



Technical Report

FlexPod Express with VMware vSphere Implementation Guide

Michael Zimmerman, David Klem, Chris Reno, and Michael Ansel, NetApp
June 2013 | TR-4107

TABLE OF CONTENTS

1	Overview.....	5
2	Audience.....	5
3	Architecture.....	5
3.1	Small Configuration.....	5
3.2	Medium Configuration.....	6
4	Hardware Details.....	7
4.1	Small Configuration.....	7
4.2	Medium Configuration.....	7
5	Software Details.....	8
6	Configuration Guidelines.....	8
7	FlexPod Express Cabling Information.....	8
7.1	Small Configuration Cabling Diagram.....	8
7.2	Small Configuration Cabling Tables.....	9
7.3	Medium Configuration Cabling Diagram.....	10
7.4	Medium Configuration Cabling Tables.....	11
8	Cisco Nexus 3048 Deployment Procedure.....	12
8.1	Initial Setup of the Cisco Nexus 3048 Switches.....	12
8.2	Software Upgrade (Optional).....	13
8.3	Features.....	13
8.4	Global Port-Channel Configuration.....	14
8.5	Global Spanning-Tree Configuration.....	14
8.6	Jumbo Frames.....	14
8.7	VLAN Definitions.....	15
8.8	Access and Management Port Descriptions.....	15
8.9	Server and Storage Management Interface Configuration.....	16
8.10	Virtual Port-Channel Global Configuration.....	16
8.11	Storage Port-Channels.....	17
8.12	Server Connections.....	18
8.13	In-Band Management SVI Configuration.....	19
8.14	Save Configuration.....	19
8.15	Uplink into Existing Network Infrastructure.....	19
9	NetApp FAS Storage Deployment Procedure.....	19

9.1	Controller FAS22xx Series.....	19
9.2	System Configuration Guides	20
9.3	Assign Controller Disk Ownership and Initialize Storage	20
9.4	Run the Setup Process	22
9.5	64-Bit Aggregates	24
9.6	IFGRP LACP.....	25
9.7	VLAN	25
9.8	IP Config	25
9.9	NFSv3.....	25
9.10	Storage Controller Active-Active Configuration.....	26
9.11	Data ONTAP SecureAdmin.....	26
9.12	Secure Shell.....	27
9.13	AutoSupport HTTPS	27
9.14	Security Best Practices	27
9.15	Enable NDMP	27
9.16	Create FlexVol Volumes	27
9.17	NFS Exports.....	28
9.18	Enable CDP	28
10	Cisco Unified Computing System C-Series Server Deployment Procedure	28
10.1	Perform Initial Cisco UCS C-Series Standalone Server CIMC Setup	28
10.2	Configure Cisco UCS C-Series RAID Configuration	30
11	VMware ESXi Deployment Procedure	33
11.1	Log in to the Cisco UCS C-Series Standalone Server CIMC Interface	33
11.2	Set Up the VMware ESXi Installation	33
11.3	Install ESXi.....	34
11.4	Set Up the ESXi Hosts' Management Networking	34
11.5	Download VMware vSphere Client and vSphere Remote Command Line	35
11.6	Log in to the VMware ESXi Hosts Using VMware vSphere Client	35
11.7	Set Up VMkernel Ports and the Virtual Switch	35
11.8	Mount Required Datastores	36
11.9	Move the VM Swap File Location.....	37
12	VMware vCenter 5.0 Deployment Procedure	37
12.1	Build a VMware vCenter VM	37
12.2	Install VMware vCenter Server	39
12.3	vCenter Setup	42

13 NetApp Virtual Storage Console Deployment Procedure	44
13.1 Install VSC 4.1 Software	44
13.2 Register VSC with vCenter Server	46
13.3 Discover and Add Storage Resources	47
13.4 Optimal Storage Settings for ESXi Hosts	50
13.5 Provisioning and Cloning Setup	51
14 Bill of Materials	52
15 Open Management Ecosystem	55
Appendix A: Cloupia Unified Infrastructure Controller Deployment Procedure	55
Import the CUIC VM into vCenter	55
Configure CUIC	55
Add NetApp Storage Controllers to CUIC	56
Add Cisco Nexus Switches to CUIC	56
Add Cisco C-Series Servers to CUIC	57
Appendix B: Cloupia Unified Infrastructure Controller Bill of Materials	57

LIST OF TABLES

Table 1) Small configuration hardware details	7
Table 2) Medium configuration hardware details	7
Table 3) Software details	8
Table 4) Cisco Nexus 3048 switch 1.	9
Table 5) Cisco Nexus 3048 switch 2.	9
Table 6) Cisco Nexus 3048 switch 1.	11
Table 7) Cisco Nexus 3048 switch 2.	11
Table 8) Controller FAS22XX series prerequisites.	19
Table 9) Small configuration components.	52
Table 10) Medium configuration components	53
Table 11) Cloupia components for small configuration	57
Table 12) Cloupia components for medium configuration.	58

LIST OF FIGURES

Figure 1) FlexPod Express small configuration.	6
Figure 2) FlexPod Express medium configuration	7
Figure 3) Small configuration cabling.	9
Figure 4) Medium configuration cabling	10

1 Overview

The small and medium FlexPod[®] Express configurations are low-cost, standardized infrastructure solutions developed to meet the needs of small and midsize businesses. The configurations have been built and tested to deliver a cost-effective, high-value, and best-practice architecture. 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 business demands.

2 Audience

This document describes the architecture and deployment procedures for both the small and medium FlexPod Express configurations. The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy FlexPod Express.

3 Architecture

Both the small and medium FlexPod Express configurations leverage Cisco[®] Unified Computing System[™] (UCS[®]) C-Series servers, Cisco Nexus[®] switches, and NetApp[®] FAS storage. Although FlexPod Express supports an open ecosystem of virtualization and management software solutions, the architecture described in this document specifically includes VMware[®] vSphere[®] virtualization and the Cloupia Unified Infrastructure Controller (CUIC). NetApp strongly recommends virtualization software and infrastructure management software as part of every FlexPod Express deployment. Each configuration leverages the best practices for and between each component to enable a reliable, enterprise-class infrastructure.


3.1 Small Configuration

The small configuration as validated with VMware vSphere includes the following components:

- Cisco Nexus 3048 switches
- Cisco UCS C220 M3 servers
- NetApp FAS2220 storage controllers
- VMware vSphere 5.0 Update 1 virtualization software
- Cloupia infrastructure management software

Figure 1 highlights the physical topology of the small FlexPod Express configuration.

Figure 1) FlexPod Express small configuration.



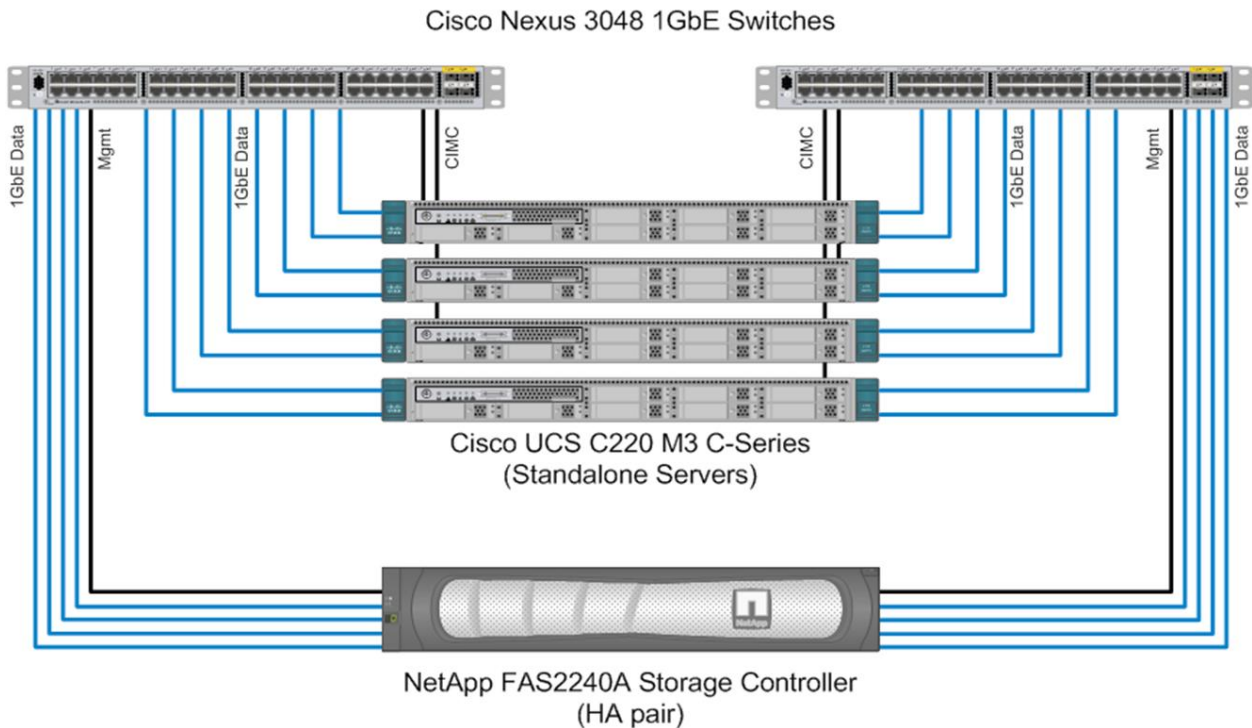
3.2 Medium Configuration

The medium configuration as validated with VMware vSphere includes the following components:

- Cisco Nexus 3048 switches
- Cisco UCS C220 M3 servers
- NetApp FAS2240 storage controllers
- VMware vSphere 5.0 Update 1 virtualization software
- Cloupia infrastructure management software

Figure 2 highlights the physical topology of the medium FlexPod Express configuration.

Figure 2) FlexPod Express medium configuration.



4 Hardware Details

4.1 Small Configuration

Table 1) Small configuration hardware details.

Layer	Component	Quantity
Compute	Cisco UCS C-Series C220 M3 servers (standalone)	2
Network	Cisco Nexus 3048 switches	2
Storage	NetApp FAS2220A (HA pair) (w/ Qty. 12 x 600GB 10K SAS HDDs)	1

4.2 Medium Configuration

Table 2) Medium configuration hardware details.

Layer	Component	Quantity
Compute	Cisco UCS C-Series C220 M3 servers (standalone)	4
Network	Cisco Nexus 3048 switches	2
Storage	NetApp FAS2240A (HA pair) (w/ Qty. 24 x 600GB 10K SAS HDDs)	1

5 Software Details

It is important to note the software versions used in this document. Table 3 details the software versions used throughout this document.

Table 3) Software details.

Layer	Component	Version or Release	Details
Compute	Cisco UCS C-Series C220 M2 standalone servers	1.4(6d)	CIMC software
Network	Cisco Nexus 3048 GbE switches	5.0(3)U4(1)	NX-OS software
Storage (Small Configuration)	NetApp FAS2220A	8.1.1 operating in 7-Mode	NetApp Data ONTAP® software
Storage (Medium Configuration)	NetApp FAS2240A	8.1.1 operating in 7-Mode	Data ONTAP software
Software	VMware vSphere	5.0 Update 1	Virtualization hypervisor suite
	NetApp Virtual Storage Console (VSC)	4.0	NetApp plug-in for VMware vCenter™
	CUIC	3.3	Orchestration and Management Software

6 Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for an FlexPod Express unit. Therefore, reference is made to which component is being configured with each step, either 1 or 2. For example, Controller 1 and Controller 2 are used to identify the two NetApp storage controllers that are provisioned; Switch 1 and Switch 2 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. Finally, to indicate that you should include information pertinent to your environment in a given step, <<var_text>> appears as part of the command structure. See the following example for the `vlan create` command.

```
controller1>vlan create vif0 <<var_mgmt_vlan>>
```


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.

7 FlexPod Express Cabling Information

7.1 Small Configuration Cabling Diagram

Each port used on each component in the small configuration is designated with a box and an associated number. Port connections are defined by matching numbers. For example, Cisco Nexus 3048 Switch 1 port Eth1/1 is labeled with a “1” and is connected to NetApp FAS2240 Storage Controller 1 port e0a, which is also labeled with a “1.”

Figure 3) Small configuration cabling.



7.2 Small Configuration Cabling Tables

Table 4) Cisco Nexus 3048 switch 1.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 3048 Switch 1	Eth1/1	NetApp FAS2220 Storage Controller 1	e0a
	Eth1/2	NetApp FAS2220 Storage Controller 2	e0a
	Eth1/3	NetApp FAS2220 Storage Controller 1	e0c
	Eth1/4	NetApp FAS2220 Storage Controller 2	e0c
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	e1a
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	e1b
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 2	e1a
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	e1b
	Eth1/25	Cisco Nexus 3048 Switch 2	Eth1/25
	Eth1/26	Cisco Nexus 3048 Switch 2	Eth1/26
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 1	Management Port
	Eth1/39	NetApp FAS2220 Storage Controller 1	Management Port


Table 5) Cisco Nexus 3048 switch 2.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 3048	Eth1/1	NetApp FAS2220 Storage Controller 1	e0b

Local Device	Local Port	Remote Device	Remote Port
Switch 2	Eth1/2	NetApp FAS2220 Storage Controller 2	e0b
	Eth1/3	NetApp FAS2220 Storage Controller 1	e0d
	Eth1/4	NetApp FAS2220 Storage Controller 2	e0d
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	e1c
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	e1d
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 2	e1c
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	e1d
	Eth1/25	Cisco Nexus 3048 Switch 1	Eth1/25
	Eth1/26	Cisco Nexus 3048 Switch 1	Eth1/26
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 2	Management Port
	Eth1/39	NetApp FAS2220 Storage Controller 2	Management Port

7.3 Medium Configuration Cabling Diagram

Figure 4) Medium configuration cabling.



7.4 Medium Configuration Cabling Tables

Table 6) Cisco Nexus 3048 switch 1.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 3048 Switch 1	Eth1/1	NetApp FAS2220 Storage Controller 1	e0a
	Eth1/2	NetApp FAS2220 Storage Controller 2	e0a
	Eth1/3	NetApp FAS2220 Storage Controller 1	e0c
	Eth1/4	NetApp FAS2220 Storage Controller 2	e0c
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	e1a
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	e1b
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 2	e1a
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	e1b
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 3	e1a
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 3	e1b
	Eth1/19	Cisco UCS C220 C-Series Standalone Server 1	e1a
	Eth1/20	Cisco UCS C220 C-Series Standalone Server 1	e1b
	Eth1/25	Cisco Nexus 3048 Switch 2	Eth1/25
	Eth1/26	Cisco Nexus 3048 Switch 2	Eth1/26
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 1	Management Port
	Eth1/38	Cisco UCS C220 C-Series Standalone Server 3	Management Port
	Eth1/39	NetApp FAS2220 Storage Controller 1	Management Port

Table 7) Cisco Nexus 3048 switch 2.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 3048 Switch 2	Eth1/1	NetApp FAS2220 Storage Controller 1	e0b
	Eth1/2	NetApp FAS2220 Storage Controller 2	e0b
	Eth1/3	NetApp FAS2220 Storage Controller 1	e0d
	Eth1/4	NetApp FAS2220 Storage Controller 2	e0d
	Eth1/13	Cisco UCS C220 C-Series Standalone Server 1	e1c
	Eth1/14	Cisco UCS C220 C-Series Standalone Server 1	e1d
	Eth1/15	Cisco UCS C220 C-Series Standalone Server 2	e1c
	Eth1/16	Cisco UCS C220 C-Series Standalone Server 2	e1d
	Eth1/17	Cisco UCS C220 C-Series Standalone Server 3	e1c
	Eth1/18	Cisco UCS C220 C-Series Standalone Server 3	e1d

Local Device	Local Port	Remote Device	Remote Port
	Eth1/19	Cisco UCS C220 C-Series Standalone Server 1	e1c
	Eth1/20	Cisco UCS C220 C-Series Standalone Server 1	e1d
	Eth1/25	Cisco Nexus 3048 Switch 1	Eth1/25
	Eth1/26	Cisco Nexus 3048 Switch 1	Eth1/26
	Eth1/37	Cisco UCS C220 C-Series Standalone Server 2	Management Port
	Eth1/38	Cisco UCS C220 C-Series Standalone Server 4	Management Port
	Eth1/39	NetApp FAS2220 Storage Controller 2	Management Port

8 Cisco Nexus 3048 Deployment Procedure

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

8.1 Initial Setup of the Cisco Nexus 3048 Switches

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

The first major decision involves the configuration of the management network for the switches themselves. For FlexPod Express, there are two main options for configuring the mgmt0 interfaces. The first involves configuring and cabling the mgmt0 interfaces into an already existing out-of-band network. In this instance, in which a management network already exists, all that is needed are valid IP addresses and the netmask configuration for this network and a connection from the mgmt0 interfaces into 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 may be configured on each mgmt0 interface as long as they are on 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, but it will be used for the virtual port-channel (vPC) peer keep-alive traffic. To enable SSH management access to the switch, the configuration of the in-band interface-vlan IP address on a switched virtual interface (SVI) is addressed later in this guide.

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

Switch 1 and 2

```

Abort Power On Auto Provisioning and continue with normal setup ?(yes/no) [n]: yes

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

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

---- Basic System Configuration Dialog ----

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

```

```

Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_switch_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
  Mgmt0 IPv4 address : <<var_mgmt0_ip_address>>
  Mgmt0 IPv4 netmask : <<var_mgmt0_netmask>>
Configure the default gateway for mgmt? (yes/no) [y]: n
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
  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]:
Configure CoPP System Policy Profile ( default / 12 / 13 ) [default]:

The following configuration will be applied:
  switchname <<var_switch_hostname>>
interface mgmt0
ip address <<var_mgmt0_ip_address>><<var_mgmt0_netmask>>
no shutdown
  no telnet server enable
  ssh key rsa 1024 force
  ssh server enable
  policy-map type control-plane copp-system-policy ( default )

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

```

8.2 Software Upgrade (Optional)

NetApp recommends performing any required software upgrades on the switch at this point in the configuration. Download and install the latest available NX-OS software for the Cisco Nexus 3048 switch from the Cisco software download site. There are multiple methods to transfer both the kickstart and system images for NX-OS to the switch. The most straightforward procedure leverages the on-board USB port on the switch. Download the NX-OS kickstart and system files to a USB drive and plug the USB drive into the external USB port on the Nexus 3048 switch.

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

Switch 1 and 2

```

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

```

2. The switch will then install the updated NX-OS files and reboot.

8.3 Features

Certain advanced features need to be enabled within NX-OS to allow additional configuration options. The interface-vlan feature is only required if using the back-to-back mgmt0 option described in Section 8.1. This will allow an IP address to be assigned to the interface VLAN (SVI), which enables in-band management communication to the switch such as SSH.

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

Switch 1 and 2

```

feature interface-vlan
feature lacp
feature vpc

```

8.4 Global Port-Channel Configuration

The default port-channel load balancing hash uses the source and destination IP to determine the load-balancing algorithm across the interfaces in the port-channel. Better distribution across the members of the port-channels can be achieved by providing more inputs to the hash algorithm beyond the source and destination IP. For this reason, NetApp highly recommends adding the source and destination TCP port to the hash algorithm.

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

Switch 1 and 2

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

8.5 Global Spanning-Tree Configuration

The Cisco Nexus platform leverages a new protection feature called bridge assurance. Bridge assurance helps to 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 a few states depending on the platform, including “network” and “edge.”

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

This mode will force the network administrator to visit the configuration of each port and can help reveal the most common configuration errors such as nonidentified edge ports or bridge assurance not enabled on a neighbor. Also, it is safer to have spanning tree block many ports than not enough, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding additional servers, storage, or uplink switches, especially if they do not support bridge assurance. In those cases, you might be required to change the port type for the ports to become active.

BPDU guard is enabled on edge ports by default as another layer of protection. This feature will shut down the port if BPDUs from another switch are seen on this interface to prevent loops in the network.

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.

Switch 1 and 2

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

8.6 Jumbo Frames

Jumbo frames should be configured throughout the network to allow any applications or operating systems to transmit these larger frames without fragmentation. Note that both endpoints and all interfaces between the endpoints (L2 and L3) must support and be configured for jumbo frames to realize 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.

Switch 1 and 2

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

```
service-policy type network-qos jumbo
```

8.7 VLAN Definitions

Before configuring individual ports with different VLANs, those L2 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 describe the L2 VLANs.

Switch 1 and 2

```
vlan <<var_nfs_vlan_id>>
  name NFS-VLAN
vlan <<var_vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<var_vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<var_mgmt_vlan_id>>
  name MGMT-VLAN
```

8.8 Access and Management Port Descriptions

Similar to assigning names to the L2 VLAN, setting proper descriptions on all of the interfaces can help with both provisioning and troubleshooting.

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

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

Switch 1

```
int eth1/1
  description FAS-1:e0a
int eth1/2
  description FAS-2:e0a
int eth1/3
  description FAS-1:e0c
int eth1/4
  description FAS-2:e0c
int eth1/13
  description Server-1:port1
int eth1/14
  description Server-1:port2
int eth1/15
  description Server-2:port1
int eth1/16
  description Server-2:port2
int eth1/17
  description Server-3:port1
int eth1/18
  description Server-3:port2
int eth1/19
  description Server-4:port1
int eth1/20
  description Server-4:port2
int eth1/25
  description vPC peer-link SwB:1/25
int eth1/26
  description vPC peer-link SwB:1/26
int eth1/37
  description Server-1:mgmt
int eth1/38
  description Server-3:mgmt
```

```
int eth1/39
  description FAS-1:mgmt
```

Switch 2

```
int eth1/1
  description FAS-1:e0b
int eth1/2
  description FAS-2:e0b
int eth1/3
  description FAS-1:e0d
int eth1/4
  description FAS-2:e0d
int eth1/13
  description Server-1:port3
int eth1/14
  description Server-1:port4
int eth1/15
  description Server-2:port3
int eth1/16
  description Server-2:port4
int eth1/17
  description Server-3:port3
int eth1/18
  description Server-3:port4
int eth1/19
  description Server-4:port3
int eth1/20
  description Server-4:port4
int eth1/25
  description vPC peer-link:1/25
int eth1/26
  description vPC peer-link:1/26
int eth1/37
  description Server-2:mgmt
int eth1/38
  description Server-4:mgmt
int eth1/39
  description FAS-2:mgmt
```

8.9 Server and Storage Management Interface Configuration

The management interfaces for both the server and storage typically utilize only a single VLAN. Because of this, the management interface ports are configured as access ports. Define both the management VLAN for each, as well as 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.

Switch 1 and 2

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

8.10 Virtual Port-Channel Global Configuration

The vPC feature requires an initial setup between the two Nexus switches to function properly. If using the back-to-back mgmt0 configuration, use the addresses defined on the interfaces, and verify that they can communicate by using the `ping <<var_mgmt0_ip_address>> vrf management` command.

From configuration mode (`config t`), type the following commands to configure the vPC global configuration for Switch 1.

Switch 1

```
vpc domain 1
  role priority 10
  peer-keepalive destination <<var_mgmt0_ip_address[of switch2]>> source
<<var_mgmt0_ip_address[of switch1]>>

int eth1/25-26
  channel-group 10 mode active

int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vmtraffic_vlan_id>>, <<var_mgmt_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
```

From configuration mode (`config t`), type the following commands to configure the vPC global configuration for Switch 2.

Switch 2

```
vpc domain 1
  role priority 10
  peer-keepalive destination <<var_mgmt0_ip_address[of switch1]>> source
<<var_mgmt0_ip_address[of switch2]>>

int eth1/25-26
  channel-group 10 mode active

int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vmtraffic_vlan_id>>, <<var_mgmt_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
```

8.11 Storage Port-Channels

The NetApp storage controllers allow an active-active connection to the network using LACP. Using LACP is preferred because it adds additional negotiation between the switches in addition to logging. Because the network is set up for vPC, this allows us 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 IFGRP.

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

Switch 1 and 2, FAS-1 Configuration

```
int eth1/1,eth1/3
  channel-group 11 mode active

int Po11
  description vPC to FAS-1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,<<var_nfs_vlan_id>>,<<var_vmtraffic_vlan_id>>
  spanning-tree port type edge trunk
```

```
vpc 11
no shut
```

Switch 1 and 2, FAS-2 Configuration

```
int eth1/2,eth1/4
  channel-group 12 mode active

int Po12
  description vPC to FAS-2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1, <<var_nfs_vlan_id>>,<<var_vmtraffic_vlan_id>>
  spanning-tree port type edge trunk
  vpc 12
  no shut
```

8.12 Server Connections

The 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 out across multiple switches enables the server to survive even with a complete switch failure.

For the small configuration, you only need to configure **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.

Switch 1 and 2, Server-1 Configuration

```
int eth1/13-14
  switchport mode trunk
  switchport trunk allowed vlan 1,<<var_vlan_nfs>>,<<var_vlan_vmotion>>, <<var_vlan_vmdata>>,
<<var_vlan_mgmt>>
  spanning-tree port type edge trunk
  no shut
```

Switch 1 and 2, Server-2 Configuration

```
int eth1/15-16
  switchport mode trunk
  switchport trunk allowed vlan 1,<<var_vlan_nfs>>, <<var_vlan_vmotion>>, <<var_vlan_vmdata>>,
<<var_vlan_mgmt>>
  spanning-tree port type edge trunk
  no shut
```

Note: Server-3 and Server-4 configurations, below, are required only for FlexPod Express medium configurations.

Switch 1 and 2, Server-3 Configuration

```
int eth1/17-18
  switchport mode trunk
  switchport trunk allowed vlan 1,<<var_vlan_nfs>>,<<var_vlan_vmotion>>, <<var_vlan_vmdata>>,
<<var_vlan_mgmt>>
  spanning-tree port type edge trunk
  no shut
```

Switch 1 and 2, Server-4 Configuration

```
int eth1/19-20
  switchport mode trunk
```

```
switchport trunk allowed vlan 1,<<var_vlan_nfs>><<var_vlan_vmotion>>, <<var_vlan_vmdata>>,
<<var_vlan_mgmt>>
spanning-tree port type edge trunk
no shut
```

8.13 In-Band Management SVI Configuration

In-band management through SSH in the FlexPod Express environment is handled by an SVI. To configure the in-band management on each of the switches, an IP address must be configured on the interface-vlan and a default gateway must be set up.

From configuration mode (`config t`), type the following commands to configure the SVI L3 interface for management purposes.

Switch 1 and 2

```
int Vlan <<var_mgmt_vlan_id>>
  ip address <<var_inband_mgmt_ip_address>>/<<var_inband_mgmt_netmask>>
  no shut

ip route 0.0.0.0/0 <<var_inband_mgmt_net_gateway>>
```

8.14 Save Configuration

Save the configuration on both switches for configuration persistence.

Switch 1 and 2

```
copy run start
```

8.15 Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod Express environment. If an existing Cisco Nexus environment is present, NetApp recommends using virtual port-channels to uplink the Cisco Nexus 3048 switches included in the FlexPod Express environment into the infrastructure. Make sure to type `copy run start` to save the configuration on each switch after the configuration is completed.

9 NetApp FAS Storage Deployment Procedure

9.1 Controller FAS22xx Series

Table 8) Controller FAS22XX series prerequisites.

Requirement	Reference	Comments
Physical site where storage system will be installed must be ready	Site Requirements Guide	Refer to the “Site Preparation” section.
Storage system connectivity requirements	Site Requirements Guide	Refer to the “System Connectivity Requirements” section.
Storage system general power requirements	Site Requirements Guide	Refer to the “Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements” section.
Storage system model-specific requirements	Site Requirements Guide	Refer to the “FAS22xx Series Systems” section.

9.2 System Configuration Guides

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

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

Controller 1 and 2

Follow the installation procedures for the controllers from the [Installation and Setup Instructions for FAS2220/FAS2240-2 System](#) guide at the NetApp [Support](#) site.

9.3 Assign Controller Disk Ownership and Initialize Storage

These steps provide details for assigning disk ownership and disk initialization and verification.

Controller 1

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

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

2. If the system is at the LOADER prompt, enter the `autoboot` command to boot Data ONTAP.
3. During system boot, press Ctrl-C when prompted for the Boot Menu.

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

Note: If 8.1.1 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.1 is the version being booted, go to step 14, Maintenance mode boot.

4. To install new software, first select option 7.
5. Answer `yes` for performing a nondisruptive upgrade.
6. Select `e0M` for the network port you want to use for the download.
7. Select `yes` to reboot now.
8. Enter the IP address: `<<var_contoller1_e0m_ip>>`, netmask: `<<var_contoller1_e0m_mask>>`, and default gateway: `<<var_contoller1_e0m_gateway>>` for `e0M` in their respective places.
9. Enter the URL: `<<var_url_boot_software>>` where the software can be found.

Note: This Web server must be pingable.

10. Press Enter for the user name, indicating no user name.
11. Select `yes` to set the newly installed software as the default to be used for subsequent reboots.
12. Select `yes` to reboot the node.
13. Press Ctrl-C when you see `Press Ctrl-C for Boot Menu`.
14. To enter Maintenance mode boot, select option 5.
15. When prompted `"Continue to Boot?"` answer `yes`.

16. To verify the HA status of your environment, use the command `ha-config show`.

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

```
ha-config modify controller ha
ha-config modify chassis ha
```

17. Use the `disk show -n` command to view how many disks are unowned.

Note: The remaining number of disks should be shown.

18. Use the `disk assign -n<<var##_of_disks>>` command to assign disks to controller 1.

Note: For the small FlexPod Express configuration, <<var##_of_disks>> should equal 9 for controller 1.

Note: For the medium FlexPod Express configuration, <<var##_of_disks>> should equal 21 for controller 1.

19. Reboot the controller by using the `halt` command.

20. At the LOADER prompt type `autoboot`.

21. Press Ctrl-C for the Boot Menu when prompted.

22. Select option 4) `Clean configuration and initialize all disks`.

23. Answer `yes` to zero disks, reset config, and install a new file system.

24. Enter `yes` to erase all the data on the disks.

Note: The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After the initialization is complete, the storage system launches setup.

Controller 2

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

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

2. If the system is at the LOADER prompt, enter the `autoboot` command to boot Data ONTAP.

3. During system boot, press Ctrl-C when prompted for the Boot Menu.

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

Note: If 8.1.1 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.1 is the version being booted, then proceed with step 14, Maintenance mode boot.

4. To install the new software first, select option 7.

5. Answer `yes` for performing a nondisruptive upgrade.

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

7. Select `yes` to reboot now.

8. Enter the IP address: <<var_controller2_e0m_ip>>, netmask: <<var_controller2_e0m_mask>>, and default gateway: <<var_controller2_e0m_gateway>> for `e0M` in their respective places.

9. Enter the URL: <<var_url_boot_software>> where the software can be found.

Note: This Web server must be pingable.

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

11. Select `yes` to set the newly installed software as the default to be used for subsequent reboots.
12. Select `yes` to reboot the node.
13. Press Ctrl-C when you see “Press Ctrl-C for Boot Menu.”
14. To enter Maintenance mode boot, select option 5.
15. When prompted “Continue to Boot?” answer `yes`.
16. To verify the HA status of your environment, use the `ha-config show` command.

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

```
ha-config modify controller ha
ha-config modify chassis ha
```

17. Use the `disk show -a` command to view the number of unowned disks.

Note: The remaining disks should be shown.

18. Use the `disk assign -n<<var_#_of_disks>>` command to assign disks to controller 2.

Note: For both the small and medium FlexPod Express configurations, `<<var_#_of_disks>>` should equal 3 for controller 2.

19. Reboot the controller using the `halt` command.
20. At the LOADER prompt type `autoboot`.
21. Press Ctrl-C for the Boot Menu when prompted.
22. Select option 4) `Clean configuration and initialize all disks`.
23. Answer `yes` to zero disks, reset configuration, and install a new file system.
24. Enter `yes` to erase all the data on the disks.

Note: The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After the initialization is complete, the storage system launches setup.

9.4 Run the Setup Process

When Data ONTAP is installed on your new storage system, the following files are not populated:

- `/etc/rc`
- `/etc/exports`
- `/etc/hosts`
- `/etc/hosts.equiv`

Controller 1

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
2. Enter the following information:

```
Please enter the new hostname []:<<var_controller1>>

Do you want to enable IPv6? [n]:

Do you want to configure interface groups? [n]:
Please enter the IP address for Network Interface e0a []:
```

Note: Press Enter to accept the blank IP address.


```

Should interface e0a take over a partner IP address during failover? [n]:
Please enter the IP address for the Network Interface e0b []:
Should interface e0b take over a partner IP address during failover? [n]:
Please enter the IP address for the Network Interface e0c []:
Should interface e1a take over a partner IP address during failover? [n]:
Please enter the IP address for the Network Interface e0d []:
Should interface e1b take over a partner IP address during failover? [n]:

Please enter the IP address for Network Interface e0M []: <<var_controller2_e0m_ip>>
Please enter the netmask for the Network Interface e0M [255.255.255.0]:
<<var_controller2_e0m_mask>>

Should interface e0M take over a partner IP address during failover? [n]: y
Please enter the IPv4 address or interface name to be taken over by e0M []: e0M

```

4. Enter the following information:

```

Would you like to continue setup through the Web interface? [n]:

Please enter the name or IP address of the IPv4 default gateway: <<var_controller2_e0m_gateway>>

The administration host is given root access to the storage system's / etc files for system
administration. To allow /etc root access to all NFS clients enter RETURN below.
Please enter the name or IP address for administrative host: <<var_adminhost_ip>>

Please enter timezone [GMT]: <<var_timezone>>

```

Note: Example time zone: America/New_York

```

Where is the filer located? <<var_location>>
Enter the root directory for HTTP files [home/http]:
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:

```

Note: Optionally, enter up to three name server IP addresses.

```

Do you want to run NIS client? [n]:
Press the Return key to continue through AutoSupport message
would you like to configure SP LAN interface [y]: enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_sp_ip>>
Please enter the netmask for the SP []: <<var_sp_mask>>
Please enter the IP address for the SP gateway: <<var_sp_gateway>>
Please enter the name or IP address of the mail host [mailhost]: <<var_mailhost>>
Please enter the IP address for <<var_mailhost>> []: <<var_mailhost_ip>>
New password: <<var_admin_passwd>>
Retype new password <<var_admin_passwd>>

```

5. Enter the admin password to log in to controller 2.

9.5 64-Bit Aggregates

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

Controller 1

1. Execute the following command to create a new aggregate:

```
aggr create aggr1 -B 64 <<var_#_of_disks>>
```

Note: For the small FlexPod Express configuration, <<var_#_of_disks>> should equal 5.

Note: For the medium FlexPod Express configuration, <<var_#_of_disks>> should equal 17.

Note: aggr1 is not required on controller 2 because it is set up as an HA pair.

9.6 IFGRP LACP

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

Controller 1 and Controller 2

1. Run the following command on the command line and also add it to the `/etc/rc` file so it is activated upon boot.

```
ifgrp create lacp ifgrp0 -b ip e0a e0b e0c e0d  
wrfile -a /etc/rc "ifgrp create lacp ifgrp0 -b ip e0a e0b e0c e0d"
```

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

9.7 VLAN

Controller 1 and Controller 2

1. Run the following commands to create a VLAN interface for NFS data traffic.

```
vlan create ifgrp0 <<var_nfs_vlan_id>>  
wrfile -a /etc/rc "vlan create ifgrp0 <<var_nfs_vlan_id>>"
```

9.8 IP Config

Controller 1

1. Run the following commands from the command line:

```
ifconfig ifgrp0-<<var_nfs_vlan_id>><<var_controller1_nfs_ip>> netmask  
<<var_controller1_nfs_mask>> mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>  
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_nfs_vlan_id>><<var_controller1_nfs_ip>> netmask  
<<var_controller1_nfs_mask>> mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>"
```

Controller 2

1. Run the following commands from the command line:

```
ifconfig ifgrp0-<<var_nfs_vlan_id>> mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>  
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_nfs_vlan_id>> mtusize 9000 partner ifgrp0-  
<<var_nfs_vlan_id>>"
```

9.9 NFSv3

Controller 1 and Controller 2

1. Add a license for NFS.

```
license add <<var_nfs_license>>
```

2. Set the following recommended options to enable NFS version 3.

```
options nfs.tcp.enable on  
options nfs.udp.enable off  
options nfs.v3.enable on
```

3. Enable NFS.

```
nfs on
```

9.10 Storage Controller Active-Active Configuration

Controller 1 and Controller 2

Enable two storage controllers in an active-active configuration.

1. Enter the cluster license on both nodes.

```
license add <<var_cf_license>>
```

2. Reboot each storage controller.

```
reboot
```

3. Log back in to both controllers

Controller 1

1. Enable failover on controller 1, if it is not already enabled.

```
cf enable
```

9.11 Data ONTAP SecureAdmin

Secure API access to the storage controller must be configured.

Controller 1

1. Issue the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup -f -q ssl t US "<<var_state>>"  
<<var_city>><<var_org>><<var_unit>><<var_controller1_fqdn>><<var_admin_email>><<var_key_length>>
```

Note: The format for this command is `secureadmin setup -q ssl domestic<t/f> country state locality org unit fqdn email [keylen] [days until expires]`.

Note: Parameters that require more than one word should be placed in quotation marks (" ").

After the initialization, the CSR is available in the file `/etc/keymgr/csr/secureadmin_tmp.pem`.

2. Configure and enable SSL and HTTPS for API access using the following options:

```
optionshttpd.access none  
optionshttpd.admin.enable on  
optionshttpd.admin.ssl.enable on  
optionsssl.enable on
```

Controller 2

1. Use the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup -f -q ssl t US "<<var_state>>"  
<<var_city>><<var_org>><<var_unit>><<var_controller2_fqdn>><<var_admin_email>><<var_key_length>>
```

The format for this command is `secureadmin setup -q ssl domestic<t/f> country state locality org unit fqdn email [keylen] [days until expires]`.

Parameters that need more than one word should be placed in quotation marks (" ").

After the initialization, the CSR is available in the file `/etc/keymgr/csr/secureadmin_tmp.pem`.

2. Configure and enable SSL and HTTPS for API access using the following options:

```
options httpd.access none  
options httpd.admin.enable on
```

```
options httpd.admin.ssl.enable on
options ssl.enable on
```

9.12 Secure Shell

SSH must be configured and enabled.

Controller 1 and Controller 2

1. Use the following one-time command to generate host keys:

```
secureadmin disable ssh
secureadmin setup -f -q ssh 768 512 1024
```

2. Use the following options to configure and enable SSH:

```
options ssh.idle.timeout 60
options autologout.telnet.timeout 5
```

9.13 AutoSupport HTTPS

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

Controller 1 and Controller 2

1. Execute the following commands to configure AutoSupport:

```
options autosupport.noteto <<var_admin_email>>
```

9.14 Security Best Practices

Note: Apply the following commands according to local security policies.

Controller 1 and Controller 2

1. Run the following commands to enhance security on the storage controller:

```
options rsh.access none
options webdav.enable off
options security.passwd.rules.maximum 14
options security.passwd.rules.minimum.symbol 1
options security.passwd.lockout.numtries 6
options autologout.console.timeout 5
```

9.15 Enable NDMP

Controller 1 and Controller 2

1. Run the following commands to enable NDMP.

```
options ndmpd.enable on
```

9.16 Create FlexVol Volumes

Controller 1

1. Create two FlexVol® volumes on controller 1 using the following commands:

```
vol create infra_swap -s none aggr1 100g
snap reserve infra_swap 0
snap sched infra_swap 0 0 0

vol create infra_datastore_1 -s none aggr1 500g
```

```
snap reserve infra datastore_1 0
sis on /vol/infra_datastore_1
```

9.17 NFS Exports

Use the following steps to create NFS exports on each controller.

Controller 1

```
exportfs -p
sec=sys,rw=<<var_esxi1_nfs_ip>>:<<var_esxi2_nfs_ip>>:<<var_esxi3_nfs_ip>>:<<var_esxi4_nfs_ip>>,
root=<<var_esxi1_nfs_ip>>:<<var_esxi2_nfs_ip>>:<<var_esxi3_nfs_ip>>:<<var_esxi4_nfs_ip>>,nosuid
/vol/infra_swap

exportfs -p
sec=sys,rw=<<var_esxi1_nfs_ip>>:<<var_esxi2_nfs_ip>>:<<var_esxi3_nfs_ip>>:<<var_esxi4_nfs_ip>>,
root=<<var_esxi1_nfs_ip>>:<<var_esxi2_nfs_ip>>:<<var_esxi3_nfs_ip>>:<<var_esxi4_nfs_ip>>,nosuid
/vol/infra_datastore_1
```

9.18 Enable CDP

Use the following steps to enable CDP on controller 1 and controller 2.

Controller 1 and Controller 2

1. Enable CDP.

```
options cdpd.enable on
```

10 Cisco Unified Computing System C-Series Server Deployment Procedure

This section provides the detailed procedure for configuring a Cisco Unified Computing System C-Series standalone server for use in either small or medium FlexPod Express configurations.

10.1 Perform Initial Cisco UCS C-Series Standalone Server CIMC Setup

These steps describe the setup of the initial Cisco UCS C-Series standalone server.

All Servers

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



Press <F2> to enter setup, <F6> Boot Menu, <F8> CIMC Config, <F12> Network Boot

Bios Version: C220M3.1.4.4c.0.022220121951
Platform ID: C220M3

CIMC IP Address : 10.61.186.94
CIMC MAC Address : 50:3D:E5:9E:32:C2

Processor(s) Intel(R) Xeon(R) CPU E5-2609 0 @ 2.40GHz
Total Memory = 64 GB Effective Memory = 64 GB
Memory Operating Speed 1066 Mhz

92

3. From the CIMC Configuration Utility, set the following options.
 - a. NIC mode:
 - Dedicated ☒
 - b. IPV4 (Basic):
 - DHCP enabled: ☐
 - CIMC IP: <<var_cimc_ip>>
 - Subnet mask: <<var_cimc_mask>>
 - Gateway: <<var_cimc_mask>>
 - c. VLAN (Advanced): Leave this option cleared to disable VLAN tagging.
 - d. NIC redundancy: None
 - e. Factory Defaults: Leave this option cleared.
 - f. Default User (Basic):
 - Default password: <<var_admin_passwd>>
 - Reenter password: <<var_password>>

```

CIMC Configuration Utility  Version 1.5  Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]
Shared LOM:     [ ]
Shared LOM 10G: [ ]
Cisco Card:     [ ]
NIC redundancy
None:           [X]
Active-standby: [ ]
Active-active:  [ ]
IPV4 (Basic)
DHCP enabled:   [ ]
CIMC IP:        10.61.186.94
Subnetmask:     255.255.255.0
Gateway:        10.61.186.1
Factory Defaults
CIMC Factory Default:[ ]
Default User (Basic)
Default password:
Reenter password:
VLAN (Advanced)
VLAN enabled:   [ ]
VLAN ID:        1
Priority:        0

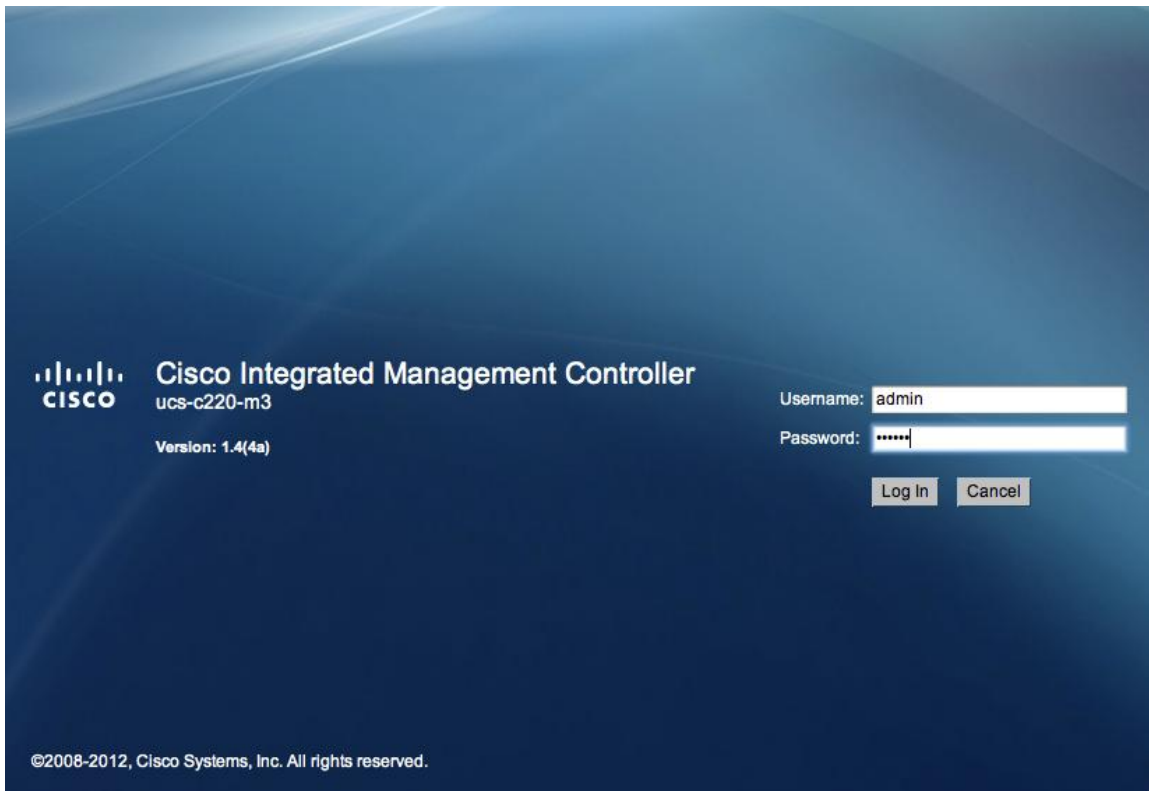
*****
<Up/Down arrow> Select items    <F10> Save    <Space bar> Enable/Disable
<F5> Refresh                    <ESC> Exit

```


4. Press F10 to save the CIMC interface configuration.
5. After the configuration is saved, press Esc to exit.

10.2 Configure Cisco UCS C-Series RAID Configuration


1. Open a Web browser and browse to the CIMC interface IP address.
2. Log in to the CIMC interface using the default user name `admin` and the admin password: `<<var_admin_passwd>>` set in the CIMC interface setup.




3. After you are successfully logged in, click the Server tab and choose Summary. Select Launch KVM Console.
4. The virtual KVM window opens. Select Virtual Media at the top of the window.
5. Click Add Image.
6. Browse to the location of the Server Configuration Utility ISO image and select it. Click Open.
7. Select the Mapped checkbox next to the selected ISO image to map the image to the server.



8. Return to the CIMC interface browser page (do not close the virtual KVM window), click the Server tab, and choose BIOS.
9. Select Configure Boot Order and click OK.
10. Add both the CDROM and HDD options to the Boot Order field. Click Apply.



11. Click the Server tab and select Summary. Select Power Cycle Server.
12. Return to the virtual KVM window. Click the KVM tab at the top of the window.
13. The server should now boot into the Server Configuration Utility.
14. Click the Server Configuration tab in the left pane.
15. Select RAID Configuration.
16. In the upper-right corner, click the Configure button. 
17. From the RAID Level drop-down menu, select Automatic setup with redundancy. Click Create Array.
18. After the RAID configuration completes, close the virtual KVM window.
19. Return to the CIMC interface browser window. Click the Server tab and then select 0. Select Power Off Server.

11 VMware ESXi Deployment Procedure

This section provides detailed procedures for installing VMware ESXi™ 5.1 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 ESXi in such an environment. This procedure highlights using the virtual KVM console and virtual media features within the Cisco UCS C-Series CIMC interface to map remote installation media to each individual server.

11.1 Log in to the Cisco UCS C-Series Standalone Server CIMC Interface

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

All Hosts

1. Navigate to a Web browser and enter the IP address for the Cisco C-Series CIMC interface. This will launch the CIMC GUI application.
2. Log in to the CIMC GUI with the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.

11.2 Set Up the VMware ESXi Installation

This section details the steps required to prepare the server for OS installation.

All Hosts

1. From the virtual KVM Console, select the Virtual Media tab.
2. Select Add Image in the right pane.
3. Browse to the ESXi 5.1 installer ISO image file and click Open.
4. Map the image that you just added by selecting Mapped.
5. To boot the server, select the KVM tab.
6. Select Power On Server in the CIMC interface Summary tab, and then click OK.

11.3 Install ESXi

The following steps describe the installation of VMware ESXi to each host's local RAID drive.

All Hosts

1. On boot, the machine detects the presence of the ESXi installation media.
2. Select the ESXi Installer from the menu that appears.
3. After the installer is finished loading, press Enter to continue with the installation.
4. After reviewing the EULA, accept and continue with the installation by pressing F11.
5. Select the local RAID drive that was set up previously as the installation location for ESXi. 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 ESXi installation image in the Virtual Media tab of the KVM Console so that the server reboots into ESXi and not the installer by clearing the Mapped checkbox.
10. The Virtual Media window might warn you that it is preferable to eject the media from the guest. Because we cannot do this (and the media is read-only), unmap the image anyway by clicking Yes.
11. Back in the KVM tab, press Enter to reboot the server.

11.4 Set Up the ESXi Hosts' Management Networking

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

All Hosts

1. After the server is done rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the user name and the root password previously entered during installation.
3. Select the option to Configure Management Network.
4. Select Network Adapters and press Enter.
5. There should be four ports that show Connected under the Status column. These ports should correspond to ports 1, 2, 3, and 4 of the quad-port Broadcom PCI-e adapter. Select all four ports and press Enter.
6. Select VLAN (optional) and press Enter.
7. Enter `<<var_mgmt_vlan_id>>` and press Enter.
8. From the Configure Management Network menu, configure the IP address of the management interface by selecting the IP Configuration option and pressing Enter.
9. Use the space bar to select the set static IP address and network configuration.
10. Enter the IP address for managing the ESXi host `<<var_esxi_mgmt_ip>>`.
11. Enter the subnet mask for the ESXi host `<<var_esxi_mgmt_mask>>`.
12. Enter the default gateway for the ESXi host `<<var_esxi_mgmt_gateway>>`.
13. Press Enter to accept the changes to the IP configuration.
14. Enter the menu to configure the DNS settings.

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

15. Enter the primary DNS server's IP address <<var_nameserver_ip>>.
16. (Optionally) Enter the secondary DNS server's IP address.
17. Enter the Fully Qualified Domain Name (FQDN) for the first ESXi host <<var_esxi_host_fqdn>>.
18. Press Enter to accept the changes to the DNS configuration.
19. Exit the Configure Management Network submenu by pressing Esc.
20. Confirm the changes made and return to the main menu by pressing Y.
21. Log out of the VMware Console by pressing Esc.

11.5 Download VMware vSphere Client and vSphere Remote Command Line

The following steps provide details for downloading the VMware vSphere 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 ESXi hosts.
2. Download and install both the vSphere Client and the Windows® version of the vSphere Remote Command Line.

11.6 Log in to the VMware ESXi Hosts Using VMware vSphere Client

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

All Hosts

1. Open the recently downloaded VMware vSphere Client and enter the IP address of the host to which you are trying to connect <<var_esxi_mgmt_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

11.7 Set Up VMkernel Ports and the Virtual Switch

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

All Hosts

1. In the vSphere client, select the host from the left pane.
2. Select the Configuration tab.
3. Select the Networking link in the Hardware box.
4. Select the Properties link in the right pane on vSwitch0.
5. Select the vSwitch configuration and click Edit.
6. Under the General tab, change the MTU to 9000.
7. Under the NIC Teaming tab, change all adapters to 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. Verify that the Management Traffic checkbox is selected.
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 MGMT-Network and enter the <<var_mgmt_vlan_id>> for 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 <<var_nfs_vlan_id>> for the 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 <<var_nfs_vlan_id>> for the VLAN ID (Optional) field.
23. Continue with the NFS VMkernel creation by clicking Next.
24. Enter the <<var_nfs_vlan_id_ip>> and the <<var_nfs_vlan_id_mask>> for the NFS VLAN Interface for the host.
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. Enter the <<var_vmotion_vlan_id>> for the VLAN ID (Optional) field.
33. Select the “Use this port group for vMotion” checkbox.
34. Continue with the vMotion-VMkernel creation by clicking Next.
35. Enter the <<var_vmotion_vlan_id_ip>> and the <<var_vmotion_vlan_id_mask>> for the vMotion® VLAN Interface for the host.
36. Continue with the vMotion VMkernel creation by clicking Next.
37. Finalize the creation of the 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 VMkernel-vMotion Network by clicking OK.
41. Close the dialog box to finalize the ESXi host networking setup.

11.8 Mount Required Datastores

This step provides details for mounting the required datastores.

All Hosts

1. In each 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 the Network File System button and click Next.
6. Enter <<var_nfs_controller1_ip>> as the Server IP address.
7. Enter the path for the NFS export: /vol/infrastructure_datastore_1.
8. Verify that the “Mount NFS read only” checkbox is left cleared.
9. Enter the Datastore Name: infrastructure_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.
13. The Add Storage wizard appears. Select the Network File System button and click Next.
14. Enter <<var_nfs_ctrl1_ip>> as the Server IP address.
15. Enter the path for the NFS export: /vol/infrastructure_swap.
16. Verify that the “Mount NFS read only” checkbox is left cleared.
17. Enter the Datastore Name: infrastructure_swap.
18. Continue with the NFS datastore creation by clicking Next.
19. Finalize the creation of the NFS datastore by clicking Finish.

11.9 Move the VM Swap File Location

These steps provide details for moving the VM swap file location.

All Hosts

1. Select the host on the left pane within the vSphere Client.
2. Go to the Configuration tab to enable configurations.
3. Click the Virtual Machine Swapfile Location link in the Software box.
4. In the right pane, click Edit.
5. Select the “Store the swapfile in a swapfile datastore selected below” button.
6. Select infrastructure_swapdatastore.
7. Finalize the moving of the swap file location by clicking OK.

12 VMware vCenter 5.0 Deployment Procedure

The following sections provide detailed procedures for installing VMware vCenter 5.0 within an FlexPod Express configuration.

12.1 Build a VMware vCenter VM

One Server Only

1. Log in to an ESXi host by using the VMware vSphere Client.
2. In vSphere Client, select the host on the left pane.
3. Right-click the host and select New Virtual Machine.
4. Select the Custom button and click Next.
5. Name the Virtual Machine vCenter_Server and click Next.
6. Select infrastructure_datastore_1 and click Next.
7. Select the Virtual Machine Version: 8 button and click Next.

8. Verify that the Windows button and the Microsoft® Windows Server 2008 R2 (64-bit) Version option are selected and click Next.
9. Select two virtual sockets and one core per virtual socket and click Next.
10. Verify that 4GB of memory is selected and click Next.
11. Select two NICs total.
12. For NIC 1, select MGMT-Network and the VMXNET 3 Adapter.
13. For NIC 2, select NFS-Network and the VMXNET 3 Adapter.
14. Click Next.
15. Leave the LSI Logic SAS SCSI Controller selected and click Next.
16. Leave Create a new virtual disk selected and click Next.
17. Set the disk size to at least 40GB and click Next.
18. Click Next.
19. Select the Edit the virtual machine settings checkbox before completion and click Continue.
20. Select the Options tab.
21. Select Boot Options.
22. On the right, select the Force BIOS Setup checkbox.
23. Click Finish.
24. In the left pane, expand the host field by clicking the “+” sign.
25. Right-click the newly created vCenter_Server virtual machine and click Open Console.
26. Click the third button (green right-arrow) to power on the VM.
27. Click the ninth button (CD with a Wrench) to map the Windows Server® 2008 R2 SP1 ISO and select Connect to ISO image on local disk.
28. Navigate to the Windows Server 2008 R2 SP1 ISO, select it, and click Open.
29. Back in the BIOS Setup Utility window, use the Right Arrow key to move to the Boot menu. Use the Down Arrow key to highlight CD-ROM Drive. Use the + key two times to move CD-ROM Drive to the top of the list. Press F10 and Enter to Save and Exit the BIOS Setup Utility.
30. The Windows Installer will boot. Select the appropriate language, time and currency format, and keyboard, and click Next. Click Install Now. Verify that Windows Server 2008 R2 Standard (Full Installation) is selected and click Next. Accept the license terms and click Next. Select Custom (advanced). Verify that Disk 0 Unallocated Space is selected and click Next. Windows installation will complete.
31. After Windows Installation is complete and the VM has rebooted, click OK to enter the Administrator password. Enter and confirm the Administrator password and click the Blue Arrow to log in. Click OK to confirm the Password Change.
32. After you are logged in to the VM desktop, in the VM console window select the VM menu. Under Guest, select Install/Upgrade VMware Tools. Click OK.
33. If prompted to eject the Windows installation media prior to running setup for VMware tools, click OK.
34. In the popup window, select Run `setup64.exe`.
35. In the VMware Tools installer window, click Next.
36. Verify that Typical is selected and click Next.
37. Click Install.
38. Click Finish.
39. Click Yes to restart the VM.
40. After the reboot is complete, select the VM menu. Under Guest, select Send Ctrl+Alt+Del. Then enter the password to log back into the VM.


41. In Server Manager, click the “+” sign next to Diagnostics in the left-hand pane.
42. Click Device Manager. In the center pane, double-click Display Adapters.
43. Right-click Standard VGA Graphics Adapter and select Update Driver Software.
44. Click Browse my computer for driver software.
45. In the drop-down text box, type C:\Program Files\Common Files\VMware\Drivers\wddm_video. Verify that the Include Subfolders checkbox is selected.
46. Click Next.
47. Verify that Windows has successfully installed the VMware SVGA 3D video driver. Click Close.
48. Click Yes to restart the Guest OS.
49. After logging in, set the VM's time zone, IP address, gateway, and host name. If necessary, activate Windows.
50. Install all available Microsoft Windows updates by navigating to Start > All Programs > Windows Update.

Note: A restart may be required.

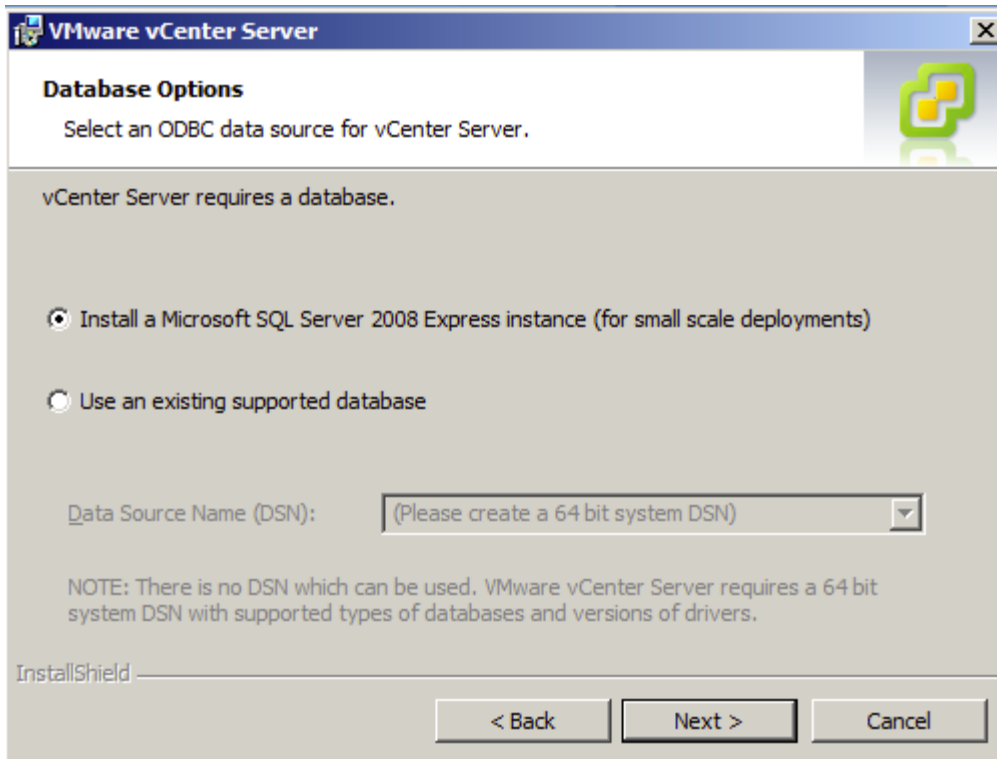
12.2 Install VMware vCenter Server

vCenter Server VM


1. In the vCenter Server VMware console, click the ninth button from the left (CD with a wrench) to map the VMware vCenter ISO and select Connect to ISO image on local disk.
2. Navigate to the VMware vCenter 5.0 Update 1 (VIM Setup) ISO, select it, and click Open.
3. In the popup, click Run `autorun.exe`.
4. In the VMware vCenter Installer window, verify that vCenter Server is selected and click Install.



5. Select the appropriate language and click OK to continue.
6. Click Next.
7. Click Next.
8. Accept the license terms and click Next.
9. Enter the user name, organization, and vCenter license key. Click Next.
10. Select Install a Microsoft SQL_Server 2008 Express Instance and click Next.



11. Click Next.



12. Note the warning and click OK.

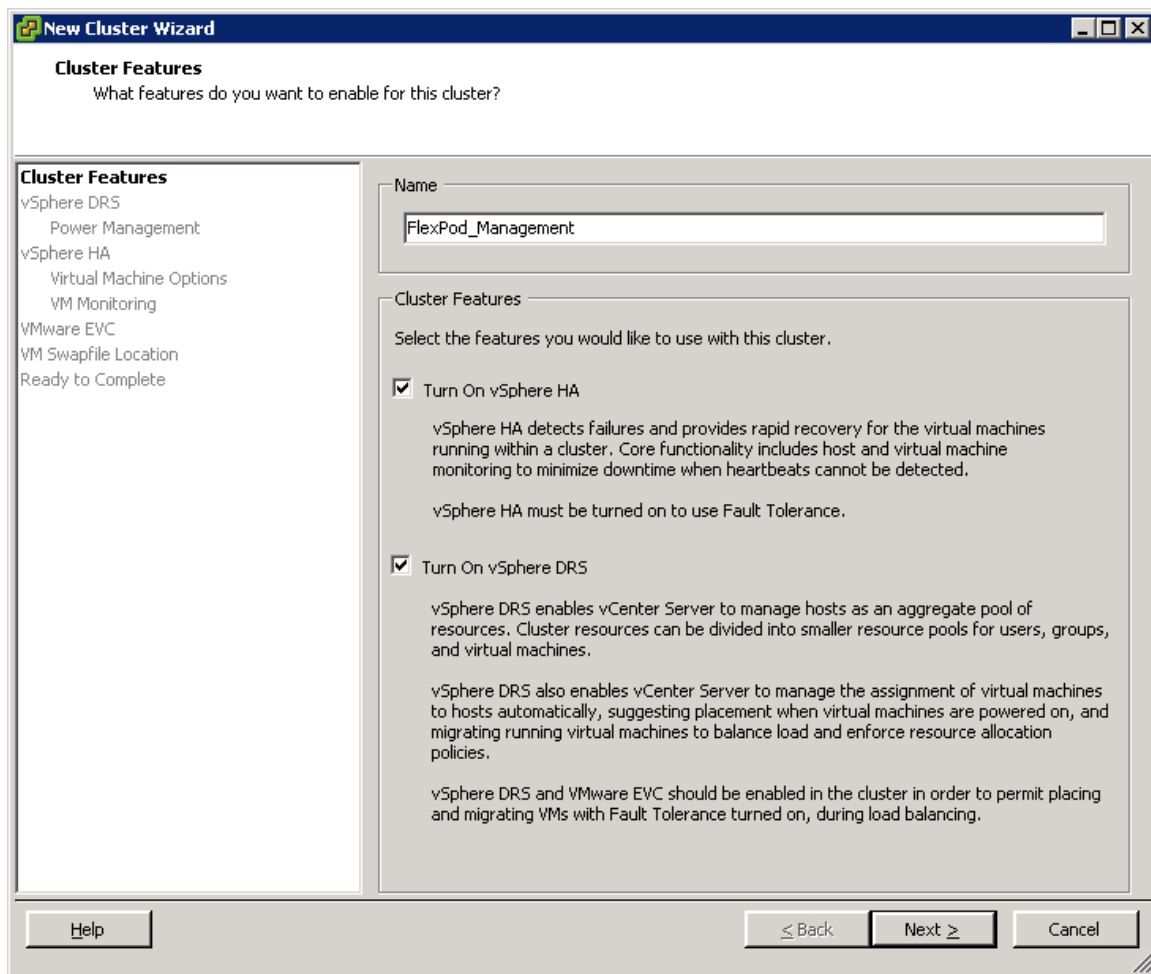
13. Click Next.

14. Click Next.
15. Verify that Create a standalone VMware vCenter Server instance is selected and click Next.
16. Click Next.
17. Click Next.
18. Select Small (less than 100 hosts or 1,000 virtual machines), click Next.
19. Click Install.
20. Click Finish.
21. Highlight vSphere Client in the VMware vCenter Installer window and click Install.
22. Click OK.
23. Click Next.
24. Click Next.
25. Accept the license agreement, click Next.
26. Enter the user name and organization, click Next.
27. Click Next.
28. Click Install.
29. Click Exit in the VMware vCenter Installer page.
30. Disconnect the VMware vCenter ISO from the vCenter VM.

12.3 vCenter Setup

vCenter Server VM

1. Using the vSphere Client, log in to the vCenter Server just created as Administrator.
2. In the center of the window, click Create a data center.
3. Enter FlexPod Express_DC as the data center name.
4. Right-click the newly created FlexPod Express_DC, and select New Cluster.
5. Name the cluster FlexPod Express_Cluster. Select the Turn On vSphere HA and Turn on vSphere DRS checkboxes. Click Next.



6. Accept the defaults for vSphere DRS and click Next.
7. Accept the defaults for Power Management and click Next.
8. Accept the defaults for vSphere HA and click Next.
9. Accept the defaults for Virtual Machine Options and click Next.
10. Accept the defaults for VM Monitoring and click Next.
11. Accept the defaults for VMware EVC and click Next.
12. Select Store the swapfile in the datastore specified by the host and click Next.
13. Click Finish.
14. Right-click the newly created FlexPod Express_Cluster and select Add Host...
15. In the Host field, enter the IP address of one of the ESXi hosts. Enter "root" as the user name and the associated password for this host.
16. Click Next.
17. Click Yes.
18. Click Next.
19. Select Assign a new license key to the host. Click Enter Key. Enter a vSphere license key and click OK. Click Next.
20. Click Next.

21. Click Next.
22. Click Finish. The host should now be added to the cluster.
23. Using the instructions above, add the remaining individual ESXi hosts to the FlexPod Express_Cluster cluster.

Note: There will be two ESXi hosts added to the cluster for the small FlexPod Express configuration. There will be four ESXi hosts added to the cluster for the medium FlexPod Express configuration.


13 NetApp Virtual Storage Console Deployment Procedure

This section provides the detailed procedures for installing the NetApp Virtual Storage Console. The deployment procedures that follow are customized to include the environment variables discussed previously. By the end of this section, a VSC will be a configured and operational plug-in with VMware vCenter.


13.1 Install VSC 4.1 Software

To install VSC 4.1 software, complete the following steps:


1. Log in to the Microsoft Windows Server VM that is running VMware vCenter.
2. Configure jumbo frames on the network adapter in the NFS-Network. Open Server Manager and click View Network Connections. Right-click the network connection in the <<var_nfs_vlan_id>> VLAN and select Properties. Click Configure. Select the Advanced tab. Select the Jumbo Packet property and use the pull-down menu to select Jumbo 9000. Click OK. Close the Network Connections window and close Server Manager.
3. From the Virtual Machine Console, download the NetApp Virtual Storage Console 4.1 to the VM from the [NetApp Support site](#).
4. To install the VSC plug-in, double-click the file.
5. In the Installation wizard landing page, click Next.
6. Click Next.



7. Select the location where VSC will be installed and click Next.



8. Make a note of the registration URL. This URL is needed to register the VSC plug-in with the vCenter Server after the installation. Click Install.




13.2 Register VSC with vCenter Server

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

1. Open a Web browser.

Note: A browser window with the URL shown in the previous figure automatically opens when the installation phase is complete. However, some browser settings might interfere with this function. If the browser window does not open automatically, open a browser window manually and enter the URL.

2. When the browser is running on the computer where VSC is installed, enter the URL provided by the Installation wizard (<https://localhost:8143/Register.html>). Otherwise, replace `localhost` with the host name or IP address (`<<var_vsc_server_ip>>`) of the VSC server.
3. In the Plug-in Service Information section, from the drop-down list, select the IP address used by the vCenter Server to access the VSC server. This IP should be in the `<<var_mgmt_vlan_id>>` VLAN.
4. In the vCenter Server Information section, enter the host name or IP address, port, user name, and password. Click Register to complete the registration.

 vSphere Plugin Registration

vSphere Plugin Registration

The Virtual Storage Console is registered as specified below. If you need to change the registration settings, update the fields below and then click "Register".

If you specify a new vCenter Server IP address, the Virtual Storage Console will unregister with the previously specified vCenter Server and then register with the newly specified vCenter Server.

Plugin service information

Host name or IP Address:

10.61.186.100

vCenter Server information

Host name or IP Address:

10.61.186.100

Port:

443

User name:

administrator

User password:

••••••••

Register


The registration process has completed successfully!

5. Close the Web browser and click Finish in the Installation Wizard window.

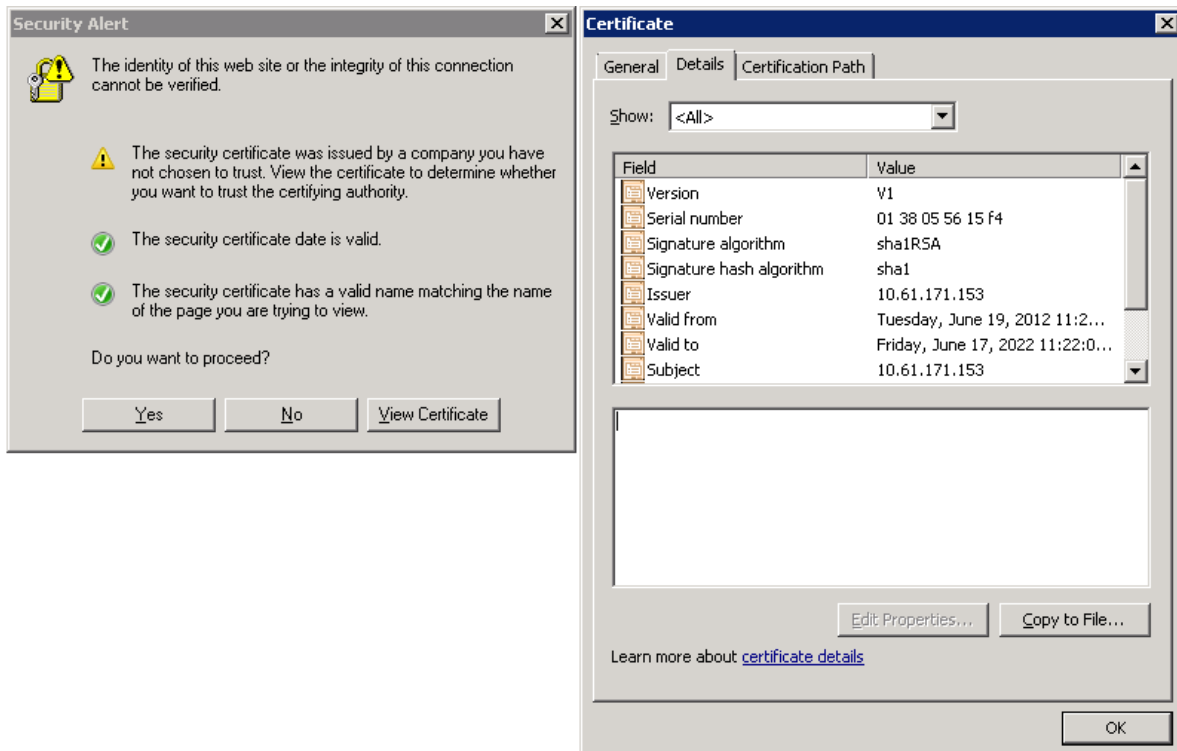
13.3 Discover and Add Storage Resources

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

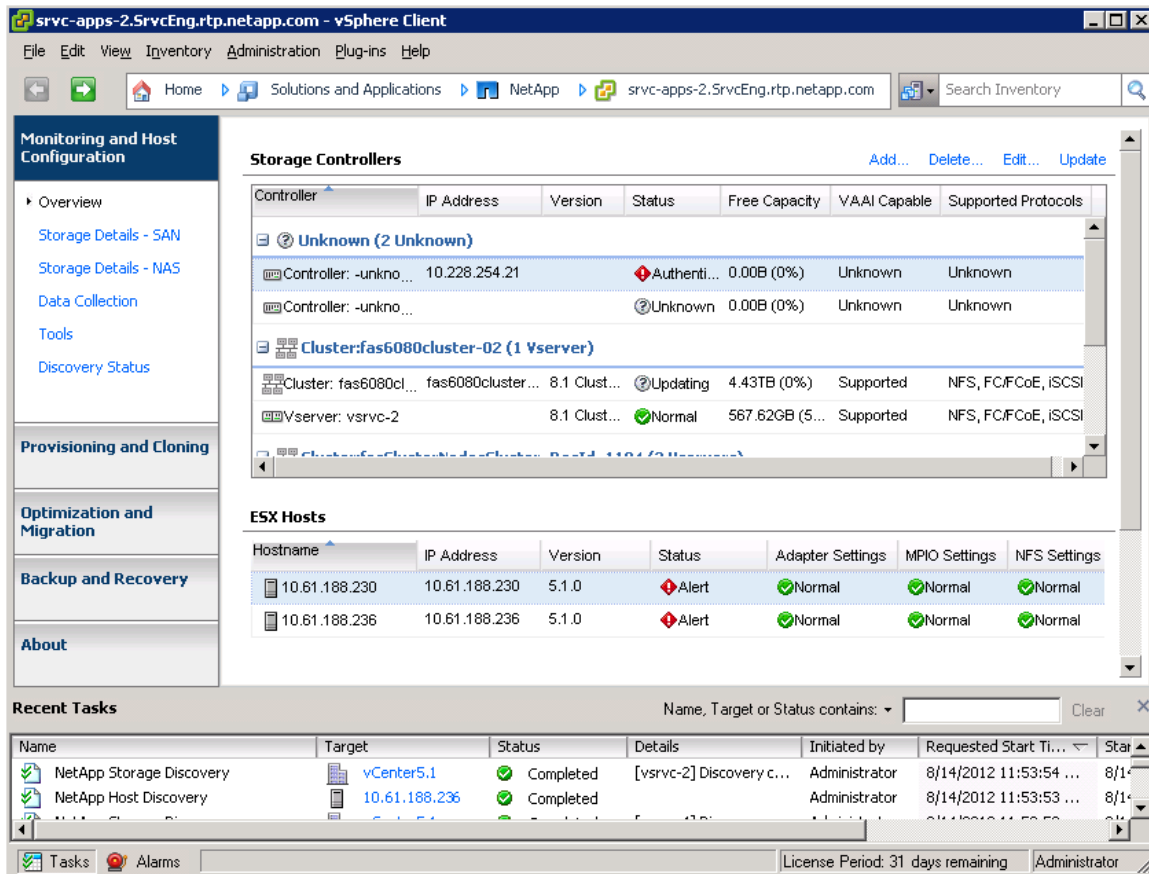
1. Log in to the vCenter Server using the vSphere Client. If you already have vSphere Client open, close and reopen it.
2. Click the Home tab in the upper left portion of the window.
3. In the Solutions and Applications section, click NetApp.



4. Click Yes when the certificate warning is displayed. Click View Certificate to see the certificate.



5. Click Monitoring and Host Configuration, if not selected by default.



If the discovery process does not start automatically, click Update in the Overview pane of the Monitoring and Host Configuration screen.

6. Add credentials for the two storage controllers.
7. Right-click the discovered storage controller.
8. Select Modify Credentials.
9. Make sure that the storage controller IP address is in the <<var_nfs_vlan_id>> VLAN.
10. Enter the root user name in the storage system configuration and its associated password.
11. Select the Use SSL checkbox.
12. Click OK.
13. Click OK.
14. Right-click in the white space under Storage Controllers and select Add Storage System.
15. Enter the Target Hostname (IP address on the NFS VLAN for controller 2). Enter root as the user name and the associated password.
16. Click OK.
17. Click OK.
18. Click Update to force discovery of storage controllers.

Storage Controllers

[Add...](#) [Delete...](#) [Edit...](#) [Update](#)

Controller	IP Address	Version	Status	Free Capacity	VAAI Capable	Supported Protocols
HA Pair:FAS2220-1/FAS2220-2						
Controller: FAS2220-1	192.168.10.90	8.1.1 7-...	Normal	1.70TB (78%)	Enabled	NFS
Controller: FAS2220-2	192.168.10.91	8.1.1 7-...	Alert	89.32GB (16...	Enabled	NFS

13.4 Optimal Storage Settings for ESXi Hosts

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

1. Select individual or multiple ESXi hosts and right-click to display the drop-down menu.
2. Select Set Recommended Values for these hosts.

Storage Controllers

[Add...](#) [Delete...](#) [Edit...](#) [Update](#)

Controller	IP Address	Version	Status	Free Capacity	VAAI Capable	Supported Protocols
HA Pair:FAS2220-1/FAS2220-2						
Controller: FAS2220-1	192.168.10.90	8.1.1 7-...	Normal	1.70TB (78%)	Enabled	NFS
Controller: FAS2220-2	192.168.10.91	8.1.1 7-...	Alert	89.32GB (16...	Enabled	NFS


ESX Hosts

Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
10.61.186.95	10.61.186.95	5.0.0	Alert	Alert	Normal	Alert
10.61.186.97	10.61.186.97	5.0.0	Alert	Alert	Normal	Alert

Set Recommended Values...
Show Details...
Skip Host...

3. Check the settings to apply to the selected vSphere hosts.

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).



Depending on the changes that have been made, servers may require a restart for network-related parameter changes to take effect. If no reboot is required, the Status value is set to Normal. If a reboot is required, the Status value is set to Pending Reboot. If a reboot is required, the ESX or ESXi servers should be placed into Maintenance Mode, evacuated (if necessary), and restarted before proceeding.

13.5 Provisioning and Cloning Setup

Provisioning and cloning in VSC 4.1 help administrators to provision both VMFS and NFS datastores at the data center, datastore cluster, or host level in VMware environments.

1. In a vSphere Client connected to vCenter, select Home > Solutions and Applications > NetApp. Select the Provisioning and Cloning tab on the left. Select Storage controllers.
2. In the main part of the window, right-click <<var_controller1>> and select Resources.
3. In the <<var_controller1>> resources window, use the arrows to move the `infrastructure_datastore_1, ifgrp0-<<var-nfs_vlan_id>>` and `aggr1` to the right. Select the Prevent further changes checkbox as shown in the following figure.

FAS2220-1 resources ✕

Configure the interfaces, volume, and aggregates you would like to use for provisioning and cloning below. The lists on the left contain all of the available interfaces, volumes, and aggregates respectively. The list on the right contains the interfaces, volumes, and aggregates that can be used for provisioning or cloning.

Interfaces: <div style="border: 1px solid black; height: 60px; width: 100%;"></div>	<div style="border: 1px solid black; padding: 2px;"> ifgrp0-10 - 192.168.10.90 </div>
Volumes: <div style="border: 1px solid black; padding: 2px;"> infrastructure_swap vol0 </div>	<div style="border: 1px solid black; padding: 2px;"> infrastructure_datastore_1 </div>
Aggregates: <div style="border: 1px solid black; padding: 2px;"> aggr0 </div>	<div style="border: 1px solid black; padding: 2px;"> aggr1 </div>

☒ Prevent further changes

Save Cancel

4. Click Save.

Note: Controller 2 is not configured for provisioning and cloning at this time because it is set up strictly as an HA partner to controller 1. Controller 2 is not configured with any usable aggregates or flexible volumes.

14 Bill of Materials

The following information details the hardware and software components used in validating both the small and medium FlexPod Express configurations included in this document.

Small Configuration

Table 9) Small configuration components.

Part Number	Product Description	Quantity Required
Cisco Components		
UCSC-EPOD-C220E-S	UCS Smart Play Bundle – FlexPod Express Small	1
N3K-UCS3048-F	Nexus 3048 for UCS Smart Play	2
N3KUK9-503U3.2	NX-OS Release 5.0(3)U3(2)	2
N2200-PAC-400W	N2K/N3K AC Power Supply, Std airflow (port side exhaust)	4
N3K-C3048-FAN	Nexus 3048 Fan Module, Port-side Exhaust	2

Part Number	Product Description	Quantity Required
N3K-C3064-ACC-KIT	Nexus 3064PQ Accessory Kit	2
UCS-SP5-C220E	UCS C220 M3 SFF w/ 2650 8x8GB 5709 1Gb 2PS	2
N20-BBLKD	UCS 2.5 inch HDD blanking panel	16
N2XX-ABPCI03-M3	Broadcom 5709 Quad Port 1Gb w/TOE iSCSI for M3 Servers	2
UCS-CPU-E5-2650	2.00 GHz E5-2650/95W 8C/20MB Cache/DDR3 1600MHz	4
UCS-MR-1X082RY-A	8GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	16
UCSC-HS-C220M3	Heat Sink for UCS C220 M3 Rack Server	4
UCSC-PCIF-01H	Half height PCIe filler for UCS	2
UCSC-PSU-650W	650W power supply for C-series rack servers	4
UCSC-RAID-11-C220	Cisco UCS RAID SAS 2008M-8i Mezz Card for C220 (0/1/10/5/50)	2
UCSC-RAIL1	Rail Kit for C220 C22 C24 rack servers	2
UCS-SP-600GB-HDD	600GB 6Gb SAS 10K RPM SFF HDD	4
CON-UCW5-SP5C220E	UCS HW 8X5XNBDOS UCS C220 M3 Rack Server	2
CON-SNTP-UCS3048F	SMARTNET 24X7X4 Nexus 3048 for UCS Smart Play	2
NetApp Components		
FAS2220-R6		1
FAS2220A-12X600-R6	FAS2220,HA,12x600GB,10k,Dual CNTLR	1
FAS2220A-HA-SW-R6	FAS2220A,HA CFO Software	2
SW-2220A-ONTAP8-P	SW, Data ONTAP Essentials,2220A,-P	2
SW-CIFS-C	SW,CIFS,-C	2
SW-FCP-C	SW,FCP,-C	2
SW-ISCSI-C	SW, iSCSI,-C	2
SW-NFS-C	SW,NFS,-C	2
X5518A-R6	Rack Mount Kit,FAS2020/40,R6	1
X800-42U-R6	Cabinet Component Power Cable,R6	2
Use NetApp QuoteEdge to determine part number	SupportEdge Standard, Premium or equivalent service from an authorized support services partner ¹	1

Medium Configuration

Table 10) Medium configuration components.

Part Number	Product Description	Quantity Required
Cisco Components		
UCSC-EPOD-C220E-M	UCS Smart Play Bundle - FlexPod Express Medium	1
N3K-UCS3048-F	Nexus 3048 for UCS Smart Play	2

¹ SupportEdge Premium required for Cooperative Support

Part Number	Product Description	Quantity Required
N3KUK9-503U3.2	NX-OS Release 5.0(3)U3(2)	2
N2200-PAC-400W	N2K/N3K AC Power Supply, Std airflow (port side exhaust)	4
N3K-C3048-FAN	Nexus 3048 Fan Module, Port-side Exhaust	2
N3K-C3064-ACC-KIT	Nexus 3064PQ Accessory Kit	2
UCS-SP5-C220E	UCS C220 M3 SFF w/ 2650 8x8GB 5709 1Gb 2PS	4
N20-BBLKD	UCS 2.5 inch HDD blanking panel	32
N2XX-ABPCI03-M3	Broadcom 5709 Quad Port 1Gb w/TOE iSCSI for M3 Servers	4
UCS-CPU-E5-2650	2.00 GHz E5-2650/95W 8C/20MB Cache/DDR3 1600MHz	8
UCS-MR-1X082RY-A	8GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	32
UCSC-HS-C220M3	Heat Sink for UCS C220 M3 Rack Server	8
UCSC-PCIF-01H	Half height PCIe filler for UCS	4
UCSC-PSU-650W	650W power supply for C-series rack servers	8
UCSC-RAID-11-C220	Cisco UCS RAID SAS 2008M-8i Mezz Card for C220 (0/1/10/5/50)	4
UCSC-RAIL1	Rail Kit for C220 C22 C24 rack servers	4
UCS-SP-600GB-HDD	600GB 6Gb SAS 10K RPM SFF HDD	8
CON-UCW5-SP5C220E	UCS HW 8X5XNBDOS UCS C220 M3 Rack Server	4
CON-SNTP-UCS3048F	SMARTNET 24X7X4 Nexus 3048 for UCS Smart Play	2
NetApp Components		
FAS2240-2-R5		1
F2240A-2-24X600-R5	FAS2240-2,HA,24x600GB,Dual CTL	1
FAS2240A-HA-SW-R5	FAS2240A, HA CFO Software, R5	2
SW-2240A-ONTAP8-P	SW, Data ONTAP Essentials, 2240A,-P	2
SW-CIFS-C	SW,CIFS,-C	2
SW-FCP-C	SW,FCP,-C	2
SW-ISCSI-C	SW, iSCSI,-C	2
SW-NFS-C	SW,NFS,-C	2
X5526A-R6	Rackmount Kit,4-Post,Universal,R6	1
X6557-R6	Cable, SASCntlr-Shelf/Shelf-Shelf/HA,0.5m	2
X6560-R6	Cable,Ethernet,0.5m RJ45 CAT6	1
X800-42U-R6	Cabinet Component Power Cable,R6	2
Use NetApp QuoteEdge to determine part number	SupportEdge Standard, Premium or equivalent service from an authorized support services partner ²	1

² SupportEdge Premium required for Cooperative Support

15 Open Management Ecosystem

FlexPod Express configurations support an open ecosystem of management and orchestration partners. See Appendix A for information on the initial setup of third-party management and orchestration software.

Appendix A: Cloupia Unified Infrastructure Controller Deployment Procedure

Import the CUIC VM into vCenter

1. Open the VMware vSphere client from a management workstation and connect to the VMware vCenter Server with the <<var_vcenter_ip>>, <<var_vcenter_username>>, <<var_admin_passwd>>.
2. Select File in the top left corner, and then choose Deploy OVF Template.
3. In the Deploy OVF Template page, click Browse and navigate to the location of the CUIC OVF file. Select the OVF file and click Open.
4. Click Next.
5. Click Next.
6. Read the terms of the End User License Agreement, and click Accept.
7. Click Next.
8. Enter `Cloupia_CUIC` as the VM name and choose `FlexPod_Express_DC` as the Inventory Location.
9. Click Next.
10. Choose `infrastructure_datastore_1`.
11. Click Next.
12. Click Next.
13. Choose the MGMT-Network network.
14. Click Next.
15. Click Next.
16. Click Finish. The import begins and the progress of the import is displayed on the screen.

Configure CUIC

1. Power on the CUIC VM within the vCenter Server.
2. During boot, enter `y` when prompted `Do you want to configure static IP.`
3. Enter <<var_cuic_vm_ip>> for the IP address, <<var_cuic_vm_gateway>> for the gateway, and <<var_cuic_vm_mask>> for the netmask.
4. Press Enter.
5. Enter `y` when asked `Do you want to continue [y/n]?`
The appliance should finish the boot process automatically
6. Open a Web browser and navigate to the <<var_cuic_vm_ip>>.
7. Log in to CUIC with default user name: `admin` and password: `admin`.
8. Select Administration and then System Administration.
9. Click the License tab.
10. Click Update License.

11. Copy and paste the license key text into the provided textbox, and click Save to continue.
12. Click OK.
13. Select Administration and then Physical Accounts.
14. Click the Data Centers tab.
15. Click Add.
16. Specify <<var_cuic_datacenter>> for the name of the data center.
17. Click Add.

Add NetApp Storage Controllers to CUIC

Controller 1 and 2

1. Select Administration and then Physical Accounts.
2. Click Add.
3. Select Storage for Category Type.
4. Select NetApp (NetApp ONTAP) for the Account Type.
5. Enter <<var_controller_hostname>> in the Account Name field.
6. Select <<var_cuic_datacenter>> in the Data Center field.
7. Enter <<var_controller_e0m_ip>> in the Server Address field.
8. Enter root in the User ID field.
9. Enter <<var_admin_passwd>> in the Password field.
10. Select HTTP for the Transport Type.
11. Enter 80 in the Port field.
12. Optionally, add a description, contact e-mail, location, and service provider.
13. Click Add.
14. After the account has been added, select the newly added account from the list and choose Test Connect.
15. A window displaying Connection Successful appears. Click Close.

Add Cisco Nexus Switches to CUIC

Switch 1 and 2

1. Select Administration and then Physical Accounts.
2. Click the Manage Network Elements tab.
3. Click Add Network Element.
4. Select <<car_cuic_datacenter>> in the Data Center field.
5. Enter <<var_mgmt0_ip_address>> in the Device IP field.
6. Select SSH for the Protocol.
7. Enter 22 in the Port field.
8. Enter admin in the Login field.
9. Enter <<var_admin_passwd>> in the Password field.
10. Leave the Enable Password field blank.
11. Select Nexus OS for the Device Category.

12. Click Submit.
13. After the account has been added select the newly added account from the list and choose Test Connection.
14. A window displaying `Connection Successful` appears. Click Close.

Add Cisco C-Series Servers to CUIIC

All Servers

1. Select Administration and then Physical Accounts.
2. Click Add.
3. Select Compute for Category Type.
4. Select Cisco Standalone Rack Server in the Account Type field.
5. Enter `<<var_server_hostname>>` in the Account Name field.
6. Select `<<var_cuic_datacenter>>` in the Data Center field.
7. Enter `<<var_server_mgmt_ip>>` in the Server Address field.
8. Enter `admin` for the User ID.
9. Enter `<<var_admin_passwd>>` in the Password field.
10. Select `https` for the Transport Type.
11. Enter `443` for the Port.
12. Optionally, add a description, contact e-mail, location, and service provider.
13. Click Add.
14. After the account has been added, select the newly added account from the list and choose Test Connect.
15. A window displaying `Connection Successful` appears. Click Close.

Appendix B: Cloupia Unified Infrastructure Controller Bill of Materials

Cloupia Components for Small FlexPod Express Configuration

Table 11) Cloupia components for small configuration.

Part Number	Product Description	Quantity Required
Cloupia Components		
CUIC-V3-START-XPSM-STD	CUIC Starter Kit - Small FlexPod Express, Phys & Virt, Std Feat Set	1
CUIC-V3-BASE-STD^	Bundle:CUIC Base Lic - Standard Feature Set	1
CUIC-V3-PLAT-XPSM^	Bundle:CUIC Platform Lic - Small FlexPod Express	1
CUIC-V3-CONN-UCS-C^	Bundle:CUIC Base Lic - Connector for Cisco UCS C-Series Features/Platforms Only	1
CUIC-V3-CONN-CSCO-XPOD^	Bundle:CUIC Base Lic - Connector for Cisco FlexPod Express Features/Platforms Only	1
CUIC-V3-CONN-NTAP-XPOD^	Bundle:CUIC Base Lic - Connector for NetApp FlexPod Express Features/Platforms Only	1
CUIC-V3-SERV-R-STD-	Bundle:CUIC Resource Lic - Std Feat Set, One Rack Svr, For	2

Part Number	Product Description	Quantity Required
Cloupia Components		
XPOD^	FlexPod Express Platforms	
CUIC-V3-NETW-STD^	Bundle:CUIC Resource Lic - Std Feat Set For One Network Device	2
CUIC-V3-STOR-STD^	Bundle:CUIC Resource Lic - Std Feat Set For One Storage Controller	2
CUIC-V3-VM-STD-25PK-XPOD^	Bundle:CUIC Resource Lic - Std Feat Set For 25 VMs, For FlexPod Express Platforms	1

Cloupia Components for Medium FlexPod Express Configuration

Table 12) Cloupia components for medium configuration.

Part Number	Product Description	Quantity Required
Cloupia Components		
CUIC-V3-START-XPMD-STD	CUIC Starter Kit - Medium FlexPod Express, Phys & Virt, Std Feat Set	1
CUIC-V3-BASE-STD^	Bundle:CUIC Base Lic - Standard Feature Set	1
CUIC-V3-PLAT-XPMD^	Bundle:CUIC Platform Lic - Medium FlexPod Express	1
CUIC-V3-CONN-UCS-C^	Bundle:CUIC Base Lic - Connector for Cisco UCS C-Series Features/Platforms Only	1
CUIC-V3-CONN-CSCO-XPOD^	Bundle:CUIC Base Lic - Connector for Cisco FlexPod Express Features/Platforms Only	1
CUIC-V3-CONN-NTAP-XPOD^	Bundle:CUIC Base Lic - Connector for NetApp FlexPod Express Features/Platforms Only	1
CUIC-V3-SERV-R-STD-XPOD^	Bundle:CUIC Resource Lic - Std Feat Set, One Rack Svr, For FlexPod Express Platforms	4
CUIC-V3-NETW-STD^	Bundle:CUIC Resource Lic - Std Feat Set For One Network Device	2
CUIC-V3-STOR-STD^	Bundle:CUIC Resource Lic - Std Feat Set For One Storage Controller	2
CUIC-V3-VM-STD-25PK-XPOD^	Bundle:CUIC Resource Lic - Std Feat Set For 25 VMs, For FlexPod Express Platforms	2

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®