



NetApp Verified Architecture

FlexPod Express for VMware vSphere 8.0 with Cisco UCS C-Series Standalone and NetApp AFF NVA

Kavyashree Mahadevaiah, NetApp

July 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.

TABLE OF CONTENTS

Program Summary	4
Solution Overview	4
FlexPod Converged Infrastructure program.....	4
NetApp Verified Architecture Program	5
Solution Technology	5
Use-case Summary	6
Technology Requirements	6
Hardware Requirements	6
Software Requirements	7
Solution Design	7
Cisco Unified Computing System	7
Cisco Nexus Switching Fabric	9
Cisco Intersight	10
NetApp AFF Storage.....	10
VMware.....	11
Physical Infrastructure	12
Cabling Information.....	12
Deployment Procedures	14
Cisco Nexus 93180YC-FX Deployment Procedure	15
NetApp Storage Deployment Procedure.....	23
Cisco UCS C220 Standalone Rack Server Deployment Procedure	44
VMware vSphere 8.0 Deployment Procedure.....	64
VMware vCenter Server 8.0 Deployment Procedure.....	73
NetApp ONTAP Tools 9.12 Deployment.....	93
Ansible Automation for Solution Deployment	103
Conclusion	104
Acknowledgment.....	104
Where to find additional information	104
Version History	105
Copyright Information	105

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.

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 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.

Hardware components

- Cisco UCS C-series C220 Standalone
- Cisco Nexus 93180YC-FX
- NetApp AFF A250


Software components

- Cisco NXOS Firmware 10.2(4)
- Cisco CIMC 4.2(3d)
- NetApp ONTAP® 9.12.1 Latest
- VMware vSphere 8.0
- Cisco Intersight
- ONTAP Tools 9.12 for VMware vSphere

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 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 scale up 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

Solution Design

Cisco Unified Computing System

The Cisco UCS C220 M6 Rack Server is a 2-socket, 1-Rack-Unit (1RU) rack server offering industry-leading performance and expandability. It delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. Cisco UCS C-Series M6 Rack Servers can be deployed as standalone servers, as part of a Cisco Unified Computing System (Cisco UCS) managed environment, and now with Cisco Intersight to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' Total Cost of Ownership (TCO) and increase business agility.


Key features

- Supports the third-generation Intel Xeon Scalable CPU, with up to 40 cores per socket
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane DC Persistent Memory
- Support for Cisco UCS VIC 1400 Series adapters as well as third-party options
- Up to 10 SAS/SATA or NVMe disk drives
- M.2 boot options

Cisco UCS Virtual Interface Cards (VICs)


CISCO UCS VIC 1455

Figure 3) CISCO UCS VIC 1455



The Cisco UCS VIC 1455 is a quad-port Small Form-Factor Pluggable (SFP28) half-height PCIe card designed for Cisco UCS C-Series M5 and M6 Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Figure 4) CISCO UCS VIC 1455 Infrastructure




CISCO UCS VIC 1467

Figure 5) CISCO UCS VIC 1467



The Cisco UCS VIC 1467 is a quad-port Small Form-Factor Pluggable (SFP28) mLOM (modular LAN On Motherboard) card designed for Cisco UCS C-Series M6 Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBA.

Figure 6) CISCO UCS VIC 1467 Infrastructure



Cisco Nexus Switching Fabric

The Cisco Nexus 9300 Series Switches offer both modular and fixed 1/10/25/40/100/400 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of non-blocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

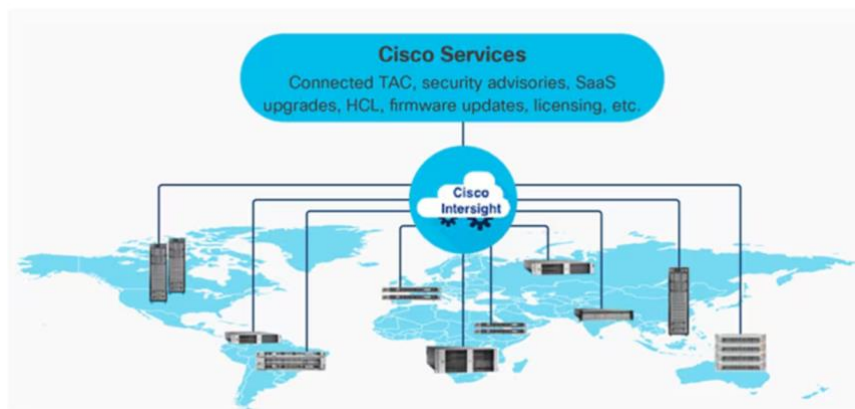
The Cisco Nexus 9000 series switch featured in this design is the Cisco Nexus 93180YC-FX configured in NX-OS standalone mode. NX-OS is a purpose-built data-center operating system designed for performance, resiliency, scalability, manageability, and programmability at its foundation. It provides a robust and comprehensive feature set that meets the demanding requirements of virtualization and automation.

The Cisco Nexus 93180YC-FX Switch is a 1RU switch that supports 3.6 Tbps of bandwidth and 1.2 bpps. The 48 downlink ports on the 93180YC-FX2 can support 1-, 10-, or 25-Gbps Ethernet or 16- or 32-Gbps Fibre Channel ports, offering deployment flexibility and investment protection. The 6 uplink ports can be configured as 40- or 100-Gbps Ethernet, offering flexible migration options. If more scale or 100 Gbps ports are needed, the Nexus 93360YC-FX2 or 9336C-FX2-E switches can be used.

Cisco Intersight

The Cisco Intersight platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. The Cisco Intersight platform is designed to be modular, so customers can adopt services based on their individual requirements. The platform significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open API design that natively integrates with third-party platforms and tools. In this solution, we derive multiple Intersight Server Profiles from a single Intersight Server Profile Template, that help achieve consistency and uniformity among multiple servers configurations

Figure 7) CISCO Intersight



NetApp AFF Storage

NetApp AFF A-Series controller lineup provides industry-leading performance while continuing to provide a full suite of enterprise-grade data services for a shared environment across on-premises data centers and the cloud. Powered by NetApp ONTAP data management software, NetApp AFF A-Series systems deliver the industry's highest performance, superior flexibility, and best-in-class data services and cloud integration to help you accelerate, manage, and protect business-critical data across your hybrid clouds. As the first enterprise-grade storage systems to support both FC-NVMe and NVMe-TCP, AFF A-Series systems boost performance with modern network connectivity.

NetApp AFF C-Series All Flash controller lineup systems built primarily for sustainability deliver industry-leading efficiency, superior flexibility, and best-in-class data services and cloud integration to help scale the IT infrastructure, simplify data management, and reduce storage cost and power consumption. Powered by ONTAP 9.12.1P1 or higher, C250 platform is best suited for mid-size business and enterprises as a capacity option with lower latency expectation and best value.

NetApp AFF A250

NetApp offers a wide range of AFF-A series controllers to meet the varying demands of the field. The entry-level, budget friendly AFF A250, provides 40% more performance and 33% more efficiency at no extra cost compared with its predecessor.

The NetApp AFF A250 offers full end-to-end NVMe support. The frontend FC-NVMe connectivity makes it possible to achieve optimal performance from an all-flash array for workloads that include artificial intelligence, machine learning, and real-time analytics as well as business-critical databases. The frontend NVMe-TCP connectivity enables customers to take advantage of NVMe technology over existing ethernet infrastructure for faster host connectivity. On the back end, the A250 supports both serial-attached SCSI (SAS) and NVMe-attached SSDs, offering the versatility for current customers to move up from their legacy A-Series systems and satisfying the increasing interest that all customers have in NVMe-based storage.

The NetApp AFF A250 offers greater port availability, network connectivity, and expandability. The NetApp AFF A250 has 4 PCIe Gen3 slots per high-availability pair. The NetApp AFF A250 offers 10GbE-T, 25GbE and 100GbE ports for IP based transport and 32Gb ports for FC and FC-NVMe traffic. This model was created to keep up with changing business needs and performance and workload requirements by merging the latest technology for data acceleration and ultra-low latency in an end-to-end NVMe storage system.

VMware

vSphere 8.0

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructures (resources including CPUs, storage, and networking) as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere 8.0 has many improvements and simplifications including, but not limited to:

- Limits with vSphere 8 have been increased including number of GPU devices is increased to 8, the number of ESXi hosts that can be managed by Lifecycle Manager is increased from 400 to 1000, the maximum number of VMs per cluster is increased from 8,000 to 10,000, and the number of VM DirectPath I/O devices per host is increased from 8 to 32.
- Security improvements include adding an SSH timeout on ESXi hosts, a TPM Provisioning policy allowing a vTPM to be replaced when cloning VMs, and TLS 1.2 as the minimum supported TLS version.
- Lifecycle Management improvements including VMware vSphere Configuration Profiles as a new alternative to VMware Host Profiles, staging cluster images and remediating up to 10 ESXi hosts in parallel instead of one at a time.
- New Virtual Hardware in VM hardware version 20 supporting the latest guest operating systems, including Windows 11.
- Distributed Resource Scheduler and vMotion improvements

For more information about VMware vSphere and its components, refer:

<https://www.vmware.com/products/vsphere.html>.

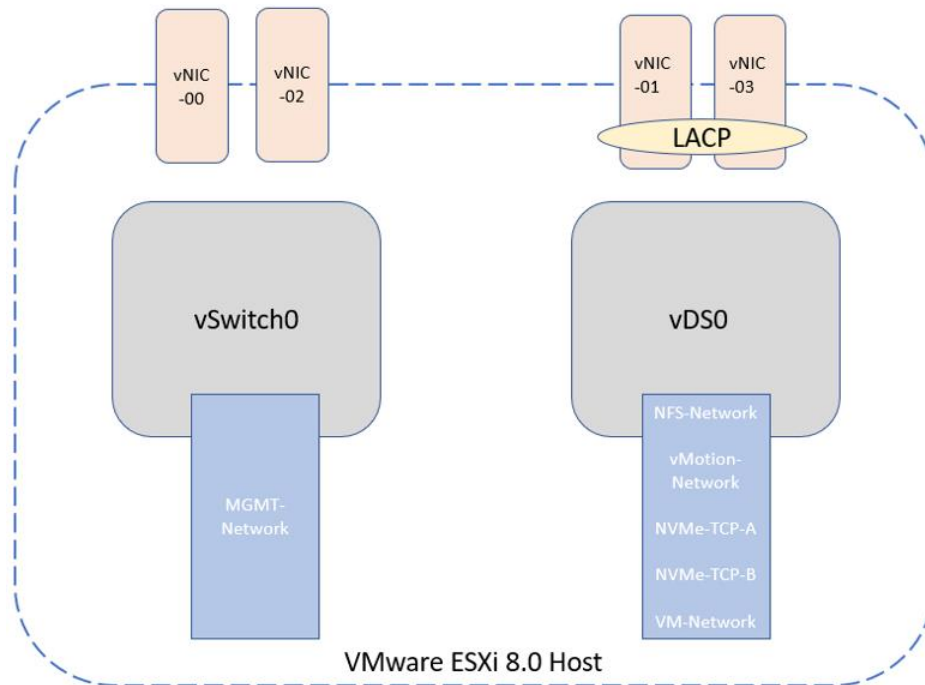
VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment. For each physical port, a single vNIC is assigned thus eliminating the need for VN-Tagging and allowing more switches to be supported.

ESXi Host Virtual Network Interface Card layout

Cisco UCS VIC 1455 has four physical ports. This solution validation includes these four physical ports in using the ESXi host. Two ports are added as uplink to standard switch (vSwitch) and the remaining two to distributed switch (vDS). Link Aggregation Group (LAG) is configured with Link Aggregation Control Protocol (LACP) mode active on the vDS to form port-channel between the C220 servers and Nexus switches.

Figure 8) ESXi Host Virtual Network Interface Card Layout





Physical Infrastructure

Cabling Information

The reference validation cabling details are shown in Figure 9 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 9) Reference validation components and cabling.





AFF A250

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-01	e1a
	Eth1/2	AFF-01	e1b
	Eth1/3	AFF-02	e1a
	Eth1/4	AFF-02	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-01	e1c
	Eth1/2	AFF-01	e1d
	Eth1/3	AFF-02	e1c
	Eth1/4	AFF-02	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

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

```

Enable advanced features

You must enable certain advanced features in Cisco NX-OS to provide additional configuration options.

1. SSH to nexus switches

2. To enable the appropriate features on Cisco Nexus switch A and switch B, run the following commands:

```
config terminal
feature interface-vlan
feature lacp
feature lldp
feature uddl
feature vpc
```

Perform global configuration

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ip name-server <dns-server-1> <dns-server-2>
ip domain-name <dns-domain-name>
ip domain-lookup
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
clock timezone <timezone> <hour-offset> <minute-offset>
For Example: clock timezone EST -5 0)
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day>
<end-month> <end-time> <offset-minutes>
(For Example: clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60)
copy run start
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
```

Note: For more information on configuring the timezone and daylight savings time or summer time, see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 10.2\(x\)](#)

Define VLANs

Before individual ports with different VLANs are configured, you must define the layer-2 VLANs on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From the configuration mode (config terminal), run the following commands to define and describe the layer-2

VLANs on Cisco Nexus switch A and B

```
vlan <native_vlan_id>
  name NATIVE-VLAN
vlan <ib_mgmt_vlan_id>
  name IB-MGMT-VLAN
vlan <nfs_vlan_id>
  name NFS
vlan <vmotion_vlan_id>
  name vMotion-VLAN
vlan <vm_traffic_vlan_id>
  name VM-Traffic-VLAN
vlan <NVMe_TCP_A_vlan_id>
  name NVMe_TCP_A-VLAN
vlan <NVMe_TCP_B_vlan_id>
  name NVMe_TCP_B-VLAN
exit
```

Add NTP distribution interface

Cisco Nexus switch A

From the global configuration mode, execute the following commands.

```
interface vlan < ib_mgmt_vlan_id>
ip address < switch_ntp_ip_a>/< mgmt_vlan_netmask_length>
no shutdown
exit
ntp peer <switch_ntp_ip_b> use-vrf management
```

Cisco Nexus switch B

From the global configuration mode, execute the following commands.

```
interface vlan <ib_mgmt_vlan_id>
ip address <switch_ntp_ip_b>/<ib_mgmt_vlan_netmask_length>
no shutdown
exit
ntp peer <switch_ntp_ip_a> use-vrf management
```

Configure port descriptions

As is the case with assigning names to the layer-2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (config t) in each of the switches, enter the following port descriptions for the FlexPod Express configuration:

Cisco Nexus switch A

```
int eth1/1
description AFF-01:e1a
int eth1/2
description AFF-01:e1b
int eth1/3
description AFF-02:e1a
int eth1/4
description AFF-02:e1b
int eth1/11
description UCS-1-VIC-P0-Standalone
int eth1/12
description UCS-1-VIC-P1-PC
int eth1/13
description UCS-2-VIC-P0- Standalone
int eth1/14
description UCS-2-VIC-P1-PC
int eth1/51
description vPC peer-link 93180YC-FX-B
int eth1/52
description vPC peer-link 93180YC-FX-B
int eth1/54
description uplink-switch
```

Cisco Nexus switch B

```
int eth1/1
description AFF-01:e1c
int eth1/2
description AFF-01:e1d
int eth1/3
description AFF-02:e1c
int eth1/4
description AFF-02:e1d
int eth1/11
description UCS-1-VIC-P2-Standalone
int eth1/12
description UCS-1-VIC-P3-PC
int eth1/13
```

```

description UCS-2-VIC-P2-Standalone
int eth1/14
description UCS-2-VIC-P3-PC
int eth1/51
description vPC peer-link 93180YC-FX-A
int eth1/52
description vPC peer-link 93180YC-FX-A
int eth1/54
description uplink-switch

```

Perform virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer-2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree-protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability.

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly.

From configuration mode (`config t`), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus switch A

```

vpc domain 1
role priority 10
peer-keepalive destination <switch_ip_b> source <switch_ip_a>
peer-switch
delay restore 150
peer-gateway
auto-recovery
ip arp synchronize
!
int Po10
description vPC peer-link
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<vmotion_vlan_id>,<vmtraffic_vlan_id>,<
NVMe_TCP_A_vlan_id>,< NVMe_TCP_B_vlan_id>
spanning-tree port type network
vpc peer-link
no shutdown
!
int eth1/51-52
channel-group 10 mode active
no shutdown
!
int Po11
description vPC ucs- 1
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<vmotion_vlan_id>,<vmtraffic_vlan_id>,<
NVMe_TCP_A_vlan_id>,< NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216

```

```

vpc 11
no shutdown
!
int eth1/12
fec rs-fec
channel-group 11 mode active
no shutdown
!
interface Po12
description vPC ucs-2
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<vmotion_vlan_id>,<vmtraffic_vlan_id>,<
NVMe_TCP_A_vlan_id>,<NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 12
no shutdown
!
int eth1/14
fec rs-fec
channel-group 12 mode active
no shutdown
!
interface Po21
description A250-1
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<
NVMe_TCP_A_vlan_id>,<NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 21
no shutdown
!
int eth1/1-2
channel-group 21 mode active
no shutdown
!
interface Po22
description A250-2
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<
NVMe_TCP_A_vlan_id>,<NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 22
no shutdown
!
int eth1/3-4
channel-group 22 mode active
no shutdown
!
interface Ethernet1/11
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>
mtu 9216
fec rs-fec
vpc orphan-port suspend
no shutdown
!
interface Ethernet1/13
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>
mtu 9216
fec rs-fec
vpc orphan-port suspend
no shutdown
exit

```

NOTE: C220 servers standalone ports connected to the nexus switch will be configured as orphan-port suspend as they are non-vPC ports.

Cisco Nexus switch B

```
vpc domain 1
peer-switch
role priority 20
peer-keepalive destination <switch_ip_a> source <switch_ip_b>
delay restore 150
peer-gateway auto-recovery
ip arp synchronize
!
int Po10
description vPC peer-link
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<vmotion_vlan_id>,<vmtraffic_vlan_id> ,
<NVMe_TCP_A_vlan_id>,< NVMe_TCP_B_vlan_id>
spanning-tree port type network vpc peer-link
no shutdown
!
int eth1/51-52
channel-group 10 mode active
no shutdown
!
int Po11
description vPC ucs-A
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan>,<nfs_vlan_id>,<vmotion_vl an_id>,<vmtraffic_vlan_id>,
<NVMe_TCP_A_vlan_id>,< NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 11
no shutdown
!
int eth1/12
fec rs-fec
channel-group 11 mode active
no shutdown
!
int Po12
description vPC ucs-2
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan>,<nfs_vlan_id>,<vmotion_vlan_id>,<vmtraffic_vlan_id>,
<NVMe_TCP_A_vlan_id>,< NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 12
no shutdown
!
int eth1/12
fec rs-fec
channel-group 12 mode active
no shutdown
!
interface Po21
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<NVMe_TCP_A_vlan_id>,<NVMe_TCP_B_vlan_id
>
spanning-tree port type edge trunk
mtu 9216
vpc 21
no shutdown
!
int eth1/1-2
channel-group 21 mode active
no shutdown
!
```

```

interface Po22
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>,<NVMe_TCP_A_vlan_id>,<NVMe_TCP_B_vlan_id>
spanning-tree port type edge trunk
mtu 9216
vpc 22
no shutdown
!
int eth1/3-4
channel-group 22 mode active
no shutdown
!
interface Ethernet1/11
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>
mtu 9216
fec rs-fec
vpc orphan-port suspend
no shutdown
!
interface Ethernet1/13
switchport mode trunk
switchport trunk native vlan <native_vlan_id>
switchport trunk allowed vlan <ib_mgmt_vlan_id>,<nfs_vlan_id>
mtu 9216
fec rs-fec
vpc orphan-port suspend
no shutdown
exit

```

Note: Configuring the Eth1/11 and Eth1/13 standalone interfaces as orphan virtual Port Channel (vPC) ports

Note: In this solution validation, a maximum transmission unit (MTU) of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components result in packets being dropped.

Uplink into existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 9000 switches included in the FlexPod environment into the infrastructure. The uplinks can be 100GbE uplinks for a 10GbE infrastructure solution or 25GbE for a 1GbE infrastructure solution, if required.

For this deployment guide, a single 100GbE uplink to existing network is provided for the in-band management network from each switch.

```

interface Po54
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <vmtraffic_vlan_id>,<ib_mgmt_vlan_id>
spanning-tree port type network
mtu 9216
vpc 54
no shutdown
!
int eth1/54
channel-group 54 mode active
no shutdown

```

Save switch configuration

After the configuration is completed on the switches, be sure to exit the configuration mode and run copy start to save the configuration.

```
copy running-config startup-config
```

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.

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

After a few minutes the cluster setup will be completed. Login to ONTAP cluster when prompted.

1. From the Dashboard click the Cluster menu and click Overview.
2. Click the More ellipsis button in the Overview pane and click Edit.

ONTAP System Manager

Search actions, objects, and pages

<>

DASHBOARD

INSIGHTS

STORAGE

Overview

Volumes

Consistency Groups

Storage VMs

Tiers

NETWORK

EVENTS & JOBS

PROTECTION

CLUSTER

Overview

Hardware

Overview

Overview

NAME

g1434-a250

MANAGEMENT INTERFACES

172.22.29.15

VERSION

NetApp Release 9.12.1P2: Wed Apr 05 19:57:43 UTC 2023

DATE AND TIME

May 10, 2023, 11:50 AM Etc/UTC

UUID

cae8de12-ef27-11ed-b10c-d039ea4099f3

More

ONTAP Update

Rename

Edit

Login Banner Message

Download Configuration

Nodes

Show / Hide

Nodes	Name	Serial Number	Up...	Utilization	Management IP	Service Pr...	System ID
g1434-a250-02 / g1434-a250-01							

3. Add additional cluster configuration details and click Save to make the changes persistent:

- Cluster location
- DNS Domain
- DNS server IP addresses
- NTP server IP addresses

Note: DNS and NTP server IP addresses can be added individually or with a comma separate list on a single line.

Edit Cluster Details



NAME

g1434-a250

LOCATION

NetApp RTP, Building 1, Eng 1/G/14

DNS DOMAINS

fpmc.sa

[+ Add](#)

NAME SERVERS

10.61.176.251

10.61.176.252

[+ Add](#)

NTP SERVERS

10.61.176.251

10.61.176.252

10.61.176.16

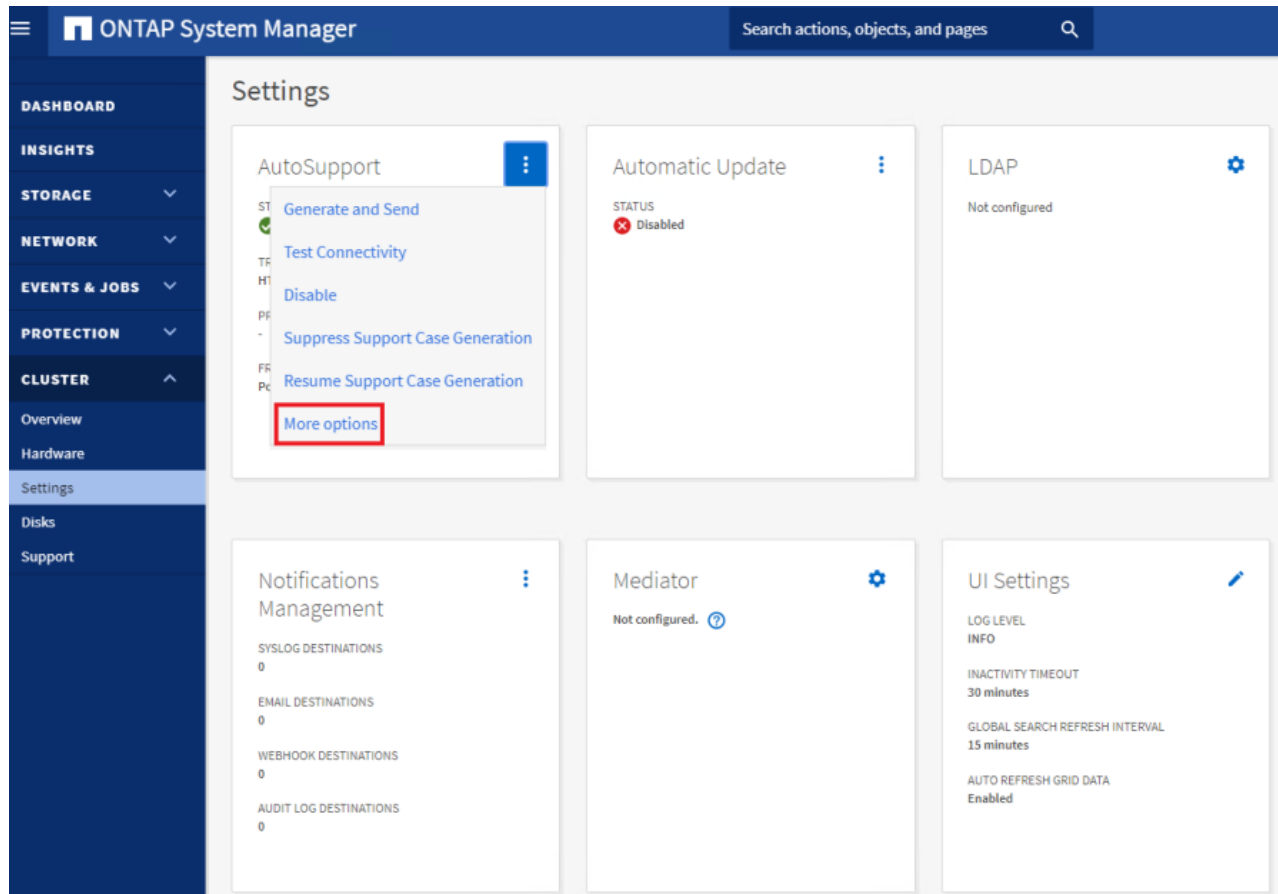
[+ Add](#)

☐ Add cluster management interface

Save

[Cancel](#)

4. Click Save. Select the Settings menu under the Cluster menu and configure AutoSupport.
5. Click the ellipsis in the AutoSupport tile and select More options



6. Click Edit to change the transport protocol, add a proxy server address and a mail host as needed.
7. Click Save to enable the changes.
8. In the Email tile to the right, click Edit and enter the desired email information:
 - a. Email send from address
 - b. Email recipient addresses
 - c. Recipient Category
9. Click Save when Complete.

AutoSupport Cluster Settings

☒ Enabled
 [More](#)

Connections [Edit](#)

TRANSPORT PROTOCOL
HTTPS

PROXY SERVER

MAIL HOST
mailhost

Email

EMAIL SEND FROM
g1434-a250@fpmc.sa

EMAIL RECIPIENTS

Email Address	Recipient Category
admin-fxp@fpmc.sa	General

[+ Add](#)

[Save](#) [Cancel](#)

10. Select CLUSTER > Settings at the top left of the page to return to the cluster settings page.

11. Locate the Licenses tile on the right and click the detail arrow.

ONTAP System Manager Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE

NETWORK

EVENTS & JOBS

PROTECTION

CLUSTER

Overview

Hardware

Settings

Disks

Support

Settings

AutoSupport

STATUS Enabled

TRANSPORT PROTOCOL
HTTPS

PROXY SERVER
-

FROM EMAIL ADDRESS
Postmaster

Automatic Update

STATUS Disabled

LDAP

Not configured

Licenses

None

[→](#)

Notifications Management

SYSLOG DESTINATIONS
0

EMAIL DESTINATIONS
0

WEBHOOK DESTINATIONS
0

AUDIT LOG DESTINATIONS
0

Mediator

Not configured

UI Settings

LOG LEVEL
INFO

INACTIVITY TIMEOUT
30 minutes

GLOBAL SEARCH REFRESH INTERVAL
15 minutes

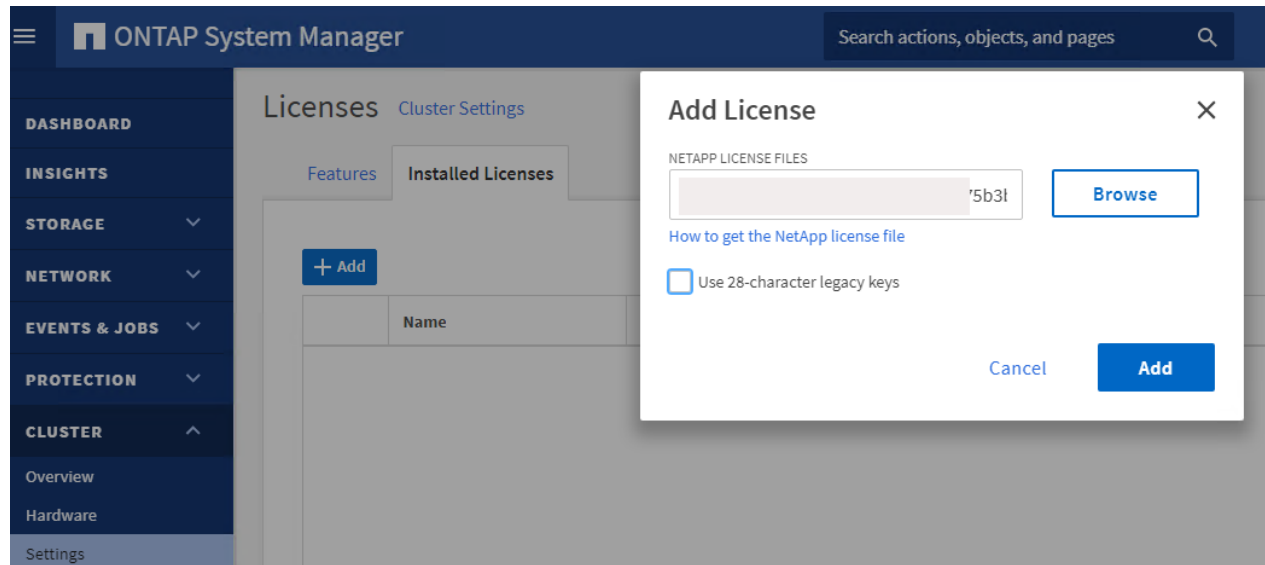
AUTO REFRESH GRID DATA
Enabled

Cloud Connections

No cloud connections added.
[Learn about cloud connectivity.](#)

[+ Add](#)

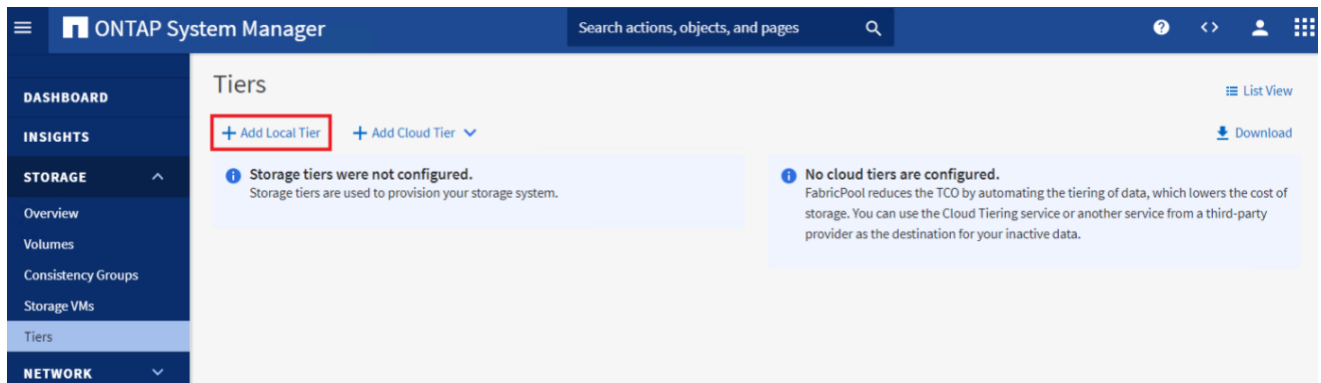
12. Add the desired licenses to the cluster by clicking Add and uploading the License files



NOTE: Depending on the environment, customers can either choose to add License keys or NetApp License File (NLF).

NOTE: NetApp ONTAP 9.10.1 and later for FAS/AFF storage systems uses a new file-based licensing solution to enable per-node NetApp ONTAP features. The new license key format is referred to as a NetApp License File, or NLF. For more information, refer to this URL: [NetApp ONTAP 9.10.1 and later Licensing Overview – NetApp](#)

13. Configure storage aggregates by selecting the Storage menu on the left and selecting Tiers.
14. Click Add Local Tier and allow NetApp ONTAP System Manager to recommend a storage aggregate configuration.



15. NetApp ONTAP will use best practices to recommend an aggregate layout. Click the Recommended details link to view the aggregate information.
16. Optionally, enable NetApp Aggregate Encryption (NAE) by checking the box for Configure Onboard Key Manager for encryption.
17. Enter and confirm the passphrase and save it in a secure location for future use.

Add Local Tier



Storage Recommendation

69.2 TiB

USABLE

2 local tiers can be added on nodes g1434-a250-02 and g1434-a250-01.

Recommendation details

LOCAL TIER DETAILS

Local Tier	Node Name	Usable Size	Type	Disks
g1434_a250_01_NVME_SSD_1	g1434-a25...	34.6 TiB	SSD	23 X 1.74 TiB (NVMe SSD partition RAID-DP)
g1434_a250_02_NVME_SSD_1	g1434-a25...	34.6 TiB	SSD	23 X 1.74 TiB (NVMe SSD partition RAID-DP)

SPARE DISKS

Node Name	Spare Disks	Type	Is Partition
g1434-a250-01	1 X 1.74 TiB	NVMe SSD	Yes
g1434-a250-02	1 X 1.74 TiB	NVMe SSD	Yes

Not sure about the recommendation? [Switch to Manual Local Tier Creation](#)

Encryption

[Considerations](#)

☒ Configure Onboard Key Manager for encryption

.....

.....

i Save the passphrase for future use. You will need the passphrase if the system needs to be recovered.
After the Onboard Key Manager is configured, back up the key database for future use.

[Cancel](#)

[Save](#)

Login to the cluster

1. Open an SSH connection to either the cluster IP or the host name.
2. Log into the admin user with the password you provided earlier.

Verify Storage Failover

Verify the status of the storage failover.

```
storage failover show
```

Note: Both <st-node01> and <st-node02> must be capable of performing a takeover. Continue with step 2 if the nodes can perform a takeover.

1. Enable failover on one of the two nodes if it was not completed during the installation.

```
storage failover modify -node <st-node01> -enabled true
```

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

2. Verify the HA status for a two-node cluster.

```
cluster ha show
```

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

3. If HA is not configured use the below commands. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

Set Auto-Revert Parameter on Cluster Management Interface

Run the following command:

```
network interface modify -vserver <clustername> -lif cluster_mgmt_lif -auto-revert true
```

Note: A storage virtual machine (SVM) is referred to as a Vserver or vservers in the GUI and CLI.

Zero all spare disks

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

```
disk zerosparses
```

Set Up Service Processor Network Interface

To assign a static IPv4 address to the Service Processor on each node, run the following commands:

```
system service-processor network modify -node <st-node01> -address-family IPv4 -enable true -dhcp none  
-ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>  
  
system service-processor network modify -node <st-node02> -address-family IPv4 -enable true -dhcp none  
-ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```

Note: The Service Processor IP addresses should be in the same subnet as the node management IP addresses.

Create Manual Provisioned Aggregates (Optional)

An aggregate containing the root volume is created during the NetApp ONTAP setup process. To manually create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain. Options for disk type include SAS, SSD, and SSD-NVM.

To create new aggregates, run the following commands:

```
storage aggregate create -aggregate <aggr1_node01> -node <st-node01> -diskcount <num-disks> -disktype  
SSD-NVM  
  
storage aggregate create -aggregate <aggr1_node02> -node <st-node02> -diskcount <num-disks> -disktype  
SSD-NVM
```

Note: Customer should have the minimum number of hot spare disks for the recommended hot spare disk partitions for their aggregate.

Note: For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

Note: In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all, but one remaining disk (spare) assigned to the controller.

Note: The aggregate cannot be created until disk zeroing completes. Run the **storage aggregate show** command to display the aggregate creation status. Do not proceed until both aggr1_node01 and aggr1_node02 are online.

Remove Default Broadcast Domains

By default, all network ports are included in separate default broadcast domain. Network ports used for data services (for example, e1a, e1b, etc.,) should be removed from their default broadcast domain and that broadcast domain should be deleted.

To perform this task, run the following commands:

```
network port broadcast-domain delete -broadcast-domain <Default-N> -ipspace Default
network port broadcast-domain show
```

Note: Delete the Default broadcast domains with Network ports (Default-1, Default-2, and so on). This does not include Cluster ports and management ports.

Disable Flow Control on 25/100GbE Data Ports

Run the following command to configure the ports on node 01:

```
network port modify -node <st-node01> -port e1a,e1b,e1c,e1d -flowcontrol-admin none
```

Run the following command to configure the ports on node 02:

```
network port modify -node <st-node02> -port e1a,e1b,e1c,e1d -flowcontrol-admin none
```

Note: Disable flow control only on ports that are used for data traffic.

Enable Cisco Discovery Protocol

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

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

Enable Link-layer Discovery Protocol on all Ethernet ports

Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches by running the following command. This command enables LLDP on all ports of all nodes in the cluster.

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

Configure login banner for the NetApp ONTAP Cluster

To create login banner for the NetApp ONTAP cluster, run the following command:

```
security login banner modify -message "Access restricted to authorized users" -vserver <clustername>
```

Note: If the login banner for the cluster is not configured, users will observe a warning in AIQUM stating "Login Banner Disabled."

Enable FIPS Mode on the NetApp ONTAP Cluster (Optional)

NetApp ONTAP is compliant in the Federal Information Processing Standards (FIPS) 140-2 for all SSL connections. When SSL FIPS mode is enabled, SSL communication from NetApp ONTAP to external client or server components outside of NetApp ONTAP will use FIPS compliant crypto for SSL.

To enable FIPS on the NetApp ONTAP cluster, run the following commands:

```
set -privilege advanced
security config modify -interface SSL -is-fips-enabled true
```

Note: If you are running NetApp ONTAP 9.8 or earlier manually reboot each node in the cluster one by one. Beginning in NetApp ONTAP 9.9.1, rebooting is not required.

Note: If FIPS is not enabled on the NetApp ONTAP cluster, the users will observe a warning in AIQUM stating “FIPS Mode Disabled.”

Remove insecure ciphers from the NetApp ONTAP Cluster

Ciphers with the suffix CBC are considered insecure. To remove the CBC ciphers, run the following NetApp ONTAP command:

```
security ssh remove -vserver <clustername> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

Note: If the users do not perform the above task, they will see a warning in AIQUM saying “SSH is using insecure ciphers.”

Configure Timezone

To configure time synchronization on the cluster, follow these steps:

Set the time zone for the cluster.

```
timezone -timezone <timezone>
```

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

Configure Simple Network Management Protocol

Note: If users have enabled FIPS then please look at the following points while configuring SNMP.

- The SNMP users or SNMP traphosts that are non-compliant with FIPS will be deleted automatically. “Configure SNMP traphosts” configuration will be non-compliant with FIPS.
- The SNMPv1 user, SNMPv2c user (After configuring SNMP community) or SNMPv3 user (with none or MD5 as authentication protocol or none or DES as encryption protocol or both) is non-compliant with FIPS.

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

```
snmp contact <snmp-contact>
snmp location <snmp-location>
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as an Active IQ® Unified Manager server or another fault management system.

Note: This step works when FIPS is disabled.

Note: An SNMPv1 traphost or SNMPv3 traphost (configured with an SNMPv3 user non-compliant to FIPS) is non-compliant to FIPS.

```
snmp traphost add <oncommand-um-server-fqdn>
```

3. Configure SNMP community.

Note: This step works when FIPS is disabled.

Note: SNMPv1 and SNMPv2c are not supported when cluster FIPS mode is enabled.

```
system snmp community add -type ro -community-name <snmp-community> -vserver <clustername>
```

Note: In new installations of NetApp ONTAP, SNMPv1 and SNMPv2c are disabled by default. SNMPv1 and SNMPv2c are enabled after you create an SNMP community.

Note: NetApp ONTAP supports read-only communities.

Configure SNMPv3 Access

SNMPv3 offers advanced security by using encryption and passphrases. The SNMPv3 users can run SNMP utilities from the traphost using the authentication and privacy settings that they specify.

Note: When FIPS is enabled, below are the supported/compliant options for authentication and privacy protocol:

- Authentication Protocol: sha, sha2-256
- Privacy protocol: aes128

To configure SNMPv3 access, run the following commands:

```
security login create -user-or-group-name <<snmp-v3-usr>> -application snmp -authentication-method usm

Enter the authoritative entity's EngineID [local EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256) [none]: <<snmp-v3-auth-
proto>>

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128) [none]: <<snmpv3-priv-proto>>

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:
```

Note: Refer to the [SNMP Configuration Express Guide](#) for additional information when configuring SNMPv3 security users.

Create Management Broadcast Domain

If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces by running the following command:

```
network port broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
```

Create NFS Broadcast Domain

To create an NFS data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands in NetApp ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra-NFS -mtu 9000
```

Create NVMe/TCP Broadcast Domains (Required only for NVMe/TCP configuration)

To create NVMe-TCP-A and NVMe-TCP-B data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands in NetApp ONTAP:


```
network port broadcast-domain create -broadcast-domain Infra-NVMe-TCP-A -mtu 9000
network port broadcast-domain create -broadcast-domain Infra-NVMe-TCP-B -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 25GbE data interfaces, run the following commands:

```
network port ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port el1
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port el2
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port el3
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port el4
```

```
network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port el1
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port el2
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port el3
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port el4
```

Change MTU on Interface Groups

To change the MTU size on the base interface-group ports before creating the VLAN ports, run the following commands:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

Create VLANs

1. Create the management VLAN ports and add them to the management broadcast domain.

```
network port vlan create -node <st-node01> -vlan-name a0a-<ib-mgmt-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<ib-mgmt-vlan-id>

network port broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <st-node01>:a0a-<ib-mgmt-vlan-id>,<st-node02>:a0a-<ib-mgmt-vlan-id>
```

2. Create the NFS VLAN ports and add them to the Infra-NFS broadcast domain.

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-NFS -ports <st-node01>:a0a-<infra-nfs-vlan-id>,<st-node02>:a0a-<infra-nfs-vlan-id>
```

3. If configuring NVMe/TCP, create NVMe/TCP VLAN ports for the NVMe/TCP LIFs on each storage controller and add them to the corresponding broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nvme-tcp-a-vlan-id>
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nvme-tcp-b-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nvme-tcp-a-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nvme-tcp-b-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-NVMe-TCP-A -ports <st-node01>:a0a-<infra-nvme-tcp-a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-NVMe-TCP-B -ports <st-node01>:a0a-<infra-nvme-tcp-b-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-NVMe-TCP-A -ports <st-node02>:a0a-<infra-nvme-tcp-a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-NVMe-TCP-B -ports <st-node02>:a0a-<infra-nvme-tcp-b-vlan-id>
```

Create SVM (Storage Virtual Machine)

1. Run the vservers create command.

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

2. Add the required data protocols to the SVM:

```
vserver add-protocols -protocols nfs,nvme -vserver Infra-SVM
```

Note: For NVMe/TCP configuration, add “nvme” protocol to the SVM.

3. Remove the unused data protocols from the SVM:

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

Note: It is recommended to remove iSCSI or FCP protocols if the protocol is not in use.

4. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp ONTAP Tools.

```
vserver modify -vserver Infra-SVM -aggr-list <aggr1_node01>,<aggr1_node02>
```

5. Enable and run the NFS protocol in the Infra-SVM.

```
vserver nfs create -vserver Infra-SVM -udp disabled -v3 enabled -v4.1 enabled -vstorage enabled
```

Note: If the NFS license was not installed during the cluster configuration, make sure to install the license before starting the NFS service.

6. Verify that the NFS vstorage parameter for the NetApp NFS VAAI plug-in was enabled.

```
g1434-a250::> vserver nfs show -fields vstorage
```

```
vserver    vstorage
```

```
-----
```

```
Infra-SVM enabled
```

Vserver Protocol Verification

1. Verify the required protocols are added to the Infra-SVM vservers.

```
g1434-a250::> vserver show-protocols -vserver Infra-SVM
```

```
Vserver: Infra-SVM
Protocols: nfs, nvme
```

2. If a protocol is not present, use the following command to add the protocol to the vservers:

```
vserver add-protocols -vserver <infra-data-svm> -protocols <nvme>
```

Create Load-Sharing Mirrors of SVM Root Volume

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra-SVM -volume infra_svm_root_m01 -aggregate <aggr1_node01> -size 1GB -type DP
volume create -vserver Infra-SVM -volume infra_svm_root_m02 -aggregate <aggr1_node02> -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m01 -type LS -schedule 15min
```

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Infra-SVM:infra_svm_root
```

Create NVMe Service (required only NVMe/TCP configuration)

NOTE: Make sure NVMe capable adapters are installed in the cluster. This can be checked by using the command **network fcp adapter show -data-protocols-supported fc-nvme**

1. Make sure that the “nvme” protocol is added to the SVM.

```
g1434-a250::> vserver show-protocols -vserver Infra-SVM

Vserver: Infra-SVM
Protocols: nfs, nvme
```

2. Create NVMe service.

```
vserver nvme create -vserver Infra-SVM -status-admin up
```

3. To verify:

```
g1434-a250::> vserver nvme show -vserver Infra-SVM

Vserver Name: Infra-SVM
Administrative Status: up
Discovery Subsystem NQN: nqn.1992-
08.com.netapp:sn.90e9cb71515311ed978d00a098e217cb:discovery
```

Note: If the NVMe license was not installed during the cluster configuration, make sure to install the license before creating the NVMe service.

Configure HTTPS access

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

```
set -privilege diag

Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, the <serial-number>) by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS fully qualified domain name (FQDN) of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -
serial <serial-number>
```

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

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

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-
country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -
email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-
SVM
```

5. To obtain the values for the parameters required in step 6 (<cert-ca> and <cert-serial>), run the **security certificate show** command.
6. Enable each certificate that was just created by using the -server-enabled true and -client-enabled false parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -
serial <cert-serial> -common-name <cert-common-name>
```

7. Disable HTTP cluster management access.

```
network interface service-policy remove-service -vserver <clustername> -policy default-management -
service management-http
```

Note: It is normal for some of these commands to return an error message stating that the entry does not exist.

Note: The command **system services firewall policy delete** is deprecated and may be removed in a future NetApp ONTAP release. So, use the above command **network interface service-policy remove-service** instead.

8. Change back to the normal admin privilege level and verify that the system logs are available in a web browser.

```
set -privilege admin
https://<node01-mgmt-ip>/spi
https://<node02-mgmt-ip>/spi
```

Set password for SVM vsadmin user and unlock the user

Set a password for the SVM vsadmin user and unlock the user using the following commands:

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

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

Configure login banner for the SVM.

To create login banner for the SVM, run the following command:

```
security login banner modify -vserver Infra-SVM -message "This Infra-SVM is reserved for authorized
users only!"
```

Note: If the login banner for the SVM is not configured, users will observe a warning in AIQUM stating “Login Banner Disabled.”

Remove insecure ciphers from the SVM.

Ciphers with the suffix CBC are considered insecure. To remove the CBC ciphers from the SVM, run the following NetApp ONTAP command:

```
security ssh remove -vserver Infra-SVM -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

Note: If the users do not perform the above task, they will see a warning in AIQUM saying “SSH is using insecure ciphers.”

Configure Export Policy Rule

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -
clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys -allow-suid true
```

2. Assign the FlexPod Export Policy to the infrastructure SVM root volume.

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

Create FlexVol® Volumes

The following information is required to create a NetApp FlexVol volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

1. To create FlexVols for datastores, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore -aggregate <aggr1_node02> -size 1TB -state
online -policy default -junction-path /infra_datastore -space-guarantee none -percent-snapshot-space 0
```

Note: If you are going to setup and use SnapCenter to backup the `infra_datastore` volume, add “`-snapshot-policy none`” to the end of the volume create command for the `infra_datastore` volume.

2. To create swap volumes, run the following command:

```
volume create -vserver Infra-SVM -volume infra_swap -aggregate <aggr1_node01> -size 200GB -state
online - policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0 -
snapshot-policy none
```

3. Create vCLS datastores to be used by the vSphere environment to host vSphere Cluster Services (vCLS) VMs using the command below:

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

4. To configure NVMe datastores, run the following commands:

```
volume create -vserver Infra-SVM -volume NVMe_Datastore_01 -aggregate <aggr1_node01> -size 500G -state
online -policy default -space-guarantee none -percent-snapshot-space 0
```

Note: If you are going to setup and use SnapCenter to backup the volume, add “`-snapshot-policy none`” to the end of the volume create command for the `NVMe_Datastore_01` volume

Note: To Configure NVMe Datastores for vSphere 8, enable the NVMe protocol on an existing SVM or create a separate SVM for NVMe workloads. In this deployment, Infra-SVM was used for NVMe datastore configuration.

Note: NVMe datastores created above can be utilized for NVMe/TCP configurations.

5. Update set of load-sharing mirrors using the command below:

```
snapmirror update-ls-set -source-path Infra-SVM:infra_svm_root
```

Disable Volume Efficiency on swap volume

On NetApp AFF systems, deduplication is enabled by default. To disable the efficiency policy on the `infra_swap` volume, run the following command:

```
volume efficiency off -vserver Infra-SVM -volume infra_swap
```

Create NFS LIFs

To create NFS LIFs, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs-lif-01 -service-policy default-data-files -home-
node <st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs-lif-01-ip> -netmask <node01-
nfs-lif-01-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

```
network interface create -vserver Infra-SVM -lif nfs-lif-02 -service-policy default-data-files -home-
node <st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs-lif-02-ip> -netmask <node02-
nfs-lif-02-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

To verify:

```
g1434-a250::> network interface show -vserver Infra-SVM -service-policy default-data-files
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Infra-SVM	nfs-lif-01	up/up	172.22.30.11/24	g1434-a250-01	a0a-2230	true
	nfs-lif-02	up/up	172.22.30.12/24	g1434-a250-02	a0a-2230	true

2 entries were displayed.

Note: For the tasks using network interface create command, the `-role` and `-firewall-policy` parameters have been deprecated and may be removed in a future version of NetApp ONTAP. Use the `-service-policy` parameter instead.

Create NVMe/TCP LIFs (required only for NVMe/TCP configuration)

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

```
network interface create -vserver Infra-SVM -lif nvme-tcp-lif-01a -service-policy default-data-nvme-tcp -home-node <st-node01> -home-port a0a-<infra-nvme-tcp-a-vlan-id> -address <st-node01-infra-nvme-tcp-a-ip> -netmask <infra-nvme-tcp-a-mask> -status-admin up
```

```
network interface create -vserver Infra-SVM -lif nvme-tcp-lif-01b -service-policy default-data-nvme-tcp -home-node <st-node01> -home-port a0a-<infra-nvme-tcp-b-vlan-id> -address <st-node01-infra-nvme-tcp-b-ip> -netmask <infra-nvme-tcp-b-mask> -status-admin up
```

```
network interface create -vserver Infra-SVM -lif nvme-tcp-lif-02a -service-policy default-data-nvme-tcp -home-node <st-node02> -home-port a0a-<infra-nvme-tcp-a-vlan-id> -address <st-node02-infra-nvme-tcp-a-ip> -netmask <infra-nvme-tcp-a-mask> -status-admin up
```

```
network interface create -vserver Infra-SVM -lif nvme-tcp-lif-02b -service-policy default-data-nvme-tcp -home-node <st-node02> -home-port a0a-<infra-nvme-tcp-b-vlan-id> -address <st-node02-infra-nvme-tcp-b-ip> -netmask <infra-nvme-tcp-b-mask> -status-admin up
```

To verify:

```
g1434-a250::> network interface show -vserver Infra-SVM -service-policy default-data-nvme-tcp
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Infra-SVM	nvme-tcp-lif-01a	up/up	172.22.33.11/24	g1434-a250-01	a0a-2233	true
	nvme-tcp-lif-01b	up/up	172.22.34.11/24	g1434-a250-01	a0a-2234	true
	nvme-tcp-lif-02a	up/up	172.22.33.12/24	g1434-a250-02	a0a-2233	true
	nvme-tcp-lif-02b	up/up	172.22.34.12/24	g1434-a250-02	a0a-2234	true

4 entries were displayed.

Create SVM management LIF (Add Infrastructure SVM Administrator)

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif svm-mgmt -service-policy default-management -home-node <st-node01> -home-port a0a-<ib-mgmt-vlan-id> -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

2. Create a default route that enables the SVM management interface to reach the outside world.

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

To verify:

```
g1434-a250::> network route show -vserver Infra-SVM
```

Vserver	Destination	Gateway	Metric
---------	-------------	---------	--------

```
-----
Infra-SVM          0.0.0.0/0      172.22.29.254      20
```

Note: A cluster serves data through at least one and possibly several SVMs. These steps have been created for a single data SVM. Customers can create additional SVMs depending on their requirement.

Configure Auto-Support

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

```
system node autosupport modify -node * -state enable -mail-hosts <mailhost> -transport https -support
enable -noteto <storage-admin-email>
```

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

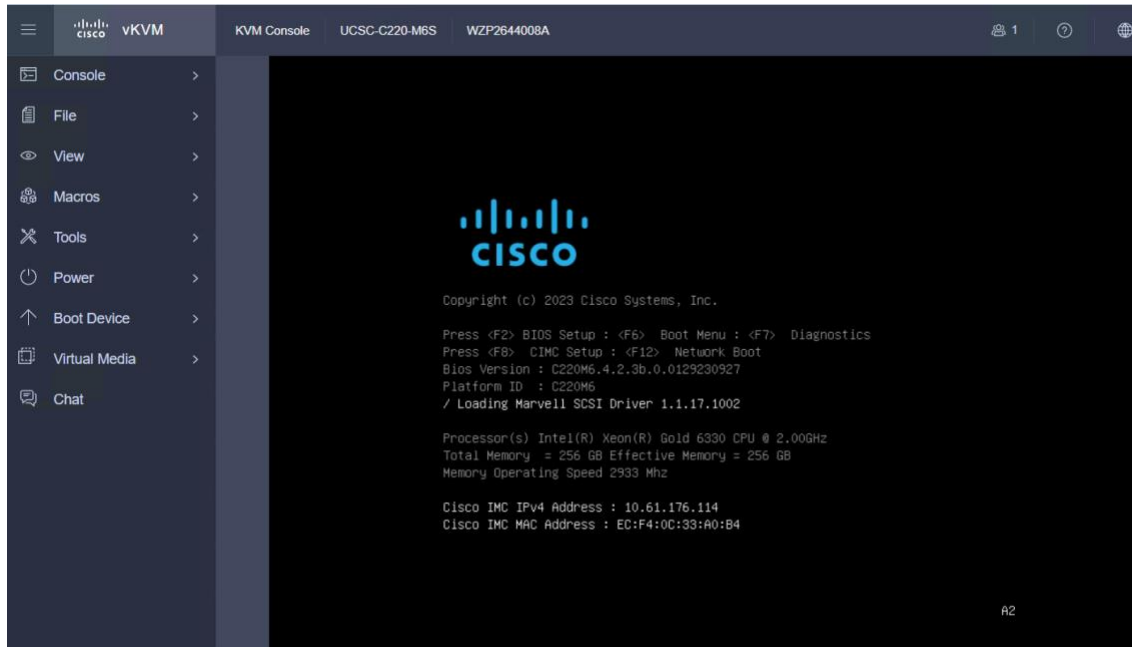
The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

Table 16) Information required for 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' interface in Cisco Intersight. 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 form is divided into two main 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 banner that reads: '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

Add Filter

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

1 of 1

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.

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

Search

Compute / Fabric


☒ Cisco UCS Server (Standalone)
☐ Cisco UCS Domain (Intersight Managed)
☐ HPE OneView

☐ Cisco UCS Domain (UCSM Managed)
☐ Cisco UCS C890
☐ Redfish Server

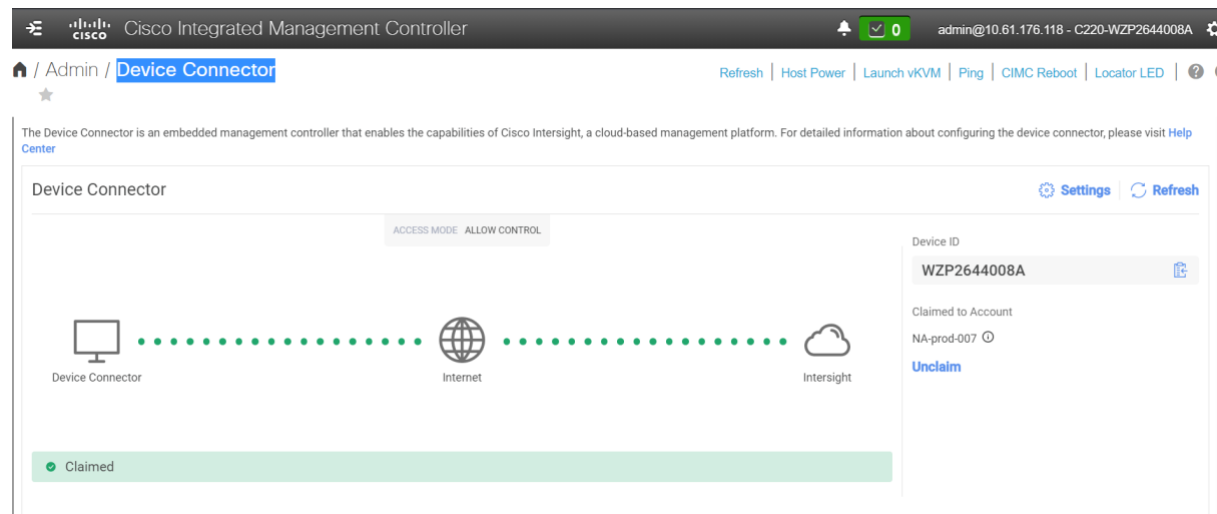
Platform Services

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.

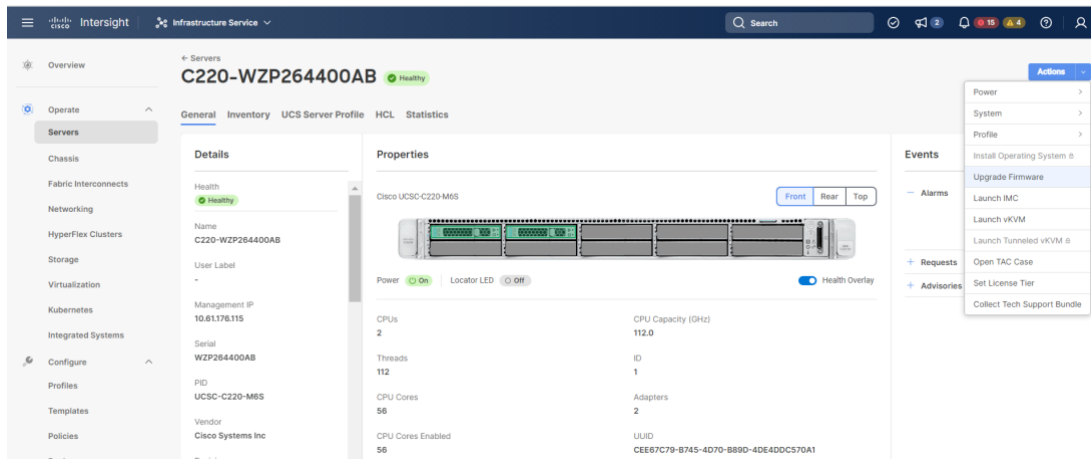


Verify on CIMC, that the server is successfully claimed into Intersight.



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](#)



Create and Deploy Server Profile Template

A server profile template enables resource management by simplifying policy alignment and server configuration. A server profile template is created using the server profile template wizard. The server profile template wizard groups the server policies into the following four categories to provide a quick summary view of the policies that are attached to a profile:

- Compute policies: BIOS, boot order, and virtual media policies, and UUID pool
- Network policies: adapter configuration, LAN connectivity, and SAN connectivity policies
- The LAN connectivity policy requires you to create Ethernet network policy, Ethernet adapter policy, and Ethernet QoS policy
- Storage policies: Local Disk identification and configuration policies
- Management policies: Local user, Network Connectivity, Virtual KVM Policy, NTP Policy

Some of the characteristics of the server profile template for FlexPod are as follows:

- BIOS policy is created to specify various server parameters in accordance with FlexPod best practices and [Cisco UCS Performance Tuning Guides](#).
- Boot order policy defines Local disk identification for M.2 Local Boot, virtual media (KVM mapped DVD), UEFI Shell, a CIMC mapped DVD for OS installation.
- Local user policy is used to enable KVM-based user access.
- Storage policy is used to enable virtual drive creation by identifying the RAID controller.

To Create UCS Server Profile Template, go to Templates under Configure, Click Create UCS Server Profile Template.

- Select the Organization accordingly.
- Enter the Name for Profile Template
- Select Target platform as UCS Server(Standalone)
- Add an optional Description.

The screenshot shows the 'Create UCS Server Profile Template' wizard in the Cisco Intersight interface. The left sidebar has a 'Templates' section highlighted. The main panel is titled 'General' and contains the following fields:

- Organization ***: FlexPod-Express
- Name ***: C220ServerProfileTemplate
- Target Platform**: ☒ UCS Server (Standalone) ☐ UCS Server (FI-Attached)
- Set Tags**: (empty text field)
- Description**: C220-FlexpodExpress

At the bottom right, there is a 'Next' button.

Create Policies for various Infrastructure Services as depicted below.

The screenshot shows the 'Create UCS Server Profile Template' wizard in the Cisco Intersight interface, now on the 'Compute Configuration' tab. The left sidebar shows the 'Compute Configuration' step is selected. The main panel is titled 'Compute Configuration' and contains the following table:

Policy Name	Action
BIOS	
Boot Order	
Firmware	
Persistent Memory	
Virtual Media	

At the bottom right, there are 'Back' and 'Next' buttons.

Create Compute Policies related to Processor

Create BIOS Policy

- Click Select Policy next to BIOS and in the pane on the right, click Create New.
- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220BIOSPolicy).
- Click Next.

On the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M6 BIOS:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html>

Set the parameters below and leave all other parameters set to “platform-default.” Set the below parameters and leave the rest as “default.”


- Memory > NVM Performance Setting: Balanced Profile
- Power and Performance > Enhanced CPU Performance: Auto
- Processor > Energy Efficient Turbo: enabled
- Processor > Processor C1E: enabled
- Processor > Processor C6 Report: enabled
- Server Management > Consistent Device Naming: enabled

Click Create.

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

 The BIOS settings will be applied only on next host reboot.

+ Boot Options

+ Intel Directed IO

+ LOM And PCIe Slots

+ Main

+ Memory

+ PCI

+ Power And Performance

A final alternative, if you have AMD-based UCS C225 or C245 servers, create a BIOS policy named AA02-AMD-M6-Virt-BIOS with the following parameters:

- Memory > NUMA Nodes per Socket: NPS4
- Processor > APBDIS: 1
- Processor > Fixed SOC P-State: P0
- Processor > ACPI SRAT L3 Cache As NUMA Domain: enabled
- Server Management > Consistent Device Naming: enabled

Create BOOT Order Policy

In the current project, we use the Local Boot option with M.2 Drives.

For the Local boot configuration, LAN connectivity policy is used to create four virtual network interface cards (vNICs); two for virtual switch (vSwitch0) and two for Virtual Distributed Switch (vDS0), various policies can also be created for the vNIC configuration.

- Click Select Policy next to Boot Order and then, in the pane on the right, click Create New.

- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220BootOrderPolicy).Click Next.
- For Configured Boot Mode, select Unified Extensible Firmware Interface (UEFI).
- Turn on Enable Secure Boot.
- Click Add Boot Device drop-down list and select Local Disk.
- Provide the Device Name as it appears in the Inventory Tab of the server. Provide the correct PCIe slot details where the Local disk is interfaced to the server. Click Create
- Click Add Boot Device drop-down list and select Virtual Media.
- Provide a device name (for example, KVM-Mapped-ISO) and then, for the subtype, select KVM Mapped DVD. Click Create

Policy Details

Add policy details

All Platforms
UCS Server (Standalone)
UCS Server (FI-Attached)

Configured Boot Mode

☒ Unified Extensible Firmware Interface (UEFI)
☐ Legacy

☒ Enable Secure Boot

Add Boot Device

Local Disk (CiscoBootoptimizedM2Raidcont)

Device Name *

CiscoBootoptimizedM2Raidcont

Slot

MSTOR-RAID

Bootloader Name

Bootloader Description

Bootloader Path

Virtual Media (KVM-Mapped-DVD)

Device Name *

KVM-Mapped DVD

Sub-Type

KVM MAPPED DVD

- Verify the order of the boot policies and adjust the boot order as necessary using arrows next to the Delete button.
- Click Create.

Policy Details

Add policy details

Configured Boot Mode ⓘ

☒ Unified Extensible Firmware Interface (UEFI) ☐ Legacy

☐ Enable Secure Boot ⓘ

[Add Boot Device](#)

+ Local Disk (CiscoBootoptimizedM2Raidcont)	<input checked="" type="checkbox"/> Enabled			
+ Virtual Media (DVD)	<input checked="" type="checkbox"/> Enabled			

Create Virtual Media Policy

Create policy in Intersight to map the ESXi8.0 ISO file onto the KVM-mapped DVD.

- Click Select Policy next to Virtual Media and then, in the pane on the right, click Create New.
- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220VirtualMediaPolicy).
- Turn on Enable Virtual Media, Enable Virtual Media Encryption, and Enable Low Power USB.
- Do not Add Virtual Media at this time, but the policy can be modified and used to map and ISO for a KVM Mapped DVD.
- Click Create.
- Click Next to move to Management Configuration

Policy Details

Add policy details

All Platforms

UCS Server (Standalone)

UCS Server (FI-Attached)

Configuration

☒ Enable Virtual Media ⓘ

☒ Enable Virtual Media Encryption ⓘ

☒ Enable Low Power USB ⓘ

[Add Virtual Media](#)

0 items found10 per page

0 of 0

<input type="checkbox"/>	Name	Type	Protocol	File Location
NO ITEMS AVAILABLE				

0 of 0

Management Policy

Create policies for IPMI Over LAN, Local User, Network connectivity, NTP server information and virtual KVM console.

Create or select existing Management policies that you want to associate with this template.

Device Connector	
IPMI Over LAN	
LDAP	
Local User	● C220LocalUserPolicy
Network Connectivity	● C220NetworkConnectivityPolicy
NTP	● C220NTPPolicy
Serial Over LAN	
SMTP	
SNMP	
SSH	
Syslog	
Virtual KVM	● C220_VirtualKVMPolicy

Close Back Next

Create IPMI Over LAN Policy

- Click Select Policy next to IPMI Over LAN and then, in the pane on the right, click Create New.
- Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220IPMIOverLANPolicy).
- On the right, select UCS Server (Standalone)
- Turn on Enable IPMI Over LAN.
- From the Privilege Level drop-down list, select admin
- Enter 00 into the Encryption Key field.

Policies > IPMI Over LAN

Create

General

2 Policy Details

Policy Details

Add policy details

▼ All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

☒ Enable IPMI Over LAN

Privilege Level
admin

Encryption Key
**

Cancel Back Create

Create Local User Policy

- Click Select Policy next to Local User and then, in the pane on the right, click Create New.

- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220LocalUserPolicy).
- Verify that UCS Server (Standalone) is selected.
- Verify that Enforce Strong Password is selected.
- Click Add New User and then click + next to the New User
- Provide the username (for example, admin), select a role for example, admin), and provide a password.

Note: The username and password combination defined here will be used as an alternate to log in to KVMs.

- Click Create to finish configuring the user.
- Click Create to finish configuring the local user policy.

Policies > Local User > C220LocalUserPolicy

Edit

General

Policy Details

Password Properties

☒ Enforce Strong Password ☐ Enable Password Expiry

0 Password History ☐ Always Send User Password

Local Users

This policy will remove existing user accounts other than the ones configured with this policy. However, the default admin user account is not deleted from the endpoint device. You can only enable/disable or change account password for the admin account by creating a user with the user name and role as 'admin'. If there are no users in the policy, only the admin user account will be available on the endpoint device. By default, IPMI support is enabled for all users

[Add New User](#)

Username	Role	Enable
admin (admin)	admin	<input checked="" type="checkbox"/>

Username * admin Role admin

Password Password Confirmation

[Cancel](#) [Back](#) [Save](#) [Save & Deploy](#)

Create Network Connectivity Policy

- Click Select Policy next to Network Connectivity and then, in the pane on the right, click Create New.
- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220NetworkConnectivityPolicy).
- Click Next.
- Provide DNS server IP addresses for Cisco UCS (for example, 10.61.176.251 and 10.61.176.252)

Policies > Network Connectivity > C220NetworkConnectivityPolicy

Edit

General

Policy Details

Policy Details

Add policy details

All Platforms
UCS Server (Standalone)
UCS Domain

Common Properties

☐ Enable Dynamic DNS ⓘ

IPv4 Properties

☐ Obtain IPv4 DNS Server Addresses from DHCP ⓘ

Preferred IPv4 DNS Server

10.61.176.251 ⓘ

Alternate IPv4 DNS Server

10.61.176.252 ⓘ

☐ Enable IPv6 ⓘ

Cancel
Back Save Save & Deploy

Create NTP Servers Policy

- Click Select Policy next to NTP and then, in the pane on the right, click Create New.
- Verify correct organization is selected from the drop-down list (for example, FlexPod Express) and provide a name for the policy (for example, C220NTPPolicy).
- Click Next.
- Enable NTP, provide the first NTP server IP address, and select the time zone from the drop-down list.
- Add a second NTP server by clicking + next to the first NTP server IP address.

Note: The NTP server IP addresses should be Nexus switch management IPs. NTP distribution was configured in the Cisco Nexus switches.

Policies > NTP > C220NTPPolicy

Edit

General

Policy Details

Policy Details

Add policy details

All Platforms
UCS Server (Standalone)
UCS Domain

☒ Enable NTP ⓘ

NTP Servers *

10.61.176.251 ⓘ ⓘ

NTP Servers *

10.61.176.252 ⓘ ⓘ +

Timezone

Pacific/Nue ⓘ

Cancel
Back Save Save & Deploy

Create Virtual KVM console Policy

- Click Select Policy next to Virtual KVM and then, in the pane on the right, click Create New.
- Verify correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220_VirtualKVMPolicy).

- Verify that UCS Server (Standalone) is selected.
- Turn on “Allow Tunneled vKVM.”
- Click Create.

Note: To fully enable Tunneled KVM, once the Server Profile Template has been created, go to System > Settings > Security and Privacy and click Configure. Turn on “Allow Tunneled vKVM Launch” and “Allow Tunneled vKVM Configuration.”

- Click Next to move to Storage Configuration.

Policy Details

Add policy details

▼ All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

☒ Enable Virtual KVM ⓘ

Max Sessions *

4 ⓘ ⓘ
1 ~ 4

Remote Port *

2068 ⓘ ⓘ
1 ~ 65535

☒ Enable Video Encryption ⓘ

☒ Enable Local Server Video ⓘ

☒ Allow Tunneled vKVM ⓘ

Storage Policy

Storage Policy is configured in the Service Profile to identify the Local disk drive for use.

- Click Select Policy next to Storage and then, in the pane on the right, click Create New.
- Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220StoragePolicy). Click Next.
- Enable M.2 RAID Configuration and Select the correct slot of M.2 RAID controller from the drop-down list.
- Click Create

Policies > Storage > C220_StoragePolicy

Edit

General

2 Policy Details

General Configuration

☐ Use JBOD drives for Virtual Drive creation ⓘ

Unused Disks State
 No Change ⓘ

Default Drive State
 Unconfigured Good ⓘ

Secure JBOD Disk Slots ⓘ

☒ M.2 RAID Configuration

Slot of the M.2 RAID controller for virtual drive creation
 M2TOR-RAID-1 (M2TOR-RAID) ⓘ

☐ MRAID/RAID Controller Configuration

☐ MRAID/RAID Single Drive RAID0 Configuration

Network Policy

Create Adapter Policy

The VIC card is interfaced to the server at PCIe Slot 1. Hence Create VIC Adapter Configuration to discover the VIC Card and its interfaces.

Intersight Infrastructure Service

Policies > Adapter Configuration > C220AdapterConfigurationPolicy

Edit

Overview

Operate

Servers

Chassis

Fabric Interconnects

Networking

HyperFlex Clusters

Storage

Virtualization

Kubernetes

Integrated Systems

Configure

Profiles

Templates

Policies

Pools

General

2 Policy Details

Policy Details

Add policy details

This policy is applicable only for UCS Servers (Standalone)

Adapter Configurations

[Add VIC Adapter Configuration](#)

PCI Slot	LLDP	FIP	Port Channel	
<input checked="" type="checkbox"/> 1	Enabled	Disabled	Disabled	...

Selected 1 of 1 [Show Selected](#) [Unselect All](#)

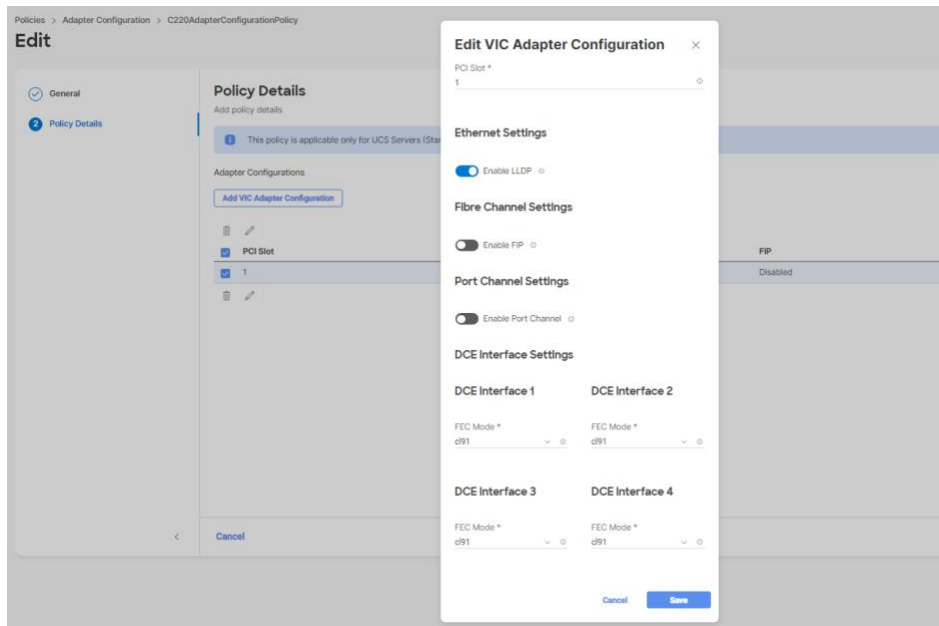
[Export](#) ⓘ

Cancel [Back](#) [Save](#) [Save & Deploy](#)

Activate Windows
Go to Settings to activate Windows.

To create the VIC Adapter Configuration, follow the below steps:

- Click Select Policy next to Adapter Configuration and then, in the pane on the right, click Create New.
- Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220AdapterPolicy). Click Next.
- Add VIC Adapter Configuration
- Add the PCIe Slot details.



- Under Ethernet Settings, Enable LLDP
- Under Fibre Channel Settings, Enable FIP
- Under Port Channel Settings, Disable Port Channel.

NOTE: Enabling Port Channel here, enables the Hardware Port-channel that causes Port Binding. Due to port binding, only one port is made available to Nexus switches. Since we need total of 4 Eth ports, we will configure software based Port-channel from VMware side with the Nexus ports.

- Under DCE Interface Settings, Leave the FEC Mode set to cl91

Note: FEC mode settings need to be aligned with switch where cl91 is equal to rs-fec and cl74 equals to the fc-fec.

Create LAN Connectivity Policy

Policy to identify the Ethernet ports of the VIC card is added in the LAN connectivity policy.

- Click Select Policy next to LAN Connectivity and then, in the pane on the right, click Create New.
- Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220LANConnectivityPolicy). Click Next.
- Select UCS Server (Standalone) for Target Platform, Click Next
- Click on Add vNIC.
- Enter the name, appropriate slot ID, uplink Port, PCI Link and PCI Order.

vNIC Name	Slot ID	Uplink Port	PCI Order
eth0	1	0	0
eth1	1	1	1
eth2	1	2	2
eth3	1	3	3

- Also add Consistent Device Naming (CDN), Select User Defined from the drop down and provide the

below value.

General Name	CDN Value
eth0	00-vSwitch0-A
eth1	01-vDS0-A
eth2	02-vSwitch0-B
eth3	03-vDS0-B

- Add Ethernet Network, Ethernet QoS, and Ethernet Adapter Policies for the interfaces accordingly.
- To add Ethernet Network Policy

Create two Ethernet Network Policy, one for Standalone vNICs that are part of Standard vSwitch and one for vNICs that are part of Distributed Switch (vDS0)

 - For Standalone vNICs (eth0 and eth2)
 - Click on Select Policy, Create New
 - Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220EthNetworkPolicyStandalone). Click Next.
 - Select VLAN Mode Trunk and Enter <ib_mgmt_vlan_id> under default VLAN, Click Create.
 - For vDS vNICs (eth1 and eth3)
 - Click on Select Policy, Create New
 - Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220EthNetworkPolicyvDS). Click Next.
 - Select VLAN Mode Trunk and Enter Native VLAN ID under default VLAN, Click Create.
- To add Ethernet QoS Policy
 - Click on Select Policy, Create New
 - Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220EthQoSPolicy). Click Next.
 - Enter 9000 for MTU, Bytes field, Click create
- To add Ethernet Adapter Policy
 - Click on Select Policy, Create New
 - Verify the correct organization is selected from the drop-down list (for example, FlexPodExpress) and provide a name for the policy (for example, C220EthAdpaterPolicy). Click Next.
 - Transmit Ring Size set to 4096
 - Receive Ring Size set to 4096
 - Tx Queue set to 1
 - Rx Queue set to 16
 - Complete Queues set to 17
 - Interrupts set to 19
 - Enable Receive Side Scaling

Edit

✓ General

2 Policy Details

Policy Details

Add policy details

i At a minimum two vNICs are required named eth0 and eth1. Learn more at [Help Center](#)

Add vNIC

	Name	Slot ID	Uplink Port	PCI Order	
<input type="checkbox"/>	eth0	1	0	0	...
<input type="checkbox"/>	eth1	1	1	1	...
<input type="checkbox"/>	eth2	1	2	2	...
<input type="checkbox"/>	eth3	1	3	3	...

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

Overview Templates

C220ServerProfileTemplate

Details

Name: C220ServerProfileTemplate

Target Platform: UCS Server (Standalone)

Last Update: Apr 17, 2023 10:16 PM

Description: C220 Server Profile Standalone

Organization: G14C220M6_Standalones

Server Assignment

Tags: No Tags

Configuration

Configuration Usage

Compute

BIOS: C220_BIOSPolicy

Boot Order: C220_BOOTPolicy

Virtual Media: G14C220_VirtualMediaPolicy

Management

Local User: C220LocalUserPolicy

Network Connectivity: C220NetworkConnectivityPolicy

NTP: C220NTPPolicy

Virtual KVM: C220_VirtualKVMPolicy

Network

Actions: Edit, Clone, Delete, Derive Profiles

Overview Server Profiles

C220ServerProfile_DERIVED-2

General Server Inventory

Details

Status: OK

Name: C220ServerProfile_DERIVED-2

Target Platform: UCS Server (Standalone)

Template Name: C220ServerProfileTemplate

Last Update: Apr 17, 2023 9:50 PM

Description: C220 Server Profile Standalone

Organization: G14C220M6_Standalones

Server Assignment

Configuration

General Identifiers Connectivity

Adapter Configuration: C220AdapterConfigurationPolicy

BIOS: C220_BIOSPolicy

Boot Order: C220_BOOTPolicy

LAN Connectivity: C220LANConnectivityPolicy

Local User: C220LocalUserPolicy

Network Connectivity: C220NetworkConnectivityPolicy

NTP: C220NTPPolicy

Storage: C220_StoragePolicy

Virtual KVM: C220_VirtualKVMPolicy

Virtual Media: G14C220_VirtualMediaPolicy

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.

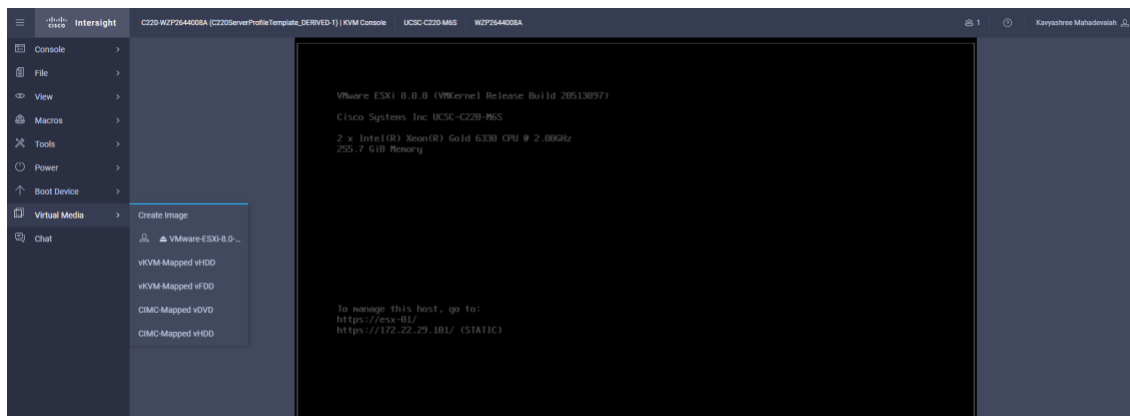
Set up VMware ESXi Installation

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.



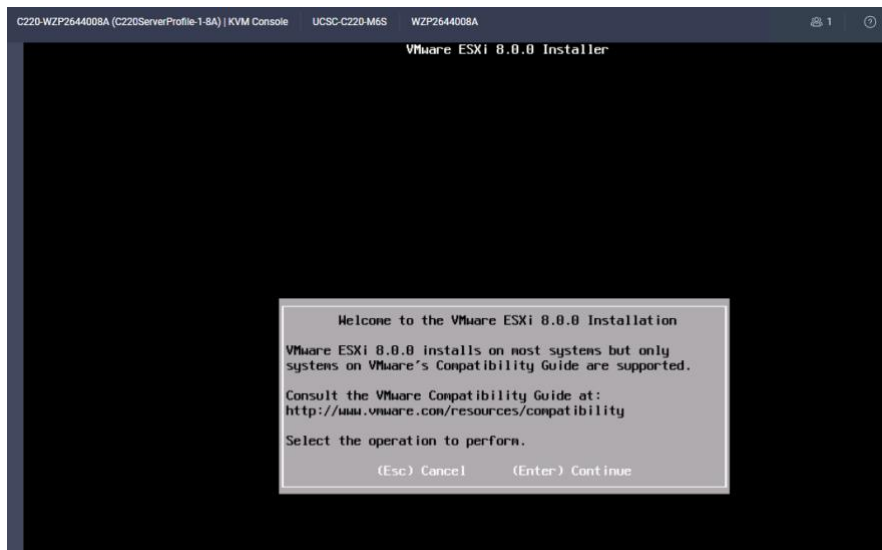
Install ESXi

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:

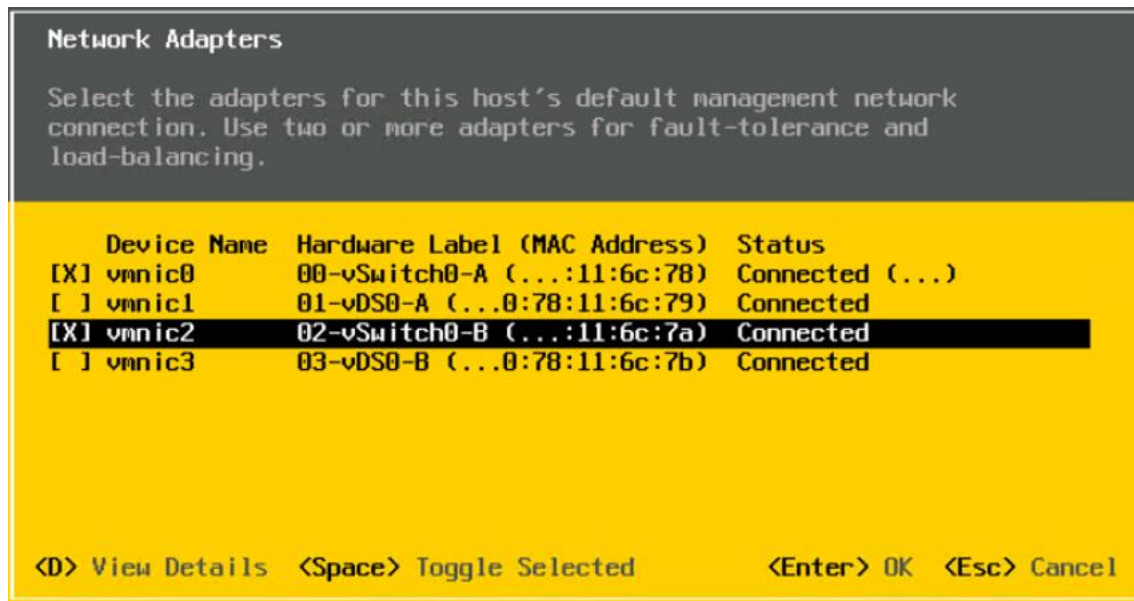
ESXi esxi-01 and esxi-02

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-vmknic -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-vmknic -d "Management Network"; esxcfg-vmknic -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-vmknic -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.

Log in to VMware ESXi hosts by using VMware host client

ESXi Host esxi-01

To log in to the esxi-01 ESXi host by using the VMware Host Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the esxi-01 management IP address.
2. Click Open the VMware Host Client.
3. Enter `root` for the username.
4. Enter the root password.
5. Click Login to connect.
6. Repeat this process to log in to esxi-02 in a separate browser tab or window.

Install Cisco VIC Drivers and NetApp NFS Plug-in for VAAI

ESXi Hosts esxi-01 and esxi02

1. Download Drivers to the Management Workstation

Download and extract where necessary the following drivers to the Management Workstation

Note: Consult the [Cisco UCS Hardware Compatibility List](#) and the [NetApp Interoperability Matrix Tool](#) to determine latest supported combinations of firmware and software.

- Follow the link [VIC driver](#) to download the nenic driver for Cisco VIC cards of version 1.0.45.0

Note: The VMware ESXi 8.0 Cisco Custom ISO contains the nenic driver version 1.0.45.0. It is not necessary to download or update the **nenic driver**, but the commands are left here to be used for future updates.

Note: Refer the document here for more details on [Downloading and installing Cisco VIC Drivers](#)

- Download the latest NetApp NFS Plug-in for VMware VAAI from the NetApp site <https://mysupport.netapp.com/site/products/all/details/nfsplugin-vmware-vaai>. Get the NetAppNasPlugin2.0.1.zip file extracted from the downloaded zip.

2. To install the Drivers on the ESXi hosts, follow these steps:

- Login to first ESXi host and copy the extracted zip files into the NFS mapped datastore using an SCP program such as WinSCP, copy the offline bundle referenced above to the /tmp directory on each ESXi host.
- Using a secure shell (SSH) tool such as PuTTY, SSH to each VMware ESXi host. Log in as root with the root password.
- Enter `cd /tmp`.
- Run the following commands on each host

```
esxcli software vib update -d /tmp/Cisco-nenic_1.0.45.0-10EM.700.1.0.15843807-offline_bundle-16216785.zip

esxcfg-advcfg -s 0 /Misc/HppManageDegradedPaths

esxcli software vib install -d /tmp/NetAppNasPlugin2.0.1.zip
```

Note: HppManageDegradedPaths configuration is needed for NVMe-TCP configuration.

- Use the following commands to ensure the correct Driver version are installed.

```
esxcli software component list | grep nenic

esxcfg-advcfg -g /Misc/HppManageDegradedPaths

esxcli software vib list | grep NetApp
```

Set Up VMkernel ports and Virtual Switch

ESXi Host esxi-01 and esxi-02

To set up the VMkernel ports and the virtual switches on the ESXi hosts, complete the following steps.

1. From the Host Client, select Networking on the left.
2. In the center pane, select the Virtual Switches tab.
3. Select vSwitch0.
4. Select Edit settings.

5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover Order section, select the vmnic that has the Standby status and click Mark Active. Verify that vmnic now have the status of Active.
8. Click Save.

Edit standard virtual switch - vSwitch0

Add uplink

MTU	9000
Uplink 1	vmnic0 - Up, 25000 Mbps ✕
Uplink 2	vmnic2 - Up, 25000 Mbps ✕
> Link discovery	Click to expand
> Security	Click to expand
> NIC teaming	Click to expand
> Traffic shaping	Click to expand

CANCEL **SAVE**

9. Select Networking on the left.
10. Go to Networking > Port Groups.
11. In the center pane, right-click VM Network and select Edit settings.
12. Change the Name of the port group to IB-MGMT Network and leave the VLAN ID set to 0.

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 vNICs. Because of this, the IB-MGMT VLAN should stay set to 0.

13. Click Save to finalize the edits for the IB-MGMT Network port group.
14. At the top, select the VMkernel NICs tab.
15. Click Add VMkernel NIC.
16. For New Port Group, enter `nfs`

+ Add port group - nfs

Name	nfs
VLAN ID	2230
Virtual switch	vSwitch0
> Security	Click to expand

CANCEL **ADD**

17. For Virtual Switch, select vSwitch0 Selected.

NOTE: Once vCenter VM is deployed, we will migrate nfs VMkernel port onto Distributed switch. For now, the port group is added into Standard switch.

18. Enter <infra_nfs_vlan_id> for the VLAN ID

19. Change the MTU to 9000.

20. Select Static IPv4 settings and expand IPv4 settings.

21. Enter the ESXi host infrastructure NFS IP address and network mask.

22. Do not select any of the services.

23. Click Create.

+ Add VMkernel NIC

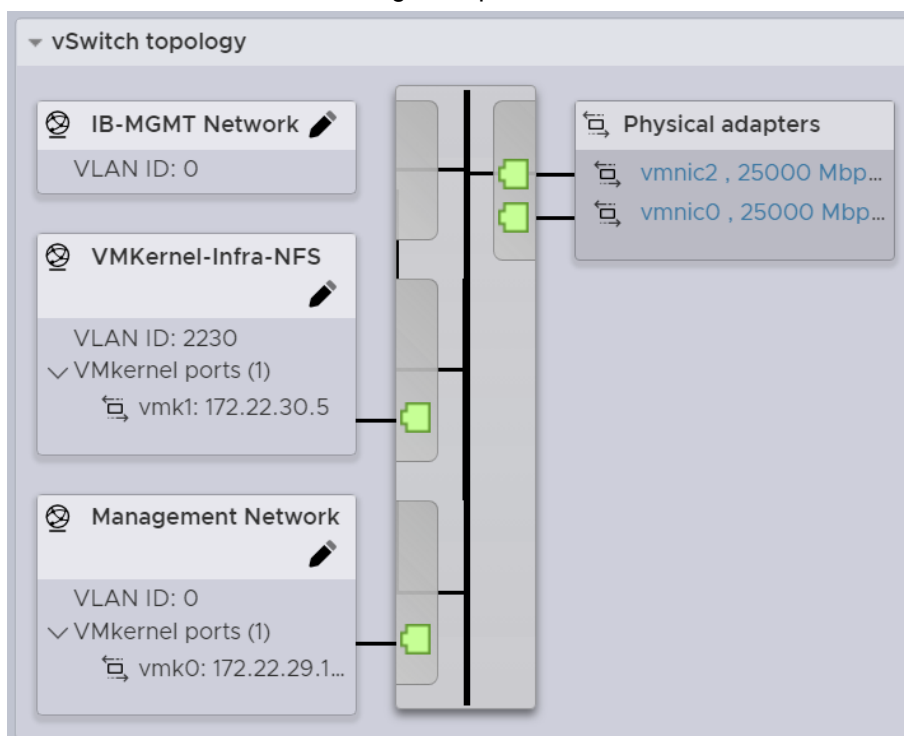
Port group	nfs
MTU	9000
IP version	IPv4 only
IPv4 settings	<p>Configuration: <input type="radio"/> DHCP <input checked="" type="radio"/> Static</p> <p>Address: 172.22.30.2</p> <p>Subnet mask: 255.255.255.0</p> <p>TCP/IP stack: Default TCP/IP stack</p> <p>Services: <input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication</p>

CANCEL **CREATE**

24. Again, select Networking.
25. In the center pane, select Port groups.
26. Right click on nfs port group and click on edit settings.
27. Expand NIC teaming section and verify that both vmnic0 and vmnic2 are set to active.

Failover order		
<input checked="" type="checkbox"/> Mark standby <input type="checkbox"/> Mark unused <input type="button" value="↑ Move up"/> <input type="button" value="↓ Move down"/>		
Name	Speed	Status
vmnic0	25000 Mbps, full duplex	Active
vmnic2	25000 Mbps, full duplex	Active

28. Click Save.
29. Select the Virtual Switches tab, then select vSwitch0. The properties for vSwitch0 VMkernel NICs should be like the following example



Mount required Datastores.

ESXi hosts esxi-01 and esxi-02

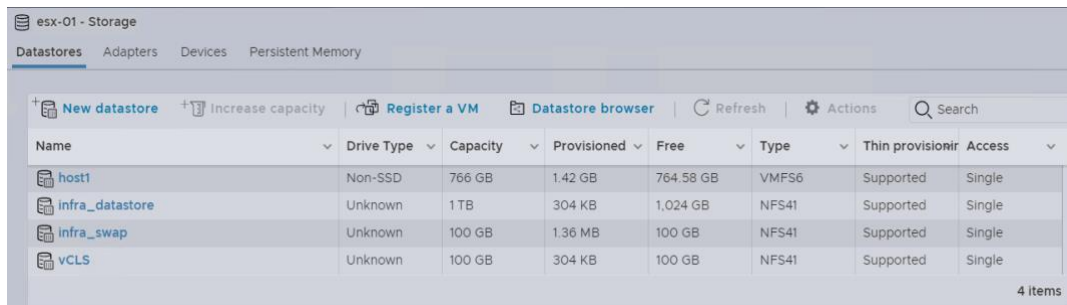
To mount the required datastores, complete the following steps on each ESXi host:

From the Host Client, select **Storage** on the left.

1. In the center pane, select **Datastores** tab.
2. Click the **New Datastore** icon to add a new datastore.
3. In the New Datastore dialog box, select **Mount NFS Datastore** and click **Next**.
 On the provide NFS Mount Details page, complete these steps:
 - a. Enter `infra_datastore` for the datastore name.
 - b. Enter the IP address for the `nfs-lif-01` LIF for the NFS server.

Note: Use the NFS LIF that resides on the same node as the NFS volume being accessed for the NFS client to have direct access to the NFS volume.

- c. Enter `/infra_datastore` for the NFS share.
 - d. Leave the NFS version set at NFS 4.
 - e. Click Next.
 - f. Click **Finish**. The datastore should now appear in the datastore list.
4. In the center pane, click the **New Datastore** icon to add a new datastore.
 5. In the New Datastore dialog box, select **Mount NFS Datastore** and click **Next**.
On the Provide NFS Mount Details page, complete these steps:
 - a. Enter `infra_swap` for the datastore name.
 - b. Enter the IP address for the `nfs-lif-01` LIF for the NFS server.
 - c. Enter `/infra_swap` for the NFS share.
 - d. Leave the NFS version set at NFS 4.
 - e. Click Next.
 - f. Click **Finish**. The datastore should now appear in the datastore list.
 6. In the center pane, click the **New Datastore** icon to add a new datastore.
 7. In the New Datastore dialog box, select **Mount NFS Datastore** and click **Next**.
On the Provide NFS Mount Details page, complete these steps:
 - a. Enter `vCLS` for the datastore name.
 - b. Enter the IP address for the `nfs-lif-01` LIF for the NFS server.
 - c. Enter `/vCLS` for the NFS share.
 - d. Leave the NFS version set at NFS 4.
 - e. Click Next.
 - f. Click **Finish**. The datastore should now appear in the datastore list.



The screenshot shows the 'esx-01 - Storage' view in vSphere. The 'Datastores' tab is selected. A table lists four datastores: host1, infra_datastore, infra_swap, and vCLS. The table columns are Name, Drive Type, Capacity, Provisioned, Free, Type, Thin provisioning, and Access.

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
host1	Non-SSD	766 GB	1.42 GB	764.58 GB	VMFS6	Supported	Single
infra_datastore	Unknown	1 TB	304 KB	1,024 GB	NFS41	Supported	Single
infra_swap	Unknown	100 GB	1.36 MB	100 GB	NFS41	Supported	Single
vCLS	Unknown	100 GB	304 KB	100 GB	NFS41	Supported	Single

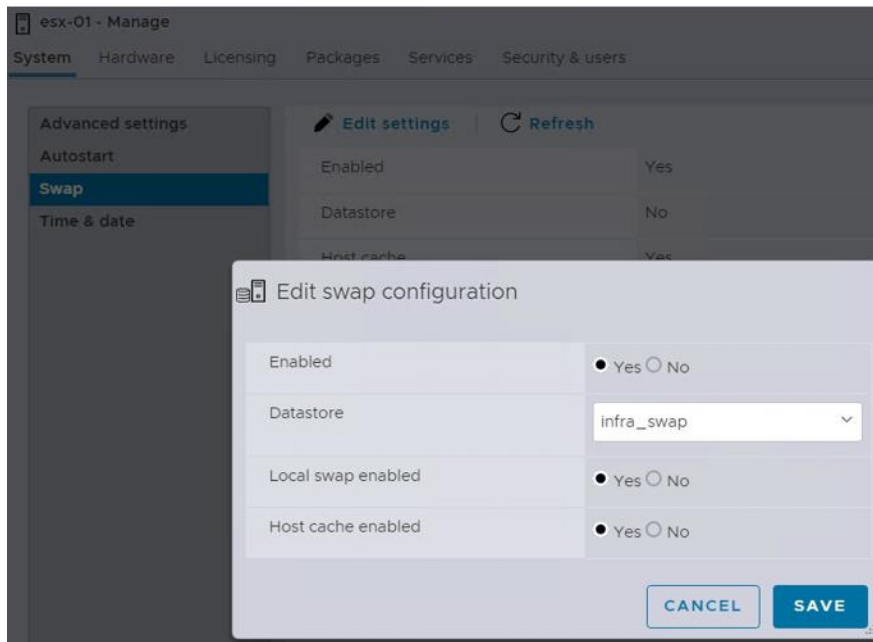
4 items

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.

Configure NTP on ESXi hosts

ESXi hosts esxi-01 and esxi-02

To configure NTP on the ESXi hosts, complete the following steps on each host:

1. From the host client, select Manage on the left.
2. In the center pane, select the Time & Date under the System tab.
3. Click Edit NTP Settings.
4. Use Network Time Protocol (Enable NTP Client) is selected.
5. Use the drop-down menu to select Start and Stop with Host.
6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.
7. Click Save to save the configuration changes.
8. In the center pane, click Services tab.
9. Click to select the ntpd service row.
10. Click Start to start the service.
11. Wait for the screen to refresh and check to make sure that the ntpd service status indicates Running.
12. Go back to the System tab and select Time & Date.
13. Click Refresh and verify that NTP service is running, and the current date and time is accurate.

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).

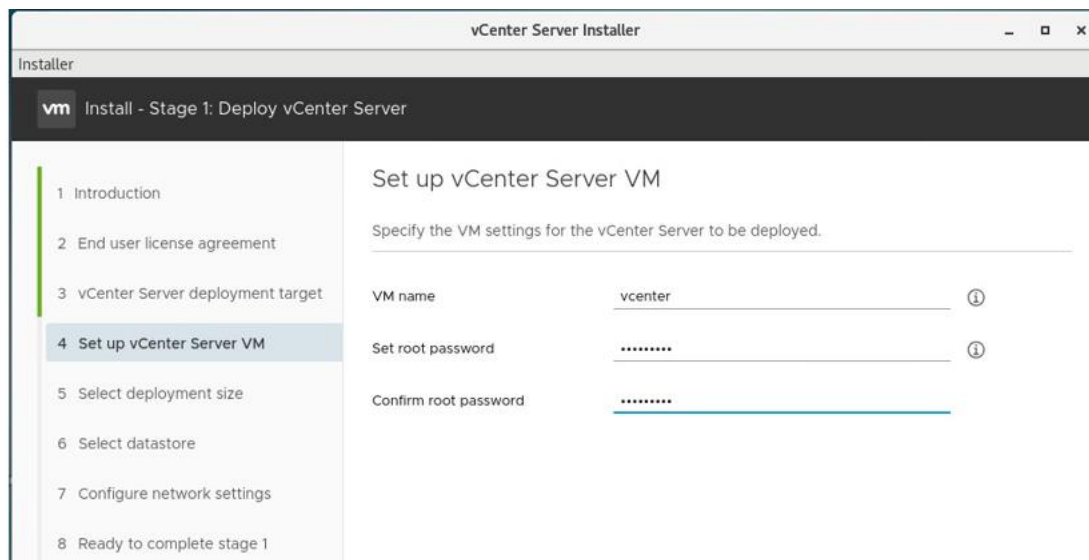
Install VMware vCenter server appliance.

To install VCSA, complete the following steps:

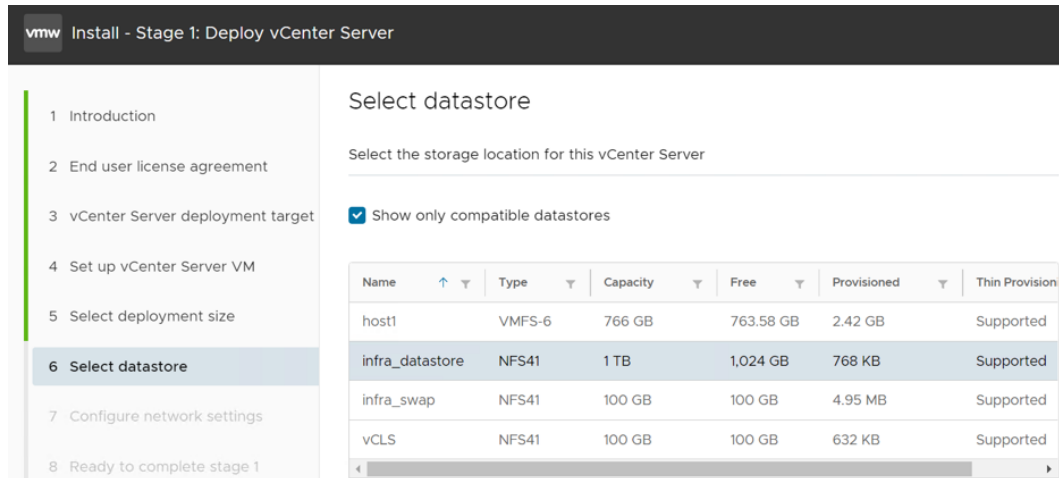
1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.
2. Download the VCSA from the VMware site.
3. Mount the ISO image on your management workstation.
4. Navigate to the installer appropriate for your environment.
5. For installing from Windows, 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.



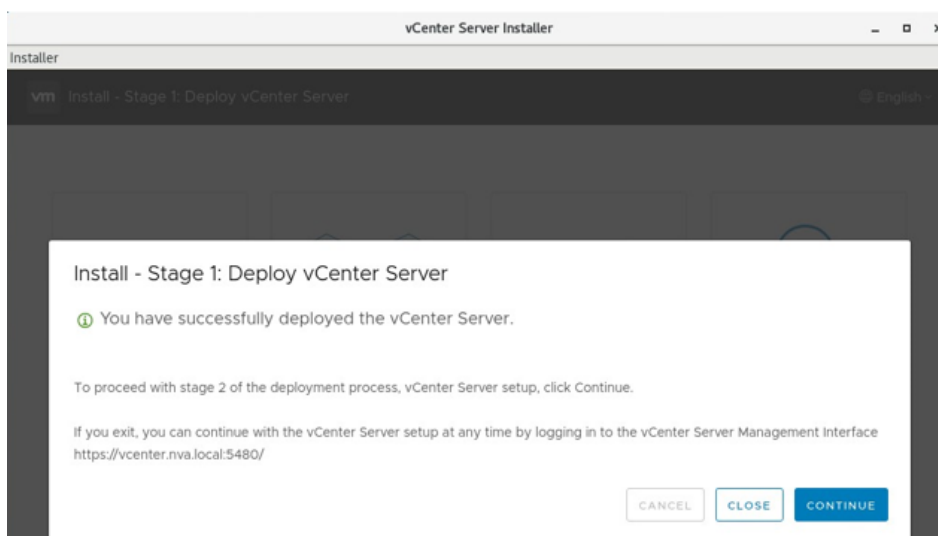
13. Select the deployment size and storage size that are suitable for your deployment. For example, choose Tiny and Default.
14. Click Next.



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.

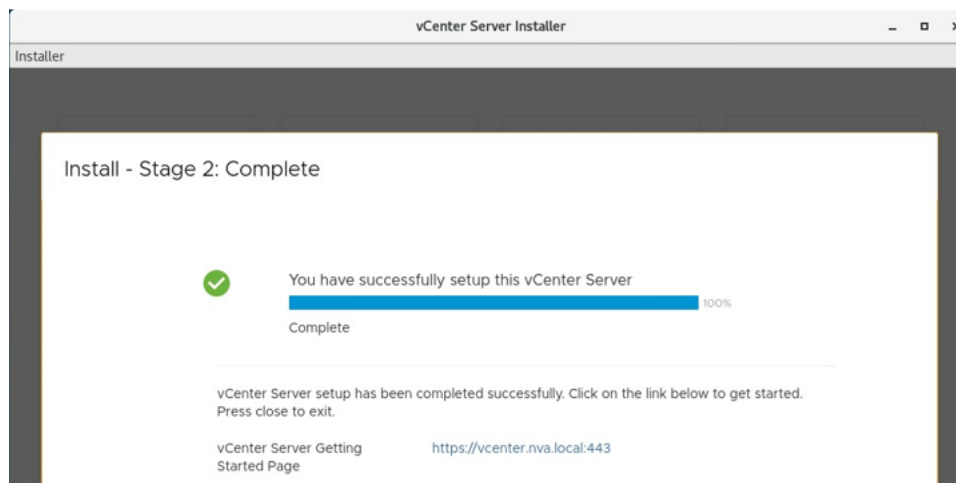
Note: The VCSA installation process takes several minutes.

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.

Configure VMware vCenter Server 8.0 and vSphere Clustering.

To configure VMware vCenter Server 8.0 and vSphere clustering, complete the following steps:

1. Go to <https://<FQDN or IP of vCenter>>.
2. Click Launch vSphere Client (HTML5).
3. Log in with the username [administrator@<SSO domainname>](#) and the SSO password you entered during the vCenter server setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center.
6. Click OK.

Create Cluster

To create a vSphere cluster, complete the following steps:

1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Select and enable DRS and vSphere HA options. Do not turn on vSAN. Click Next.
4. Select Import image from a new host for setting up the Cluster image. Enter the host FQDN, username

and password details and accept the certificate message. Observe the Image on selected host output.

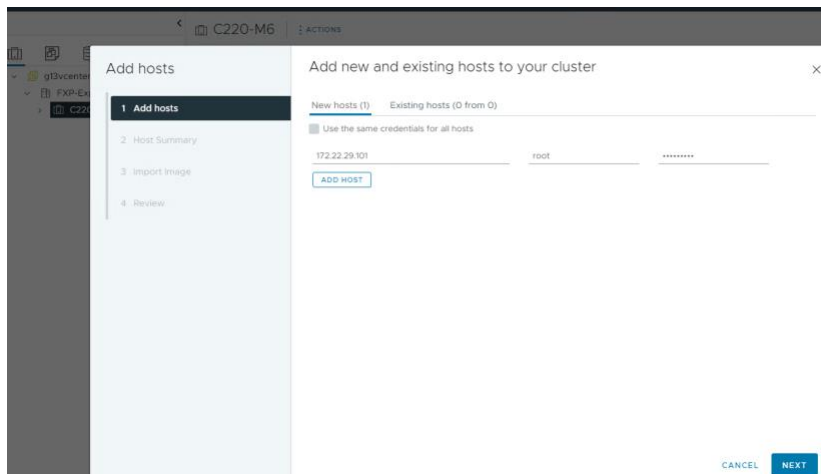
Leave the “Also move selected host to cluster” option selected. Click Next

5. Review the Cluster settings and Click Finish.
6. Expand the FlexPod Express datacenter, right click the newly added cluster and select Settings.
7. 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.
8. Select Datastore Specified by Host Option.
9. Click OK.

Add ESXi Hosts to Cluster

To add ESXi hosts to the cluster, complete the following steps:

1. Select Add Host in the Actions menu of the cluster.
2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter IP or FQDN of the new host.
 - b. Check Use the Same Credentials for All Hosts, if desired.
 - c. Enter the root username and password for the host.
 - d. Enter IP or FQDN of additional hosts.
 - e. Click Next



3. In the Host Summary dialog, expand the first host to see the warning about the host having powered on VM.
4. Click Next.
5. Review the information in the Review and finish dialog.
6. Click Finish.
7. Expand the cluster to see the hosts added to the cluster.
8. You can suppress the warning on ESXi shell and SSH being enabled in the summary tab.

Take host out of maintenance mode

After the hosts are added to the cluster, they might be placed into maintenance mode. As a result, vCenter could report a warning indicating insufficient vSphere HA failover resources.

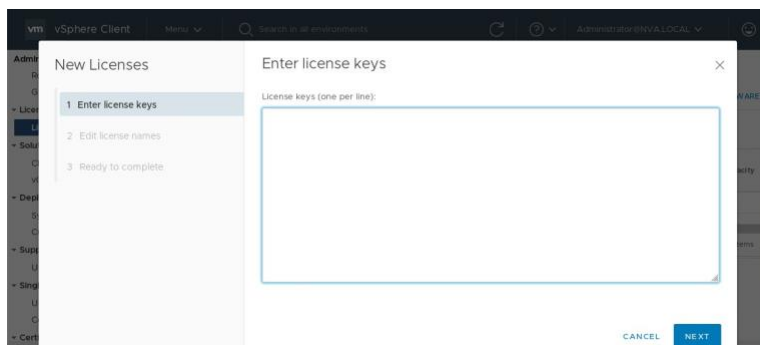
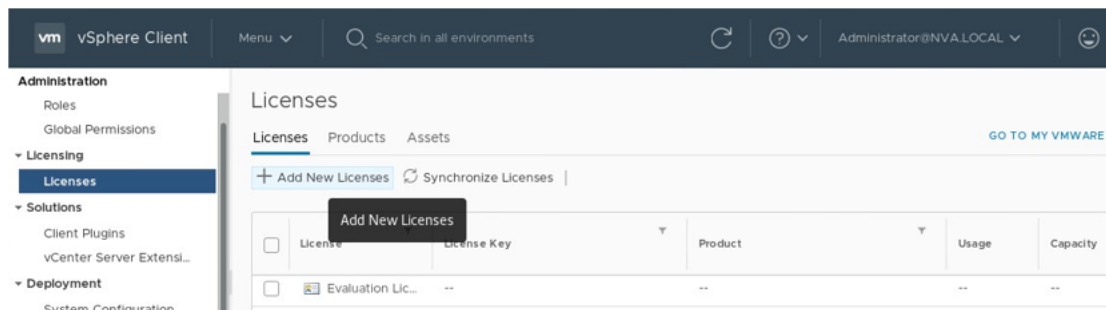
To take an ESXi host out of maintenance mode, complete the following steps:

1. Right click on the host in the cluster and select Exit Maintenance Mode under the Maintenance Mode menu.
2. After a few minutes, the alarm, indicated as a red dot on the data center, should be cleared.

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.




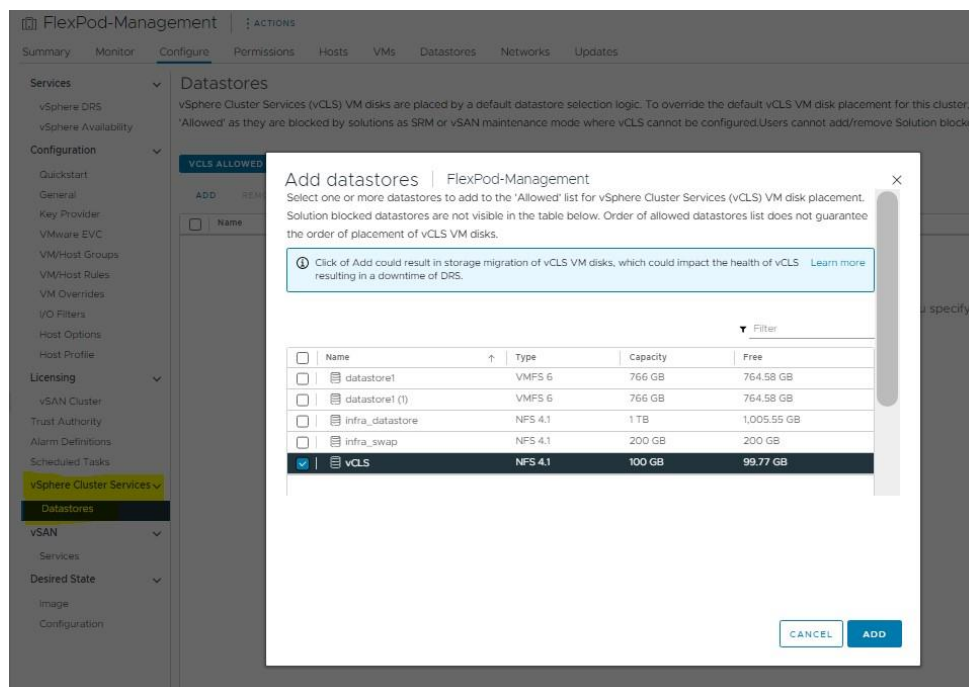
5. Type in the license keys, one per line, in the dialog box and click Next.
6. Edit License Name, if needed, and click Next.
7. Click Finish to complete entering licenses.
8. Right-click the vCenter server under Hosts and Clusters and select Assign License.
9. Select the vCenter license and click OK.
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.

Configure ESXi Host Swap and Mount Required Datastores

To mount the required datastores, follow these steps on the ESXi host(s):

1. Right-click the added ESXi host(s) and click Settings.

2. In the center pane under Virtual Machines, click Swap File location.
3. On the right, click EDIT.
4. Select infra_swap datastore and click OK.
5. Repeat steps 1-4 to set the swap file location for each configured ESXi host.
6. From the vCenter Home screen, choose  symbol to go to Storage.
7. Located on the left, expand FlexPodExpress-DC.
8. Located on the left, right-click infra_datastore and choose Mount Datastore to Additional Hosts.
9. Choose the ESXi host(s) and click OK
10. Repeat this process to mount the infra_swap and vCLS datastores are mounted to all the ESXi hosts.
11. Right-click the cluster and select Settings. In the center pane under vSphere Cluster Services, select Datastores. In the center of the window, click ADD. Select the vCLS datastore and click ADD.




12. In the center pane, choose Hosts. Verify the ESXi hosts now have the datastores mounted.

Create and Configure VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for setting up VMware vDS in vCenter. NFS, vMotion, NVMe over TCP A/B, and VM-Traffic port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would require changes in the Cisco Nexus 9K switches, and possibly the NetApp storage cluster. All VLANs are allowed in Intersight. Link Aggregation Group will be configured on vDS to form port channel between the servers and Nexus switches.

In this document, the infrastructure ESXi management VMkernel ports, the In-Band management interfaces including the vCenter management interface are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down.

Configure the VMware vDS in vCenter

1. After logging into the VMware vSphere HTML5 Client, select Inventory under the top-level menu.
2. Click , the fourth icon at the top, to go to Networking.

3. Expand the vCenter and right-click the FlexPod-Express datacenter and click Distributed Switch > New Distributed Switch.
4. Give the Distributed Switch a descriptive name (for example, vDS0) and click NEXT.
5. Make sure version 8.0.0 - ESXi 8.0 and later is selected and click NEXT.
6. Change the Number of uplinks to 2, Disable Network I/O Control. Enter VM-Traffic for the Port group name. Click NEXT.
7. Review the information and click FINISH to complete creating the vDS.

Configure LACP on vDS

New Link Aggregation Group
×

Name
lag1

Number of ports
2

Mode
Active

Load balancing mode
Source and destination IP address and TCP/UD

Timeout mode
Slow

Port policies

You can apply VLAN and NetFlow policies on individual LAGs within the same uplink port group. Unless overridden, the policies defined at uplink port group level will be applied.

VLAN trunk range
☐ Override
0-4094

NetFlow
☐ Override
Disabled

CANCEL

OK

1. On the vDS0, Select Configure > Settings and click on LACP. Create NEW LAG.
 - a. Provide name
 - b. Enter the number of ports 2
 - c. Select Mode to Active
 - d. Load balancing mode to Source and destination IP address and TCP/UDP port
 - e. Leave the rest and click OK
2. Expand the FlexPod-DC datacenter and the newly created vDS. Click the newly created vDS.
3. Right-click the vDS and click Settings > Edit Settings.
4. In the Edit Settings window, click the Advanced tab.
5. Change the MTU to 9000.
6. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

Distributed Switch - Edit Settings | vDS01 X

General
Advanced
Uplinks

MTU (Bytes)
9000

Multicast filtering mode
IGMP/MLD snooping

Discovery protocol

Type
Link Layer Discovery Protocol

Operation
Both

Administrator contact

Name

Other details

CANCEL
OK

Create port groups on vDS

1. To create the NFS port group, right-click the vDS, and select Distributed Port Group > New Distributed Port Group.
2. Enter NFS as the name and click NEXT.
3. Set the VLAN type to VLAN, enter the VLAN ID used for NFS (for example, 2230), check the box for Customize default policies configuration, and click NEXT.
4. Leave the Security options set to Reject and click NEXT.
5. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.
6. Move UP lag1 to the Active uplinks and click MOVE DOWN twice to place Uplink 1 and Uplink 2 in the list of Unused uplinks. This will use LACP. Click OK

7. To create the vMotion port group, right-click the vDS, select Distributed Port Group > New Distributed Port Group.
8. Enter vMotion as the name and click NEXT.
9. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion (for example, 2232), check the box for Customize default policies configuration, and click NEXT.
10. Leave the Security options set to Reject and click NEXT.
11. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.
12. Move UP lag1 to the Active uplinks and click MOVE DOWN twice to place Uplink 1 and Uplink 2 in the list of Unused uplinks. This will use LACP.
13. Click NEXT.
14. Leave NetFlow disabled and click NEXT.
15. Leave Block all ports set as No and click NEXT.
16. Confirm the options and click FINISH to create the port group.
17. To create the NVMe-TCP port groups, right-click the vDS, select Distributed Port Group > New Distributed Port Group.
18. Enter NVMe-TCP-A as the name and click NEXT.
19. Set the VLAN type to VLAN, enter the VLAN ID used for NVMe-TCP-A (for example, 2233), check the box for Customize default policies configuration, and click NEXT.
20. Leave the Security options set to Reject and click NEXT.
21. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.
22. Move UP lag1 to the Active uplinks and click MOVE DOWN twice to place Uplink 1 and in the list of Unused uplinks. This will use LACP
23. Repeat steps 17 to 22 for the NVMe-TCP-B port group with VLAN 2234.
24. To create the VM-Traffic port group, right-click the vDS, select Distributed Port Group > New Distributed Port Group.
25. Enter VM-Traffic as the name and click NEXT.
26. Right-click the VM-Traffic port group and click Edit Settings.
27. Select VLAN.
28. Select VLAN for VLAN type and enter the VM-Traffic VLAN ID (for example, 2031). Click OK.

Add ESXi hosts to vDS

1. Right-click the vDS and click Add and Manage Hosts.
2. Make sure Add Hosts is selected and click NEXT.
3. Click SELECT ALL to select all ESXi hosts. Click NEXT.

Note: If all hosts had an alignment in the ESXi console screen between vmnic numbers and vNIC numbers, leave Adapters on all hosts selected.

Manage physical adapters



Add or remove physical network adapters to this distributed switch.

Adapters on all hosts Adapters per host

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

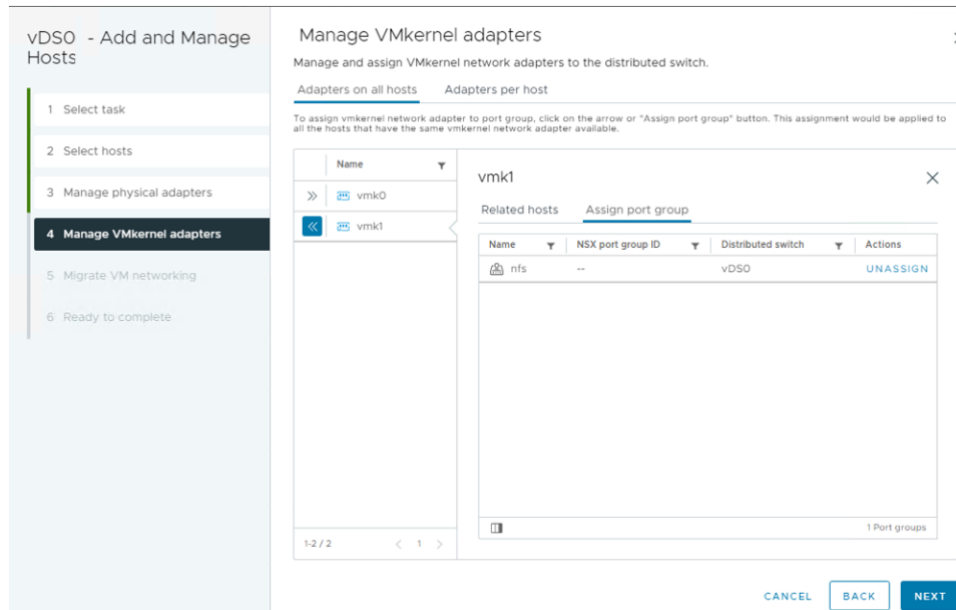
	Physical network adapters ▼	In use by switch	Assign uplink ▼
>>	vmnic0	1 host / 1 switch	None
>>	vmnic1	This switch	lag1-0
>>	vmnic2	1 host / 1 switch	None
>>	vmnic3	This switch	lag1-1
4 physical network adapters			

CANCEL

BACK

NEXT

4. To the right of vmnic1, use the pulldown to select lag1-0 and for vmnic3, use the pulldown to select lag1-1 and Click NEXT.
- Note:** It is important to assign the uplinks as shown above. This allows the port groups to use LAG instead of individual ink.
5. 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.

6. Do not migrate any virtual machine networking ports. Click NEXT.
7. Click FINISH to complete adding the ESXi host to the vDS.
8. 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.

Add VMkernel ports for each of the hosts in vDS

1. Select Hosts and Clusters and select the first ESXi host. In the center pane, select the Configure tab.
2. In the list under Networking, select VMkernel adapters.
3. Select ADD NETWORKING.
4. In the Add Networking window, ensure that VMkernel Network Adapter is selected and click NEXT.
5. Ensure that select an existing network is selected.
6. Select vMotion and click OK.
7. Click NEXT.
8. From the MTU drop-down list, select Custom and ensure the MTU is set to 9000.
9. From the TCP/IP stack drop-down list, select default.
10. Under Enabled Services, Enable vMotion and Provisioning services. Click NEXT.

Add Networking

1 Select connection type

2 Select target device

3 Port properties

4 IPv4 settings

5 Ready to complete

Port properties

Specify VMkernel port settings.

Network label

vMotion (vDSO)

MTU

Custom

9000

TCP/IP stack

Default

Available services

Enabled services

☒ vMotion

☒ Provisioning

☐ Fault Tolerance logging

☐ Management

☐ vSphere Replication

☐ vSphere Replication NFC

☐ vSAN

☐ vSphere Backup NFC

☐ NVMe over TCP

☐ NVMe over RDMA

CANCEL

BACK

NEXT

11. Select Use static IPv4 settings and fill in the IPv4 address and Subnet mask for the first ESXi host's vMotion IPv4 address and Subnet mask. Click NEXT.
12. Review the information and click FINISH to complete adding the vMotion VMkernel port.
13. Again, Select ADD NETWORKING
14. In the Add Networking window, ensure that VMkernel Network Adapter is selected and click NEXT.
15. Ensure that select an existing network is selected and click BROWSE.
16. Select NVMe-TCP-A and click OK.
17. Click NEXT.
18. From the MTU drop-down list, select Custom and ensure the MTU is set to 9000.
19. Under available services select NVMe Over TCP and Click NEXT.

Add Networking

1 Select connection type

2 Select target device

3 Port properties

4 IPv4 settings

5 Ready to complete

Port properties

Specify VMkernel port settings.

Network label

NVMe-TCP-A (vDSO)

MTU

Custom

9000

TCP/IP stack

Default

Available services

Enabled services

☐ vMotion

☐ Provisioning

☐ Fault Tolerance logging

☐ Management

☐ vSphere Replication

☐ vSphere Replication NFC

☐ vSAN

☐ vSphere Backup NFC

☒ NVMe over TCP

☐ NVMe over RDMA

CANCEL

BACK

NEXT
















20. Review the information and click FINISH to complete adding the NVMe over TCP VMkernel port.

21. Do the same to Add NVMe over TCP VMkernel port for NVMe-TCP-B side.

VMkernel adapters

ADD NETWORKING...

REFRESH

		Device	Network Label	Switch	IP Address	TCP/IP Stack
⋮	>>	 vmk0	 Management Network	 vSwitch0	172.22.29.101	Default
⋮	>>	 vmk1	 NFS	 vDSO	172.22.30.10	Default
⋮	>>	 vmk2	 vMotion	 vDSO	172.22.32.10	Default
⋮	>>	 vmk3	 NVMe-TCP-A	 vDSO	172.22.33.10	Default
⋮	>>	 vmk4	 NVMe-TCP-B	 vDSO	172.22.34.10	Default

Repeat the above steps for other configured ESXi hosts.

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

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:

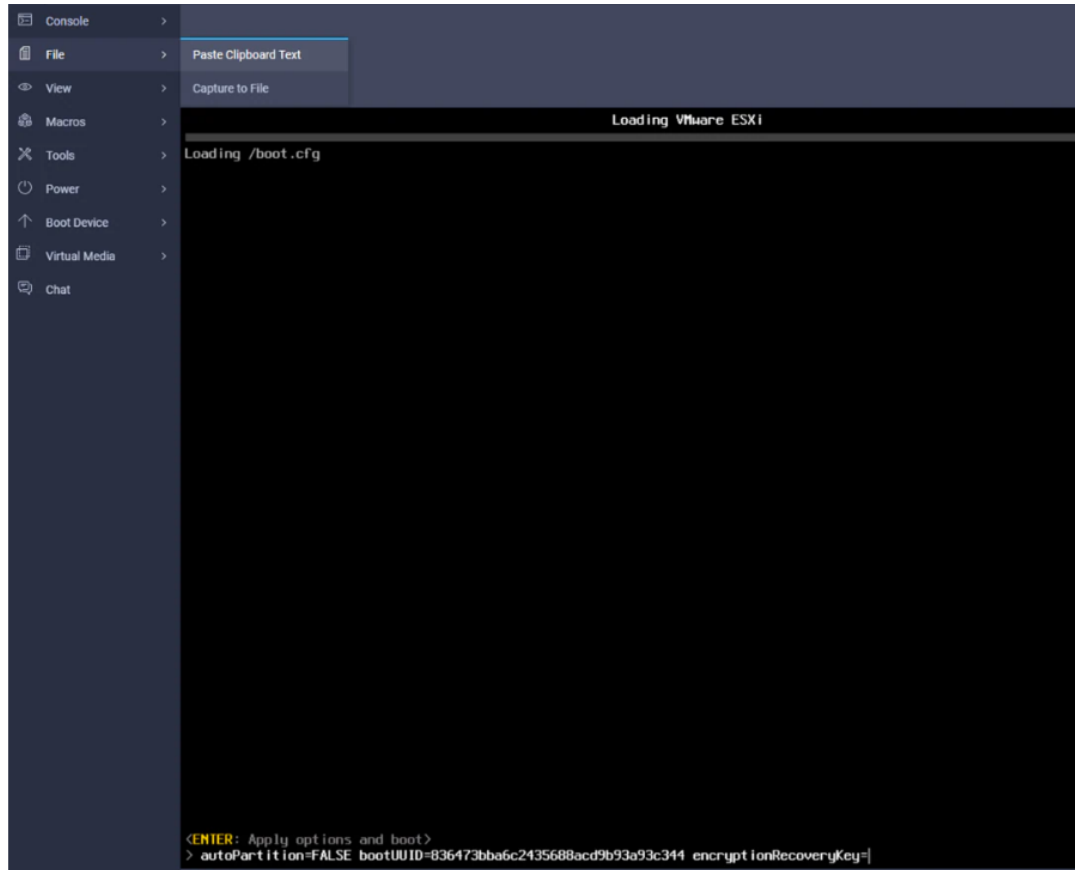
```
[root@esxi-01:~] esxcli system settings encryption recovery list
```

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, 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>

NetApp ONTAP NVMe Setup and Finalizing Storage Configuration

1. Create NVMe namespace.

```
vserver nvme namespace create -vserver <SVM_name> -path <namespace_path> -size <size_of_namespace> -ostype <OS_type>
```

```
g1434-a250::> vserver nvme namespace create -vserver Infra-SVM -path /vol/NVMe_Datastore_01/NVMe_namespace_01 -size 500G -ostype vmware
```

2. Create NVMe subsystem.

```
vserver nvme subsystem create -vserver <SVM_name> -subsystem <name_of_subsystem> -ostype <OS_type>
g1434-a250::> vserver nvme subsystem create -vserver Infra-SVM -subsystem fp-esxi-hosts -ostype vmware
```

3. Verify the subsystem was created.

```
g1434-a250::> vserver nvme subsystem show -vserver Infra-SVM

Vserver Subsystem      Target NQN
-----
Infra-SVM fp-esxi-hosts nqn.1992-08.com.netapp:sn.90e9cb71515311ed978d00a098e217cb:subsystem.fp-esxi-hosts
```

VMware vSphere NVMe Configuration

Note: Steps 1 and 2 have already been completed in the VMware ESXi Manual Configuration section of this document. Just run Step 2 and if HppManageDegradedPaths is 0, avoid the reboot and go to Step 3.

Configure NVMe-TCP on ESXi Host

1. Reboot the Host. After reboot, verify that the HppManageDegradedPaths parameter is now disabled.

```
[root@esxi-01:~] esxcfg-advcfg -g /Misc/HppManageDegradedPaths
Value of HppManageDegradedPaths is 0
```

2. Get the ESXi host NQN string and add this to the corresponding subsystem on the NetApp ONTAP array.

```
[root@esxi-01:~] esxcli nvme info get
Host NQN: nqn.2014-08.com.vmware:nvme:esxi-01
```

3. Add the host NQN(s) obtained in the last step to the **NetApp ONTAP subsystem** one by one.

```
g1434-AFF::> vserver nvme subsystem host add -vserver Infra-SVM -subsystem fp-esxi-hosts -host-nqn nqn.2014-08.com.vmware:nvme:esxi-01

g1434-AFF::> vserver nvme subsystem host add -vserver Infra-SVM -subsystem fp-esxi-hosts -host-nqn nqn.2014-08.com.vmware:nvme:esxi-02
```

Note: It is important to add the host NQNs using separate commands as shown above. NetApp ONTAP will accept a comma separated list of host NQNs without generating an error message however the ESXi hosts will not be able to map the namespace.

4. Verify the host NQNs were added successfully.

```
g1434-AFF::> vserver nvme subsystem host show -vserver Infra-SVM
Vserver Subsystem Host NQN
-----
Infra-SVM fp-esxi-hosts
          nqn.2014-08.com.vmware:nvme:esxi-01
          nqn.2014-08.com.vmware:nvme:esxi-02
2 entries were displayed.
```

5. Map the Namespace to the subsystem.

```
g1434-AFF:> vservers nvme subsystem map add -vservers Infra-SVM -subsystem fp-exsi-hosts -
path /vol/NVMe_Datastore_01/NVMe_namespace_01
```

6. Verify the Namespace is mapped to the subsystem.

```
g1434-AFF:> vservers nvme subsystem map show -vservers Infra-SVM -instance

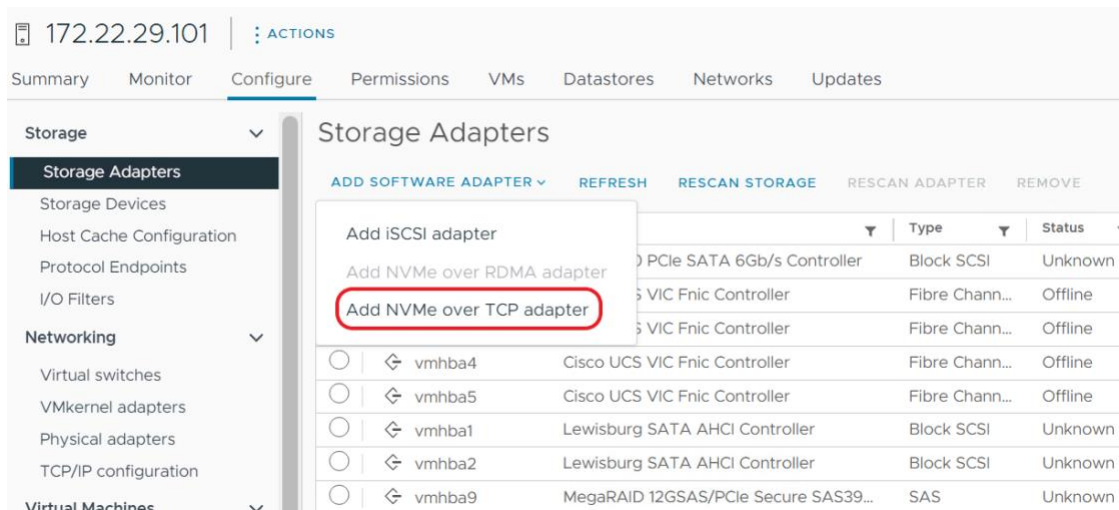
Vserver Name: Infra-SVM
  Subsystem: fp-exsi-hosts
    NSID: 00000001h
Namespace Path: /vol/NVMe_Datastore_01/NVMe_namespace_01
Namespace UUID: af6c06fb-f64a-4502-9432-8db2012e4e19
```

7. 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

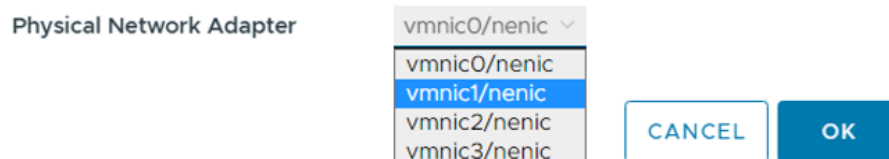
1. Select Hosts and Clusters and select the first ESXi host. In the center pane, select the Configure tab.
2. In the list under Storage, select Storage Adapters.
3. Select ADD SOFTWARE ADAPTER and click on Add NVMe over TCP adapter



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



Enable software NVMe adapter on the selected physical network adapter.



5. Again, select ADD SOFTWARE ADAPTER and click on Add NVMe over TCP adapter
6. In the Add Software NVMe over TCP adapter window, ensure that **vmnic3** Network Adapter is selected and click OK.
7. 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**.
8. 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

- 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**.
- 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.
- Repeat steps 3-10 for other ESXi host.
- 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
```

- For adding NVMe datastore, Right-click the host under **Hosts and Clusters** and select **Storage > New Datastore**. Leave VMFS selected and click **NEXT**.

14. 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 nvme_datastore

	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**

15. Leave VMFS 6 selected and click **NEXT**.

16. Leave all Partition configuration values at the default values and click **NEXT**.

17. Review the information and click **FINISH**.

18. Select **Storage** and select the new NVMe datastore. In the center pane, select **Hosts**. Ensure all the NVMe hosts have mounted the datastore.

g13vcenter.fpmc.sa

FXP-Express

host1

host2

infra_datastore

infra_swap

nvme_datastore

vCLS

nvme_datastore

ACTIONS

Summary Monitor Configure Permissions Files **Hosts** VMs

	Name	State	Status	Cluster
	172.22.29.101	Connected	✓ Normal	C220-M6
	172.22.29.102	Connected	✓ Normal	C220-M6

Finalize the NetApp ONTAP Storage Configuration

Make the following configuration changes to finalize the NetApp controller configuration.

1. Configure DNS for infrastructure SVM

To configure DNS for the Infra-SVM, run the following command

```
dns create -vsriver <vsriver-name> -domains <dns-domain> -nameserve <dns-servers>
```

```
Example:
dns create -vserver Infra-SVM -domains flexpodb4.cisco.com -nameservers 10.102.1.151,10.102.1.152
```

2. Delete the residual default broadcast domains with ifgroups (Applicable for 2-node cluster only)

To delete the residual default broadcast domains that are not in use, run the following commands:

```
broadcast-domain delete -broadcast-domain <broadcast-domain-name>

broadcast-domain delete -broadcast-domain Default-1
broadcast-domain delete -broadcast-domain Default-2
```

3. Test Auto Support

To test the Auto Support configuration by sending a message from all nodes of the cluster, run the following command:

```
autosupport invoke -node * -type all -message "FlexPod Express ONTAP storage configuration completed"
```

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).

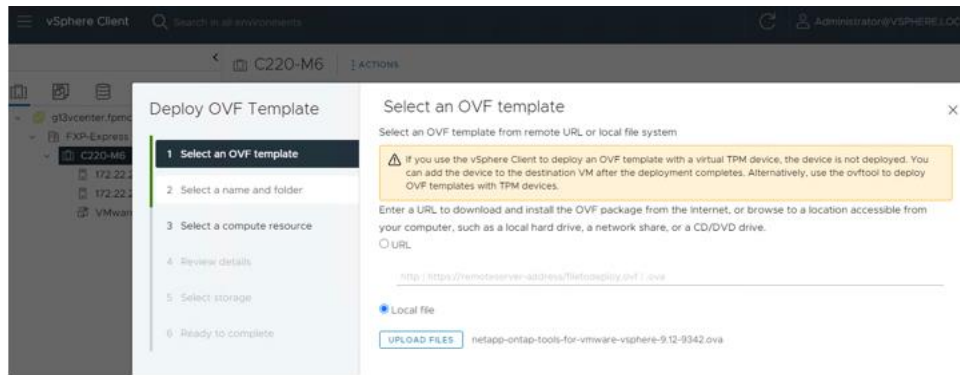
The Backup and Recovery capability has been integrated with SnapCenter and requires additional licenses for SnapCenter to perform backup and recovery of virtual machines and applications.

Note: The requirements for deploying NetApp ONTAP Tools are listed here.

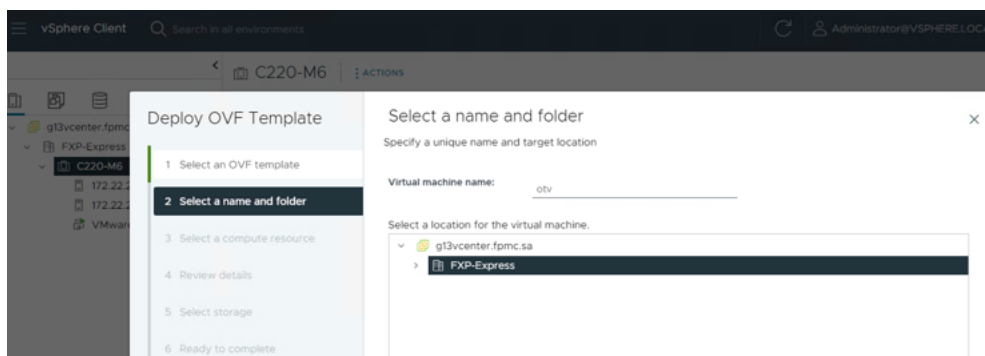
Install ONTAP Tools 9.12

To install the NetApp ONTAP tools 9.12 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

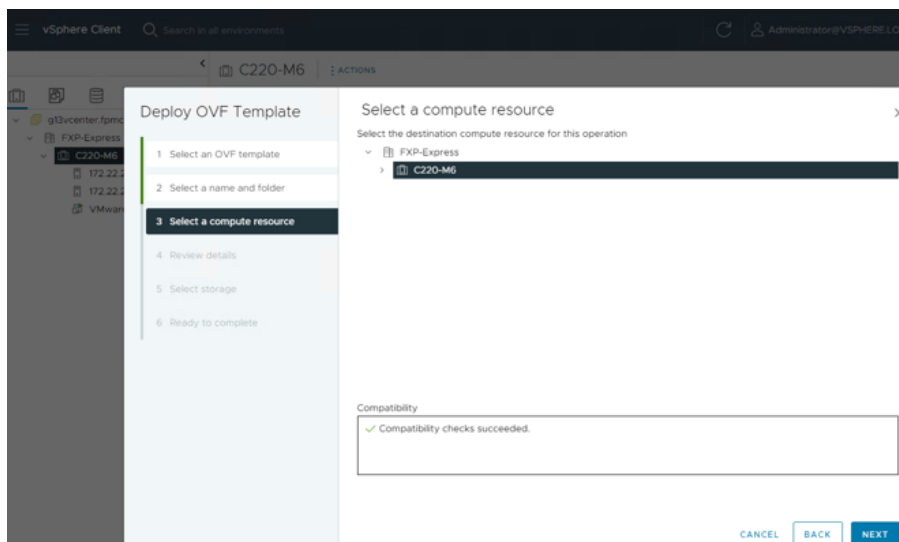
1. Go to vCenter Server > Host and Clusters > Deploy OVF Template.
2. Enter a URL for the package and click Next or browse locally to select the NetApp ONTAP tools OVA file downloaded from NetApp Support site and click Open and then click Next.



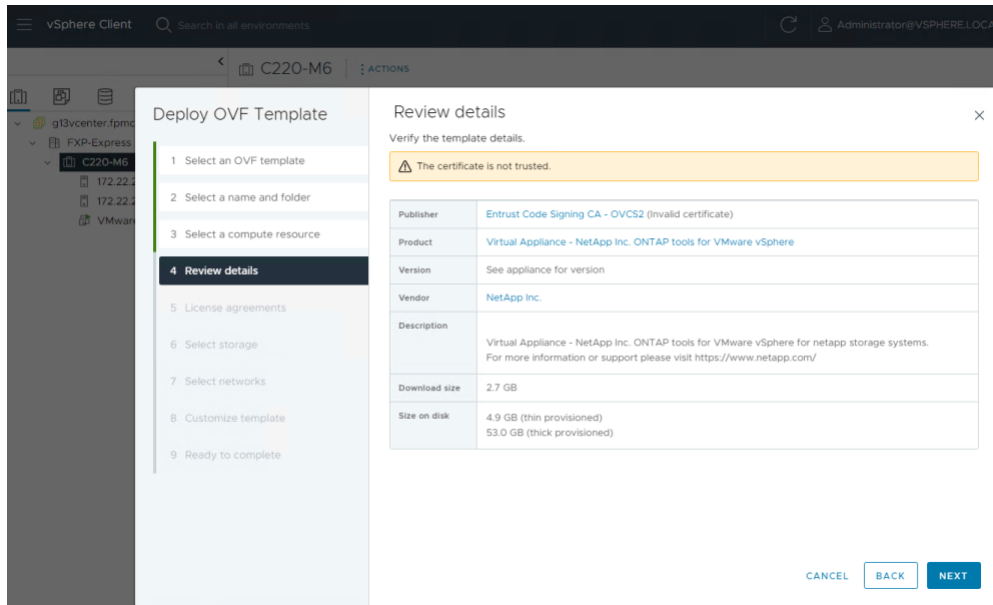
3. Enter the VM name and select the FlexPod® Express datacenter to deploy and click Next.



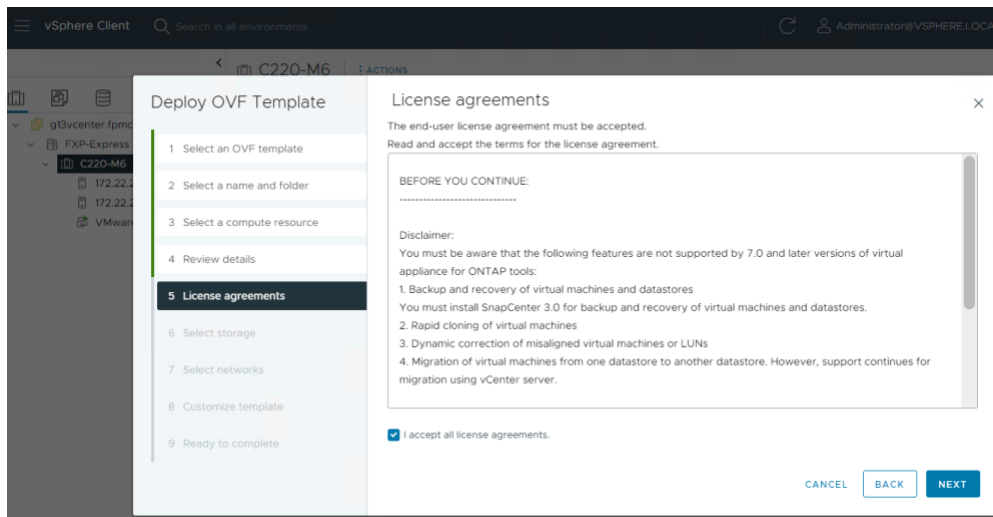
4. Select a compute resource for the deployment and click Next



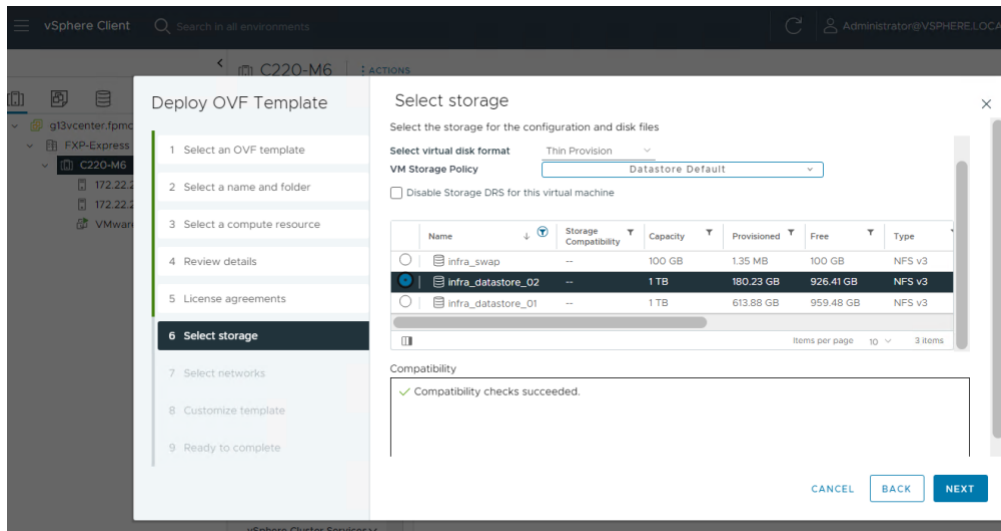
5. Review template details and click Next.



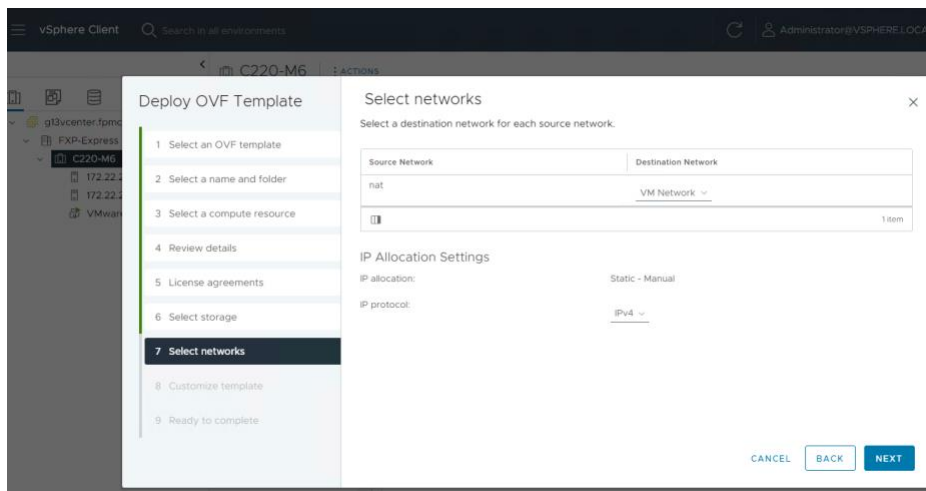
6. Accept license and click Next



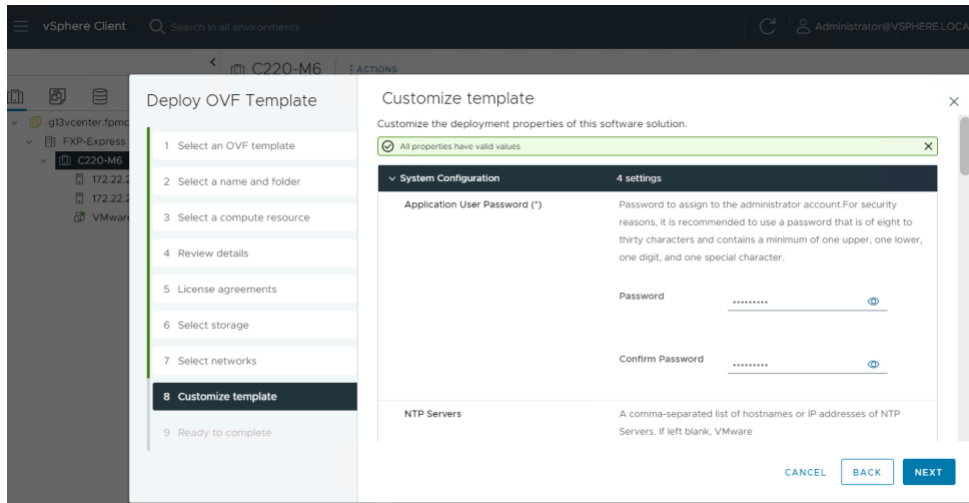
7. Select the Thin Provision virtual disk format and one of the NFS datastores. Click Next.



8. Choose a destination network, configure IP allocation setting, and click Next.



9. From Customize Template, Enter the NetApp ONTAP tools administrator password, NTP server, NetApp ONTAP tools maintenance user password, Derby Database password, vCenter server information, and network configuration details and Click Next.



10. Review the configuration details entered and click Finish to complete the deployment of NetApp ONTAP tools VM.

```
ONTAP tools for VMware vSphere

System IP addresses:
  IPv4 address: 172.22.31.50

Log in to the Appliance in a web browser using

  https://172.22.31.50:9083/
  https://otv-01:9083/

Support bundles are found under the /support directory at

  sftp://172.22.31.50

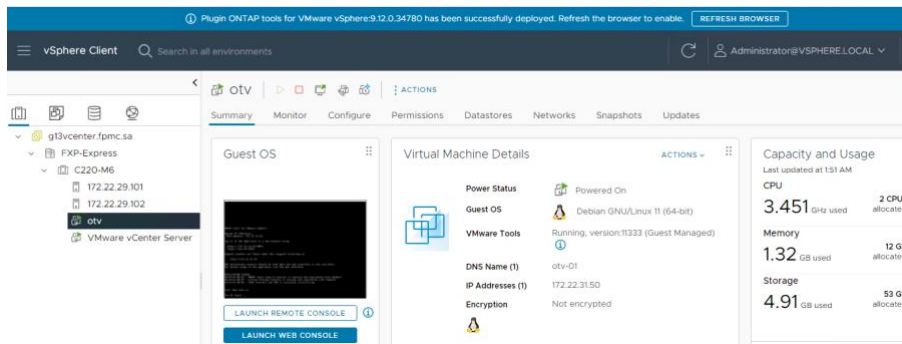
The maintenance console should be used when the web interface is not available.
For normal usage of the Appliance, use the web interface.

APPLICATION STATUS:
04/25/23 08:49 : ONTAP tools plug-in service is running and registered with vSphere
04/25/23 08:49 : Virtual Storage Console is running and registered with vSphere
04/25/23 08:48 : VASA Provider and SRA is currently initializing

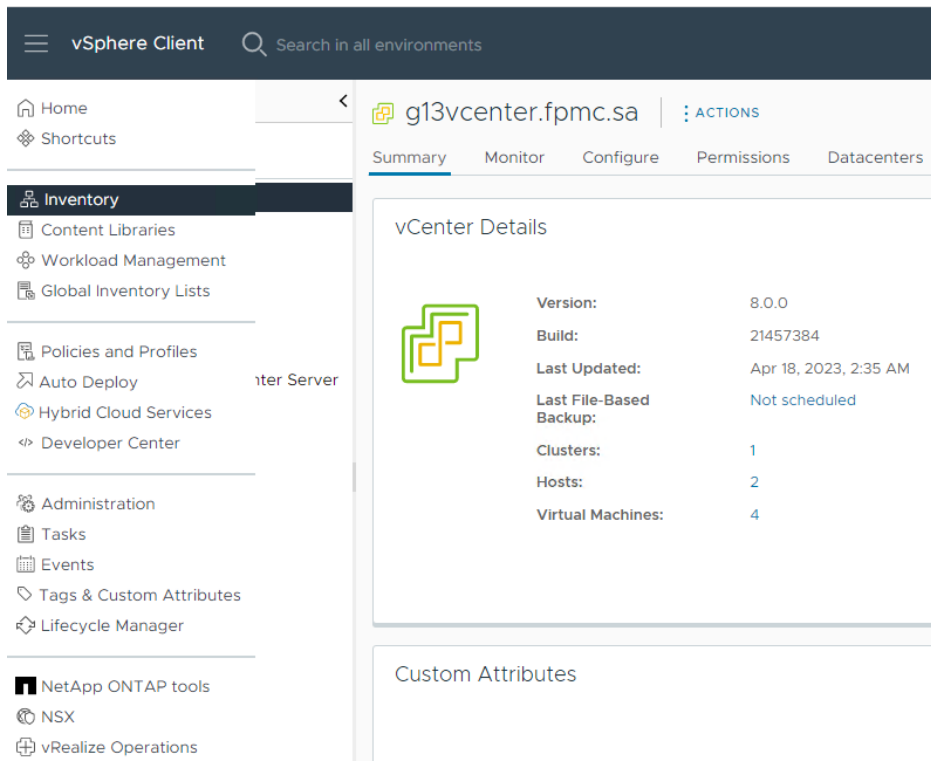
Hint: Num Lock on

otv-01 login: _
```

11. Power on the NetApp ONTAP tools and open the VM console to confirm NetApp ONTAP tools started up properly.
12. On the vCenter GUI, it will indicate that NetApp ONTAP tools had been installed and the page should be refreshed to enable. Click on Refresh Browser to enable NetApp ONTAP Tools.



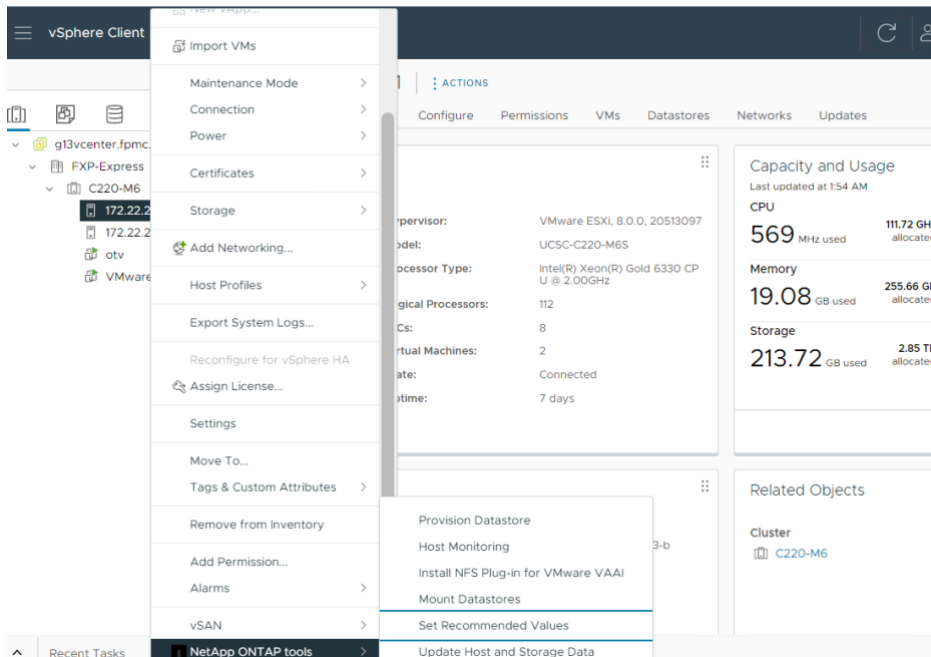
13. From the Home menu, confirm that the NetApp ONTAP Tools is listed, and it shows up in the Installed Plugins list.



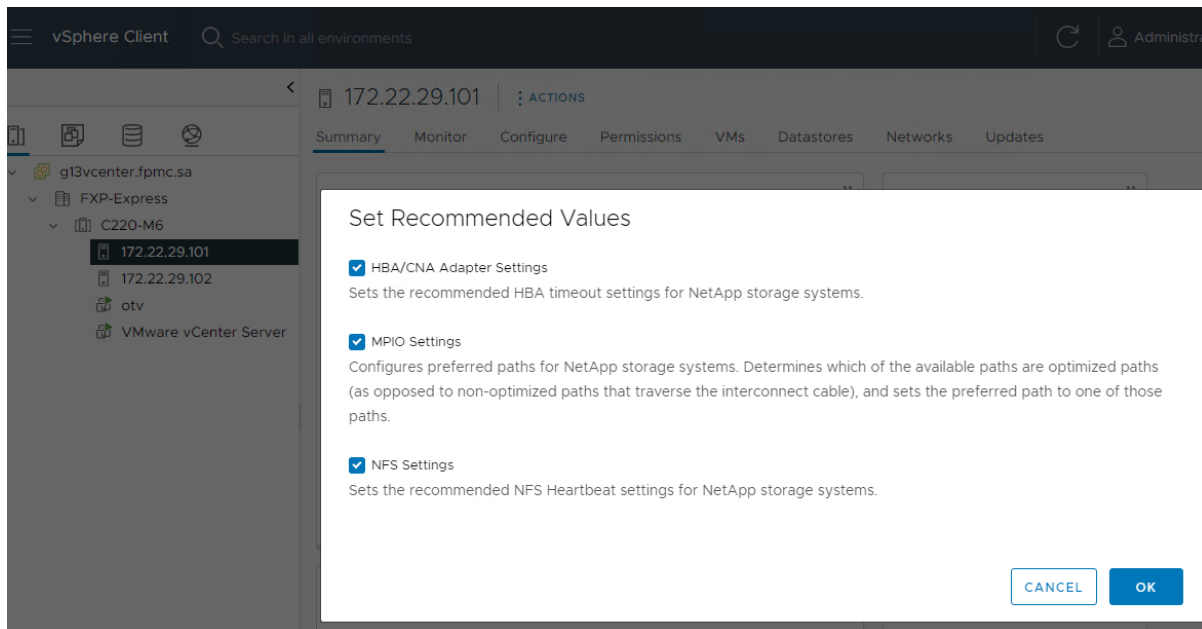
Optimal storage settings for ESXi hosts

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.



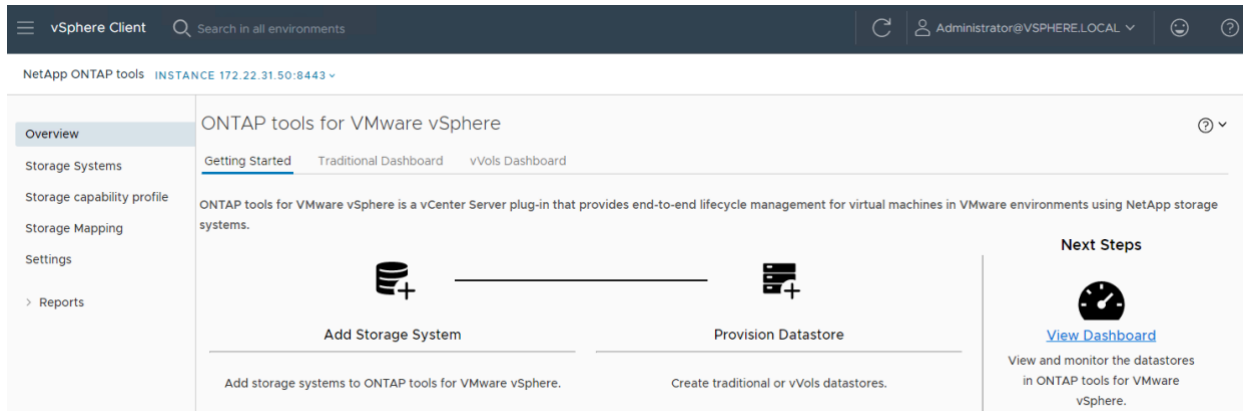
Discover and add storage resources

Note: From ONTAP tools 9.12 release onwards all ONTAP storage systems communication happens through certificate-based authentication.

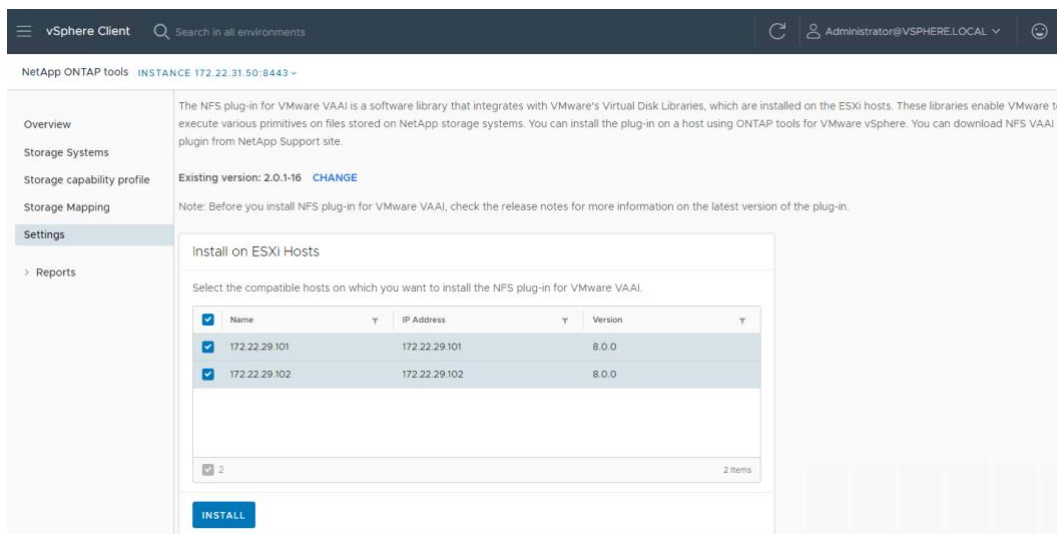
To add storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, follow these steps:

1. Log in to the vCenter Server.

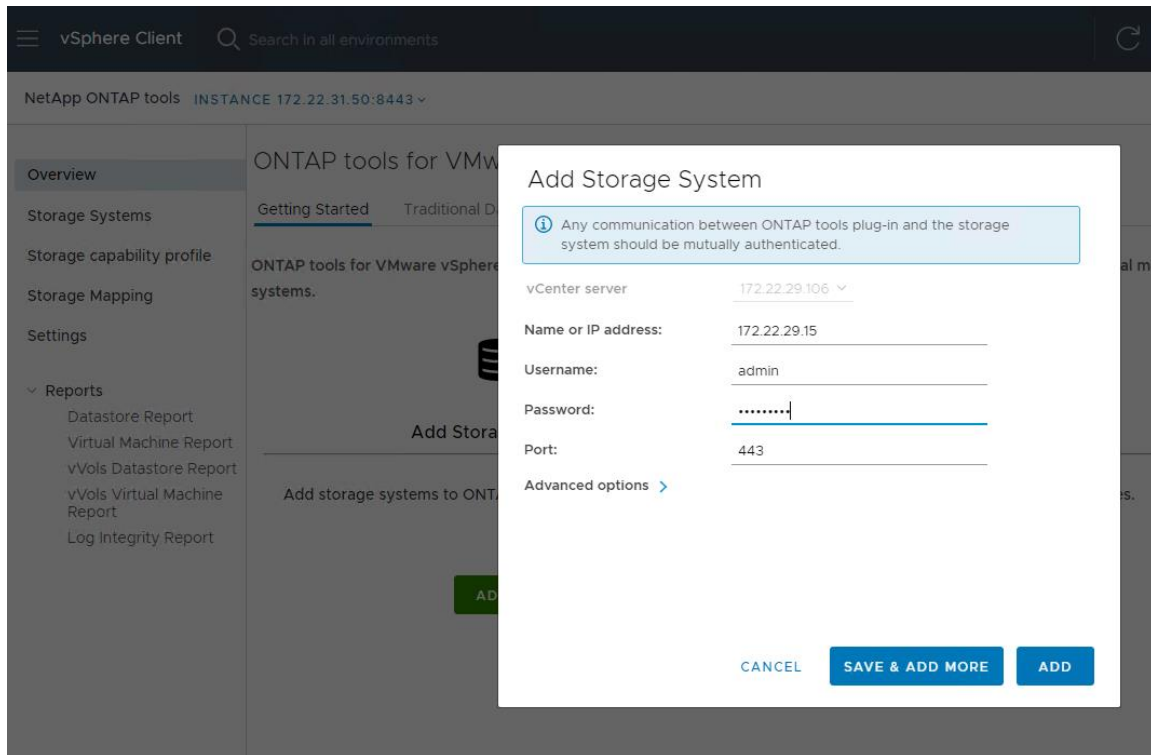
2. In the Home screen, click the Home tab and click NetApp ONTAP Tools.
3. Go to Storage Systems > Add.
4. Go to Overview > Getting Started, and then click Add under Add Storage System.

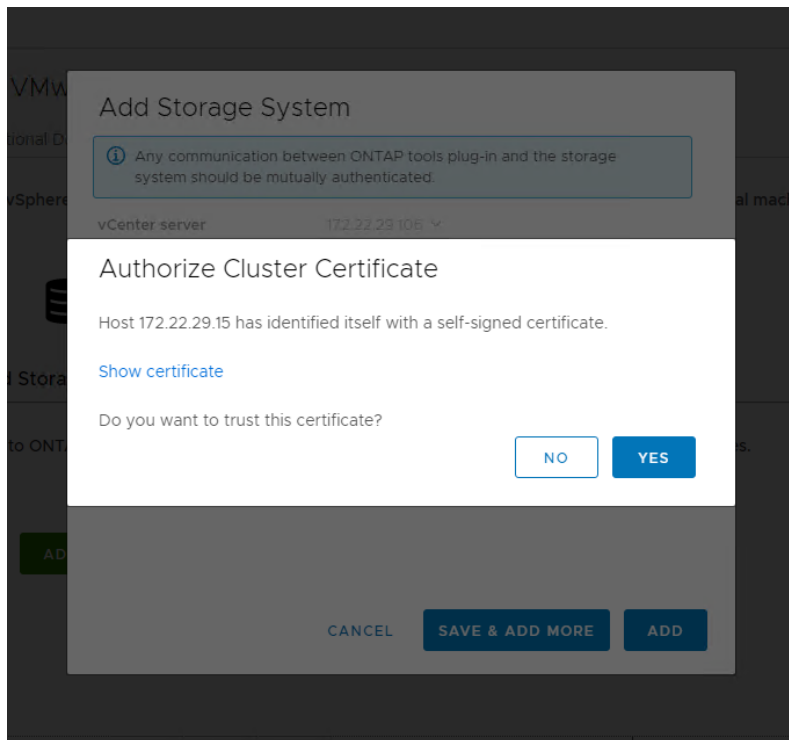


5. Specify the vCenter server instance where the storage will be located.

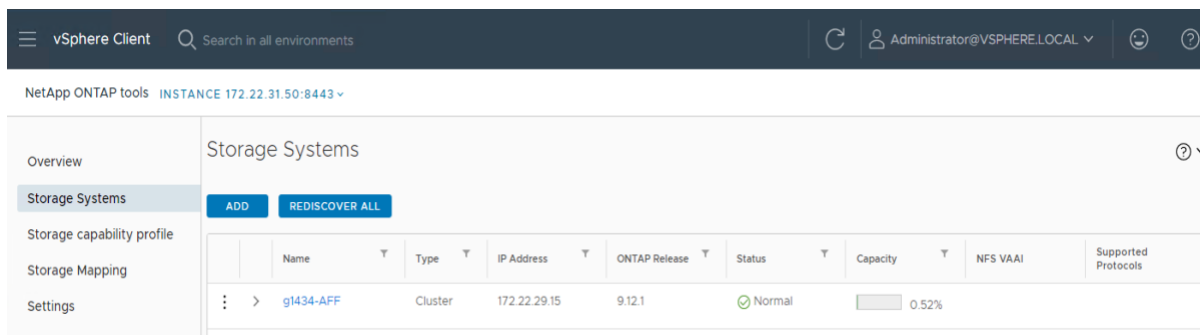


6. In the Name or IP Address field, enter the storage cluster management
7. Enter admin for the username and the admin password for password.
8. Confirm using port 443 to connect to this storage system.
9. Click Add to add the storage system

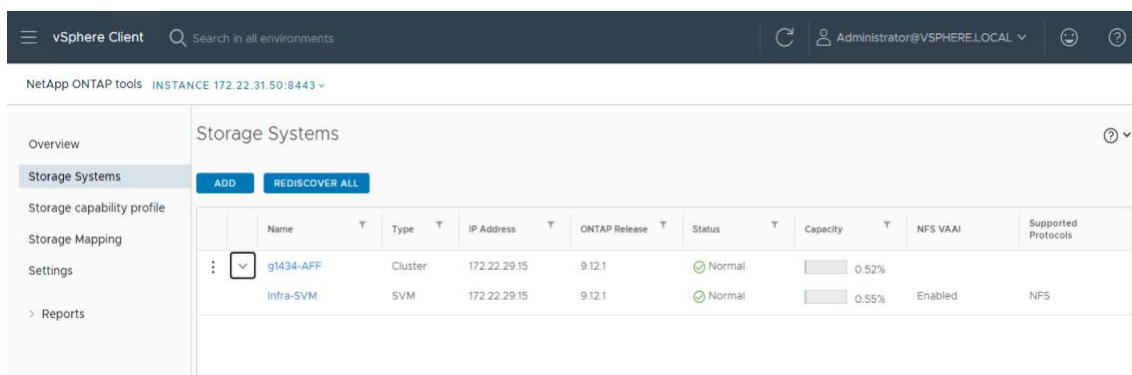




10. Select Storage System on the left pane to verify the storage system had been properly

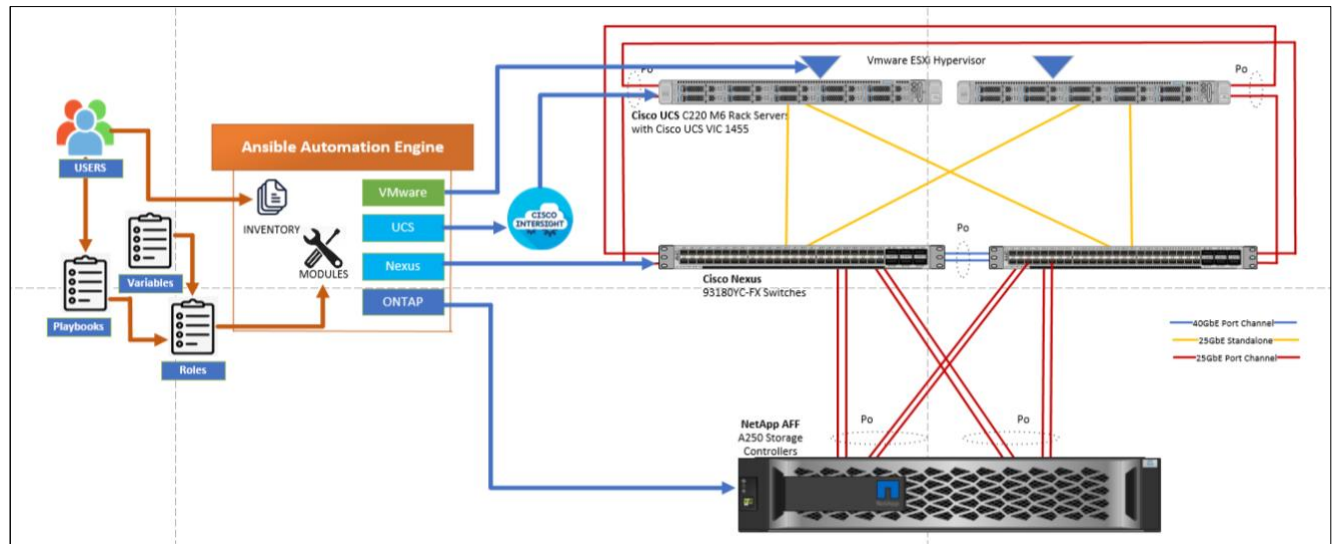


11. Expand the arrow next to the cluster name to see the SVM level information.



Note: It is best practice to use ONTAP Tools to provision new datastores after it is installed and configured.

This section specifies the Ansible automation support for the complete end-to-end deployment of this solution. The Red Hat Ansible integration with the FlexPod solution automates the deployment of complete FlexPod infrastructure enabling customers to take advantage of programming and automating the infrastructure at scale with agility, extending the benefits of automation to the entire stack.



Note: A guide for installing and getting started with Ansible can be found at: https://docs.ansible.com/ansible_community.html

To download the Ansible playbooks for configuring the infrastructure, the management workstation needs a working installation of Git as well as access to the public GitHub repository. Customers can also manually download the repository and copy the files to the management workstation. The Ansible playbooks for this solution can be found at the following link:

Note: Before executing the Ansible playbooks to set up various devices, several variables must be updated based on the customer-specific implementation. These variables contain values such as the interfaces, interface numbering, pools, policies and ports on Cisco UCS, IP addresses and interfaces for storage, etc.

To copy the GitHub repository for the project, clone the collection to a new (empty) folder on the management workstation. Cloning the repository creates a local copy, which is then used to modify and run the playbooks. To clone the GitHub repository, follow these steps:

1. From the management workstation, create a new folder for the project. The GitHub collection will be cloned to this new folder.
2. Open a command-line or console interface on the management workstation and change directories to the newly created folder.

3. Clone the GitHub collection using the following command:

```
cd <new folder>
git clone https://github.com/ucs-compute-solutions/FlexPod-Express-Intersight.git
```

Run the Ansible Playbooks

To run the Ansible playbook after updating the necessary lab environment variables specific to the deployment, execute the below command:

```
ansible-playbook -i inventory <playbook yaml file>
```

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 deployment of the converged infrastructure solution. 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	July 2023	Initial release.
Version 2.0	September 2023	Update for Automation

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.