



NetApp Verified Architecture

# FlexPod Express with VMware vSphere 6.0: Small and Medium Configurations NVA Deployment Guide

Authors: Karthick Radhakrishnan and Arvind Ramakrishnan, NetApp  
Reviewers: Jeffrey Fultz and Chris O'Brien, Cisco Systems, Inc.

October 2015 | NVA-0020-DEPLOY | Version 1.0

## TABLE OF CONTENTS

<b>1</b>	<b>Solution Overview</b>	<b>4</b>
1.1	Solution Technology	4
1.2	Use Case Summary	5
<b>2</b>	<b>Technology Requirements</b>	<b>5</b>
2.1	Hardware Requirements	5
2.2	Software Requirements	6
<b>3</b>	<b>FlexPod Express Cabling Information</b>	<b>6</b>
3.1	FlexPod Express Small Configuration	6
3.2	FlexPod Express Medium Configuration	9
<b>4</b>	<b>Deployment Procedures</b>	<b>13</b>
4.1	Cisco Nexus 3048 Deployment Procedure	14
4.2	NetApp FAS Storage Deployment Procedure	21
4.3	Cisco UCS C-Series Rack Server Deployment Procedure	37
4.4	VMware vSphere 6.0 Deployment Procedure	42
4.5	VMware vCenter 6.0 Deployment Procedure	50
4.6	NetApp Virtual Storage Console 6.0 Deployment Procedure	67
<b>5</b>	<b>Bill of Materials</b>	<b>79</b>
5.1	Small Configuration	79
5.2	Medium Configuration	81
<b>6</b>	<b>Conclusion</b>	<b>82</b>
	<b>References</b>	<b>82</b>

## LIST OF TABLES

Table 1)	FlexPod Express small configuration hardware requirements	5
Table 2)	FlexPod Express medium configuration hardware requirements	6
Table 3)	Software requirements	6
Table 4)	Cabling information for the FlexPod Express small configuration	7
Table 5)	Cabling information for the FlexPod Express medium configuration	10
Table 6)	Required VLANs	14
Table 7	lists VMware virtual machines (VMs) created	14
Table 8)	VMware virtual machines created	14
Table 9)	Controller FAS25XX series prerequisites	22
Table 10)	Clustered Data ONTAP software installation prerequisites	22
Table 11)	Cluster create in clustered Data ONTAP prerequisites	27

Table 12) Prerequisites for cluster join in clustered Data ONTAP .....	29
Table 13) Small configuration components. ....	79
Table 14) Medium configuration components.....	81

**LIST OF FIGURES**

Figure 1) Physical topology of FlexPod Express small configuration. ....	4
Figure 2) Physical topology of FlexPod Express medium configuration. ....	5
Figure 3) FlexPod Express small configuration cabling diagram. ....	7
Figure 4) FlexPod Express medium configuration cabling diagram.....	10

# 1 Solution Overview

FlexPod Express is a suitable platform for running a variety of virtualization hypervisors as well as bare-metal operating systems and enterprise workloads. FlexPod Express delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements. The small and medium FlexPod Express configurations are low-cost, standardized infrastructure solutions developed to meet the needs of small and midsize businesses. Each configuration provides a standardized base platform capable of running a number of business-critical applications while providing scalability options to enable the infrastructure to grow with the demands of the business.

FlexPod Express:

- Combines all application and data needs into one platform
- Suitable for small-midsize organizations, remote and departmental deployments
- Provides easy infrastructure scaling
- Reduces cost and complexity

## 1.1 Solution Technology

The small and medium FlexPod Express configurations use Cisco UCS C-Series Rack-Mount Servers, Cisco Nexus Switches (1GbE), and NetApp FAS storage systems (NetApp clustered Data ONTAP: switchless). This document describes the implementation of VMware vSphere 6.0 on the small and medium FlexPod Express offerings. The configurations are based on best practices for each component in the solution architecture to enable a reliable, enterprise-class infrastructure.

Figures 1 and 2 depict the topology of the FlexPod Express small and medium offerings.

Figure 1) Physical topology of FlexPod Express small configuration.

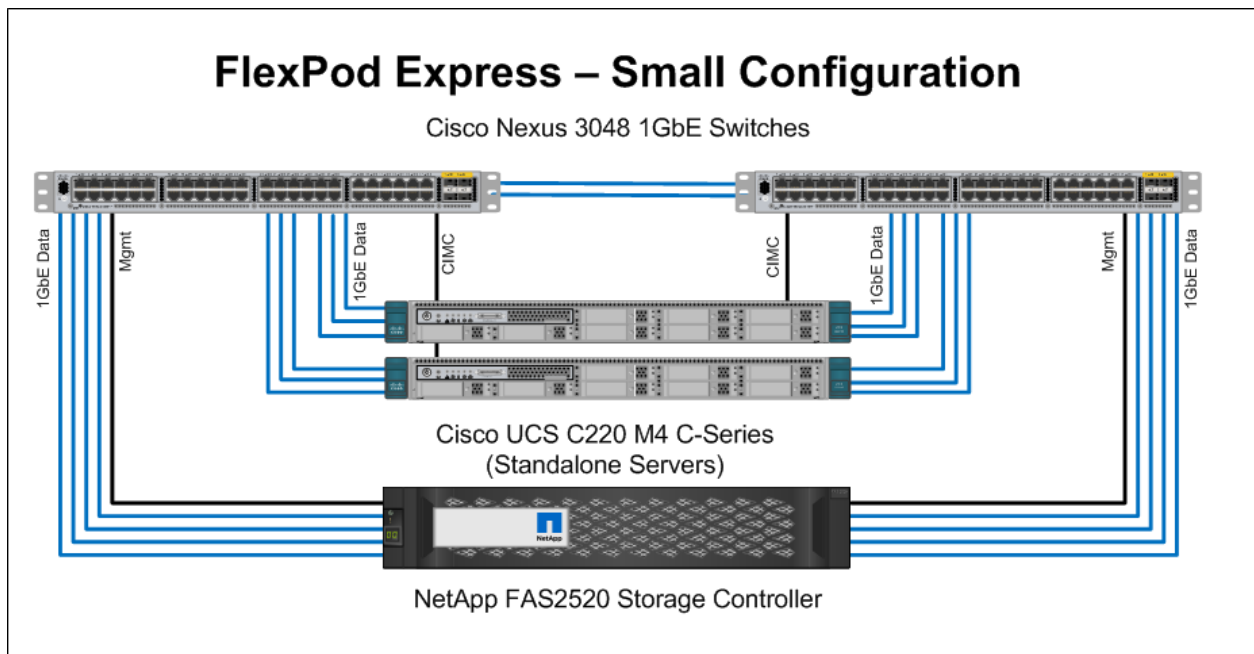



Figure 2) Physical topology of FlexPod Express medium configuration.



## 1.2 Use Case Summary

This document describes the deployment procedures and best practices to set up a FlexPod Express small and/or medium with VMware vSphere 6.0 as the workload. The server operating system/hypervisor is VMware vSphere ESXi, and an instance of VMware vCenter Server is installed to manage the ESXi instances. The whole infrastructure is supported by NetApp FAS storage systems that serve data over storage area network (SAN) and network-attached storage (NAS) protocols.

## 2 Technology Requirements

The hardware and software components required to implement the FlexPod Express small and medium configurations are detailed in this section.

### 2.1 Hardware Requirements

Table 1 lists the hardware components required to implement the FlexPod Express small configuration solution.

Table 1) FlexPod Express small configuration hardware requirements.

Layer	Hardware	Quantity
Compute	Cisco UCS C220 M4 rack-mount servers (standalone)	2
Network	Cisco Nexus 3048 switches	2
Storage	NetApp FAS2520 (high-availability pair)	1

Layer	Hardware	Quantity
Disks	900GB, 10,000 rpm SAS w/ Advanced Drive Partitioning	12

Table 2 lists the hardware components required to implement the FlexPod Express medium configuration solution.

Table 2) FlexPod Express medium configuration hardware requirements.

Layer	Hardware	Quantity
Compute	Cisco UCS C220 M4 rack-mount servers (standalone)	4
Network	Cisco Nexus 3048 switches	2
Storage	NetApp FAS2520 (high-availability pair)	1
Disks	900GB, 10,000 rpm SAS w/ Advanced Drive Partitioning	12

## 2.2 Software Requirements

Table 3 lists the software components required to implement the FlexPod Express small and medium configurations.

Table 3) Software requirements.

Layer	Component	Version or Release	Details
Compute	Cisco UCS C220 M4 rack-mount servers	2.0(3)	Cisco Integrated Management Controller (IMC) software
Network	Cisco Nexus 3048 Gigabit Ethernet switches	6.0(2)U6(1)	Cisco NX-OS software
Storage	NetApp FAS2520 high-availability storage	8.3	NetApp Data ONTAP software
Hypervisor	VMware vSphere	6.0	VMware Virtualization Hypervisor Suite
Other software	NetApp Virtual Storage Console (VSC)	6.0	NetApp Plug-In for VMware vCenter

## 3 FlexPod Express Cabling Information

### 3.1 FlexPod Express Small Configuration

Figure 3 provides a cabling diagram for the FlexPod Express small configuration. Table 4 provides cabling information.

Figure 3) FlexPod Express small configuration cabling diagram.

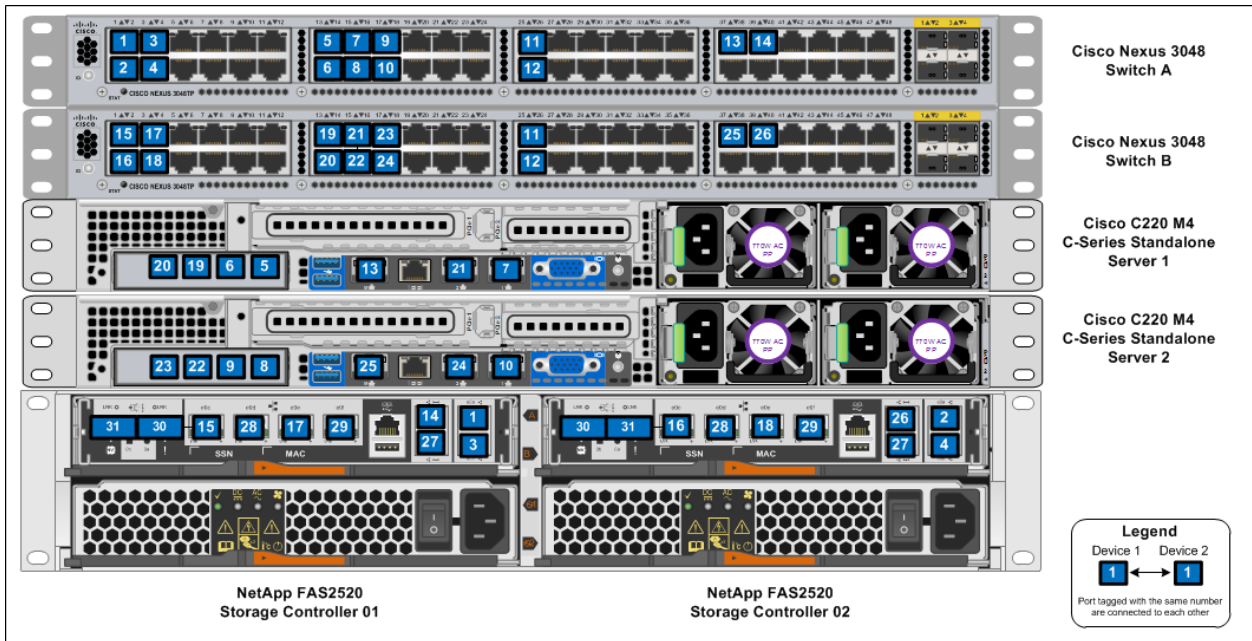


Table 4) Cabling information for the FlexPod Express small configuration.

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3048 Switch A	Eth1/1	NetApp FAS2520 Storage Controller 01	e0a	<b>1</b>
	Eth1/2	NetApp FAS2520 Storage Controller 02	e0a	<b>2</b>
	Eth1/3	NetApp FAS2520 Storage Controller 01	e0b	<b>3</b>
	Eth1/4	NetApp FAS2520 Storage Controller 02	e0b	<b>4</b>
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 1	<b>5</b>
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 2	<b>6</b>
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 1	LOM1	<b>7</b>
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 1	<b>8</b>
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 2	<b>9</b>
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 2	LOM1	<b>10</b>
Eth1/25	Cisco Nexus 3048 Switch B	Eth1/25	<b>11</b>	

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
	Eth1/26	Cisco Nexus 3048 Switch B	Eth1/26	12
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 1	Cisco IMC	13
	Eth1/39	NetApp FAS2520 Storage Controller 01	e0M	14

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3048 Switch B	Eth1/1	NetApp FAS2520 Storage Controller 01	e0c	15
	Eth1/2	NetApp FAS2520 Storage Controller 02	e0c	16
	Eth1/3	NetApp FAS2520 Storage Controller 01	e0e	17
	Eth1/4	NetApp FAS2520 Storage Controller 02	e0e	18
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 3	19
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 4	20
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 1	LOM2	21
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 3	22
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 4	23
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 2	LOM2	24
	Eth1/25	Cisco Nexus 3048 Switch A	Eth1/25	11
	Eth1/26	Cisco Nexus 3048 Switch A	Eth1/26	12
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 2	Cisco IMC	25
	Eth1/39	NetApp FAS2520 Storage Controller 02	e0M	26



Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2520 Storage Controller 01	e0d	NetApp FAS2520 Storage Controller 02	e0d	28
	e0f	NetApp FAS2520 Storage Controller 02	e0f	29
	ACP	NetApp FAS2520 Storage Controller 02	ACP	27
	SAS 0a	NetApp FAS2520 Storage Controller 02	SAS 0b	30
	SAS 0b	NetApp FAS2520 Storage Controller 02	SAS 0a	31

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2520 Storage Controller 02	e0d	NetApp FAS2520 Storage Controller 01	e0d	28
	e0f	NetApp FAS2520 Storage Controller 01	e0f	29
	ACP	NetApp FAS2520 Storage Controller 01	ACP	27
	SAS 0a	NetApp FAS2520 Storage Controller 01	SAS 0b	31
	SAS 0b	NetApp FAS2520 Storage Controller 01	SAS 0a	30

### 3.2 FlexPod Express Medium Configuration

Figure 4 provides a cabling diagram for the FlexPod Express small configuration. Table 5 provides cabling information.

Figure 4) FlexPod Express medium configuration cabling diagram.




Table 5) Cabling information for the FlexPod Express medium configuration.

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3048 Switch A	Eth1/1	NetApp FAS2520 Storage Controller 01	e0a	1
	Eth1/2	NetApp FAS2520 Storage Controller 02	e0a	2
	Eth1/3	NetApp FAS2520 Storage Controller 01	e0b	3
	Eth1/4	NetApp FAS2520 Storage Controller 02	e0b	4
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 1	5
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 2	6
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 1	LOM1	7
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 1	8

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 2	9
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 2	LOM1	10
	Eth1/19	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 1	11
	Eth1/20	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 2	12
	Eth1/21	Cisco UCS C220 C-Series Standalone Server 3	LOM1	13
	Eth1/22	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port1	14
	Eth1/23	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port 2	15
	Eth1/24	Cisco UCS C220 C-Series Standalone Server 4	LOM1	16
	Eth1/25	Cisco Nexus 3048 Switch B	Eth1/25	17
	Eth1/26	Cisco Nexus 3048 Switch B	Eth1/26	18
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 1	Cisco IMC	19
	Eth1/38	Cisco UCS C220 C-Series Standalone Server 3	Cisco IMC	20
	Eth1/39	NetApp FAS2520 Storage Controller 01	e0M	21

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3048 Switch B	Eth1/1	NetApp FAS2520 Storage Controller 01	e0c	22
	Eth1/2	NetApp FAS2520 Storage Controller 02	e0c	23
	Eth1/3	NetApp FAS2520 Storage Controller 01	e0e	24
	Eth1/4	NetApp FAS2520 Storage Controller 02	e0e	25
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 3	26

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 4	27
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 1	LOM2	28
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 3	29
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 4	30
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 2	LOM2	31
	Eth1/19	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 3	32
	Eth1/20	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 4	33
	Eth1/21	Cisco UCS C220 C-Series Standalone Server 3	LOM2	34
	Eth1/22	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port3	35
	Eth1/23	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port 4	36
	Eth1/24	Cisco UCS C220 C-Series Standalone Server 4	LOM2	37
	Eth1/25	Cisco Nexus 3048 Switch A	Eth1/25	17
	Eth1/26	Cisco Nexus 3048 Switch A	Eth1/26	18
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 2	Cisco IMC	38
	Eth1/38	Cisco UCS C220 C-Series Standalone Server 4	Cisco IMC	39
	Eth1/39	NetApp FAS2520 Storage Controller 02	e0M	40

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2520 Storage	e0d	NetApp FAS2520 Storage Controller 02	e0d	42
	e0f	NetApp FAS2520 Storage Controller 02	e0f	43

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Controller 01	ACP	NetApp FAS2520 Storage Controller 02	ACP	41
	SAS 0a	NetApp FAS2520 Storage Controller 02	SAS 0b	44
	SAS 0b	NetApp FAS2520 Storage Controller 02	SAS 0a	45

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2520 Storage Controller 02	e0d	NetApp FAS2520 Storage Controller 01	e0d	42
	e0f	NetApp FAS2520 Storage Controller 01	e0f	43
	ACP	NetApp FAS2520 Storage Controller 01	ACP	41
	SAS 0a	NetApp FAS2520 Storage Controller 01	SAS 0b	45
	SAS 0b	NetApp FAS2520 Storage Controller 01	SAS 0a	44

## 4 Deployment Procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either Component 01 or Component 02. For example, Controller 01 and Controller 02 identify the two NetApp storage controllers that are provisioned in this document, and Switch A and Switch B identify the pair of Cisco Nexus switches that are configured.

Additionally, this document details steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: Server-1, Server-2, and so on.

To indicate that you should include information pertinent to your environment in a given step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<ib_mgmt_vlan_id>>
```

This document is intended to enable you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes. Table 6 describes the VLANs necessary for deployment as outlined in this guide. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

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

Table 6) Required VLANs.

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Management VLAN	VLAN for management interfaces	186
Native VLAN	VLAN to which untagged frames are assigned	2
Network File System (NFS) VLAN	VLAN for NFS traffic	3011
VMware vMotion VLAN	VLAN designated for the movement of virtual machines from one physical host to another	3012
Virtual machine traffic VLAN	VLAN for virtual machine application traffic	3013

Table 7 lists VMware virtual machines (VMs) created.

Table 8) VMware virtual machines created.

Virtual Machine Description	Host Name
VMware vCenter Server	
NetApp Virtual Storage Console	

## 4.1 Cisco Nexus 3048 Deployment Procedure

The following section details the Cisco Nexus 3048 switch configuration for use in a FlexPod Express environment.

### Cisco Nexus 3048 Switch Initial Setup

Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup, and defines the control-plane policing policy.

The first major decision involves the configuration of the management network for the switches. For FlexPod Express, there are two main options for configuring the mgmt0 interfaces. The first involves configuring and cabling the mgmt0 interfaces into an existing out-of-band network. In this instance, when a management network already exists, all you need are valid IP addresses and the netmask configuration for this network and a connection from the mgmt0 interfaces to this network.

The other option, for installations without a dedicated management network, involves cabling the mgmt0 interfaces of each Cisco Nexus 3048 switch together in a back-to-back configuration. Any valid IP address and netmask can be configured on each mgmt0 interface as long as they are in the same network. Because they are configured back to back with no switch or other device in between, no default gateway configuration is needed, and they should be able to communicate with each other. This link cannot be used for external management access such as SSH access, but it will be used for the virtual PortChannel (vPC) peer keepalive traffic. To enable SSH management access to the switch, you need to configure the in-band interface VLAN IP address on an SVI, as discussed later in this document.

1. Power on the switch and follow the onscreen prompts as illustrated here for the initial setup of both switches, substituting the appropriate values for the switch-specific information.

### Cisco Nexus Switch A and Switch B

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and
basic configuration, no - continue with Power On Auto Provisioning] (yes/skip/no)[no]: yes

---- System Admin Account Setup ----
```

```

Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin":<<admin_password>>
Confirm the password for "admin":<<admin_password>>

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

Please register Cisco Nexus 3000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus devices must be registered to receive entitled
support services.

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

Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]:Enter
Configure read-write SNMP community string (yes/no) [n]:Enter
Enter the switch name : <<switch_A/B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:Enter
Mgmt0 IPv4 address : <<switch_A/B_mgmt0_ip_addr>>
Mgmt0 IPv4 netmask : <<switch_A/B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]:n
Enable the telnet service? (yes/no) [n]:Enter
Enable the ssh service? (yes/no) [y]:Enter
Type of ssh key you would like to generate (dsa/rsa) : rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]:Enter
Configure default interface layer (L3/L2) [L2]:Enter
Configure default switchport interface state (shut/noshut) [noshut]:Enter
Configure CoPP System Policy Profile ( default / 12 / 13 ) [default]:Enter

The following configuration will be applied:
switchname <<switch_A/B_hostname>>
interface mgmt0
ip address <<switch_A/B_mgmt0_ip_addr>> <<switch_A/B_mgmt0_netmask>>
no shutdown
no telnet server enable
ssh key rsa 1024 force
ssh server enable
system default switchport
no system default switchport shutdown
policy-map type control-plane copp-system-policy ( default )

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

```

## Upgrading NX-OS (Optional)

You should perform any required software upgrades on the switch at this point in the configuration process. Download and install the latest available Cisco NX-OS software for the Cisco Nexus 3048 switch from the [Cisco software download site](#). There are multiple ways to transfer both the kickstart and system images for Cisco NX-OS to the switch. The most straightforward procedure uses the onboard USB port on the switch. Download the Cisco NX-OS kickstart and system files to a USB drive and plug the USB drive into the external USB port on the Cisco Nexus 3048 switch.

**Note:** Cisco NX-OS software release 6.0(2)U6(1) is used in this solution

1. Copy the files to the local bootflash memory and update the switch by following the procedure shown below.

### Cisco Nexus Switch A and Switch B

```
copy usb1:<<kickstart_image_file>> bootflash:
```

```
copy usb1:<<system_image_file>> bootflash:
install all kickstart bootflash:<<kickstart_image_file>> system bootflash:<<system_image_file>>
```

**Note:** The switch will install the updated Cisco NX-OS files and reboot.

## Enabling Advanced Features

Certain advanced features need to be enabled in Cisco NX-OS to provide additional configuration options.

**Note:** The `interface-vlan` feature is required only if you are using the back-to-back `mgmt0` option described throughout this document. This feature allows an IP address to be assigned to the interface VLAN (SVI), which enables in-band management communication to the switch, such as through SSH.

1. Enter configuration mode using the (`config t`) command and type the following commands to enable the appropriate features on each switch.

### Cisco Nexus Switch A and Switch B

```
feature interface-vlan
feature lacp
feature vpc
```

## Performing Global PortChannel Configuration

The default PortChannel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the PortChannel. Better distribution across the members of the PortChannels can be achieved by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For that reason, adding the source and destination TCP ports to the hash algorithm is highly recommended.

From configuration mode (`config t`) type the following commands to configure the global PortChannel load-balancing configuration on each switch.

### Cisco Nexus Switch A and Switch B

```
port-channel load-balance ethernet source-dest-port
```

## Performing Global Spanning-Tree Configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure and a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network and edge, depending on the platform.

The recommended setting for bridge assurance is to consider all ports to be network ports by default.

This setting will force the network administrator to review the configuration of each port and will help reveal the most common configuration errors, such as unidentified edge ports or a neighbor that does not have bridge assurance enabled. Also, it is safer to have spanning tree block too many ports than not enough, allowing the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In those cases, you might need to change the port type to make the ports active.

Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature will shut down the port if BPDUs from another switch are seen on this interface.



From configuration mode (`config t`) type the following commands to configure the default spanning-tree options, including the default port type and BPDU guard on each switch.

### Cisco Nexus Switch A and Switch B

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

## Configuring Jumbo Frames

Jumbo frames should be configured throughout the network to allow any applications and operating systems to transmit these larger frames without fragmentation. Note that both endpoints and all interfaces between the endpoints (Layer 2 and Layer 3) must support and be configured for jumbo frames to achieve the benefits and to prevent performance problems by fragmenting frames.

From configuration mode (`config t`) type the following commands to enable jumbo frames on each switch.

### Cisco Nexus Switch A and Switch B

```
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9000
system qos
  service-policy type network-qos jumbo
exit
```

## Defining VLANs

Before configuring individual ports with different VLANs, those Layer 2 VLANs must be defined on the switch. It's also good practice to name the VLANs to help with any troubleshooting in the future.

From configuration mode (`config t`) type the following commands to define and give descriptions to the Layer 2 VLANs.

### Cisco Nexus Switch A and Switch B

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<ib_mgmt_vlan_id>>
  name IB-MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

## Configuring Access and Management Port Descriptions

As with the assignment of names to the Layer 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

For the small configuration, the descriptions for the management ports and data ports associated with Server-3 and Server-4 are not required because the FlexPod Express small configuration contains only two servers.

From configuration mode (`config t`) in each switch, type the following commands to set up the port descriptions.

## FlexPod Express Small Configuration

Enter the following port descriptions for the FlexPod Express small configuration.

### Cisco Nexus Switch A

### Cisco Nexus Switch B

<pre>int eth1/1   description Controller-01:e0a int eth1/2   description Controller-02:e0a int eth1/3   description Controller-01:e0b int eth1/4   description Controller-02:e0b int eth1/13   description Server-1:MLOM Port1 int eth1/14   description Server-1:MLOM Port2 int eth1/15   description Server-1:LOM Port1 int eth1/16   description Server-2:MLOM Port1 int eth1/17   description Server-2:MLOM Port2 int eth1/18   description Server-2:LOM Port1 int eth1/25   description vPC peer-link NX3048-B:1/25 int eth1/26   description vPC peer-link NX3048-B:1/26 int eth1/37   description Server-1:mgmt int eth1/39   description Controller-01:mgmt</pre>	<pre>int eth1/1   description Controller-01:e0c int eth1/2   description Controller-02:e0c int eth1/3   description Controller-01:e0e int eth1/4   description Controller-02:e0e int eth1/13   description Server-1: Server-1:MLOM Port3 int eth1/14   description Server-1:MLOM Port4 int eth1/15   description Server-1:LOM Port2 int eth1/16   description Server-2:MLOM Port3 int eth1/17   description Server-2:MLOM Port4 int eth1/18   description Server-2:LOM Port2 int eth1/25   description vPC peer-link NX3048-A:1/25 int eth1/26   description vPC peer-link NX3048-A:1/26 int eth1/37   description Server-2:mgmt int eth1/39   description Controller-02:mgmt</pre>
---	---

## FlexPod Express Medium Configuration

Enter the following port descriptions for the FlexPod Express medium configuration.

### Cisco Nexus Switch A

### Cisco Nexus Switch B

<pre>int eth1/19   description Server-3: MLOM Port1 int eth1/20   description Server-3: MLOM Port12 int eth1/21   description Server-3: LOM Port1 int eth1/22   description Server-4: MLOM Port1 int eth1/23   description Server-4: MLOM Port2 int eth1/24   description Server-4: LOM Port1 int eth1/38   description Server-3:mgmt</pre>	<pre>int eth1/19   description Server-3: MLOM Port3 int eth1/20   description Server-3: MLOM Port4 int eth1/21   description Server-3: LOM Port2 int eth1/22   description Server-4: MLOM Port3 int eth1/23   description Server-4: MLOM Port4 int eth1/24   description Server-4: LOM Port2 int eth1/38   description Server-4:mgmt</pre>
---	--

## Configuring Server and Storage Management Interfaces

The management interfaces for both the server and storage typically use only a single VLAN. Therefore, you should configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (`config t`) type the following commands to configure the port settings for the management interfaces of both the servers and storage.

## Cisco Nexus Switch A and Switch B

```
int eth1/37-39
  switchport access vlan <<ib_mgmt_vlan_id>>
  spanning-tree port type edge
exit
```

## Performing Virtual PortChannel Global Configuration

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you are using the back-to-back mgmt0 configuration, be sure to use the addresses defined on the interfaces and verify that they can communicate by using the `ping <<switch_A/B_mgmt0_ip_addr>>vrf management` command.

From configuration mode (`config t`) type the following commands to configure the vPC global configuration for switch A.

### Cisco Nexus Switch A

```
vpc domain 1
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source <<switch_A_mgmt0_ip_addr>> vrf
management
  peer-gateway
  auto-recovery
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<ib_mgmt_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start
```

From configuration mode (`config t`), type the following commands to configure the vPC global configuration for switch B.

### Cisco Nexus Switch B

```
vpc domain 1
  role priority 20
  peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source <<switch_B_mgmt0_ip_addr>> vrf
management
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<ib_mgmt_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start
```

## Configuring Storage PortChannels

The NetApp storage controllers allow an active-active connection to the network using Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach allows you to have active-active connections from the storage to completely separate physical switches. Each controller will have two links to each switch, but all four are part of the same vPC and interface group (IFGRP).

From configuration mode (`config t`) type the following commands on each switch to configure the individual interfaces and the resulting PortChannel configuration for the ports connected to the NetApp FAS controller.

### Cisco Nexus Switch A and Switch B and Controller-01 Configuration

```
int eth1/1, eth1/3
  channel-group 11 mode active
int Po11
  description vPC to Controller-01
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<ib_mgmt_vlan_id>>
  spanning-tree port type edge trunk
  vpc 11
  no shut
```

### Cisco Nexus Switch A and Switch B and Controller-02 Configuration

```
int eth1/2, eth1/4
  channel-group 12 mode active
int Po12
  description vPC to Controller-02
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<ib_mgmt_vlan_id>>
  spanning-tree port type edge trunk
  vpc 12
  no shut
exit
copy run start
```

## Configuring Server Connections

The Cisco UCS servers have multiple Ethernet interfaces that can be configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

For the small configuration, you need to configure only Server-1 and Server-2 because only two servers are used in the small FlexPod Express configuration.

From configuration mode (`config t`) type the following commands to configure the port settings for the interfaces connected to each server.

### FlexPod Express Small Configuration

#### Cisco Nexus Switch A and Switch B, Server-1 and Server-2 Configuration

```
int eth1/13-18
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,
<<ib_mgmt_vlan_id>>
```

```
spanning-tree port type edge trunk
no shut
exit
copy run start
```

## FlexPod Express Medium Configuration

### Cisco Nexus Switch A and Switch B, Server-3 and Server-4 Configuration

```
int eth1/19-24
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
  <<native_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>, <<ib_mgmt_vlan_id>>
  spanning-tree port type edge trunk
  no shut
exit
copy run start
```

### Performing In-Band Management SVI Configuration

In-band management using SSH in the FlexPod Express environment is handled by an SVI. To configure the in-band management on each switch, you must configure an IP address on the interface VLAN and set up a default gateway.

From configuration mode (`config t`), type the following commands to configure the Layer 3 SVI for management purposes.

#### Cisco Nexus Switch A

```
int Vlan <<ib_mgmt_vlan_id>>
ip address <<inband_mgmt_ip_address_A>>/<<inband_mgmt_netmask>>
no shut
ip route 0.0.0.0/0 <<inband_mgmt_gateway>>
```

#### Cisco Nexus Switch B

```
int Vlan <<ib_mgmt_vlan_id>>
ip address <<inband_mgmt_ip_address_B>>/<<inband_mgmt_netmask>>
no shut
ip route 0.0.0.0/0 <<inband_mgmt_gateway>>
```

## 4.2 NetApp FAS Storage Deployment Procedure

This section describes the NetApp FAS storage deployment procedure.

### Controller FAS25xx Series

#### NetApp Hardware Universe

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

1. Make sure that the hardware and software components are supported with the version of Data ONTAP that you plan to install by checking the [NetApp Hardware Universe](#) at the [NetApp Support](#) site.
2. Access the [Hardware Universe](#) Application to view the System Configuration guides. Click the “Controllers” tab to view the compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.

3. Alternatively, to compare components by storage appliance, click “Compare Storage Systems.”

Table 9) Controller FAS25XX series prerequisites.

Controller FAS25XX Series Prerequisites
<p>To plan the physical location of the storage systems, refer to the <a href="#">NetApp Hardware Universe</a>. Refer the following sections:</p> <ul style="list-style-type: none"> <li>• Electrical requirements</li> <li>• Supported power cords</li> <li>• Onboard ports and cables</li> </ul> <p>Refer <a href="#">site requirements guide replacement tutorial</a> for finding NetApp FAS platform information using Hardware Universe.</p>

## Storage Controllers

Follow the physical installation procedures for the controllers in the [FAS25xx documentation](#) available at the NetApp Support site.

## NetApp Clustered Data ONTAP 8.3

### Complete the Configuration Worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [Clustered Data ONTAP 8.3 Software Setup Guide](#) at the [NetApp Support](#) site.

**Note:** This system will be set up in a two-node switchless cluster configuration.

Table 10) Clustered Data ONTAP software installation prerequisites.

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<<var_node01_mgmt_ip>>
Cluster node 01 netmask	<<var_node01_mgmt_mask>>
Cluster node 01 gateway	<<var_node01_mgmt_gateway>>
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3 URL	<<var_url_boot_software>>

### Node 01

To configure node 01, complete the following steps:

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

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

2. Set boot monitor defaults.

```
Set-defaults
```

3. Allow the system to boot up.

```
autoboot
```

4. Press Ctrl-C when prompted.

**Note:** If Data ONTAP 8.3 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3 is the version being booted, select option 8 and `yes` to reboot the node and go to step 14.

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

```
7
```

6. Answer `yes` to perform an upgrade.

```
y
```

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

```
e0M
```

8. Select `yes` to reboot now.

```
y
```

9. After reboot, enter the IP address, netmask, and default gateway for e0M in their respective places.

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

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

**Note:** This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

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

```
y
```

13. Enter `yes` to reboot the node.

```
y
```

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

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

```
Ctrl - C
```

15. Select option 5 to enter into maintenance mode.

```
5
```

16. Remove the disk ownership. Enter `Y` to remove the disk ownership and offline the existing volumes/aggregates.

```
disk remove_ownership
```

```
All disks owned by system ID 536902178 will have their ownership information removed.  
Do you wish to continue? y
```

```
Volumes must be taken offline. Are all impacted volumes offline(y/n)?? y  
Removing the ownership of aggregate disks may lead to partition of aggregates between high-  
availability pair.
```

```
Do you want to continue(y/n)? y
```

17. Halt the node. The node will enter Loader prompt.

```
halt
```

18. Boot Data ONTAP.

```
autoboot
```

19. Press Ctrl-C for Boot Menu:

```
Ctrl - C
```

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

```
4
```

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

```
y
```

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

```
y
```

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

## Node 02

To configure node 02, complete the following steps:

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

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

2. Set boot monitor defaults.

```
set-defaults
```

3. Allow the system to boot up.

```
autoboot
```

4. Press Ctrl-C when prompted.

```
Ctrl-C
```

**Note:** If Data ONTAP 8.3 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3 is the version being booted, select option 8 and yes to reboot the node and go to step 14.

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

```
7
```

6. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

8. Select yes to reboot now.

```
y
```

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



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

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

**Note:** This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

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

```
y
```

13. Select yes to reboot the node.

```
y
```

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

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

```
Ctrl - C
```

15. Select option 5 to enter into maintenance mode.

```
5
```

16. Remove the disk ownership. Enter Y to remove the disk ownership and offline the existing volumes/aggregates.

```
disk remove_ownership
```

```
All disks owned by system ID 536902178 will have their ownership information removed. Do you wish to continue? y
```

```
Volumes must be taken offline. Are all impacted volumes offline(y/n)? y
```

```
Removing the ownership of aggregate disks may lead to partition of aggregates between high-availability pair.
```

```
Do you want to continue(y/n)? y
```

17. Halt the node. The node will enter Loader prompt.

```
halt
```

18. Boot Data ONTAP.

```
autoboot
```

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

```
Ctrl - C
```

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

```
4
```

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

```
y
```

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

```
y
```

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

## Node Setup in Clustered Data ONTAP

From a console port program attached to the storage controller A (Node 01) console port, execute the node setup script. This script will appear when Data ONTAP 8.3 first boots on a node.

### 1. Follow the prompts below:

```
Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.

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

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

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
Enter the node management interface netmask: <<var_node01_mgmt_mask>>
Enter the node management interface default gateway: <<var_node01_mgmt_gateway>>
A node management interface on port e0M with IP address <<var_node01_mgmt_ip>> has been created.

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

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

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

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

### 2. Press Return and log in to the node using the admin user ID and no password to get a node command prompt.

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

::> system node reboot

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

### 3. After reboot, go through the node setup procedure with preassigned values.

```
Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.
```

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

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

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

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

For System Setup, this node's management address is: <<var\_node01\_mgmt\_ip>>.

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

4. Log in to the node with the admin user and no password.
5. Repeat this entire procedure for node 2 of the storage cluster.

## Cluster Create in Clustered Data ONTAP

Table 11) Cluster create in clustered Data ONTAP prerequisites.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Clustered Data ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node01 IP address	<<var_node01_mgmt_ip>>
Cluster node01 netmask	<<var_node01_mgmt_mask>>
Cluster node01 gateway	<<var_node01_mgmt_gateway>>

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01.

Using the console session to node 01 the Cluster Setup wizard is brought up by typing `cluster setup`.

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

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

To create a new cluster, complete the following steps:

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

```
create
```

2. Type `no` for single node cluster option.

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

3. Type `no` for cluster network using network switches.

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

4. The system defaults are displayed. Enter `yes` to use the system defaults. Use the following prompts to configure the cluster ports.

```
Existing cluster interface configuration found:
```

Port	MTU	IP	Netmask
e0d	9000	169.254.128.103	255.255.0.0
e0f	9000	169.254.52.249	255.255.0.0

```
Do you want to use this configuration? {yes, no} [yes]:
```

5. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_nfs_license>>
```

**Note:** The cluster is created. This can take a minute or two.

**Note:** For this validated architecture NetApp recommends installing license keys for NetApp SnapRestore®, NetApp FlexClone®, and NetApp SnapManager® Suite. Additionally, install all required storage protocol licenses. After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0a]: e0M
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

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

**Note:** If you have more than one name server IP address, separate the IP addresses with a comma.

7. Set up the node.

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

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

**Note:** The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet.

## Cluster Join in Clustered Data ONTAP

Table 12) Prerequisites for cluster join in clustered Data ONTAP.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster node02 IP address	<<var_node02_mgmt_ip>>
Cluster node02 netmask	<<var_node02_mgmt_mask>>
Cluster node02 gateway	<<var_node02_mgmt_gateway>>

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01, and the node joining the cluster in this example is node 02.

To join the cluster, complete the following steps from the console session of node 02:

1. If prompted, enter `admin` in the login prompt.

```
admin
```

2. The Cluster Setup wizard is brought up by typing `cluster setup`.

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

**Note:** If a login prompt is displayed instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings, and then enter the `cluster setup` command.

3. Run the following command to join a cluster:

```
join
```

4. Data ONTAP detects the existing cluster and agrees to join the same cluster. Follow the below prompts to join the cluster.

```
Existing cluster interface configuration found:

Port    MTU    IP                Netmask
e0d     9000   169.254.144.37   255.255.0.0
e0f     9000   169.254.134.33   255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:
```

5. The steps to join a cluster are displayed.

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

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

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node02_netmask>>]: Enter
Enter the node management interface default gateway [<<var_node02_gw>>]: Enter
```

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

**Note:** The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet.

## Logging in to the Cluster

Open an SSH connection using the cluster IP or host name and log in as the admin user with the password provided during setup.

## Zeroing All Spare Disks

To zero all spare disks in the cluster, complete the following step:

1. Run the following command:

```
disk zerospares
```

## Changing Disk Ownership

Data ONTAP 8.3 has the Advanced Drive Partitioning (ADP) feature, which increases storage efficiency. This document covers the creation of one data aggregate in active/passive configuration. However, if desired, there can be multiple data aggregates in active/active configuration. Complete the following steps to configure ADP active/passive configuration:

1. Disable the disk auto-assign.

```
storage disk option modify -autoassign off -node <<var_node01>>, <<var_node02>>
```

2. Find the data partitions owned by node02.

```
nbice-fpel::> storage disk show -data-owner <<var_node02>>
Disk          Usable      Disk      Container
Size Shelf  Bay Type      Type      Name      Owner
-----
1.0.0          546.9GB    0    0 SAS      shared    aggr0_nbice_fpel_02_0
                                     nbice-fpel-02
1.0.2          546.9GB    0    2 SAS      shared    aggr0_nbice_fpel_02_0
                                     nbice-fpel-02
1.0.4          546.9GB    0    4 SAS      shared    aggr0_nbice_fpel_02_0
                                     nbice-fpel-02
1.0.6          546.9GB    0    6 SAS      shared    aggr0_nbice_fpel_02_0
                                     nbice-fpel-02
1.0.8          546.9GB    0    8 SAS      shared    aggr0_nbice_fpel_02_0
                                     nbice-fpel-02
1.0.10        546.9GB    0   10 SAS      shared    -          nbice-fpel-02
6 entries were displayed.
```

3. Increase the privilege level.

```
nbice-fpel::> set -privilege advanced
```

```
Warning: These advanced commands are potentially dangerous; use them only when directed to do so
by NetApp personnel.
Do you want to continue? {y|n}: y
```

#### 4. Remove the disk ownership for the disks listed in the previous command.

```
storage disk removeowner -data true -disk <Disk Name>
```

Warning: Disks may be automatically assigned to the node because the disk's auto-assign option is enabled. If the affected volumes are not offline, the disks may be auto-assigned during the remove owner operation, which will cause unexpected results. To verify that the volumes are offline, abort this command and use "volume show".

```
Do you want to continue? {y|n}: y  
6 entries were acted on.
```

**Note:** To remove the ownership of multiple disks, the <Disk Name> values can be comma separated in the previous command.

#### 5. Assign all the unassigned data partition disks to node 01.

```
storage disk assign -data true -disk <Disk Name> -owner <<var_node01>>
```

**Note:** The disk assign command should be executed one at a time for each disk.

#### 6. Decrease the privilege level.

```
set -privilege admin
```

## Enabling Cisco Discovery Protocol in Clustered Data ONTAP

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

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

#### 1. Run the following command:

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

## Setting Auto-Revert on Cluster Management

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

#### 1. Run the following command:

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

## Setting Up Service Processor Network Interface Setup

To assign a static IPv4 address to the service processor on each node, complete the following step:

#### 1. Run the following commands:

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

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

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

## Enabling Storage Failover in Clustered Data ONTAP

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

#### 1. Verify the status of storage failover.

```
storage failover show
```

**Note:** Both the nodes <<var\_node01>> and <<var\_node02>> must be capable of performing a takeover. Go to step 3, if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

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

**Note:** Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Go to step 6 if high availability is configured.
5. Enable HA mode only for the two-node cluster.

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

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

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

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

## Creating Jumbo Frame MTU Broadcast Domain in Clustered Data ONTAP

To create a Data broadcast domain with an MTU of 9000, complete the following step:

1. Create broadcast domain on Data ONTAP.

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
```

## Removing Data Ports from the Default Broadcast Domain

```
broadcast-domain remove-ports -broadcast-domain Default -ports <<var_node01>>:e0a,  
<<var_node01>>:e0b, <<var_node01>>:e0c, <<var_node01>>:e0e, <<var_node02>>:e0a,  
<<var_node02>>:e0b, <<var_node02>>:e0c, <<var_node02>>:e0e
```

## Configuring IFGRP LACP in Clustered Data ONTAP

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

1. From the cluster prompt, complete the following steps:

```
ifgrp create -node <<controller01>> -ifgrp a0a -distr-func port -mode multimode_lacp  
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0a  
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0b  
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0c  
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0e  
  
ifgrp create -node <<controller02>> -ifgrp a0a -distr-func port -mode multimode_lacp  
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0a  
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0b  
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0c  
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0e
```



## Configuring Jumbo Frames in Clustered Data ONTAP

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

```
nbice-fpel::> network port modify -node <<var_node01>> -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

nbice-fpel::> network port modify -node <<var_node02>> -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
```

## Creating VLAN in Clustered Data ONTAP

1. Create NFS VLAN ports and add to the data broadcast domain.

```
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>>
```

2. Create IB-MGMT-VLAN ports and add to the default broadcast domain.

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

broadcast-domain add-ports -broadcast-domain Default -ports <<var_node01>>: a0a-
<<ib_mgmt_vlan_id>>, <<var_node01>>: a0a-<<ib_mgmt_vlan_id>>
```

## Creating Aggregates in Clustered Data ONTAP

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 it will contain.

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <<var_node01>> -diskcount <<var_num_disks>>
```

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

**Note:** Start with five disks initially; you can add disks to an aggregate when additional storage is required. Note that in this configuration with a FAS2520, it may be desirable to create an aggregate with all but one remaining disk (spare) assigned to the controller.

**Note:** The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until `aggr1_node1` is online.

2. Disable NetApp Snapshot<sup>®</sup> copies for the data aggregate recently created.

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

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

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

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>>
```

## Configuring NTP in Clustered Data ONTAP

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

1. To set the time zone for the cluster, run the following command:

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

**Note:** For example, in the eastern United States, the time zone is `America/New_York`.

2. To set the date for the cluster, run the following command:

```
date <<ccyymmddhhmm.ss>>
```

**Note:** The format for the date is `<[Century] [Year] [Month] [Day] [Hour] [Minute] . [Second]>`; for example, `201505181453.17`

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

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

## Configuring SNMP in Clustered Data ONTAP

To configure SNMP, complete the following steps:

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

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

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

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

## Configuring SNMPv1 in Clustered Data ONTAP

To configure SNMPv1, complete the following step:

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

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

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

## Configuring SNMPv3 in Clustered Data ONTAP

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

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

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

3. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol, when prompted.
5. Select `des` as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol, when prompted.

## Configuring AutoSupport HTTPS in Clustered Data ONTAP

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

1. Run the following command:

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

## Creating Storage Virtual Machine (Vserver)

To create an infrastructure Vserver, complete the following steps:

1. Run the `vserver create` command.

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

2. Select the Vserver data protocols to configure, leaving NFS.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc,iscsi
```

3. Enable and run the NFS protocol in the Infra-SVM Vserver.

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

4. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI Plug-in.

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

## Configuring HTTPS Access in Clustered Data ONTAP

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

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

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

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each Vserver shown, the certificate common name should match the DNS FQDN of the Vserver. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a Certificate Authority (CA). To delete the default certificates, run the following commands:

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

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -  
type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for Infra-SVM and the cluster Vserver. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.ciscorobo.com -type server -size  
2048 -country US -state "California" -locality "San Jose" -organization "Cisco" -unit "UCS" -
```

```
email-addr "abc@cisco.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

5. To obtain the values for the parameters that would be required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again use TAB completion.

```
security ssl modify [TAB] ...  
Example: security ssl modify -vserver clus -server-enabled true -client-enabled false -ca  
clus.ciscorobo.com -serial 55243646 -common-name clus.ciscorobo.com
```

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

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

8. It is normal for some of these commands to return an error message stating that the entry does not exist.
9. Revert to the regular admin privilege level and set up to allow the Vserver logs to be available by web.

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

## Configuring NFSv3 in Clustered Data ONTAP

To configure NFS on the Vserver, run all commands.

1. Create a new rule for each ESXi host in the default export policy.

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

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol  
nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid  
false  
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 2 -protocol  
nfs -clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid  
false  
vserver export-policy rule show
```

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

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

## Creating FlexVol Volumes in Clustered Data ONTAP

To create a NetApp FlexVol<sup>®</sup> volume, complete the following step:

1. The following information is required to create a FlexVol volume: the volume's name, size, and the aggregate on which it will exist. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -size 500GB -  
state online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-  
snapshot-space 0  
  
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_node01 -size 100GB -state  
online -policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0  
-snapshot-policy none
```

## Enabling Deduplication in Clustered Data ONTAP

To enable deduplication on appropriate volumes, complete the following step:

1. Run the following commands:

```
volume efficiency on -vserver Infra-SVM -volume infra_datastore_1
```

## Creating NFS LIF in Clustered Data ONTAP

1. Create an NFS logical interface (LIF).

```
network interface create -vserver Infra-SVM -lif nfs_infra_swap -role data -data-protocol nfs -
home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> -address
<<var_node01_nfs_lif_infra_swap_ip>> -netmask <<var_node01_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 -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 show
```

**Note:** It is recommended to create a new LIF for each datastore.

## Adding Infrastructure Vserver Administrator

To add the infrastructure Vserver administrator and Vserver administration logical interface in the out-of-band management network, complete the following step:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data -data-protocol none -home-node
<<var_node01>> -home-port a0a-<<ib_mgmt_vlan_id>> -address <<var_vserver_mgmt_ip>> -netmask
<<var_vserver_mgmt_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-
policy mgmt -auto-revert true
```

**Note:** The Vserver management IP here should be in the same subnet as the storage cluster management IP.

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

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

network route show
```

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

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

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

## 4.3 Cisco UCS C-Series Rack Server Deployment Procedure

The following section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in either the small or medium FlexPod Express configuration.

### Performing Initial Cisco UCS C-Series Standalone Server Setup for Cisco IMC

These steps provide details for the initial setup of the Cisco IMC interface for Cisco UCS C-Series standalone servers.

## All Servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM 1.port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.
2. Power on the server and press F8 when prompted to enter the Cisco IMC configuration.



3. In the Cisco IMC configuration utility, set the following options:
  - Network Interface Card (NIC) Mode:
    - Dedicated
  - IP (Basic):
    - IPV4:
    - DHCP enabled:
    - CIMC IP:<<cimc\_ip>>
    - Prefix/Subnet:<<cimc\_netmask>>
    - Gateway: <<cimc\_gateway>>
  - VLAN (Advanced): Leave cleared to disable VLAN tagging.
    - NIC Redundancy
    - None:

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated:      [ ]                   None:           [ ]
Shared LOM:     [X]                   Active-standby: [X]
Cisco Card:     [ ]                   Active-active:  [ ]
  Riser1:       [ ]                   VLAN (Advanced)
  Riser2:       [ ]                   VLAN enabled:   [ ]
  MLom:         [ ]                   VLAN ID:        1
Shared LOM Ext: [ ]                   Priority:        0
IP (Basic)
IPV4:           [X]                   IPV6:           [ ]
DHCP enabled    [ ]
CIMC IP:        192.168.50.18
Prefix/Subnet: 255.255.255.0
Gateway:        192.168.50.1
Pref DNS Server: 10.61.186.19

*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings

```

4. Press F1 to see additional settings.
  - Common Properties:
    - Host name: <<esxi\_host\_name>>
    - Dynamic DNS: [ ]
    - Factory Defaults: Leave cleared.
  - Default User (Basic):
    - Default password: <<admin\_password>>
    - Reenter password: <<admin\_password>>
    - Port Properties: Use default values.
    - Port Profiles: Leave cleared.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
Common Properties
Hostname: icee1-ucs2-cimc
Dynamic DNS: [ ]
DDNS Domain:
FactoryDefaults
Factory Default: [ ]
Default User(Basic)
Default password:
Reenter password:
Port Properties
Auto Negotiation: [ ]
Speed[1000/100 Mbps]: 100
Duplex mode[half/full]: full
Port Profiles
Reset: [ ]
Name:
-no_pp
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F2>PreviousPage

```

5. Press F10 to save the Cisco IMC interface configuration.
6. After the configuration is saved, press Esc to exit.

**Note:** Upgrade the C-Series rack-mount server software to the latest version. This document covers 2.0(3j).

## Configuring Cisco UCS C-Series Servers ESXi Installation on FlexFlash Cards

Some Cisco UCS C-Series rack-mount servers support an internal secure digital (SD) memory card for storage of server software tools and utilities. The SD card is hosted by the Cisco flexible flash (FlexFlash) storage adapter. Users can also install operating systems that have small storage footprints on these cards.

This section describes the steps for installing the VMware ESXi operating system on these cards.

### FlexFlash Card RAID Configuration

Cisco UCS C-Series rack servers that support Cisco FlexFlash cards come with an internal Cisco FlexFlash controller that is responsible for RAID functions.

When a single card is installed on the server, the RAID configuration will be in a degraded mode.

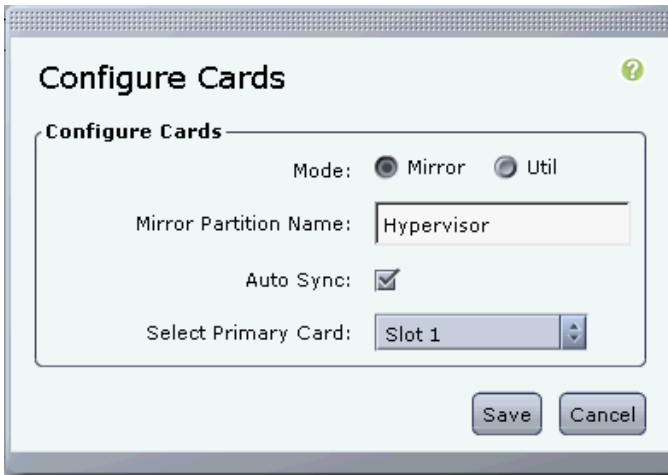
RAID 1 (mirror) can be set up by adding a card (if an additional card is not already present) to the server. The two cards will need to be synchronized for RAID 1 to become operational.

To synchronize RAID on the two cards, complete the following steps:

1. From the Cisco IMC interface browser window, click the Storage tab and choose Cisco FlexFlash.
2. In the right pane, click Controller Info tab. Under Actions, click Configure Cards.
3. In the Configure Cards window, make the following changes:
  - Mode: Mirror
  - Mirror Partition Name: Type the <<mirror\_partition\_name>>
  - Auto Sync: select the checkbox

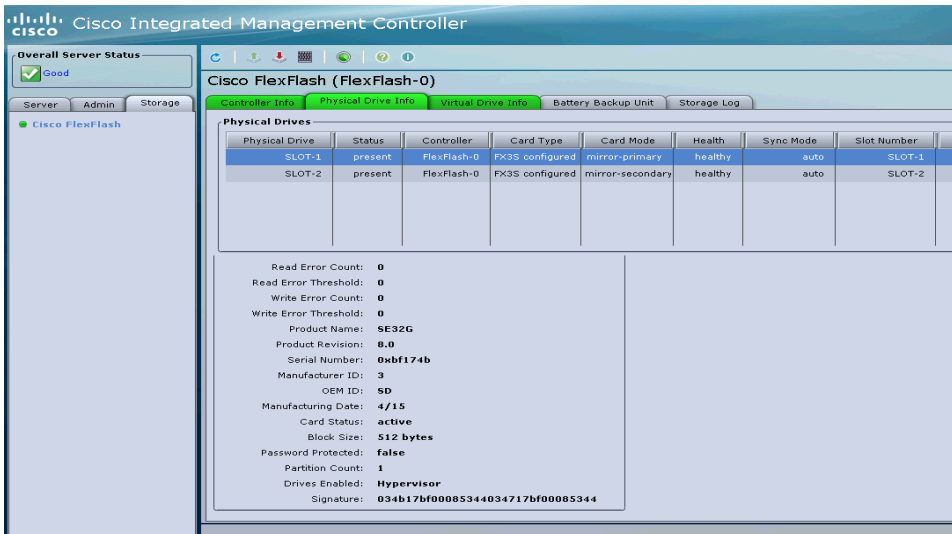


- Primary Card: select Slot 1
4. Click Save.



**Note:** The sync process will take a few minutes to complete and for the FlexFlash cards to show a healthy state.

5. Click Physical Drive Info tab to make sure the cards are healthy.



6. Click Virtual Drive Info and select the newly created virtual drive.
7. Under Actions, click Enable/Disable Virtual Drive(s).
8. Enable the checkbox on the virtual drive.
9. Click Save.
10. Under Actions, click Erase Virtual Drive(s).

**Note:** It is recommended to format the virtual drive to remove any existing data and partition information.

11. Enable the checkbox on the virtual drive.
12. Click Save.
13. Repeat steps 1 through 12 for all the servers.

## Boot Order Configuration

1. From the Cisco IMC interface browser window (do not close the virtual KVM window), click the Server tab and select BIOS.
2. Select Configure Boot Order and click OK.
3. In the Boot Order section, remove all the entries and configure the following:
  - Add Virtual Media
    - Name: KVM-CD-DVD
    - Sub Type: KVM MAPPED DVD
    - State: EnabledClick Add Device.
  - Add SD Card
    - Name: FlexFlash
    - State: Enabled
    - Order: 2
4. Click Add Device.
5. Click Save. Click Close.
6. Click Save Changes.

## 4.4 VMware vSphere 6.0 Deployment Procedure

This section provides detailed procedures for installing VMware ESXi 6.0 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Multiple methods exist for installing VMware ESXi in such an environment. This procedure uses the virtual KVM console and virtual media features of the Cisco IMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.

### Logging in to Cisco IMC Interface for Cisco UCS C-Series Standalone Servers

The following steps detail the method for logging in to the Cisco IMC interface for Cisco UCS C-Series standalone servers. You must log in to the Cisco IMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

#### All Hosts

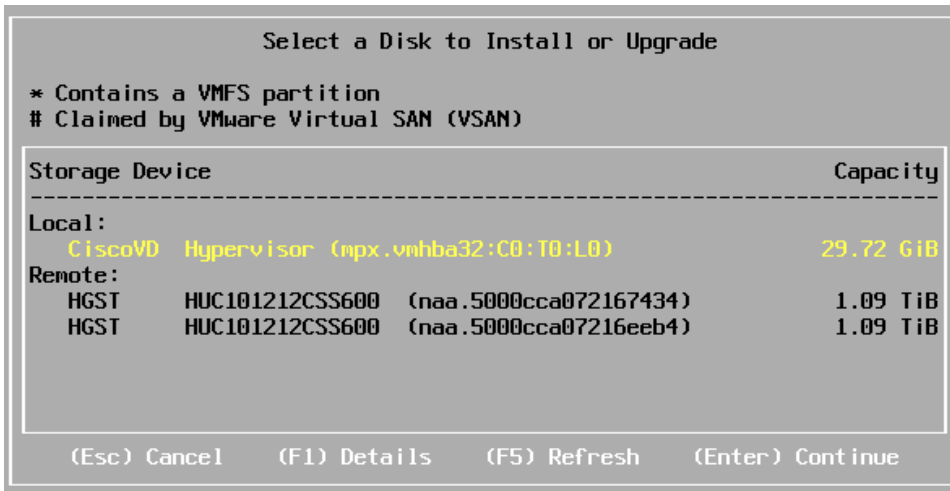
1. Obtain a copy of the Cisco Custom Image for ESXi 6.0 from <https://my.vmware.com/web/vmware/details?downloadGroup=OEM-ESXI60GA-CISCO&productId=491>.
2. Navigate to a web browser and enter the IP address for the Cisco IMC interface for the Cisco UCS C-Series. This step launches the Cisco IMC GUI application.
3. Log in to the Cisco IMC GUI using the admin user name and credentials.
4. In the main menu, select the Server tab.
5. Click Launch KVM Console.
6. From the virtual KVM console, select the Virtual Media tab.
7. Select Map CD/DVD.
8. Browse to the VMware ESXi 6.0 Custom Image ISO file and click Open. Click Map Device.
9. Select the Power menu and choose `Power Cycle System (cold boot)`. Click Yes.

## Installing VMware ESXi 6.0

The following steps describe how to install VMware ESXi on each host's Cisco FlexFlash card.

### All Hosts

1. When the system boots, the machine detects the presence of the VMware ESXi installation media.
2. Select the VMware ESXi installer from the menu that appears.
3. After the installer is finished loading, press Enter to continue with the installation.
4. After reading the end-user license agreement (EULA), accept it and continue with the installation by pressing F11.
5. Select the virtual drive that was set up previously as the installation location for VMware ESXi and press Enter to continue with the installation.



6. Select the appropriate keyboard layout and press Enter to continue.
7. Enter and confirm the root password and press Enter to continue.
8. The installer will warn you that existing partitions will be removed on the volume. Continue with the installation by pressing F11.
9. After the installation is complete, be sure to unmap the VMware ESXi installation image on the Virtual Media tab of the KVM console to help make sure that the server reboots into VMware ESXi and not the installer.
10. The Virtual Media window might warn you that it is preferable to eject the media from the guest. Because you cannot do this in this example (and the media is read-only), unmap the image anyway by selecting Yes.
11. Repeat the steps 1 through 10.
12. On the KVM tab, press Enter to reboot the server.


## Setting Up VMware ESXi Host Management Networking

The following steps describe how to add the management network for each VMware ESXi host.

### All Hosts

1. After the server finishes rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the login name and the root password previously entered during the installation process.

3. Select the Configure Management Network option.
4. Select Network Adapters and press Enter.
5. Six ports should be listed as Connected in the Status column that is displayed. These ports should correspond to two onboard LAN-on-motherboard (LOM) ports and ports 1, 2, 3, and 4 of the quad-port Broadcom PCI Express (PCIe) adapter. Select all ports and press Enter.



6. Select VLAN (optional) and press Enter.
7. Enter the VLAN ID: <<ib\_mgmt\_vlan\_id>>. Press Enter.
8. From the Configure Management Network menu, configure the IP address of the management interface by selecting the IP Configuration option. Press Enter.
9. Use the space bar to select set static IP address and network configuration.
10. Enter the IP address for managing the VMware ESXi host: <<esxi\_host\_mgmt\_ip>>.
11. Enter the subnet mask for the VMware ESXi host: <<esxi\_host\_mgmt\_netmask>>.
12. Enter the default gateway for the VMware ESXi host: <<esxi\_host\_mgmt\_gateway>>.
13. Press Enter to accept the changes to the IP configuration.
14. Enter the IPv6 configuration menu.
15. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
16. Enter the menu to configure the DNS settings.
17. Because the IP address is assigned manually, the DNS information must also be entered manually.
18. Enter the primary DNS server's IP address: <<nameserver\_ip>>.
19. (Optional) Enter the secondary DNS server's IP address.
20. Enter the FQDN for the VMware ESXi host: <<esxi\_host\_fqdn>>.
21. Press Enter to accept the changes to the DNS configuration.
22. Exit the Configure Management Network submenu by pressing Esc.
23. Press Y to confirm the changes and reboot the server.
24. Log out of the VMware Console by pressing Esc.

## Downloading VMware vSphere Client and vSphere Remote Command Line

The following steps provide details for downloading the VMware vSphere Web Client and installing the remote command line.

1. Open a web browser on a management workstation and navigate to the management IP address of one of the VMware ESXi hosts.
2. Download and install both the VMware vSphere Client and the Microsoft Windows version of the VMware vSphere remote command line.

## Logging in to VMware ESXi Hosts Using the VMware vSphere Client

This step provides details for logging into each VMware ESXi host using the VMware vSphere Client.

### All Hosts

1. Open the recently downloaded VMware vSphere Web Client and enter the IP address of the host to which you want to connect: <<esxi\_host\_mgmt\_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click the Login button to connect.

## Setting up the VMkernel Ports and Virtual Switch

The following steps provide details for setting up VMkernel ports and virtual switches.

### All Hosts

1. In the VMware vSphere Client, select the host on the left pane.
2. Select the Configuration tab.
3. Select the Networking link from the Hardware section.
4. Select the Properties link in the field to the right of vSwitch0.
5. Select the vSwitch configuration and click Edit.
6. On the General tab, change the MTU to 9000.
7. On the NIC Teaming tab, change all adapters so that they are active adapters by clicking each individual adapter and using the Move Up button to the right.
8. Close the properties for vSwitch0 by clicking OK.
9. Select the Management Network configuration and click Edit.
10. Change the network label to `VMkernel-MGMT` and select the Management Traffic checkbox.
11. Finalize the edits for the management network by clicking OK.
12. Select the VM Network configuration and click Edit.
13. Change the network label to `IB-MGMT Network` and enter <<var\_ib-mgmt\_vlan\_id>> in the VLAN ID (Optional) field.
14. Finalize the edits for the VM network by clicking OK.
15. Click Add to add a network element.
16. Select Virtual Machine.
17. Enter `NFS-Network` for the network label and enter the VLAN ID: <<nfs\_vlan\_id>>.
18. Click Next.
19. Click Finish.
20. Click Add to add a network element.

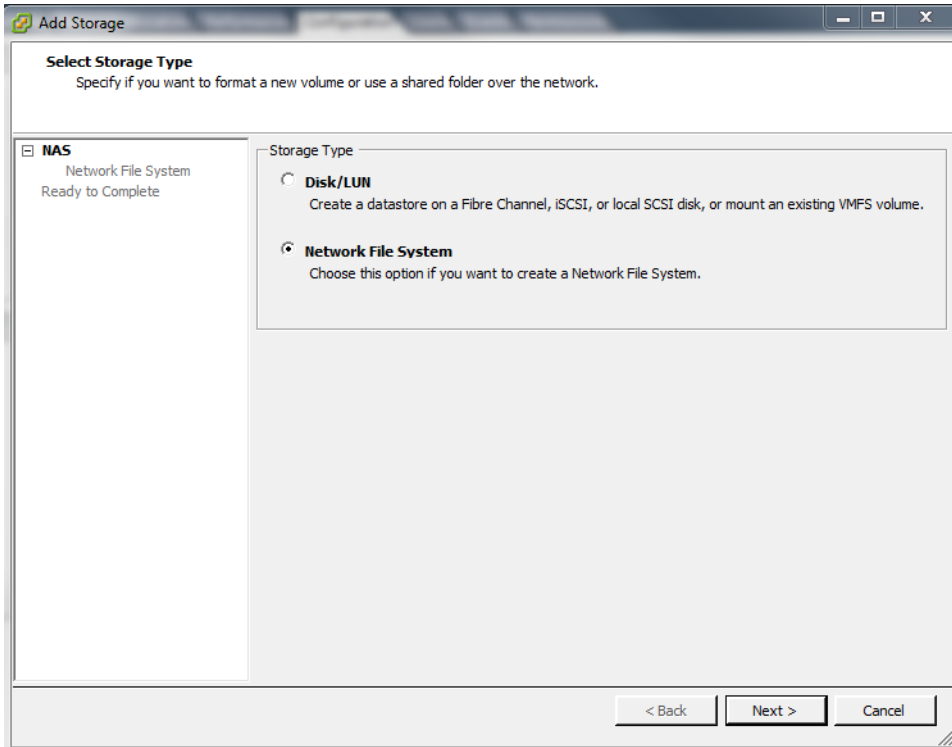
21. Select the VMkernel button and click Next.
22. Change the network label to VMkernel-NFS and enter the VLAN ID (optional): <<nfs\_vlan\_id>>.
23. Continue with the NFS VMkernel creation by clicking Next.
24. For the NFS VLAN interface for the host, enter <<esxi\_host\_nfs\_ip>>  
<<esxi\_host\_nfs\_netmask>>.
25. Continue with the NFS VMkernel creation by clicking Next.
26. Finalize the creation of the NFS VMkernel interface by clicking Finish.
27. Select the VMkernel-NFS configuration and click Edit.
28. Change the MTU to 9000.
29. Finalize the edits for the VMkernel NFS network by clicking OK.
30. Click Add to add a network element.
31. Select the VMkernel button and click Next.
32. Change the network label to VMkernel-vMotion and enter the VLAN ID (optional):  
<<vmotion\_vlan\_id>>.
33. Select the checkbox to use this port group for VMware vMotion.
34. Continue with the VMware vMotion VMkernel creation by clicking Next.
35. For the VMware vMotion VLAN interface for the host, enter: <<esxi\_host\_vmotion\_ip>>  
<<esxi\_host\_vmotion\_netmask>>.
36. Continue with the VMware vMotion VMkernel creation by clicking Next.
37. Finalize the creation of the VMware vMotion VMkernel interface by clicking Finish.
38. Select the VMkernel vMotion configuration and click Edit.
39. Change the MTU to 9000.
40. Finalize the edits for the VMware vMotion VMkernel network by clicking OK.
41. Click Add to add a network element.
42. Leave the virtual machine connection type selected and click Next.
43. Change the network label to VM-Network and enter the VLAN ID (optional) :  
<<vmtraffic\_vlan\_id>>.
44. Click Next.
45. Click Finish.
46. Close the dialog box to finalize the VMware ESXi host networking setup.

## Mounting Required Datastores

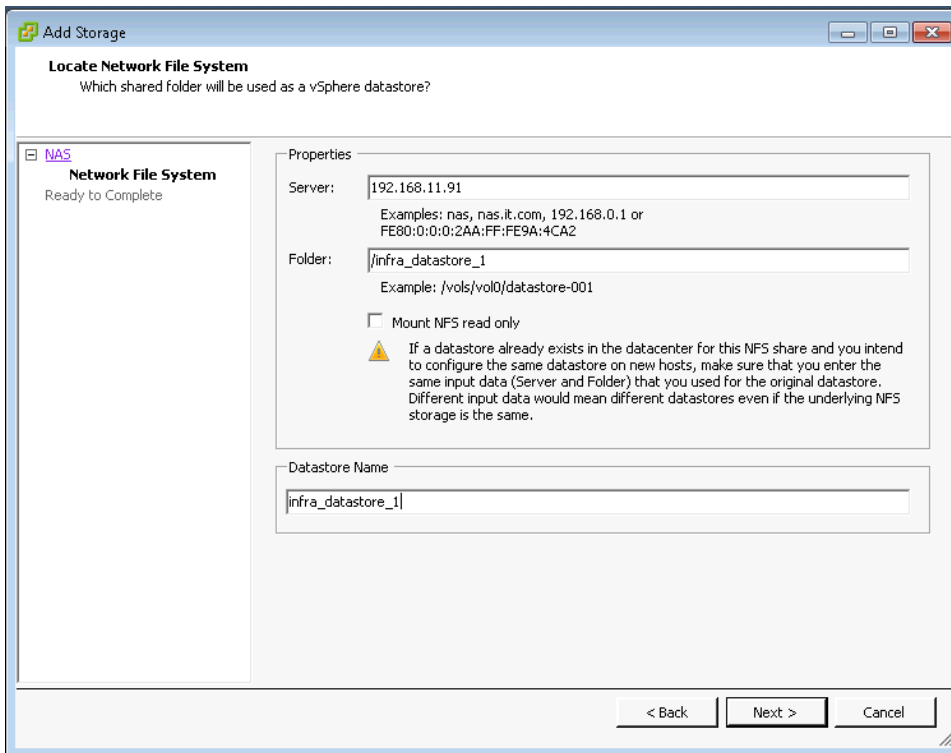
This step provides details for mounting the required datastores.

### All Hosts

1. In each VMware vSphere Client, select the host on the left pane.
2. Go to the Configuration tab to enable configurations.
3. Click the Storage link in the Hardware box.
4. In the right pane, in the Datastore section, click Add Storage.




5. The Add Storage wizard appears. Select Network File System and click Next.
6. Enter the server IP address: <<nfs\_infra\_datastore\_1 lif ip>>.
7. Enter the path for the NFS export: /infra\_datastore\_1.
8. Make sure that the Mount NFS read only checkbox is left unchecked.
9. Enter the datastore name: infra\_datastore\_1.

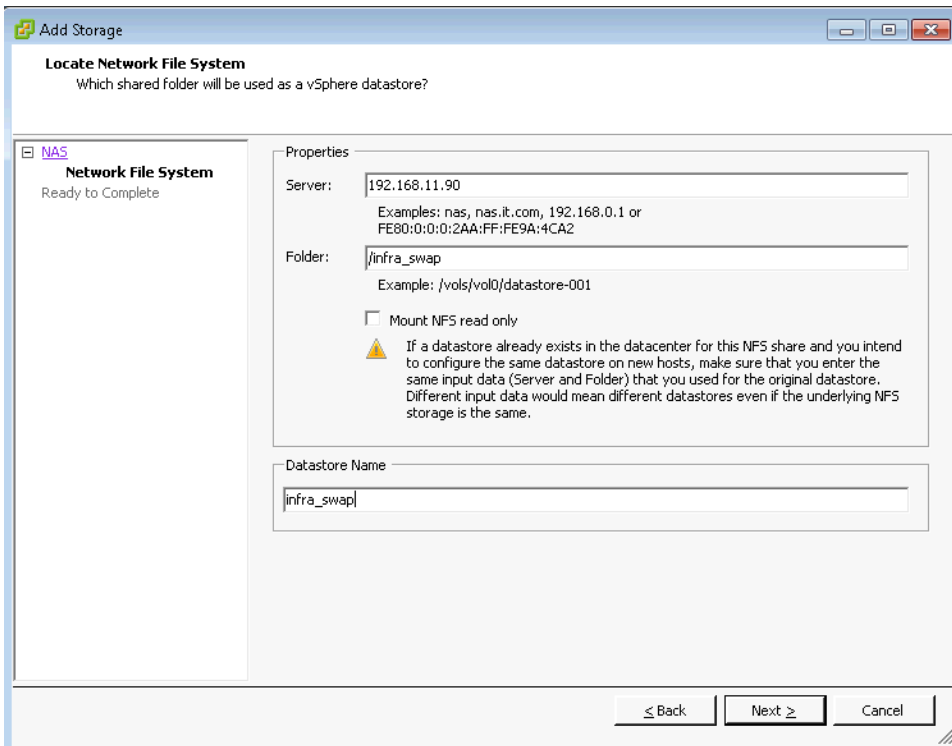


10. Continue with the NFS datastore creation by clicking Next.
11. Finalize the creation of the NFS datastore by clicking Finish.
12. In the right pane, in the Datastore section, click Add Storage.  
The Add Storage wizard appears.
13. Select Network File System and click Next.





14. Enter the server IP address: <<nfs\_infra\_swap lif ip>>.
15. Enter the path for the NFS export: /infra\_swap.
16. Make sure the Mount NFS read only checkbox is left unchecked.
17. Enter the datastore name: infra\_swap.



18. Continue with the NFS datastore creation by clicking Next.
19. Finalize the creation of the NFS datastore by clicking Finish.

## Moving the Virtual Machine Swap-File Location

These steps provide details for moving the virtual machine swap-file location.

### All Hosts

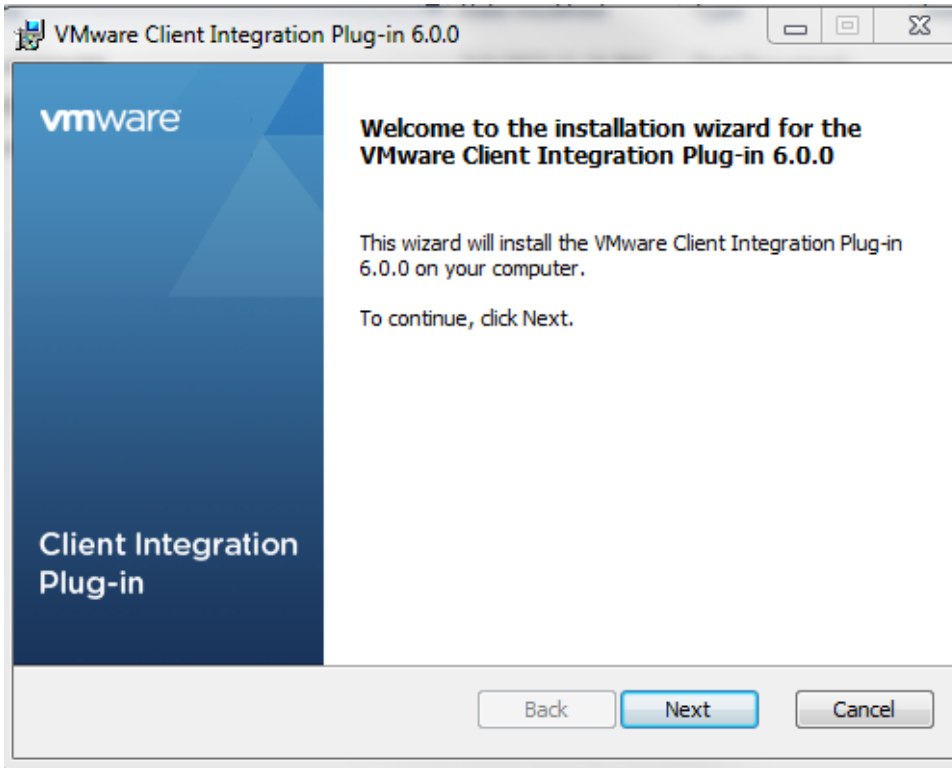
1. Select the host in the left pane in the VMware vSphere Client.
2. Go to the Configuration tab to enable configuration.
3. Click the Virtual Machine Swapfile Location link in the Software box.
4. In the right pane, click Edit.
5. Select the `Store the swap file in a swap file datastore selected below` button.
6. Select the `infra_swap` datastore.
7. Finalize the movement of the swap-file location by clicking OK.

## 4.5 VMware vCenter 6.0 Deployment Procedure

The procedures in the following subsections provide detailed instructions for installing VMware vCenter 6.0 in a FlexPod Express environment. A VMware vCenter Server will be configured after the procedures are completed.

### Installing the Client Integration Plug-in

1. Download the .iso installer for the vCenter Server Appliance and Client Integration Plug-in.
2. Mount the ISO image to the Windows virtual machine or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter Server appliance.
3. In the software installer directory, navigate to the `vcsa` directory and double-click `VMware-ClientIntegrationPlugin-6.0.0.exe`. The Client Integration Plug-in installation wizard appears.

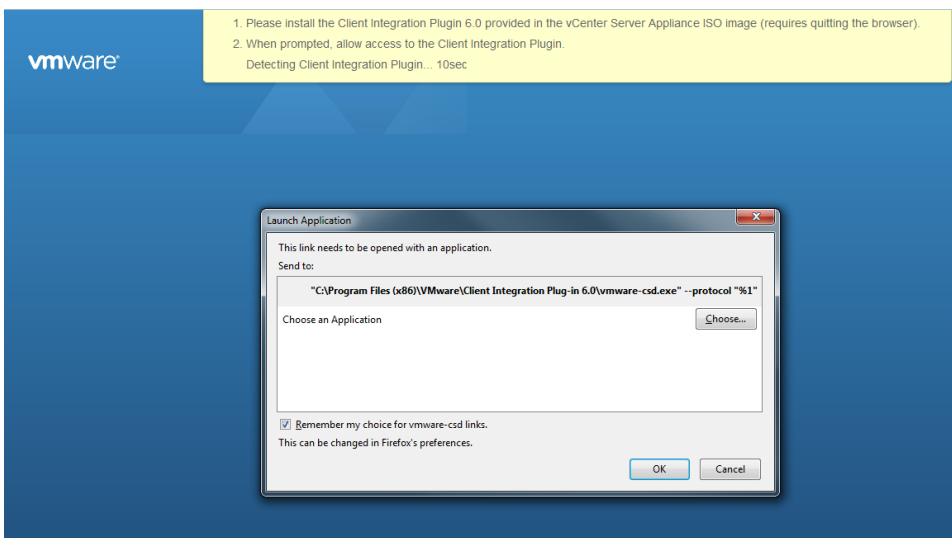


4. On the Welcome page, click Next.
5. Read and accept the terms in the EULA and click Next.
6. Click Next.
7. Click Install.


## Building the VMware vCenter Virtual Machine

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

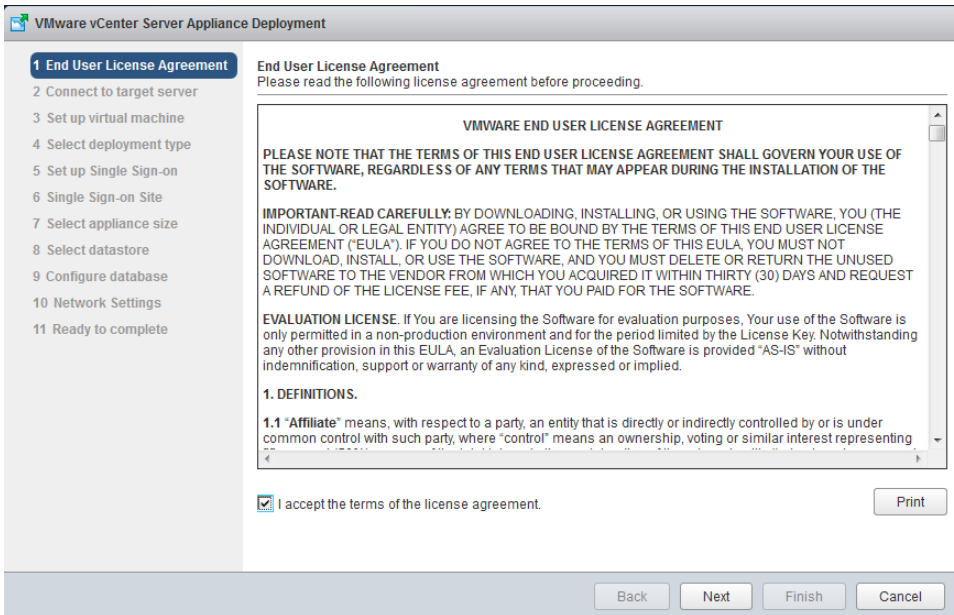
1. In the software installer directory, double-click `vcsa-setup.html`.
2. Allow the plug-in to run on the browser when prompted.



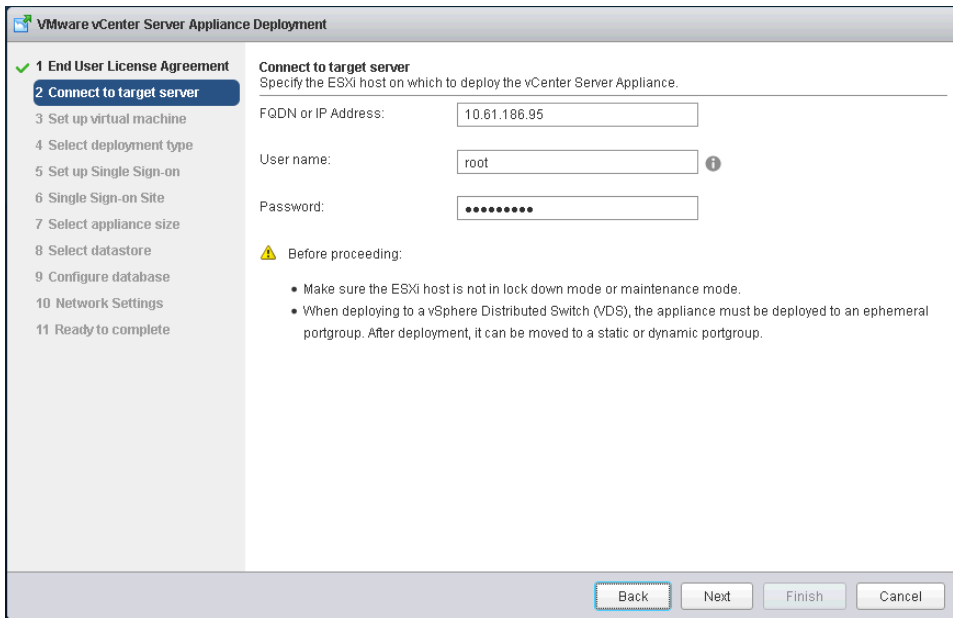
3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.



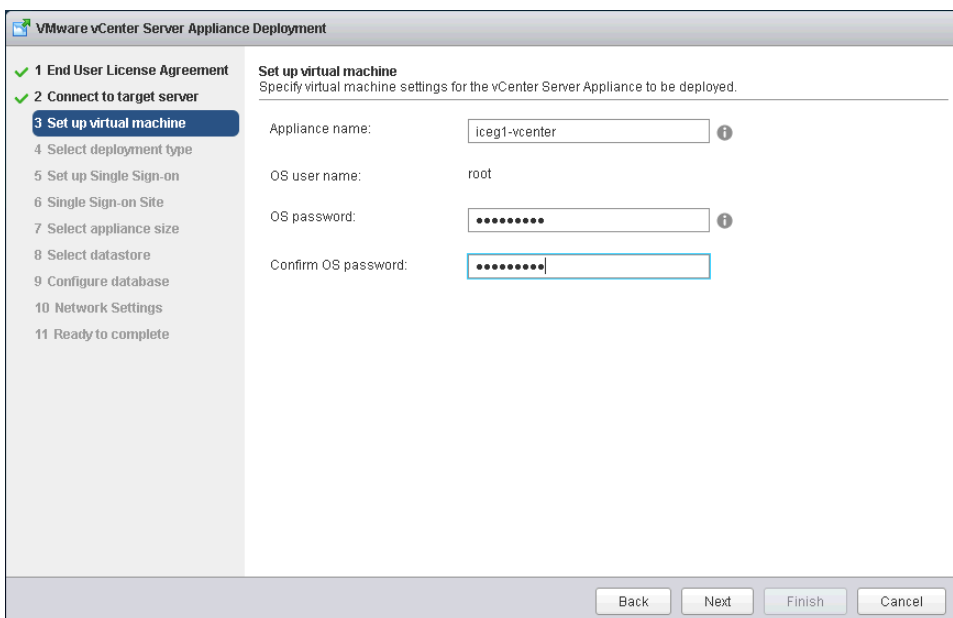
4. Read and accept the EULA, and click Next.



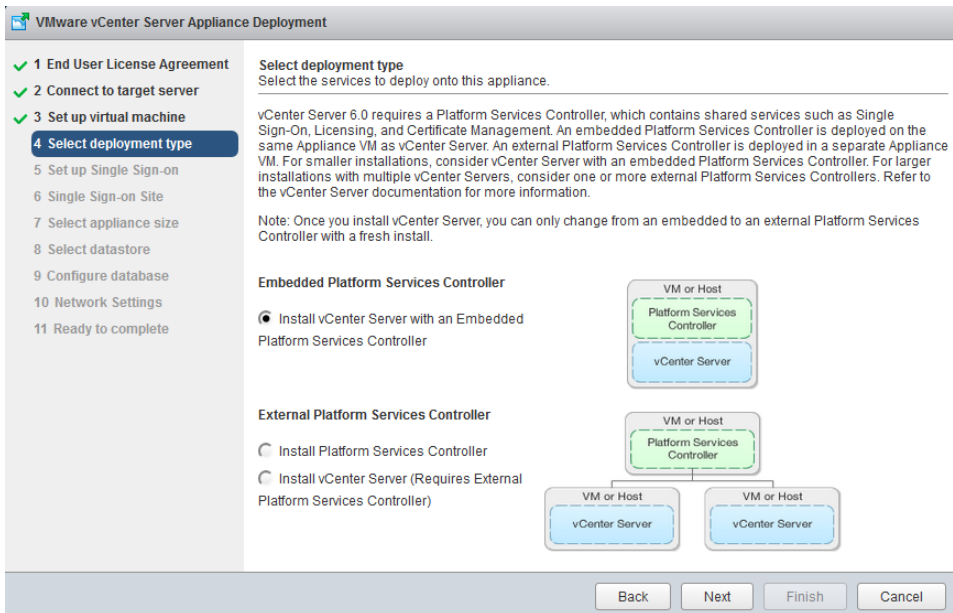
5. Enter the ESXi host name, user name, and password.



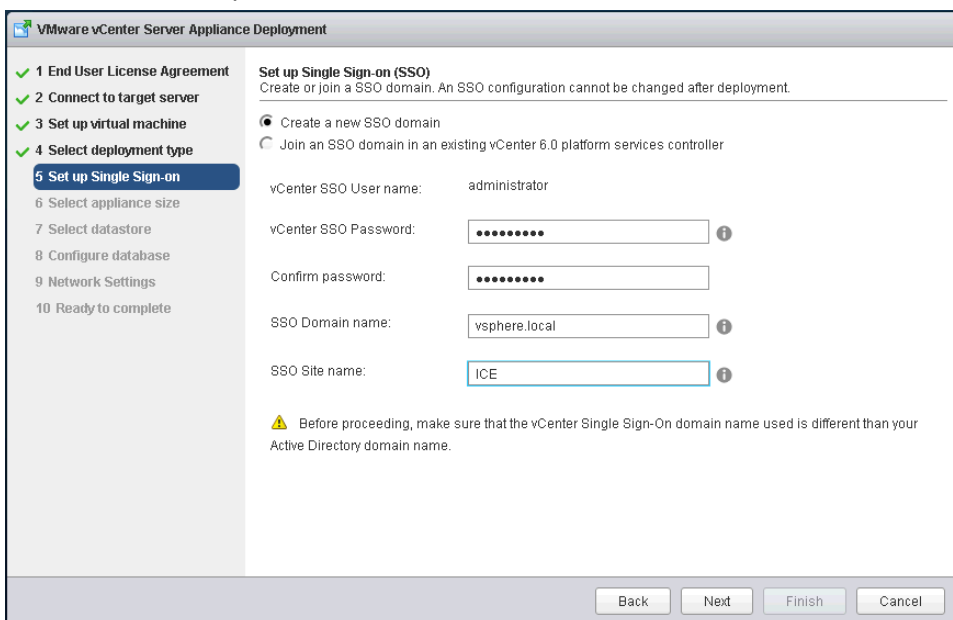
6. Click Yes to accept the certificate.
7. In the “Set up virtual machine” page, enter the appliance name and password details.



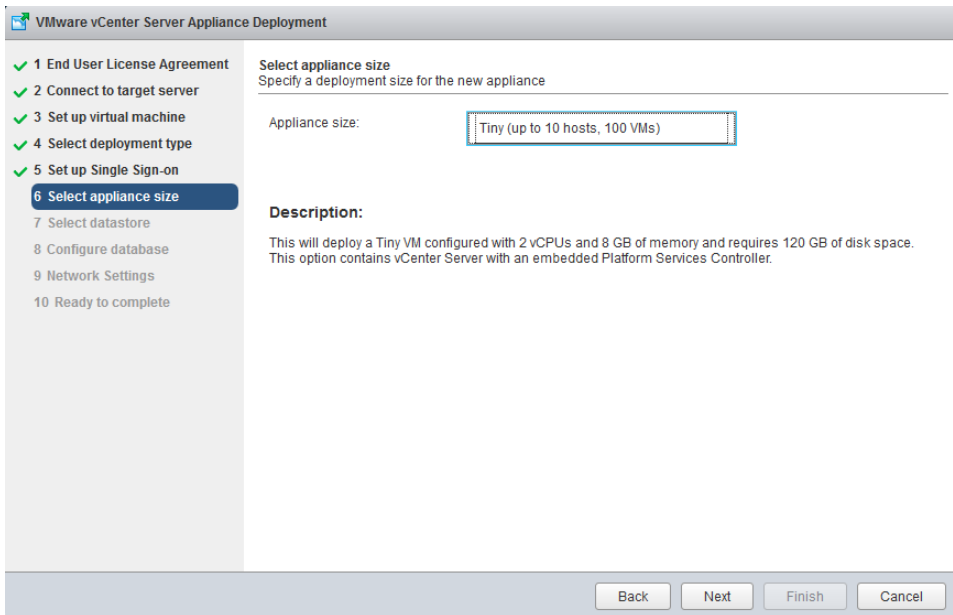
8. In the Select deployment type page, choose “Install vCenter Server with an embedded Platform Services Controller” and click Next.



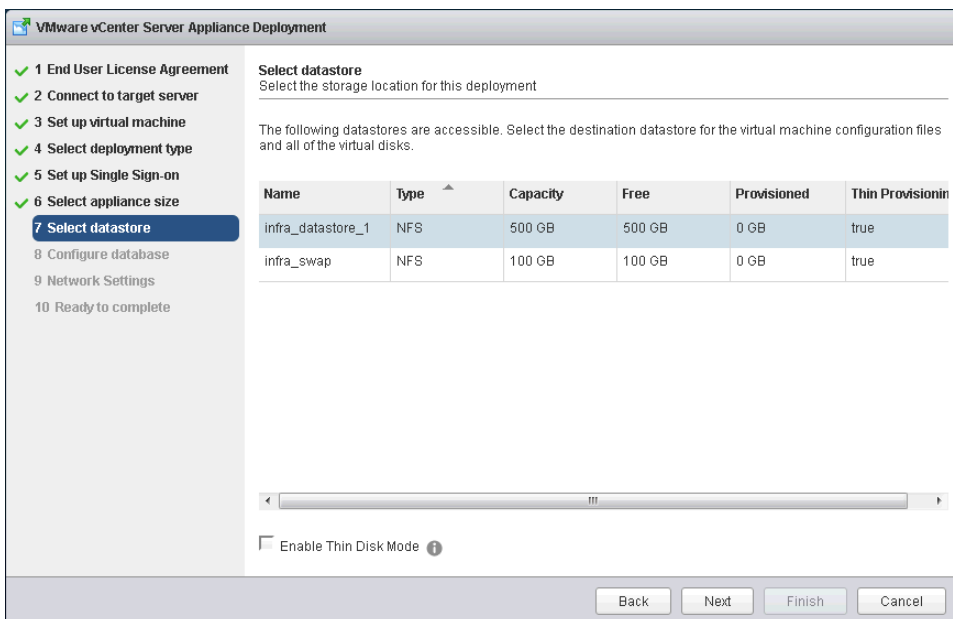
9. In the Set up Single Sign-On page, select **Create a new SSO domain**.
10. Enter the SSO password, domain name, and site name. Click **Next**.



11. Select the appliance size and click **Next**. For example, **Tiny** (up to 10 hosts, 100 VMs).

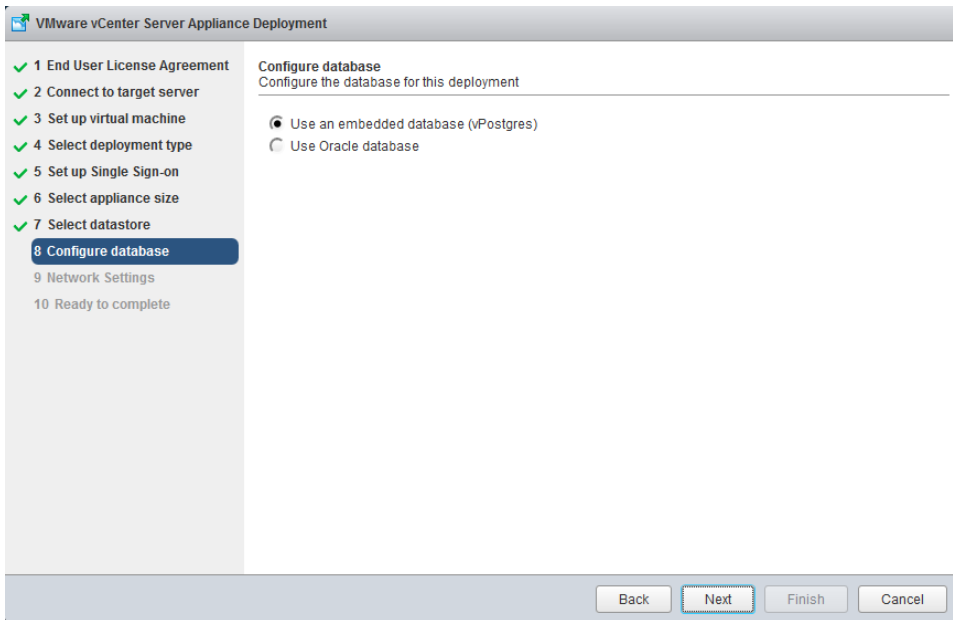


In the Select datastore page, choose `infra_datastore_1`.



12. Click Next.

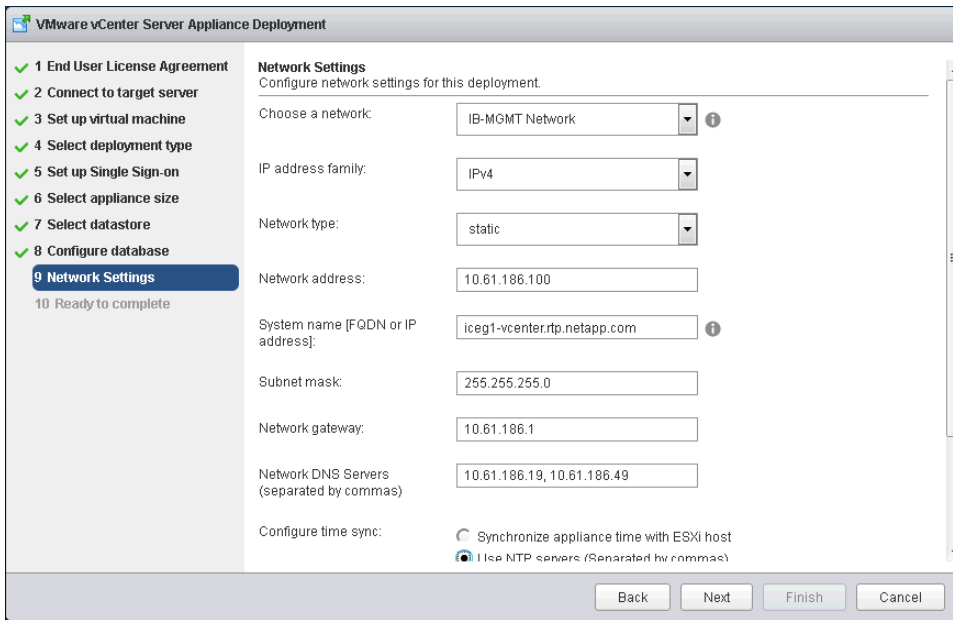
13. In the Configure database page, select the embedded database option and click Next.



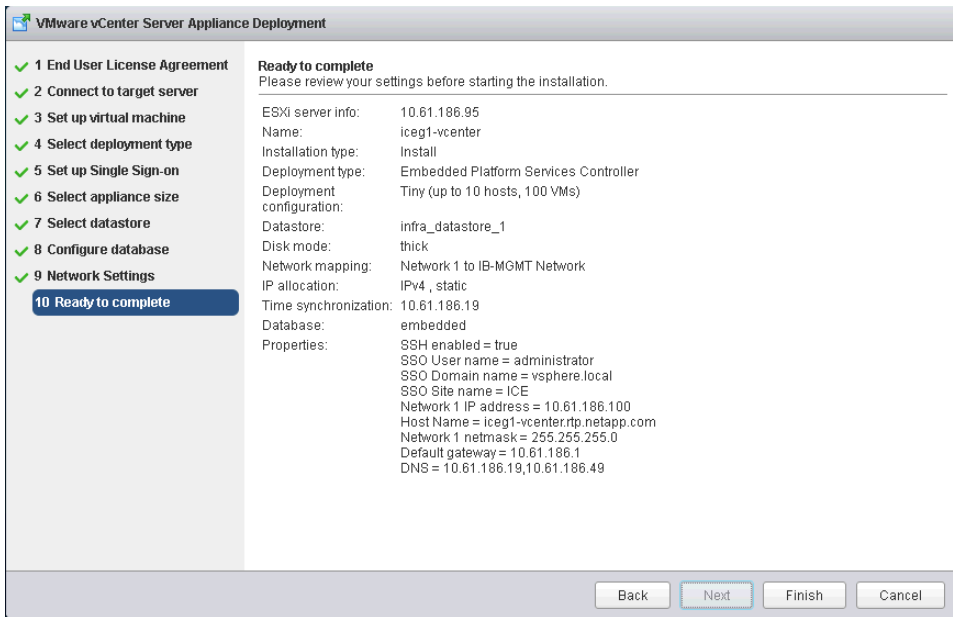
14. In the Network Settings page, configure the following settings:

- Choose a Network: IB-MGMT-Network
- IP address family: IPV4
- Network type: static
- Network address: <<var\_vcenter\_ip>>
- System name: <<var\_vcenter\_fqdn>>
- Subnet mask: <<var\_vcenter\_subnet\_mask>>
- Network gateway: <<var\_vcenter\_gateway>>
- Network DNS Servers: <<var\_dns\_server>>
- Configure time sync: Use NTP servers
- (Optional). Enable SSH

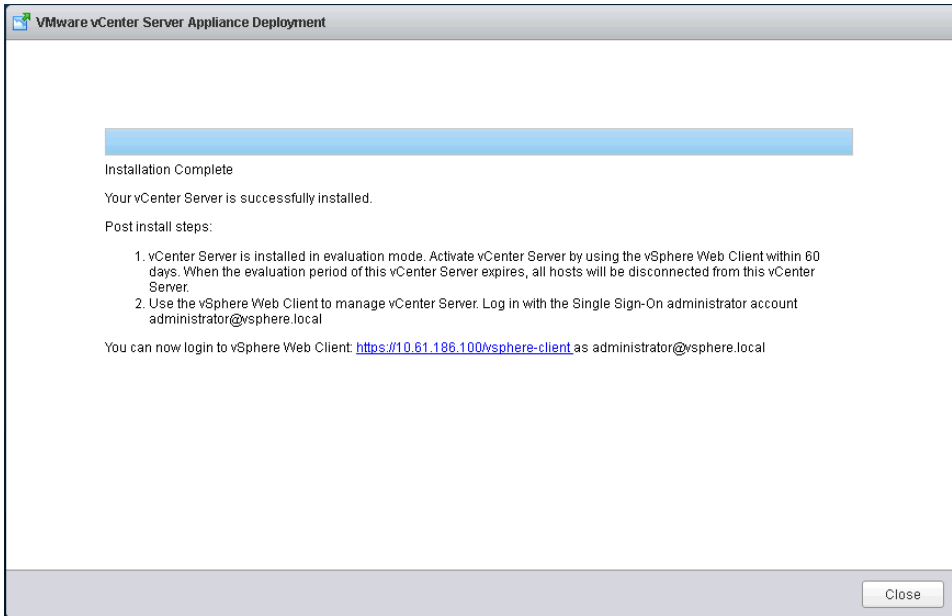




15. Review the configuration and click Finish.



16. The vCenter appliance installation will take few minutes to complete.



## Configuring ESXi Core Dump Collector

It is recommended to configure a core dump location for the ESXi hosts to store the state of working memory in the event of host failure. The dumps can be used for debugging purposes.

The below steps describe the procedure to setup a Core Dump collector:

1. From the management workstation, open the VMware vSphere CLI command prompt.
2. Set each ESXi host to core dump to the ESXi Dump Collector by running the following commands:

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump network set --
interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --server-port 6500

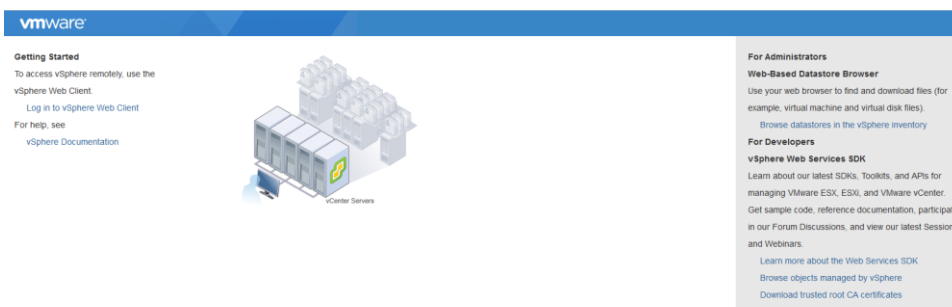
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump network set --
enable true

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump network get

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump network check
```

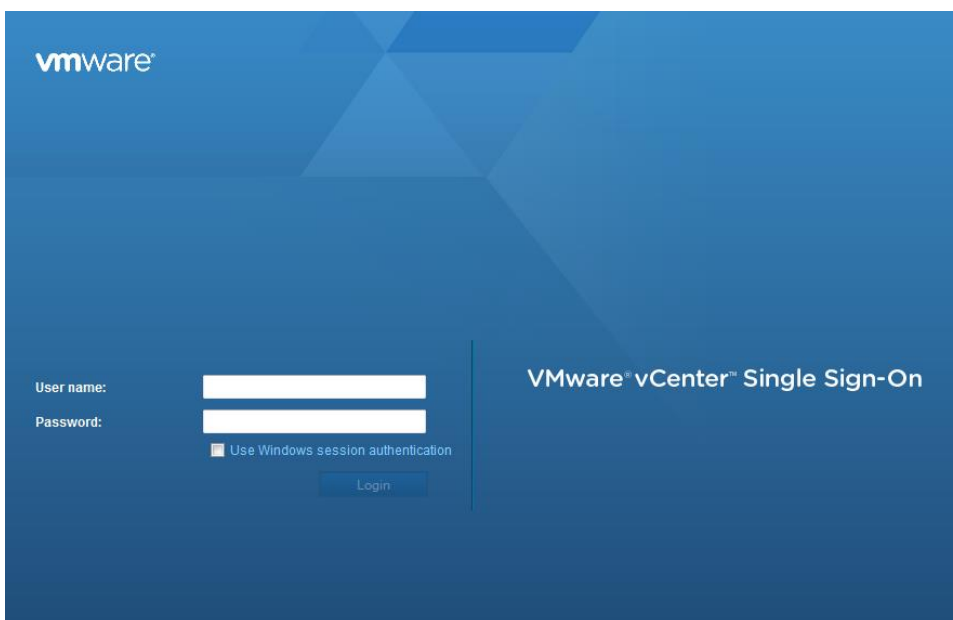

## Setting Up VMware vCenter Server

1. Using a web browser, navigate to [https://<<var\\_vcenter\\_ip>](https://<<var_vcenter_ip>).

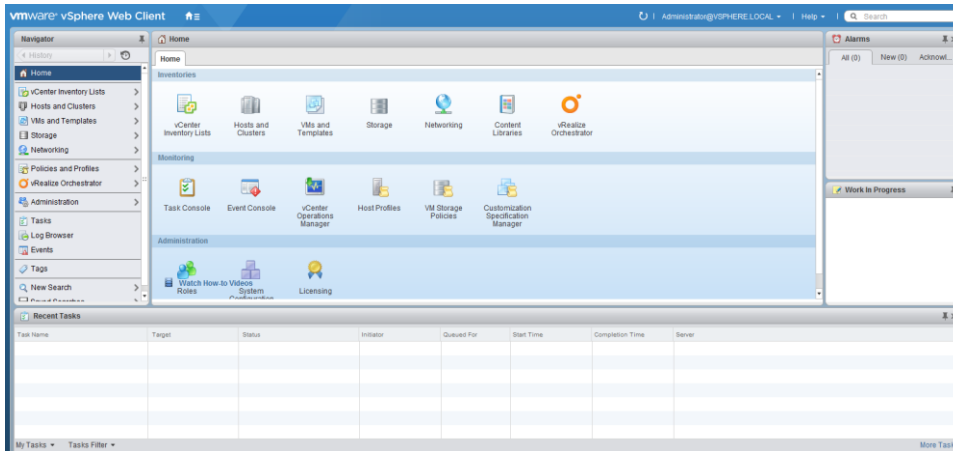


2. Click Log in to vSphere Web Client.

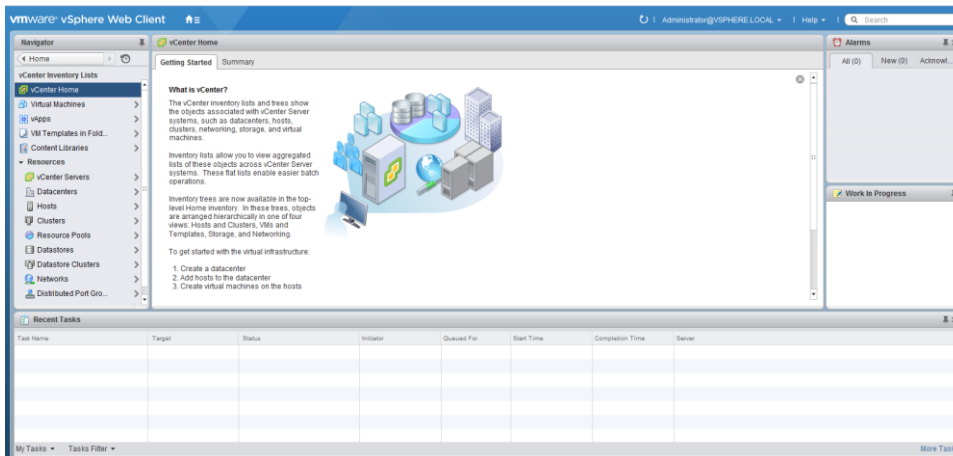
3. Click OK if the Launch Application window appears.



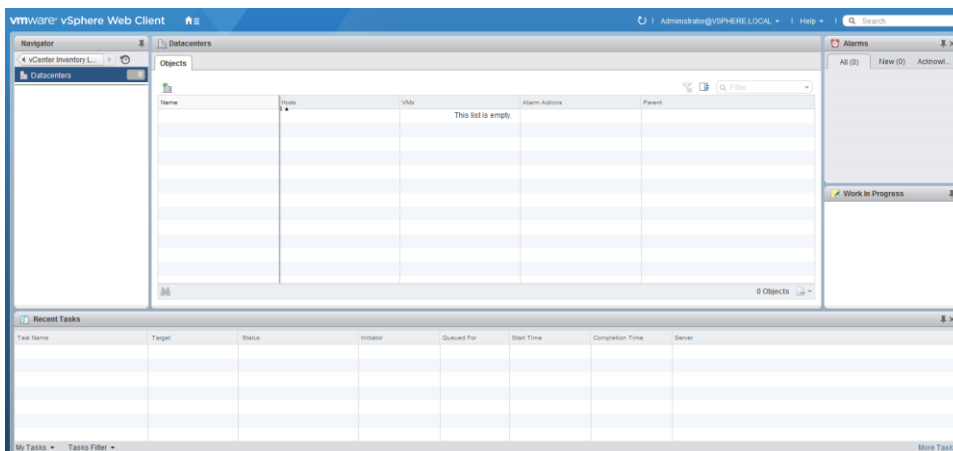
4. Log in using single sign-on user name and password created during the vCenter installation.



5. Navigate to vCenter Inventory Lists on the left pane.



6. Under Resources, click Datacenters in the left pane.

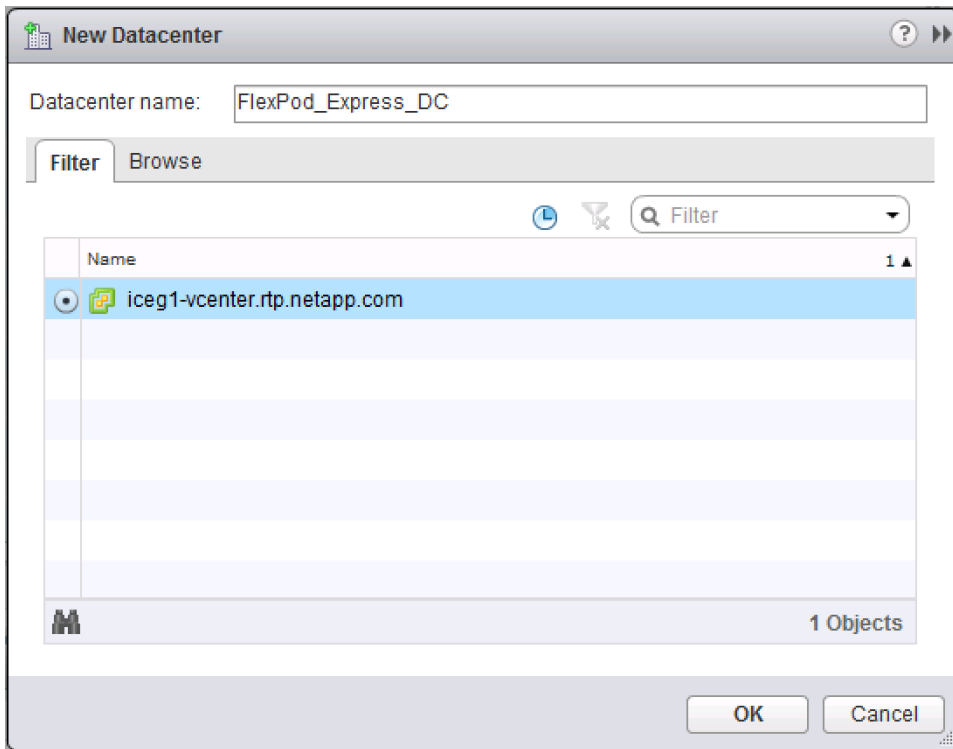


7. To create a data center, click the icon in the center pane that has a green plus symbol above it.

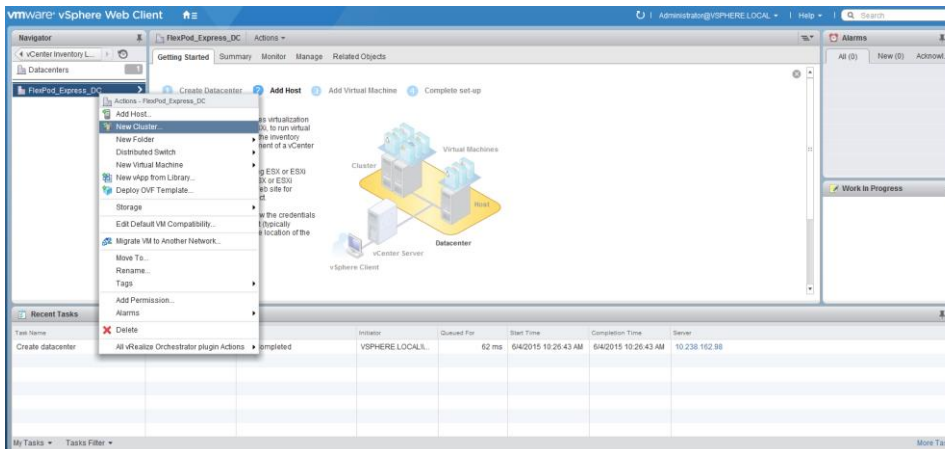
8. Type FlexPod\_Express\_DC in the Datacenter name field.

9. Select the vCenter name/IP option.

10. Click OK.



11. Right-click the data center FlexPod\_Express\_DC from the list in the center pane. Click New Cluster.



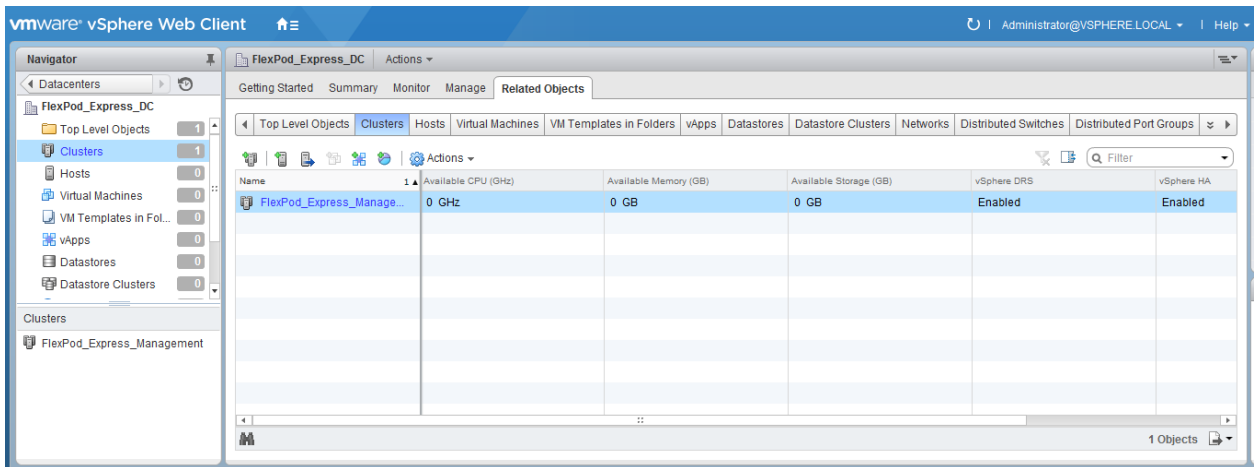
12. Name the cluster FlexPod\_Express\_Management.

13. Check the box beside DRS. Leave the default values.

14. Check the box beside vSphere HA. Leave the default values.

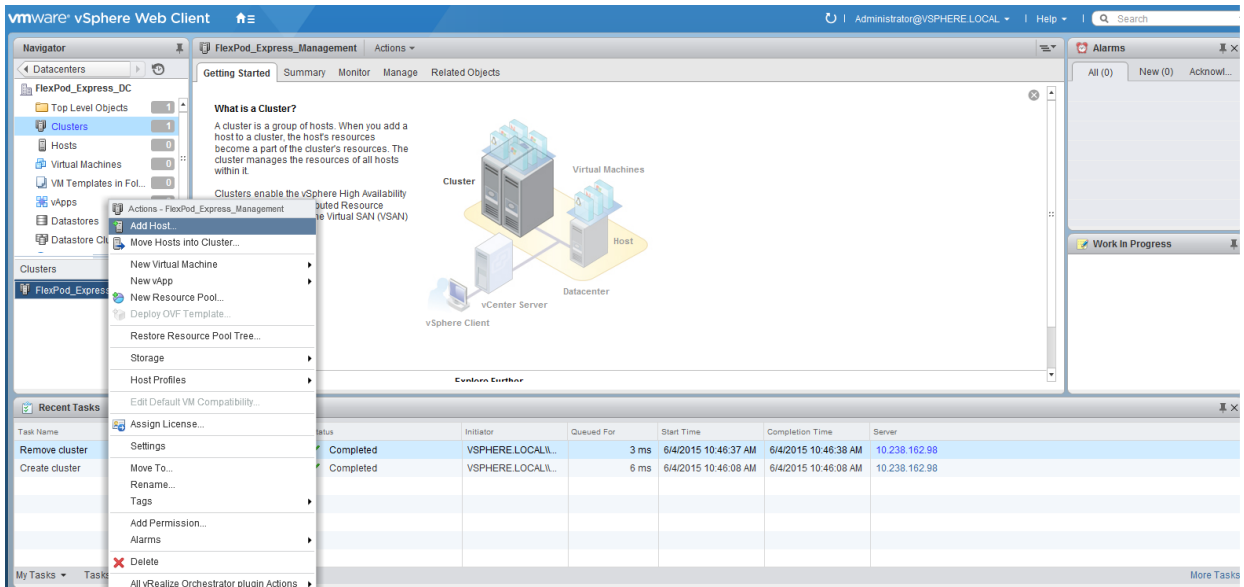
Name	FlexPod_Express_Management
Location	FlexPod_Express_DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable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

15. Click OK to create the new cluster.
16. On the left pane, double-click FlexPod\_Express\_DC.
17. Click Clusters.



18. Under the Clusters pane, right-click the FlexPod\_Express\_Management.

19. Click Add Host.



20. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.

21. Type root as the user name and the root password. Click Next to continue.

22. Click Yes to accept the certificate.

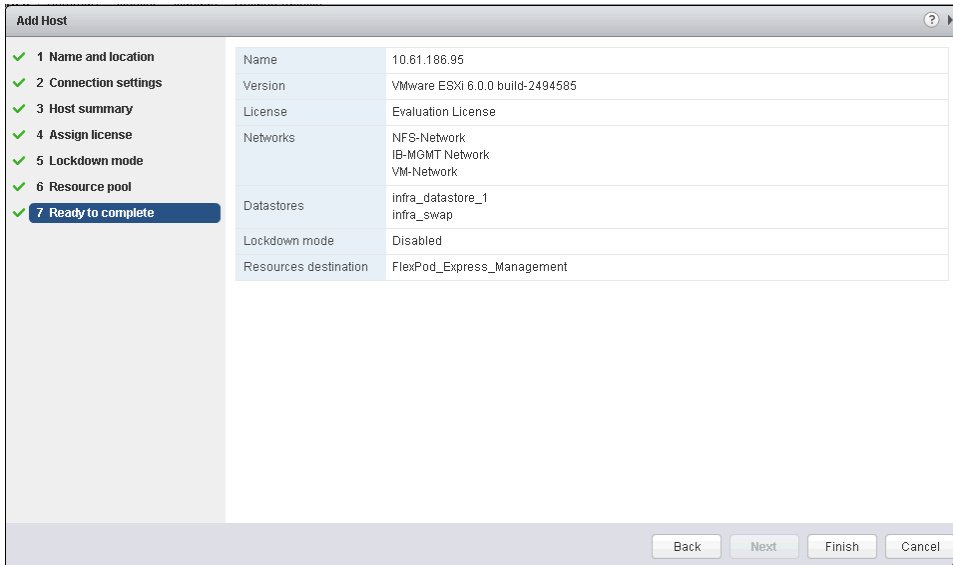
23. Review the host details and click Next to continue.

24. Assign a license and click Next to continue.

25. Click Next to continue.

26. Click Next to continue.

27. Review the configuration parameters. Click Finish to add the host.



28. Repeat the steps 18 through 27 to add the remaining VMware ESXi hosts to the cluster.

**Note:** Two VMware ESXi hosts will be added to the cluster for the FlexPod Express small configuration.

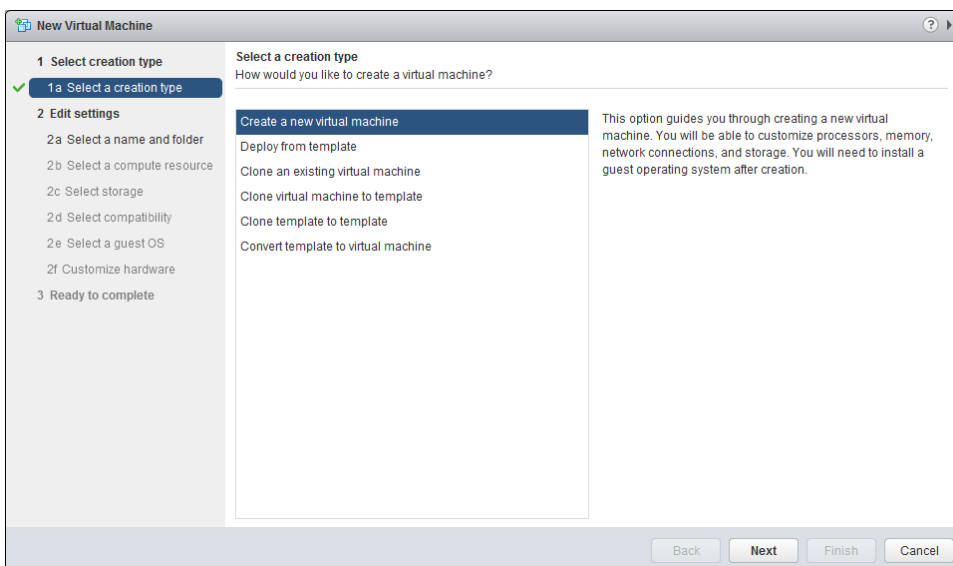
**Note:** Four VMware ESXi hosts will be added to the cluster for the FlexPod Express medium configuration.

## Setting Up a Microsoft Windows Template

To create a Microsoft Windows template, complete the following steps:

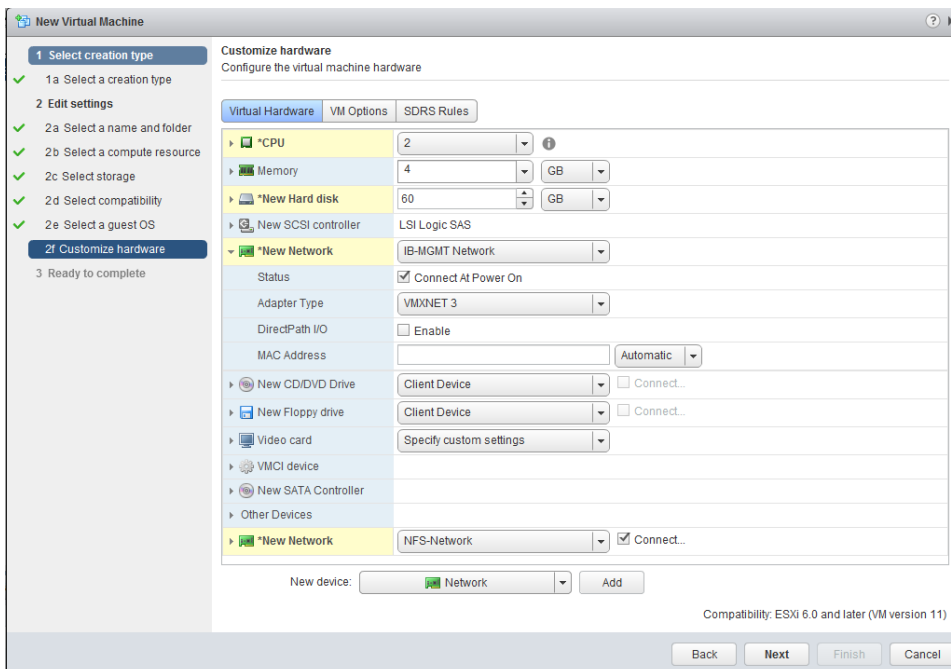
**Note:** Download Microsoft Windows 2012 R2 (64-bit) and upload it to a datastore.

1. Log in to the VMware vCenter Server by using the VMware vSphere Client.
2. Navigate to vCenter Inventory Lists > Clusters > FlexPod\_Express\_Management.
3. Right-click the cluster and select New Virtual Machine.
4. Select Create a new virtual machine and click Next.

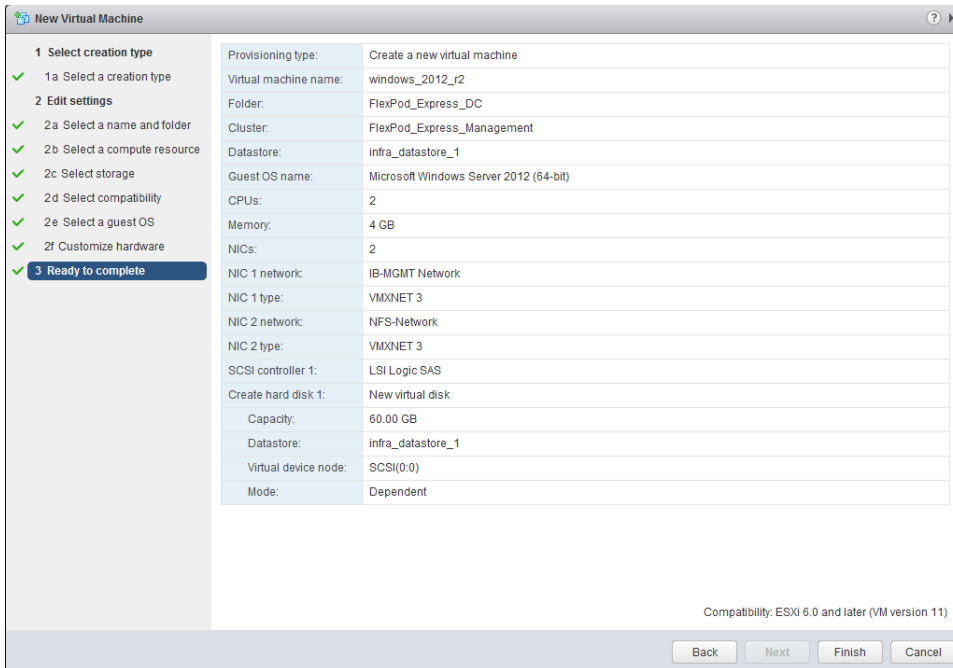




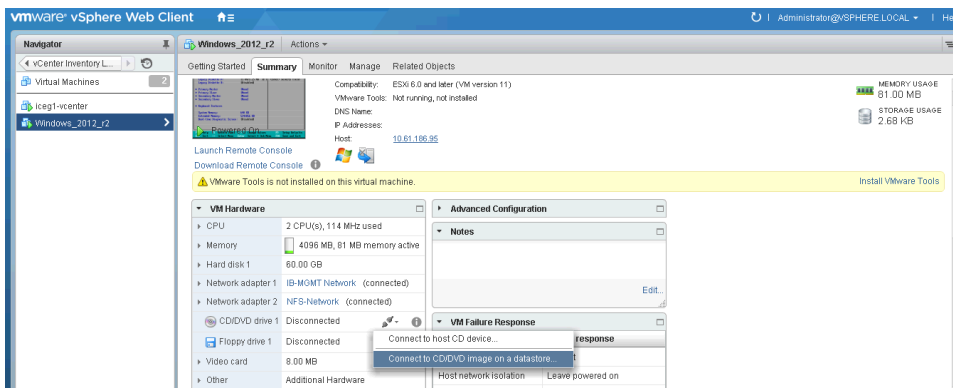
5. Enter a name for the virtual machine and select `FlexPod_Express_DC`. Click Next.
6. Make sure that the `FlexPod_Express_Management` cluster is selected and click Next.
7. Select `infra_datastore_1` and click Next.
8. In the Select compatibility page, select ESXi6.0 and later. Click Next.
9. Verify that the Microsoft Windows option and the Microsoft Windows Server 2012 (64-bit) version are selected. Click Next.
10. Select the Virtual Hardware tab and customize the hardware as follows:
  - CPU: 2
  - Memory: 4GB.
  - New Hard disk: 60GB.
  - New Network: IB-MGMT Network
  - Select Connect At Power on.
  - Adapter Type: VMXNET 3
11. From the New device list, select Network and click Add.
  - New Network: NFS-Network
  - Select Connect At Power on.
  - Adapter Type: VMXNET 3



12. Select the VM Options tab. Under Boot Options, select Force BIOS setup.
13. Click Next.
14. Review the virtual machine settings and click Finish.

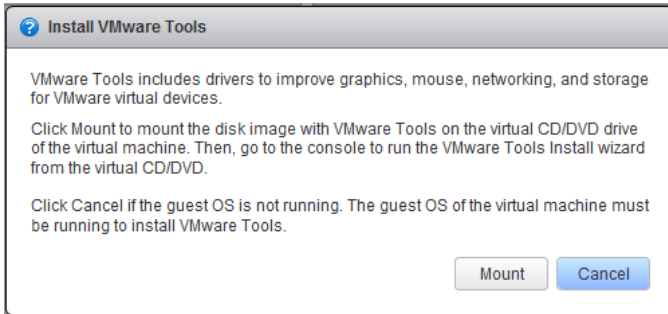


15. Navigate to vCenter Inventory Lists > Clusters > FlexPod\_Express\_Management.
16. In the Virtual machines pane, select the newly created VM. In the center pane, click the Summary tab.
17. Right-click the VM and click Power on.
18. Right-click the virtual machine and select Open Console.
19. In the Summary tab, expand the VM Hardware section.
20. Click the plug icon next to CD/DVD drive 1 and select Connect to CD/DVD image on a datastore.



21. Navigate to the Microsoft Windows 2012 ISO image.
22. From the VM console, click Send Ctrl+Alt+Delete.
23. The Microsoft Windows installer boots. Select the appropriate language, time and currency format, and keyboard. Click Next.
24. Click Install Now. Enter the product license key and click Next.
25. Select "Windows Server 2012 R2 Standard (Server with a GUI)" and click Next.
26. Read and accept the license terms and click Next.

27. Select Custom (advanced). Make sure that Disk 0 Unallocated Space is selected. Click Next to allow the Microsoft Windows installation to complete.
28. After the Microsoft Windows installation is complete and the virtual machine has rebooted, enter and confirm the administrator password. Click Finish.
29. Log in into the VM desktop.
30. From the vSphere Web Client, click `Install VMware Tools` in the VM Summary tab.
31. Click Mount.



32. If prompted to eject the Microsoft Windows installation media before running the setup for the VMware tools, click OK. Then click OK again.
  33. From the connected CD drive, run `setup64.exe`.
  34. In the VMware Tools installer window, click Next.
  35. Make sure that Typical is selected and click Next.
  36. Click Install.
  37. If prompted to trust software from VMware, select the checkbox to always trust and click Install.
  38. Click Finish.
  39. Click Yes to restart the virtual machine.
  40. After the reboot is complete, select `Send Ctrl+Alt+Del` and then enter the password to log in to the virtual machine.
  41. Set the time zone for the virtual machine and the IP address, gateway, and host name.
- Note:** A reboot is required.
42. Log back in to the virtual machine and download and install all required Microsoft Windows updates.
- Note:** This process requires several reboots.
43. Right-click the virtual machine in VMware vCenter and click Clone to Template.
  44. Enter the name `windows_2012_r2_template` for the clone.
  45. Select the data center `FlexPod_Express_DC`. Click Next.
  46. Select the cluster `FlexPod_Express_Management` as the target cluster to host the template. Click Next.
  47. Select `infra_datastore_1`. Click Next.
  48. Click Finish.

## 4.6 NetApp Virtual Storage Console 6.0 Deployment Procedure

This section provides detailed instructions to deploy NetApp Virtual Storage Console (VSC) 6.0.

## NetApp VSC 5.0 Pre-Installation Considerations

The following licenses are required to run NetApp VSC on storage systems that run clustered Data ONTAP 8.3

- Protocol licenses (NFS)
- NetApp SnapRestore (for backup and recovery)
- NetApp SnapManager suite

## Installing NetApp VSC 6.0

To install the VSC 6.0 software, complete the following steps:

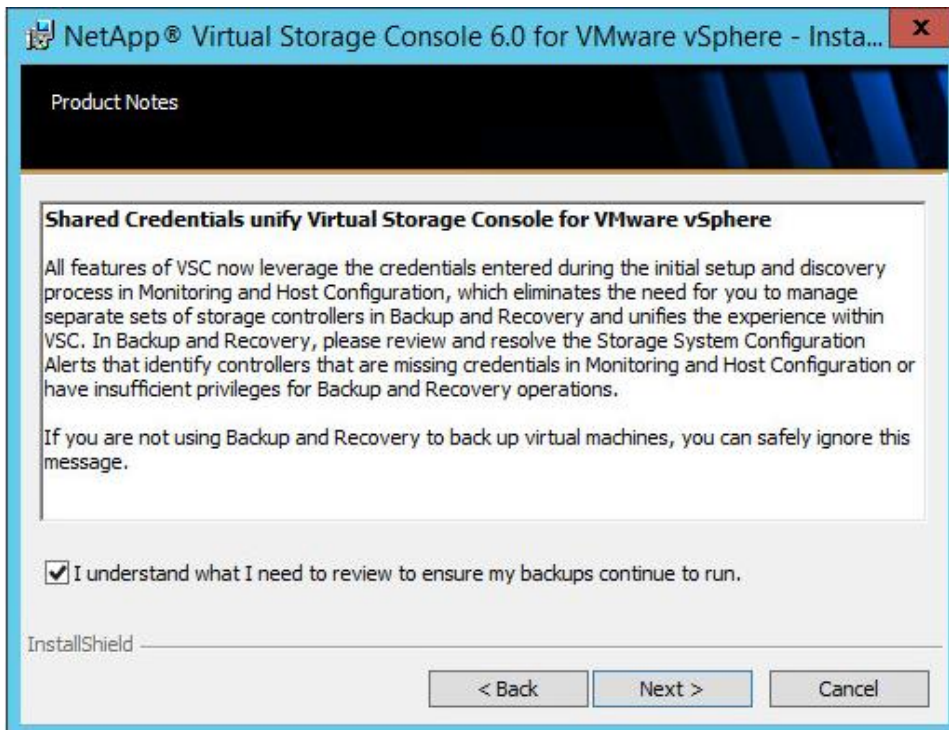
1. Log into the vCenter server using the vSphere Web Client.
2. On the right pane, click VMs and Templates.
3. Select the `windows_2012_r2_template` and right-click it.
4. Select New VM from this Template.
5. Provide a name for the VSC VM, select `FlexPod_Express_DC` as the location for the VM. Click Next.
6. Select the `FlexPod_Express_Management` cluster, click Next.
7. Select `infra_datastore_1`, click Next.
8. Click Next.
9. Review the VM settings and click Finish.
10. Select the newly created VM and select Power on the virtual machine on the right pane.
11. Right-click the VM and select Open Console.
12. Login into the VM as Administrator, assign an IP addresses, and join the machine to the Active Directory domain. Install the current version of Adobe Flash Player on the VM. Install all Windows updates on the VM.
13. Download the x64 version of the [Virtual Storage Console 6.0](#) from the [NetApp Support](#) site.
14. Right-click the file downloaded and select Run As Administrator.
15. On the Installation wizard Welcome page, select the language and click OK.




16. Click Next.



17. Select the checkbox to accept the message, click Next.

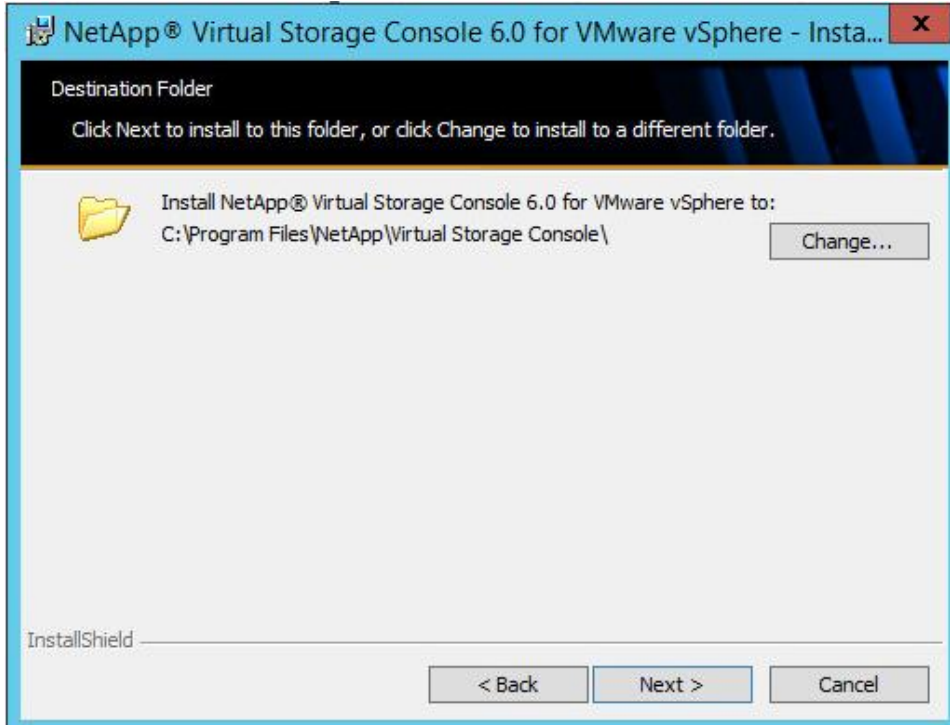


18. Select the backup and recovery capability. Click Next.



**Note:** The backup and recovery capability requires an additional license.

19. Click Next to accept the default installation location.



20. Click Install.

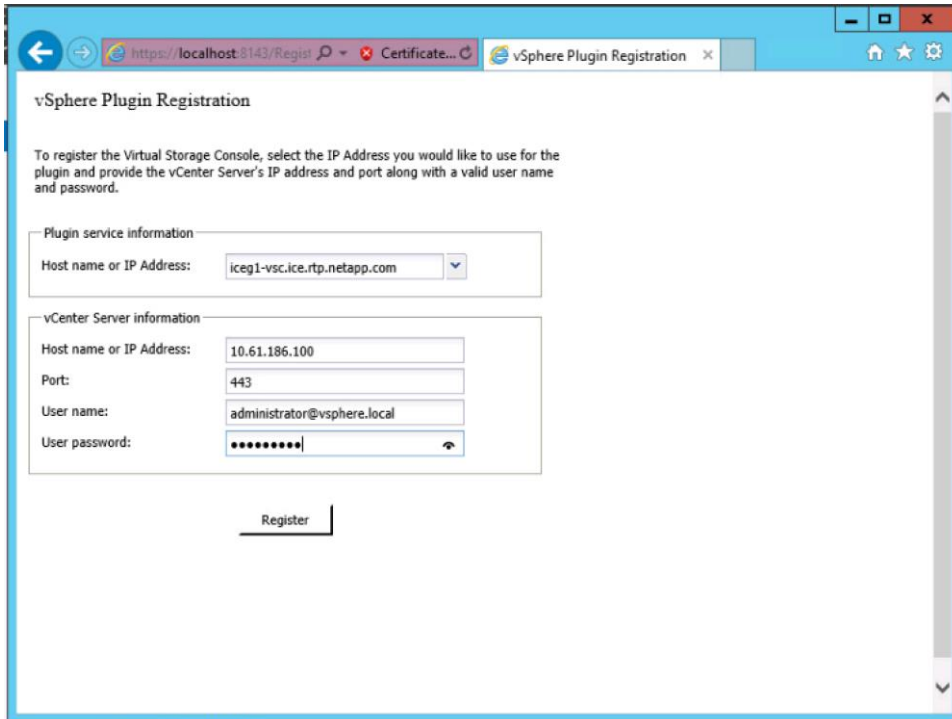


21. Click Finish.

## Registering VSC with vCenter Server

To register the VSC with the vCenter Server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If not, open a browser on the VSC VM and navigate to <https://localhost:8143/Register.html>.
2. Click Continue to this website (not recommended).
3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.
4. In the vCenter Server Information section, enter the host name or IP address, user name (SSO user name), and user password for the vCenter Server. Click Register to complete the registration.



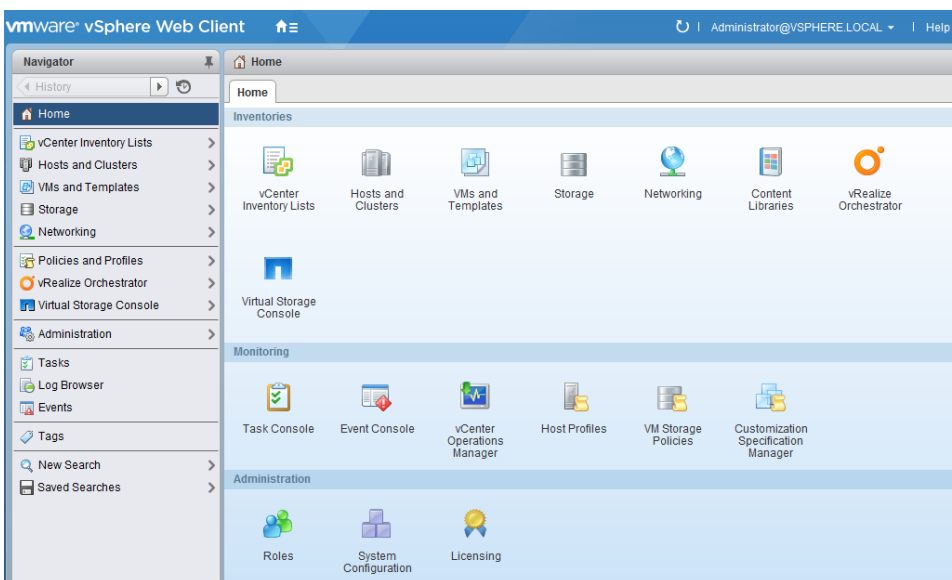
5. Upon successful registration, the storage controllers are discovered automatically.

**Note:** Storage discovery process will take some time to complete.

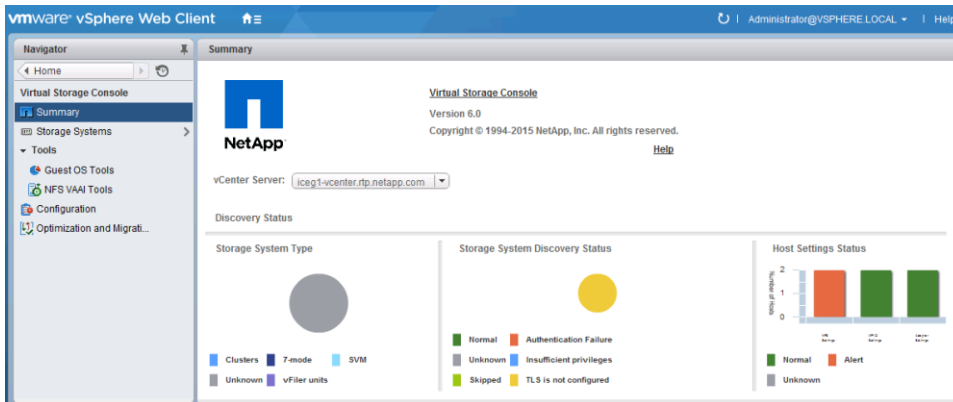
## Updating Credentials for Storage Resources

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

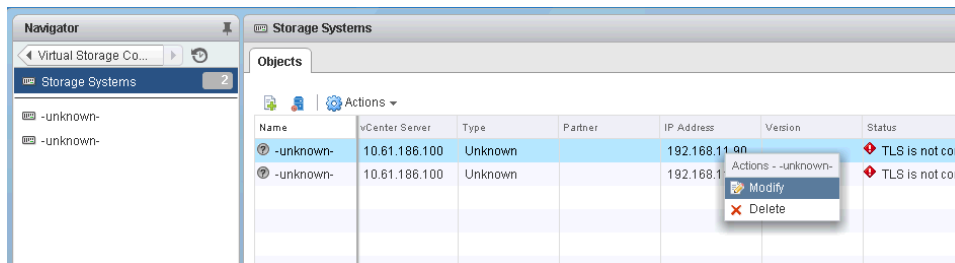
1. Using the vSphere Client, log in to the vCenter Server. If the vSphere Client was previously opened, close it and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.



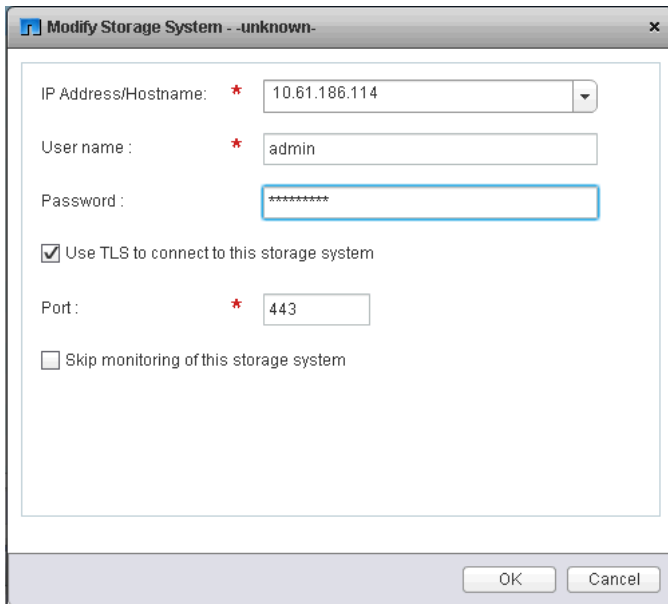




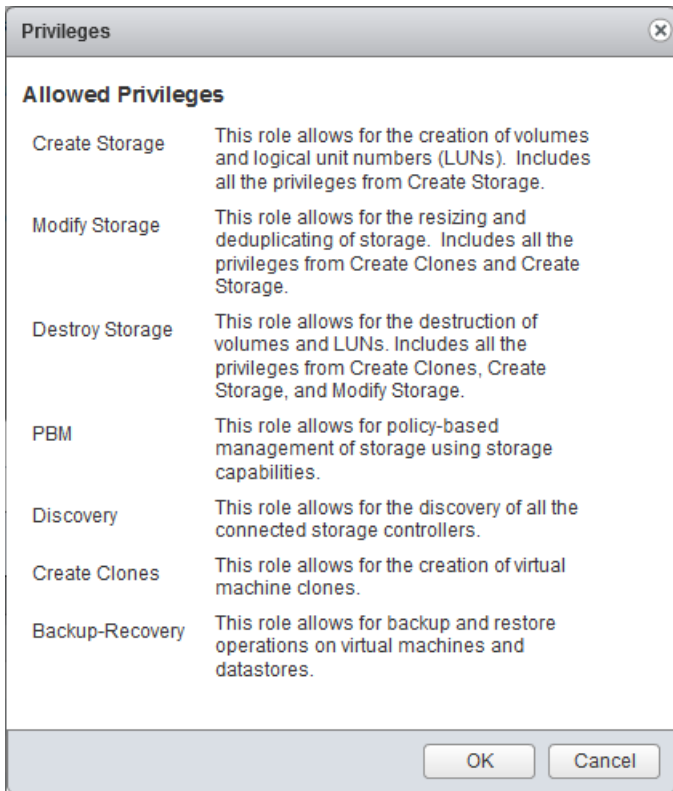
3. In the navigation pane, select Storage systems, if it is not selected by default.



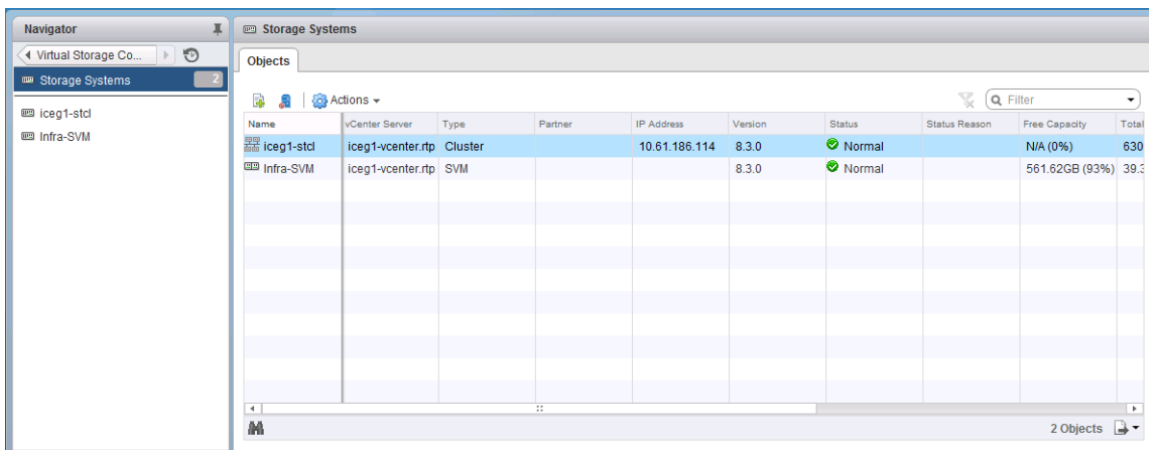
4. Right-click the unknown controller and select Modify.



5. Enter the storage cluster management IP address in the Management IP address field. Enter admin for the user name, and the admin password for password. Make sure that Use TLS is selected. Click OK.



6. Click OK to accept the controller privileges.
7. Refresh the vSphere Web Client to view the updated information.

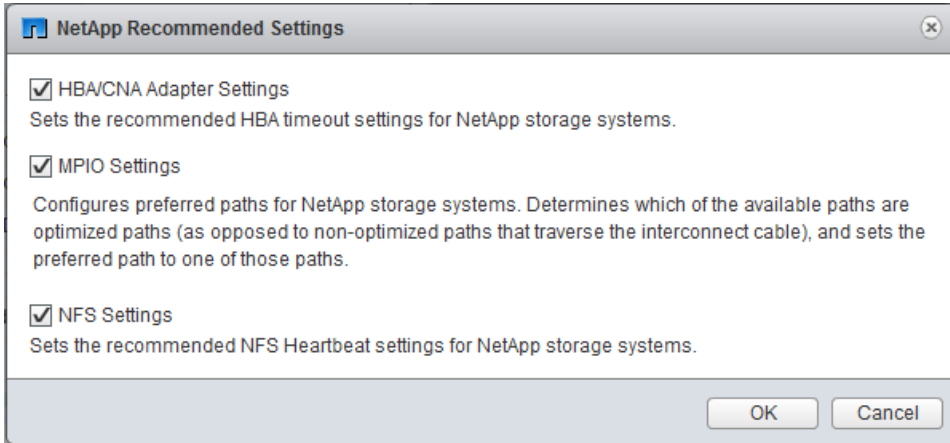


## Optimal Storage Settings for ESXi Hosts

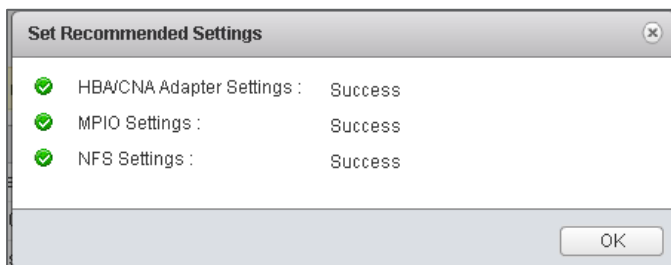
VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the home screen, click Hosts and Clusters. For each ESXi hosts, right-click and select “NetApp VSC > Set Recommended Values” for these hosts.
2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.

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



3. Click OK.



## VSC 5.0 Backup and Recovery

### Prerequisites to use Backup and Recovery Capability

Before you begin using the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must make sure that the storage systems that contain the datastores and virtual machines for which you are creating backups have valid storage credentials.

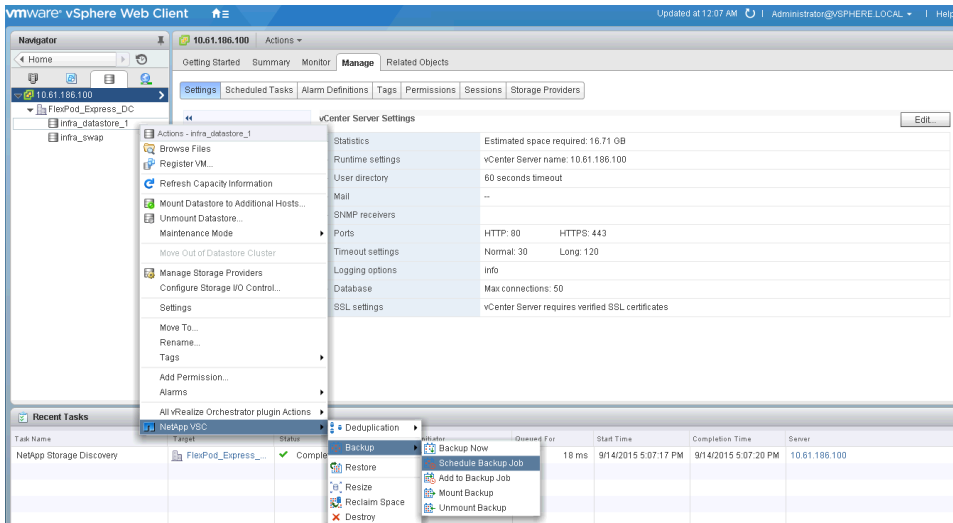
If you are planning to leverage the SnapMirror update option, add all the destination storage systems with valid storage credentials.

### Backup and Recovery Configuration

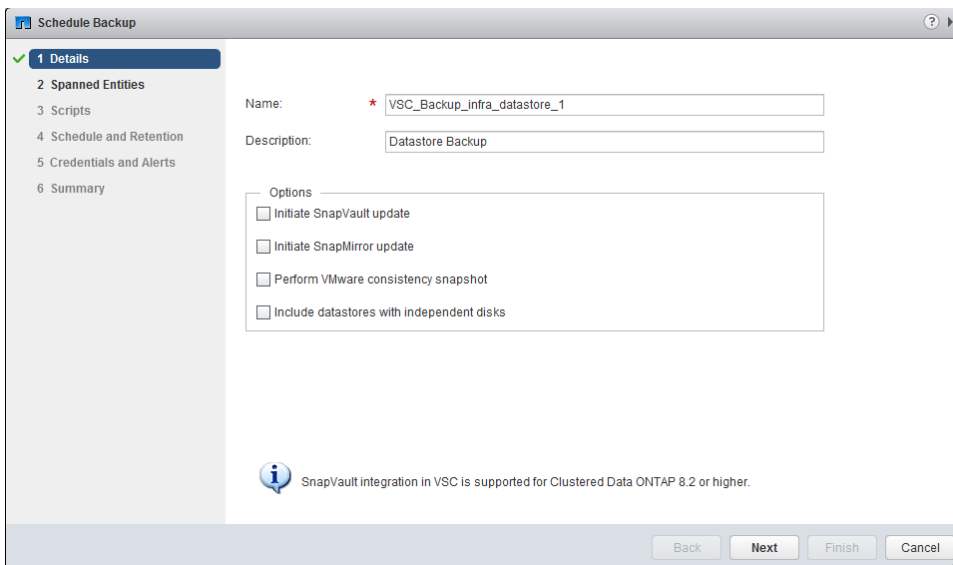
The following steps detail the procedure to configure a backup job for a datastore.

1. From the Home screen, click Storage.
2. Right-click the datastore which you need to backup. Select NetApp VSC > Backup > Schedule Backup Job.

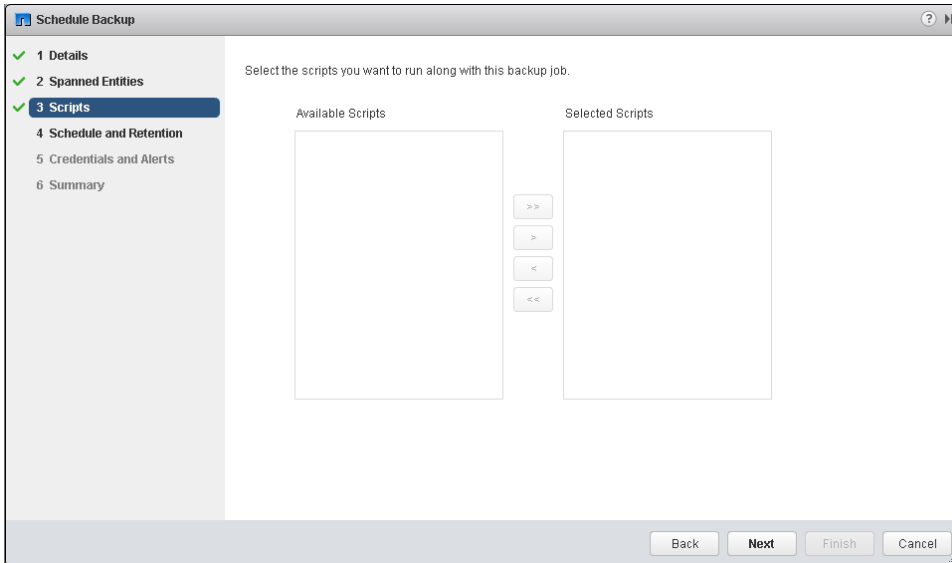
**Note:** If you prefer one-time backup, then choose `Backup Now` instead of `Schedule Backup`.



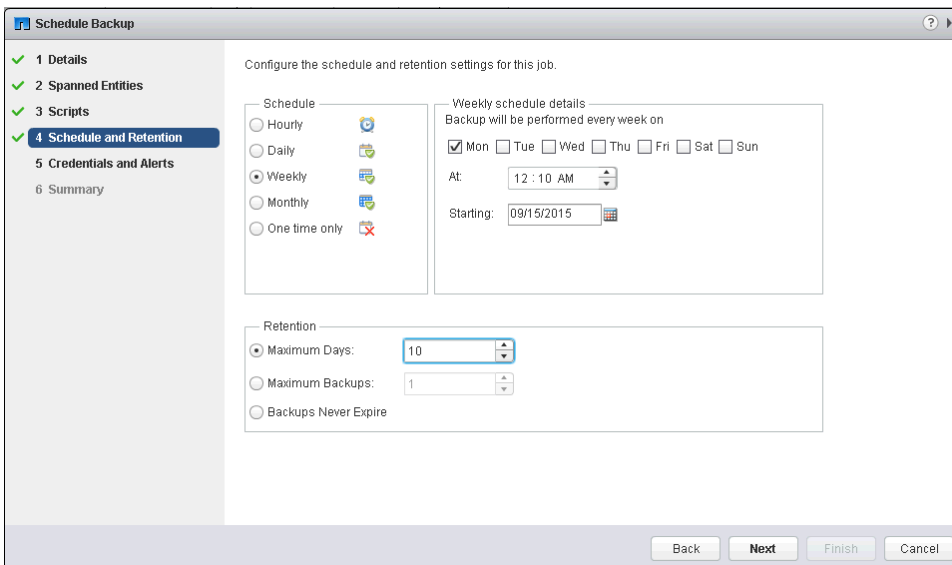
### 3. Type a backup job name and description.



4. Click Next.
5. Click Next.
6. Select one or more backup scripts if available and click Next.



7. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.



8. Use the default vCenter credentials or type the user name and password for the vCenter Server and click Next.
9. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate e-mail addresses. Click Next.

**Schedule Backup**

- 1 Details
- 2 Spanned Entities
- 3 Scripts
- 4 Schedule and Retention
- 5 Credentials and Alerts**
- 6 Summary

**Credentials**

Use default vCenter credentials

Use the following user name and password

Username:

Password:

**Email Alerts**

Notify On:

Source Email Address:

Destination Email Address(s):

SMTP Host:

10. Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.

**Schedule Backup**

- 1 Details
- 2 Spanned Entities
- 3 Scripts
- 4 Schedule and Retention
- 5 Credentials and Alerts
- 6 Summary**

The Backup Job will be created with the following options:

Name: VSC\_Backup\_infra\_datastore\_1

Description: Datastore Backup

Perform this backup: Every week on Mon starting 9/15/2015 at 00:10.

Backup retention: Maximum of 10 days

Email notification will be sent on: Always


Email notification will be sent from: vsc\_backup@example.com

Email notification will be sent to:

Email notification SMTP host: smtp.example.com

Run Job Now

11. Click OK.



12. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by typing the command:

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

13. To delete any existing automatic Snapshot copies that have been created on the volume, type the following command:

```
volume snapshot show -volume infra_datastore_1
volume snapshot delete -volume infra_datastore_1 -snapshot *
Press Y to confirm deletion.
```

## 5 Bill of Materials

This section details the hardware and software components used in validating both the small and medium FlexPod Express configurations included in this document.

### 5.1 Small Configuration

Table 13) Small configuration components.

Part Number	Product Description	Quantity Required
<b>Cisco Components</b>		
<b>Network Switches</b>		
N3K-C3048-FA-L3	Cisco Nexus 3048 Std Airflow (port side exhaust) AC P/S LAN Ent	2
N2200-PAC-400W	N2K/N3K AC Power Supply Std airflow (port side exhaust)	4
CAB-C13-C14-AC	Power cord C13 to C14 (recessed receptacle) 10A	4
N3K-C3048-BAS1K9	Cisco Nexus 3048 Base License	2
N3K-C3048-LAN1K9	Cisco Nexus 3048 LAN Enterprise License	2
N3K-C3048-FAN	Cisco Nexus 3048 Fan Module Port-side Exhaust	2
N3K-C3064-ACC-KIT	Cisco Nexus 3064PQ Accessory Kit	2
N3KUK9-602U2.3	Cisco NX-OS Release 6.0(2)U2(3)	2
CON-SNT-48FAL3	Cisco SMARTNET®8X5XNBD Nexus 3048 Std Airflow AC P/S LAN Ent	2
<b>Cisco UCS Compute</b>		
UCSC-C220-M4L	UCS C220 M4 LFF w/o CPU mem HD PCIe PSU rail kit	2
UCS-CPU-E52640D	2.60 GHz E5-2640 v3/90W 8C/20MB Cache/DDR4 1866MHz	4
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	16
UCSC-MLOM-IRJ45	Intel i350 quad-port MLOM NIC	2
CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	4
UCSC-PSU1-770W	770W AC Hot-Plug Power Supply for 1U C-Series Rack Server	4

Part Number	Product Description	Quantity Required
UCSC-BBLKD-L	3.5-inch HDD Blanking Panel	16
UCSC-HS-C220M4	Heat sink for UCS C220 M4 rack servers	4
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	2
UCS-SD-32G-S	32GB SD Card for UCS servers	4
C1UCS-OPT-OUT	Cisco ONE Data Center Compute Opt Out Option	2
CON-OSP-C220M4L	SN7C-24X7X4OS UCS C220 M4 LFF w/o CPU, mem, HD (Service Duration: 36 months)	2
NetApp Components		
FAS2520A-001-R6	FAS2520 High Availability System	2
X80102A-R6-C	Bezel,FAS2520,R6,-C	1
FAS2520-111-R6-C	FAS2520,12x900GB,10K,-C	1
X1558A-R6-C	Power Cable,In-Cabinet,48-IN,C13-C14,-C	2
SVC-FLEXPOD-SYSTEMS	Systems Used in FlexPod Solution, Attach PN	1
X6560-R6-C	Cable,Ethernet,0.5m RJ45 CAT6,-C	1
X6561-R6	Cable,Ethernet,2m RJ45 CAT6	2
X6557-EN-R6-C	Cbl,SAS Cntrl-Shelf/Shelf-Shelf/HA,0.5m,EN,-C	2
DOC-2520-C	Documents,2520,-C	1
X5518A-R6-C	Kit,FAS2XXX,-C,R6	1
OS-ONTAP-CAP2-1P-C	OS Enable,Per-0.1TB,ONTAP,Perf-Stor,1P,-C	108
SWITCHLESS	2-Node Switchless Cluster	1
SW-2-2520A-SMGR-C	SW-2,SnapManager Suite,2520A,-C	2
SW-2-2520A-SRESTORE-C	SW-2,SnapRestore,2520A,-C	2
SW-2-2520A-FLEXCLN-C	SW-2,FlexClone,2520A,-C	2
SW-2-2520A-ISCSI-C	SW-2,iSCSI,2520A,-C	2
SW-ONTAP8.2.2-CLM	SW,Data ONTAP 8.2.2,Cluster-Mode	2

<sup>1</sup>SupportEdge Premium is required for cooperative support.



## 5.2 Medium Configuration

Table 14) Medium configuration components.

Part Number	Product Description	Quantity Required
<b>Cisco Components</b>		
<b>Network Switches</b>		
N3K-C3048-FA-L3	Cisco Nexus 3048 Std Airflow (port side exhaust) AC P/S LAN Ent	2
N2200-PAC-400W	N2K/N3K AC Power Supply Std airflow (port side exhaust)	4
CAB-C13-C14-AC	Power cord C13 to C14 (recessed receptacle) 10A	4
N3K-C3048-BAS1K9	Cisco Nexus 3048 Base License	2
N3K-C3048-LAN1K9	Cisco Nexus 3048 LAN Enterprise License	2
N3K-C3048-FAN	Cisco Nexus 3048 Fan Module Port-side Exhaust	2
N3K-C3064-ACC-KIT	Cisco Nexus 3064PQ Accessory Kit	2
N3KUK9-602U2.3	Cisco NX-OS Release 6.0(2)U2(3)	2
CON-SNT-48FAL3	Cisco SMARTNET@8X5XNBD Nexus 3048 Std Airflow AC P/S LAN Ent	2
<b>Cisco UCS Compute</b>		
UCSC-C220-M4L	UCS C220 M4 LFF w/o CPU mem HD PCIe PSU rail kit	4
UCS-CPU-E52640D	2.60 GHz E5-2640 v3/90W 8C/20MB Cache/DDR4 1866MHz	8
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	32
UCSC-MLOM-IRJ45	Intel i350 quad-port MLOM NIC	4
CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	8
UCSC-PSU1-770W	770W AC Hot-Plug Power Supply for 1U C-Series Rack Server	8
UCSC-BBLKD-L	3.5-inch HDD Blanking Panel	16
UCS-SD-32G-S	32GB SD Card for UCS servers	8
UCSC-HS-C220M4	Heat sink for UCS C220 M4 rack servers	8
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	4
C1UCS-OPT-OUT	Cisco ONE Data Center Compute Opt Out Option	4
CON-OSP-C220M4L	SNTC-24X7X4OS UCS C220 M4 LFF w/o CPU, mem, HD (Service Duration: 36 months)	4

Part Number	Product Description	Quantity Required
NetApp Components		
FAS2520A-001-R6	FAS2520 High Availability System	2
X80102A-R6-C	Bezel,FAS2520,R6,-C	1
FAS2520-111-R6-C	FAS2520,12x900GB,10K,-C	1
X1558A-R6-C	Power Cable,In-Cabinet,48-IN,C13-C14,-C	2
SVC-FLEXPOD-SYSTEMS	Systems Used in FlexPod Solution, Attach PN	1
X6560-R6-C	Cable,Ethernet,0.5m RJ45 CAT6,-C	1
X6561-R6	Cable,Ethernet,2m RJ45 CAT6	2
X6557-EN-R6-C	Cbl,SAS Cntrl-Shelf/Shelf-Shelf/HA,0.5m,EN,-C	2
DOC-2520-C	Documents,2520,-C	1
X5518A-R6-C	Kit,FAS2XXX,-C,R6	1
OS-ONTAP-CAP2-1P-C	OS Enable,Per-0.1TB,ONTAP,Perf-Stor,1P,-C	108
SWITCHLESS	2-Node Switchless Cluster	1
SW-2-2520A-SMGR-C	SW-2,SnapManager Suite,2520A,-C	2
SW-2-2520A-SRESTORE-C	SW-2,SnapRestore,2520A,-C	2
SW-2-2520A-FLEXCLN-C	SW-2,FlexClone,2520A,-C	2
SW-2-2520A-ISCSI-C	SW-2,iSCSI,2520A,-C	2
SW-ONTAP8.2.2-CLM	SW,Data ONTAP 8.2.2,Cluster-Mode	2

<sup>1</sup>SupportEdge Premium is required for cooperative support.

## 6 Conclusion

FlexPod Express is the optimal shared infrastructure foundation on which to deploy a variety of IT workloads. Cisco and NetApp created a platform that is both flexible and scalable for multiple use cases and applications. One common use case is to deploy VMware vSphere as the virtualization solution, as described in this document. The flexibility and scalability of FlexPod also enable customers to start out with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements.

## References

This report refers to the following documents and resources:

- NetApp FAS2500 Storage  
<http://www.netapp.com/in/products/storage-systems/fas2500/>
- Cisco UCS C-Series Rack Servers  
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html>

- VMware vSphere  
<http://www.vmware.com/in/products/vsphere>
- Interoperability Matrix Tools
  - VMware and Cisco UCS  
<http://www.vmware.com/resources/compatibility/search.php>
  - NetApp, Cisco UCS, and VMware  
<http://support.netapp.com/matrix>

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 © 1994–2015 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).

## Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmplist.aspx>.

NVA-0020-DEPLOY-1015