



NetApp Verified Architecture

FlexPod Express for VMware vSphere 8.0 with Cisco UCS C-Series Standalone and NetApp AFF using Infrastructure as Code (IaC) NVA

Ruchika Lahoti, NetApp

September 2023 | NVA-1171

In partnership with:



Abstract

This NetApp Verified Architecture (NVA) deployment guide provides the detailed steps needed to configure the infrastructure components to deploy VMware vSphere 8.0 with Cisco UCS C-series standalone servers and the associated tools to create a highly reliable and highly available FlexPod Express-based virtual infrastructure using ansible.

TABLE OF CONTENTS

Program Summary	3
Solution Overview	3
FlexPod Converged Infrastructure program.....	3
NetApp Verified Architecture Program	4
Solution Technology	4
Use-case Summary	5
Technology Requirements	5
Hardware Requirements	5
Software Requirements	6
Physical Infrastructure	6
Cabling Information.....	6
Ansible Automation Workflow and Solution Deployment	8
Prerequisites	9
Prepare Management Workstation (Control Machine).....	9
Clone GitHub Collection.....	10
Deployment Procedures	10
Cisco Nexus 93180YC-FX Deployment Procedure	12
NetApp Storage Deployment Procedure.....	13
Cisco UCS C220 Standalone Rack Server Deployment Procedure.....	19
VMware vSphere 8.0 Deployment Procedure	28
VMware vCenter Server 8.0 Deployment Procedure	33
vCenter and ESXi Ansible Setup	36
Ansible ONTAP Storage Configuration Part 2	39
NetApp ONTAP Tools 9.12 Deployment	42
Conclusion	44
Acknowledgment.....	44
Where to find additional information	44
Version History	45
Copyright Information	46

Program Summary

FlexPod® Express with Cisco UCS C-series Standalone Rack Servers and NetApp AFF is a predesigned, best practice architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the same set of tools with which they are familiar. New FlexPod Express customers can easily scale and manage their FlexPod solutions as they scale and grow their environment.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses who are looking for an easy-to-manage infrastructure that is suitable for almost any of their workload needs.

Customers interested in understanding the manual deployment procedures for this solution should refer to is provided in this [deployment guide](#).

Solution Overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure program.

FlexPod Converged Infrastructure program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Based on customer requirements, you can update a given CVD or NVA configuration to meet customer needs as long as the changes do not create an unsupported configuration.

As depicted in Figure 1, the FlexPod program includes two solutions: FlexPod Express and FlexPod Datacenter. FlexPod Express offers customers an entry-level solution with technologies available from Cisco and NetApp. FlexPod Datacenter delivers an optimal multipurpose foundation for various workloads and applications for the data center.

Figure 1) FlexPod portfolio.



NetApp Verified Architecture Program

The NVA program offers customers an engineering solution with the following qualities:


- Thoroughly tested
- Prescriptive in nature
- Minimized deployment risks
- Accelerated time to market

This guide details the automated deployment of VMware vSphere 8.0 on FlexPod Express with UCS C-series Standalone servers and NetApp AFF storage. The following sections list the components used for the deployment of this solution.

Solution Technology

This solution leverages technologies from NetApp, Cisco, and VMware. It features NetApp AFF A250 running ONTAP 9.12.1, dual Cisco Nexus 93180YC-FX switches, and Cisco UCS C220 M6 servers that run VMware vSphere 8.0. Figure 2 shows an architecture of this validated solution and the cabling illustrations.

Figure 2) FlexPod Express for VMware vSphere 8.0 with Cisco UCS C-series Standalone and NetApp AFF architecture.



The storage data path from the vSphere 8.0 hosts, running on the UCS C220 M6 rack servers are going from the virtual NIC connected to the Nexus 93180YC-FX switches to the AFF A250 storage through 25GbE networking card. Alternatively, for solutions that do not require high storage data bandwidth, the 10GbE onboard ports on AFF A250 can be utilized for storage data path. The solution environment can be scaled to at least five UCS C-series Standalone Servers requiring two switch ports on each switch per server.

Use-case Summary

You can apply the FlexPod Express solution to several use cases, including the following:

- Remote Office / Branch Office
- Small and midsize businesses
- Edge Computing deployments
- Environments that require a dedicated and cost-effective solution
- Ideal for virtualized and mixed workloads.

Technology Requirements

A FlexPod Express system requires a combination of hardware and software components. In addition to the required hardware and software components, you can add additional hardware components to scaleup the solution. Furthermore, you can add additional software and applications to help manage the solution or provide additional functionalities.

Hardware Requirements

Depending on your business requirements, you can use different hypervisors on the same reference FlexPod Express with UCS C-series Standalone hardware configuration.

Table 1 lists the reference hardware components for a FlexPod Express with UCS C-series Standalone configuration.

Table 1) Hardware requirements for the base FlexPod Express with UCS C-series Standalone configuration.

Hardware	Quantity
AFF A250 series HA pair	1
Cisco Nexus 9000 series switches	2
Cisco UCS C220 M6 server	2
Cisco UCS Virtual Interface Card (VIC) 1455/1467 for C220 M6 rack server	2

Note: The actual hardware components that are selected for a solution implementation can vary based on customer requirements. For example, instead of using an AFF A250 HA pair, you can use an AFF A150 or AFF C250 controller HA pair to meet the cost or capacity requirements.

- The rest of this deployment guide assumes the use of an AFF A250 HA pair for storage and a pair of Cisco Nexus 93180YC-FX switches for networking.
- The management network and console connections for the FlexPod components are assumed to be connected to an existing infrastructure, which is deployment specific, and therefore not documented in this deployment guide.

Software Requirements

Table 2 lists the software components that are required to implement the FlexPod Express with UCS C-series Standalone solution.

Table 2) Software requirements for the FlexPod Express with UCS C-series Standalone implementation.

Software	Version	Details
Cisco UCS CIMC	4.2(3d)	For UCS C220 M6 servers
Cisco nenic driver	1.0.45.0	For VIC 1455 / 1467 interface cards
Cisco NX-OS	10.2(4)	For Cisco Nexus 93180YC-FX switches
NetApp ONTAP	9.12.1P2	For AFF A250 controllers
ONTAP Tools for VMware vSphere	9.12	

Table 3 lists the software that is required for a VMware vSphere implementation on FlexPod Express with UCS C-series Standalone.

Table 3) Software requirements for a VMware vSphere 8.0 implementation on the FlexPod Express with UCS C-series Standalone.

Software	Version
VMware vSphere ESXi hypervisor	8.0
VMware vCenter Server appliance	8.0

Physical Infrastructure

Cabling Information

The reference validation cabling details are shown in Figure 3 and Table 4 through Table 9. For this deployment guide, the console and management network of the FlexPod components are connected to the existing console and management network and are not documented in the cabling information below.

Figure 3) Reference validation components and cabling.




Table 4) Cabling information for Cisco Nexus 93180YC-FX switch A.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 93180YC-FX A	Eth1/54	Remote switch for in-band management network uplink	deployment specific
	Eth1/1	AFF A250 A	e1a
	Eth1/2	AFF A250 A	e1b
	Eth1/3	AFF A250 B	e1a
	Eth1/4	AFF A250 B	e1b
	Eth1/11	UCS-1-VIC-P0	Eth0
	Eth1/12	UCS-1-VIC-P1	Eth1
	Eth1/13	UCS-2-VIC-P0	Eth0
	Eth1/14	UCS-2-VIC-P1	Eth1
	Eth1/51	Cisco Nexus 93180YC-FX B	Eth1/51
	Eth1/52	Cisco Nexus 93180YC-FX B	Eth1/52

Table 5) Cabling information for Cisco Nexus 93180YC-FX switch B.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 93180YC-FX B	Eth1/54	Remote switch for in-band management network	deployment specific
	Eth1/1	AFF A250 A	e1c
	Eth1/2	AFF A250 A	e1d
	Eth1/3	AFF A250 B	e1c
	Eth1/4	AFF A250 B	e1d
	Eth1/11	UCS-1-VIC-P2	Eth2
	Eth1/12	UCS-1-VIC-P3	Eth3
	Eth1/13	UCS-2-VIC-P2	Eth2
	Eth1/14	UCS-2-VIC-P3	Eth3
	Eth1/51	Cisco Nexus 93180YC-FX A	Eth1/51
	Eth1/52	Cisco Nexus 93180YC-FX A	Eth1/52

Table 6) Cabling information for NetApp AFF A250 A.

Local Device	Local Port	Remote Device	Remote Port
NetApp AFF A250 A	e1a	Cisco NX 93180YC-FX A	Eth1/1
	e1b	Cisco NX 93180YC-FX A	Eth1/2
	e1c	Cisco NX 93180YC-FX B	Eth1/1
	e1d	Cisco NX 93180YC-FX B	Eth1/2
	e0c	NetApp AFF A250 B	e0c
	e0d	NetApp AFF A250 B	e0d

Table 7) Cabling information for NetApp AFF A250 B.

Local Device	Local Port	Remote Device	Remote Port
NetApp AFF A250 B	e1a	Cisco NX 93180YC-FX A	Eth1/3
	e1b	Cisco NX 93180YC-FX A	Eth1/4
	e1c	Cisco NX 93180YC-FX B	Eth1/3
	e1d	Cisco NX 93180YC-FX B	Eth1/4
	e0c	NetApp AFF A250 A	e0c
	e0d	NetApp AFF A250 A	e0d

Table 8) Cabling information for Cisco UCS VIC-1455 A.

Local Device	Local Port	Remote Device	Remote Port
Cisco UCSVIC-1455 A	P0	Cisco NX 93180YC-FX A	Eth1/11
	P1	Cisco NX 93180YC-FX A	Eth1/12
	P2	Cisco NX 93180YC-FX B	Eth1/11
	P3	Cisco NX 93180YC-FX B	Eth1/12

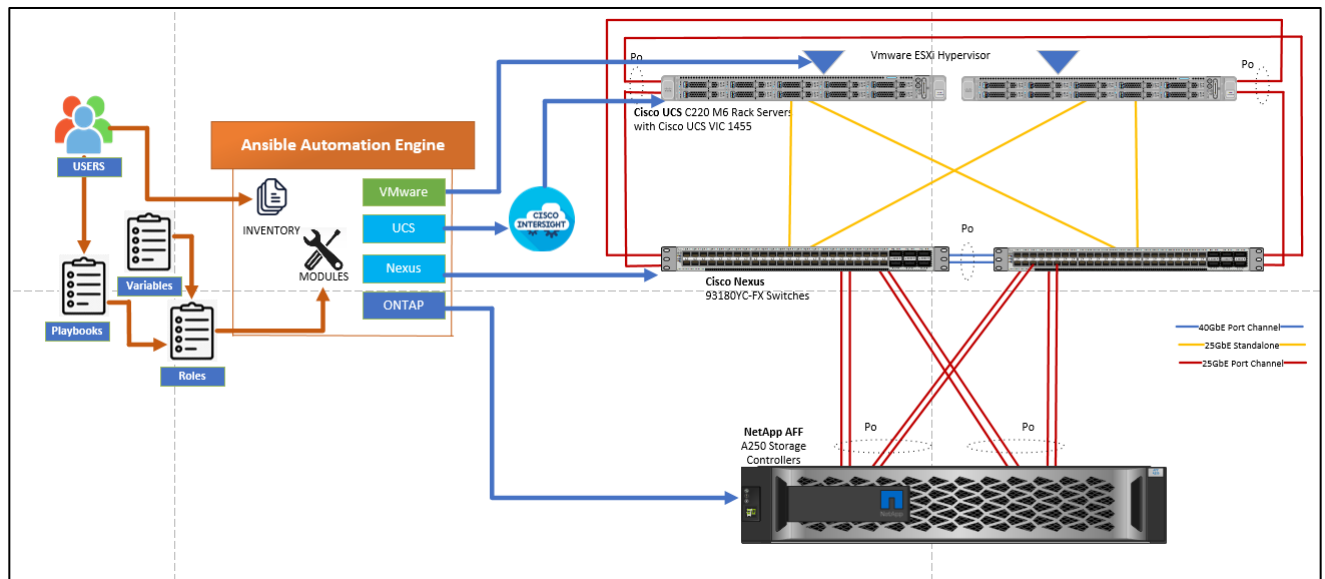
Table 9) Cabling information for Cisco UCS VIC-1455 B.

Local Device	Local Port	Remote Device	Remote Port
Cisco UCSVIC-1455 B	P0	Cisco NX 93180YC-FX A	Eth1/13
	P1	Cisco NX 93180YC-FX A	Eth1/14
	P2	Cisco NX 93180YC-FX B	Eth1/13
	P3	Cisco NX 93180YC-FX B	Eth1/14

Ansible Automation Workflow and Solution Deployment

The Ansible automated FlexPod solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, NetApp ONTAP Storage, Cisco UCS and VMware ESXi.

Below Figure illustrates the FlexPod solution implementation workflow which is explained in the following sections. The FlexPod infrastructure layers are first configured in the order illustrated.



Prerequisites

Setting up the solution begins with a management workstation or VM that has access to the Internet and with a working installation of Ansible. The management workstation commonly runs a variant of Linux for ease of use with these command-line-based tools. Instructions for installing the workstation are not included in this document, but basic installation and configuration of Ansible is covered. A guide for getting started with Ansible can be found at the following link:

- Ansible Community Documentation: https://docs.ansible.com/ansible_community.html
- To use the Ansible playbooks demonstrated in this document, the management workstation must also have a working installation of Git and access to the public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following link:
GitHub repository: <https://github.com/ucs-compute-solutions/FlexPod-Express-Intersight>
- The Cisco Nexus, NetApp Storage, and Cisco UCS must be physically racked, cabled, powered, and configured with management IP addresses before the Ansible-based installation procedure can begin as shown in the cabling diagram (Figure 2).
- Before running each Ansible Playbook to setup the Network, Storage, Cisco UCS, and VMware ESXi, various variables must be updated based on the customers environment and specific implementation with values such as the VLANs, pools and ports on Cisco UCS, IP addresses for NFS and NVMe-TCP interfaces and values needed for VMware ESXi.
- Day 2 Configuration tasks such as adding datastores or ESXi servers can be performed manually or with Cisco Intersight Cloud Orchestrator (ICO).

Prepare Management Workstation (Control Machine)

In this procedure, the installation steps are performed on the CentOS Stream 8 (install default Server with GUI) management host to prepare the host for solution deployment to support the automation of Cisco UCS, Cisco Nexus, NetApp Storage and VMware ESXi using Ansible Playbooks.

Note: The following steps were performed on a CentOS Stream 8 Virtual Machine as the root user.

1. Install the EPEL repository on the management host.

```
cd
dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

2. Install Ansible engine.

```
dnf install ansible
```

3. Verify Ansible version to make sure it is release 2.9 or later.

```
ansible --version
ansible [core 2.13.7]
  config file = /etc/ansible/ansible.cfg
  configured module search path = ['/root/.ansible/plugins/modules',
'/usr/share/ansible/plugins/modules']
  ansible python module location = /root/.local/lib/python3.9/site-packages/ansible
  ansible collection location = /root/.ansible/collections:/usr/share/ansible/collections
  executable location = /usr/bin/ansible
  python version = 3.9.17 (main, Jun  6 2023, 20:11:04) [GCC 9.4.0]
  jinja version = 3.1.2
  libyaml = False
```

4. Install sshpass.

```
dnf install sshpass
```

5. Install NetApp specific python modules.

```
pip3 install netapp-lib
```

6. Install ansible-galaxy collections and other dependencies for Cisco Nexus, NetApp ONTAP, Cisco UCS, and VMware as follows:

```
ansible-galaxy collection install cisco.intersight
ansible-galaxy collection install cisco.nxos
pip3 install ansible-pylibssh
ansible-galaxy collection install netapp.ontap
ansible-galaxy collection install community.vmware
pip3 install wheel
pip3 install --upgrade pip setuptools
pip3 install -r ~/.ansible/collections/ansible_collections/community/vmware/requirements.txt
```

Clone GitHub Collection

Note: You need to use a GitHub repository from one public location; the first step in the process is to clone the GitHub collection named FlexPod-Express-Intersight (<https://github.com/ucs-compute-solutions/FlexPod-Express-Intersight.git>) to a new empty folder on the management workstation. Cloning the repository creates a local copy, which is then used to run the playbooks that have been created for this solution.

1. From the management workstation, create a new folder for the project. The GitHub collection will be cloned in a new folder inside this one, named /root/FlexPod-Express-Intersight
2. Open a command-line or console interface on the management workstation and change directories to the new folder just created.
3. Clone the GitHub collection using the following command:

```
git clone https://github.com/ucs-compute-solutions/FlexPod-Express-Intersight.git
```

4. Change directories to the new folder named FlexPod-Express-Intersight.

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 A or component B, or -01 and -02 in naming. For example, storage controller A and

storage controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert deployment-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. Table 10 describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site information and used to implement the document configuration steps.

Note: For this validation, existing network infrastructure is used for the out-of-band management connectivity of the FlexPod components, and those details are not included in this guide.

Table 10) Required VLANs.

VLAN Name	VLAN Purpose	VLAN ID
Native VLAN	VLAN to which untagged frames are assigned	1101
In-band Management VLAN	VLAN for in-band management interfaces	2229
NFS	VLAN for NFS traffic	2230
VM-Traffic	VLAN for VM application traffic	2231
vMotion	VLAN designated for the movement of virtual machines (VMs) from one physical host to another	2232
NVMe-TCP-A	NVMe-TCP-A path when using NVMe-TCP	2233
NVMe-TCP-B	NVMe-TCP-B path when using NVMe-TCP	2234

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <xxx_vlan_id>, where xxx is the purpose of the VLAN (such as NFS). Substitute those variables with the VLAN IDs appropriate for the deployment environment.

There are various management tools and ways to manage and deploy a VMware solution. This NVA provides information on deploying the basic VMware infrastructure. Table 11 lists the Standard Virtual Switch created for this solution and Table 12 lists the Distributed Switch and its details.

Table 11) VMware standard vSwitches created for the solution.

vSwitch Name	Adapters	MTU	Failover Order
vSwitch0	vmnic0, vmnic2	9000	For the Management Network, the failover order is configured for active/active configuration.

Table 12) VMware distributed Switch created for the solution.

vSwitch Name	Adapters	MTU	Failover Order
vDS0	vmnic1, vmnic3	9000	For the NFS, vMotion, NVMe-TCP-A, NVMe-TCP-B and Virtual Machine port groups. The failover order is configured for LAG configuration.

Table 13) VMware Infrastructure VMs created for the solution.

VM Description	Host Name
VMware vCenter Server	g13vcenter.fpmc.sa
ONTAP Tools for VMware vSphere	otv.fpmc.sa

Cisco Nexus 93180YC-FX Deployment Procedure

The following section details the Cisco Nexus 93180YC-FX switch configuration used in a FlexPodExpress environment.

Initial setup of Cisco Nexus 93180YC-FX switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.

Note: This procedure assumes that you are using a Cisco Nexus 93180YC-FX running NX-OS software release 10.2(4).

Table 14) Nexus 10.2(4) configuration information.

Switch Detail	Switch Detail Value
Switch administrator password	<admin_password>
Switch A name	<switchname_a>
Switch B name	<switchname_b>
Switch A management IP address	<switch_ip_a>
Switch B management IP address	<switch_ip_b>
Switch management netmask	<switch_netmask>
Switch management gateway	<switch_gateway>
Switch NTP server	<ntp_ip>
Switch A NTP distribution interface IP	<switch_ntp_ip_a>
Switch B NTP distribution interface IP	<switch_ntp_ip_b>
In-band management VLAN netmask length	<ib_mgmt_vlan_netmask_length>

1. After 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.
2. You can configure the FlexPod Express out-of-band management network in multiple ways. In this deployment guide, the FlexPod Express Cisco Nexus 93180YC-FX switches are connected to an existing out-of-band management network. Layer 3 network connectivity is required between the out-of-band and in-band management subnets.
3. To configure the Cisco Nexus 93180YC-FX switches, power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both the switches, substituting the variables below with the appropriate information for switches A and 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

Disabling POAP.....Disabling POAP

poap: Rolling back, please wait... (This may take 5-15 minutes)
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A/B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address : <nexus-A/B-out_of_band_mgmt0-ip>
Mgmt0 IPv4 netmask : <nexus-A/B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
```

```
IPv4 address of the default gateway : <nexus-A/B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <ntp_server>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
```

4. A summary of your configuration is displayed, and you are asked if you would like to edit the configuration. If your configuration is correct, enter n.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
```

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

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

Ansible Nexus Switch Configuration

Configure the Cisco Nexus switches from the management workstation

1. Add Nexus switch ssh keys to /root/.ssh/known_hosts. Adjust known_hosts as necessary if errors occur.

```
ssh admin@<nexus-A-mgmt0-ip>
exit
ssh admin@<nexus-B-mgmt0-ip>
exit
```

2. Edit the following variable files to ensure proper Cisco Nexus variables are entered:

- FlexPod-Express-Intersight/FlexPod-Express-Intersight/inventory
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/host_vars/n9kA.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/host_vars/n9kB.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/NEXUSconfig/defaults/main.yml

Note: Switch configuration can be done one switch at a time by commenting one switch out in inventory and running the playbook. This may need to be done if the switches are shared with other FlexPod and additional configuration needs to be added between playbook runs.

1. From /FlexPod-Express-Intersight/FlexPod-Express-Intersight/, run the Setup_Nexus.yml Ansible playbook.

```
ansible-playbook ./Setup_Nexus.yml -i inventory
```

2. Once the Ansible playbook has been run on both switches, it is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the time-zone and daylight savings time or summertime, please see Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 10.2(x). Sample clock commands for the United States Eastern timezone are:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

3. ssh into each switch and execute the following commands:

```
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day>
<end-month> <end-time> <offset-minutes>
```

NetApp Storage Deployment Procedure

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF A250 installation

NetApp Hardware Universe

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

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

Access the [HWU](#) application to view the system configuration guides. Click the Products tab to select the Platforms menu to view the compatibility between different versions of the ONTAP software and the NetApp storage appliances with your desired specifications.

Alternatively, to compare components by storage appliance, click Utilities and select Compare Storage Systems.

Controller AFF A250 prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections:

Controller AFF A250 prerequisites

- Electrical Requirements
- Supported Power Cords
- Onboard Ports and Cables

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A250 Documentation](#).

NetApp ONTAP 9.12.1P2

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9 Software Setup Guide](#) (available in the [ONTAP 9 Documentation Center](#)).

Note: This system is set up in a two-node switchless cluster configuration.

Table 15) ONTAP 9.12.1P2 installation and configuration information

Cluster Detail	Cluster Detail Value
Cluster st-node01 IP address	< st-node01_mgmt_ip>
Cluster st-node01 SP address	< st-node01_sp_ip>
Cluster st-node01 netmask	< st-node01_mgmt_mask>
Cluster st-node01 gateway	< st-node01_mgmt_gateway>
Cluster st-node02 IP address	< st-node02_mgmt_ip>
Cluster st-node02 SP address	< st-node02_sp_ip>
Cluster st-node02 netmask	< st-node02_mgmt_mask>
Cluster st-node02 gateway	< st-node02_mgmt_gateway>
ONTAP 9.12.1P2 URL	<url_boot_software>
Name for cluster	<clustername>

Cluster administrator password	<clustermgmt_password>
Cluster management IP address	<clustermgmt_ip>
Cluster management gateway	<clustermgmt_gateway>
Cluster management netmask	<clustermgmt_mask>
Cluster feature license keys	<licensekeys>
Domain name	<domain_name>
DNS server IP (you can enter more than one)	<dns_server_ip>
NTP server IP (you can enter more than one)	<ntp_server_ip>
Controller location	<controller_location>

To initialize controller A (node st-node01) and controller B (node st-node02), use two serial console port program sessions to communicate with the storage controller A and controller B, respectively.

Initialize node st-node01

To initialize node st-node01, complete the following steps:

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...
```

1. Allow the system to boot.

```
autoboot
```

2. Press Ctrl-C to enter the Boot menu.

Note: If ONTAP 9.12.1P2 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.12.1P2 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 13.

3. To install new software, select option 7.
4. Enter y to perform an upgrade.
5. Select e0M for the network port you want to use for the download.
6. Enter n to skip the reboot.
7. Enter the IP address, network mask, and default gateway for e0M in their respective places.

```
< st-node01_mgmt_ip> < st-node01_mgmt_mask> < st-node01_mgmt_gateway>
```

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

Note: This web server must be pingable.

```
< url_boot_software>
```

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


```

Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} [y] y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y

Rebooting...
Files /cfcard/x86_64/freebsd/image2/VERSION and /var/VERSION differ
.
Terminated
Setting default boot image to image2...
done.
Uptime: 15m0s
System rebooting...

```

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

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

Note: When initialization and creation of the root aggregate is complete, the storage system reboots. You can continue with the configuration of node st-node02 while the initialization and creation of the root aggregate for node st-node01 is in progress.

Note: For more information about root aggregate and disk partitioning, please refer to the following ONTAP documentation on root-data partitioning.

While node st-node01 is initializing, begin the initializing procedures for node st-node02.

Initialize node st-node02

To initialize node st-node02, complete the following steps:

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. Allow the system to boot.

```
autoboot
```

3. Press Ctrl-C to enter the Boot menu.

Note: If ONTAP 9.12.1P2 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.12.1P2 is the version being booted, select option 8 and `y` to reboot the node. Then, continue with step 13.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter `n` to skip the reboot.
8. Enter the IP address, network mask, and default gateway for e0M in their respective places.

```
< st-node02_mgmt_ip> < st-node02_mgmt_mask> < st-node02_mgmt_gateway>
```

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

Note: This web server must be pingable.

```
<url_boot_software>
```

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

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} [y] y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y

Rebooting...
Files /cfcard/x86_64/freebsd/image2/VERSION and /var/VERSION differ
.
Terminated
Setting default boot image to image2...
done.
Uptime: 12m55s
System rebooting...
```

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

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

Note: When initialization and creation of root aggregate is complete, the storage system reboots.

Configure node st-node01 and create cluster

After the clean configuration and initialize all disks procedures are completed on the controller node, the node setup script appears when ONTAP 9.12.1P2 boots on the node for the first time. Proceed with the following steps when the node setup script wizards have started on both nodes.

Note: The NetApp ONTAP cluster can be configured using either ONTAP System Manager or via CLI after the basic network configuration information is provided for node st-node01, this documentation describes using the System Manager to complete the configuration.

1. Follow the prompts to setup node st-node01.

```
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: < st-node01_mgmt_ip>
Enter the node management interface netmask: < st-node01_mgmt_mask>
Enter the node management interface default gateway: < st-node01_mgmt_gateway>
A node management interface on port e0M with IP address < st-node01_mgmt_ip> has been created.

Use your web browser to complete cluster setup by accessing https://< st-node01_mgmt_ip>

Otherwise, press Enter to complete cluster setup using the command line interface:
```

2. Launch a web browser and complete the cluster setup by accessing https://< st-node01_mgmt_ip >

ONTAP System Manager

ONTAP 9.12.1 [Tips for initializing a storage system](#)

Health

2 healthy nodes were found.

AFF-A250

Initialize Storage System

STORAGE SYSTEM NAME

Enter cluster name

ADMINISTRATIVE PASSWORD

Enter new password

Confirm password

Networking

CLUSTER MANAGEMENT IP ADDRESS SUBNET MASK GATEWAY

IP Address Length IP Address

NODE SERIAL NUMBERS NODE MANAGEMENT IP ADDRESSES

791941000181 IP Address

792247000052 IP Address

☐ Use Domain Name Service (DNS)

3. Provide the required information.
 - a. Enter the cluster name and administrator password.
 - b. Complete the Networking information for the cluster and each node.
 - c. Leave the rest of the options and click on Submit to start the cluster setup.

ONTAP System Manager

2 healthy nodes were found.

AFF-A250

STORAGE SYSTEM NAME

g1434-a250

You will see this name when managing the storage system.

ADMINISTRATIVE PASSWORD

Networking

CLUSTER MANAGEMENT IP ADDRESS SUBNET MASK GATEWAY

172.22.29.15 255.255.255.0 172.22.29.1

NODE SERIAL NUMBERS NODE MANAGEMENT IP ADDRESSES

791941000181 172.22.29.16

792247000052 172.22.29.17

☐ Use Domain Name Service (DNS)

Others

☐ Use time services (NTP)

Submit

The nodes should be discovered automatically; if they are not, Refresh the browser page. By default, the cluster interfaces are created on all the new factory shipping storage controllers.

Note: If all the nodes are not discovered, then configure the cluster using the command line.

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

1. Click Submit.
2. A few minutes will pass while the cluster is configured. When prompted, login to ONTAP System Manager to continue the cluster configuration.
3. Click Prepare Storage to create data aggregates.

Note: You can use Ansible scripts at this point to configure the ONTAP Storage Configuration via Ansible.

Ansible ONTAP Storage Configuration - Part 1

1. Edit the following variable files to ensure proper ONTAP Storage variables are entered:
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/inventory
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/ontap
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/vars/ontap_main.yml
2. From /FlexPod-Express-Intersight/FlexPod-Express-Intersight/, run the Setup_ONTAP.yml Ansible playbook with the associated tag for this section:

```
ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_1
```

Use the -vvv tag to see detailed execution output log.

Cisco UCS C220 Standalone Rack Server Deployment Procedure

Cisco UCS C-series Standalone provides a high-performance, next-generation server system. It simplifies the system management and saves cost and is an ideal solution for a small-scale deployment.

The hardware and software components support Cisco's unified fabric, which runs multiple types of datacenter traffic over a single converged network adapter. It provides a high degree of workload agility and scalability.

The following section provides detailed procedures for configuring a Cisco UCS C-series Standalone in Intersight for use in the FlexPod Express configuration.

Perform initial Cisco UCS C-Series standalone server setup for Cisco Integrated Management Server

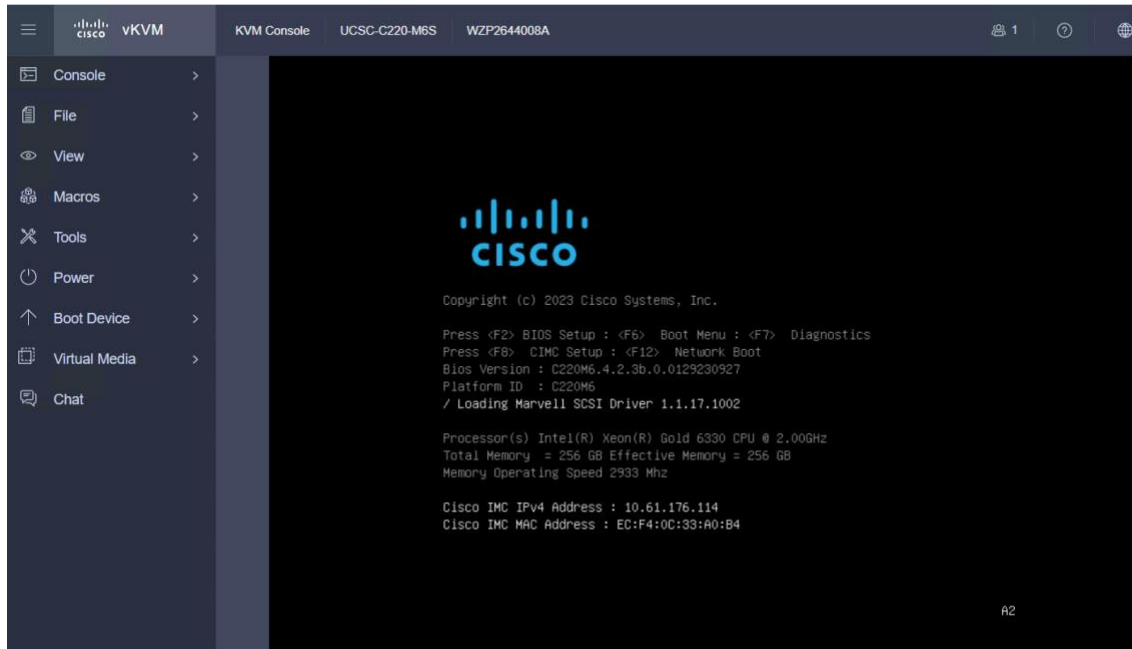
The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

Table 16) Information required to configure CIMC on server

Detail	Detail Value
CIMC IP address	<cimc_ip>
CIMC subnet mask	<cimc_netmask>
CIMC default gateway	<cimc_gateway>

All servers:

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM 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 CIMC configuration.



3. In the CIMC configuration utility, set the following options:

- NIC Properties
 - Network interface card (NIC) mode:
 - Dedicated [X]
 - NIC redundancy
 - None: [X]
 - VLAN (Advanced): Leave cleared to disable VLAN tagging.

1. IP (Basic):

- IPV4: [X]
- DHCP enabled: []
- CIMC IP: <<cimc_ip>>
- Prefix/Subnet: <<cimc_netmask>>
- Gateway: <<cimc_gateway>>

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated: [X]                         None: [X]
Shared LOM: [ ]                       Active-standby: [ ]
Cisco Card:                            Active-active: [ ]
Riser1: [ ]                           VLAN (Advanced)
Riser3: [ ]                           VLAN enabled: [ ]
MLom: [ ]                             VLAN ID: 1
Shared LOM Ext: [ ]                   Priority: 0
IP (Basic)
IPv4: [X]                             IPv6: [ ]
DHCP enabled: [ ]
CIMC IP: 10.61.176.114
Prefix/Subnet: 255.255.255.0
Gateway: 10.61.176.1
Pref DNS Server: 10.61.176.251
Smart Access USB
Enabled [ ]
*****
<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: C220-W2P2644008A
Dynamic DNS: [ ]
DDNS Domain:
FactoryDefaults
Factory Default: [ ]
Default User(Admin)
Enter New Default User password:
Re-Enter New Default User password:
Port Properties
Auto Negotiation: [X]
Admin Mode      Operation Mode
Speed[1000/100/10Mbps]: Auto      1000
Duplex mode[half/full]: Auto      full
Port Profiles
Reset: [ ]
Name:
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F2>Previous Page

```

Setup the Cisco Intersight Account and License

Setup Intersight account

- Go to <https://intersight.com> and click Create an account.
- Read and accept the license agreement. Click Next.
- Provide an Account Name and click Create.
- On successful creation of the Intersight account, following page will be displayed.



Set up Cisco Intersight Licensing

Note: When setting up a new Cisco Intersight account (as explained in this document), the account needs to be enabled for Cisco Smart Software Licensing.

- Log into the Cisco Smart Licensing portal: cisco-smart-licensing
- Verify that the correct virtual account is selected.
- Under Inventory > General, generate a new token for product registration.
- Copy this newly created token.
- In Cisco Intersight, click Select Service > System, then click Administration > Licensing.
- Under Actions, click Register.
- Enter the copied token from the Cisco Smart Licensing portal. Click Next.
- Drop-down the pre-selected Default Tier * and select the license type (for example, Essentials).
- Select Move All Servers to Default Tier.
- Click Register, then click Register again.
- When the registration is successful (takes a few minutes), the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

The Cisco Intersight platform uses a subscription-based license with multiple tiers. Customers can purchase a subscription duration of one, three, or five years and choose the required Cisco UCS server volume tier for the selected subscription duration. Each Cisco endpoint automatically includes a Cisco Intersight Base license at no additional cost when customers access the Cisco Intersight portal and claim a device. Customers can purchase any of the following higher-tier Cisco Intersight licenses using the Cisco ordering tool.

Servers in the Cisco Intersight managed mode require at least the Essentials license. For more information about the features provided in the various licensing tiers, see https://intersight.com/help/getting_started#licensing_requirements.

Essentials License is used in the current project since the UCS servers alone are claimed into Intersight for cost optimization.

Set Up Cisco Intersight Resource Group

In this procedure, a Cisco Intersight resource group is created where resources such as targets will be logically grouped. In this deployment, a single resource group is created to host all the resources, but customers can choose to create multiple resource groups for granular control of the resources.

- Log into Cisco Intersight
- At the top, select System. On the left, click Settings (the gear icon)
- Click Resource Groups in the middle panel.
- Click + Create Resource Group in the top-right corner.
- Provide a name for the Resource Group (for example, C220G13-RTP).
- Under Memberships, select Custom.
- Click Create.

The screenshot shows the 'Create Resource Group' page in the Cisco Intersight portal. At the top, there is a breadcrumb '← Resource Groups' and the title 'Create Resource Group'. Below the title is a subtitle 'Create a Resource Group to manage and access the targets.' The main form is divided into two sections: 'General' and 'Memberships'. In the 'General' section, there is a 'Name *' field with the value 'C220G13-RTP' and a 'Description' field. In the 'Memberships' section, there are two radio buttons: 'Custom' (which is selected) and 'All'. At the bottom of the form, there is a blue information bar with a question mark icon and the text 'The selected targets will be part of the Resource Group created.'

Set Up Cisco Intersight Organization

In this step, an Intersight organization is created where all Cisco Intersight managed mode configurations including policies are defined.

- Log into the Cisco Intersight portal.
- At the top, select System. On the left, click Settings (the gear icon).
- Click Organizations in the middle panel.
- Click + Create Organization in the top-right corner.
- Provide a name for the organization (for example, FlexPodExpress).
- Select the Resource Group created in the last step (for example, C220G13-RTP).
- Click Create.

← Organizations

Create Organization

General

Name *
FlexpodExpress
Description

Resource Groups

Select the Resource Groups to be associated with the Organization. Organization created will provide access to the resources in the selected Resource Groups.

2 items found 10 per page 1 of 1

Name	Used Organizations	Description
<input type="checkbox"/> default	default	The Default Resource Grou...
<input type="checkbox"/> C220G13-RTP	-	FlexPod-Express

Cancel
Create

Claim a Cisco UCS C220 Standalone Server in the Cisco Intersight Platform

After setting up the Cisco UCS C-series Standalone Server for Cisco Intersight Managed Mode, Servers can be claimed to a new or an existing Cisco Intersight account.

The Device ID and Claim ID information is obtained from the CIMC management controller under Admin -> Device Connector.

Cisco Intersight
System
Search

Settings
Admin
Targets
Software Repository
Tech Support Bundles
Audit Logs
Sessions
Licensing

← Targets

Claim a New Target

Select Target Type

Filters

☒ Available for Claiming

Categories

☒ All
☐ Cloud
☐ Compute / Fabric
☐ Unmanaged

☒
Cisco UCS Server (Standalone)

☐
Cisco UCS Domain (Intersight Managed)

☐
HPE OneView

☐
Cisco UCS Domain (UCSM Managed)

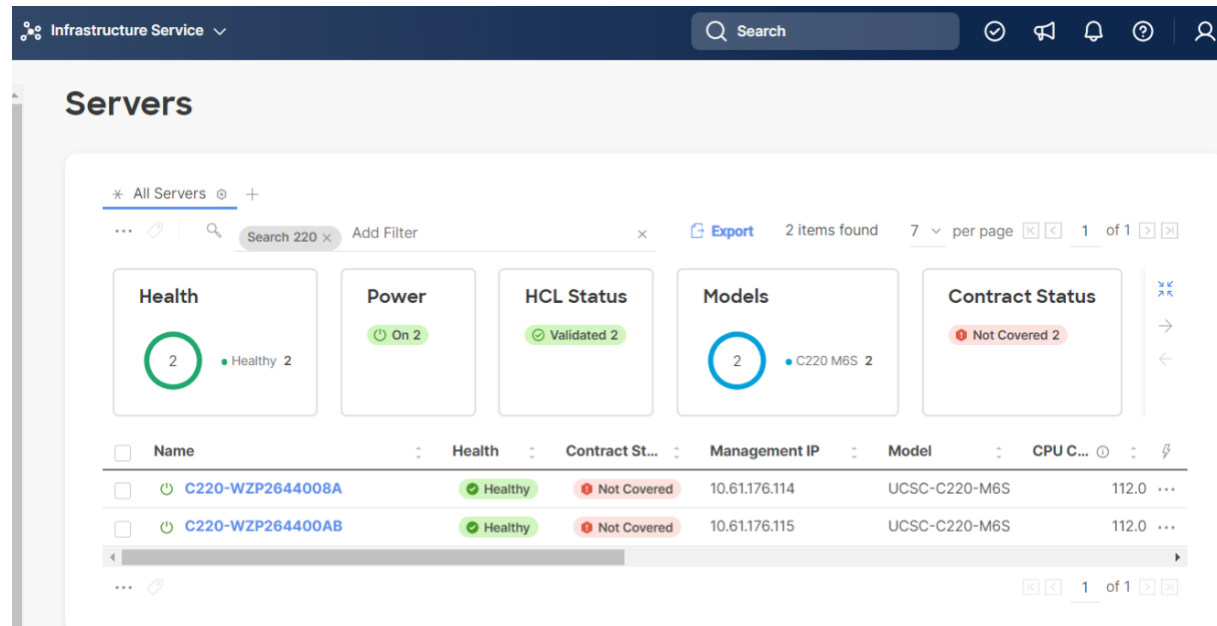
☐
Cisco UCS C890

☐
Redfish Server


Platform Services

Cancel
Start

When a Cisco UCS Server is successfully added to the Cisco Intersight platform, all future configuration steps are completed from the Cisco Intersight portal.




Verify on CIMC, that the server is successfully claimed into Intersight.



Upgrade Cisco UCS Server using Cisco Intersight (Optional)

This document assumes the use of Cisco UCS Firmware Software version 4.2(3d). To upgrade the Cisco UCS Server Firmware software, see [Intersight Configuration Guides](#)



Ansible Cisco UCS IMM Configuration

To configure the Cisco UCS from the Ansible management workstation, follow the steps in this procedure. The `group_vars/ucs.yml` file contains two important variables:

- `server_cpu_type` – Intel or AMD – the type of CPU in the server
- `vic_type` – 4G or 5G – 5G is the latest 15000-series VICs while 4G is all previous generations

To execute the playbooks against your Intersight account, you need to create an API key and save a `SecretKey.txt` file from your Cisco Intersight account.

- In Cisco Intersight, select System > Settings > API > API Keys.
- Click Generate API Key.
- Under Generate API Key, enter a Description (for example, API Key for Ansible) and select API key for OpenAPI schema version 2. Click Generate.


Generate API Key

Description
API Key for Ansible


API Key Purpose

☒ API key for OpenAPI schema version 2
☐ API key for OpenAPI schema version 3 (This is a feature in preview and for SDK developer use only)

Close
Generate

In the Generate API Key window, click the upper  icon to copy the API Key ID to the clipboard. Paste this key into the `api_key_id` variable in the `FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/ucs.yml` variable file and save it.

Using an editor, open the FlexPod-Express-Intersight/FlexPod-Express-Intersight/SecretKey.txt file and clear

all text from the file. Then click the lower  icon in the Generate API Key window and paste the Secret Key into the SecretKey.txt file and save it.

Edit the following variable files to ensure proper UCS variables are entered:

- FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/ucs.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/UCS-IMM/create_server_policies/defaults/main.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/UCS-IMM/create_server_profile_template/defaults/main.yml

The /FlexPod-Express-Intersight/FlexPod-Express-Intersight directory contains two Ansible playbooks to set up Cisco UCS IMM server profile templates: Create_IMM_Server_Policies.yml and Create_IMM_Server_Profile_Templates.yml. Both playbooks are designed to be run more than once with different combinations of server_cpu_type and vic_type. It is important that when Create_IMM_Server_Policies.yml is run that Create_IMM_Server_Profiles_Templates.yml is run before changing the server_cpu_type and vic_type variables. Many of the policies and templates will be assigned unique names according to these variables. Also, since the UCS-IMM Ansible playbooks connect to the Cisco Intersight API website instead of hardware components, that the use of the inventory file is not needed. To set up the Cisco Intersight IMM policies, and server profile templates, execute the following:

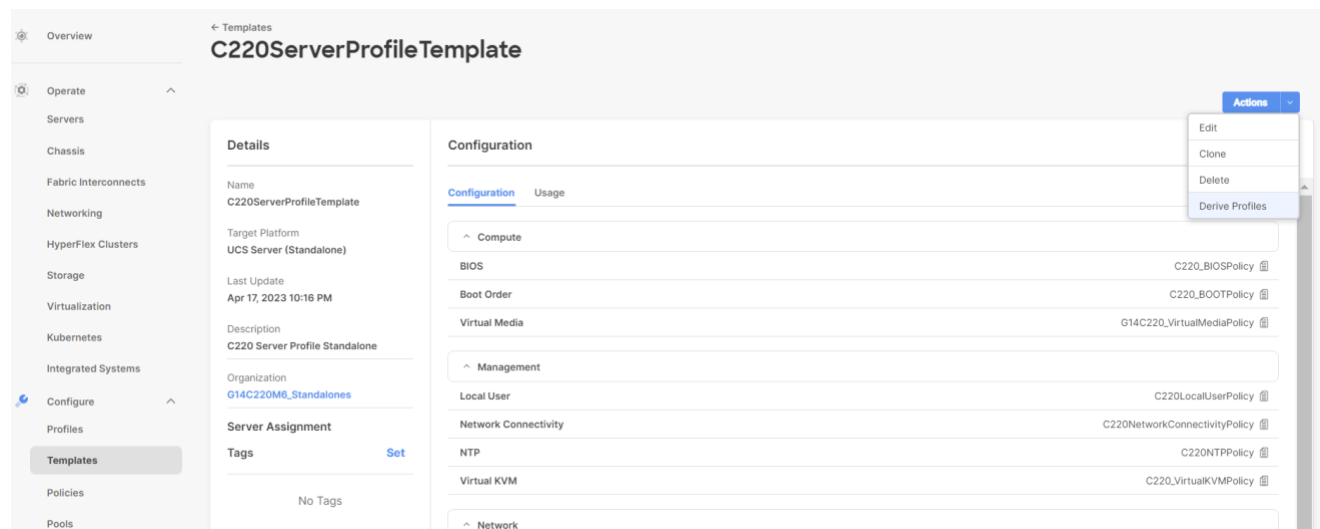
```
ansible-playbook ./Setup_IMM_Server_Policies.yml
ansible-playbook ./Setup_IMM_Server_Profile_Templates.yml
```

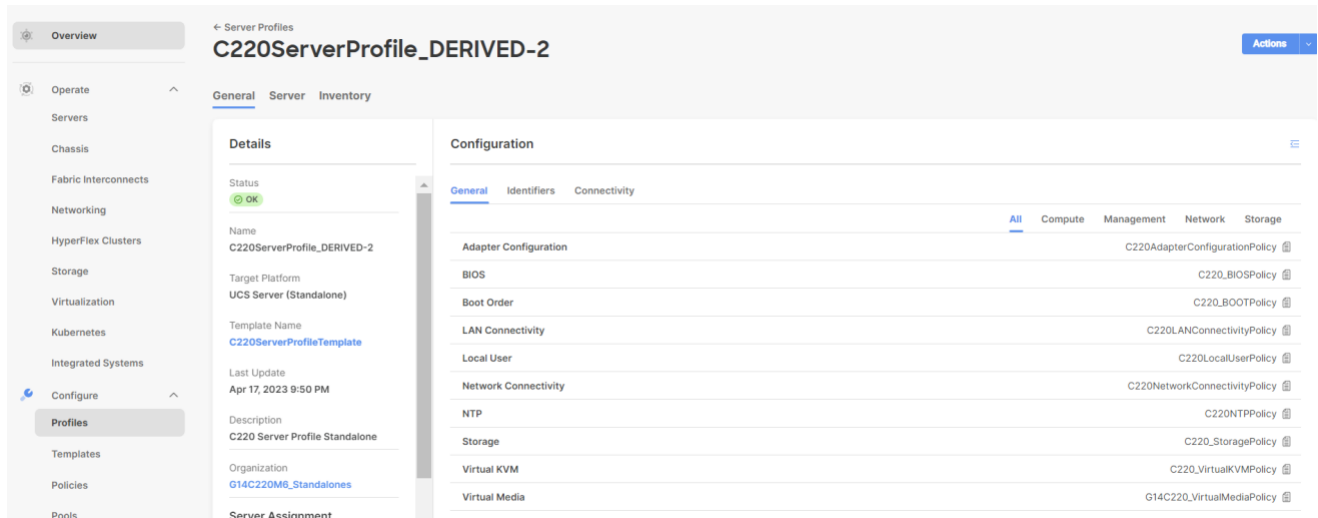
Note: Server Profiles will be generated from the Server Profile Templates and assigned to servers after the Cisco UCS IMM Manual Configuration

Cisco UCS IMM Setup Completion

Complete the following procedures whether performing an Ansible configuration or a Manual configuration of the FlexPod.

Derive and Deploy Server Profiles from the Cisco Intersight Server Profile Template





NOTE: Since this is a local boot, ONTAP Boot Storage setup is not required.

VMware vSphere 8.0 Deployment Procedure

This section provides detailed procedures for installing VMware ESXi 8.0 in a FlexPod Express configuration. After the procedures are completed, two Locally booted ESXi hosts are provisioned.

Note: VMware recommends a minimum cluster size of three servers. For this validation, the minimum supported cluster size of two servers is used. You can optionally deploy additional servers based on your solution requirements.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in KVM console and virtual media features in Cisco Intersight to map remote installation media to individual servers.

Download Cisco custom image for ESXi 8.0

If the VMware ESXi custom image has not been downloaded, complete the following steps to complete the download:

1. Go to the following link: [VMware vSphere Hypervisor \(ESXi\) 8.0](#)
2. You need a user ID and password on [vmware.com](#) to download this software.
3. Download the VMware-ESXi-8.0-20513097-Custom-Cisco-4.2.3-b.iso file.

Launch KVM console for the server

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system through remote media. Via Intersight, launch the KVM Console for each of the servers.

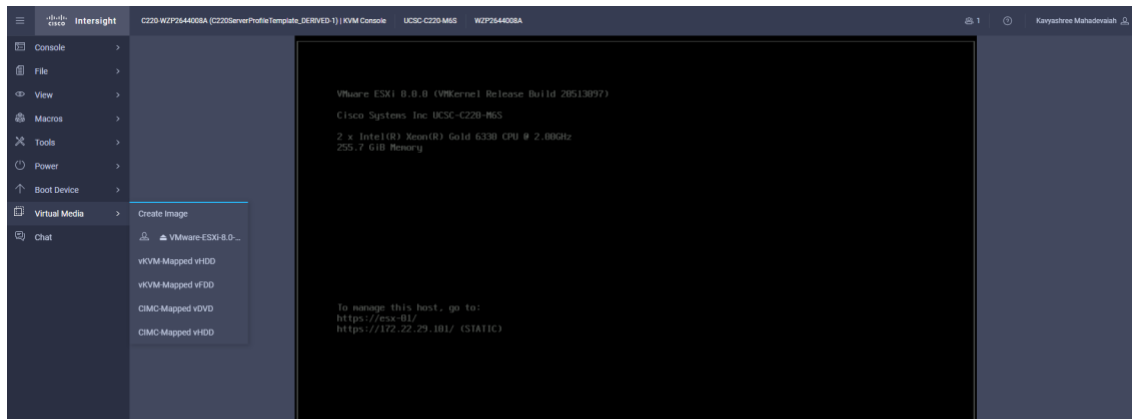
Set up VMware ESXi Installation

ESXi Hosts esxi-01 and esxi-02

Skip this section if you are using vMedia policies; the ISO file will already be connected to KVM.

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

1. In the KVM window, click Virtual Media.
2. Click vKVM mapped DVD.
3. Browse to the ESXi installer ISO image file and click Open.
4. Click Map Device.
5. Check the screen to monitor the server boot.



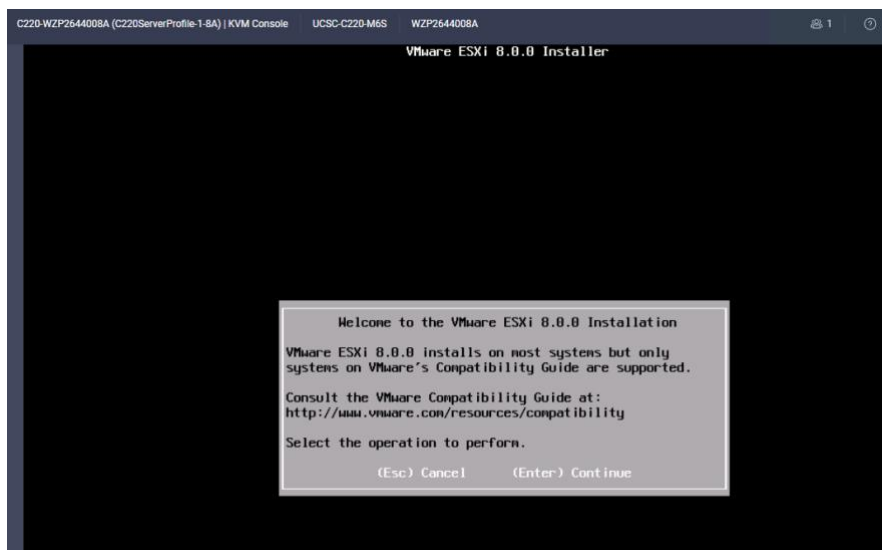
Install ESXi

ESXi Hosts esxi-01 and esxi-02

To install VMware ESXi to the local disk of the hosts, complete the following steps on each host:

1. Boot the server by selecting Boot Server and clicking OK. Then click OK again.
2. After reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the local disk as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, select the Virtual Media tab and clear the P mark next to the ESXi installation media. Click Yes.

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



10. After the installation is complete, press Enter to reboot the server.

Set up management networking for ESXi Hosts

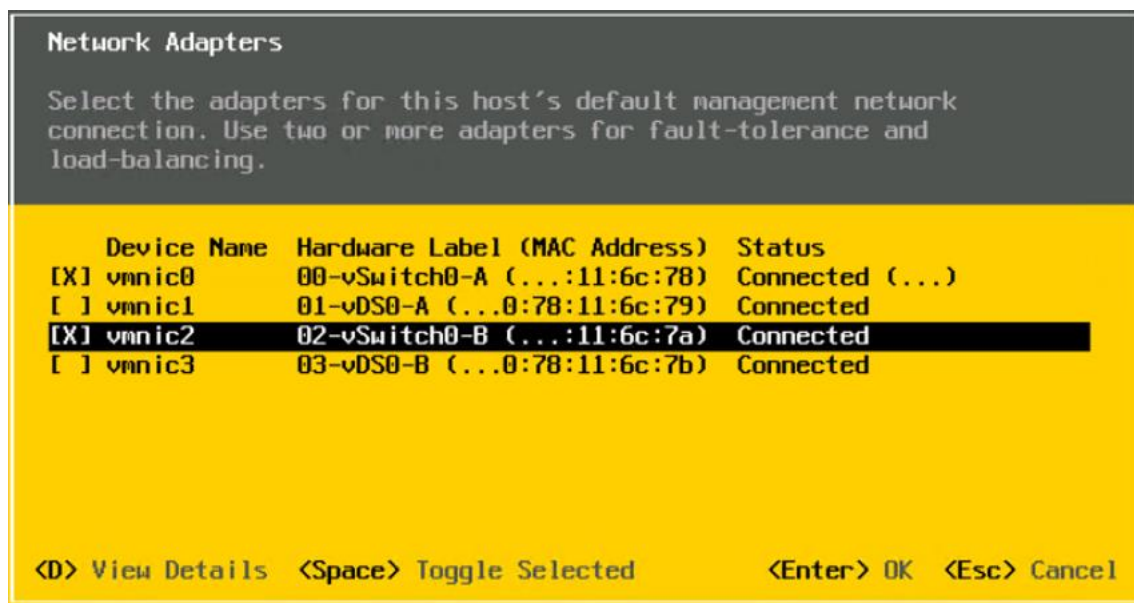
Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi esxi-01 and esxi-02

To configure each ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.
5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.

Note: Verify that the numbers in the Hardware Label field match the vmnic numbers in the DeviceName field based on CDN.



9. Use the Space bar to also select the vmnic that has the Hardware Label 02-vSwitch0-B.
10. Press Enter.

Note: In the UCS Configuration portion of this document, the IB-MGMT VLAN was set as the native VLAN on the 00-vSwitch0-A and 02-vSwitch0-B Standalone vNICs. Because of this, the IB-MGMT VLAN should not be set here and should remain Not set.

11. Select IPv4 Configuration and press Enter.

Note: When using DHCP to set the ESXi host networking configuration, setting up a manual IP address is not required.

12. Select the Set Static IPv4 Address and Network Configuration option by using the space bar.
13. Enter the IP address for managing the first ESXi host.

14. Enter the subnet mask for the first ESXi host.
15. Enter the default gateway for the first ESXi host.
16. Press Enter to accept the changes to the IP configuration.
17. Select the DNS Configuration option and press Enter.

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

18. Enter the IP address of the primary DNS server.
19. Optional: Enter the IP address of the secondary DNS server.
20. Enter FQDN for the first ESXi host.
21. Press Enter to accept the changes to the DNS configuration.
22. Press Esc to exit the Configure Management Network menu.
23. Select Test Management Network to verify that the management network is set up correctly and press Enter.
24. Press Enter to run the test, press Enter again once the test has completed, review environment if there is a failure.
25. Select the Configure Management Network again and press Enter.
26. Select the IPv6 Configuration option and press Enter.
27. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
28. Press Esc to exit the Configure Management Network submenu.
29. Press Y to confirm the changes and reboot the ESXi host

Reset VMware ESXi host VMkernel port vmk0 MAC address.

NOTE: On VMware ESXi, the MAC address of the vmk0 VMkernel port is by default the same as the MAC address of vmnic0 or your Eth0 vNIC. This is not a problem behind a pair of FIs because the FIs do not exchange MAC information. When directly connected to switches, this can cause a problem if the vSwitch attempts to send packets from vmk0 over vmnic0 or your Eth2 vNIC because the same MAC address will be seen on both switch A (vmnic0) and switch B (vmk0). To resolve this issue of the same MAC on both switches, delete vmk0, then add it back. When you add it back, it will have a random VMware MAC address.

To reset the MAC address of vmk0 to a random VMware-assigned MAC address, complete the following steps:

1. From the ESXi console menu main screen, press Ctrl-Alt-F1 to access the VMware console CLI. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of static macros.
2. Log in as root.
3. Enter **esxcfg-vmknics -l** to get a detailed listing of interface vmk0. vmk0 should be a part of the Management Network port group. Note the IP address and network mask of vmk0.
4. To ensure random VMware MAC address is assigned for vmk0 port, remove vmk0 and add vmk0 back again using the following command:

```
esxcfg-vmknics -d "Management Network"; esxcfg-vmknics -a -i <var_vmk0_ip> -n <var_vmk0_netmask>
"Management Network"
```

5. Verify that vmk0 has been added again with a random MAC address:

```
esxcfg-vmknics -l
```

6. Tag vmk0 as the management interface:

```
esxcli network ip interface tag add -i vmk0 -t Management
```

7. When vmk0 was re-added if a message popped up saying vmk1 was marked as the management interface, remove it by the following command:

```
esxcli network ip interface tag remove -i vmk1 -t Management
```


8. Enter **exit** to log out of the command line interface.
9. Press Ctrl-Alt-F2 to return to the ESXi console menu interface.

FlexPod VMware ESXi Ansible Configuration

Use Ansible to Configure All VMware ESXi Hosts from the Management Workstation

1. Edit the following variable files to ensure proper VMware variables are entered:
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/inventory
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/VMware/ESXihosts/defaults/main.yml
2. From /FlexPod-Express-Intersight/FlexPod-Express-Intersight, run the Setup_ESXi.yml Ansible playbooks:

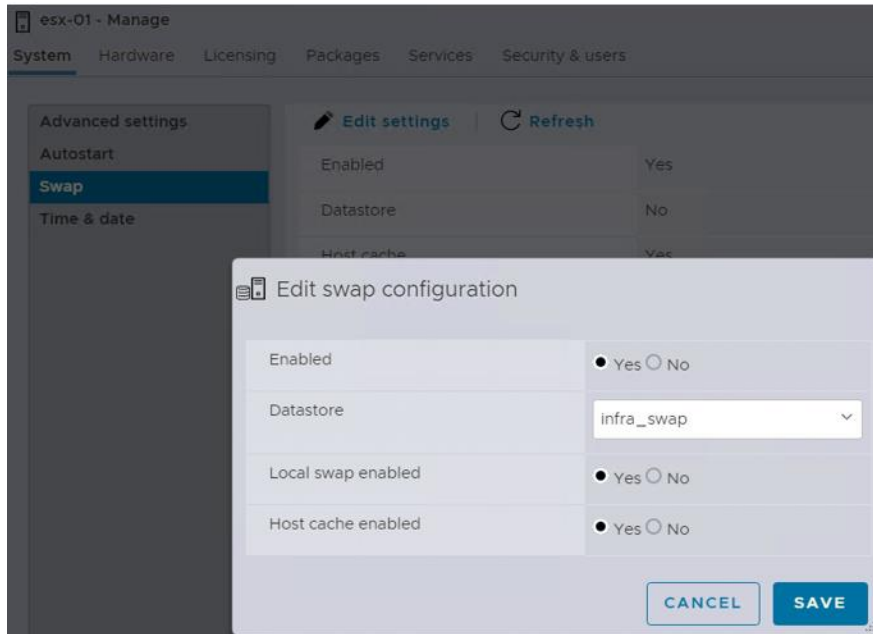
```
ansible-playbook ./Setup_ESXi.yml -i inventory
```

Configure ESXi Host Swap

ESXi hosts esxi-01 and esxi-02

To configure host swap on the ESXi hosts, follow these steps on each host:

1. Click Manage in the left navigation pane.
2. In the center pane, select Swap under the System tab.
3. Click Edit Settings. Select `infra_swap` from the Datastore options.



4. Click Save

VMware vCenter Server 8.0 Deployment Procedure

This section provides detailed procedures for installing VMware vCenter Server 8.0 in a FlexPod Express configuration.

Note: FlexPod Express uses the VMware vCenter Server Appliance (VCSA).


Install VMware vCenter server appliance.

To install VCSA, complete the following steps:

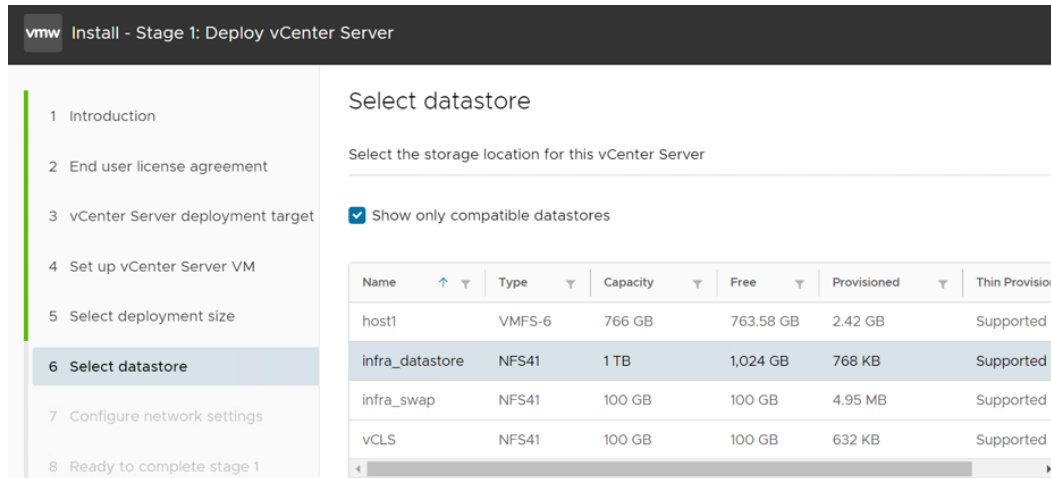
1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.
2. Download the VCSA from the VMware site.
3. Mount the ISO image on your management workstation.
4. Navigate to the installer appropriate for your environment.
5. For installing from Windows, navigate to the `vcsa-ui-installer > win32` directory and double-click `installer.exe` to start the installation. For installing from Linux, navigate to `vcsa-ui-installer > lin64` and run the installer to start the installation.

Note: Depending on the platform you use to install VCSA, the GUI screenshots might look slightly different.

6. Click Install.
7. Click Next on the Introduction page.
8. Accept the EULA and click Next.
9. Specify the vCenter server deployment target host, username, and password information. For example, enter the host name or IP address of the first ESXi host, username (root), and password.
10. Click Next. Click Yes to accept the certificate warning and continue.
11. Specify the vCenter VM name and root password.
12. Click Next.



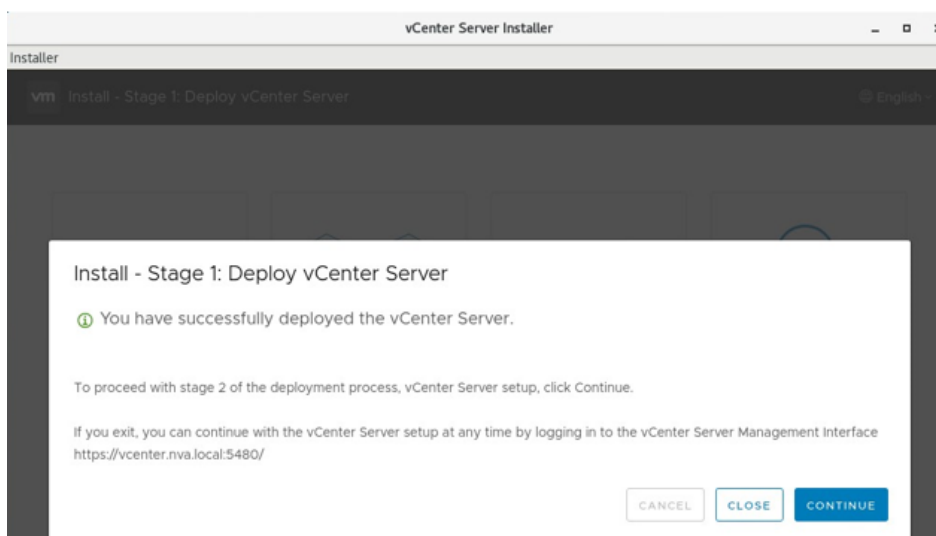
13. Select the deployment size and storage size that are suitable for your deployment. For example, choose Tiny and Default.
14. Click Next.



15. Select the storage location for the vCenter. For example, click to select infra_datastore.
16. Click Next.
17. Enter the vCenter network configuration information and click Next.
 - VM Network is selected automatically for Network when deploying vCenter to the first ESXi host.
 - Select IP version.
 - Select IP assignment method.
 - Enter the FQDN to be used for the vCenter.
 - Enter the IP address.
 - Enter the subnet mask. Enter the default gateway.
 - Enter the DNS server.
18. Click Next.
19. Review all the settings and click Finish.

Note: The VCSA installation process takes several minutes.

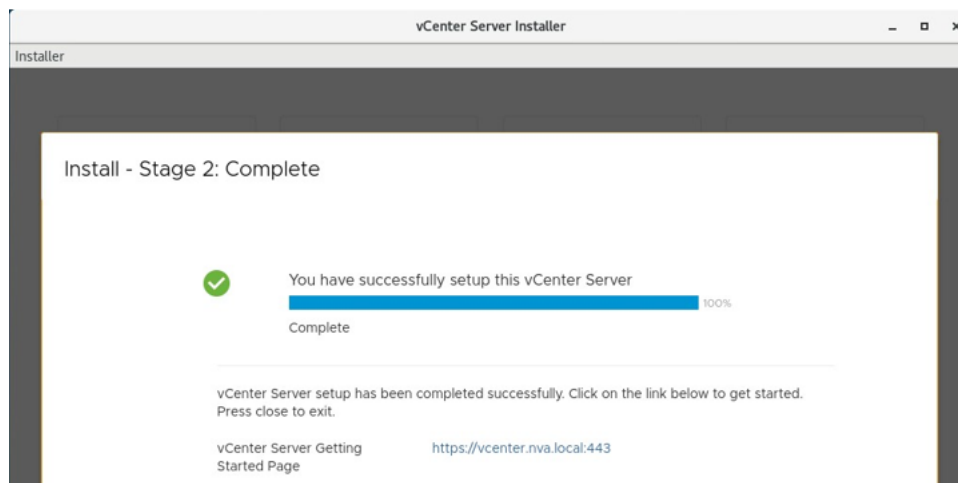
20. After install stage 1 completes, a message appears stating that it has completed.



21. Click Continue to begin stage 2 configuration.

22. On the Stage 2 Introduction page, click Next.
23. In the Appliance Configuration, configure these settings:
 - a. Time Synchronization Mode: Synchronize time with NTP servers.
 - b. NTP Servers: <nexus-a-ntp-ip>, <nexus-b-ntp-ip>
 - c. SSH access: Enabled.
24. Configure the SSO domain name and administrator password.

Note: Record these values for your reference, especially if you deviate from the vsphere.local domain name.
25. Click Next.
26. Join the VMware Customer Experience Program if desired. Click Next.
27. Review your configuration settings. Click Finish.



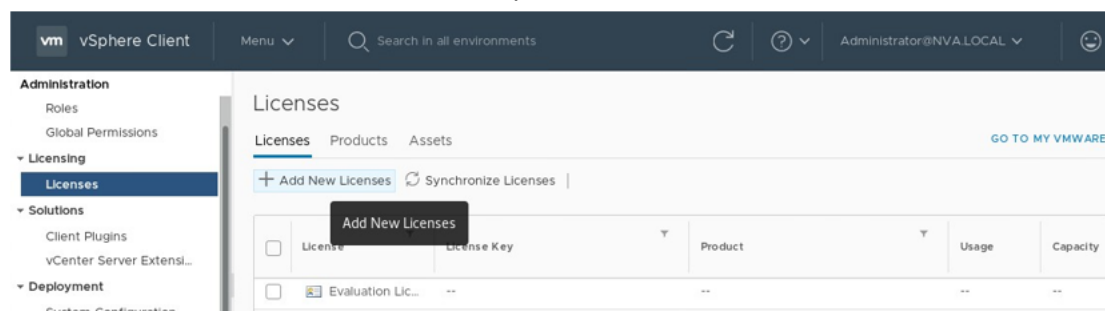
Note: The link that the installer provides to access vCenter Server is clickable.

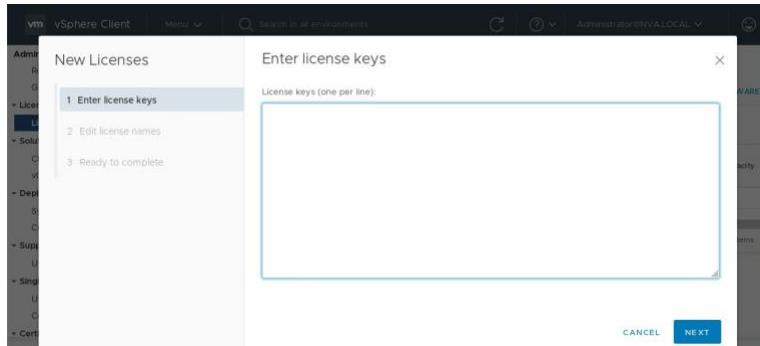
28. Click CLOSE.

Add and assign vCenter and vSphere licenses

To add and assign the vCenter and vSphere licenses, follow these steps:

1. Log into the vCenter server.
2. Under Menu, select Administration.
3. Under the Licensing group on the left pane, click Licenses.
4. Click Add New Licenses in the center pane.





5. Type in the license keys, one per line, in the dialog box and click Next.
6. Edit License Name, if needed, and click Next.
7. Click Finish to complete entering licenses.
8. Right-click the vCenter server under Hosts and Clusters and select Assign License.
9. Select the vCenter license and click OK.
10. Right-click an ESXi host under Hosts and Clusters and select Assign License.
11. Select the vSphere license and click OK to assign.
12. Repeat steps 9 and 10 for all the ESXi hosts in the cluster.

vCenter and ESXi Ansible Setup

Configure the VMware vCenter and the three management ESXi hosts

Edit the following variable files to ensure proper VMware variables are entered:

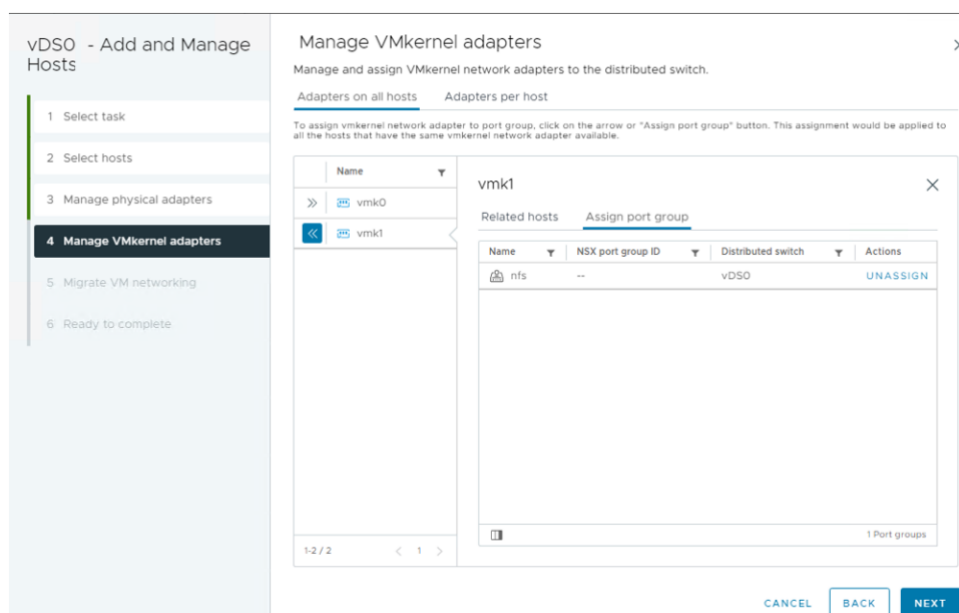
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/inventory
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/VMware/ESXlpostvC/defaults/main.yml
- FlexPod-Express-Intersight/FlexPod-Express-Intersight/roles/VMware/ESXlpostvCNVMe/defaults/main.yml

From /FlexPod-Express-Intersight/FlexPod-Express-Intersight, run the Setup_vCenter.yml Ansible playbook:

```
ansible-playbook ./Setup_vCenter.yml -i inventory
```

Note: After the playbook run is complete, complete the following manual steps to complete vCenter setup.

1. Expand the FlexPod Express datacenter, right click the newly added cluster and select Settings.
2. In the center pane, go to Configuration > General in the list located on the left and select EDIT located on the right of General to specify the swap file location.
3. Select Datastore Specified by Host Option.
4. Click OK.
5. Right-click the vDS and click Add and Manage Hosts.
6. Make sure Manage host networking is selected and click NEXT.
7. Click SELECT ALL to select all ESXi hosts. Click NEXT.
8. Click NEXT.
9. To the right of vmk1, click ASSIGN PORT GROUP and assign nfs port group to migrate NFS VMkernel ports to vDS0 and click NEXT.



Note: vmk1 was created using nfs port group on standard switch to mount data stores.

1. Do not migrate any virtual machine networking ports. Click NEXT.
2. Click FINISH to complete adding the ESXi host to the vDS.
3. Add the other hosts using the same steps as described above.

Note: Verify the port channel status on both Nexus switches. At this stage, the port channel should be up.

Finalize the vCenter and ESXi Setup

This procedure enables you to finalize the VMware installation.

VMware ESXi 8.0 TPM Attestation

Note: If your Cisco UCS servers have Trusted Platform Module (TPM) 2.0 modules installed, the TPM can provide assurance that ESXi has booted with UEFI Secure Boot enabled and using only digitally signed code. In the Cisco UCS Configuration section of this document, UEFI secure boot was enabled in the boot order policy. A server can boot with UEFI Secure Boot with or without a TPM 2.0 module. If it has a TPM, VMware vCenter can attest that the server booted with UEFI Secure Boot.

To verify the VMware ESXi 8.0 TPM Attestation, follow these steps:

For Cisco UCS servers that have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client.

1. In the vCenter HTML5 Interface, under Hosts and Clusters select the cluster.
2. In the center pane, click the Monitor tab.
3. Click Monitor > Security. Review the host's status in the Attestation column.

Note: It may be necessary to disconnect and reconnect or reboot a host from vCenter to get it to pass attestation the first time.

Avoiding Boot Failure When UEFI Secure Booted Server Profiles are Moved

Typically, hosts in FlexPod Datacenter are configured for boot from SAN. Cisco UCS supports stateless compute where a server profile can be moved from one blade or compute node to another seamlessly. When a server profile is moved from one server to another server with the following conditions, the ESXi host runs into PSOD and ESXi will fail to boot:

TPM present in the node (Cisco UCS M5 and M6 family servers)

- Host installed with ESXi 7.0 U2 or above

- Boot mode is UEFI Secure
- Error message: Unable to restore system configuration. A security violation was detected.
<https://via.vmw.com/security-violation>

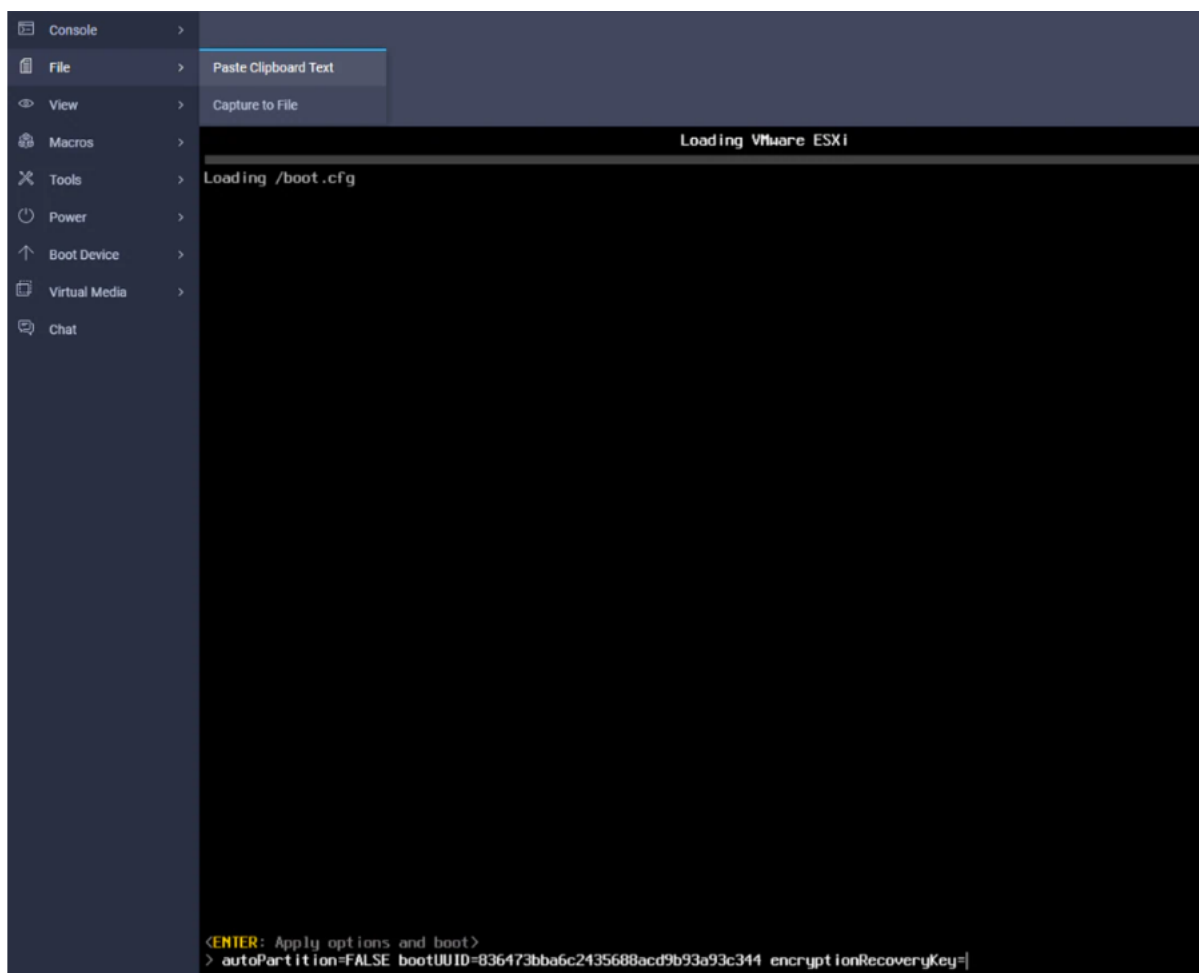
1. Log into the host using SSH.
2. Gather the recovery key using this command:

```
[root]@esxi-01:~] esxcli system settings encryption recovery list
```

Recovery ID	Key

{74AC4D68-FE47-491F-B529-6355D4AAF52C}	529012-402326-326163-088960-184364-097014-312164-590080-407316-660658-634787-601062-601426-263837-330828-197047

3. Store the keys from all hosts in a safe location.
4. After associating the Server Profile to the new compute-node or blade, stop the ESXi boot sequence by pressing Shift + O when you see the ESXi boot screen.



5. Add the recovery key using the following boot option: encryptionRecoveryKey=recovery_key. Press Enter to continue the boot process.
6. To persist the change, enter the following command at the VMware ESXi ssh command prompt:

```
/sbin/auto-backup.sh
```

Note: For more information, refer to <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-security/GUID-10F7022C-DBE1-47A2-BD86-3840C6955057.html>

Ansible ONTAP Storage Configuration Part 2

Configure the ONTAP NVMe setup and finalize ONTAP storage using Ansible

1. Edit the following variable files to ensure proper variables are entered:
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/group_vars/all.yml
 - FlexPod-Express-Intersight/FlexPod-Express-Intersight/vars/ontap_main.yml

Note: Update the “nvme_namespaces” and “nvme_subsystem” variables in vars/ontap_main.yml file. Add the NQNs from each ESXi host to the corresponding variable “nvme_nqn” in group_vars/all.yml file. The NVMe namespace will be shared by all the hosts in the nvme subsystem in this solution

Note: The ONTAP NVMe setup is only required for NVMe/TCP configurations.

2. From /FlexPod-Express-Intersight/FlexPod-Express-Intersight, invoke the ansible scripts for this section using the following command:

```
ansible-playbook ./Setup_ONTAP.yml -t ontap_config_part_2
```

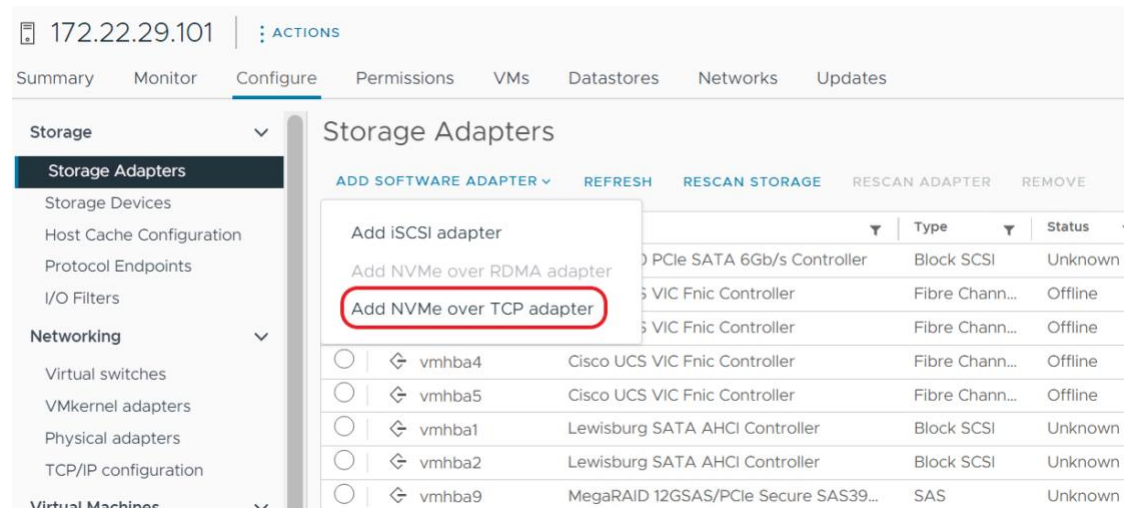
Note: After the playbook run is complete, Reboot each ESXi host.

Note: Since the ESXi hosts are not configured with the NVMe controllers yet, we will not see any output for the `esxcli nvme controller list` command. Follow the next section to configure ESXi hosts for NVMe over TCP. We cannot successfully add NVMe controller unless a namespace is mapped to the ESXi host NQN.

Configure ESXi Host NVMe over TCP Datastore

1. Select Hosts and Clusters and select the first ESXi host. In the center pane, select the Configure tab.
2. In the list under Storage, select Storage Adapters.

Select ADD SOFTWARE ADAPTER and click on Add NVMe over TCP adapter



3. In the Add Software NVMe over TCP adapter window, ensure that vmnic1 Network Adapter is selected and click OK

Add Software NVMe over TCP adapter 172.22.29.101 ×

Enable software NVMe adapter on the selected physical network adapter.

Physical Network Adapter

vmnic0/nenic ▼

vmnic0/nenic

vmnic1/nenic

vmnic2/nenic

vmnic3/nenic

CANCEL

OK

- Again, select ADD SOFTWARE ADAPTER and click on Add NVMe over TCP adapter
- In the Add Software NVMe over TCP adapter window, ensure that vmnic3 Network Adapter is selected and click OK.
- Select the first VMware NVMe over TCP Storage Adapter added (for example, vmhba64). In the middle of the window, select the Controllers tab. Click ADD CONTROLLER.
- Enter the IP address of nvme-tcp-lif-01a and click DISCOVER CONTROLLERS. Select the two controllers in the Infra-NVMe-TCP-A subnet and click OK. The two controllers should now appear under the Controllers tab.

Add controller vmhba64 ×

Automatically

Manually

Host NQN

nqn.2014-08.com.vmware:nvme:esx-01

📋 COPY

IP

172.22.33.11

☐ Central discovery controller

Enter IPv4 / IPv6 address

Port Number

Range more from 0

Digest parameter

☐ Header digest

☐ Data digest

DISCOVER CONTROLLERS

Select which controller to connect

<input type="checkbox"/>	Id	Subsystem NQN	Transport Type	IP	Port Number
<input type="checkbox"/>	65535	nqn.1992-08.com.netapp:s...	nvm	172.22.34.12	4420
<input checked="" type="checkbox"/>	65535	nqn.1992-08.com.netapp:s...	nvm	172.22.33.12	4420
<input type="checkbox"/>	65535	nqn.1992-08.com.netapp:s...	nvm	172.22.34.11	4420
<input checked="" type="checkbox"/>	65535	nqn.1992-08.com.netapp:s...	nvm	172.22.33.11	4420

8. Select the second VMware NVMe over TCP Storage Adapter added (for example, vmhba65). In the middle of the window, select the Controllers tab. Click ADD CONTROLLER.
9. Enter the IP address of nvme-tcp-lif-01b and click DISCOVER CONTROLLERS. Select the two controllers in the Infra-NVMe-TCP-B subnet and click OK. The two controllers should now appear under the Controllers tab.
10. Repeat steps 3-10 for other ESXi host.
11. Verify that the NetApp ONTAP target NVMe controllers are properly discovered on the ESXi Host.

```
[root@esxi-01:~] esxcli nvme controller list
Name
Controller Number  Adapter  Transport Type  Is Online  Is VVOL
-----
nqn.1992-08.com.netapp:sn.5f7833cfca3511edb03cd039ea4099f3:subsystem.fp-exsi-
hosts#vmhba64#172.22.33.12:4420          256  vmhba64  TCP          true      false
nqn.1992-08.com.netapp:sn.5f7833cfca3511edb03cd039ea4099f3:subsystem.fp-exsi-
hosts#vmhba64#172.22.33.11:4420          257  vmhba64  TCP          true      false
nqn.1992-08.com.netapp:sn.5f7833cfca3511edb03cd039ea4099f3:subsystem.fp-exsi-
hosts#vmhba65#172.22.34.12:4420          259  vmhba65  TCP          true      false
nqn.1992-08.com.netapp:sn.5f7833cfca3511edb03cd039ea4099f3:subsystem.fp-exsi-
hosts#vmhba65#172.22.34.11:4420          260  vmhba65  TCP          true      false
```

12. For adding NVMe datastore, Right-click the host under Hosts and Clusters and select Storage > New Datastore. Leave VMFS selected and click NEXT.
13. Name the datastore (for example, nvme_datastore) and select the NVMe Disk. Click NEXT.

New Datastore

- Type
- Name and device selection**
- VMFS version
- Partition configuration
- Ready to complete

Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

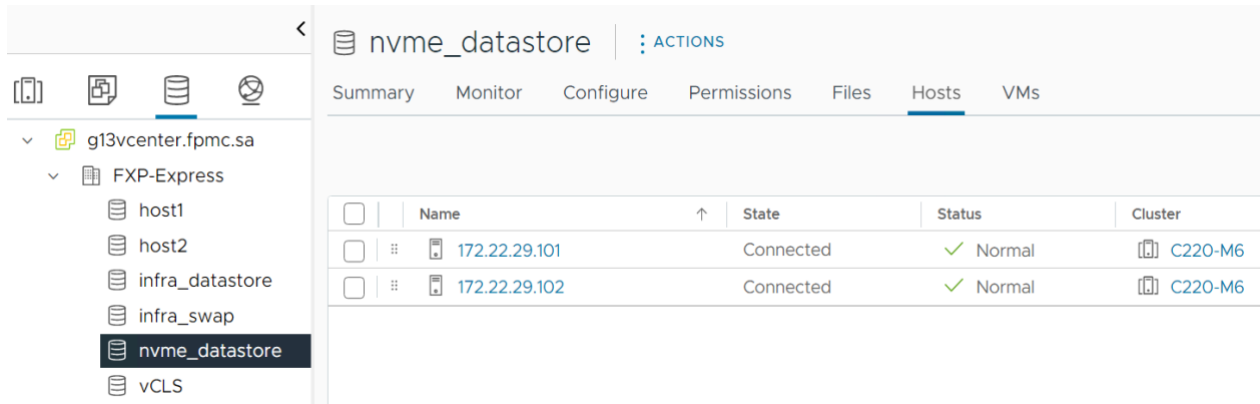
Name

Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Clust VMD Supp
NVMe TCP Disk (uuid.af6c...	0	500.00 GB	Supported	Flash	512e	No

[EXPORT](#) 1 item

CANCEL [BACK](#) [NEXT](#)

14. Leave VMFS 6 selected and click NEXT.
15. Leave all Partition configuration values at the default values and click NEXT.
16. Review the information and click FINISH.
17. Select Storage and select the new NVMe datastore. In the center pane, select Hosts. Ensure all the NVMe hosts have mounted the datastore.



NetApp ONTAP Tools 9.12 Deployment

The NetApp ONTAP tools for VMware vSphere provide end-to-end life cycle management for virtual machines in VMware environments that use NetApp storage systems. It simplifies storage and data management for VMware environments by enabling administrators to directly manage storage within the vCenter Server. This section describes the deployment procedures for the ONTAP Tools for vSphere.

NetApp ONTAP Tools for VMware vSphere 9.12

Pre-installation Considerations

The following licenses are required for NetApp ONTAP Tools on storage systems that run NetApp ONTAP 9.8 or above:

- Protocol licenses (NFS, FCP, and/or iSCSI)
- NetApp FlexClone® ((optional) Required for performing test failover operations for SRA and for vVols operations of VASA Provider.
- NetApp SnapRestore (for backup and recovery).
- The NetApp SnapManager® Suite.
- NetApp SnapMirror® or NetApp SnapVault™ (Optional - required for performing failover operations for SRA and VASA Provider when using vVols replication).

The Backup and Recovery capability has been integrated with SnapCenter and requires additional licenses for SnapCenter to perform backup and recovery of virtual machines and applications.

Note: Beginning with ONTAP 9.10.1, all licenses are delivered as NLFs (NetApp License File). NLF licenses can enable one or more ONTAP features, depending on your purchase. ONTAP 9.10.1 also supports 28-character license keys using System Manager or the CLI. However, if an NLF license is installed for a feature, you cannot install a 28-character license key over the NLF license for the same feature.

TCP Port	Requirement
443 (HTTPS)	Secure communications between VMware vCenter Server and the storage systems
8143 (HTTPS)	ONTAP Tools listens for secure communications
9083 (HTTPS)	VASA Provider uses this port to communicate with the vCenter Server and obtain TCP/IP settings
7	ONTAP tools sends an echo request to ONTAP to verify reachability and is required only when adding storage system and can be disabled later.

Note: The requirements for deploying NetApp ONTAP Tools are listed [here](#).

Install NetApp ONTAP Tools via Ansible

1. Clone the repository from <https://github.com/NetApp/ONTAP-Tools-for-VMware-vSphere>.
2. Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.
3. Update the following variable files:

```
hosts
group_vars/vcenter
vars/ontap_tools_main.yml
```

4. To invoke the ansible scripts, use the following command:

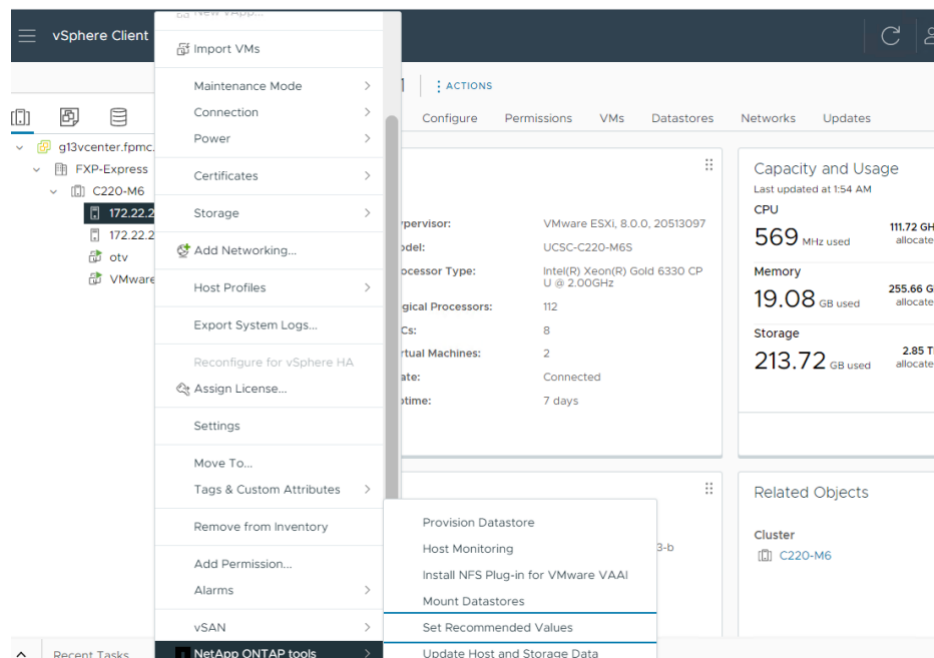
```
ansible-playbook -i hosts Setup_ONTAP_tools.yml
```

Note: The above playbook installs NetApp ONTAP Tools and registers it with VMWare vCenter. It also adds ONTAP Storage System to ONTAP tools.

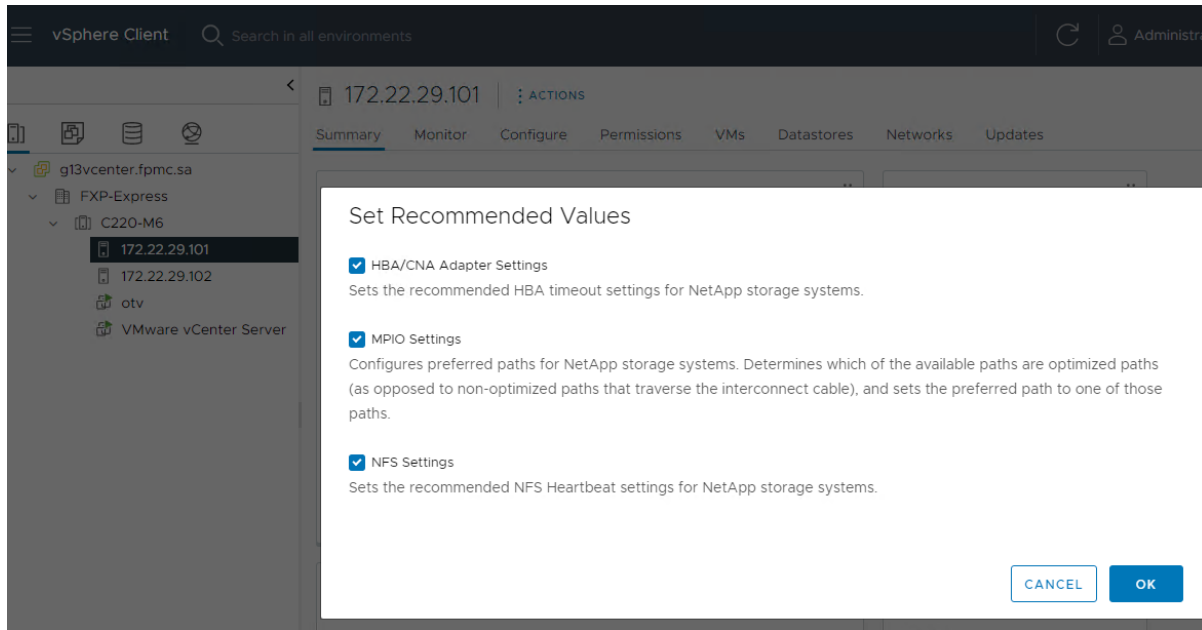
Optimal storage settings for ESXi hosts

NetApp ONTAP tools enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, follow these steps:

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



2. Check the settings that you would like to apply to the selected vSphere hosts. Click OK to apply the settings.



Conclusion

FlexPod Express with UCS C-series Standalone is designed for small to midsize businesses, remote offices, branch offices (ROBOs), and other businesses that require dedicated solutions. Cisco Intersight provides the ability to automate the policy-based configuration and deployment of rack-mount servers without the need for fabric interconnects.

This validated solution uses a combination of components from NetApp and Cisco and provides a step-by-step guide for easy adoption and automated deployment of the converged infrastructure solution using Ansible. By selecting different solution components and scaling with additional components, the FlexPod Express with UCS C-series Standalone solution can be tailored for specific business needs and can provide a reliable and flexible virtual infrastructure for application deployments.

Acknowledgment

For their support and contribution to the design, validation, and creation of this NetApp Verified Architecture, the author would like to acknowledge the significant contribution and expertise that resulted in developing this document:

Abhinav Singh – Sr. Technical Marketing Engineer, NetApp

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp Product Documentation
<https://docs.netapp.com>
- NetApp Hardware Universe
<https://hwu.netapp.com>
- NetApp Interoperability Matrix Tool (IMT)
<http://mysupport.netapp.com/matrix>

- NetApp Support Site
<https://mysupport.netapp.com>
- Cisco Intersight Configuration Guides
https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide.html
- Cisco Hardware and Software Compatibility list
<https://ucshcltool.cloudapps.cisco.com/public/>

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.

Version History

Version	Date	Document version history
Version 1.0	Sept 2023	Initial release.

Copyright Information

Copyright © 2023 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark Information

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