01
```

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

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled  
vserver nfs show
```

Create Load-Sharing Mirror 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 root volume of the infrastructure SVM on each node.

```
volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate n01_ssd01 -size 1GB -type DP  
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate n02_ssd01 -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  
-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
```

Create iSCSI Service

Create the iSCSI service on each SVM.

Note: This command also starts the iSCSI service and sets the iSCSI IQN for the SVM.

```
iscsi create -vserver Infra-SVM  
iscsi show
```

Configure HTTPS Access

Secure access to the storage controller is configured by default with self-signed certificates.

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for each ESXi host created so that each host has its own rule index. For example, the first ESXi host has rule index 1, the second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 8 clientmatch <<var_esxi_host8_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

2. Assign the default export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```

Create FlexVol Volumes

To create a NetApp FlexVol® volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate n01_ssd01 -size 1TB -state online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-snapshotspace 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate n02_ssd01 -size 100GB -state online -policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate n02_ssd01 -size 500GB -state online -policy default -space-guarantee none -percent-snapshot-space 0
snapmirror update-ls-set -source-path Infra-SVM:rootvol
```

Create Boot LUNs for ESXi Hosts

The following procedure describes the process for configuring boot LUNs on the SSD aggregates, but it could be used on any ONTAP storage system. To create boot LUNs for ESXi hosts, complete the following steps:

1. Turn off automatic Snapshot copies of the volume.

```
volume modify -vserver Infra-SVM -volume esxi_boot -snapshot-policy none
```

2. Enable deduplication on the volume

```
volume efficiency on -vserver Infra-SVM -volume esxi_boot
```

3. Create LUNs for ESXi boot partitions for infrastructure hosts.

```
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-Infra-08 -size 15GB -ostype vmware space-reserve disabled
```

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi home-node <<var_node01>> -home-port a0a-<<var_iSCSI-A_vlan_id>> -address <<var_node01_iscsi_lif01a_ip>> -netmask <<var_node01_iscsi_lif01a_mask>> -status-admin up failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi home-node <<var_node01>> -home-port a0a-<<var_iSCSI-B_vlan_id>> -address
```

```
<<var_node01_iscsi_lif01b_ip>> -netmask <<var_node01_iscsi_lif01b_mask>> -status-admin up
failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi
home-node <<var_node02>> -home-port a0a-<<var_iSCSI-A_vlan_id>> -address
<<var_node02_iscsi_lif02a_ip>> -netmask <<var_node02_iscsi_lif02a_mask>> -status-admin up
failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi
home-node <<var_node02>> -home-port a0a-<<var_iSCSI-B_vlan_id>> -address
<<var_node02_iscsi_lif02b_ip>> -netmask <<var_node02_iscsi_lif02b_mask>> -status-admin up
failover-policy disabled -firewall-policy data -auto-revert false
network interface show
```

Create NFS LIFs

To create an NFS LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_infra_node_1 -role data -data-protocol nfs
home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_node01_nfs_ip>>
netmask <<var_node01_nfs_mask>> -status-admin up -failover-policy broadcast-domain-wide firewall-
policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_infra_node_2 -role data -data-protocol nfs
home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_node02_nfs_ip>>
netmask <<var_node02_nfs_mask>> -status-admin up -failover-policy broadcast-domain-wide firewall-
policy data -auto-revert true
```

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 e0M -address <<var_svm_mgmt_ip>> -netmask <<var_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_svm_mgmt_gateway>>
network route show
```

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

```
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
```

Configure iSCSI Boot

The iSCSI IQN values required for this step are not available until the Cisco UCS service profile has been configured. Complete the steps in the section “Create Service Profile Templates” and then run the following commands using the IQN variables listed in Table 12.

1. Create igroups for LUN mapping.

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol iscsi -ostype vmware -
initiator <<var_vm_host_infra_01_iqn>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol iscsi -ostype vmware -
initiator <<var_vm_host_infra_02_iqn>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-08 -protocol iscsi -ostype vmware -
initiator <<var_vm_host_infra_08_iqn>>
```

2. Map boot LUNs to hosts.

```
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-Infra-02 -igroup VM-Host-Infra-02 -
lunid 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-08 -igroup VM-Host-Infra-08 -
lunid 0
```

Set Up SVM for Exchange and SharePoint Workload

Create SVM

To create an SVM for Exchange and SharePoint, complete the following steps:

Note: The SVM is referred to as a Vserver in the ONTAP CLI.

1. Run the `vserver create` command.

```
vserver create -vserver Work-SVM -rootvolume rootvol -aggregate n02_ssd01 -rootvolume-
securitystyle unix
```

2. Select the SVM data protocols to configure.

```
vserver remove-protocols -vserver Work-SVM -protocols fcp,cifs,ndmp
```

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

```
vserver modify -vserver Work-SVM -aggr-list n01_ssd01, n02_ssd01
```

4. Enable and run the NFS protocol in the Work-SVM.

```
nfs create -vserver Work-SVM -udp disabled
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in.

```
vserver nfs modify -vserver Work-SVM -vstorage enabled
vserver nfs show
```

Create Load-Sharing Mirror 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 root volume of the infrastructure SVM on each node.

```
volume create -vserver Work-SVM -volume rootvol_m01 -aggregate n01_ssd01 -size 1GB -type DP
volume create -vserver Work-SVM -volume rootvol_m02 -aggregate n02_ssd01 -size 1GB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path Work-SVM:rootvol -destination-path Work-SVM:rootvol_m01 -type LS
schedule 15min
snapmirror create -source-path Work-SVM:rootvol -destination-path Work-SVM:rootvol_m02 -type LS
schedule 15min
snapmirror initialize-ls-set -source-path Work-SVM:rootvol
```

Create iSCSI Service

Create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM.

```
iscsi create -vserver Work-SVM iscsi show
```

Configure HTTPS Access

Secure access to the storage controller is configured by default with self-signed certificates.

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for each ESXi host created so that each host has its own rule index. For example, the first ESXi host has rule index 1, the second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Work-SVM -policyname default -ruleindex 9 -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Work-SVM -policyname default -ruleindex 16 clientmatch <<var_esxi_host8_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

2. Assign the default export policy to the infrastructure SVM root volume.

```
volume modify -vserver Work-SVM -volume rootvol -policy default
```

Create FlexVol Volumes

To create NetApp FlexVol volumes for SQL databases and SharePoint VMs, run the following commands:

```
volume create -vserver Work-SVM -volume VM_datastore_1 -aggregate n01_ssd01 -size 1TB -state online -policy default -junction-path /VM_datastore_1 -space-guarantee none -percent-snapshotspace 0
volume create -vserver Work-SVM -volume SQL1_Data -aggregate n01_ssd01 -size 1TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL1_SharePoint -aggregate n01_ssd01 -size 8TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL1_Log -aggregate n01_ssd01 -size 2TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL1_Snapinfo -aggregate n01_ssd01 -size 2TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL2_Data -aggregate n02_ssd01 -size 1TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL2_SharePoint -aggregate n02_ssd01 -size 8TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL2_Log -aggregate n02_ssd01 -size 2TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume SQL2_Snapinfo -aggregate n02_ssd01 -size 2TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Work-SVM -volume DocAve1_Media -aggregate n01_ssd01 -size 1.2TB -state online -policy default -space-guarantee none -percent-snapshot-space 0
snapmirror update-ls-set -source-path Work-SVM:rootvol
```

Note: The Exchange volumes are created later due to their number and specificity.

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver Work-SVM -lif iscsi_lif01a -role data -data-protocol iscsi home-node <<var_node01>> -home-port a0a-<<var_iSCSI-A_vlan_id>> -address <<var_node01_iscsi_lif01a_ip>> -netmask <<var_node01_iscsi_lif01a_mask>> -status-admin up failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Work-SVM -lif iscsi_lif01b -role data -data-protocol iscsi home-node <<var_node01>> -home-port a0a-<<var_iSCSI-B_vlan_id>> -address <<var_node01_iscsi_lif01b_ip>> -netmask <<var_node01_iscsi_lif01b_mask>> -status-admin up failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Work-SVM -lif iscsi_lif02a -role data -data-protocol iscsi home-node <<var_node02>> -home-port a0a-<<var_iSCSI-A_vlan_id>> -address <<var_node02_iscsi_lif02a_ip>> -netmask <<var_node02_iscsi_lif02a_mask>> -status-admin up failover-policy disabled -firewall-policy data -auto-revert false
network interface create -vserver Work-SVM -lif iscsi_lif02b -role data -data-protocol iscsi home-node <<var_node02>> -home-port a0a-<<var_iSCSI-B_vlan_id>> -address <<var_node02_iscsi_lif02b_ip>> -netmask <<var_node02_iscsi_lif02b_mask>> -status-admin up failover-policy disabled -firewall-policy data -auto-revert false
network interface show
```


Create NFS LIFs

To create an NFS LIF, run the following commands:

```
network interface create -vserver Work-SVM -lif nfs_work_node_1 -role data -data-protocol nfs
home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_nfs_work_1>> -netmask
<<var_nfs_work_1_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy
data -auto-revert true
network interface create -vserver Work-SVM -lif nfs_work_node_2 -role data -data-protocol nfs
home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_nfs_work_2>> -netmask
<<var_nfs_work_2_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy
data -auto-revert true
```

Add Infrastructure SVM Administrator

To add the work 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 Work-SVM -lif vsmgmt -role data -data-protocol none -home-node
<<var_node02>> -home-port e0M -address <<var_svm_mgmt_ip>> -netmask <<var_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 Work-SVM -destination 0.0.0.0/0 -gateway <<var_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 Work-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
```

3.2 Cisco UCS Server Configuration

This section provides the complete configuration of the Cisco UCS server environment.

FlexPod Cisco UCS Base

Set Up Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

This section provides detailed procedures for configuring the Cisco UCS for use in a FlexPod environment. The steps explain how to provision the Cisco UCS B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS 6248 Fabric Interconnect A

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

```
Enter the configuration method. (console/gui) ? 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 password? (y/n) [y]: y
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)? (yes/no) [n]: y
Enter the 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 the DNS Server IPv4 address? (yes/no) [n]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? (yes/no) [n]: y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

2. Review the settings displayed on the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt to verify that the configuration has been saved.

Cisco UCS 6248 Fabric Interconnect B

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

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

```
Enter the configuration method. (console/gui) ? 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 of 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 confirm that the configuration has been saved.

FlexPod 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 6248 fabric interconnect cluster address.
2. Click the Launch UCS Manager link inside the HTML box to open the Cisco UCS Manager software.
3. If prompted to accept the security certificates, accept as necessary.
4. When prompted, enter admin as the username and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 3.1(3a)

This reference architecture uses Cisco UCS Manager software version 3.1(3a). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 fabric interconnect software to version 3.1(3a), refer to Cisco UCS Manager Install and Upgrade Guides.

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. This procedure assumes that the appropriate firmware has been uploaded to the Cisco UCS fabric interconnects.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter <<host_fw_pkg_name>> as the name of the host firmware package.

6. Keep the Simple option selected.
7. Select version 3.1(3a)B for the blade package.

Name :

Description :

How would you like to configure the Host Firmware Package?

☒ Simple ☐ Advanced

Blade Package :

Rack Package :

Service Pack :

8. Click OK to create the host firmware package.
9. Click OK.

Add Block of IP Addresses for Out-of-Band KVM Access

To create a block of IP addresses for server keyboard, video, and mouse (KVM) access in the Cisco UCS environment, complete the following steps:

Note: This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager. IP address blocks cannot be modified after they are created. If changes are required, the range of addresses must be deleted and then recreated.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > Root > IP Pools.
3. Right-click IP Pools and select Create IP Pool.
4. Name the pool in-band-mgmt.
5. Click Next.
6. Click Add.
7. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information.

Create IP Pool

Name	From	To	Subnet	Default Gateway	Primary DNS	Secondary DNS
[172.21.91.90-172.21.91.98]	172.21.91.90	172.21.91.98	255.255.255.0	172.21.91.254	172.21.91.71	172.21.91.72

8. Click Finish to create the IP block.
9. 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:

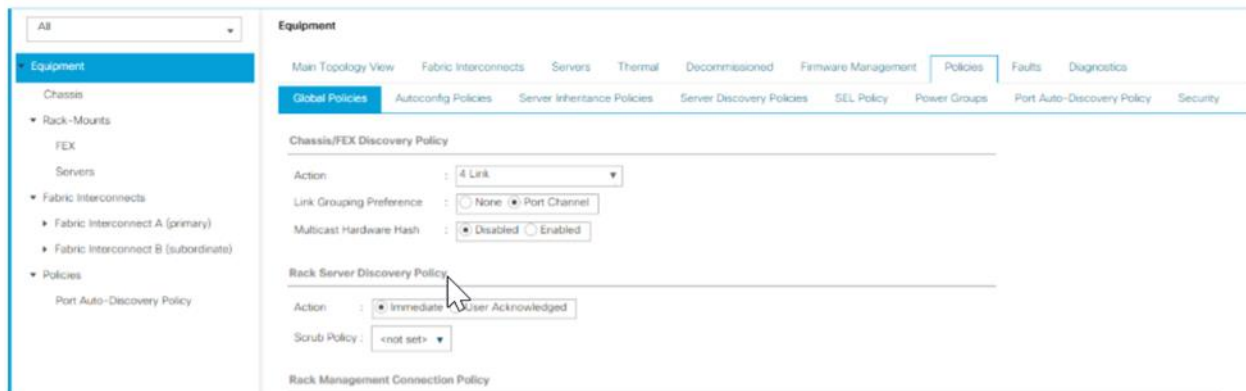
1. In the Cisco UCS Manager, click the Admin tab in the navigation pane.
2. Select All > Timezone Management.
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes and then click OK.
5. Click Add NTP Server.
6. Enter `<<var_global_ntp_server_ip>>` and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

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and Cisco UCS C-Series servers.

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 tree on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set Chassis/FEX Discovery Policy to 4 Link or set it to match the number of uplink ports that are cabled between the chassis and the fabric interconnects.
4. Set the Link Grouping Preference option to Port Channel.



5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (Primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis. Right-click the ports and select Configure as Server Port.
5. Click Yes to confirm the server ports and click OK.
6. Select ports 29, 30, 31, and 32, which are connected to the Cisco Nexus 9396PX switches. Right-click the ports and select Configure as Uplink Port.
7. Click Yes to confirm the uplink ports and click OK.
8. In the left pane, navigate to fabric interconnect A. In the right pane, navigate to Physical Ports > Ethernet Ports. Confirm that the ports are configured correctly in the If Role column.

1	0	1	54:7F:EE:23:52:28	Server	Physical	Up	Enabled
1	0	2	54:7F:EE:23:52:29	Server	Physical	Up	Enabled
1	0	3	54:7F:EE:23:52:2A	Server	Physical	Up	Enabled
1	0	4	54:7F:EE:23:52:2B	Server	Physical	Up	Enabled
1	0	29	54:7F:EE:23:52:44	Network	Physical	Up	Enabled
1	0	30	54:7F:EE:23:52:45	Network	Physical	Up	Enabled
1	0	31	54:7F:EE:23:52:46	Network	Physical	Up	Enabled
1	0	32	54:7F:EE:23:52:47	Network	Physical	Up	Enabled

9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (Subordinate) > Fixed Module.
10. Expand Ethernet Ports
11. Select the ports that are connected to the chassis. Right-click the ports and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select ports 29, 30, 31, and 32, which are connected to the Cisco Nexus 9396 switches. Right-click the ports and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

15. In the left pane, navigate to fabric interconnect B. In the right pane, navigate to Physical Ports > Ethernet Ports. Confirm that the ports are configured correctly in the If Role column.

Acknowledge Cisco UCS Chassis

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

1. In Cisco UCS Manager, click the Equipment tab in the Navigation pane.
2. Expand Chassis and select the chassis that is listed.
3. Right-click the chassis and select Acknowledge Chassis.
4. Click Yes and then click OK to complete chassis acknowledgement.

Create Uplink Port Channels to Cisco Nexus 9396PX Switches

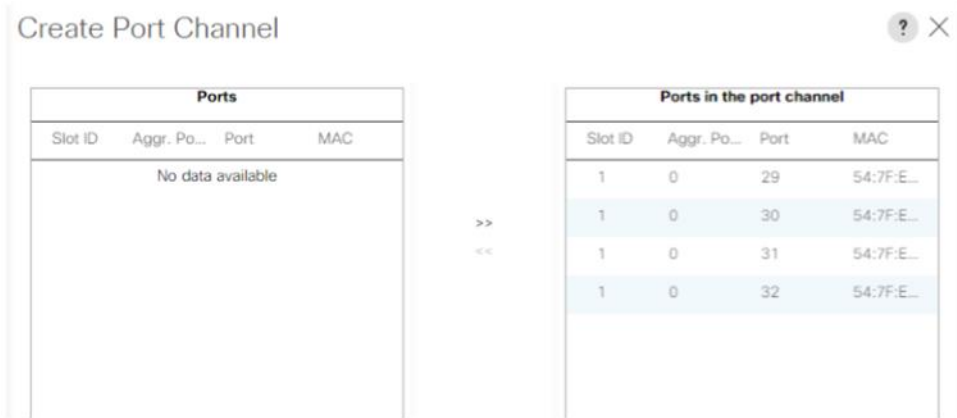
To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
Note: This procedure creates two port channels: one from fabric A to both Cisco Nexus 9396PX switches and one from fabric B to both Cisco Nexus 9396PX switches.
2. Under LAN > LAN Cloud, expand the fabric A node.



3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter `<<var_fabric_A_Portchannel_ID>>` as the unique ID of the port channel.
6. Enter `<<var_fabric_A_Portchannel_name>>` as the name of the port channel.
7. Click Next.
8. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 29
 - Slot ID 1 and port 30
 - Slot ID 1 and port 31
 - Slot ID 1 and port 32

9. Click >> to add the ports to the port channel.



10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B node.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter <<var_fabric_B_Portchannel_ID>> as the unique ID of the port channel.
16. Enter <<var_fabric_B_Portchannel_name>> as the name of the port channel.
17. Click Next.
18. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 29
 - Slot ID 1 and port 30
 - Slot ID 1 and port 31
 - Slot ID 1 and port 32
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

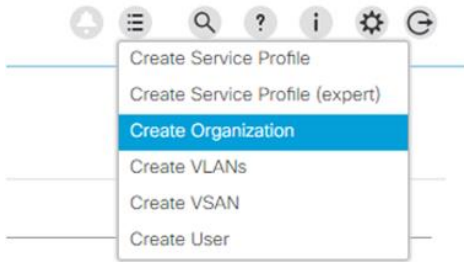
Create Organization (Optional)

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

Note: Although the use of organizations is not assumed in this document, this section covers the process of creating an organization.

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

1. In Cisco UCS Manager, from the menu in the toolbar at the top of the window, select Create Organization.



2. Enter a name for the organization.
3. Optional: Enter a description for the organization.
4. Click OK.
5. Click OK in the confirmation message.

Create MAC Address Pools

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

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

Note: This procedure creates two MAC address pools, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter `MAC_Pool_A` as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Keep the assignment order setting at Default.
8. Click Next.
9. Click Add to add a block of MAC addresses to the pool.



10. Specify the starting MAC address in the block for fabric A.

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

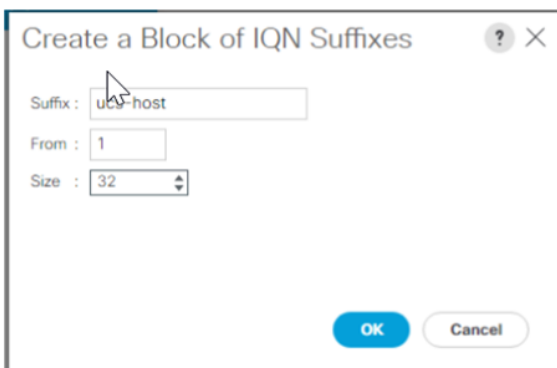
11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.

15. Right-click MAC Pools under the root organization.
16. Select Create MAC Pool to create the MAC address pool.
17. Enter `MAC_Pool_B` as the name of the MAC pool.
18. Optional: Enter a description for the MAC pool.
19. Keep the assignment order setting at Default.
20. Click Next.
21. Click Add to add a block of MAC addresses to the pool.
22. Specify the starting MAC address in the block for fabric B.
Note: For the FlexPod solution, the recommendation is to place 0B in the next-to-last octet of the starting MAC address to identify all MAC addresses in this pool as fabric B addresses.
23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

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

1. In the Cisco UCS Manager, select the SAN tab on the left.
2. Select Pools > Root.
3. Right-click IQN Pools under the root organization.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter `IQN_Pool` for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter `iqn.1992-08.com.cisco` as the prefix.
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter `ucs-host` as the suffix.
12. Enter 1 in the From field.
13. Specify a size of the IQN block sufficient to support the available server resources.



Create a Block of IQN Suffixes

Suffix :

From :

Size :

OK Cancel

14. Click OK.
15. Click Finish.
16. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

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

1. In the Cisco UCS Manager, select the LAN tab on the left.
2. Select Pools > Root.
3. Two IP pools are created, one for each switching fabric.
4. Right-click IP Pools under the root organization.
5. Select Create IP Pool to create the IP pool.
6. Enter `iSCSI_IP_Pool_A` for the name of the IP pool.
7. Optional: Enter a description of the IP pool.
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
12. Set the size to enough addresses to accommodate the servers.
13. Click OK.
14. Click Finish.
15. Right-click IP Pools under the root organization.
16. Select Create IP Pool to create the IP pool.
17. Enter `iSCSI_IP_Pool_B` for the name of the IP pool.
18. Optional: Enter a description of the IP pool.
19. Select Sequential for Assignment Order.
20. Click Next.
21. Click Add.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers.
24. Click OK.
25. Click Finish.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > Root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID_Pool` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.

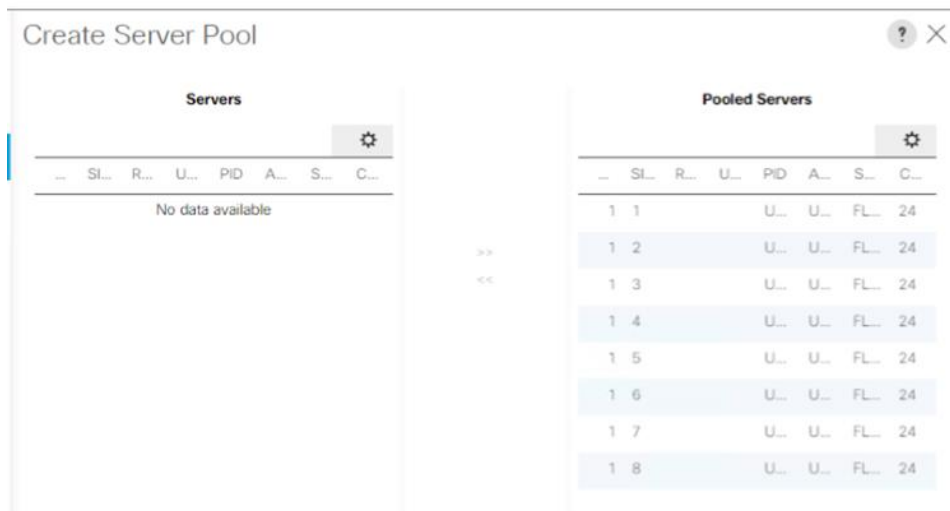
7. Keep the Prefix setting as Derived.
8. Keep the Assignment Order setting at Default.
9. Click Next.
10. Click Add to add a block of UUIDs to the pool.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool (Optional)

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > Root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `Host_Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select the servers to be used for the pool and click >> to add them to the `Host_Pool` server pool.



9. Click Finish.
10. Click OK.

Create VLANs

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.
10. Expand the list of VLANs in the navigation pane, right-click the newly created Native VLAN, and select Set as Native VLAN.
11. Click Yes and then click OK.
12. Right-click VLANs.
13. Select Create VLANs
14. Enter iSCSI-A-VLAN as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VLAN ID for the first iSCSI VLAN.
17. Click OK and then OK again.
18. Right-click VLANs.
19. Select Create VLANs
20. Enter iSCSI-B-VLAN as the name of the VLAN to be used for the second iSCSI VLAN.
21. Keep the Common/Global option selected for the scope of the VLAN.
22. Enter the VLAN ID for the second iSCSI VLAN.
23. Click OK and then OK again.
24. Right-click VLANs.
25. Select Create VLANs.
26. Enter `Mgmt-VLAN` as the name of the VLAN to be used for management traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the in-band management VLAN ID.
29. Keep the Sharing Type as None.
30. Click OK and then click OK again.
31. Right-click VLANs.
32. Select Create VLANs.
33. Enter NFS-VLAN as the name of the VLAN to be used for NFS.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the NFS VLAN ID.
36. Keep the Sharing Type as None.
37. Click OK and then click OK again.
38. Right-click VLANs.
39. 39. Select Create VLANs.
40. Enter `vMotion` as the name of the VLAN to be used for vMotion.
41. Keep the Common/Global option selected for the scope of the VLAN.

42. Enter the vMotion VLAN ID.
43. Keep the Sharing Type as None.
44. Click OK and then click OK again.

VLANs

Advanced Filter Export Print

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	No	None
VLAN IB-MGMT (3347)	3347	Lan	Ether	No	None
VLAN INFRA-NFS (3349)	3349	Lan	Ether	No	None
VLAN iSCSI-A-VLAN (3350)	3350	Lan	Ether	No	None
VLAN iSCSI-B-VLAN (3351)	3351	Lan	Ether	No	None
VLAN Native-VLAN (2)	2	Lan	Ether	Yes	None
VLAN vMotion-VLAN (3348)	3348	Lan	Ether	No	None

Create VLAN Group and Assign In-band Profile

A VLAN group is required to set up in-band KVM access. To create a VLAN group, complete the following steps:

1. In the Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLAN Groups and select Create VLAN Group.
4. Name the VLAN group FlexPod and select all VLANs.
5. Select the radio button next to the Native VLAN and click Next.
6. Click Next.
7. Select the two uplink port channels and use the >> button to add them to the VLAN group.
8. Select LAN > LAN Cloud. Then select the Global Policies tab.
9. In the In-band VLAN Group box, select FlexPod VLAN Group, select MGMT-VLAN Network as the network, and select `in-band-mgmt` as the IP pool name.
10. Select Save Changes and then select OK.

Set Jumbo Frames in Cisco UCS Fabric

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

1. In the 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.

LAN /

General Events FSM

Actions Properties

Use Global Owner : Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

5. Click Save Changes.
6. Click OK.

Create Local Disk Configuration Policy (Optional)

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

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter iSCSI-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.
7. Keep the FlexFlash State and FlexFlash RAID Reporting State settings at Disable.
8. Click OK to create the local disk configuration policy.
9. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables CDP on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > Root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable_CDP` as the policy name.
6. For CDP, select the Enabled option.

7. Click OK to create the network control policy.
8. Click OK.

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

Create Power Control Policy

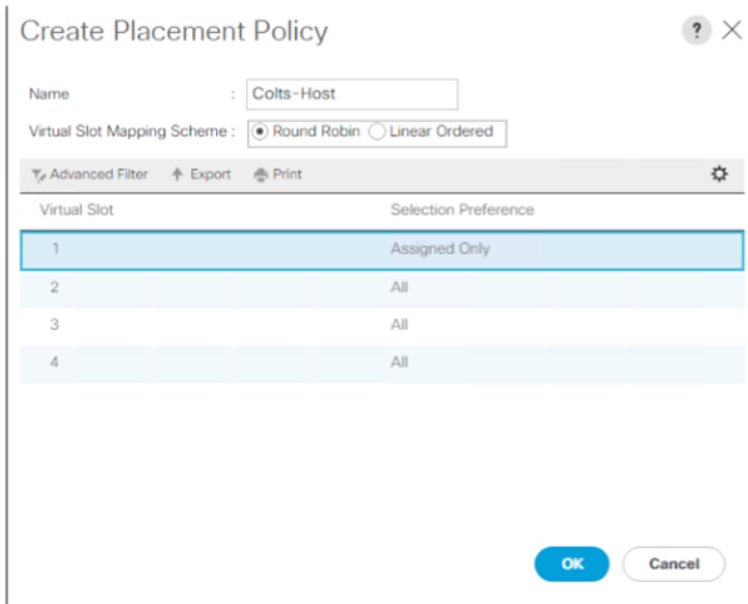
To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Create vNIC/vHBA Placement Policy for VM Infrastructure Hosts

To create a virtual network interface card/virtual host bus adapter (vNIC/vHBA) placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter Colts-Host as the name of the placement policy.
6. Click 1 and select Assigned Only under the Selection Preference column.



Create Placement Policy

Name :

Virtual Slot Mapping Scheme : ☒ Round Robin ☐ Linear Ordered

Advanced Filter Export Print

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

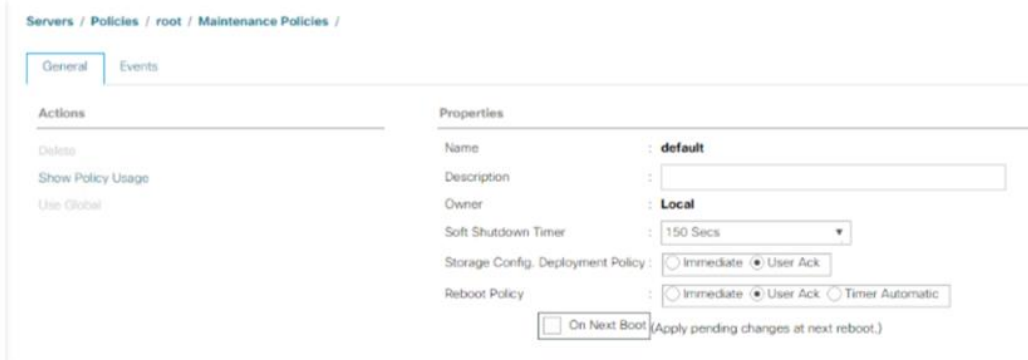
OK Cancel

- Click OK and then click OK again.

Update Default Maintenance Policy

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

- In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Select Policies > Root.
- Select Maintenance Policies > Default.
- Change the reboot policy setting to User Ack.



Servers / Policies / root / Maintenance Policies /

General Events

Actions

- Delete
- Show Policy Usage
- Use Global

Properties

Name : default

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : ☐ Immediate ☒ User Ack

Reboot Policy : ☐ Immediate ☒ User Ack ☐ Timer Automatic

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

- Click Save Changes.
- Click OK to acknowledge the change.

Create vNIC Templates

This section describes how to create the required vNICs.

Create Data vNICs

To create multiple vNIC templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > Root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. For fabric ID, select Fabric A.
 - Note:** Do not select the Enable Failover option.
 - Note:** Under Target, do not select the VM option.
7. Select Updating Template as the template type.
8. Under VLANs, select the checkboxes for the following VLANs:
 - MGMT-VLAN
 - Native-VLAN
 - vMotion-VLAN
 - NFS-VLAN
9. Set Native-VLAN as the native VLAN.
10. For MTU, enter 9000.
11. In the MAC Pool list, select MAC_Pool_A.
12. In the Network Control Policy list, select Enable_CDP.

Create vNIC Template

Select	Name	Native VLAN
<input type="checkbox"/>	ISCSI-A-VLAN	<input type="radio"/>
<input type="checkbox"/>	ISCSI-B-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	MGMT-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

Create VLAN

CDN Source : ☒ vNIC Name ☐ User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

13. Click OK to create the vNIC template.
14. Click OK.
15. In the navigation pane, click the LAN tab.
16. Select Policies > Root.
17. Right-click vNIC Templates.
18. Select Create vNIC Template.

19. Enter `vNIC_Template_B` as the vNIC template name.
20. Select Fabric B.
Note: Do not select the Enable Failover checkbox.
Note: Under Target, do not select the VM checkbox.
21. Select Updating Template as the template type.
22. Under VLANs, select the checkboxes for the following VLANs:
 - MGMT-VLAN
 - Native-VLAN
 - vMotion-VLAN
 - NFS-VLAN
23. Set Native-VLAN as the native VLAN.
24. For MTU, enter 9000.
25. In the MAC Pool list, select `MAC_Pool_B`.
26. In the Network Control Policy list, select `Enable_CDP`.
27. Click OK to create the vNIC template.
28. Click OK.

Create iSCSI vNICs

To create iSCSI vNICs, complete the following steps:

1. Select the LAN tab on the left.
2. Select Policies > Root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `iSCSI_Template_A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM option is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select iSCSI-A-VLAN.
10. Set iSCSI-A-VLAN as the native VLAN.
11. Under MTU, enter 9000.
12. From the MAC Pool list, select `MAC_Pool_A`.
13. From the Network Control Policy list, select `Enable_CDP`.
14. Click OK to complete creating the vNIC template.
15. Click OK.
16. Select the LAN tab on the left.
17. Select Policies > Root.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter `iSCSI_Template_B` as the vNIC template name.
21. Select Fabric B. Do not select the Enable Failover option.
22. Under Target, make sure that the VM option is not selected.

23. Select Updating Template for Template Type.
24. Under VLANs, select iSCSI-B-VLAN.
25. Set iSCSI-B-VLAN as the native VLAN.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select MAC_Pool_B.
28. From the Network Control Policy list, select Enable_CDP.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif02a and iscsi_lif02b). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

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

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

Create 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

Add Local Disk

Add Local LUN

Add Local JBOD

Add SD Card

Add Internal USB

Add External USB

Add Embedded Local LUN

Add Embedded Local Disk

Add CD/DVD

Boot Order

+ - Advanced Filter Export Print

Name	Or...	vNIC/vHBA/L...	Type	WWN	LUN ...	Slot N...	Boot ...	Boot ...	Descr...
Remote CD/DVD	1								
iSCSI 2									
iSCSI		iSCSI-A-vNIC	Primary						
iSCSI		iSCSI-B-vNIC	Seco...						

Move Up Move Down Delete

Create Service Profile Templates

In this procedure, one service profile template for 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.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template.
6. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
7. Select the Updating Template option.
8. Under UUID, select UUID_Pool as the UUID pool.
9. Click Next.

Configure Storage Provisioning

To configure the storage policy, complete the following steps:

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

Configure Networking Options

To configure the networking options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Expert option to configure LAN connectivity.

3. Click the upper Add button to add a vNIC to the template.
4. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
5. Select the Use vNIC Template option.
6. In the vNIC Template list, select vNIC_Template_A.
7. In the Adapter Policy list, select VMware.
8. Click OK to add this vNIC to the template.
9. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.
10. In the Create vNIC box, enter vNIC-B as the name of the vNIC.
11. Select the Use vNIC Template option.
12. In the vNIC Template list, select vNIC_Template_B.
13. In the Adapter Policy list, select VMware.
14. Click OK to add the vNIC to the template.
15. Click the upper Add button to add a vNIC to the template.
16. In the Create vNIC dialog box, enter iSCSI-A-vNIC as the name of the vNIC.
17. Select the Use vNIC Template checkbox. In the vNIC Template list, select iSCSI_Template_A.
18. In the Adapter Policy list, select VMware.
19. Click OK to add this vNIC to the template.
20. Click the upper Add button to add a vNIC to the template.
21. In the Create vNIC dialog box, enter iSCSI-B-vNIC as the name of the vNIC.
22. Select the Use vNIC Template checkbox.
23. In the vNIC Template list, select iSCSI_Template_B.
24. In the Adapter Policy list, select VMware.
25. Click OK to add this vNIC to the template.
26. Expand the iSCSI vNICs section (if it is not already expanded).
27. Select ign-pool under Initiator Name Assignment.
28. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
29. Enter iSCSI-A-vNIC as the name of the vNIC.
30. Select iSCSI-A-vNIC for Overlay vNIC.
31. Set the iSCSI Adapter Policy to default.
32. Set the VLAN to iSCSI-A-VLAN.
33. Leave the MAC Address set to None.
34. Click OK.
35. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
36. Enter iSCSI-B-vNIC as the name of the vNIC.
37. Set the Overlay vNIC to iSCSI-B-vNIC.
38. Set the iSCSI Adapter Policy to default.
39. Set the VLAN to iSCSI-B-VLAN.
40. Leave the MAC Address set to None.
41. Click OK.
42. Click OK.
43. Review the table in the Networking page to make sure that all vNICs were created.

44. Click Next.

Create Service Profile Template

Optionally specify LAN configuration information.

How would you like to configure LAN connectivity?

☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

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

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

Configure Storage Options

To configure the storage options, complete the following steps:

1. Select the No vHBAs option for the How Would You Like to Configure SAN Connectivity? option.
2. Click Next.

Configure Zoning Options

To configure the zoning options, complete the following step:

1. Set no zoning options and click Next.

Configure vNIC/HBA Placement

To configure vNIC/HBA placement, complete the following steps:

1. Set the vNIC/vHBA placement options.
2. In the Select Placement list, select the Colts-Host placement policy.
3. Select vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - vNIC-A
 - vNIC-B
 - iSCSI-vNIC-A
 - iSCSI-vNIC-B
4. Review the table to verify that all vNICs and vHBAs were assigned to the policy in the appropriate order.

Create Service Profile Template

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: Colts-Host Create Placement Policy

vNICs **vHBAs**

Name

No data available

>> assign >>

<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Selection Preference
▼ vCon 1		Assigned Only
vNIC iSCSI-A...	3	
vNIC iSCSI-B...	4	
vNIC vNIC-A	1	
vNIC vNIC-B	2	
vCon 2		All

Move Up Move Down

- Click Next.

Configure vMedia Policy

To configure the vMedia policy, complete the following step:

- Click Next.

Configure Server Boot Order

To configure the server boot order, complete the following steps:

- Select Boot-Fabric-A for Boot Policy.
- In the Boot Order pane, select iSCSI-A-vNIC.
- Click the Set iSCSI Boot Parameters button.
- Leave the Set iSCSI Boot Parameters dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
- Set `iSCSI_IP_Pool_A` as the Initiator IP address policy.
- Keep the iSCSI Static Target Interface button selected and click the Add button.
- Log in to the storage cluster management interface and run the `iscsi show` command.
- Note or copy the iSCSI target name for infra-SVM.
- In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from infra-SVM.
- Enter the IP address of `iscsi_lif02a` for the IPv4 Address field.
- Click OK to add the iSCSI static target.
- Keep the iSCSI Static Target Interface option selected and click the Add button.
- In the Create iSCSI Static Target window, paste the iSCSI target node name from infra-SVM into the iSCSI Target Name field.
- Enter the IP address of `iscsi_lif01a` in the IPv4 Address field.
- Click OK.

16. Click OK.
17. In the Boot Order pane, select iSCSI-vNIC-B.
18. Click the Set iSCSI Boot Parameters button.
19. In the Set iSCSI Boot Parameters dialog box, set the Initiator Name Assignment to <not set>.
20. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.
21. Keep the iSCSI Static Target Interface option selected and click the Add button.
22. In the Create iSCSI Static Target window, paste the iSCSI target node name from infra-SVM into the iSCSI Target Name field (same target name as earlier).
23. Enter the IP address of `iscsi_lif02b` in the IPv4 address field.
24. Click OK to add the iSCSI static target.
25. Keep the iSCSI Static Target Interface option selected and click Add.
26. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from infra-SVM into the iSCSI Target Name field.
27. Enter the IP address of `iscsi_lif01b` in the IPv4 Address field.
28. Click OK.
29. Click OK.
30. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
31. Click Next to continue to the next section.

Configure Maintenance Policy

To configure the maintenance policy, complete the following steps:

1. Select the default Maintenance Policy.
2. Click Next.

Configure Server Assignment

To configure the server assignment, complete the following steps:

1. In the Pool Assignment list, select Host_Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Expand Firmware Management at the bottom of the page and select Colts-Host from the Host Firmware list.
4. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.
2. Click Finish to create the service profile template.
3. Click OK in the confirmation message.

Create Service Profiles

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > Root > Service Template VM-Host-Infra-Fabric-A.

3. Right-click VM-Host-Infra-Fabric-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 1 in the Name Suffix Starting Number field.
6. Enter 8 as the number of instances to create.
7. Click OK to create the service profiles for infrastructure hosts.
8. Click OK in the confirmation message.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment must have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into the following tables (Table 6 and Table 7).

Table 6) iSCSI LIFs for iSCSI IQN.

SVM	iSCSI Target IQN
Infra-SVM	

Note: To gather the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

Table 7) vNIC iSCSI IQNs for fabric A and fabric B.

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-Infra-01		<< var_vm_host_infra_01_iqn >>
VM-Host-Infra-02		<< var_vm_host_infra_02_iqn >>
VM-Host-Infra-08		<< var_vm_host_infra_08_iqn >>

Note: To gather the vNIC IQN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > Root. Click each service profile and then click the iSCSI vNICs tab on the right. Note the initiator name displayed at the top of the page under Service Profile Initiator Name.

3.3 Cisco Nexus Storage Networking Configuration

This section describes how to configure the Cisco Nexus switches for use in the FlexPod environment.

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. Before configuring the switches, make sure that they are running Cisco Nexus NX-OS 7.0(3)I4(6) or later.

Set Up Initial Configuration

Cisco Nexus 9396PX A

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

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter power on autoprovisioning.

```
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
Number of rsa key bits <1024-2048> : 1024
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 the Cisco Nexus B switch on <<var_nexus_B_hostname>>, complete the following steps:

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter power on autoprovisioning.

```
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
Number of rsa key bits <1024-2048> : 1024
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
```

FlexPod Cisco Nexus iSCSI Storage vSphere on ONTAP

This section describes how to configure the FlexPod environment for iSCSI.

Enable Licenses

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

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

1. Log in as an administrator.
2. Run the following commands:

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

Set Global Configurations

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To create the necessary VLANs, run the following commands from the global configuration mode on both switches:

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_iscsi-a_vlan_id>>
name iSCSI-A-VLAN
exit
vlan <<var_iscsi-b_vlan_id>>
name iSCSI-B-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

From the global configuration mode, run the following commands:

```
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <var_switch_ntp_ip>/<var_ib-mgmt_vlan_netmask_length>
no shutdown
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 9396PX A

To add individual port descriptions for troubleshooting activity and verification for switch A, run the following commands from the global configuration mode:

```
interface Eth1/1
description <<var_ucs_clustername>>-A:1/31
exit
interface Eth1/2
description <<var_ucs_clustername>>-B:1/31
exit
interface Eth1/3
description <<var_ucs_clustername>>-A:1/29
exit
interface Eth1/4
description <<var_ucs_clustername>>-B:1/29
exit
interface Eth1/5
description <<var_node01>>:e0e
exit
interface Eth1/6
description <<var_node01>>:e0g
exit
interface Eth1/7
description <<var_node02>>:e0e
exit
interface Eth1/8
description <<var_node02>>:e0g
exit
interface Eth1/47
description <<var_nexus_B_hostname>>:1/47
exit
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, run the following commands from the global configuration mode:

```
interface Eth1/1
description <<var_ucs_clustername>>-A:1/32
exit
interface Eth1/2
description <<var_ucs_clustername>>-B:1/32
exit
interface Eth1/3
description <<var_ucs_clustername>>-A:1/30
exit
interface Eth1/4
description <<var_ucs_clustername>>-B:1/30
exit
interface Eth1/5
description <<var_node01>>:e0f
exit
interface Eth1/6
description <<var_node01>>:e0h
exit
interface Eth1/7
```

```

description <<var_node02>>:e0f
exit
interface Eth1/8
description <<var_node02>>:e0h
exit
interface Eth1/47
description <<var_nexus_A_hostname>>:1/47
exit
interface Eth1/48
description <<var_nexus_A_hostname>>:1/48
exit

```

Create Port Channels

Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To create the necessary port channels between devices, run the following commands from the global configuration mode:

```

interface Po10
description vPC peer-link
exit
interface Eth1/47-48
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth1/5-6
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth1/7-8
channel-group 12 mode active
no shutdown
exit
interface Po13
description <<var_ucs_clustername>>-A
exit
interface Eth1/1
channel-group 13 mode active
no shutdown
exit
interface Eth1/3
channel-group 13 mode active
no shutdown exit
interface Po14
description <<var_ucs_clustername>>-B
exit
interface Eth1/2
channel-group 14 mode active
no shutdown
exit
interface Eth1/4
channel-group 14 mode active
no shutdown
exit
copy run start

```

Configure Port Channels for Cisco Nexus 9396PX A and Cisco Nexus 9396PX B

To configure port channel parameters, run the following commands from the global configuration mode on both switches:

```

interface Po10
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type network
exit
interface Po11
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>,
<<var_iscsib_vlan_id>> spanning-tree port type edge trunk
mtu 9216
exit
interface Po12
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>,
<<var_iscsib_vlan_id>> spanning-tree port type edge trunk
mtu 9216
exit
interface Po13
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po14
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
copy run start

```

Configure Virtual Port Channels for Cisco Nexus 9396PX A

To configure virtual port channels (vPCs) for switch A, run the following commands from the global configuration mode:

```

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
delay restore 150
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start

```

Configure Virtual Port Channels for Cisco Nexus 9396PX B

To configure vPCs for switch B, run the following commands from the global configuration mode:

```
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
delay restore 150
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
```

Note: Remember to run `copy run start` to permanently save the switch configurations.

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 9396PX switches included in the FlexPod environment into the infrastructure. The previous 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.

4 Extensibility

In this section, we will discuss the extensibility to the cloud and multitenancy.

4.1 Multitenant


Servers on FlexPod can be deployed quickly because FlexPod allows server characteristics to be abstracted into service profiles. Service profiles can be mapped very quickly to physical servers, in contrast to traditional approaches that require manual configuration for every change.

This abstraction also allows FlexPod to operate as a highly efficient private cloud platform. Servers can be provisioned on a project-by-project basis and return to the server pool rapidly after their use. The provisioning of these servers can be controlled by the administrator role, and the use of the platform can be metered for chargeback purposes. FlexPod's capability to abstract server, network, and storage characteristics into profiles also allows multiple workloads to share hardware while maintaining separation of application and data. This approach results in greater utilization of resources compared to traditional server deployments.

Note: In the past, because customers needed at least six hours to deploy code out to our Web servers, they were only able to make changes once a day. Today, customers use Cisco UCS service profiles, NetApp rapid cloning, and our custom orchestration tool to deploy or update as many as 60 virtual servers in less than 30 minutes. And it's all automated—we've replaced a complicated, mistake-prone 100-step runbook process with push-button code deployments that are bulletproof and complete 92% faster.

FlexPod enables its computing, networking, and storage resources to be securely partitioned through the use of service profiles and APIs. This feature allows multiple enterprise clients and applications to share FlexPod in a mutually exclusive way (Figure 3).

Figure 3) FlexPod multitenant capability.



Note: The FlexPod validated data center solution built on a flexible, shared infrastructure is a much better way of doing IT. It's cost effective and makes it easy to add storage, plug in additional blades, assign new server profiles, and bring up new multitenant environments. End-to-end secure multitenancy also dramatically simplifies compliance

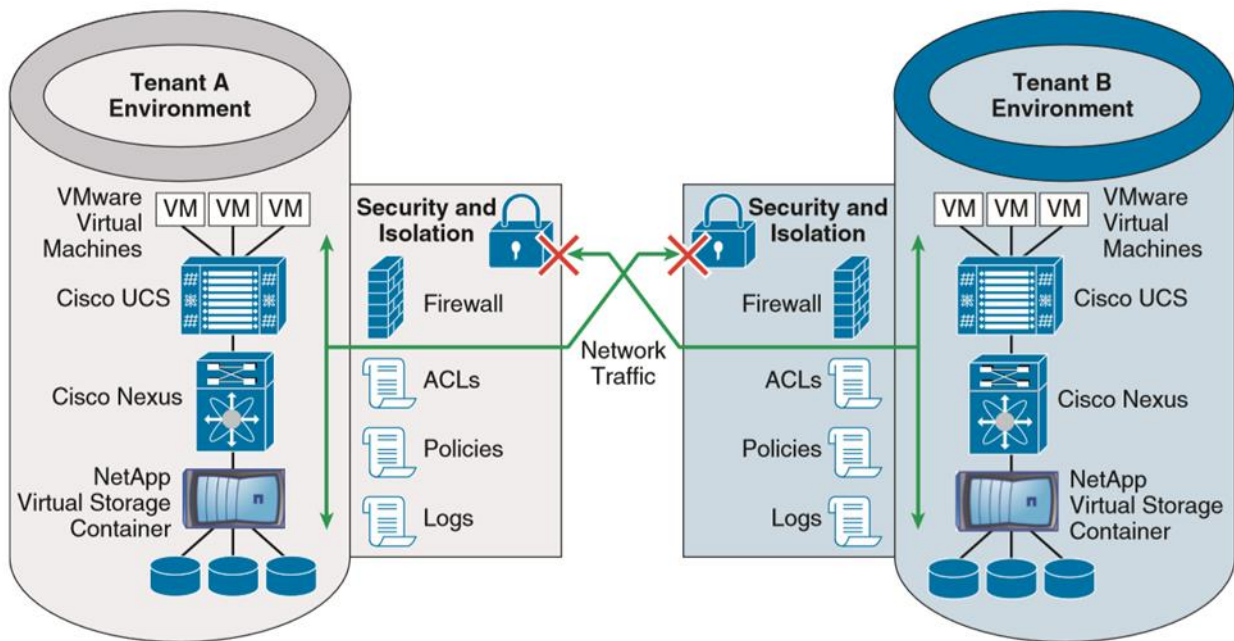
Many data centers are designed to run virtualized servers to gain application workload agility and better hardware resource utilization. Other data centers need to support non-virtualized servers. Either design can greatly benefit from FlexPod's service profiles, unified virtual networking, and APIs. The FlexPod management ecosystem of solution developers delivers solutions that streamline the deployment of FlexPod, orchestrate service delivery, and enable self-service IT.

4.2 Architecture Overview

One of the essential characteristics of a cloud architecture is the ability to pool resources. The provider's compute, network, and storage resources are pooled to serve multiple customers using a multitenant model, with different physical and virtual resources dynamically assigned and reassigned according to customer demand. There is a sense of location of independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (for example, country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Each tenant subscribed to compute, network, and storage resources in a cloud is entitled to a given SLA. One tenant may have higher SLA requirements than another based on business model or organizational hierarchy. For example, tenant A may have higher compute and network bandwidth requirements than tenant B, while tenant B may have a higher storage capacity requirement. The main design objective is to ensure that tenants within this environment properly receive their subscribed SLA's while their data, communication, and application environments are securely separated, protected, and isolated from other tenants (Figure 4).

Figure 4) Architecture overview.

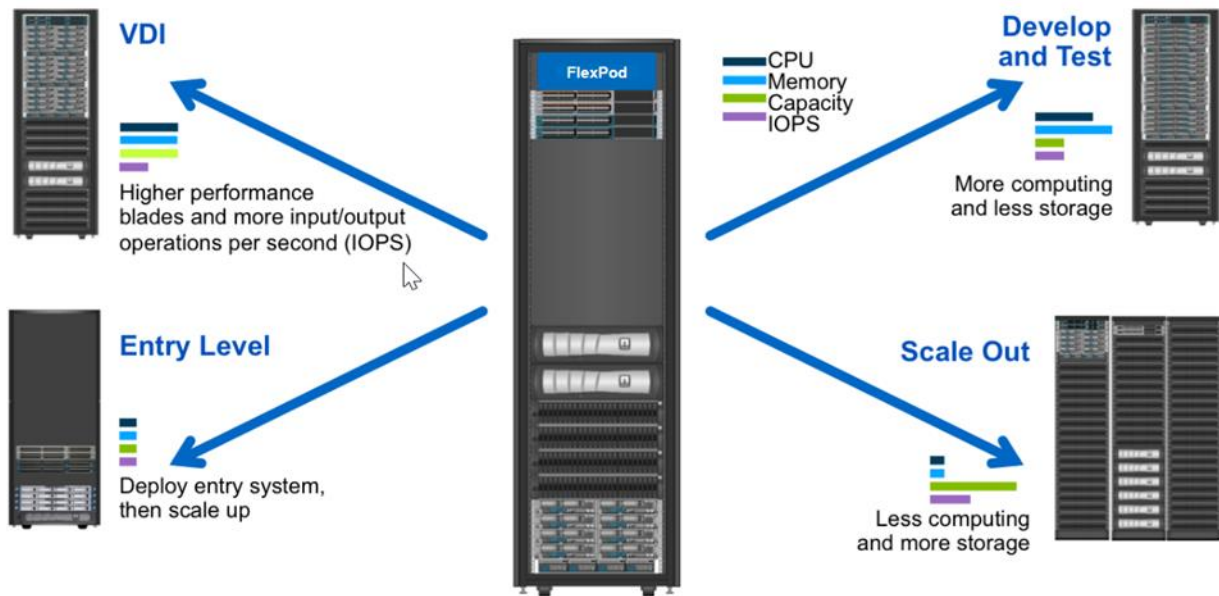


4.3 FlexPod Scale Up and Scale Out

FlexPod is a validated infrastructure solution built on the Cisco Unified Computing System (Cisco UCS), Cisco Nexus data center switches, NetApp FAS storage components, and software from a range of partners. FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. A FlexPod design provides a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different business solutions. This document describes how to build several different solutions on top of FlexPod.

FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. Figure 5 shows a few FlexPod scaling options. FlexPod is a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different use cases.

Figure 5) Example of scaling a FlexPod configuration.



The FlexPod architecture and accompanying collateral, delivered by Cisco and NetApp, aid customers, partners, and field personnel with new FlexPod deployments. For customers who want standard scalable configurations, FlexPod is an excellent choice for their infrastructures because environmental requirements are defined.

4.4 NetApp Cloud Volumes ONTAP

NetApp Cloud Volumes ONTAP is a software-only storage service that runs NetApp ONTAP storage software. It provides non-disruptive, secure and proven NFS, CIFS, and iSCSI data management for the cloud. OnCommand Cloud Manager is required to deploy and manage Cloud Volumes ONTAP. Cloud Manager provides a simple interface for your ONTAP-based storage in the cloud and on premises.

Cloud Volumes ONTAP runs in a hyperscaler cloud environment, bringing intelligence and Data Fabric connectivity to hyperscaler storage volumes. Cloud Volumes ONTAP is partially limited by the performance of the underlying volumes managed by the cloud provider. The result is more manageable storage, and, in some cases, the caching capability of Cloud Volumes ONTAP offers a performance improvement. However, there are always some limitations in terms of input/output operations per second (IOPS) and latency due to the reliance on public cloud provider storage.

Cloud Volumes ONTAP Performance Considerations with Focus on Amazon Web Services Platform

Disk

Amazon Web Services (AWS) uses Elastic Block Store (EBS) volumes, which can be described as the type of disk used to back storage for Cloud Volumes ONTAP.

- General purpose SSD (gp2): SSD with scalable IOPS (standard + burst rate)
 - Most commonly selected option
 - 3 x IOPS/GB: scalable performance, 500GB = 1,500 IOPS; 1TB = 3,000 IOPS
 - Larger disk size enables better performance
- Provisioned IOPS SSD (io1): SSD with fixed IOPS

- Guaranteed IOPS
- High performance workloads
- More costly than other options
- Throughput optimized HDD (st1):
 - Ideal for streaming data workloads, such as NetApp SnapMirror® data replication software
- EBS magnetic/cold HDD (sc1): traditional spinning media:
 - Not recommended or supported

Instance Type

AWS instance types can be described as the hardware dedicated to the Cloud Volumes ONTAP install, including CPU, RAM, and network bandwidth:

- General purpose: balanced resources
 - T2: low cost; burst performance oriented
 - M4: balanced resources for many workloads
- Compute optimized: better for workloads that need more CPU resources (deduplication, compression, and compaction):
 - C4: latest generation Intel Xeon processors
- Memory optimized: better for workloads with large working sets (high file count, complex directory structure, and database workloads):
 - X1: optimized for large-scale, enterprise-class, and in-memory applications
 - R4: optimized for memory-intensive applications

Other Considerations

- Cloud Volumes ONTAP HA functionality:
 - Cloud Volumes ONTAP supports HA configurations

Note: There might be performance issues for some workloads; however, there are continuous improvements being made to this feature.
- Write speed for non-HA (single-node):
 - Normal. Data is written to NVRAM prior to being committed to disk. This is the safest option and should be used in most cases.
 - High. Data is simply left in active memory buffers and committed to disk. Data is not written to NVRAM; therefore, if an unplanned shutdown occurs, data might be lost. This option is not recommended and should only be used for workloads with transient data that can safely be lost.
- Usage profile – configurable options when you deploy Cloud Volumes ONTAP:
 - Highest performance. This option is recommended for applications requiring lowest latency.
 - Performance with efficiency. This is a good performance option with ONTAP storage efficiency (deduplication).
- Shared tenancy versus dedicated hardware:
 - Shared tenancy is the most common option; however, you might experience "noisy neighbor" symptoms from other AWS workloads. Dedicated hardware can prevent this symptom, but it is costs more.

4.5 FlexPod Managed Private Cloud

With FlexPod Managed Private Cloud, you get the advantage of traditional managed hosting, combined with all the benefits of cloud in a secure and dedicated environment, with data sovereignty. Automated

self-service provisioning through role-based access and real-time usage dashboards enable you to control resources and utilization. While you maintain ownership, a FlexPod expert partner manages and operates the infrastructure.

As a fully managed offering, FlexPod Managed Private Cloud avoids many of the technical and organizational obstacles that can impede private cloud adoption while offering alternatives to some of the security and regulatory concerns that many organizations face with public cloud solutions. This solution offers the advantages of traditional managed hosting, including cost reduction, customization, and faster enablement, then enhances them with cloud's rapid provisioning, upfront capex avoidance, better flexibility, and self-service capability. It serves all these benefits in a secure and dedicated environment, providing an attractive, lower risk path forward to the cloud.

Leveraging market-leading FlexPod converged infrastructure for a managed private cloud provides your organization with powerful options when advancing their cloud journey, can reduce gaps in IT skills and resources, and combines the benefits of a managed cloud with the high-performance advantage of FlexPod infrastructure. You can advance your cloud journey with multiple options for on-premises, off-premises, and collocated cloud platforms. Reduce gaps in skills and resources by leveraging partner skillsets and expertise, and reduce hassle with support for software updates, installations, and repair support. Drive superior performance for your business-critical applications with a trusted, high-performance converged infrastructure solution (Figure 6).

Figure 6) Place cloud at the source of demand with a FlexPod managed private cloud.



The new Managed Private Cloud solution enables select delivery partners to offer a managed service solution through FlexPod. This new solution enables channel partners to deploy business-critical applications and manage on-premises private clouds without IT infrastructure administrators having to adopt a new platform or learn new skills. With Managed Private Cloud, you can have FlexPod in your data center or in remote locations, but you manage it remotely, advancing cloud-capabilities for both partners and their customers. Work with a partner of your choice for options with customer provisioning and deployment of applications onto your managed platform.

5 Conclusion

NetApp and Cisco combined their technologies to develop FlexPod, a new solution for virtual computing. FlexPod is a defined set of hardware and software products that serves as an integrated infrastructure stack for all virtualization solutions. Combining best-in-class compute, network, and storage elements,

FlexPod offers an uncompromising infrastructure solution for customers who want to deploy virtualized, nonvirtualized, and hybrid infrastructure solutions for a variety of enterprise applications.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- FlexPod Datacenter
<https://www.netapp.com/us/media/tr-3884.pdf>
- FlexPod Datacenter Specifications
<https://www.netapp.com/us/media/tr-4036.pdf>
- FlexPod Datacenter with ACI Solutions Guide
<https://www.netapp.com/us/media/tr-4399.pdf>

Version History

Version	Date	Document Version History
Version 1.0	July 2019	Initial release.

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 © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

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.