



NetApp Verified Architecture

# FlexPod Datacenter FedRAMP Readiness with VMware vSphere 6.0, HyTrust CloudControl and DataControl

NVA Design and Deployment

Arvind Ramakrishnan, Karthick Radhakrishnan, NetApp  
December 2016 | NVA-0031 | Version 1.0

Reviewed by



## TABLE OF CONTENTS

<b>1</b>	<b>Program Summary</b> .....	<b>3</b>
<b>2</b>	<b>Solution Overview</b> .....	<b>3</b>
2.1	Use Case Summary .....	6
<b>3</b>	<b>Secure Multi-Tenancy and Data Management</b> .....	<b>6</b>
<b>4</b>	<b>Technology Requirements</b> .....	<b>7</b>
4.1	Hardware Requirements .....	7
4.2	Software Requirements .....	7
<b>5</b>	<b>Deployment Procedures</b> .....	<b>8</b>
5.1	Cisco Nexus Configuration.....	8
5.2	NetApp Storage Configuration—Part I .....	10
5.3	Cisco UCS Configuration .....	17
5.4	NetApp Storage Configuration—Part II .....	37
5.5	VMware vSphere 6.0 Setup .....	38
5.6	Cisco Nexus 1000v VSM Configuration .....	42
5.7	HyTrust CloudControl Installation and Configuration .....	45
5.8	HyTrust DataControl Installation and Configuration .....	54
5.9	Set VM Restart Priority .....	58
<b>6</b>	<b>FedRAMP Security Controls</b> .....	<b>59</b>
<b>7</b>	<b>Conclusion</b> .....	<b>59</b>
	<b>References</b> .....	<b>59</b>

## LIST OF TABLES

Table 1)	Hardware requirements .....	7
Table 2)	Software requirements. ....	7
Table 3)	iSCSI LIFs for iSCSI IQN. ....	37
Table 4)	vNIC iSCSI IQNs for fabric A and fabric B.....	37
Table 5)	VMKernel ports.....	38
Table 6)	FedRAMP moderate impact security controls. ....	59

## LIST OF FIGURES

Figure 1)	FlexPod Datacenter architecture.....	4
Figure 2)	FlexPod Datacenter design with HyTrust CloudControl. ....	5
Figure 3)	FlexPod Datacenter design with HyTrust DataControl. ....	6

## 1 Program Summary

FlexPod® Datacenter is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), Cisco Nexus family of switches, and NetApp® fabric-attached storage (FAS) systems. The FlexPod Datacenter solution is tailored to be the infrastructure backbone of various public/private and hybrid cloud environments.

The Federal Risk and Authorization Management Program (FedRAMP) is a United States government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. This approach uses a “do once, use many times” framework that saves an estimated 30% to 40% of government costs and the time and staff required to conduct redundant agency security assessments. FedRAMP is the result of close collaboration with cybersecurity and cloud experts from the General Services Administration, National Institute of Standards and Technology, Department of Homeland Security, Department of Defense, National Security Agency, Office of Management and Budget, the Federal Chief Information Officer Council and its working groups, and private industry.

For more information about FedRAMP, go to [www.fedramp.gov](http://www.fedramp.gov).

The FlexPod Datacenter solution was assessed for FedRAMP readiness. This document provides a detailed overview of the information system that was audited as part of the program.

## 2 Solution Overview

FlexPod Datacenter lets you consolidate several siloed or independent workloads and host them on the same physical infrastructure. Although this capability reduces the overall cost of implementing a data center, it comes with the added challenges of secure management of data belonging to different workloads and tenants.

The FlexPod Datacenter solution described in this document addresses these challenges. The base infrastructure is built using the following guides:

- [FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1 Design Guide](#)
- [FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1 Deployment Guide](#)

The additional steps to implement secure multi-tenancy are covered in this document.

The assessment report of FlexPod Datacenter for FedRAMP readiness is available at [www.coalfire.com/netapp-whitepaper](http://www.coalfire.com/netapp-whitepaper).

### FlexPod Datacenter

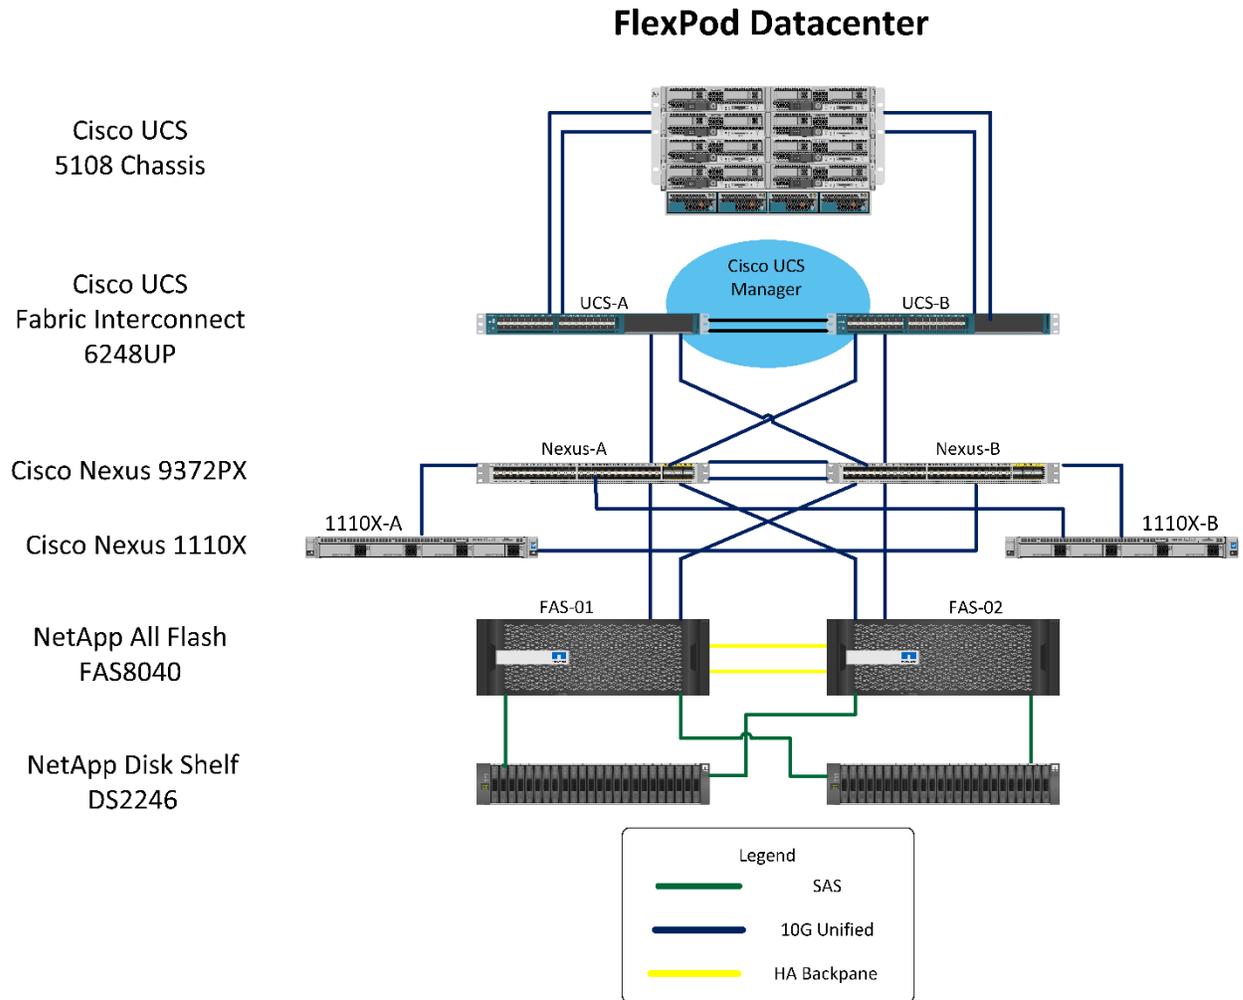
The FlexPod Datacenter solution combines NetApp storage systems, Cisco Unified Computing System (Cisco UCS) servers, and Cisco Nexus fabric into a single flexible architecture. The FlexPod integrated infrastructure leads in efficiency and flexibility, scaling and flexing as needed, with validated designs that reduce deployment time, project risk, and the cost of IT.

In this deployment, the FlexPod Datacenter solution is treated as the core infrastructure-as-a-service component. In addition, the HyTrust CloudControl and HyTrust DataControl software suites enable FlexPod readiness for FedRAMP environments.

More information about the FlexPod Datacenter design is available in the [Design Guide](#).

Figure 1 represents the FlexPod Datacenter solution that was used in this report.

Figure 1) FlexPod Datacenter architecture.



## HyTrust CloudControl

The HyTrust CloudControl (HTCC) appliance is a secure hardened operating system built on the CentOS platform. HTCC serves as a proxy to the VMware vCenter management platform and enhances the platform with forensic grade logging and advanced administrative control. With HTCC's granular role-based access control (RBAC), administrative functions can be easily set to control permissions on a virtual object level. HTCC applies smart labels to enable further segregation of virtual objects by constraining object access based on certain labels. HTCC offers two-person integrity for destructive actions on virtual machines through the secondary approval function.

HTCC offers automated compliance validation and implementation for VMware ESXi hosts. Variables can be set and then applied to each host so that the host security posture complies with the required baseline of the standard(s). HTCC can use Intel Trusted Execution Technology (TXT) to enable trusted compute pools by labeling hosts and configuring virtual machines to run only on a host that has the correct label.

HTCC is deployed in the mapped mode and as a cluster configuration. In the mapped mode, all the hosts that need to be protected by HTCC are configured with a published IP (PIP). This PIP is used by users and clients to access the hosts.

HTCC is deployed as a transparent proxy and sits between the users and all management interfaces to the protected hosts. From this central vantage point, HTCC intercepts and logs all administrative requests

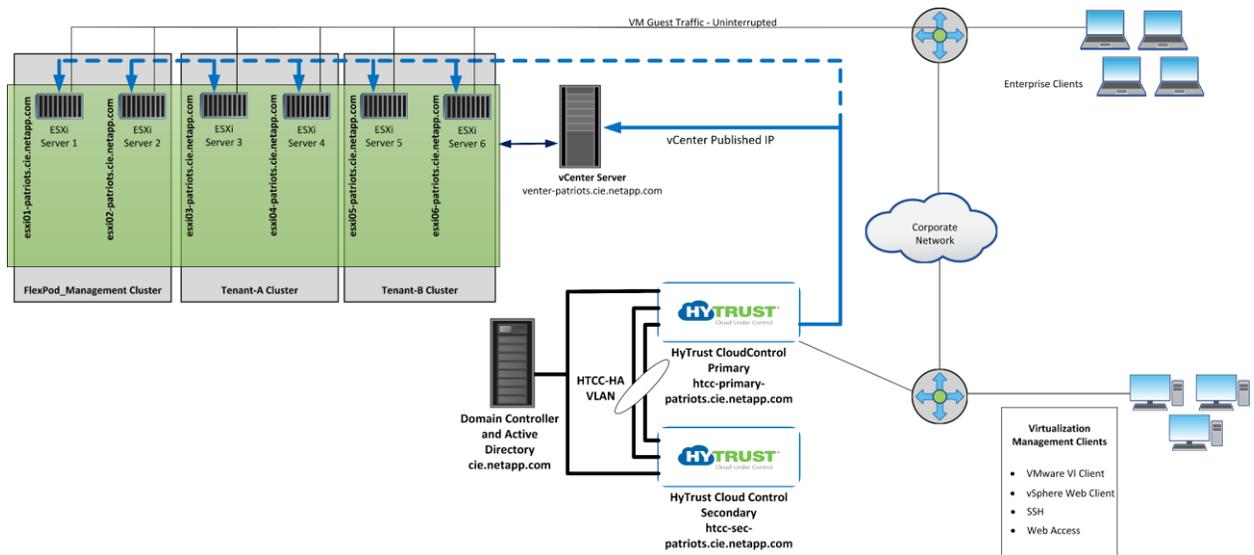
coming through the PIP and enforces role- and resource-based policies that protect workloads from unauthorized access.

A private cluster network is set up on a dedicated VLAN for the HTCC cluster nodes to communicate with each other.

HTCC is integrated with an Active Directory/Domain instance to apply the user identities and privileges extended to each user. Also, HTCC provides a set of access controls to users that can be configured to have specific privileges in the virtual infrastructure space.

Figure 2 is a representation of HyTrust CloudControl integrated with the VMware virtual infrastructure, which is deployed on FlexPod.

Figure 2) FlexPod Datacenter design with HyTrust CloudControl.



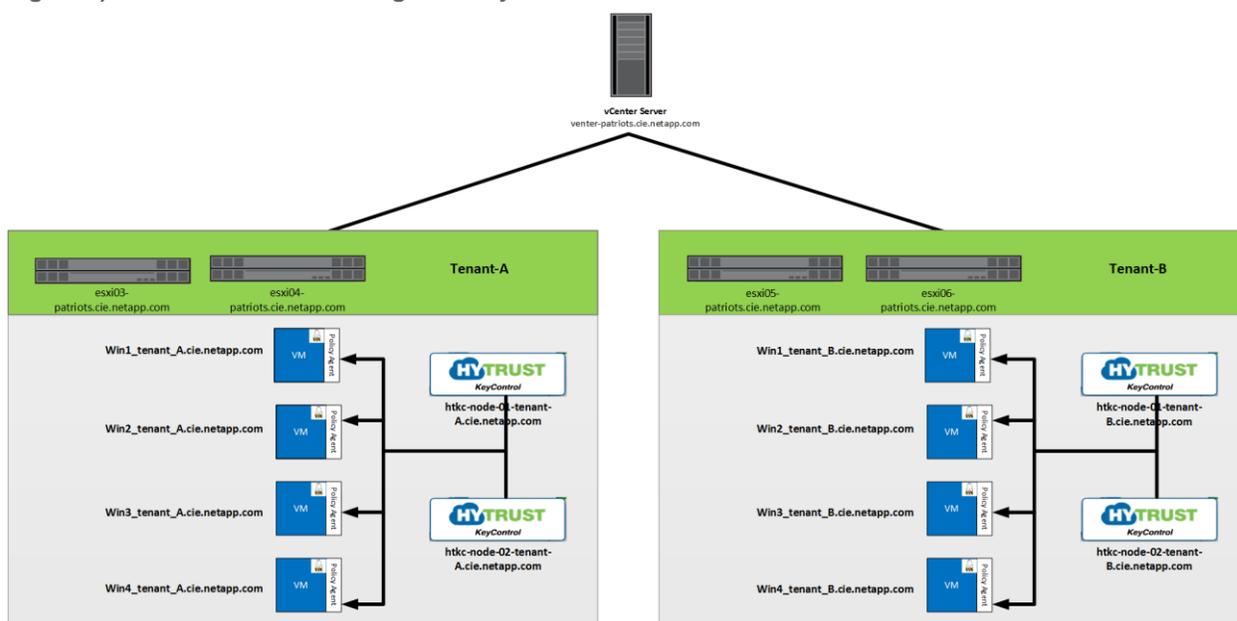
## HyTrust DataControl

HyTrust DataControl (HTDC) provides encryption of virtual machine data while it is in motion and at rest. HTDC is deployed as a virtual appliance in a high-availability configuration. The solution includes three critical components: Key Control, Policy Engine, and Policy Agent.

Administrators can configure or modify encryption policies through the Policy Engine; the Policy Engine then collects the rules for the Key Controller. The Key Controller in turn makes sure that the Policy Agent (which resides in the VM/workload) executes on these policies by managing encryption key creation, renewals, and destruction.

Figure 3 illustrates how HTDC protects the data of the VMs running on various tenants within the FlexPod environment.

Figure 3) FlexPod Datacenter design with HyTrust DataControl.



## 2.1 Use Case Summary

The following use cases were identified as the most significant and essential requirements in a cloud service provider scenario and were implemented using the FlexPod Datacenter solution and HyTrust:

- Secure multi-tenancy
- Secure data management

## 3 Secure Multi-Tenancy and Data Management

The FlexPod Datacenter solution provides secure multi-tenancy and data management capabilities. This capability is achieved by implementing logical separation and access control within each component of the solution. A brief description of how this capability was achieved within each layer of the FlexPod stack follows.

### Storage

Multiple logical storage controllers were created by using storage virtual machines (SVMs) to cater to the storage needs of tenants. Each tenant was mapped to an SVM for all its storage requirements, and the resources that were assigned to an SVM were not shared with any other coexisting SVM. The SVMs also had their own set of users and administrators. Multiple logical interfaces (LIFs) were configured for each SVM to handle iSCSI and NFS data traffic on dedicated VLANs.

### Network

Each tenant within the FlexPod platform was provided with dedicated VLANs for all management and data traffic. These VLANs are not routable and therefore there is no communication between any two VLANs in the infrastructure. All the traffic using the data VLANs was kept private to the infrastructure and each tenant was provided with a management VLAN for external access. This network configuration is set up on both Cisco Nexus 9000 and 1000v switches.

### Compute

The Cisco UCS infrastructure is split into multiple organizations and each organization corresponds to a tenant. The network interface configurations for the vNICs and iSCSI vNICs were set up as per the

VLANs assigned to each tenant. Thus, network segregation is also achieved at the Cisco UCS layer. The Cisco UCS organizations are also configured to have their own MAC pools, iSCSI initiator pools, boot policies, BIOS policies, and so on. This configuration provides complete isolation from other tenants and their resources.

### VMware vSphere

Built on top of the above defined infrastructure, the VMware vSphere environment is configured to use the assigned resources. Dedicated ESXi clusters are created for each tenant and the hosts within each tenant have their own datastores and network resources. VM-to-host affinity rules are configured to make sure that the VMs do not accidentally attempt to move to a host assigned to a different tenant. Even if such a move is attempted, the operation will fail because of unavailability or lack of connectivity to required port groups and datastores.

### Secure Data Management

The secure data management capabilities are provided by NetApp Data ONTAP® software and HyTrust DataControl. Data ONTAP provides secure multi-tenancy for the storage resources and DataControl encrypts the virtual machine data of all the virtual machines in a tenant. In addition to DataControl, NetApp Storage Encryption drives can also provide drive encryption for data at rest.

## 4 Technology Requirements

This section covers the technology requirements for the FlexPod Datacenter FedRAMP Readiness with VMware vSphere 6.0, HyTrust CloudControl and DataControl solution.

### 4.1 Hardware Requirements

Table 1 lists the hardware components required to implement the solution.

Table 1) Hardware requirements.

Layer	Hardware	Quantity
Compute	Cisco UCS 6248UP fabric interconnect	2
	Cisco UCS 5108 chassis	1
	Cisco UCS B200 M4 blades with VIC 1240	6
Network	Cisco Nexus 9372PX	2
	Cisco Nexus 1110-x	2
Storage	All Flash FAS8040	1 HA pair
	Disk shelf: DS2246 with 24 x 800GB SSD	2

### 4.2 Software Requirements

Table 2 lists the software components required to implement the solution.

Table 2) Software requirements.

Layer	Device	Version
Compute	Cisco UCS fabric interconnects 6200 series, Cisco UCS B-200 M4	3.1(1h)

Layer	Device	Version
	Cisco eNIC	2.3.0.7
	Cisco fNIC	1.6.0.25
Network	Cisco Nexus 9000 NX-OS	7.0(3)I1(3)
	Cisco Nexus 1000V	5.2(1)SV3(1.5b)
	Cisco Nexus 1110-X	5.2(1)SP1(7.3a)
Storage	NetApp AFF8040	Data ONTAP 8.3.2P2
Software	VMware vSphere ESXi	6.0U1b
	VMware vCenter	6.0U1b
	NetApp Virtual Storage Console (VSC)	6.2
	HyTrust CloudControl	5.0
	HyTrust DataControl	3.2.1

## 5 Deployment Procedures

The base FlexPod Datacenter deployment is performed as described in [http://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_esxi60u1\\_n9k.html](http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi60u1_n9k.html).

This deployment guide focuses only on the base deployment of the FlexPod solution. The procedures that follow focus on the additional configuration that is necessary to implement the use cases.

In this deployment, three tenants—management, tenant-A, and tenant-B—are deployed.

The management tenant hosts all the VMs that are necessary to manage the entire infrastructure; for example, vCenter, HyTrust CloudControl, and so on. Tenants A and B host VMs belonging to them and each has its instance of HyTrust DataControl, which encrypts the VMs within that tenant only.

### 5.1 Cisco Nexus Configuration

#### Create VLANs

##### Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

1. Create additional VLANs for handling the data and management traffic of tenants.

```

config t
vlan <<var_htcc_ha_vlan>>
name HTCC-HA-VLAN
exit
vlan <<var_vm_traffic_vlan_tenant_a>>
name VM-Traffic-VLAN-Tenant-A
exit
vlan <<var_iscsi_a_vlan_tenant_a>>
name iSCSI-A-VLAN-Tenant-A
exit
vlan <<var_iscsi_b_vlan_tenant_a>>
name iSCSI-B-VLAN-Tenant-A
exit
vlan <<var_nfs_vlan_tenant_a>>
name NFS-VLAN-Tenant-A
exit
vlan <<var_vmotion_vlan_tenant_a>>

```

```

name vMotion-VLAN-Tenant-A
exit
vlan <<var_vm_traffic_vlan_tenant_b>>
name VM-Traffic-VLAN-Tenant-B
exit
vlan <<var_iscsi_a_vlan_tenant_b>>
name iSCSI-A-VLAN-Tenant-B
exit
vlan <<var_iscsi_b_vlan_tenant_b>>
name iSCSI-B-VLAN-Tenant-B
exit
vlan <<var_nfs_vlan_tenant_b>>
name NFS-VLAN-Tenant-B
exit
vlan <<var_vmotion_vlan_tenant_b>>
name vMotion-VLAN-Tenant-B
exit

copy run start

```

## Configure Port Channel Parameters

### Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

1. Add the previously created VLAN to the existing port channels.

```

interface port-channel 10
switchport trunk allowed vlan add
<<var_htcc_ha_vlan_id>>,<<var_vm_traffic_vlan_tenant_a>>,<<var_iscsi_a_vlan_tenant_a>>,<<var_iscsi_b_vlan_tenant_a>>,<<var_nfs_vlan_tenant_a>>,<<var_vmotion_vlan_tenant_a>>,<<var_vm_traffic_vlan_tenant_b>>,<<var_iscsi_a_vlan_tenant_b>>,<<var_iscsi_b_vlan_tenant_b>>,<<var_nfs_vlan_tenant_b>>,<<var_vmotion_vlan_tenant_b>>
exit

interface port-channel 11
switchport trunk allowed vlan add
<<var_iscsi_a_vlan_tenant_a>>,<<var_iscsi_b_vlan_tenant_a>>,<<var_nfs_vlan_tenant_a>>,<<var_iscsi_a_vlan_tenant_b>>,<<var_iscsi_b_vlan_tenant_b>>,<<var_nfs_vlan_tenant_b>>

interface port-channel 12
switchport trunk allowed vlan add
<<var_iscsi_a_vlan_tenant_a>>,<<var_iscsi_b_vlan_tenant_a>>,<<var_nfs_vlan_tenant_a>>,<<var_iscsi_a_vlan_tenant_b>>,<<var_iscsi_b_vlan_tenant_b>>,<<var_nfs_vlan_tenant_b>>

interface port-channel 111
switchport trunk allowed vlan add
<<var_htcc_ha_vlan_id>>,<<var_vm_traffic_vlan_tenant_a>>,<<var_iscsi_a_vlan_tenant_a>>,<<var_iscsi_b_vlan_tenant_a>>,<<var_nfs_vlan_tenant_a>>,<<var_vmotion_vlan_tenant_a>>,<<var_vm_traffic_vlan_tenant_b>>,<<var_iscsi_a_vlan_tenant_b>>,<<var_iscsi_b_vlan_tenant_b>>,<<var_nfs_vlan_tenant_b>>,<<var_vmotion_vlan_tenant_b>>
exit

interface port-channel 112
switchport trunk allowed vlan add
<<var_htcc_ha_vlan_id>>,<<var_vm_traffic_vlan_tenant_a>>,<<var_iscsi_a_vlan_tenant_a>>,<<var_iscsi_b_vlan_tenant_a>>,<<var_nfs_vlan_tenant_a>>,<<var_vmotion_vlan_tenant_a>>,<<var_vm_traffic_vlan_tenant_b>>,<<var_iscsi_a_vlan_tenant_b>>,<<var_iscsi_b_vlan_tenant_b>>,<<var_nfs_vlan_tenant_b>>,<<var_vmotion_vlan_tenant_b>>
exit

copy run start

```

**Note:** Perform a shut and no shut on the port-channel interfaces if the VLAN additions do not show up in the show vpc brief output.

## 5.2 NetApp Storage Configuration—Part I

In addition to the procedures described in the [Cisco Valid Designs \(CVD\)](#) document, complete the procedures in this section to set up the storage system. Three SVMs will be created, one for each tenant.

### Create Aggregates

An aggregate containing the root volume is created during the Data 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 will contain. You can also create multiple aggregates and allocate them to different SVMs. In this deployment, the same aggregates are shared across SVMs.

To create the aggregates required for this solution, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_fas_01 -nodes <<var_node01>> -diskcount 15
aggr create -aggregate aggr1_fas_02 -nodes <<var_node02>> -diskcount 15
```

**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 `aggr1_fas_01` and `aggr1_fas_02` are online.

2. Disable NetApp Snapshot<sup>®</sup> copies for the two data aggregates that you created in step 1.

```
system node run -node <<var_node01>> aggr options aggr1_fas_01 nosnap on
system node run -node <<var_node02>> aggr options aggr1_fas_02 nosnap on
```

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

```
system node run -node <<var_node01>> snap delete -A -a -f aggr1_fas_01
system node run -node <<var_node02>> snap delete -A -a -f aggr1_fas_01
```

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>>
```

### Set Up Management Broadcast Domain

To set up the default broadcast domain for the management network interfaces, complete the following step:

1. Run the following commands:

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

### Create Broadcast Domains in Clustered Data ONTAP

1. To create a data broadcast domain with an MTU of 9,000 and a management broadcast domain with an MTU of 1,500, complete the following steps:

```
broadcast-domain create -broadcast-domain IB_MGMT -mtu 1500
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_NFS_Tenant_A -mtu 9000
broadcast-domain create -broadcast-domain Infra_NFS_Tenant_B -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
```

```

broadcast-domain create -broadcast-domain Infra_iSCSI-A_Tenant_A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A_Tenant_B -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B_Tenant_A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B_Tenant_B -mtu 9000

```

## Create LACP Interface Groups

The ifgrp interface group 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, complete the following step:

1. Run the following commands:

```

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

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

ifgrp show

```

**Note:** All interfaces must be in the down state 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 a clustered Data ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), complete the following step:

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

```

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, the ports 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 will 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 for all the tenants and add them to their respective broadcast domains, complete the following step:

1. Run the following commands:

```

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_IB_MGMT_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_IB_MGMT_vlan_id>>
broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <<var_node01>>:a0a-
<<var_IB_MGMT_vlan_id>>, <<var_node02>>:a0a-<<var_IB_MGMT_vlan_id>>

```

```

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 Infra_NFS -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_NFS_vlan_id_Tenant_A>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_NFS_vlan_id_Tenant_A>>
broadcast-domain add-ports -broadcast-domain Infra_NFS_Tenant_A -ports <<var_node01>>:a0a-
<<var_NFS_vlan_id_Tenant_A>>, <<var_node02>>:a0a-<<var_NFS_vlan_id_Tenant_A>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_NFS_vlan_id_Tenant_B>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_NFS_vlan_id_Tenant_B>>
broadcast-domain add-ports -broadcast-domain Infra_NFS_Tenant_B -ports <<var_node01>>:a0a-
<<var_NFS_vlan_id_Tenant_B>>, <<var_node02>>:a0a-<<var_NFS_vlan_id_Tenant_B>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-A_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-A_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_node01>>:a0a-
<<var_iSCSI_A_vlan_id>>, <<var_node02>>: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-B_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_node01>>:a0a-
<<var_iSCSI_B_vlan_id>>, <<var_node02>>:a0a-<<var_iSCSI-B_vlan_id>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-A_vlan_id_Tenant_A>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-A_vlan_id_Tenant_A>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A_Tenant_A -ports <<var_node01>>:a0a-
<<var_iSCSI_A_vlan_id_Tenant_A>>, <<var_node02>>:a0a-<<var_iSCSI-A_vlan_id_Tenant_A>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-B_vlan_id_Tenant_A>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-B_vlan_id_Tenant_A>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B_Tenant_A -ports <<var_node01>>:a0a-
<<var_iSCSI_B_vlan_id_Tenant_A>>, <<var_node02>>:a0a-<<var_iSCSI-B_vlan_id_Tenant_A>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-A_vlan_id_Tenant_B>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-A_vlan_id_Tenant_B>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A_Tenant_B -ports <<var_node01>>:a0a-
<<var_iSCSI_A_vlan_id_Tenant_B>>, <<var_node02>>:a0a-<<var_iSCSI-A_vlan_id_Tenant_B>>

network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iSCSI-B_vlan_id_Tenant_B>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iSCSI-B_vlan_id_Tenant_B>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B_Tenant_B -ports <<var_node01>>:a0a-
<<var_iSCSI_B_vlan_id_Tenant_B>>, <<var_node02>>:a0a-<<var_iSCSI-B_vlan_id_Tenant_B>>

```

## Set Up Storage Virtual Machine

### Create Storage Virtual Machine

To create an infrastructure SVM for all the tenants, complete the following steps:

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

1. Run the `vserver create` command:

```
vserver create -vserver Infra-SVM-MGMT -rootvolume rootvol -aggregate aggr1_fas_01 -rootvolume-
security-style unix
```

2. Select the SVM data protocols to configure.

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

3. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp Virtual Storage Console (VSC).

```
vserver modify -vserver Infra-SVM-MGMT -aggr-list aggr1_fas_01, aggr1_fas_02
```

#### 4. Enable and run the NFS protocol in the Infra-SVM.

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

#### 5. Enable the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in.

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

#### 6. Repeat steps 1 to 5 to create `Infra-SVM-Tenant-A` and `Infra-SVM-Tenant-B`.

**Note:** Make sure to replace the correct SVM name in the `-vserver` option in the previous steps.

### Create Load-Sharing Mirror of SVM Root Volume

To create a load-sharing mirror of an SVM root volume for all the tenants, 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-MGMT -volume rootvol_m01 -aggregate aggr1_fas_01 -size 1GB -type  
DP  
volume create -vserver Infra-SVM-MGMT -volume rootvol_m02 -aggregate aggr1_fas_02 -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-MGMT/rootvol -destination-path //Infra-SVM-  
MGMT/rootvol_m01 -type LS -schedule 15min  
snapmirror create -source-path //Infra-SVM-MGMT/rootvol -destination-path //Infra-SVM-  
MGMT/rootvol_m02 -type LS -schedule 15min
```

#### 4. Initialize the mirroring relationship.

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

#### 5. Repeat steps 1 to 4 to create a load-sharing mirror of the SVM root volume for `Infra-SVM-Tenant-A` and `Infra-SVM-Tenant-B`.

**Note:** Make sure to replace the correct SVM name in the `-vserver` option in the previous steps.

### Create iSCSI Service

To create the iSCSI service, complete the following step:

#### 1. 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 Infra-SVM-MGMT  
iscsi create -vserver Infra-SVM-Tenant-A  
iscsi create -vserver Infra-SVM-Tenant-B  
iscsi show
```

### Configure NFSv3

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

#### 1. Create a 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-MGMT -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-MGMT -policyname default -ruleindex 2 -
clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra-SVM-Tenant-A -policyname default -ruleindex 1 -
clientmatch <<var_esxi_host3_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Infra-SVM-Tenant-A -policyname default -ruleindex 2 -
clientmatch <<var_esxi_host4_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra-SVM-Tenant-B -policyname default -ruleindex 1 -
clientmatch <<var_esxi_host5_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Infra-SVM-Tenant-B -policyname default -ruleindex 2 -
clientmatch <<var_esxi_host6_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-MGMT -volume rootvol -policy default
volume modify -vserver Infra-SVM-Tenant-A -volume rootvol -policy default
volume modify -vserver Infra-SVM-Tenant-B -volume rootvol -policy default
```

## Create FlexVol Volumes

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

### 1. Run the following commands:

```
volume create -vserver Infra-SVM-MGMT -volume infra_datastore_1 -aggregate aggr1_fas_02 -size 1TB
-state online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-
snapshot-space 0

volume create -vserver Infra-SVM-MGMT -volume infra_swap -aggregate aggr1_fas_01 -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-MGMT -volume esxi_boot -aggregate aggr1_fas_01 -size 500GB -
state online -policy default -space-guarantee none -percent-snapshot-space 0

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

volume create -vserver Infra-SVM-Tenant-A -volume infra_datastore_1_tenant_A -aggregate
aggr1_fas_02 -size 1TB -state online -policy default -junction-path /infra_datastore_1_A -space-
guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM-MGMT -volume infra_swap_tenant_A -aggregate aggr1_fas_01 -size
100GB -state online -policy default -junction-path /infra_swap_tenant_A -space-guarantee none -
percent-snapshot-space 0 -snapshot-policy none

volume create -vserver Infra-SVM-MGMT -volume esxi_boot_tenant_A -aggregate aggr1_fas_01 -size
500GB -state online -policy default -space-guarantee none -percent-snapshot-space 0

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

volume create -vserver Infra-SVM-Tenant-A -volume infra_datastore_1_tenant_B -aggregate
aggr1_fas_02 -size 1TB -state online -policy default -junction-path /infra_datastore_1_B -space-
guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM-MGMT -volume infra_swap_tenant_B -aggregate aggr1_fas_01 -size
100GB -state online -policy default -junction-path /infra_swap_tenant_B -space-guarantee none -
percent-snapshot-space 0 -snapshot-policy none

volume create -vserver Infra-SVM-MGMT -volume esxi_boot_tenant_B -aggregate aggr1_fas_01 -size
500GB -state online -policy default -space-guarantee none -percent-snapshot-space 0

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

## Create Boot LUNs for ESXi Hosts

To create boot LUNs for ESXi hosts, complete the following steps:

### 1. Turn off automatic Snapshot copies on the volume.

```
volume modify -vserver Infra-SVM-MGMT -volume esxi_boot -snapshot-policy none
volume modify -vserver Infra-SVM-Tenant_A -volume esxi_boot_tenant_a -snapshot-policy none
volume modify -vserver Infra-SVM-Tenant_B -volume esxi_boot_tenant_a -snapshot-policy none
```

### 2. Enable deduplication on the volume.

```
volume efficiency on -vserver Infra-SVM-MGMT -volume esxi_boot
volume efficiency on -vserver Infra-SVM-Tenant_A -volume esxi_boot_tenant_A
volume efficiency on -vserver Infra-SVM-Tenant_B -volume esxi_boot_tenant_B
```

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

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

lun create -vserver Infra-SVM-Tenant_A -volume esxi_boot_tenant_A -lun VM-Host-Infra-01_tenant_A
-size 15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM-Tenant_A -volume esxi_boot_tenant_A -lun VM-Host-Infra-02_tenant_A
-size 15GB -ostype vmware -space-reserve disabled

lun create -vserver Infra-SVM-Tenant_B -volume esxi_boot_tenant_B -lun VM-Host-Infra-01_tenant_B
-size 15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM-Tenant_B -volume esxi_boot_tenant_B -lun VM-Host-Infra-02_tenant_B
-size 15GB -ostype vmware -space-reserve disabled
```

## Create iSCSI LIFs

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

```
network interface create -vserver Infra-SVM-MGMT -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-MGMT -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-MGMT -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-MGMT -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 create -vserver Infra-SVM-Tenant-A -lif iscsi_lif01a_tenant_A -role data -data-
protocol iscsi -home-node <<var_node01>> -home-port a0a-<<var_iSCSI-A_vlan_id_tenant_A>> -address
<<var_node01_iscsi_lif01a_ip_tenant_A>> -netmask <<var_node01_iscsi_lif01a_mask_tenant_A>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-A -lif iscsi_lif01b_tenant_A -role data -data-
protocol iscsi -home-node <<var_node01>> -home-port a0a-<<var_iSCSI-B_vlan_id_tenant_A>> -address
<<var_node01_iscsi_lif01b_ip_tenant_A>> -netmask <<var_node01_iscsi_lif01b_mask_tenant_A>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-A -lif iscsi_lif02a_tenant_A -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iSCSI-A_vlan_id_tenant_A>> -address
<<var_node02_iscsi_lif02a_ip_tenant_A>> -netmask <<var_node02_iscsi_lif02a_mask_tenant_A>> -
```

```

status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-A -lif iscsi_lif02b_tenant_A -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iSCSI-B_vlan_id_tenant_A>> -address
<<var_node02_iscsi_lif02b_ip_tenant_A>> -netmask <<var_node02_iscsi_lif02b_mask_tenant_A>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-B -lif iscsi_lif01a_tenant_B -role data -data-
protocol iscsi -home-node <<var_node01>> -home-port a0a-<<var_iSCSI-A_vlan_id_tenant_B>> -address
<<var_node01_iscsi_lif01a_ip_tenant_B>> -netmask <<var_node01_iscsi_lif01a_mask_tenant_B>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-B -lif iscsi_lif01b_tenant_B -role data -data-
protocol iscsi -home-node <<var_node01>> -home-port a0a-<<var_iSCSI-B_vlan_id_tenant_B>> -address
<<var_node01_iscsi_lif01b_ip_tenant_B>> -netmask <<var_node01_iscsi_lif01b_mask_tenant_B>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-B -lif iscsi_lif02a_tenant_B -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iSCSI-A_vlan_id_tenant_B>> -address
<<var_node02_iscsi_lif02a_ip_tenant_B>> -netmask <<var_node02_iscsi_lif02a_mask_tenant_B>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM-Tenant-B -lif iscsi_lif02b_tenant_B -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iSCSI-B_vlan_id_tenant_B>> -address
<<var_node02_iscsi_lif02b_ip_tenant_B>> -netmask <<var_node02_iscsi_lif02b_mask_tenant_B>> -
status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface show

```

## Create NFS LIFs

To create an NFS LIF for all the tenants, run the following commands:

```

network interface create -vserver Infra-SVM-MGMT -lif nfs_infra_datastore_1 -role data -data-
protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> -address
<<var_node01_nfs_lif_infra_datastore_1_ip>> -netmask
<<var_node01_nfs_lif_infra_datastore_1_mask>> -status-admin up -failover-policy broadcast-domain-
wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM-MGMT -lif nfs_infra_swap -role data -data-protocol
nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> -address
<<var_node02_nfs_lif_infra_swap_ip>> -netmask <<var_node02_nfs_lif_infra_swap_mask>> -status-
admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM-Tenant-A -lif nfs_infra_datastore_1_tenant_A -role
data -data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id_tenant_A>> -
address <<var_node01_nfs_lif_infra_datastore_1_ip_tenant_A>> -netmask
<<var_node01_nfs_lif_infra_datastore_1_mask_tenant_A>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM-Tenant-A -lif nfs_infra_swap_tenant_A -role data -
data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id_tenant_A>> -address
<<var_node02_nfs_lif_infra_swap_ip_tenant_A>> -netmask
<<var_node02_nfs_lif_infra_swap_mask_tenant_A>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM-Tenant-B -lif nfs_infra_datastore_1_tenant_B -role
data -data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id_tenant_B>> -
address <<var_node01_nfs_lif_infra_datastore_1_ip_tenant_B>> -netmask
<<var_node01_nfs_lif_infra_datastore_1_mask_tenant_B>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM-Tenant-B -lif nfs_infra_swap_tenant_B -role data -
data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id_tenant_B>> -address
<<var_node02_nfs_lif_infra_swap_ip_tenant_B>> -netmask
<<var_node02_nfs_lif_infra_swap_mask_tenant_B>> -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 for all the tenants, complete the following steps:

1. Run the following commands:

```
network interface create -vserver Infra-SVM-MGMT -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-MGMT -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 Infra-SVM-MGMT
Enter a new password: <<var_password>>
Enter it again: <<var_password>>

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

4. Repeat the previous steps for all other tenants.

## 5.3 Cisco UCS Configuration

To configure the Cisco UCS environment to host multiple tenants, complete the procedures in this section.

### Create Organizations

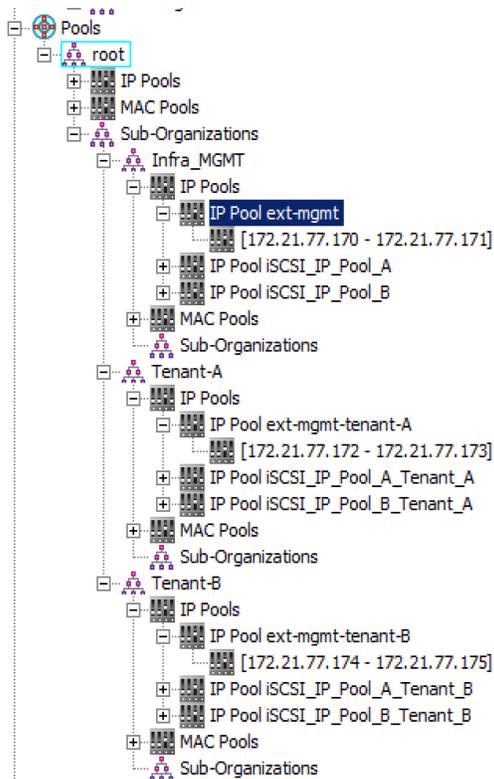
1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. From the right pane, click New and select Create Organization from the drop-down list.
3. Enter a name for the organization and provide a description

**Note:** Create a total of three organizations to host the three tenants.

### Create IP Pools for In-Band KVM Access

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Go to Pools > root> Sub Organizations and select an organization.
3. Right-click IP Pools and select Create IP Pool.
4. Enter a name for the pool and click Next.
5. Click Add.
6. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information. Click Next.
7. Click Finish to create the IP block.

**Note:** Create a total of three IPs for the three tenants within their respective organizations.



## Create MAC Address Pools

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Go to Pools > root > Sub Organizations and select the organization.

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

3. Right-click MAC Pools under the organization and select Create MAC Pool.
4. Enter a name for the MAC pool.
5. Optional: Enter a description for the MAC pool.
6. Click Next.
7. Click Add.
8. Specify a starting MAC address.

**Note:** For this FlexPod solution, it is recommended to modify the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses and the tenant to which they belong.

Tenant Name	Recommended Octet Modification
Management	0A
Tenant-A	AA
Tenant-B	BA

9. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Click OK.
10. Click Finish.
11. In the confirmation message, click OK.
12. Right-click MAC Pools under the organization and select Create MAC Pool.

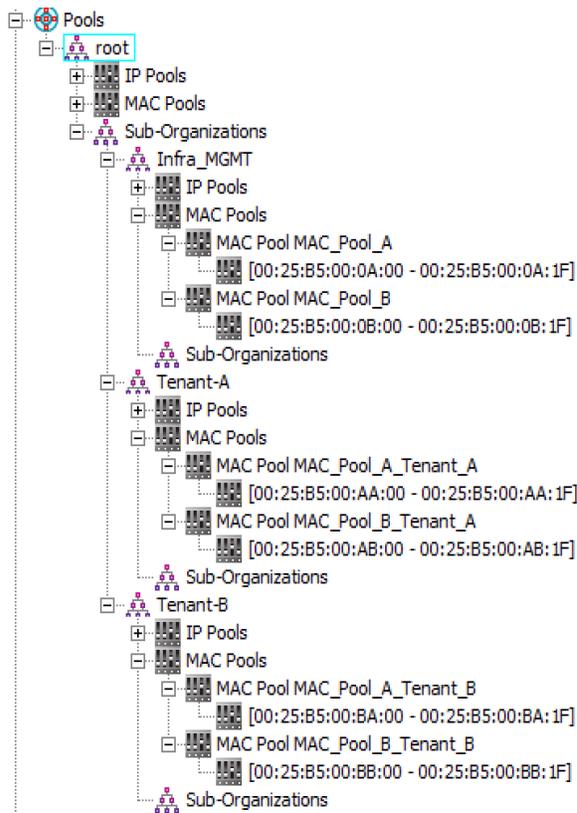
13. Enter `MAC_Pool_B` as the name of the MAC pool.
14. Optional: Enter a description for the MAC pool.
15. Click Next.
16. Click Add.
17. Specify a starting MAC address.

**Note:** For this FlexPod solution, the recommendation is to modify the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric B addresses and the tenant to which they belong.

Tenant Name	Recommended Octet Modification
Management	0B
Tenant-A	AB
Tenant-B	BB

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

**Note:** Create a total of six MAC pools within their respective organizations, two for each tenant, one each in fabric A and fabric B.





## Create IP Pools for iSCSI Boot

1. In Cisco UCS Manager, click the LAN tab on the left.

2. Go to Pools > root > Sub Organizations and select the organization.

**Note:** Two IP pools are created per tenant, one for each switching fabric.

3. Right-click IP Pools under the tenant/organization and select Create IP Pool.

4. Enter a name for the IP pool for fabric A.

5. Optional: Enter a description of the IP pool.

6. Select Sequential for Assignment Order.

7. Click Next.

8. Click Add.

9. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.

**Note:** In this deployment, the IP ranges for the iSCSI boot for each tenant and fabric are on a dedicated subnet and VLAN.

10. Set the size to enough addresses to accommodate the servers.

11. Click OK.

12. Click Finish.

13. Right-click IP Pools under the tenant/organization and select Create IP Pool.

14. Enter a for the name of the IP pool for fabric B.

15. Optional: Enter a description of the IP pool.

16. Select Sequential for Assignment Order.

17. Click Next.

18. Click Add.

19. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.

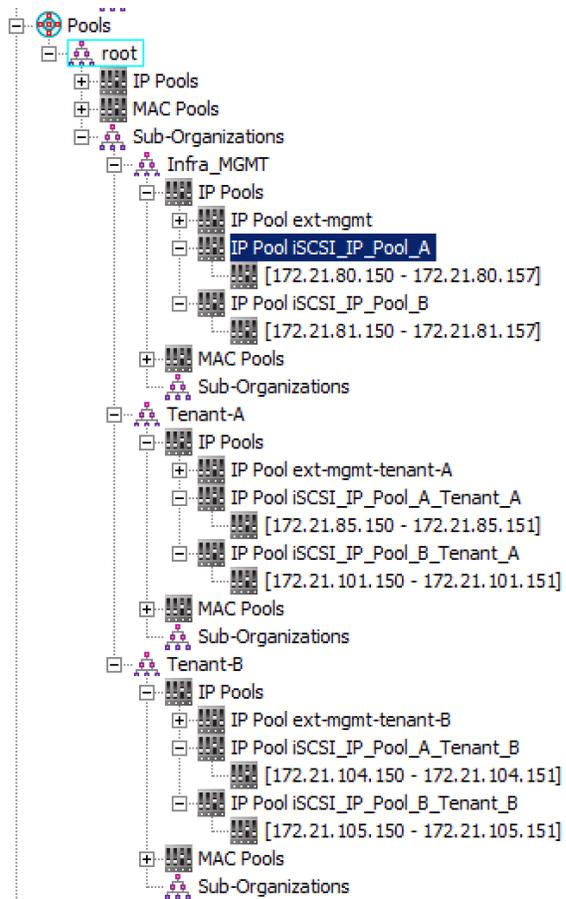
**Note:** In this deployment, the IP ranges for the iSCSI boot for each tenant and fabric are on a dedicated subnet and VLAN.

20. Set the size to enough addresses to accommodate the servers.

21. Click OK.

22. Click Finish.

**Note:** Create a total of six iSCSI IP pools within their respective organizations, two for each tenant, one each in fabric A and fabric B.



## Create UUID Suffix Pool

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Go to Pools > root > Sub Organizations and select the organization.
3. Right-click UUID Suffix Pools and select Create UUID Suffix Pool.
4. Enter a name for the UUID suffix pool.
5. Optional: Enter a description for the UUID suffix pool.
6. Keep the prefix at the derived option.
7. Click Next.
8. Click Add to add a block of UUIDs.
9. Keep the From field at the default setting.

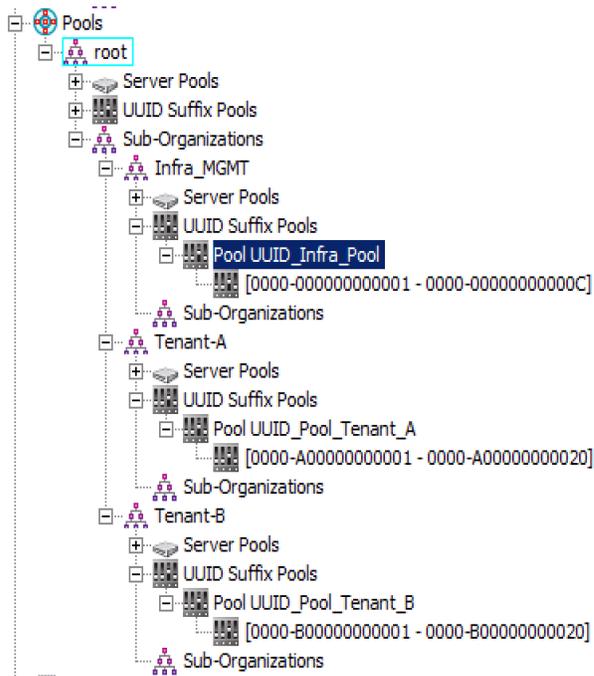
For this FlexPod solution, it is recommended to modify the fifth digit of the UUID to identify all the UUIDs based on their tenants.

Tenant Name	UUID recommendation
Management	0000- <b>0</b> ....
Tenant-A	0000- <b>A</b> ....
Tenant-B	0000- <b>B</b> ....

10. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
11. Click OK.
12. Click Finish.

13. Click OK.

**Note:** Create a total of three UUID pools within their respective organizations.



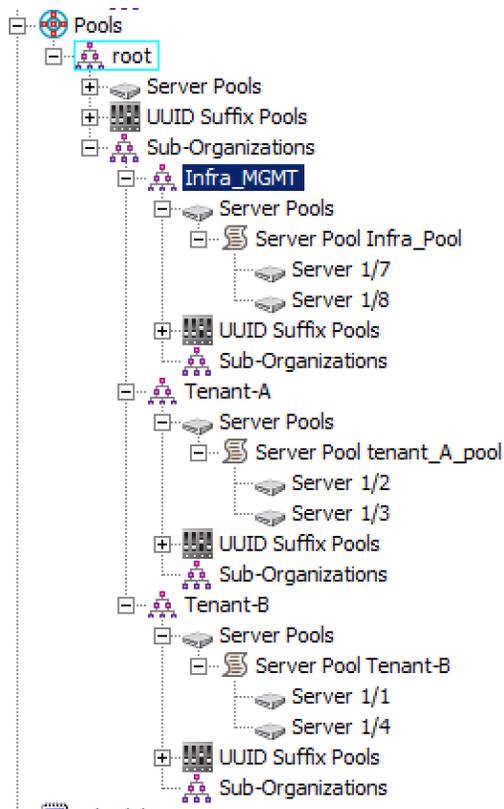
## Create Server Pool

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Go to Pools > root > Sub Organizations and select the organization.
3. Right-click Server Pools and select Create Server Pool.
4. Enter a name for the server pool.
5. Optional: Enter a description for the server pool.
6. Click Next.
7. Select two (or more) servers to be used for the cluster/tenant and click >> to add them to the server pool.

**Note:** Perform this step for the tenant A and tenant B clusters.

8. Click Finish.
9. Click OK.

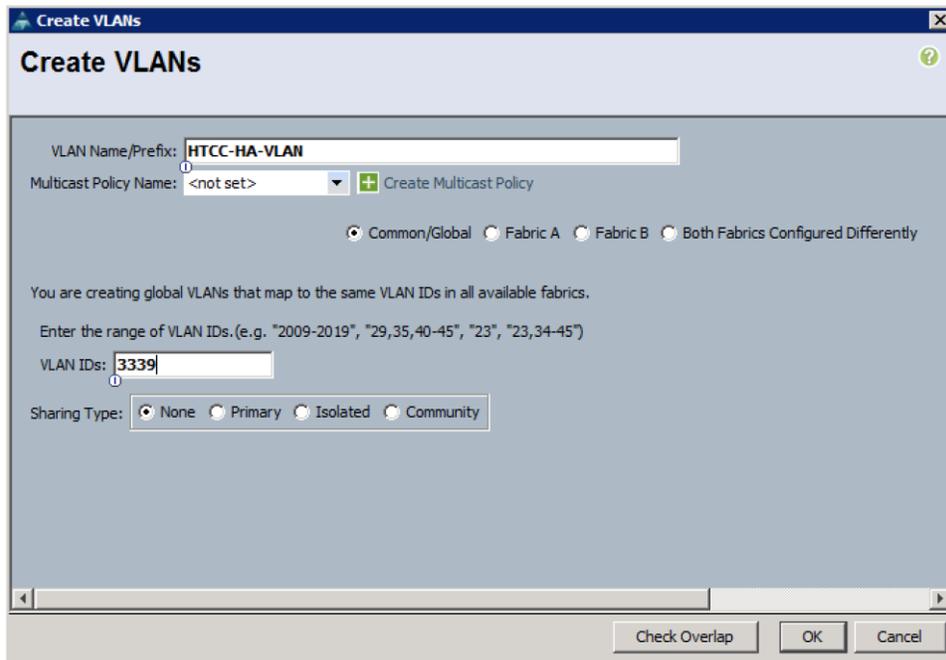
**Note:** In total, three server pools need to be created for the three tenants within their respective organizations.



## Create VLANs

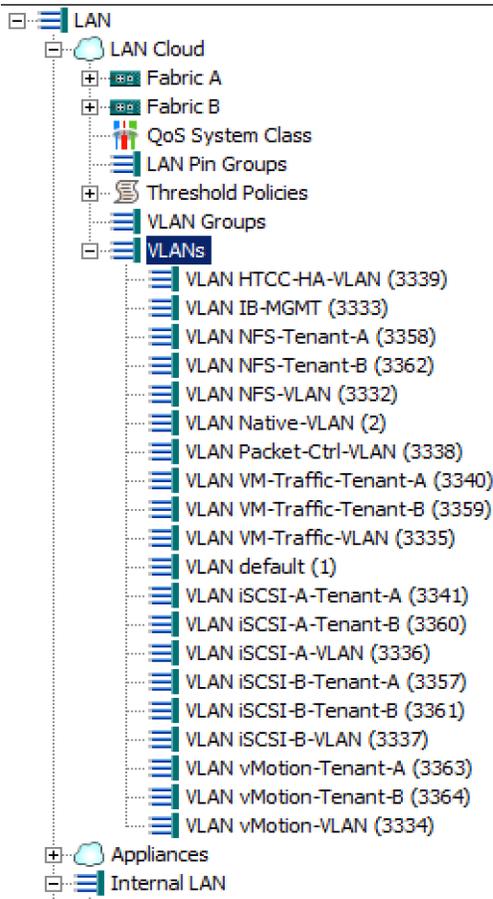
To create additional VLANs, complete the following steps.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs and select Create VLANs.
4. Enter `HTCC-HA-VLAN` as the name of the VLAN.
5. Keep the `Common/Global` option selected for the scope of the VLAN.
6. Enter the VLAN ID for the HTCC-HA-VLAN.
7. Keep the sharing type as None.



8. Click OK and then click OK again.
9. Right-click VLANs and select Create VLANs.
10. Enter VM-Traffic-Tenant-A as the name of the VLAN.
11. Keep the Common/Global option selected for the scope of the VLAN.
12. Enter the VLAN ID for the VM-Traffic-Tenant-A VLAN.
13. Keep the Sharing Type as None.
14. Click OK and then click OK again.
15. Right-click VLANs and select Create VLANs.
16. Enter iSCSI-A-Tenant-A as the name of the VLAN.
17. Keep the Common/Global option selected for the scope of the VLAN.
18. Enter the VLAN ID for the iSCSI-A-Tenant-A VLAN.
19. Keep the Sharing Type as None.
20. Click OK and then click OK again.
21. Right-click VLANs and select Create VLANs.
22. Enter iSCSI-B-Tenant-A as the name of the VLAN.
23. Keep the Common/Global option selected for the scope of the VLAN.
24. Enter the VLAN ID for the iSCSI-B-Tenant-A VLAN.
25. Keep the Sharing Type as None.
26. Click OK and then click OK again.
27. Right-click VLANs and select Create VLANs.
28. Enter NFS-Tenant-A as the name of the VLAN.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the VLAN ID for the NFS-Tenant-A VLAN.
31. Keep the Sharing Type as None.

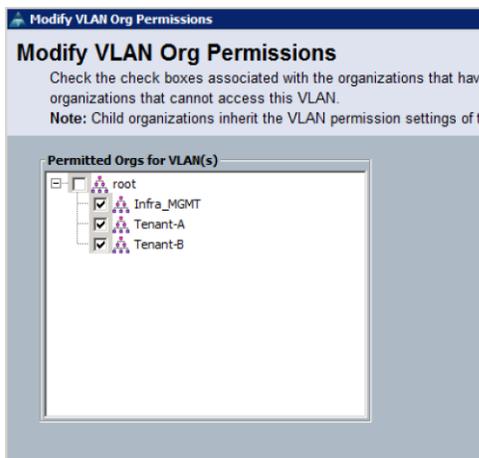
32. Click OK and then click OK again.
33. Right-click VLANs and select Create VLANs.
34. Enter `vMotion-Tenant-A` as the name of the VLAN.
35. Keep the `Common/Global` option selected for the scope of the VLAN.
36. Enter the VLAN ID for the vMotion-Tenant-A VLAN.
37. Keep the Sharing Type as None.
38. Click OK and then click OK again.
39. Right-click VLANs and select Create VLANs.
40. Enter `VM-Traffic-Tenant-B` as the name of the VLAN.
41. Keep the `Common/Global` option selected for the scope of the VLAN.
42. Enter the VLAN ID for the VM-Traffic-Tenant-B VLAN.
43. Keep the Sharing Type as None.
44. Click OK and then click OK again.
45. Right-click VLANs and select Create VLANs.
46. Enter `iSCSI-A-Tenant-B` as the name of the VLAN.
47. Keep the `Common/Global` option selected for the scope of the VLAN.
48. Enter the VLAN ID for the iSCSI-A-Tenant-B VLAN.
49. Keep the Sharing Type as None.
50. Click OK and then click OK again.
51. Right-click VLANs and select Create VLANs.
52. Enter `iSCSI-B-Tenant-B` as the name of the VLAN.
53. Keep the `Common/Global` option selected for the scope of the VLAN.
54. Enter the VLAN ID for the iSCSI-B-Tenant-B VLAN.
55. Keep the Sharing Type as None.
56. Click OK and then click OK again.
57. Right-click VLANs and select Create VLANs.
58. Enter `NFS-Tenant-B` as the name of the VLAN.
59. Keep the `Common/Global` option selected for the scope of the VLAN.
60. Enter the VLAN ID for the NFS-Tenant-B VLAN.
61. Keep the Sharing Type as None.
62. Click OK and then click OK again.
63. Right-click VLANs and select Create VLANs.
64. Enter `vMotion-Tenant-B` as the name of the VLAN.
65. Keep the `Common/Global` option selected for the scope of the VLAN.
66. Enter the VLAN ID for the vMotion-Tenant-B VLAN.
67. Keep the Sharing Type as None.
68. Click OK and then click OK again.



## Modify VLAN Organization Permissions

Assign the VLANs to the desired tenant/organizations; some of the VLANs will need to be shared with other tenants and most of the VLANs are restricted to a single organization.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click LAN > LAN Cloud > VLANs.
3. Select a VLAN from the list and on the right pane click Modify VLAN Org Permissions.
4. Select the desired tenant/organization(s) and click OK.

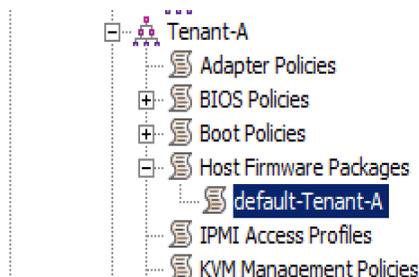


5. Refer to the following table for the VLAN assignment.

VLAN Name	Organization(s)
IB-MGMT	Management, Tenant-A, Tenant-B
Native-VLAN	Management, Tenant-A, Tenant-B
Packet-Ctrl-VLAN	Management, Tenant-A, Tenant-B
HTCC-HA-VLAN	Management
NFS-VLAN	Management
VM-Traffic-VLAN	Management
iSCSI-A-VLAN	Management
iSCSI-B-VLAN	Management
vMotion-VLAN	Management
NFS-Tenant-A	Tenant-A
VM-Traffic-Tenant-A	Tenant-A
iSCSI-A-Tenant-A	Tenant-A
iSCSI-B-Tenant-A	Tenant-A
vMotion-Tenant-A	Tenant-A
NFS-Tenant-B	Tenant-B
VM-Traffic-Tenant-B	Tenant-B
iSCSI-A-Tenant-B	Tenant-B
iSCSI-B-Tenant-B	Tenant-B
vMotion-Tenant-B	Tenant-B

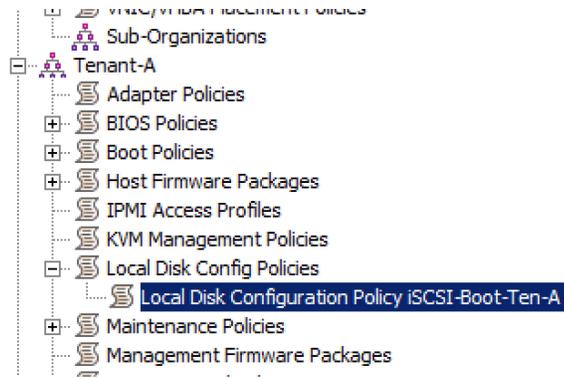
## Create Host Firmware Package

1. Create a host firmware package within each of the three tenants by referring to the [FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1 CVD](#).



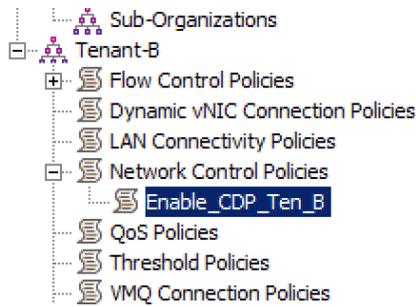
## Create Local Disk Configuration Policy

1. Create a local disk configuration policy within each of the three tenants/organizations by referring to the [CVD](#).



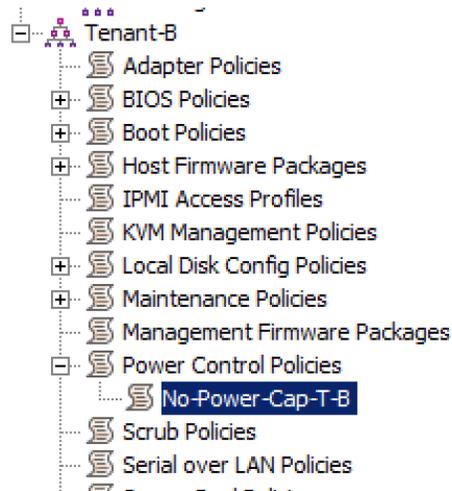
## Create Network Control Policy for Cisco Discovery Protocol (CDP)

1. Create a network control policy within each of the three tenants/organizations by referring to the [CVD](#).



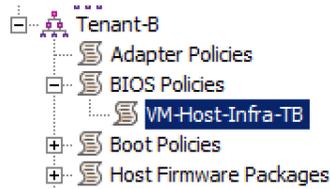
## Create Power Control Policy

1. Create a power control policy within each of the three tenants/organizations by referring to the [CVD](#).



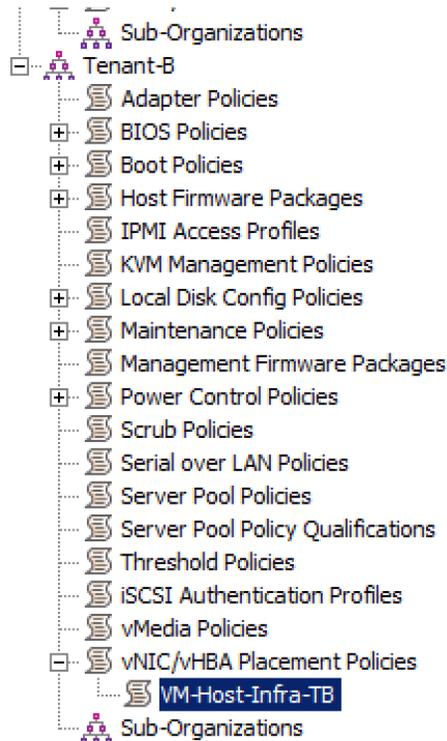
## Create Server BIOS Policy

1. Create a server BIOS policy within each of the three tenants/organizations by referring to the [CVD](#).



## Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

1. Create a vNIC/vHBA placement policy within each of the three tenants/organizations by referring to the [CVD](#).



## Update the Default Maintenance Policy

1. Update the Default Maintenance Policy for each of the three tenants/organizations by referring to the [CVD](#).

## Create vNIC Templates

**Note:** If vNIC templates are already created, you can modify the templates to match the configuration detailed in this section.

## Create Data vNIC Templates

**Note:** Two data vNICs need to be created for each tenant, one in each fabric.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Go to Policies > root > Sub Organizations and select an organization.
3. Right-click vNIC Templates and select Create vNIC Template.

4. Enter a name for the vNIC template.
5. Keep Fabric A selected.
6. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template as the template type.
9. Under VLANs, select the required VLANs by referring to the following table.

Tenant	VLANs for vNIC
Management	HTCC-HA-VLAN, IB-MGMT, NFS-VLAN, Native-VLAN, VM-Traffic-Tenant-A, VM-Traffic-Tenant-B, VM-Traffic-VLAN, vMotion-VLAN
Tenant-A	IB-MGMT, NFS-Tenant-A, Native-VLAN, VM-Traffic-Tenant-A, vMotion-Tenant-A
Tenant-B	IB-MGMT, NFS-Tenant-B, Native-VLAN, VM-Traffic-Tenant-B, vMotion-Tenant-B

10. Set Native-VLAN as the native VLAN.
11. For MTU, enter 9000.
12. In the MAC Pool list, select the MAC pool for the tenant and fabric A.
13. Select the Network Control Policy list created for the tenant/organization.
14. Click OK to create the vNIC template.
15. Click OK.
16. In the navigation pane, select the LAN tab.
17. Go to Policies > root > Sub Organizations. Select the same organization in which the previous vNIC template for fabric A was created.
18. Right-click vNIC Templates and select Create vNIC Template.
19. Enter a name for the vNIC template.
20. Select Fabric B.
21. Do not select the Enable Failover checkbox.
22. Under Target, make sure that the VM checkbox is not selected.
23. Select Updating Template as the template type.
24. Under VLANs, select the necessary VLANs by referring to the following table:

Tenant	VLANs for vNIC
Management	HTCC-HA-VLAN, IB-MGMT, NFS-VLAN, Native-VLAN, VM-Traffic-Tenant-A, VM-Traffic-Tenant-B, VM-Traffic-VLAN, vMotion-VLAN
Tenant-A	IB-MGMT, NFS-Tenant-A, Native-VLAN, VM-Traffic-Tenant-A, vMotion-Tenant-A
Tenant-B	IB-MGMT, NFS-Tenant-B, Native-VLAN, VM-Traffic-Tenant-B, vMotion-Tenant-B

25. Set default as the native VLAN.
  26. For MTU, enter 9000.
  27. In the MAC Pool list, select the MAC pool for the tenant and fabric B.
  28. Select the Network Control Policy list created for the tenant/organization.
  29. Click OK to create the vNIC template.
  30. Click OK.
- Repeat the previous steps for the other tenants/organizations.

## Create iSCSI vNIC Templates

**Note:** Two iSCSI vNICs need to be created for each tenant, one in each fabric.

1. Select the LAN tab on the left.
2. Go to Policies > root > Sub Organizations and select an organization.
3. Right-click vNIC Templates and select Create vNIC Template.
4. Enter a name for the iSCSI vNIC template.
5. Leave Fabric A selected. Do not select the Enable Failover checkbox.
6. Under Target, make sure that the VM checkbox is not selected.
7. Select Updating Template for the template type.
8. Under VLANs, select the iSCSI VLAN for fabric A. Refer to the following table:

Tenant	VLAN for iSCSI vNIC
Management	iSCSI-A-VLAN
Tenant-A	iSCSI-A-Tenant-A
Tenant-B	iSCSI-A-Tenant-B

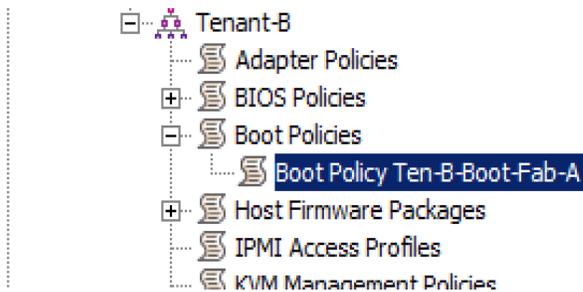
9. Set the selected VLAN as the native VLAN.
10. Under MTU, enter 9000.
11. From the MAC Pool list, select the MAC pool for the tenant and fabric A.
12. Select the Network Control Policy list created for the tenant/organization.
13. Click OK to complete creating the vNIC template.
14. Click OK.
15. Select the LAN tab on the left.
16. Select Policies > root > Sub Organizations. Select the same organization in which the previous vNIC template for fabric A was created.
17. Right-click vNIC Templates and select Create vNIC Template.
18. Enter a name for the vNIC template.
19. Select Fabric B. Do not select the Enable Failover checkbox.
20. Under Target, make sure that the VM checkbox is not selected.
21. Select Updating Template for the template type.
22. Under VLANs, select the iSCSI VLAN for fabric B. Refer to the following table.

Tenant	VLAN for iSCSI vNIC
Management	iSCSI-B-VLAN
Tenant-A	iSCSI-B-Tenant-A
Tenant-B	iSCSI-B-Tenant-B

23. Set the selected VLAN as the native VLAN.
24. Under MTU, enter 9000.
25. From the MAC Pool list, select the MAC pool for the tenant and fabric-B.
26. Select the Network Control Policy list created for the tenant/organization.
27. Click OK to complete creating the vNIC template.
28. Click OK.

## Create Boot Policies

1. Create a boot BIOS policy within each of the three tenants by referring to the [CVD](#).



## Create Service Profile Template

To create service profile templates for each tenant/organization, complete the following steps.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Go to Policies > root > Sub Organizations and select an organization.
3. Right-click the organization and select Create Service Profile Template.
4. Enter a name for the service profile template. The procedure described in the steps below configures the service profile template to boot from storage node 1 on fabric A.
5. Select the Updating Template option.
6. Under UUID, select the UUID pool created in the organization and click Next.

## Configure Storage Provisioning

1. If there are servers with no physical disks, select the iSCSI-Boot local storage policy created in Organization. Otherwise, select the default Local Storage Policy.
2. Click Next.

## Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Expert option to configure the LAN connectivity.
3. Click the upper Add button to add a vNIC to the template.
4. In the Create vNIC dialog box, enter a name for the vNIC in fabric A.
5. Select the Use vNIC Template checkbox.
6. From the vNIC Template list, select the vNIC template created within the tenant/organization for fabric A.
7. From 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 a name for the vNIC in fabric B.
11. Select the Use vNIC Template checkbox.
12. From the vNIC Template list, select the vNIC template created within the tenant/organization for fabric B.
13. From 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 a name for the iSCSI vNIC in fabric A.
17. Select the Use vNIC Template checkbox.
18. From the vNIC Template list, select the iSCSI vNIC template created within the tenant/organization for fabric A.
19. From the Adapter Policy list, select VMWare.
20. Click OK to add this vNIC to the template.
21. Click the upper Add button to add a vNIC to the template.
22. In the Create vNIC dialog box, enter a name for the iSCSI vNIC in fabric B.
23. Select the Use vNIC Template checkbox.
24. From the vNIC Template list, select the iSCSI vNIC template created within the tenant/organization for fabric B.
25. From the Adapter Policy list, select VMWare.
26. Click OK to add this vNIC to the template.
27. Expand the iSCSI vNICs section (if not already expanded).
28. Select IQN-Pool under Initiator Name Assignment.
29. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
30. Enter a name for the iSCSI vNIC in fabric A.
31. Select the iSCSI vNIC created in step 16 for fabric A as the overlay vNIC.
32. Set the iSCSI Adapter Policy to default.
33. Select the iSCSI-A VLAN created for the tenant/organization.
34. Leave the MAC Address set to None.
35. Click OK.
36. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
37. Enter a name for the iSCSI vNIC in fabric B.
38. Select the iSCSI vNIC created in step 22 for fabric B as the overlay vNIC.
39. Set the iSCSI Adapter Policy to the default.
40. Select the iSCSI-B VLAN created for the tenant/organization.
41. Leave MAC Address set to None.
42. Click OK.
43. Click OK.
44. Review the table in the Networking page to make sure that all vNICs are created and click Next.

### **Configure Storage Options**

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

### **Configure Zoning Options**

1. Set no zoning options and click Next.

### Configure vNIC/vHBA Placement

1. From the Select Placement list, select the VM-Host-Infra placement policy created for the tenant/organization.
2. Select vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
  - a. vNIC-A
  - b. vNIC-B
  - c. iSCSI-vNIC-A
  - d. iSCSI-vNIC-B
3. Review the table to verify that all vNICs were assigned to the policy in the appropriate order and click Next.

### Configure vMedia Policy

1. Do not configure a vMedia Policy at this point.
2. Click Next.

### Configure Server Boot Order

1. Select the boot policy created previously for the organization.
2. In the Boot Order pane, select iSCSI vNIC for fabric A.
3. Click the Set iSCSI Boot Parameters button.
4. In the Set iSCSI Boot Parameters dialog box, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
5. Leave the Initiator Name Assignment dialog box to <not set> to use the single service profile initiator name defined in the previous steps.
6. Select the iSCSI IP pool created for fabric A as the initiator IP address policy.
7. Keep the iSCSI Static Target Interface button selected and click the  button at the bottom right.
8. Log in to the storage cluster management interface and run the `iscsi show` command.
9. Note or copy the iSCSI target name for the SVM that corresponds to this organization.

Tenant	Storage Virtual Machine
Management	Infra-SVM
Tenant-A	Infra-SVM-Tenant-A
Tenant-B	Infra-SVM-Tenant-B

10. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name of the appropriate SVM into the iSCSI Target Name field.
11. Enter the IP address of the LIF `iscsi_lif02a` of the corresponding tenant in the IPv4 Address field.
12. Click OK to add the iSCSI static target.
13. Keep the iSCSI Static Target Interface option selected and click the  button.
14. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name of the appropriate SVM (from step 10) into the iSCSI Target Name field.
15. Enter the IP address of the LIF `iscsi_lif01a` of the corresponding tenant in the IPv4 Address field.
16. Click OK.
17. Click OK.
18. In the Boot Order pane, select iSCSI vNIC for fabric B.

19. Click the Set iSCSI Boot Parameters button.
20. In the Set iSCSI Boot Parameters dialog box, set Leave Initiator Name Assignment to <not set>.
21. Select the iSCSI IP Pool created for fabric B for the Initiator IP Address Policy.
22. Keep the iSCSI Static Target Interface option selected and click the  button at the bottom right.
23. In the Create iSCSI Static Target window, paste the iSCSI target node name of the appropriate SVM (from step 10) into the iSCSI Target Name field.
24. Enter the IP address of the LIF iscsi\_lif02b of the corresponding tenant for the IPv4 address field.
25. Click OK to add the iSCSI static target.
26. Keep the iSCSI Static Target Interface option selected and click the  button.
27. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name of the appropriate SVM (from step 10) into the iSCSI Target Name field
28. Enter the IP address of iscsi\_lif01b in the IPv4 Address field.
29. Click OK.
30. Click OK.
31. 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.
32. Click Next to continue to the next section.

### **Configure Maintenance Policy**

1. Leave the default maintenance policy of the organization/tenant selected.
2. Click Next.

### **Configure Server Assignment**

1. From the Pool Assignment list, select the server pool created for the organization.
2. Optional: Select a server pool qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management at the bottom of the page and select default from the Host Firmware list.
5. Click Next.

### **Configure Operational Policies**

1. Select the BIOS policy created for the organization.
2. Expand Power Control Policy Configuration and select the power control policy created for the organization.
3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

**Note:** At least three service profile templates need to be created, one for each tenant/organization.

**Note:** An extra boot policy to boot from fabric B can also be created within each tenant/organization.

### **Create Service Profiles**

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organization. Select a service profile template.
3. Right-click the service profile template and select Create Service Profiles from Template.
4. Enter a prefix for the service profile.

5. Enter a numerical value as Name Suffix Starting Number.
6. Enter the number of instances of service profiles to be deployed.
7. Click OK to create the service profile(s).
8. Click OK in the confirmation message.

**Note:** Service Profiles need to be deployed within each organization using the service profile templates created within each organization.

## Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 3 and Table 4.

Table 3) iSCSI LIFs for iSCSI IQN.

Vserver	Target: WWPN (FC) or IQN (iSCSI)
Infra-SVM	
Infra-SVM-Tenant-A	
Infra-SVM-Tenant-B	

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

Table 4) vNIC iSCSI IQNs for fabric A and fabric B.

Cisco UCS Service Profile Name	Initiator: IQN (iSCSI)	Variables
VM-Host-Infra-01		<<var_vm_host_infra_01_iqn>>
VM-Host-Infra-02		<<var_vm_host_infra_02_iqn>>
VM-Host-Infra-Tenant-A-01		<<var_vm_host_infra_tenant_a_01_iqn>>
VM-Host-Infra-Tenant-A-02		<<var_vm_host_infra_tenant_a_02_iqn>>
VM-Host-Infra-Tenant-B-01		<<var_vm_host_infra_tenant_b_01_iqn>>
VM-Host-Infra-Tenant-B-02		<<var_vm_host_infra_tenant_b_02_iqn>>

**Note:** To obtain the iSCSI vNIC IQN information in the Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the iSCSI vNICs tab on the right. The initiator name is displayed at the top of the page under Service Profile Initiator Name.

## 5.4 NetApp Storage Configuration—Part II

### iSCSI Boot Configuration

1. Create igroups for LUN mapping.

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

igroup create -vserver Infra-SVM-Tenant-A -igroup VM-Host-Infra-01_tenant_A -protocol iscsi -
ostype vmware -initiator <<var_vm_host_infra_01_iqn_tenant_A>>
igroup create -vserver Infra-SVM-Tenant-A -igroup VM-Host-Infra-02_tenant_A -protocol iscsi -
```

```

ostype vmware -initiator <<var_vm_host_infra_02_iqn_tenant_A>>

igroup create -vserver Infra-SVM-Tenant-B -igroup VM-Host-Infra-01_tenant_B -protocol iscsi -
ostype vmware -initiator <<var_vm_host_infra_01_iqn_tenant_B>>
igroup create -vserver Infra-SVM-Tenant-B -igroup VM-Host-Infra-02_tenant_B -protocol iscsi -
ostype vmware -initiator <<var_vm_host_infra_02_iqn_tenant_B>>

```

**Note:** The initiator IQNs are available in Table 2.

## 2. Map boot LUNs to hosts.

```

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

lun map -vserver Infra-SVM-Tenant-A -volume esxi_boot -lun VM-Host-Infra-01_tenant_A -igroup VM-
Host-Infra-01_tenant_A -lun-id 0
lun map -vserver Infra-SVM-Tenant-A -volume esxi_boot -lun VM-Host-Infra-02_tenant_A -igroup VM-
Host-Infra-02_tenant_A -lun-id 0
lun map -vserver Infra-SVM-Tenant-B -volume esxi_boot -lun VM-Host-Infra-01_tenant_B -igroup VM-
Host-Infra-01_tenant_B -lun-id 0
lun map -vserver Infra-SVM-Tenant-B -volume esxi_boot -lun VM-Host-Infra-02_tenant_B -igroup VM-
Host-Infra-02_tenant_B -lun-id 0

```

## 5.5 VMware vSphere 6.0 Setup

In this deployment, a total of six ESXi hosts were deployed and were added to three clusters: management, tenant A, and tenant B. Each cluster was allocated two ESXi servers and each cluster represents a tenant.

Tenant	ESXi Hosts
Management	esxi01-patriots.cie.netapp.com esxi02-patriots.cie.netapp.com
Tenant-A	esxi03-patriots.cie.netapp.com esxi04-patriots.cie.netapp.com
Tenant-B	esxi05-patriots.cie.netapp.com esxi06-patriots.cie.netapp.com

Follow the procedures as described in the [CVD](#) to install ESXi 6.0 on the Cisco UCS blades. Additional configurations that are required for various sections are described in the sections that follow. For all sections not listed below, repeat the procedure as described in the [CVD](#) for all the ESXi hosts.

### Set Up VMkernel Ports and Virtual Switches

You must create the necessary VMkernel ports and vSwitches for the ESXi servers within each tenant.

Table 5 lists the VMkernel ports that need to be created in the ESXi hosts for each tenant using the corresponding VLAN IDs.

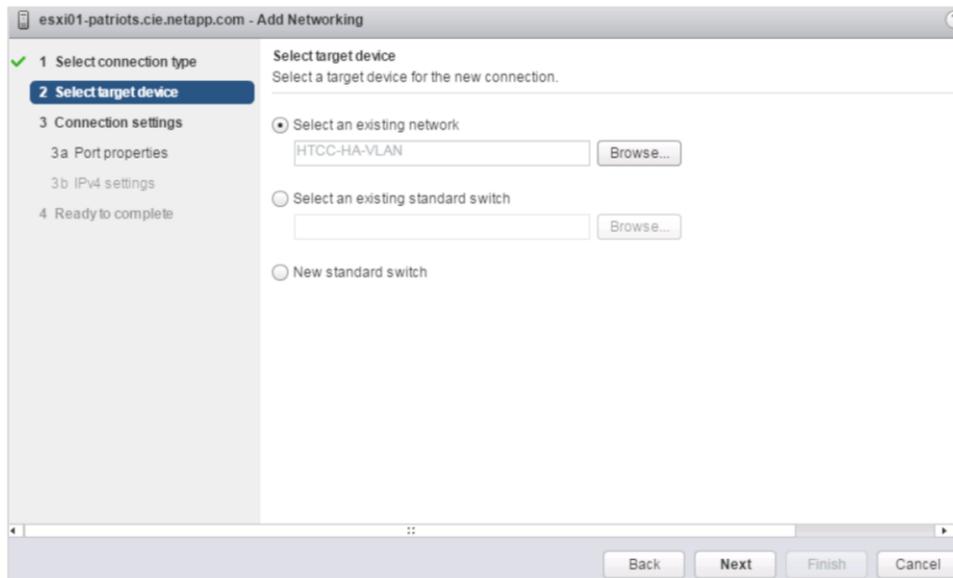
Table 5) VMKernel ports.

Tenant	VMkernel and vSwitch Details		
Management	vSwitch0	iScsiBootvSwitch	vSwitch1
	VMkernel-MGMT VMkernel-NFS VMkernel-vMotion VMkernel-HTCC-HA VMkernel-VMTraffic	VMkernel-iSCSI-A	VMkernel-iSCSI-B

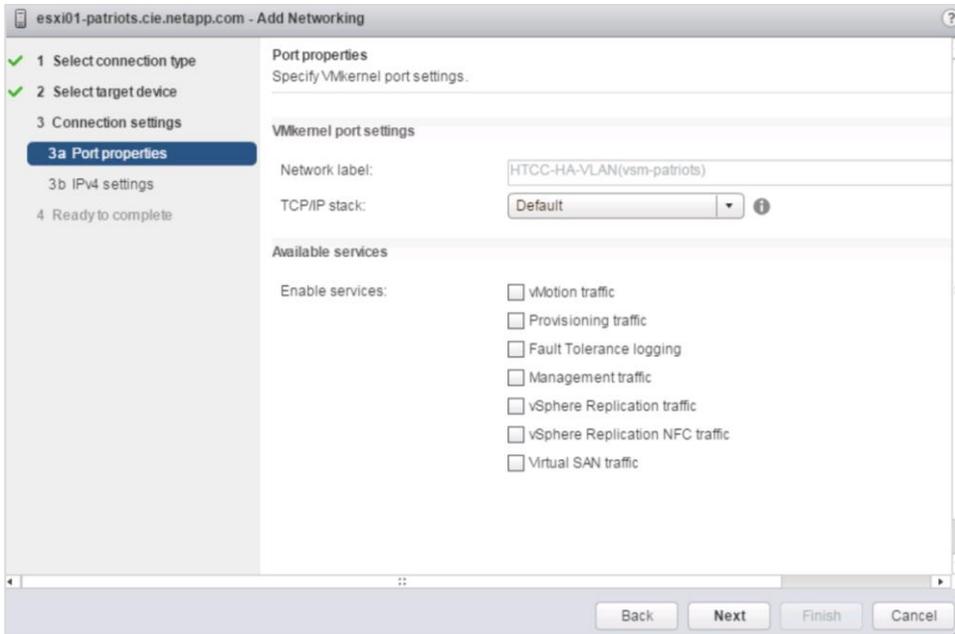
Tenant	VMkernel and vSwitch Details		
	VMkernel-VMTraffic-Tenant-A VMkernel-VMTraffic-Tenant-B		
Tenant-A	VMkernel-MGMT VMkernel-NFS-Tenant-A VMkernel-vMotion-Tenant-A VMkernel-VMTraffic-Tenant-A	VMkernel-iSCSI-A-Tenant-A	VMkernel-iSCSI-B-Tenant-A
Tenant-B	VMkernel-MGMT VMkernel-NFS-Tenant-B VMkernel-vMotion-Tenant-B VMkernel-VMTraffic-Tenant-B	VMkernel-iSCSI-A-Tenant-B	VMkernel-iSCSI-B-Tenant-B

The following procedure describes the addition of an HTCC VMkernel port to vSwitch0 of an ESXi host that will be placed in the management cluster later.

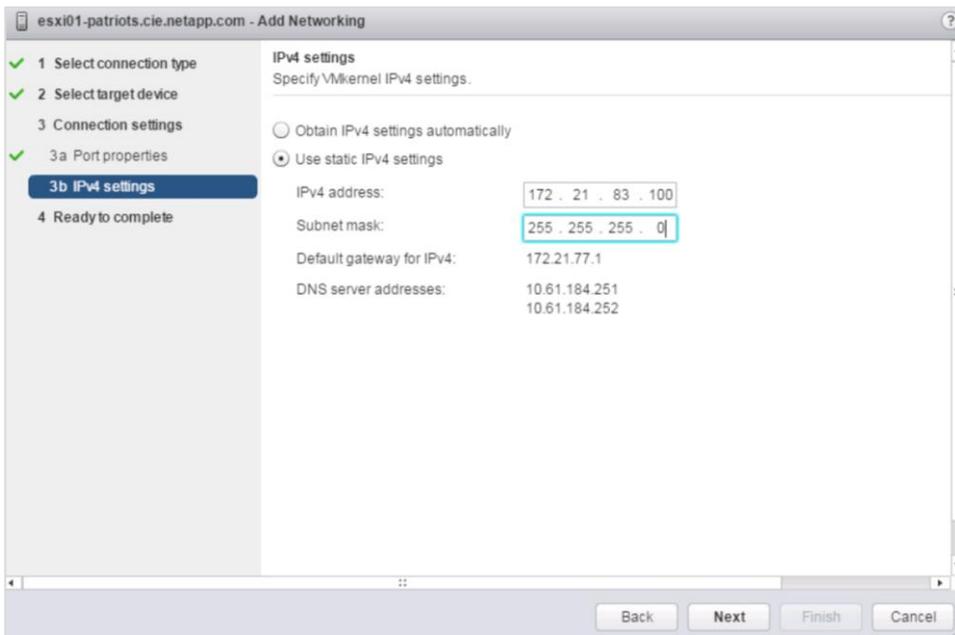
1. From the Home menu of vSphere Web Client, select Hosts and Clusters under the Inventories section.
2. Select an ESXi host from the list and click the Manage tab.
3. Click the Networking tab and select VMkernel Adapters.
4. Click the first icon to add host networking.
5. Select VMkernel Network Adapter and click Next.
6. Click the Select an Existing Network button and click Browse.
7. Select the HTCC-HA-VLAN from the list and click OK. Click Next.



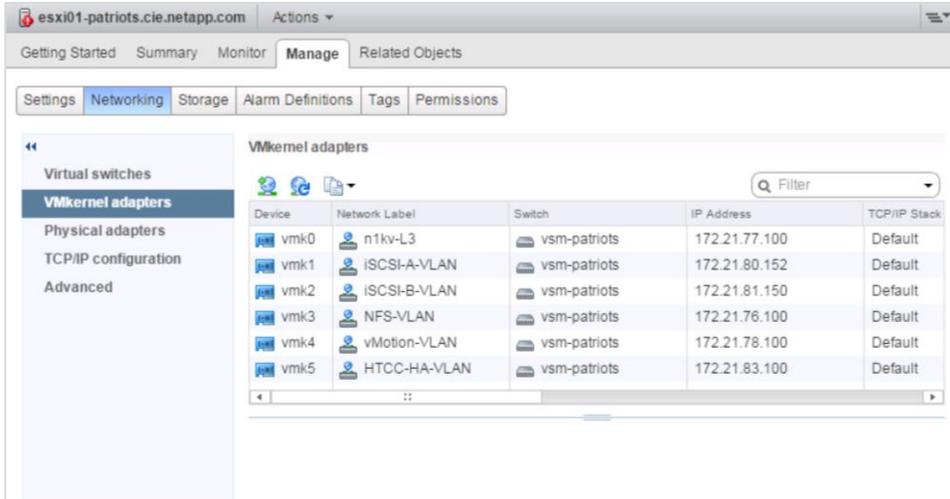
8. Use the default network label and click Next.



9. Select Use Static IPv4 Settings Automatically. Enter the IP address and the subnet mask for the HTCC-HA-VLAN interface for the selected host.



10. Review the settings and click Finish.
11. A new VMkernel adapter is created for the HTCC-HA-VLAN network.



12. Repeat steps 2 to 11 for the remaining ESXi hosts that need to be placed in the management cluster.

**Note:** The HTCC-HA-VLAN is required in the management cluster only to host HyTrust CloudControl.

Complete the previous steps to create the VMkernel ports for all the ESXi servers across all clusters/tenants.

## Set Up iSCSI Multipathing

To set up iSCSI multipathing, complete the following steps:

1. Follow the steps to set up iSCSI multipathing as described in the [CVD](#).
2. When providing the iSCSI LIF IP addresses for Dynamic Discovery, make sure to provide the LIF IPs belonging to the appropriate tenant.

## Mount Required Datastores

1. Using the procedure described in the [CVD](#), mount the following listed datastores/volumes to the ESXi hosts.

ESXi Hosts	Datastores/ Volumes
esxi01-patriots.cie.netapp.com esxi02-patriots.cie.netapp.com	Infra_datastore_1 Infra_swap
esxi03-patriots.cie.netapp.com esxi04-patriots.cie.netapp.com	Infra_datastore_1_tenant_A Infra_swap_tenant_A
esxi05-patriots.cie.netapp.com esxi06-patriots.cie.netapp.com	Infra_datastore_1_tenant_B Infra_swap_tenant_B

**Note:** When mounting the datastores, make sure to provide the NFS LIF IP from the corresponding tenant.

## Move VM Swap File Location

Using the procedures described in the [CVD](#), store the VM Swap File on the infra\_swap datastore mounted on each ESXi.

## VMware vCenter 6.0

Follow the procedures described in the [CVD](#) to deploy the vCenter Server Appliance. The vCenter Server Appliance will need to be deployed on an ESXi host that will eventually be part of the management cluster.

### Create Clusters and Add Hosts

Using the procedures described in the [CVD](#), create additional clusters for the tenants within the same FlexPod\_DC data center.

After the clusters are created, add the ESXi hosts to them. The cluster-to-ESXi server association is as follows:

Cluster	ESXi Hosts
Management	esxi01-patriots.cie.netapp.com esxi02-patriots.cie.netapp.com
Tenant-A	esxi03-patriots.cie.netapp.com esxi04-patriots.cie.netapp.com
Tenant-B	esxi05-patriots.cie.netapp.com esxi06-patriots.cie.netapp.com

## 5.6 Cisco Nexus 1000v VSM Configuration

### Register Cisco Nexus 1000v with vCenter

Cisco Nexus 1000v can be registered to the vCenter server using the Virtual Switch Update Manager or by manually registering the 1000v using the PowerShell script available at <http://VSM-IP-Address/vcplugin/registerVCPlugin.ps1>.

### Update the Primary VSM

This step assumes that the base configuration of the Cisco Nexus 1110-X and the primary and secondary VSM is completed as described in the [CVD](#).

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin and run the following commands:

```
config t
vlan <<var_htcc_ha_vlan_id>>
name HTCC-HA-VLAN
exit
vlan <<var_vm_traffic_vlan_tenant_a>>
name VM-Traffic-VLAN-Tenant-A
exit
vlan <<var_iscsi_a_vlan_tenant_a>>
name iSCSI-A-VLAN-Tenant-A
exit
vlan <<var_iscsi_b_vlan_tenant_a>>
name iSCSI-B-VLAN-Tenant-A
exit
vlan <<var_nfs_vlan_tenant_a>>
name NFS-VLAN-Tenant-A
exit
vlan <<var_vmotion_vlan_tenant_a>>
name vMotion-VLAN-Tenant-A
exit
vlan <<var_vm_traffic_vlan_tenant_b>>
name VM-Traffic-VLAN-Tenant-B
exit
vlan <<var_iscsi_a_vlan_tenant_b>>
```

```

name iSCSI-A-VLAN-Tenant-B
exit
vlan <<var_iscsi_b_vlan_tenant_b>>
name iSCSI-B-VLAN-Tenant-B
exit
vlan <<var_nfs_vlan_tenant_b>>
name NFS-VLAN-Tenant-B
exit
vlan <<var_vmotion_vlan_tenant_b>>
name vMotion-VLAN-Tenant-B
exit

port-profile type ethernet system-uplink
switchport trunk allowed vlan add <<var_htcc_ha_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled

port-profile type ethernet system-uplink-tenant-a
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_tenant_a>>,
<<var_vmotion_vlan_tenant_a>>, <<var_vm-traffic_vlan_tenant_a>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
system mtu 9000
state enabled

port-profile type ethernet system-uplink-tenant-b
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_tenant_b>>,
<<var_vmotion_vlan_tenant_b>>, <<var_vm-traffic_vlan_tenant_b>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
system mtu 9000
state enabled

port-profile type ethernet iscsi-a-uplink-tenant-a
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_a_vlan_tenant_a>>
switchport trunk allowed vlan <<var_iscsi_a_vlan_tenant_a>>
no shutdown
system vlan <<var_iscsi_a_vlan_tenant_a>>
system mtu 9000
state enabled

port-profile type ethernet iscsi-b-uplink-tenant-a
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_tenant_a>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_tenant_a>>
no shutdown
system vlan <<var_iscsi_b_vlan_tenant_a>>
system mtu 9000
state enabled

port-profile type ethernet iscsi-a-uplink-tenant-b
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_a_vlan_tenant_b>>
switchport trunk allowed vlan <<var_iscsi_a_vlan_tenant_b>>
no shutdown
system vlan <<var_iscsi_a_vlan_tenant_b>>
system mtu 9000
state enabled

```

```

port-profile type ethernet iscsi-b-uplink-tenant-b
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_tenant_b>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_tenant_b>>
no shutdown
system vlan <<var_iscsi_b_vlan_tenant_b>>
system mtu 9000
state enabled

port-profile type vethernet HTCC-HA-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_htcc_ha_vlan_id>>
no shutdown
system vlan <<var_htcc_ha_vlan_id>>
state enabled
exit

port-profile type vethernet NFS-VLAN-tenant-a
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_tenant_a>>
no shutdown
system vlan <<var_nfs_vlan_tenant_a>>
state enabled

port-profile type vethernet vMotion-VLAN-tenant-a
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_tenant_a>>
no shutdown
system vlan <<var_vmotion_vlan_tenant_a>>
state enabled

port-profile type vethernet VM-Traffic-VLAN-tenant-a
vmware port-group
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_tenant_a>>
no shutdown
system vlan <<var_vm-traffic_vlan_tenant_a>>
state enabled

port-profile type vethernet iSCSI-A-VLAN-tenant-a
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_tenant_a>>
no shutdown
system vlan <<var_iscsi_a_vlan_tenant_a>>
state enabled

port-profile type vethernet iSCSI-B-VLAN-tenant-a
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_tenant_a>>
no shutdown
system vlan <<var_iscsi_b_vlan_tenant_a>>
state enabled

port-profile type vethernet NFS-VLAN-tenant-b
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_tenant_b>>
no shutdown
system vlan <<var_nfs_vlan_tenant_b>>
state enabled

port-profile type vethernet vMotion-VLAN-tenant-b
vmware port-group
switchport mode access

```

```

switchport access vlan <<var_vmotion_vlan_tenant_b>>
no shutdown
system vlan <<var_vmotion_vlan_tenant_b>>
state enabled

port-profile type vethernet VM-Traffic-VLAN-tenant-b
vmware port-group
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_tenant_b>>
no shutdown
system vlan <<var_vm-traffic_vlan_tenant_b>>
state enabled

port-profile type vethernet iSCSI-A-VLAN-tenant-b
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_tenant_b>>
no shutdown
system vlan <<var_iscsi_a_vlan_tenant_b>>
state enabled

port-profile type vethernet iSCSI-B-VLAN-tenant-b
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_tenant_b>>
no shutdown
system vlan <<var_iscsi_b_vlan_tenant_b>>
state enabled

exit

copy run start

```

## Add VMware ESXi Hosts to Cisco Nexus 1000v

The ESXi hosts can be added to the 1000v using the Virtual Switch Update Manager as described in the [CVD. These hosts can also be added](#) by manually installing the Cisco Nexus 1000v VEM on each ESXi by downloading it from [http://VSM-IP-Address/cross\\_cisco-vem-v199-5.2.1.3.1.5b.0-6.0.1.vib](http://VSM-IP-Address/cross_cisco-vem-v199-5.2.1.3.1.5b.0-6.0.1.vib).

## 5.7 HyTrust CloudControl Installation and Configuration

HyTrust CloudControl (HTCC) offers system managers and administrators an end-to-end virtualization security platform to manage access, standardize and control configuration, and protect a virtual infrastructure within a customer's environment.

### Network Architecture and Topology

HyTrust CloudControl can be deployed in two network configurations: Mapped Mode or Router Mode.

In Mapped Mode, HTCC works as a proxy server and does not require making any architectural changes to the virtual infrastructure network.

In Router Mode, HTCC joins two IPv4 networks, passing information from one network to the other. This mode also requires changes to the existing routing infrastructure.

In this deployment, HTCC will be installed in a HA configuration in Mapped Mode. To facilitate the HA configuration, HTCC requires a dedicated private network between the two HTCC instances. The HTCC HA VLAN is used for this purpose.

## Install HyTrust CloudControl in High-Availability Configuration

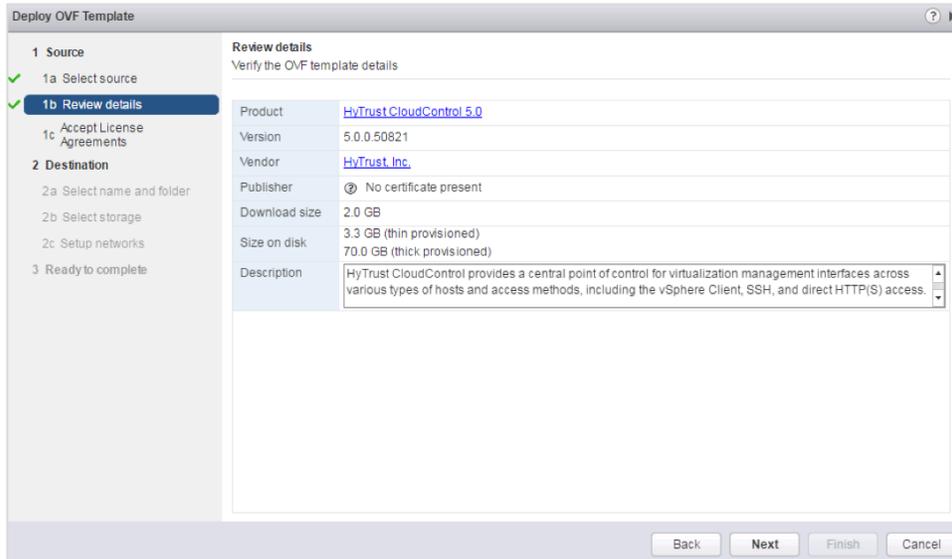
The HyTrust CloudControl will be installed in the management cluster within the VMware environment.

### Obtaining the Software

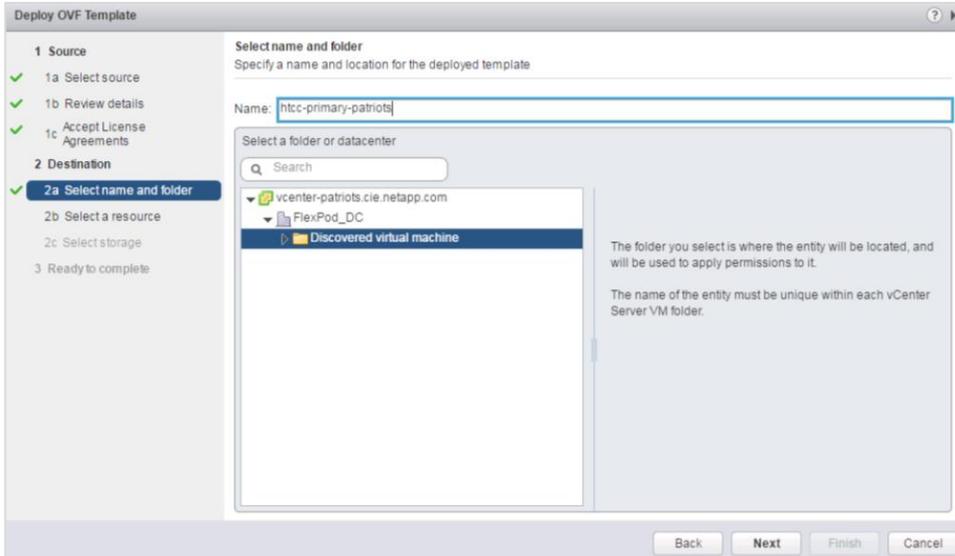
1. Log in to the HyTrust website or follow the directions you received from HyTrust Support to obtain the download URL of the HTCC OVF file.

### Install Primary HTCC Appliance

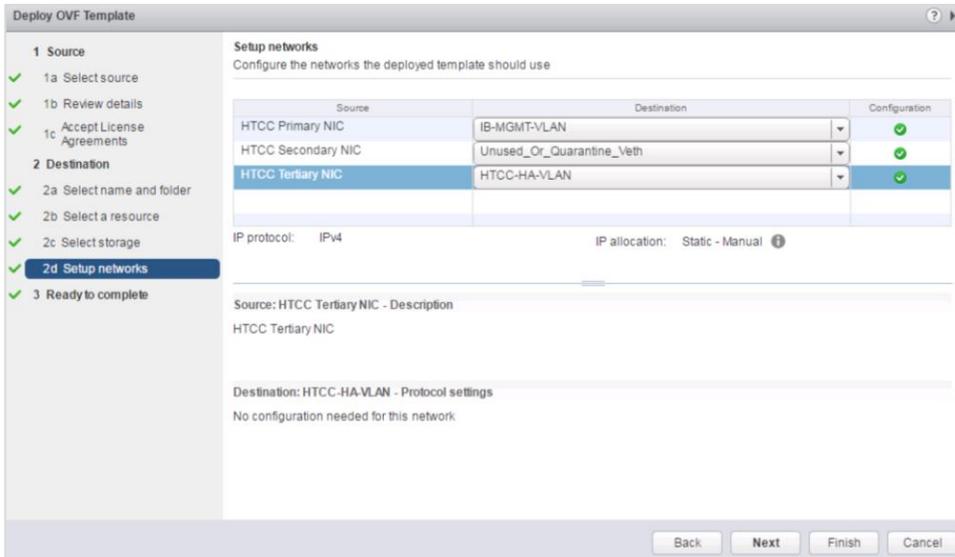
1. From the Home menu in vSphere Web Client, select VMs and Templates under the Inventories section.
2. Right-click the FlexPod\_DC data center or VM Folder and select Deploy OVF Template.
3. Click Allow.
4. Browse to the OVF file of the HyTrust CloudControl appliance and click Open.
5. Click Next.
6. Review the OVF template details and click Next.



7. Accept the license agreement and click Next.
8. Enter a name for the HTCC primary appliance and select a folder or data center where it should reside.



9. Click Next.
10. Select the FlexPod\_Management cluster and click Next.
11. Select infra\_datastore\_1 as the storage and click Next.
12. Assign the Appliance NICs as follows:
  - a. HTCC Primary NIC (eth0) → IB-MGMT-VLAN
  - b. HTCC Secondary NIC (eth1) → Unused\_Or\_Quarantine\_Veth
  - c. HTCC Tertiary NIC (eth2) → HTCC-HA-VLAN

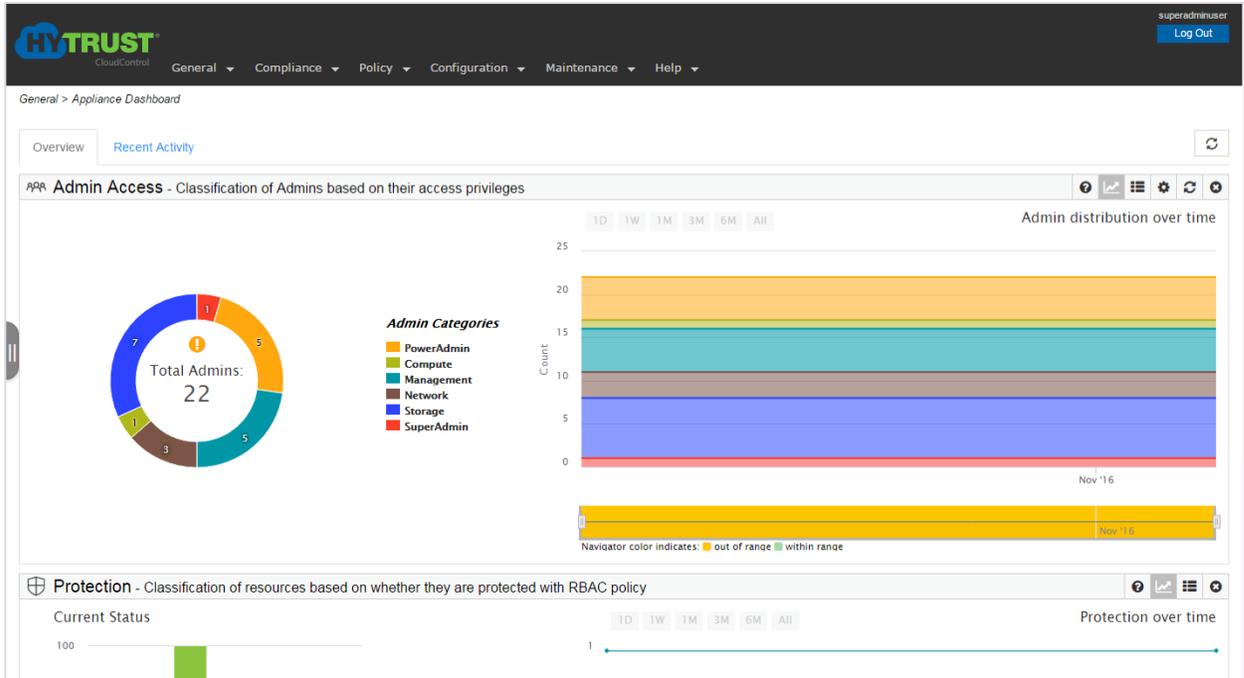


13. Click Next.
14. Review the settings and click Finish.
15. Wait for the OVF template to be deployed.
16. Select the HTCC primary virtual machine from the Inventory pane and from the Summary tab click Launch Remote Console.

17. Click Launch Application if prompted.
18. Click the green button (second from left) to power on the VM.

## Configure Primary HTCC Management Network Interface

1. In the console window of the HTCC VM, log in as `ascadminuser` with the password `Pa$$w0rd123!`.
2. Enter the current password, `Pa$$w0rd123!`.
3. Assign a new password for the `ascadminuser` and reenter the password to confirm.
4. Start the setup procedure by running the `setup` command.
5. Enter `n` when asked to configure a virtual management IP address.
6. Enter the IPv4 address for the management network connection (`eth0`) interface, `<<var_htcc_pri_mgmt_ip>>`.
7. Enter the netmask `<<var_htcc_pri_netmask>>`.
8. Enter the gateway `<<var_htcc_pri_gw>>`.
9. Enter the DNS server IP addresses.
10. Review the network settings and confirm.
11. Log out after the network settings are updated.
12. Open a web browser and navigate to `https://<<var_htcc_pri_mgmt_ip>>/asc`.
  - Note:** FQDN is not supported until the installation wizard completes.
  - Note:** Use the IPv4 address.
  - Note:** NetApp recommends using Mozilla Firefox as the browser.
13. Allow the security exceptions if prompted.
14. Log in using the default user name `superadminuser` and password `Pa$$w0rd123!`.
15. Accept the license agreement and click Next.
16. Upload the license file with the `.xml` extension and click Next.
17. In the Network Configuration page, assign a host name for the HTCC primary appliance and review the network settings.
18. Update the (comma separated) list of DNS servers if necessary.
  - Note:** Provide only IP addresses for DNS servers.
19. Select the Enable NTP Servers checkbox and enter the NTP server IP addresses (comma separated).
  - Note:** Provide only IP addresses for NTP servers.
20. Click Next.
21. Click Finish to complete the installation wizard.
  - Note:** The Finish button is not enabled until the Installation wizard completes.
22. Upon successful installation, the HTCC Management Console Appliance Dashboard appears.



23. From the vSphere Web Client, connect to the console of the HTCC primary virtual machine.
24. Log in as `ascadminuser`.
25. Start the HA setup procedure by running the `hasetup` command.
26. At the `Please specify network settings for the Connection 1 (eth0) interface` prompt, confirm the settings assigned to the primary HTCC. Enter `n` to skip reconfiguring the network settings.
27. At the `Deploy as primary (production) or secondary (standby) (pri/sec)` prompt, type `pri`.
28. Enter `y` to configure a private network for high availability.
29. At the `Please specify network settings for High Availability services on Connection 3 (eth2) interface` prompt, enter the primary HTCC connection 3 (eth2) details.
 

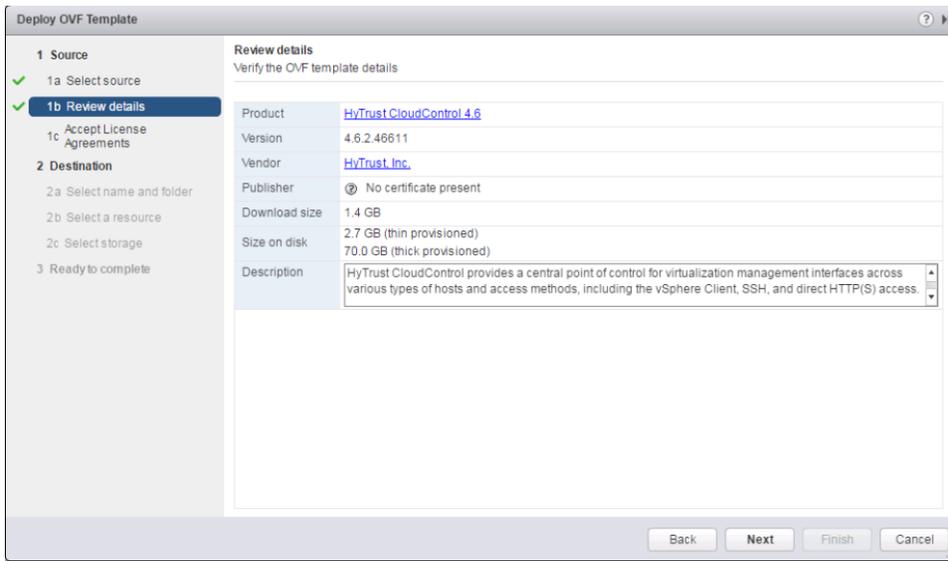
**Note:** The `<<var_htcc_pri_ha_ip>>` and `<<var_htcc_pri_ha_netmask>>` network parameters defined for the HTCC-HA-VLAN need to be used.
30. Enter `y` when prompted to save the settings.
31. Enter `n` when asked to configure a virtual management IP address.

The HA setup for primary HTCC is now complete. Next, you must install and configure a second HTCC instance and join the two HTCCs to create an HTCC-HA cluster.

## Install Secondary HTCC Appliance

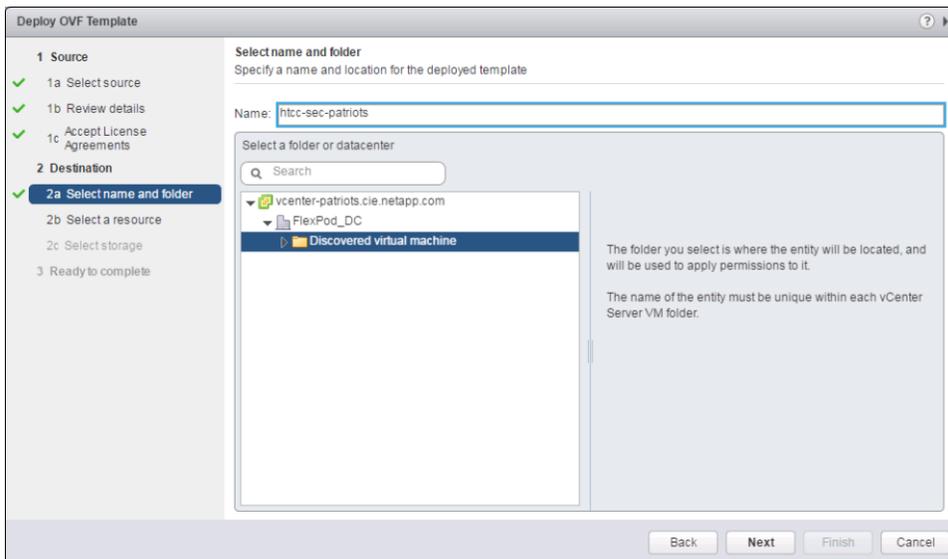
1. From the Home menu in the vSphere Web Client, select VMs and Templates under the Inventories section.
2. Right-click the FlexPod\_DC data center or the VM folder and select Deploy OVF Template.
3. Click Allow.
4. Browse to the OVF file of the HyTrust CloudControl appliance and click Open.
5. Click Next.

6. Review the OVF template details and click Next.



7. Accept the license agreement and click Next.

8. Enter a name for the HTCC secondary appliance and select a folder or data center where it should reside.



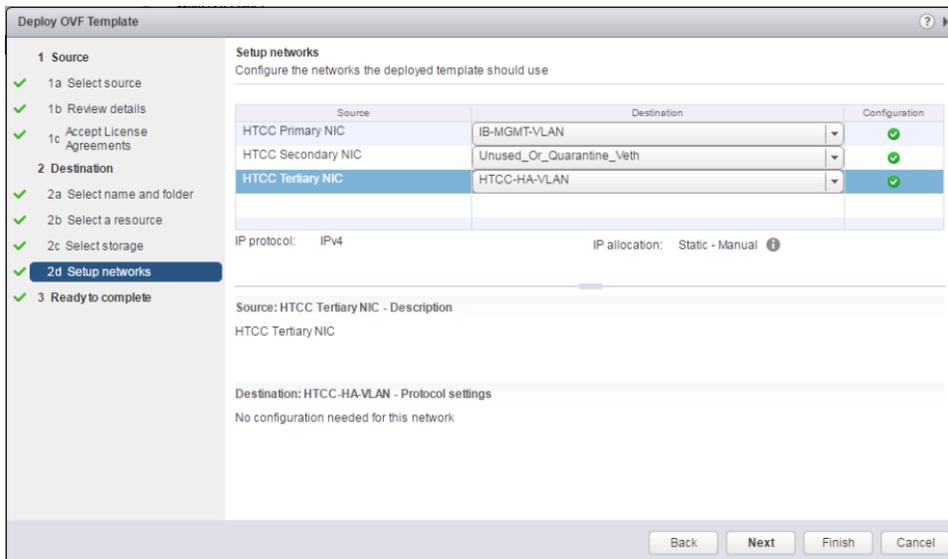
9. Click Next.

10. Select the FlexPod\_Management cluster and click Next.

11. Select infra\_datastore\_1 as the storage and click Next.

12. Assign the appliance NICs as follows:

- a. HTCC Primary NIC (eth0) → IB-MGMT-VLAN
- b. HTCC Secondary NIC (eth1) → Unused\_Or\_Quarantine\_Veth
- c. HTCC Tertiary NIC (eth2) → HTCC-HA-VLAN

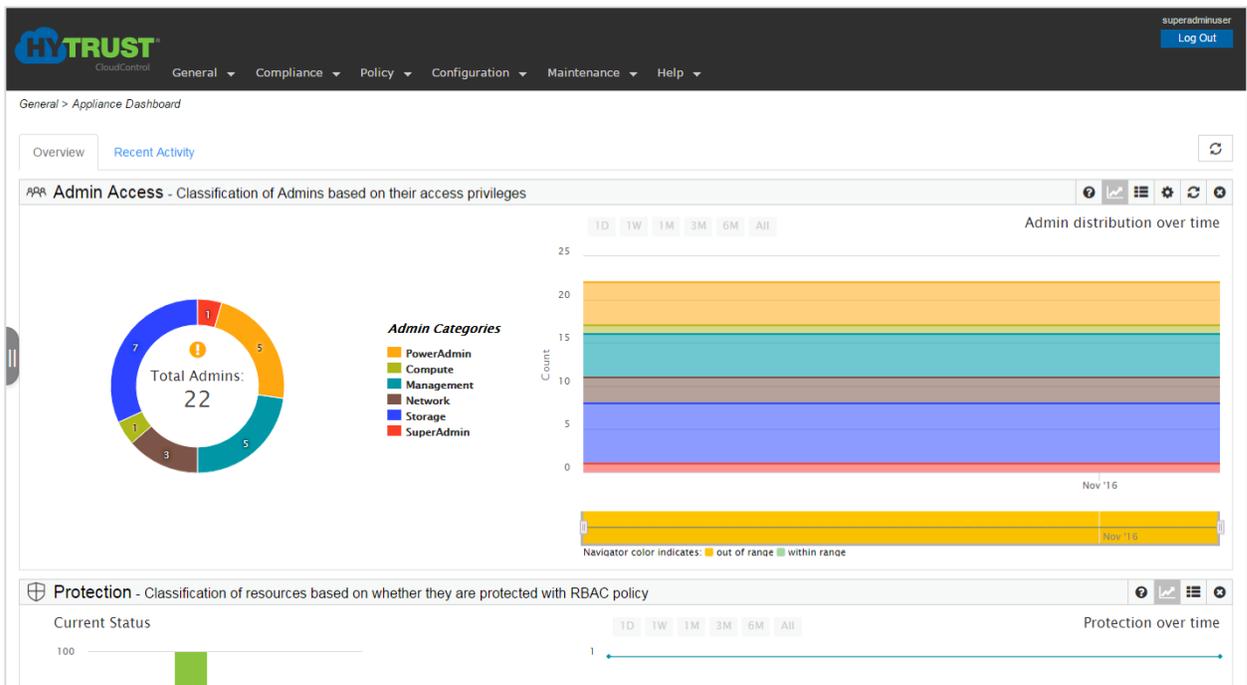


13. Click Next.
14. Review the settings and click Finish.
15. Wait for the OVF template to be deployed.
16. Select the HTCC secondary virtual machine from the Inventory pane and from the Summary tab click Launch Remote Console.
17. Click Launch Application if prompted.
18. Click the green button (second from left) to power on the VM.

## Configure Secondary HTCC Management Network Interface

1. In the console window of the HTCC VM, log in as `ascadminuser` with the password `Pa$$w0rd123!`.
2. Enter the current password `Pa$$w0rd123!`.
3. Assign a new password for the `ascadminuser` and reenter the password to confirm.
4. Start the setup procedure by running the `setup` command.
5. Enter `n` when asked to configure a virtual management IP address.
6. Enter the IPv4 address for the management network connection (eth0) interface, `<<var_htcc_sec_ip >>`.
7. Enter the netmask `<<var_htcc_sec_netmask>>`.
8. Enter the gateway `<<var_htcc_sec_gw>>`.
9. Enter the DNS server IP addresses.
10. Review the network settings and confirm.
11. Log out after the network settings have been updated.
12. Open a web browser and navigate to `https://<<var_htcc_sec_ip>>/asc`.
  - Note:** FQDN is not supported until the installation wizard completes.
  - Note:** Use IPv4 address.
  - Note:** NetApp recommends using Mozilla Firefox as the browser.
13. Allow the security exceptions if prompted.

14. Log in using the default user name `superadminuser` and the password `Pa$$w0rd123!`.
15. Accept the license agreement and click Next.
16. Upload the license file with the `.xml` extension and click Next.
17. In the Network Configuration page, assign a host name for the HTCC secondary appliance and review the network settings.
18. Update the list of (comma separated) DNS servers if necessary.
  - Note:** Provide only IP addresses for DNS servers.
19. Click the Enable NTP Servers checkbox and enter the NTP server IP addresses (comma separated).
  - Note:** Provide only IP addresses for NTP servers.
20. Click Next.
21. Click Finish to complete the installation wizard.
  - Note:** The Finish button is not enabled until the installation wizard completes.
22. The HTCC Management Console Appliance Dashboard appears on successful installation.



23. From the vSphere Web Client, connect to the console of the HTCC secondary virtual machine.
24. Log in as `ascadminuser`.
25. Start the HA setup procedure by running the `hasetup` command.
26. At the Please specify network settings for the Connection 1 (eth0) interface prompt, confirm the settings assigned to the secondary HTCC. Enter `n` to skip reconfiguring the network settings.
27. At the Deploy as primary (production) or secondary (standby) (pri/sec) prompt, type `sec`.
28. Enter `y` to configure a private network for high availability.

29. At the `Please specify network settings for High Availability services on Connection 3 (eth2) interface` prompt, enter the secondary HTCC Connection 3 (eth2) details.

**Note:** The `<var_htcc_pri_ha_ip>>` network parameter defined for the HTCC-HA-VLAN should be used.

**Note:** This process might take several minutes as the secondary HTCC establishes communication with the primary HTCC.

After this process completes, the secondary HTCC updates and displays the HyTrust high-availability (HA) system status as `Enabled` and the mode as `Secondary`. The HA status is also updated on the primary HTCC and shows the mode as `Primary` after the CLI command window is refreshed.

## Configure HTCC to Directory Service Mode

Configure HyTrust CloudControl to perform authentication against a Microsoft Active Directory service for a streamlined access policy to the HTCC appliance.

### Create a Service Account

1. Log in to the Windows machine running the Active Directory server using credentials that have sufficient privileges to create accounts.
2. In Active Directory, add a new user to serve as the HTCC service account.
  - Full name: `HtaServiceAccount`
  - User login name: `htaserviceaccount`

### Create Security Groups

The default HTCC rules are created by mapping existing user groups in Active Directory to default roles in HTCC when HTCC is converted to Directory Service mode.

Refer to the “HyTrust CloudControl Administration Guide” to create the necessary Security Groups.

### Integrate HTCC with Active Directory

**Note:** Converting HTCC to Directory Service mode for authentication and authorization cannot be reversed.

1. Browse to the HTCC Management Console.
2. Click the Configuration tab and from the drop-down click Authentication.
3. Select the Directory Service button and click Apply.
4. Enter the domain name.
5. Enter the service account name created earlier and enter the password.
6. Select Automated Discovery under the Configuration Method and click Next.
7. Select the View Active Directory Advanced Settings checkbox and click Next.
8. Review the preferred global catalog, domain details, user search context, and group search context. Make any necessary changes and click Next.
9. Map the HTCC roles to the Active Directory security groups created and click Next.
10. Review the settings and click Finish.
11. After the conversion is complete, log in to the HTCC Management console with the Active Directory credentials.

**Note:** Before logging in to the HTCC Management Console, the security groups in Active Directory must be populated with the required users.

## Add HTCC Protected Hosts

**Note:** NetApp recommends using Fully Qualified Domain Names in place of IP addresses wherever possible for the host IP and published IP.

### vCenter

1. From the HTCC Management Console, click Compliance. From the drop-down, click Hosts.
2. Click Add. The Add Host wizard appears.
3. Select vCenter vSphere Web Client Server and VMware NSX. Then click Next.  
**Note:** No NSX will be added during this step.
4. Enter the vCenter host name/IP followed by the user ID and password. Then click Next.
5. Enter a description for the vCenter host (optional).
6. Verify that the Protected checkbox is selected and click Next.
7. Enter the published IP (PIP) and the published IP mask and click Next.
8. Enter the vSphere Web Client sever host name/IP followed by the user ID and password.
9. Enter the published IP and netmask for the vSphere Web Client server. Click Next.
10. Click Next on the Authentication Mode Configuration section without making any changes.
11. Click Finish.
12. After the vCenter discovery is completed, select an ESXi host by clicking it.
13. In the General tab, enter the user ID and password.
14. Click the Published IP tab and enter the published IP and netmask for the ESXi. Then click OK.
15. Repeat steps 13 to 15 for all the remaining ESXi hosts.

### Configure SAML Data Provider

1. Log in to the console of the HyTrust CloudControl primary appliance as the `ascadminuser`.
2. Enter the following command in the console:

```
asc certs -b
```

3. Enter `y` to import certificates for all the hosts.
4. After all certificates are imported, log in to the HyTrust CloudControl web interface with `SuperAdmin` privileges.
5. Click the Compliance tab and select Hosts.
6. Select the checkbox beside the vSphere Web Client Server and click Download SAML Metadata.
7. Log in to the vSphere Web Client. From the Home menu, click Administration.
8. Under Single Sign-On, select Configuration.
9. Select SAML Service Providers in the right pane and click Import.
10. Click Import from File and navigate to the downloaded SAML metadata.
11. Click Import.

## 5.8 HyTrust DataControl Installation and Configuration

HyTrust DataControl (HTDC) provides encryption and key management for virtual machines. Its major components are HyTrust KeyControl and HyTrust DataControl Policy Agent.

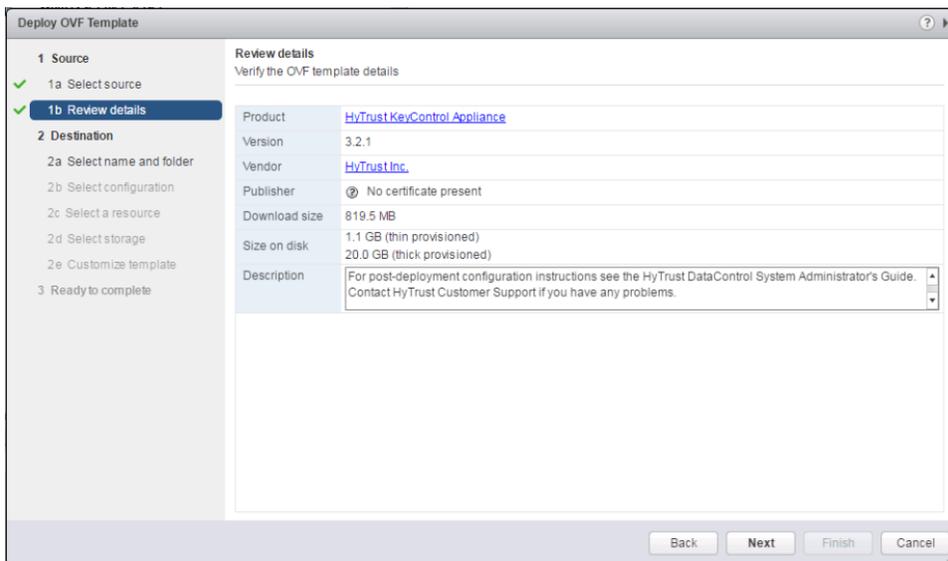
The HyTrust DataControl installation procedure includes installing the HyTrust KeyControl nodes in a cluster configuration and the policy agents in the VMs.

A clustered instance of HyTrust DataControl is installed in tenants A and B to protect the VMs residing within the tenant/cluster.

To install HTDC in tenants A and B, complete the following procedures.

## Install First HyTrust KeyControl Node

1. Log in to vSphere Web Client.
2. From the Home menu, click Hosts and Clusters.
3. Right-click the tenant (A/B) cluster and click Deploy OVF Template.
4. Click Allow to enable the Client Integration Plugin, if prompted.
5. Browse to the HyTrust DataControl.ova file and click Open.
6. Click Next.
7. Review the details and click Next.

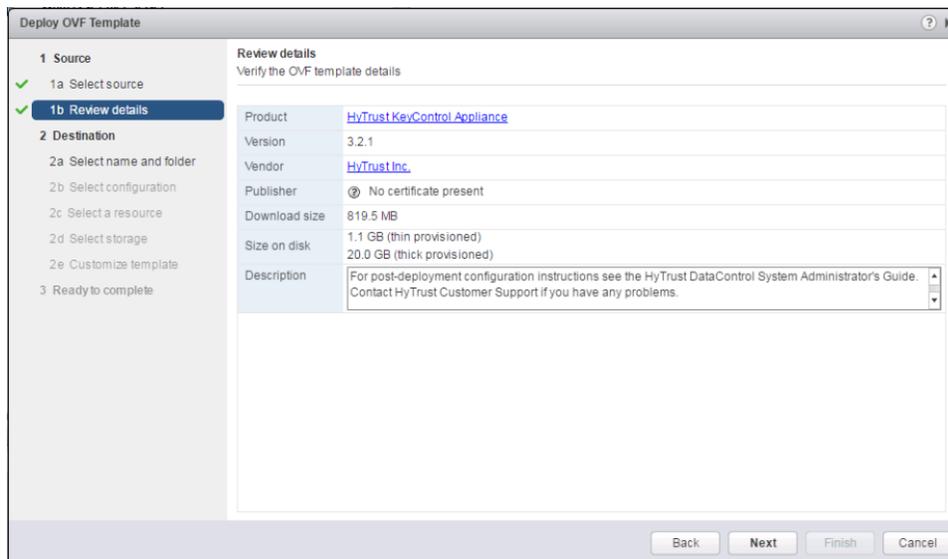


8. Enter a name for the first HyTrust KeyControl virtual machine and select a folder or data center in which it will reside. Click Next.
9. In the Configuration section, select the default recommended option and click Next.
10. Select the `infra_datastore` provisioned for that cluster and click Next.
11. In the Network selection, select the VM-Traffic-VLAN created for the respective tenant/cluster for the VM network and click Next.
12. In the Customization template, enter the following:
  - a. The first KeyControl system IP address defined in the VM-Traffic-VLAN for the tenant
  - b. The first KeyControl system host name
  - c. Domain name
  - d. Netmask
  - e. Gateway
  - f. Primary DNS server
13. Click Next.
14. Review the settings and click Finish.
15. After the HyTrust KeyControl is deployed, launch the remote console for the virtual machine.

16. Click Launch Application if prompted.
17. Click the green button (second from left) to power on the VM.
18. Enter a new password for the HyTrust KeyControl and confirm the password.
19. Select No when prompted to add this KeyControl node to a cluster.
20. Select OK.
21. Reboot the KeyControl system.
22. Open a web browser and navigate to the IP address of the first HyTrust KeyControl system.
23. Log in with user name `secroot` and password `secroot`.
24. Read and accept the license agreement.
25. Enter and confirm a new password for the WebGUI. Click Update Password.
26. Configure the e-mail and mail server settings according to your organization's standards. Click Update Mail settings.

## Install Second HyTrust KeyControl Node

1. Log in to vSphere Web Client.
2. From the Home menu, click Hosts and Clusters.
3. Right-click the tenant (A/B) cluster and click Deploy OVF Template.
4. Click Allow to enable the Client Integration Plugin, if prompted.
5. Browse to the HyTrust DataControl.ova file and click Open.
6. Click Next.
7. Review the details and click Next.



8. Enter a name for the second HyTrust KeyControl virtual machine and select a folder or data center in which it will reside. Click Next.
9. In the Configuration section, select the default recommended option and click Next.
10. Select the infra\_datastore provisioned for that cluster and click Next.
11. In the Network selection, choose the VM-Traffic-VLAN created for the respective tenant/cluster for the VM network and click Next.
12. In the Customization template, enter the following:

- a. The second KeyControl system IP address defined in the VM-Traffic-VLAN for the tenant
  - b. The second KeyControl system host name
  - c. Domain name
  - d. Netmask
  - e. Gateway
  - f. Primary DNS server
13. Click Next.
  14. Review the settings and click Finish.
  15. After the HyTrust KeyControl is deployed, launch the remote console for the VM.
  16. Click Launch Application if prompted.
  17. Click the green button (second from left) to power on the virtual machine.
  18. Enter a new password for the HyTrust KeyControl and confirm the password.
  19. Select Yes when prompted to add this KeyControl node to a cluster and click OK.
  20. Enter a description and click OK.
  21. Enter the IP address of the First HyTrust KeyControl system.
  22. Enter a passphrase for the system.  
**Note:** Remember this passphrase; you will need to provide it again.
  23. Log in to the WebGUI of the first KeyControl system.
  24. Click Cluster in the top pane and click the Servers tab.
  25. Select the second KeyControl system, click Actions, and then click Authenticate.
  26. Enter the passphrase that was entered previously and click Authenticate.
  27. After authentication completes, the KeyControl node is listed as Authenticated but Unreachable until cluster synchronization completes and the cluster is ready for use.

## Create VM Sets

All protected VMs in the HyTrust DataControl environment are managed through VM sets.

A VM set is a logical grouping of related VMs. Also, authentication between the protected VMs and the KeyControl cluster requires the use of a per-VM certificate that is used during registration of the VM with the KeyControl cluster. This process ties the VM to a specific administration group and VM set.

1. Log in to KeyControl WebGUI.
2. Click the Cloud icon.
3. Click Actions and select Create New Cloud VM Set.
4. Enter a name and provide a description. Leave Cloud Admin Group selected by default.
5. Click Create and then click Close.

## Install the HyTrust DataControl Policy Agent

Complete the following procedure to install the HyTrust DataControl Policy Agent. The DataControl Policy Agent is installed in the VMs that need to be protected by HTDC.

Repeat this procedure for installing the agent on the VMs belonging to tenant A and tenant B. Make sure that the VMs are registered to the HTDC instance running on the same cluster/tenant.

**Note:** This deployment focuses only on protecting Windows VMs. Therefore, the following procedure describes the installation of HyTrust DataControl Policy Agent on Windows VMs. To install the Policy Agent on Linux VMs, refer to the “HyTrust DataControl Administration Guide.”

1. Select the Windows VM within which you would like to install the DataControl Policy Agent.
2. Log in to the VM. Download and install [.NET Framework version 4](#).
3. Before proceeding with installation, make sure that all drives in the VMs have been assigned a drive letter.
4. Log in to the WebGUI of the KeyControl system. Click Cloud. Under Actions, click Download Policy Agent.
5. Extract the downloaded agent file and navigate to the Windows client.
6. Make sure that the Disk Defragmenter service on each client computer is enabled before installing the Policy Agent software.
7. Right-click the Windows Policy Agent Client and select Run as Administrator.
8. Click Next on the Welcome screen.
9. Accept the license agreement.
10. Choose a destination to install and click Next.
11. Verify that the HT Bootloader checkbox is selected and click Next.
12. Leave Drive Letter Assignment on Automatic.
13. Review the VM's network details and click Install.
14. Click OK when prompted to copy the key file `id_rsa` to a different machine.
15. Leave the Reboot Now button selected and click Finish.
16. After reboot, log in to the VM and navigate the installation location of the Policy Agent.
17. Copy the `id_rsa` file to another machine and keep it safe.
18. Click Start, select HyTrust GUI, and click Register.
19. Enter the following details in the Registration dialog box:
  - a. The first KeyControl IP address/host name
  - b. The second KeyControl IP address/host name
  - c. Password for the secroot WebGUI user
  - d. Name of the Cloud VM set created earlier
  - e. (Optional) Description
20. Click Register.
21. Click OK after the registration is successful.
22. In WebGUI, right-click each drive you want to encrypt/protect and select Add and Encrypt.
23. Click Yes to continue.
24. Repeat steps 1 to 23 for the VMs that you would like to protect.

## 5.9 Set VM Restart Priority

To set up the VM restart priority for the HyTrust CloudControl and KeyControl appliances, complete the following steps:

1. From the vSphere Web Client, select Hosts and Clusters.
2. Navigate to the Management cluster. On the right pane, click Manage and then select Settings.
3. Under the Configuration pane, select VM Overrides and then click Add.
4. Click the + button to add VMs. Select the HyTrust CloudControl primary and secondary VMs from the list and click OK.
5. From the VM restart priority drop-down list, select High and then click OK.
6. Repeat steps 2 to 5 for the HyTrust DataControl VMs running in the tenant clusters.

## 6 FedRAMP Security Controls

Table 6 lists the FedRAMP moderate impact security controls that were addressed by the information system.

Table 6) FedRAMP moderate impact security controls.

Control Family	Control # Addressed	Total Controls
ACCESS CONTROL	AC-1, AC-2, AC-2(2), AC-2(4), AC-2(5), AC-3, AC-3(3), AC-4, AC-5, AC-6, AC-6(2), AC-6(9), AC-6(10), AC-7, AC-8, AC-10, AC-11, AC-12, AC-17(2), AC-17(3)	20
AUDIT AND ACCOUNTABILITY	AU-2, AU-3, AU-5, AU-7, AU-7(1), AU-8, AU-8(1), AU-9, AU-12	9
CONFIGURATION MANAGEMENT	CM-2, CM-2(1), CM-2(2), CM-3, CM-5, CM-5(3), CM-6, CM-6(1), CM-7(2), CM-8, CM-8(3)	11
SECURITY ASSESSMENT AND AUTHORIZATION	CA-7, CA-9	2
CONTINGENCY PLANNING	CP-2, CP-2(2), CP-10	3
IDENTIFICATION AND AUTHENTICATION	IA-2, IA-2(11), IA-3, IA-5, IA-5(1), IA-6, IA-7, IA-8	8
MEDIA PROTECTION	MP-5, MP-5(4)	2
RISK ASSESSMENT	RA-5, RA-5(5)	2
SYSTEM AND SERVICES ACQUISITION	SA-2, SA-3, SA-4, SA-5, SA-8, SA-10, SA-11	7
SYSTEM AND COMMUNICATIONS PROTECTION	SC-2, SC-4, SC-5, SC-6, SC-8(1), SC-13, SC-20, SC-21, SC-22, SC-23, SC-28, SC-39	12
SYSTEM AND INFORMATION INTEGRITY	SI-3, SI-4(5), SI-7, SI-7(1), SI-7(7), SI-10, SI-11, SI-16	8

## 7 Conclusion

The FlexPod Datacenter solution caters to a wide variety of workloads and helps in building scalable and robust data centers. As part of the FedRAMP readiness exercise, FlexPod Datacenter can address a significant number of FedRAMP Moderate Impact Baseline controls across various control families. This exercise showcases the built-in security features of FlexPod Datacenter and the additional security features that can be implemented by integrating HyTrust CloudControl and DataControl in a FlexPod environment.

## References

This report references the following resources:

- [FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1 Design Guide](#)
- [FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1 Deployment Guide](#)

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

### Copyright Information

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

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

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

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

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

**Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-0031-1216