White Paper
Public

# FlexPod with End-to-End 100G, Cisco Intersight Managed Mode, VMware 7U3, and NetApp ONTAP 9.12

## White Paper to deploy FlexPod using Intersight Cloud Orchestrator

Published: September 2023

In partnership with:

# Contents

## Executive Summary

The FlexPod solution is a validated approach for deploying Cisco® and NetApp technologies and products to build shared private and public cloud infrastructure. Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data-center platforms. The success of the FlexPod solution is achieved through its ability to evolve and incorporate both technology and product innovations in the areas of management, compute, storage, and networking. This document covers deployment details of incorporating the new Cisco Unified Computing System™ (UCS®) 5th-generation components into the FlexPod and the ability to manage FlexPod components from the cloud using Cisco Intersight® technology. Some of the main advantages of integrating Cisco UCS 5th-generation components into the FlexPod infrastructure follow:

- A simpler programmable infrastructure: The infrastructure is delivered through a single partner open application programming interface (API) that you can integrate with third-party capabilities.
- End-to-End 100-Gbps Ethernet: The 5th-generation Cisco UCS VIC 15231 Virtual Interface Card (Cisco UCS VIC 15231), the 5th-generation Cisco UCS 6536 Fabric Interconnect (Cisco UCS 6536 FI), and the Cisco UCSX-I-9108-100G Intelligent Fabric Module (Cisco UCSX-I-9108-100G IFM) deliver 100G Ethernet from the server through the network to the storage.
- End-to-End 32-Gbps Fibre Channel: The 5th-generation Cisco UCS VIC 15231, the 5th-generation Cisco UCS 6536 FI, and the Cisco UCSX-I-9108-100G IFM deliver 32G Ethernet from the server (through 100G Fibre Channel over Ethernet (FCoE) through the network to the storage.
- Innovative cloud operations: Feature delivery is continuous, and you don't need to maintain on-premises virtual machines to support management functions.
- Built for investment protections: The solution is design-ready for future technologies such as liquid cooling and high-wattage CPUs; it also is Compute Express Link (CXL)-ready.

In addition to the compute-specific hardware and software innovations, the integration of the Cisco Intersight cloud platform with VMware vCenter and NetApp Active IQ Unified Manager delivers monitoring, orchestration, and workload optimization capabilities for different layers (virtualization and storage) of the FlexPod infrastructure. The modular nature of the Cisco Intersight platform also provides an easy upgrade path to additional services, such as workload optimization.

Customers interested in understanding the FlexPod design and deployment details, including the configuration of various elements of design and associated best practices, should refer to Cisco Validated Designs for FlexPod here: https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html.

## Solution Overview

This chapter contains the following:

## Introduction

The Cisco UCS X-Series with Intersight Managed Mode (IMM) is a modular compute system, configured and managed from the cloud. It is designed to meet the needs of modern applications and to improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco Intersight platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.

Powered by the Cisco Intersight cloud-operations platform, the Cisco UCS with the X-Series, enables the next-generation cloud-operated FlexPod infrastructure that not only simplifies data-center management, but also allows the infrastructure to adapt to the unpredictable needs of modern applications as well as traditional workloads. With the Cisco Intersight platform, you get all the benefits of SaaS delivery and full lifecycle management of Cisco Intersight connected servers and integrated NetApp storage systems across data centers, remote sites, branch offices, and edge environments.

## Audience

The intended audience of this document includes IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

## Purpose of This Document

This document provides a guided deployment integrating the Cisco Intersight managed UCS X-Series platform, using end-to-end 100-Gbps Ethernet within the FlexPod design. The document explains both configurations and best practices for a successful deployment. This deployment guide also highlights the integration of VMware vCenter and NetApp Active IQ Unified Manager to the Cisco Intersight platform. The solution enables the delivery of a true cloud-based integrated approach to infrastructure management.

## What's new in this release?

The following design elements distinguish this version of FlexPod from previous models:

- End-to-End 100G Ethernet and 32G Fibre Channel with these new Cisco UCS Components:
  - 5th-generation Cisco UCS 6536 FI integrated into the FlexPod design
  - 5th-generation Cisco UCS 15000 Series Virtual Interface Cards (VICs)
  - integrated into the FlexPod design
  - Cisco UCSX-I-9108-100G IFM for the X-Series 9508 Chassis
  - Cisco UCS C225 and C245 M6 Servers with AMD EPYC CPUs.
- Now with Non-Volatile Memory Express over Transmission Control Protocol (NVMe-TCP) storage protocol with NetApp ONTAP 9.12.1
- An integrated, more complete end-to-end Automated Day 0 configuration of the FlexPod Infrastructure
- VMware vSphere 7.0 Update 3

# Deployment Hardware and Software

This chapter contains the following:



## Design Requirements

The FlexPod Converged Solution with Cisco UCS and Intersight meets the following general design requirements:

- Resilient design across all layers of the infrastructure with no single point of failure
- Scalable design with the flexibility to add compute capacity, storage, or network bandwidth as needed
- Modular design that you can be replicated to expand and grow as the needs of your business grow
- Flexible design that can easily support customization of models of included components
- Simplified design with the ability to integrate and automate with external automation tools
- Cloud-enabled design, which you can configure, manage, and orchestrate using a graphical user interface (GUI) or APIs

The following sections describe how to connect and configure various solution components so you can deliver a solution that meets all these design requirements.

## Physical Topology

The FlexPod solution with end-to-end 100G Ethernet is built using the following hardware components:

- Cisco UCS X9508 Chassis with Cisco UCSX-I-9108-100G IFMs and up to 8 Cisco UCS X210c M6/M7 Compute Nodes with 3rd-/4th-generation Intel Xeon Scalable CPUs
- 5th-generation Cisco UCS 6536 FIs to support 25G, 100G, and 32G Fibre Channel connectivity
- Cisco UCS C225 M6 and C245 M6 Rack Mount Servers with AMD EPYC CPUs
- A high-speed Cisco Nexus® 9300 Cloud Scale switching design with 100G connectivity; note that at least (6) 100G ports per switch are required for NVMe-TCP based deployments or a minimum of (4) 100G for NVMe over Fibre Channel are required
- NetApp AFF A800/A400 with end-to-end NVMe storage over 100G or through a 32G Fibre Channel network
- Cisco MDS 32G/64G storage-area network (SAN) switches to support a Fibre Channel storage configuration for Fibre Channel Protocol (FCP) based designs; Cisco MDS and Fibre Channel connectivity are not needed when implementing an IP-based connectivity design with Internet Small Computer System Interface (iSCSI) boot, Network File System (NFS), and NVMe-TCP; alternatively, the Cisco UCS 6536 can operate in switch mode acting as a Fibre Channel switch for FCP-based designs

### Software components

- Cisco Intersight SaaS platform to deploy, maintain, and support UCS and FlexPod components
- Cisco Intersight Assist Virtual Appliance to connect NetApp ONTAP, VMware vCenter, and Cisco Nexus and MDS switches with Cisco Intersight
- NetApp Active IQ Unified Manager to monitor and manage the NetApp ONTAP integration with Cisco Intersight
- VMware vCenter to set up and manage the virtual infrastructure as well as Cisco Intersight integration

## FlexPod IP-based Storage Design

Figure 1 shows various hardware components and the network connections for the IP-based FlexPod design.

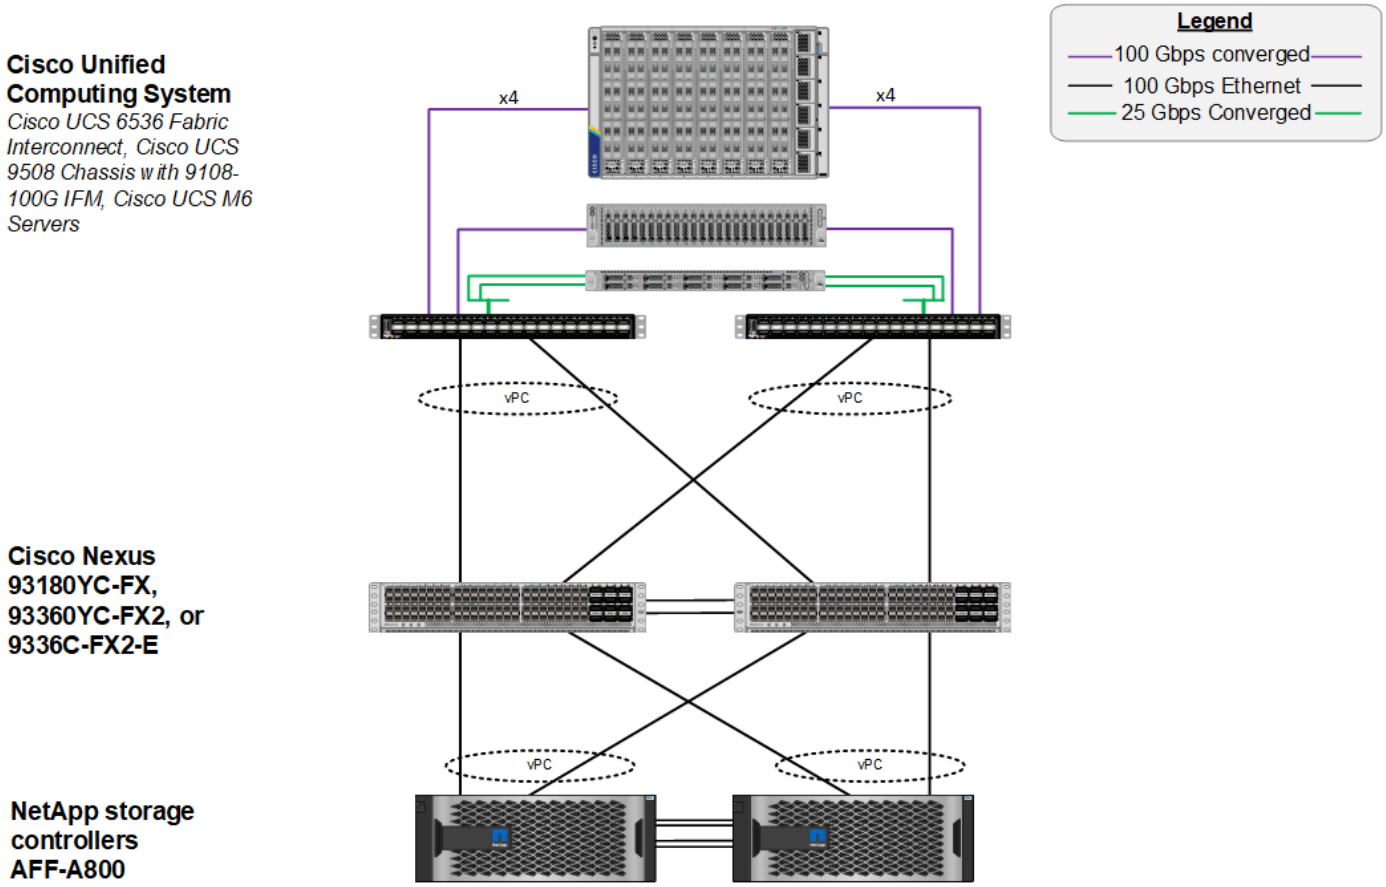


**Figure 1.**
FlexPod Physical Topology for IP-based storage access

The reference hardware configuration in Figure 1 includes:

- Two Cisco Nexus 93360YC-FX2 Cloud Scale switches running in NX-OS mode
- Two Cisco UCS 6536 FI to provide the chassis connectivity; One 100G port from each FI to each switch, configured as Virtual Port Channels

- One Cisco UCS X9508 Chassis that connects to the FIs using the Cisco UCSX 9108-100G Intelligent Fabric Modules (IFMs), where four 100G ports are used on each IFM to the corresponding FI; if additional bandwidth is required, all eight 100G ports can be used
- One NetApp AFF A800/A400 HA pair that connects to the Cisco Nexus 9300 Cloud Scale switches using two 100G ports from each controller configured as a Port Channel
- Two (one shown) Cisco UCS C245 Rack Mount Servers that connect to the FIs using two 100G ports per server
- Two (one shown) Cisco UCS C225 Rack Mount Servers that connect to the FIs with breakout using four 25G ports per server

## FlexPod FC-based Storage Design

Figure 2 shows various hardware components and the network connections for the FC-based FlexPod design.

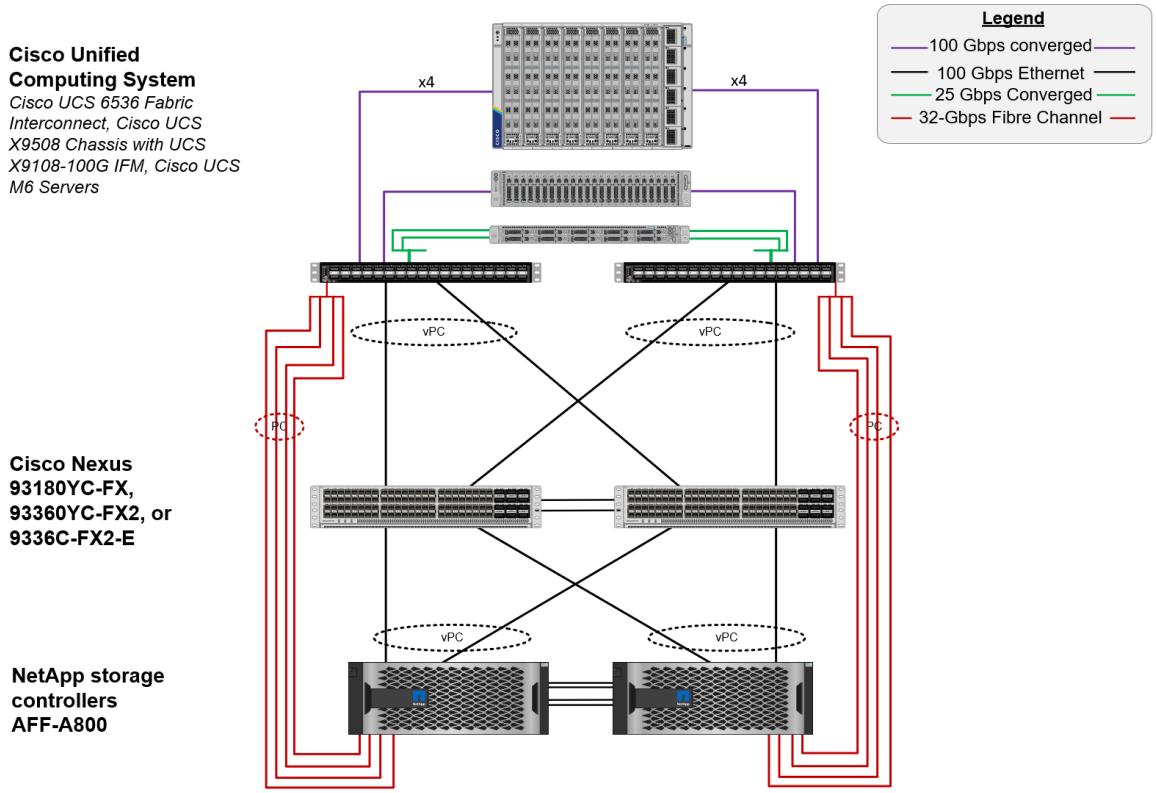


**Figure 2.**
FlexPod Physical Topology for FC-based storage access

The reference hardware configuration includes:

- Two Cisco Nexus 93360YC-FX2 Cloud Scale switches running in NX-OS mode
- Two Cisco UCS 6536 FIs to provide the chassis connectivity; one 100G port from each FI, configured as a Virtual Port Channel, connected to each Cisco Nexus 9300; four FC ports connected to the Cisco MDS 9132T switches with breakout using 32-Gbps optics configured as a single Port Channel for SAN connectivity

- One Cisco UCS X9508 Chassis that connects to FIs using Cisco UCSX 9108-100G IFMs, where four 100G ports are used on each IFM to connect to the corresponding FI; if additional bandwidth is required, all eight 100G ports can be used
- One NetApp AFF A800/A400 HA pair that connects to the Cisco Nexus 9300 Cloud Scale switches using two 100G ports from each controller configured as a Virtual Port Channel; Two 32G FC ports from each controller connected to each Cisco MDS 9132T for SAN connectivity
- Two (one shown) Cisco UCS C245 Rack Mount servers that connect to the FIs using two 100G ports
- Two (one shown) Cisco UCS C225 Rack-Mount Servers that connect to the FIs with breakout using four 25G ports per server

### NetApp storage controller connectivity

The NetApp storage controller and disk shelves should be connected according to best practices for the storage controller and disk shelves. For guidance, refer to NetApp Support: https://docs.netapp.com/us-en/ontap-systems/index.html.

## VLAN configuration

Table 1 lists Virtual LANs (VLANs) configured for setting up the FlexPod environment along with their usage.

**Table 1.** VLAN usage examples

| VLAN ID | Name | Usage | IP subnet |
|---|---|---|---|
| 1000 | OOB-MGMT | Out-of-band management VLAN to connect management ports for various devices | 198.18.0.0/24; GW: 198.18.0.1 |
| 1001 | IB-MGMT | In-band management VLAN for in-band management connectivity; for example, VMware ESXi management and Keyboard Video and Monitor (KVM) | 198.18.1.0/24; GW: 198.18.1.1 |
| 1002 | vMotion | VMware vMotion traffic | 198.18.2.0/24 ** |
| 1011* | iSCSI-A | iSCSI-A. A path for iSCSI Storage traffic including Boot from iSCSI SAN traffic | 198.18.11.0/24 ** |
| 1012* | iSCSI-B | iSCSI-B. B path for iSCSI storage traffic including Boot from iSCSI SAN traffic | 198.18.12.0/24 ** |
| 1013 | NVMe-TCP-A | NVMe-TCP-A path when using NVMe-TCP | 198.18.13.0/24 ** |
| 1014 | NVMe-TCP-B | NVMe-TCP-B path when using NVMe-TCP | 198.18.14.0/24 ** |
| 1015 | NFS | NFS VLAN for mounting datastores in ESXi servers for virtual machines (VMs) | 198.18.15.0/24 ** |
| 3001 | VM-HOST | VM Network Traffic | |

* iSCSI VLANs are not required if using FC storage for Boot.

** IP gateway is not needed because no routing is required for these subnets but can be beneficial for general troubleshooting.

**Note:** Some of the key highlights of the VLAN usage follow:

- VLAN 1000 is for out-of-band (OOB) management interfaces.
- VLAN 1001 is used for in-band management of management VMs, VMware ESXi hosts, and other infrastructure services.
- VLAN 1002 is used for VM vMotion.
- A pair of iSCSI VLANs (1011 and 1012), A side/B Side, is configured to provide access to boot logical unit numbers (LUNs) for ESXi hosts. These VLANs are not needed if the deployment is using FC-only connectivity. They are required only if a Volume is defined as using “iscsi” as the mount “protocol”. Typically, they are required only if iSCSI boot has been defined.
- A pair of NVMe-TCP VLANs (1013 and 1014), A side/B Side, are configured to provide access to NVMe datastores. They are required only if a Volume is defined as using “nvme-tcp” as the mount “protocol”.
- VLAN 1015 provides ESXi hosts access to the NFS datastores hosted on the NetApp Controllers for deploying VMs.

## Management virtual machines

Table 2 lists the infrastructure VMs necessary for deployment as outlined in this document.

**Table 2.** Management virtual machines

| Virtual machine description | VLAN | IP address | Comments |
|---|---|---|---|
| vCenter Server | 1001 | 198.18.1.100 | Hosted on pre-existing management infrastructure; can be moved to FlexPod after deployment |
| NetApp ONTAP tools | 1001 | 198.18.1.99 | Hosted on FlexPod |
| NetApp SnapCenter for vSphere | 1001 | 198.18.1.98 | Hosted on FlexPod |
| Active IQ Unified Manager | 1001 | 198.18.1.97 | Hosted on FlexPod |
| Cisco Intersight Assist | 1001 | 198.18.1.96 | Hosted on pre-existing management infrastructure; can be moved to FlexPod after deployment |
| Cisco IMM Toolkit | 1001 | 198.18.1.95 | Hosted on pre-existing management infrastructure; can be moved to FlexPod after deployment |
| Cisco IMM Transition Tool | 1001 | 198.18.1.94 | Hosted on pre-existing management infrastructure; can be moved to FlexPod after deployment |

## Software revisions

Table 3 lists the software revisions for various components of the solution.

**Table 3.** Software revisions

| Layer | Device | Image bundle | Comments |
|---|---|---|---|
| Compute | Cisco UCS | 4.2(3d) | Cisco UCS X-Series GA release for infrastructure including FIs and Intelligent Fabric Module (IFM) |
| Network | Cisco Nexus 9300 Cloud Scale NX-OS | 10.2(4)M | |
| Storage | NetApp AFF A800/A400 | NetApp ONTAP 9.12.1P2 | |
| Software | Cisco UCS X210c | 5.1(0.230054) | Cisco UCS X-series GA release for compute nodes |
| | Cisco UCS C225/245 M6 | 5.1(0.230054) | |
| | Cisco Intersight Assist appliance | 1.0.9+ | 1.0.9-342 initially installed and then automatically upgraded; the running version will continue to change as it is SaaS-delivered |
| | VMware vCenter | 7.0 Update 3h | Build 20395099 |
| | VMware ESXi | 7.0 Update 3d | Build 19482537 used in Cisco Custom ISO |
| | VMware ESXi native Fibre channel network interface card (nfnic) driver | 5.0.0.37-1 | Support for FC-NVMe |
| | VMware ESXi native ethernet network interface card (nenic) driver | 1.0.45.0-1 | |
| | NetApp ONTAP tools for VMware vSphere | 9.11 | Formerly Virtual Storage Console (VSC) |
| | NetApp NFS plug-In for VMware Array Awareness Integration (VAAI) | 2.0(1) | |
| | NetApp SnapCenter for vSphere | 4.7 | vSphere plug-in for SnapCenter included |
| | NetApp Active IQ Unified Manager | 9.12P1 | |

## FlexPod cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. A cabling diagram was used to simplify cabling requirements.

The cabling diagram in this section contains the details for the prescribed and supported configuration of the NetApp AFF 800 running NetApp ONTAP 9.12.1P2.

**Note:** For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT). Be sure to use the cabling directions in this section as a guide.

**Note:** This document assumes that OOB management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used throughout the configuration steps.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk-shelf cabling, refer to NetApp Support.

Figure 3 details the cable connections used in the validation lab for the FlexPod topology based on the Cisco UCS 6536 FI. Four 32G uplinks via breakout connections as Port Channels from each Cisco UCS FI to the MDS switches, and a total of eight 32G links connect the FIs to the NetApp AFF controllers. Also, 100G links connect the Cisco UCS FIs to the Cisco Nexus switches and the NetApp AFF controllers to the Cisco Nexus switches. Additional 1G management connections are needed for an OOB network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS FI and Cisco Nexus switch is connected to the OOB network switch, and each AFF controller has a connection to the OOB network switch. Layer 3 network connectivity is required between the OOB and In-band (IB) management subnets. This cabling diagram includes both the FC-boot and iSCSI-boot configurations.

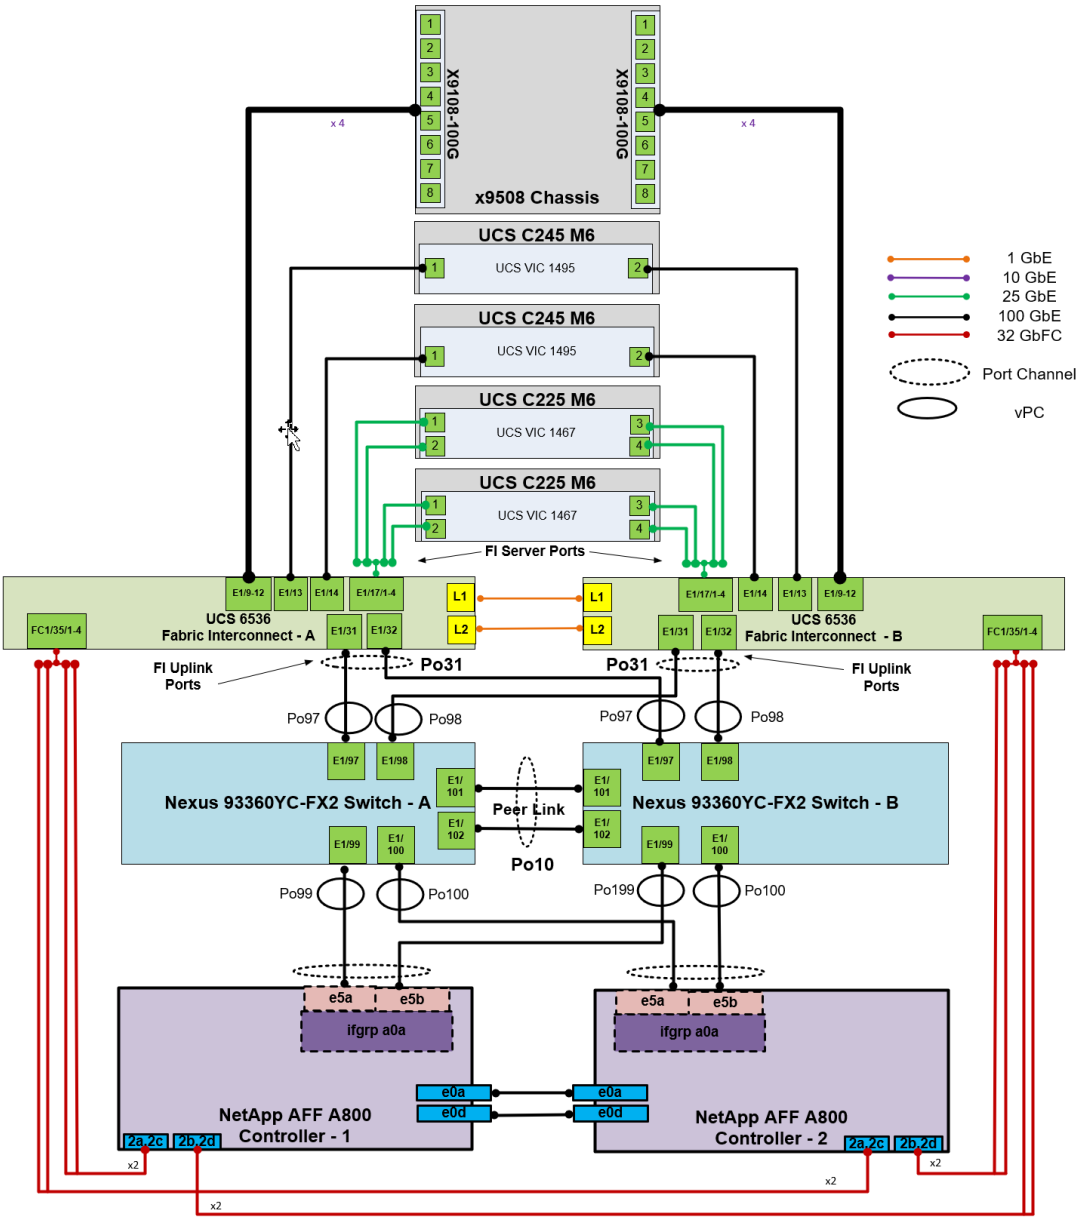


**Figure 3.**
FlexPod cabling with Cisco UCS 6536 Fabric Interconnect

# Initial configuration of network switches

This chapter contains the following:

- Physical Connectivity
- Initial Configuration
- Cisco Nexus Switch Manual Configuration

This chapter provides a detailed procedure for configuring the Cisco Nexus 9300 Cloud Scale switches for use in a FlexPod environment. This solution uses the Cisco Nexus 9300 Cloud Scale switches for LAN switching in.

**Note:** The following procedures walk you through configuration of the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 10.2(4)M.

- If you are using the Cisco Nexus 9300 Cloud Scale switches for both local area network (LAN) and storage area network (SAN) switching, please refer to section FlexPod with Cisco Nexus 9300 Cloud Scale SAN Switching Configuration in the Appendix.
- The following procedure includes the setup of Network Time Protocol (NTP) distribution on both the mgmt0 port and the in-band management VLAN. The *interface-vlan* feature and *ntp* commands are used to set it up. This procedure also assumes that the default Virtual Route Forwarding (VRF) is used to route the in-band management VLAN.
- This procedure sets up the uplink virtual port channel (vPC) with the INBAND-MGMT and OOBAND-MGMT VLANs allowed.
- This validation assumes that both switches have been reset to factory defaults by using the *write erase* command followed by the *reload* command.

## Physical connectivity

Follow the physical connectivity guidelines for FlexPod explained in section FlexPod cabling.

## Initial configuration

The following procedures describe the basic configuration of the Cisco Nexus switches for use in the FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 10.2(4)M, the Cisco suggested Cisco Nexus switch release at the time of this validation.

### Procedure 1. Initial configuration for Cisco Nexus A Switch <nexus-A-hostname> from the serial console

1. Configure the switch.

**Note:** On initial boot, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
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-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-out_of_band_mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-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]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: n
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

3. To set up the initial configuration of the Cisco Nexus B switch, repeat steps 1 and 2 with the appropriate host and IP address information.

# NetApp ONTAP storage initial configuration

This chapter contains the following:

- NetApp AFF A400/A800 Controllers
- Disk Shelves
- NetApp ONTAP 9.12.1P2

## NetApp AFF A400/A800 controllers

Refer to the section NetApp Hardware Universe for planning the physical location of the storage systems:

- Site preparation
- System connectivity requirements
- Circuit breaker, power outlet balancing, system cabinet power-cord plugs, and console pinout requirements
- AFF Series systems

### NetApp Hardware Universe

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

To confirm that the hardware and software components that you would like to use are supported with the version of NetApp ONTAP that you plan to install, follow the steps at the NetApp Support site.

**Procedure 1. Confirm hardware and software components**

1. Access the HWU application to view the System Configuration guides. Click the Platforms menu to view the compatibility between different versions of the NetApp ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

### Controllers

Follow the physical installation procedures for the controllers here: https://docs.netapp.com/us-en/ontap-systems/index.html.

## Disk shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF A400 and AFF A800 is available at the NetApp Support site.

When using serial-attached Small Computer Systems Interface (SCSI) (SAS) disk shelves with NetApp storage controllers, refer to: https://docs.netapp.com/us-en/ontap-systems/sas3/index.html for proper cabling guidelines.

When using NVMe drive shelves with NetApp storage controllers, refer to: https://docs.netapp.com/us-en/ontap-systems/ns224/index.html for installation and servicing guidelines.

## NetApp ONTAP 9.12.1P2

### Complete configuration worksheet

Before running the setup script, complete the Cluster setup worksheet in the NetApp ONTAP 9 Documentation Center. You must have access to the NetApp Support site to open the cluster setup worksheet.

### Configure NetApp ONTAP nodes

Before running the setup script, review the configuration worksheets in the Software setup section of the NetApp ONTAP 9 Documentation Center to learn about configuring NetApp ONTAP. Table 4 lists the information needed to configure two NetApp ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

**Table 4.** NetApp ONTAP Software installation prerequisites

| Cluster detail | Cluster detail value |
|---|---|
| Cluster node 01 IP address | <node01-oob-mgmt-ip> |
| Cluster node 01 netmask | <node01-oob-mgmt-mask> |
| Cluster node 01 gateway | <node01-oob-mgmt-gateway> |
| Cluster node 02 IP address | <node02-oob-mgmt-ip> |
| Cluster node 02 netmask | <node02-oob-mgmt-mask> |
| Cluster node 02 gateway | <node02-oob-mgmt-gateway> |
| ONTAP 9.12.1P2 URL (http server hosting NetApp ONTAP software) | <url-boot-software> |

### Procedure 1. Configure node01

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

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

2. Allow the system to boot up.

   ```
   autoboot
   ```

3. Press **Ctrl-C** when prompted.

**Note:** If NetApp ONTAP 9.11.1P2 is not the version of the software being booted, continue with the following steps to install new software. If NetApp ONTAP 9.12.1P2 is the version being booted, select option 8 and `y` to reboot the node, and then continue with section Set Up Node.

4. To install new software, select option 7 from the menu.
5. Enter `y` to continue the installation.
6. Select `e0M` for the network port for the download.
7. Enter `n` to skip the reboot.
8. Select option 7 from the menu: `Install new software first`

9. Enter `y` to continue the installation.
10. Enter the IP address, netmask, and default gateway for `e0M`.

```
Enter the IP address for port e0M: <node01-oob-mgmt-ip>
Enter the netmask for port e0M: <node01-oob-mgmt-mask>
Enter the IP address of the default gateway: <node01-oob-mgmt-gateway>
```

11. Enter the Uniform Resource Locator (URL) where the software can be found.

**Note:** The e0M interface should be connected to the management network, and the web server must be reachable (using ping) from node 01.

```
<url-boot-software>
```

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

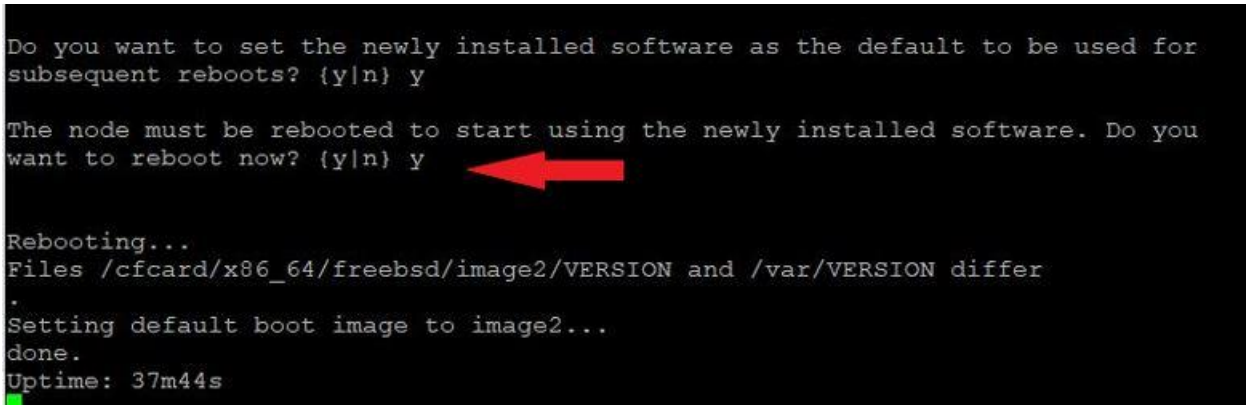


**Note:** When installing new software, the system might perform firmware upgrades to the basic input/output system (BIOS) and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

**Note:** During the NetApp ONTAP installation, a prompt to reboot the node requests a Y/N response.

15. Press **Ctrl-C** when the following message displays:

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

16. Select option 4 for Clean Configuration and Initialize All Disks.
17. Enter `y` to zero disks, reset config, and install a new file system.
18. Enter `yes` 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 02 while the initialization and creation of the root aggregate for node 01 is in progress.

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

### Procedure 2. Configure node02

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 the following message displays:

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

2. Allow the system to boot up.

   ```
   autoboot
   ```

3. Press Ctrl-C when prompted.

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

4. To install new software, select option 7.
5. Enter `y` to continue the installation.
6. Select `e0M` for the network port you want to use for the download.
7. Enter `n` to skip the reboot.
8. Select option 7: `Install new software first`
9. Enter `y` to continue the installation.
10. Enter the IP address, netmask, and default gateway for e0M.

    ```
    Enter the IP address for port e0M: <node02-oob-mgmt-ip>
    Enter the netmask for port e0M: <node02-oob-mgmt-mask>
    Enter the IP address of the default gateway: <node02-oob-mgmt-gateway>
    ```

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

**Note:** The web server must be reachable (ping) from node 02.

```
<url-boot-software>
```

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

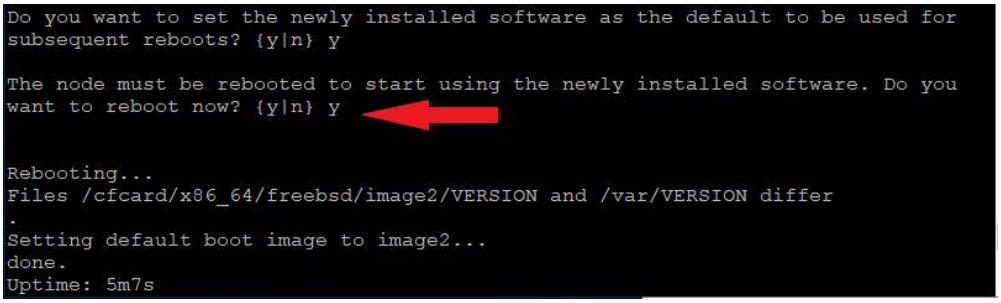


**Note:** When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-B prompt. If these actions occur, the system might deviate from this procedure.

**Note:** During the NetApp ONTAP installation, a prompt to reboot the node requests a Y/N response.

15. Press **Ctrl-C** when you see this message:

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

16. Select option 4 for Clean Configuration and Initialize All Disks.
17. Enter y to zero disks, reset config, and install a new file system.
18. Enter `yes` to erase all the data on the disks.

**Note:** Wait for the storage system reboot after initialization and creation of the root aggregate.

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

**Procedure 3. Set up node**

1. From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when NetApp ONTAP 9.12.1P2 boots on the node for the first time.
2. Follow the prompts to set up node01.

```
Welcome to the cluster setup wizard.

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

You can return to cluster setup at any time by typing “cluster setup”.
To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a
problem occur on your system.
For further information on AutoSupport, see: http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: <press Enter>
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created.

Use your web browser to complete cluster setup by accesing https://<node01-mgmt-ip>
Otherwise press Enter to complete cluster setup using the command line interface:
```

# Cisco Intersight Managed Mode domain initial configuration

This chapter contains the following:

- Cisco Intersight Managed Mode domain Set Up
- VLAN and VSAN Configuration
- Cisco UCS IMM Manual Configuration
- Cisco UCS IMM Setup Completion

The Cisco Intersight platform is a management solution delivered as a service with embedded analytics for Cisco and third-party IT infrastructures. The Cisco Intersight Managed Mode is a new architecture that manages Cisco UCS FI-attached systems through a Redfish-based standard model. Cisco Intersight Managed Mode standardizes both policy and operation management for Cisco UCS B200 M6 and Cisco UCSX X210c M6 compute nodes used in this deployment guide.

Cisco UCS C-Series M6 servers, connected and managed through Cisco UCS FIs, are also supported by Cisco Intersight Managed Mode. For a complete list of supported platforms, visit: https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide/b_intersight_managed_mode_guide_chapter_01010.html.

## Cisco Intersight Managed Mode domain setup

**Procedure 1. Set up Cisco Intersight Managed Mode on Cisco UCS FIs**

The Cisco UCS FIs need to be set up to support Cisco Intersight Managed Mode. When converting an existing pair of Cisco UCS FIs from Cisco UCS Manager mode to Cisco Intersight Managed Mode, first erase the configuration and reboot your system.

**Note:** Converting FIs to Cisco Intersight Managed Mode is a disruptive process, and configuration information will be lost. We encourage you to make a backup of your existing configuration. If a software version that supports Intersight Managed Mode (4.1(3) or later) is already installed on Cisco UCS FIs, do not upgrade the software to a recommended recent release using Cisco UCS Manager. The software upgrade should be performed using the Cisco Intersight platform to make sure Cisco UCS X-series firmware is part of the software upgrade.

1. Configure FI-A. On the Basic System Configuration Dialog screen, set the management mode to Intersight. All the remaining settings are like those for the Cisco UCS Manager managed mode (UCSM-Managed).

```
Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlexPod environment in ucsm managed mode, follow
these steps:

Connect to the console port on the first Cisco UCS fabric interconnect.

 Enter the configuration method. (console/gui) ? console


 Enter the management mode. (ucsm/intersight)? intersight


 The Fabric interconnect will be configured in the intersight managed mode. Choose (y/n) to
proceed: y
```

```
  Enforce strong password? (y/n) [y]: Enter

  Enter the password for "admin": <password>
  Confirm the password for "admin": <password>

  Enter the switch fabric (A/B) []: A

  Enter the system name: <ucs-cluster-name>

  Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>

  Physical Switch Mgmt0 IPv4 netmask : <ucs-mgmt-mask>

  IPv4 address of the default gateway : <ucs-mgmt-gateway>

  DNS IP address : <dns-server-1-ip>

  Configure the default domain name? (yes/no) [n]: y

  Default domain name : <ad-dns-domain-name>

Following configurations will be applied:

  Management Mode=intersight
  Switch Fabric=A
  System Name=<ucs-cluster-name>
  Enforced Strong Password=yes
  Physical Switch Mgmt0 IP Address=<ucsa-mgmt-ip>
  Physical Switch Mgmt0 IP Netmask=<ucs-mgmt-mask>
  Default Gateway=<ucs-mgmt-gateway>
  DNS Server=<dns-server-1-ip>
  Domain Name=<ad-dns-domain-name>

  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

2. After applying the settings, make sure you can ping the FI management IP address. When FI-A is correctly set up and is available, FI-B will automatically discover FI-A during its setup process, as shown in the next step.
3. Configure FI-B. For the configuration method, select console. FI-B will detect the presence of FI-A and will prompt you to enter the admin password for FI-A. Provide the management IP address for FI-B and apply the configuration.

```
Cisco UCS Fabric Interconnect B
Enter the configuration method. (console/gui) ? console

 Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect
will be added to the cluster. Continue (y/n) ? y

 Enter the admin password of the peer Fabric interconnect: <password>
 Connecting to peer Fabric interconnect... done
 Retrieving config from peer Fabric interconnect... done
 Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
 Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucs-mgmt-mask>

 Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4
Address

 Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

 Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

**Procedure 2. Set up Cisco Intersight account – if not already created. Skip to Procedure 3 if already created.**

1. Go to https://intersight.com and click **Create an account**.
2. Read and accept the license agreement. Click **Next**.
3. Provide an Account Name and click **Create**.
4. On successful creation of the Cisco Intersight account, the following page will be displayed:

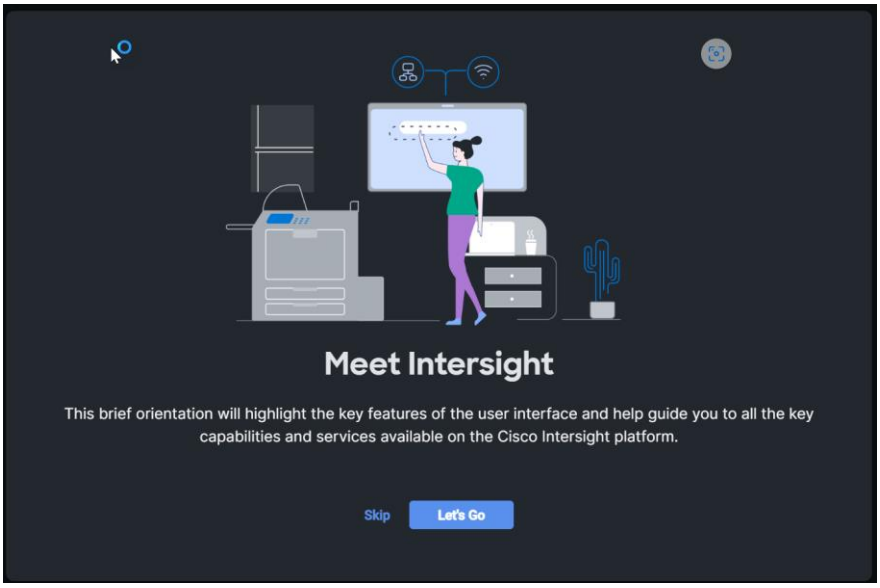


**Note:** You can also choose to add the Cisco UCS FIs to an existing Cisco Intersight account.

**Procedure 3. Set up Cisco Intersight licensing. Skip to procedure 4 if already complete.**

**Note:** When setting up a new Cisco Intersight account (as explained in this document), you need to enable the account for Cisco Smart software licensing.

1. Log into the Cisco Smart licensing portal: Smart Software Licensing
2. Verify that the correct virtual account is selected.
3. Under Inventory > General, generate a new token for product registration.
4. Copy this newly created token.

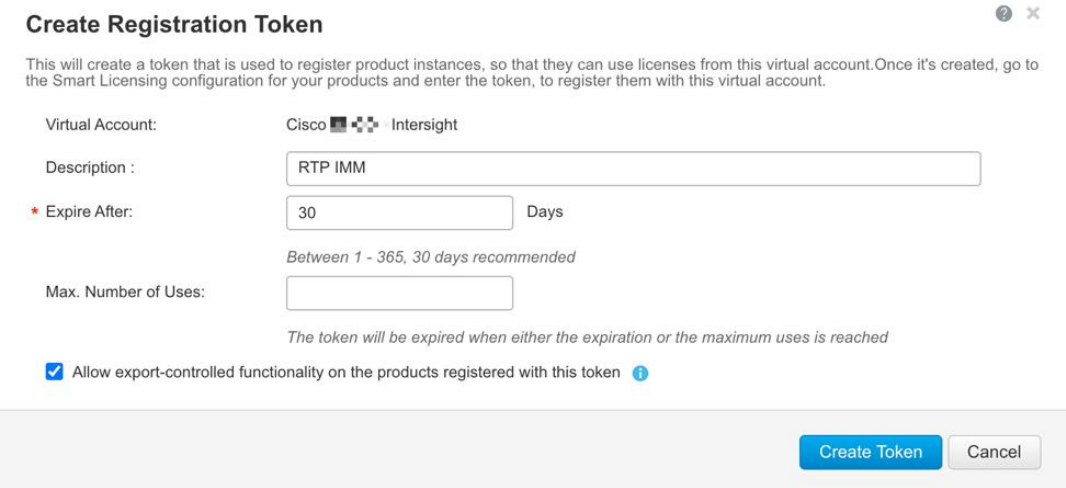


5. In the Cisco Intersight platform, click **Select Service > System**, and then click **Administration > Licensing**.
6. Under **Actions**, click **Register**.

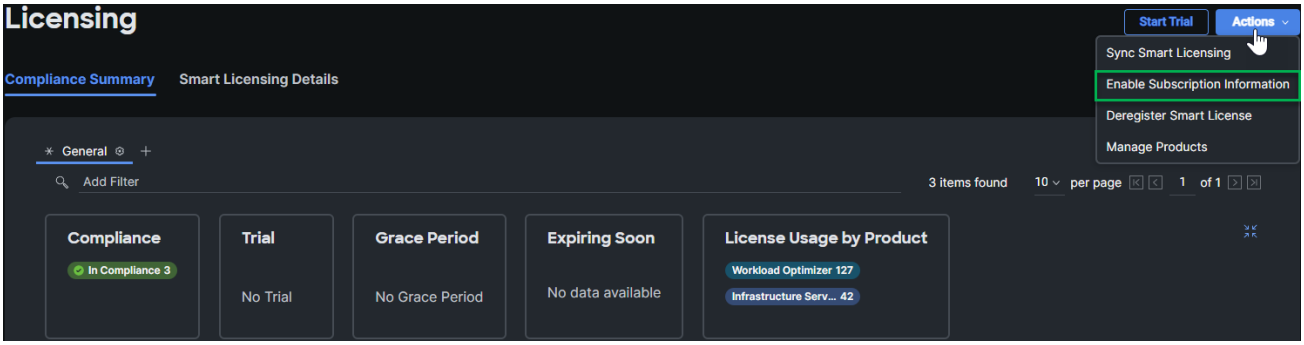


7. Enter the copied token from the Cisco Smart Licensing portal. Click **Next**.
8. Drop down the pre-selected Default Tier * and select the license type (for example, Premier).
9. Select Move All Servers to Default Tier.

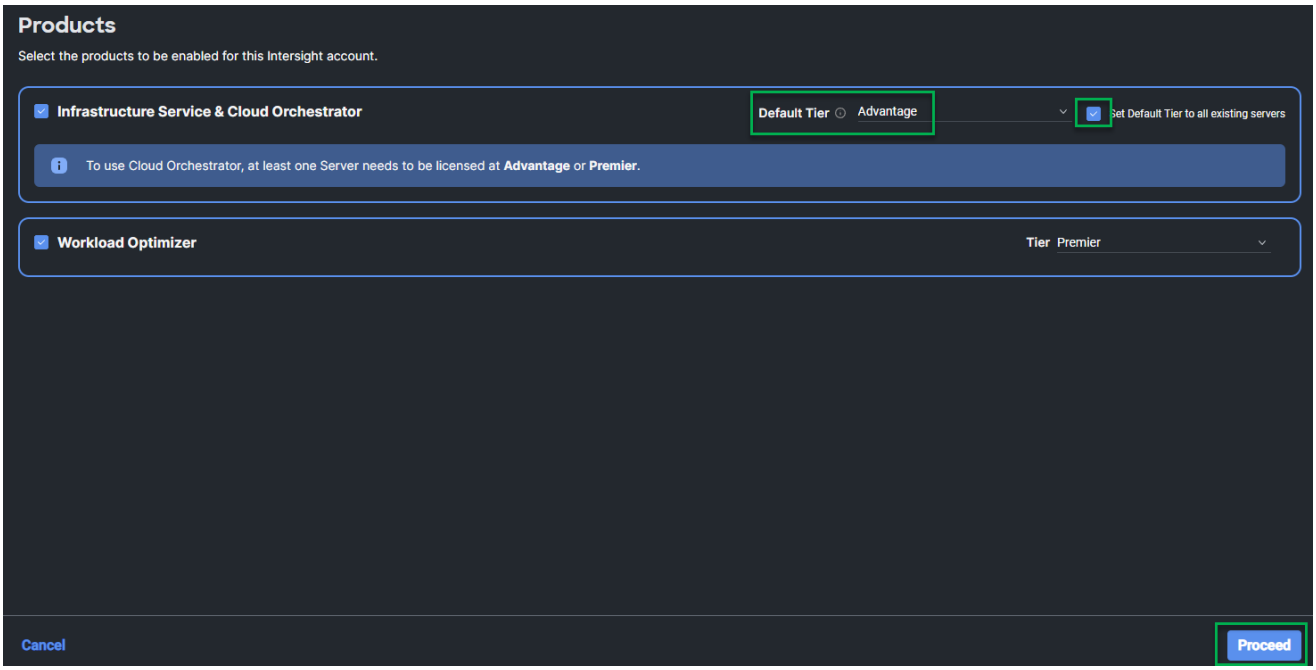


10. Click **Register**, and then click **Register** again.
11. When the registration is successful (it takes a few minutes), the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

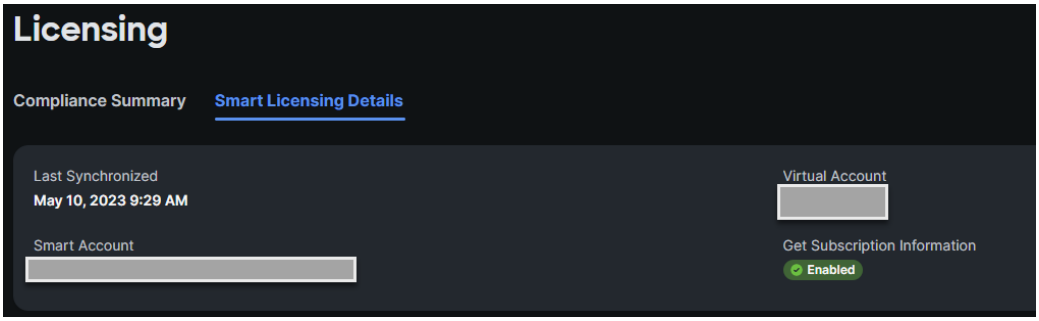


### Procedure 4. Set up Cisco Intersight Resource Group

In this procedure, you create a Cisco Intersight resource group where resources such as targets are logically grouped. In this deployment, you will create a single resource group to host all the resources. Note that the default resource group is created by default. You can use the default resource group or create a new one, doing the following: Note that the name of the resource group should be the same as the organization.

1. Log into the Cisco Intersight platform.
2. At the top, select System. On the left, click **Settings** (the gear icon).
3. Click **Resource Groups** in the middle panel.
4. Click **+ Create Resource Group** in the top-right corner.
5. Provide a name for the Resource Group (for example, RICH).

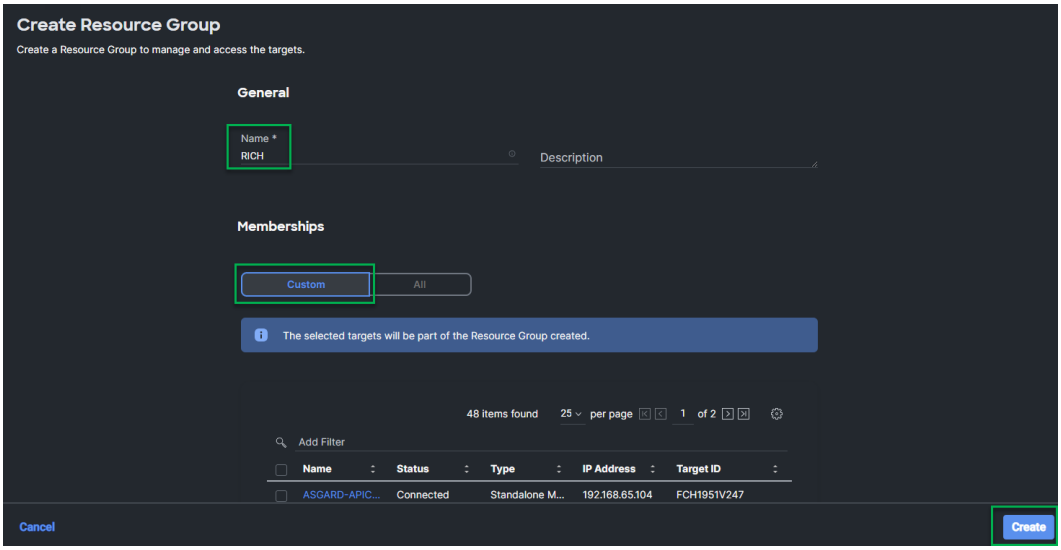


6. Under Memberships, select **Custom**.
7. Click **Create**.

### Procedure 5. Set up Cisco Intersight Organization

In this step, an Intersight Organization is created where all Cisco Intersight Managed Mode configurations, including policies, are defined.

1. Log into the Cisco Intersight portal.
2. At the top, select System. On the left, click **Settings** (the gear icon).
3. Click **Organizations** in the middle panel.
4. Click **+ Create Organization** in the top-right corner.
5. Provide a name for the organization (for example, RICH).
6. Select the Resource Group created in the last step (for example, RICH).
7. Click **Create**.

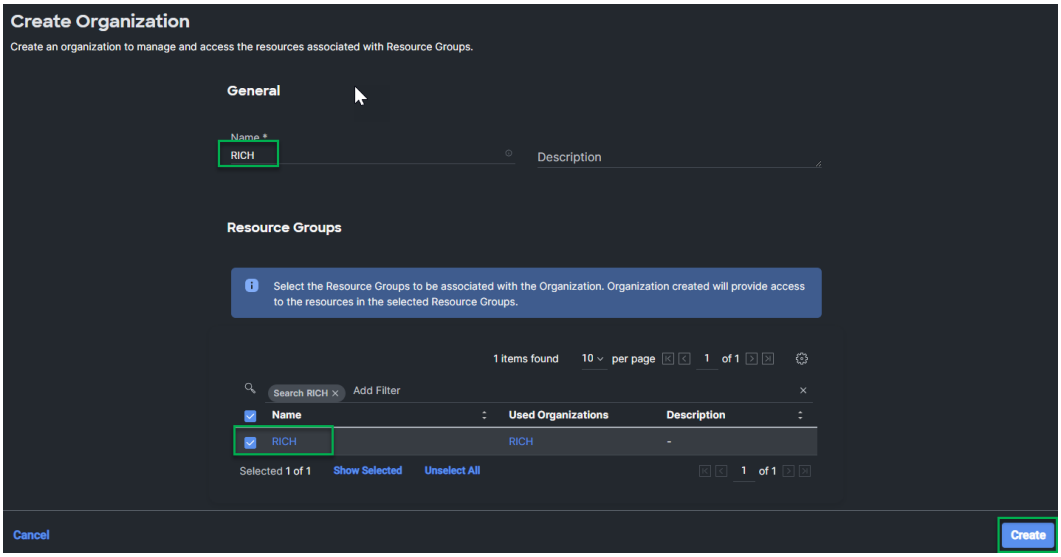

**Procedure 6. Claim Cisco UCS Fabric Interconnects/Domain in Cisco Intersight portal.**

Make sure the initial configuration for the FIs has been completed. Log into the FI-A Device Console using a web browser to capture the Cisco Intersight connectivity information.

1. Use the management IP address of FI-A to access the device from a web browser and the previously configured admin password to log into the device.
2. Under Device Connector, the current device status will show "Not claimed." Note or copy the Device ID, and Claim Code information for claiming the device in the Cisco Intersight application.

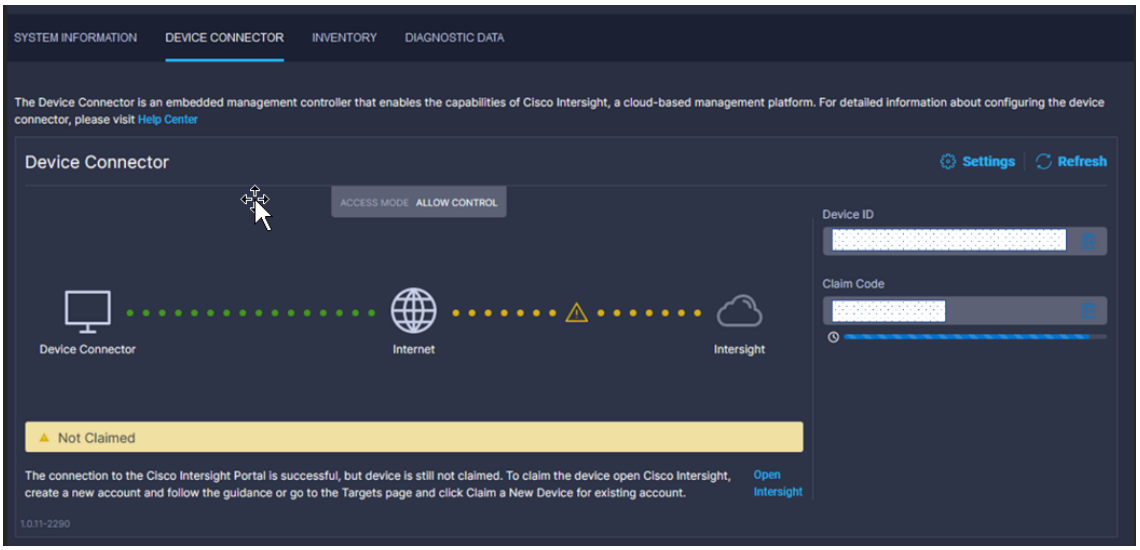


3. Log into the Cisco Intersight portal.
4. At the top, select System. On the left, click **Administration > Targets.**
5. Click Claim a New Target.
6. Select Cisco UCS Domain (Intersight Managed) and click Start.

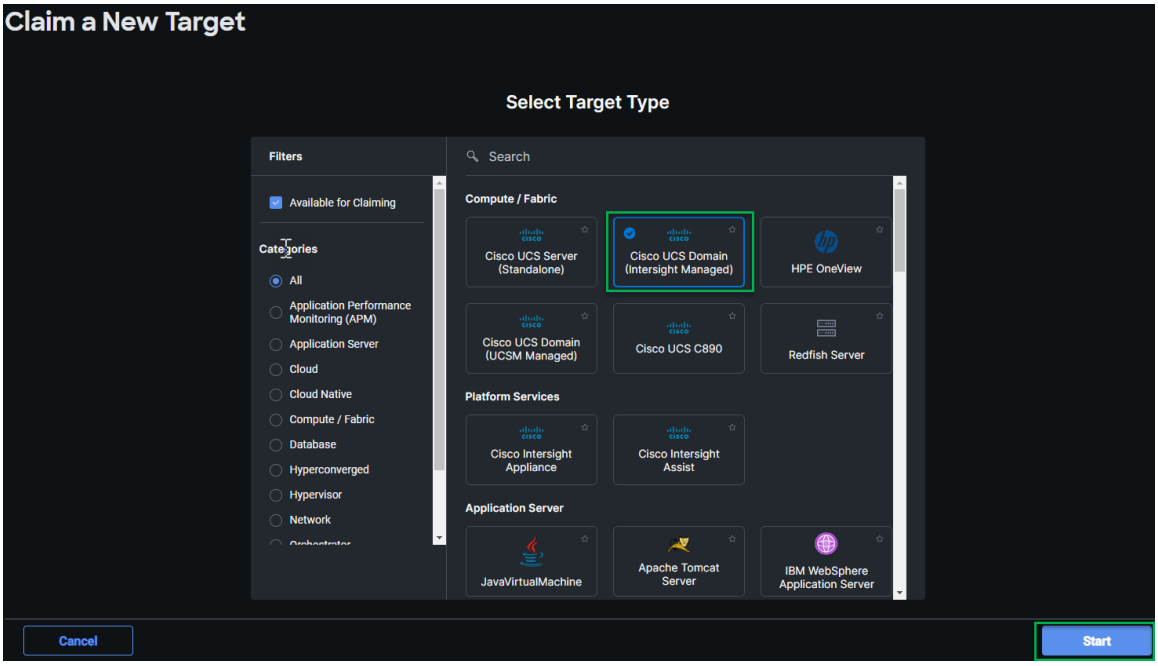

7. Copy and paste the Device ID and Claim from the Cisco UCS FI to Intersight application.
8. Select the previously created Resource Group and click **Claim**.

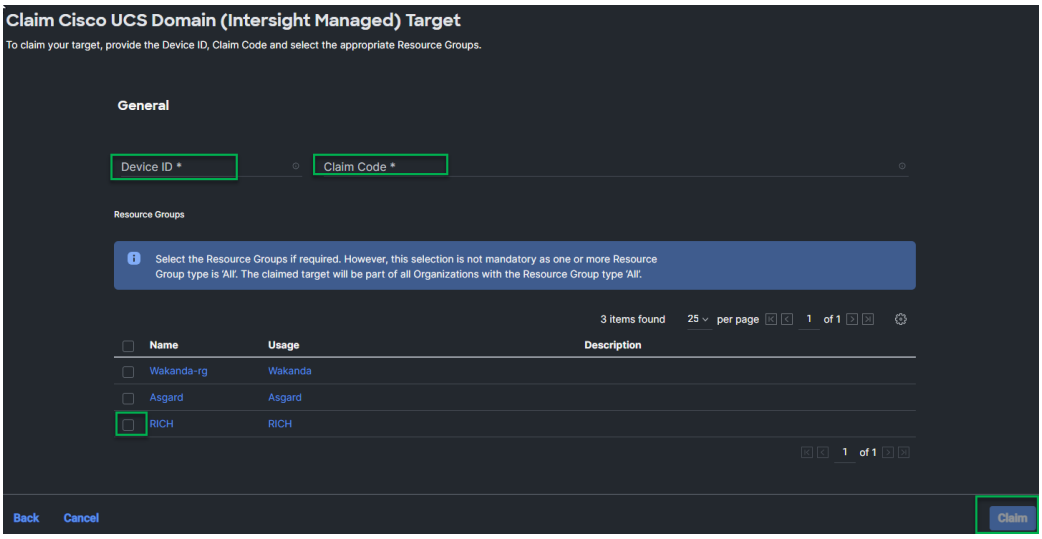


9. With a successful device claim, Cisco UCS FI should appear as a target in the Cisco Intersight application.

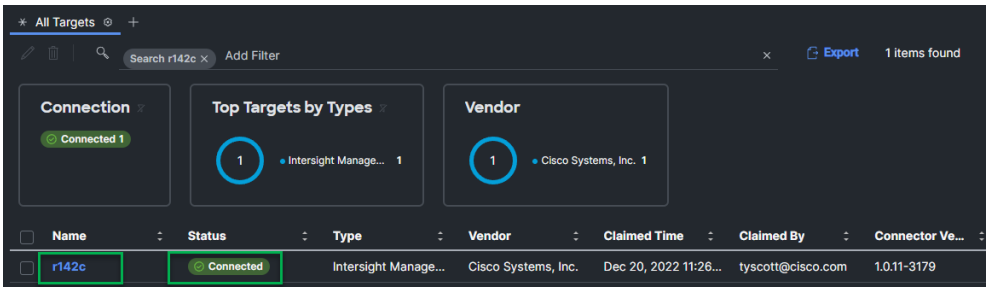


## Procedure 7. Verify addition of Cisco FIs to Cisco Intersight platform

1. Log into the web GUI of the Cisco UCS FI and click the browser refresh button.

   The FI status should now be set to **Claimed**.

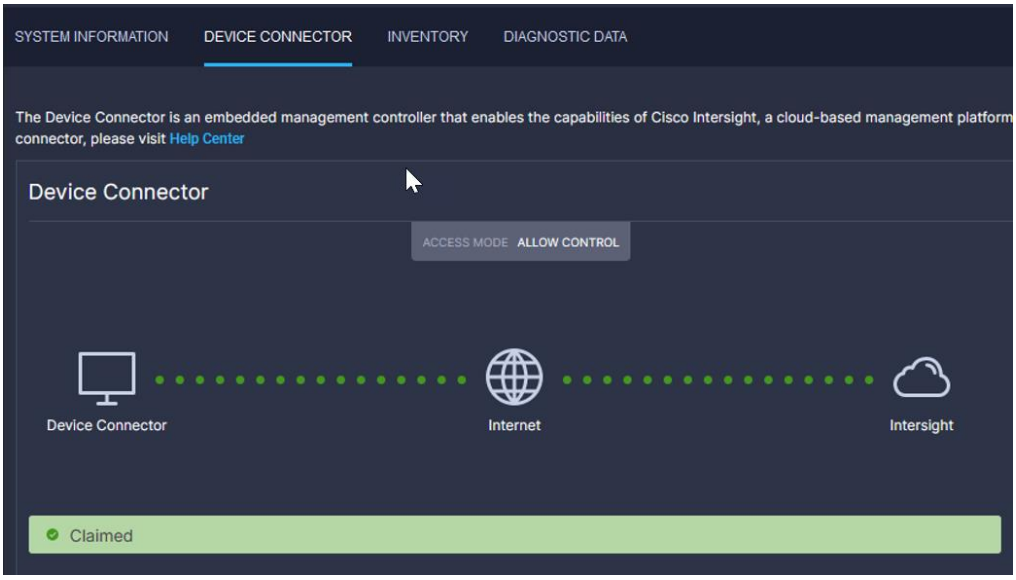

**Procedure 8. Upgrade FI firmware using Cisco Intersight platform**

**Note:** If the Cisco UCS FIs were upgraded to the latest recommended software using Cisco UCS Manager, this upgrade process is still required through the Cisco Intersight portal to install the X-Series firmware.

1. Log into the Cisco Intersight portal.
2. At the top, from the drop-down list, select **Infrastructure Service** and then select **Fabric Interconnects** under Operate on the left.
3. Click the ellipses "..." at the end of the row for either of the FIs and select **Upgrade Firmware**.
4. Click **Start**.
5. Verify the FI information and click **Next**.
6. Enable **Advanced Mode** using the toggle switch and uncheck Fabric Interconnect Traffic Evacuation.
7. Select the 4.2(3d) release from the list and click **Next**.
8. Verify the information and click **Upgrade** to start the upgrade process.
9. Watch the Request panel of the main Cisco Intersight screen because the system will ask for user permission before upgrading each FI. Click the circle with the arrow and follow the prompts on the screen to grant permission.
10. Wait for both the FIs to successfully upgrade.

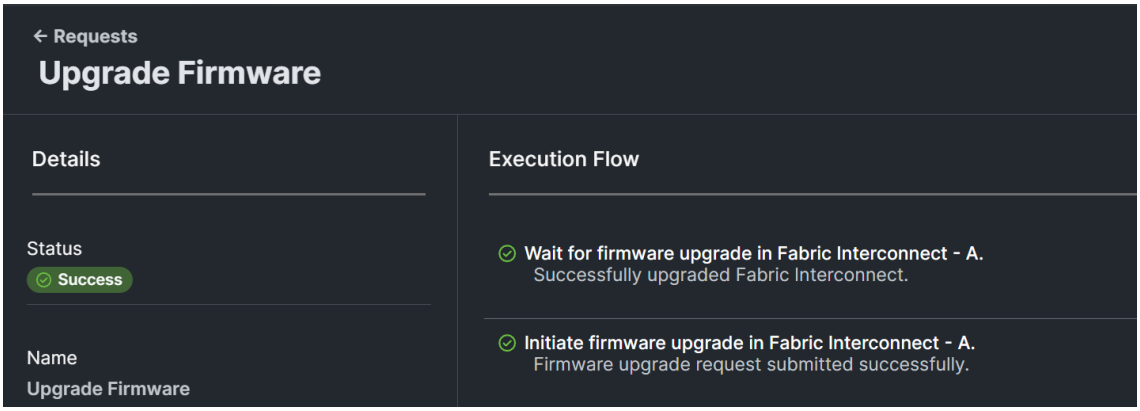


# Deploy VMware and Cisco Intersight management virtual machines

This chapter contains the following:

- Download and Install VMware vCenter 7.0U3h
- Download and Install Cisco Intersight Assist appliance
- Download and Install imm toolkit virtual machine
- Download and Install Cisco Intersight Transition Tool appliance

## Download and Install VMware vCenter 7.0U3h

The procedures in the following sections provide detailed instructions for installing the VMware vCenter 7.0U3h Server Appliance in a FlexPod environment.

### Procedure 1. Download vCenter 7.0U3h from VMware

**Note:** You will need a VMware user id and password to download this software.

1. Click this link: Download VMware vCenter Server 7.0U3h and download the **VMware vCenter Server Appliance**: VMware-VCSA-all-7.0.3-20395099.iso.
2. You will need a VMware user id and password on vmware.com to download this software.

### Procedure 2. Install the VMware vCenter Server Appliance

**Note:** The VCSA deployment consists of two stages: installation and configuration.

1. Locate and copy the **VMware-VCSA-all-7.0.3-20395099.iso** file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0 U3 vCenter Server Appliance.
2. Mount the ISO image as a disk on the management workstation. (For example, with the *Mount* command in Windows Server 2012 and later or right click the image and select **Mount**).
3. In the mounted disk directory, navigate to the **vcsa-ui-installer > win32** directory and double-click `installer.exe`. The vCenter Server Appliance Installer wizard will appear.
4. Click **Install** to start the vCenter Server Appliance deployment wizard.
5. Click **NEXT** in the Introduction section.
6. Read and accept the license agreement and click **NEXT**.
7. In the **vCenter Server deployment target** screen, enter the Fully Qualified Domain Name (FQDN) or IP address of the destination host, **User name,** and **Password**. Click **NEXT**.

**Note:** Installation of vCenter on a separate existing management infrastructure vCenter is recommended. If a separate management infrastructure is not available, you can choose the recently configured first ESXi host as an installation target. The recently configured ESXi host is shown in this deployment.

8. Click **YES** to accept the certificate.
9. Enter the Appliance **VM name** and **root password** details shown in the **Set up vCenter Server VM** section. Click **NEXT**.
10. In the **Select deployment size** section, select the Deployment size and Storage size. For example, select **Deployment size: Small** and **Storage size: Default**. Click **NEXT**.
11. **Select datastore** (for example, infra_datastore) for storage. Click **NEXT**.
12. In the **Network Settings** section, configure the following settings:
    - Select a Network: (for example, **IB-MGMT Network**).

**Note:** When the vCenter is running on FlexPod, it is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and not moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, trying to bring up vCenter on a different host than the one it was running on before the shutdown will cause problems with the network connectivity. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to a different host on vSwitch0 does not require vCenter to already be up and running.

- IP version: **IPV4**
- IP assignment: **static**
- FQDN: <vcenter-fqdn>
- IP address: <**vcenter-ip**>
- Subnet mask or prefix length: <**vcenter-subnet-mask**>
- Default gateway: <**vcenter-gateway**>
- DNS Servers: <dns-server1>,<dns-server2>

13. Click **NEXT**.
14. Review all values and click **FINISH** to complete the installation.

**Note:** The vCenter Server appliance installation will take a few minutes to complete.

15. When Stage 1, Deploy vCenter Server, is complete, Click **CONTINUE** to proceed with stage 2.
16. Click **NEXT**.
17. In the vCenter Server configuration window, configure these settings:
    - Time Synchronization Mode: Synchronize time with NTP servers.
    - NTP Servers: NTP server IP addresses from IB-MGMT VLAN
    - SSH access: **Enabled**.
18. Click **NEXT**.
19. Complete the Single Sign-On (SSO) configuration as shown in the following screenshot (or according to your organization's security policies):

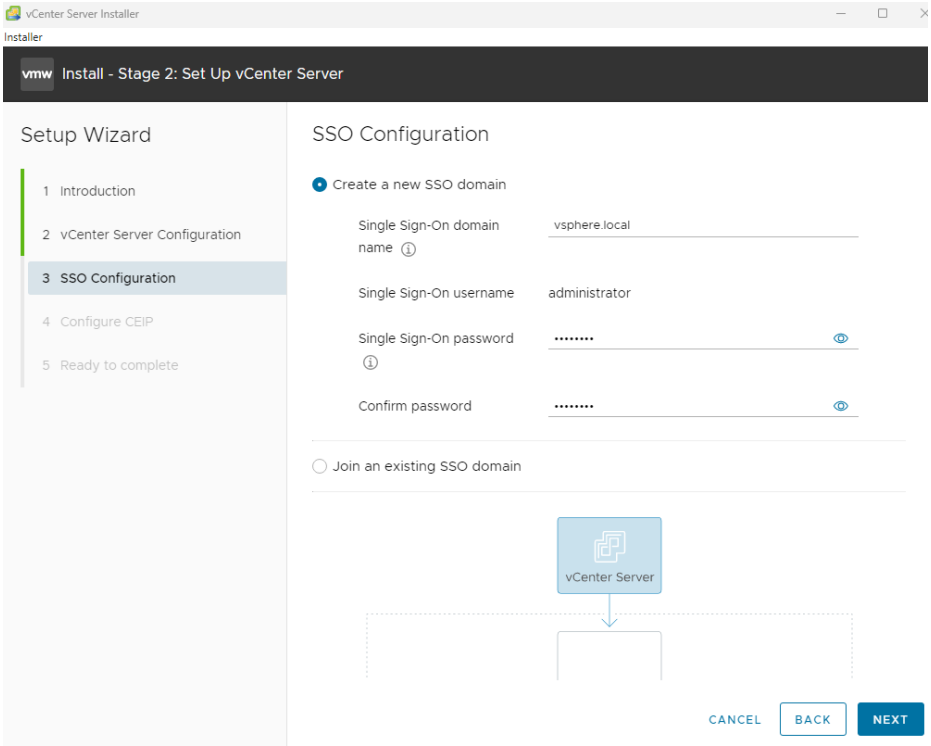

20. Click **NEXT**.
21. Decide whether to join VMware's Customer Experience Improvement Program (CEIP).
22. Click **NEXT**.
23. Review the configuration and click **FINISH**.
24. Click **OK**.

**Note:** vCenter Server setup will take a few minutes to complete and Install – Stage 2 with show Complete.

25. Click **CLOSE**. Eject or unmount the VCSA installer ISO.

### Procedure 3. Verify vCenter CPU settings

**Note:** If a vCenter deployment size of small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS X210c M6 and B200 M6 servers are 2-socket servers. During this validation, the **Small** deployment size was selected and vCenter was set up for a 4-socket server. This setup can cause problems in the VMware ESXi cluster Admission Control.

1. Open a web browser on the management workstation and navigate to the vCenter or ESXi server where the vCenter appliance was deployed and log in.
2. Click **vCenter VM**, right-click and click **Edit settings**.
3. In the **Edit settings** window, expand CPU and check the value of Sockets.
4. If the number of Sockets matches the server configuration, click **Cancel**.
5. If the number of Sockets does not match the server configuration, it will need to be adjusted:
6. Right-click the vCenter VM and click **Guest OS > Shut down**. Click **Yes** on the confirmation.
7. When vCenter is shut down, right-click the vCenter VM and click **Edit settings**.
8. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to that on the server configuration.

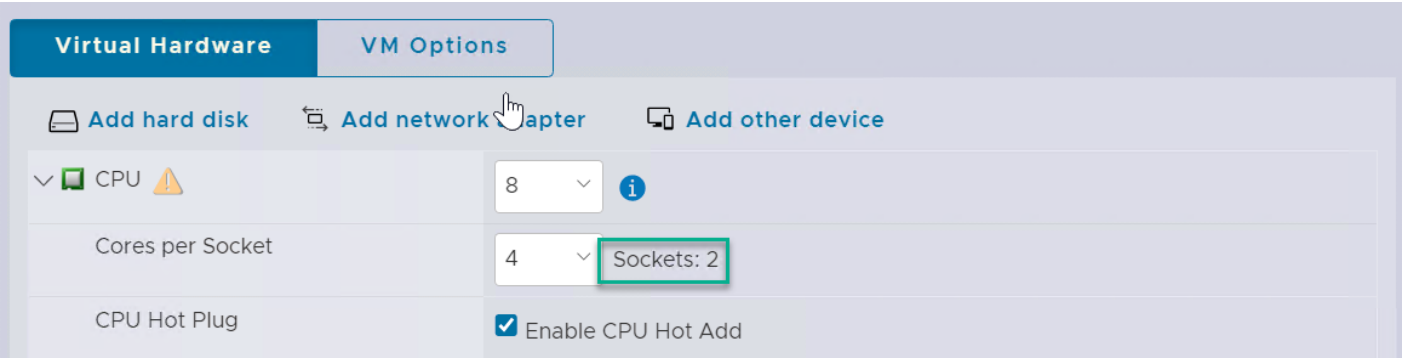


9. Click **Save**.
10. Right-click the vCenter VM and click **Power > Power on**. Wait approximately 10 minutes for vCenter to come up.

**Procedure 4. Set up VMware vCenter Server**

1. Using a web browser, navigate to https://<vcenter-ip-address>:5480. Navigate to the security screens.
2. Log into the **VMware vCenter Server Management** interface as **root** with the root password set in the vCenter installation.
3. In the menu on the left, click **Time**.
4. Click **EDIT** to the right of Time zone.
5. Select the appropriate Time zone and click **SAVE**.
6. In the menu on the left select **Administration**.
7. According to your Security Policy, adjust the settings for the root user and password.
8. In the menu on the left click **Update**.
9. Follow the prompts to stage and install any available vCenter updates.
10. In the upper right-hand corner of the screen, click **root > Logout** to logout of the Appliance Management interface.
11. Using a web browser, navigate to https://<vcenter-fqdn> and navigate through security screens.

**Note:** With VMware vCenter 7.0 and later, you must use the vCenter FQDN.

12. Select LAUNCH VSPHERE CLIENT (HTML5).

    The VMware vSphere HTML5 Client is the only option in vSphere 7. All the old clients have been deprecated.

13. Log in using the SSO username (for example administrator@vsphere.local) and password created during the vCenter installation. Dismiss the Licensing warning.

**Procedure 5. Add AD user authentication to vCenter (optional)**

1. In the **AD Infrastructure**, using the Active Directory Users and Computers tool, set up a Domain Administrator user with a username such as flexadmin (FlexPod Admin).
2. Connect to https://<vcenter-fqdn> and select LAUNCH VSPHERE CLIENT (HTML5).
3. Log in as **administrator@vsphere.local** (or the SSO user set up in vCenter installation) with the corresponding password.
4. Under the top-level menu, click **Administration**. In the list on the left, under **Single Sign On**, select **Configuration.**
5. In the center pane, under **Configuration**, select the **Identity Provider** tab.
6. In the list under **Type**, select **Active Directory Domain**.
7. Click **JOIN AD**.
8. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click **JOIN**.
9. Click Acknowledge.

10. In the list on the left under **Deployment**, click **System Configuration**. Select the radio button to select the vCenter, then click **REBOOT NODE**.
11. Input a reboot reason and click **REBOOT**. The reboot will take approximately 10 minutes for full vCenter initialization.
12. Log back into the vCenter vSphere HTML5 Client with the SSO Credentials.
13. Under the top-level menu, click **Administration**. In the list on the left, under **Single Sign On**, click **Configuration**.
14. In the center pane, under **Configuration**, click **the Identity Provider** tab. Under **Type**, select **Identity Sources**. Click **ADD**.
15. Make sure Active Directory (Integrated Windows Authentication) is selected, your Windows Domain name is listed, and Use machine account is selected. Click **ADD**.
16. In the list select the **Active Directory (Integrated Windows Authentication)** Identity source type. If desired, select SET AS DEFAULT and click **OK**.
17. On the left under Access Control, select **Global Permissions**.
18. In the center pane, click **ADD** to add a Global Permission.
19. In the **Add Permission** window, select your AD domain for the Domain.
20. On the User/Group line, enter either the FlexPod Admin username or the Domain Admins group. Leave the Role set to Administrator. Check the box for **Propagate to children**.

**Note:** The FlexPod Admin user was created in the Domain Admins group. The selection here depends on whether the FlexPod Admin user will be the only user used in this FlexPod or if additional users will be added later. By selecting the Domain Admins group, any user placed in that AD Domain group will be able to log into vCenter as an Administrator.

21. Click **OK** to add the selected User or Group. The user or group should now appear in the Global Permissions list with the Administrator role.
22. Log out and log back into the vCenter HTML5 Client as the FlexPod Admin user. You will need to add the domain name to the user; for example, flexadmin@example.com.

## Download and Install Cisco Intersight Assist appliance

This appliance works with NetApp's ONTAP storage and VMware vCenter using third-party device connectors and Cisco Nexus and MDS switches using Cisco device connectors. Since third-party infrastructure and Cisco switches do not contain any usable built-in Cisco Intersight device connector, Cisco Intersight Assist virtual appliance enables the appliance to communicate with these devices.

**Note:** A single Cisco Intersight Assist virtual appliance can support NetApp ONTAP storage, VMware vCenter, and Cisco Nexus and MDS switches.

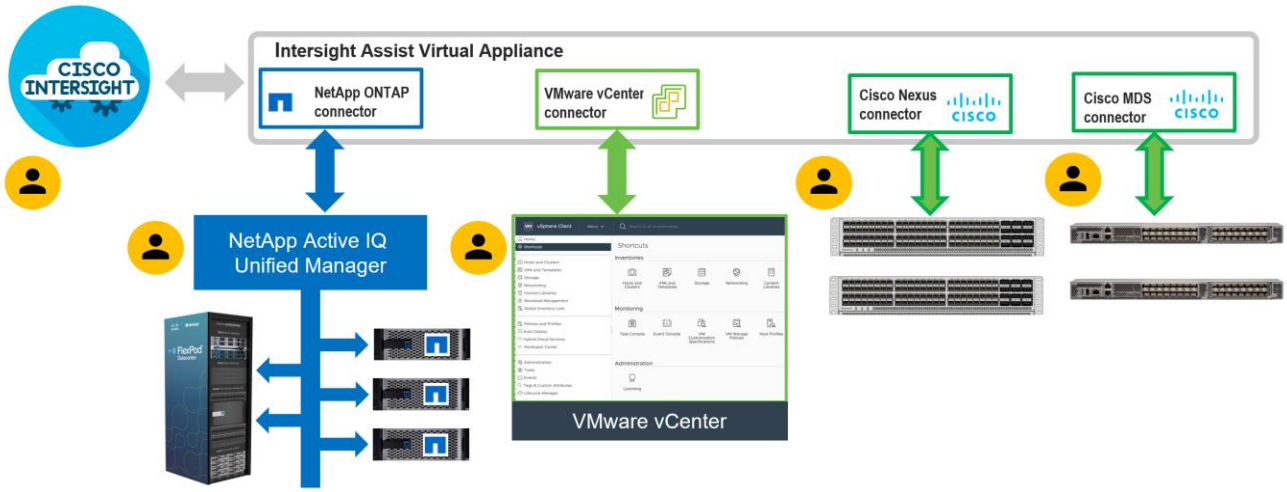


**Figure 4.**
Managing NetApp and VMware vCenter through Cisco Intersight using Cisco Intersight Assist

### Procedure 1. Install Cisco Intersight Assist

1. To install Cisco Intersight Assist from an Open Virtual Appliance (OVA), download the latest release of the Cisco Intersight Virtual Appliance for vSphere from Cisco Software Download.

**Note:** It is important to install Release 1.0.9-499 at a minimum.

### Procedure 2. Set up DNS entries.

1. Setting up Cisco Intersight Virtual Appliance requires an IP address and two hostnames for that IP address. The hostnames must be in the following formats:

- **myhost.mydomain.com**: A hostname in this format is used to access the GUI. It must be defined as an A record and Pointer record (PTR) record in DNS. The PTR record is required for reverse lookup of the IP address. If an IP address resolves to multiple hostnames, the first one in the list is used.
- **dc-myhost.mydomain.com:** The dc- must be prepended to your hostname. This hostname must be defined as the CNAME of myhost.mydomain.com. The appliance uses hostnames in this format internally by to manage device connections.

2. In this lab deployment, the following information was used to deploy a Cisco Intersight Assist Virtual Machine:

- **Hostname:** intersight-assist.example.com
- **IP address**: 198.18.1.96
- **DNS Entries** (Windows AD/DNS):
  - A Record and CNAME:

| | | |
|---|---|---|
| intersight-assist | Host (A) | 198.18.1.96 |
| dc-intersight-assist | Alias (CNAME) | intersight-assist.example.com |

  - PTR (reverse lookup):

| | | |
|---|---|---|
| 198.18.1.96 | Pointer (PTR) | intersight-assist.example.com. |

For more information, refer to: Cisco Intersight Virtual Appliance and Intersight Assist Getting Started Guide

**Procedure 3. Deploy Cisco Intersight OVA.**

**Note:** Ensure that the appropriate entries of type A, CNAME, and PTR records exist in the DNS, as explained in the previous section. Log into the vSphere Client and select **Hosts and Clusters.**

1. From Hosts and Clusters, right-click the cluster and click **Deploy OVF Template**.
2. Select Local file and click **UPLOAD FILES**. Browse to and select the intersight-appliance-installer-vsphere-1.0.9-342.ova or the latest release file and click **Open**. Click **NEXT**.
3. Name the Intersight Assist Virtual Machine and select the location. Click **NEXT**.
4. Select the cluster and click **NEXT**.
5. Review details, click **Ignore All**, and click **NEXT**.
6. Select a deployment configuration. If you need only the Cisco Intersight Assist functions, you can use a deployment size of **Tiny**. If you are using Cisco Intersight Workload Optimizer (IWO) in this Cisco Intersight account, use the **Small** deployment size. Click **NEXT**.
7. Select the appropriate datastore for storage and select the **Thin Provision** virtual disk format. Click **NEXT**.
8. Select an appropriate management network (for example, IB-MGMT Network) for the OVA. Click **NEXT**.

**Note:** The Cisco Intersight Assist Virtual Machine must be able to access both the IB-MGMT network on FlexPod and Intersight.com. Select and configure the management network appropriately. If you are selecting the IB-MGMT network on FlexPod, make sure the routing and firewall are set up correctly to access the Internet.

9. Fill in all values to customize the template. Click **NEXT**.
10. Review the deployment information and click **FINISH** to deploy the appliance.
11. When the OVA deployment is complete, right-click the Intersight Assist VM and click **Edit Settings**.
12. Expand CPU and verify the socket configuration. For example, in the following deployment, on a 2-socket system, the VM was configured for 16 sockets:

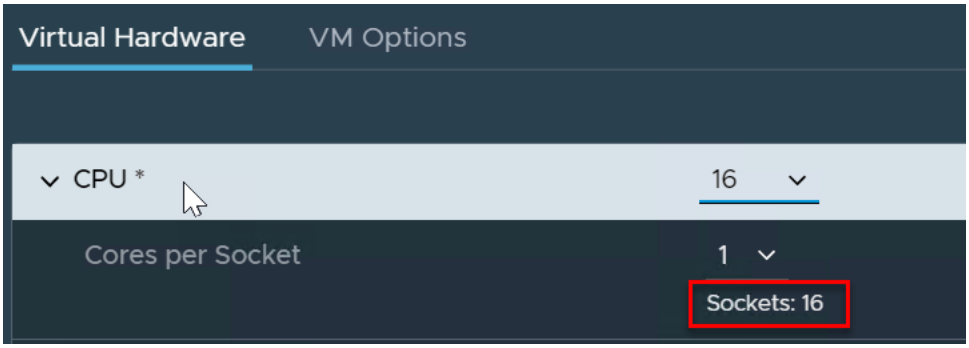


13. Adjust the Cores per Socket so that the number of sockets matches the server CPU configuration (2 sockets in this deployment):

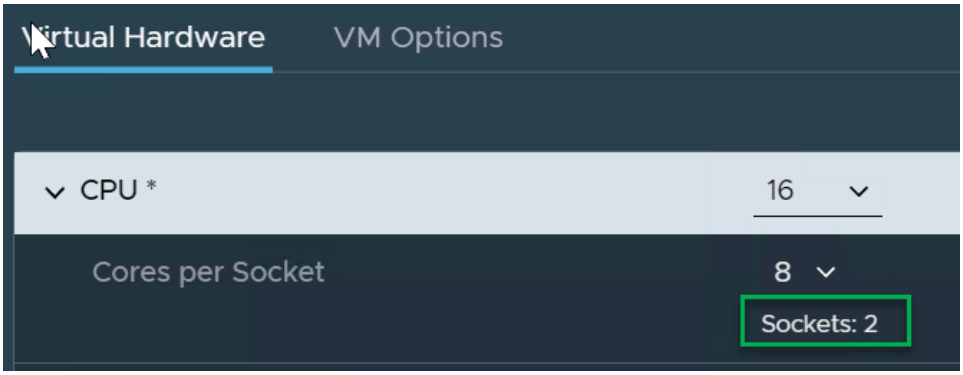


14. Click **OK**.
15. Right-click the Cisco Intersight Assist Virtual Machine and select **Power > Power On**.
16. When the virtual machine powers on and the login prompt is visible (use remote console), connect to https://intersight-assist-fqdn.

**Note:** It may take a few minutes for https://intersight-assist-fqdn to respond.

17. Navigate the security prompts and select **Intersight Assist**. Click **Start**.

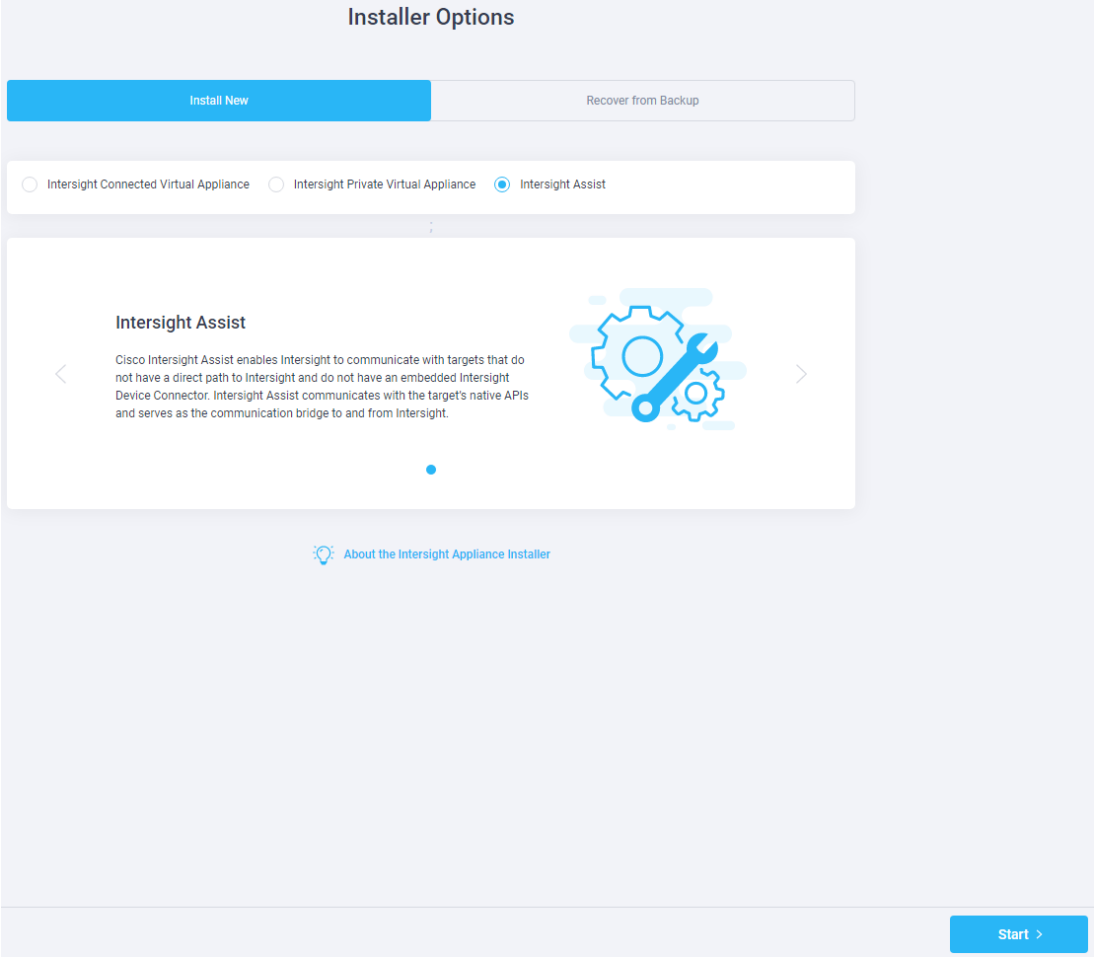


18. Cisco Intersight Assist Virtual Machine needs to be claimed in the Cisco Intersight appliance using the Device ID and Claim Code information visible in the GUI.

19. Log into the Cisco Intersight appliance and connect to the appropriate account.
20. From the Cisco Intersight appliance, at the top select **System**, then click **Administration > Targets**.
21. Click **Claim a New Target**. Select Cisco Intersight Assist and click **Start**.

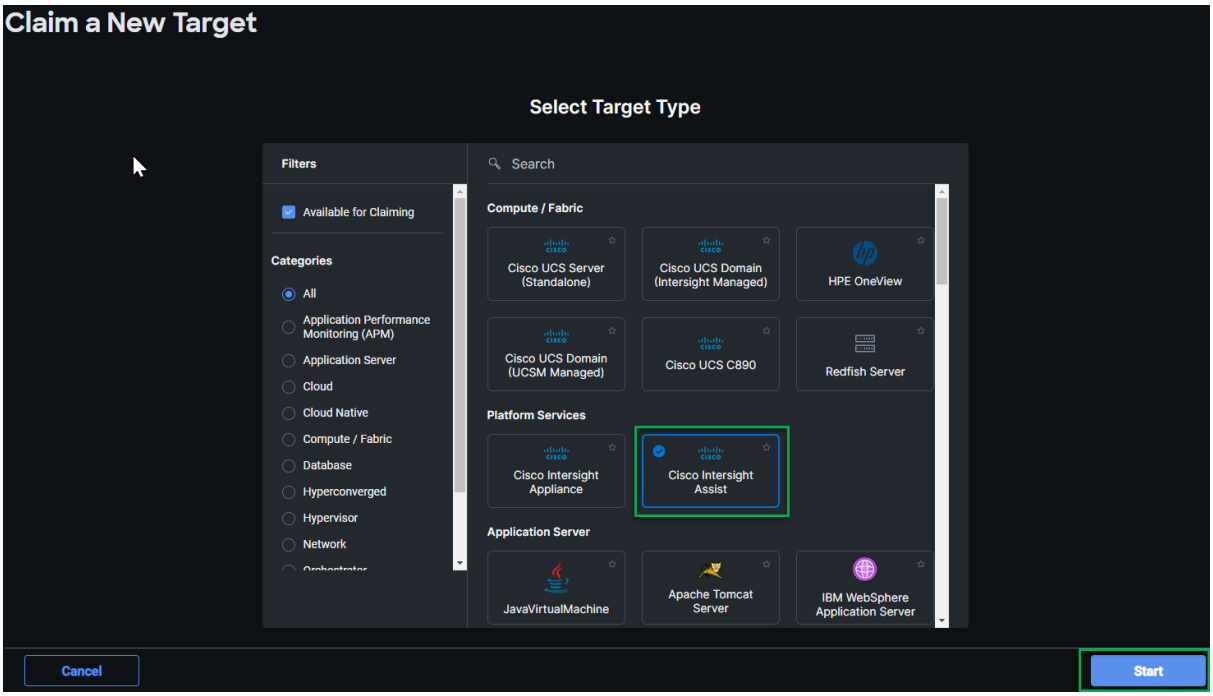


22. Copy and paste the Device ID and Claim Code shown in the Cisco Intersight Assist web interface to the Cisco Intersight Device Claim window.
23. We recommend putting Cisco Intersight Assist into the “default” Resource Group/Organization. If adding to the “default” Resource Group, no Resource Group selection is required.

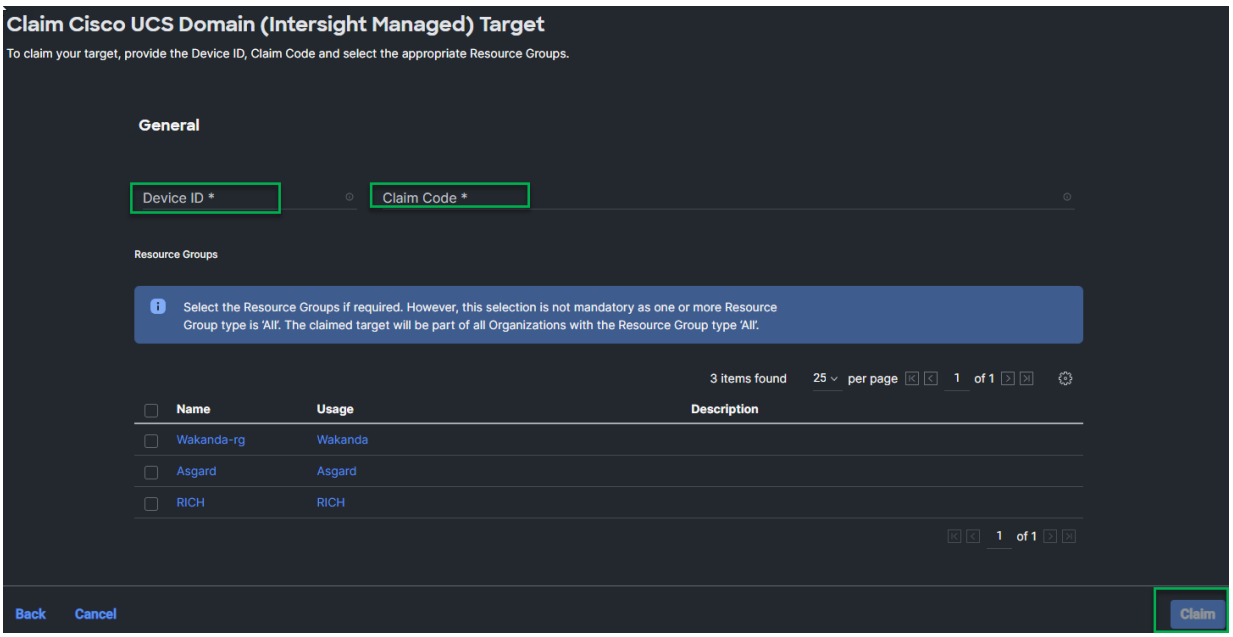


24. Cisco Intersight Assist will now appear as a claimed device.

25. In the Cisco Intersight Assist web interface, verify that the appliance is connected successfully, and click **Continue**.

**Note:** The Cisco Intersight Assist software will now be downloaded and installed into the Cisco Intersight Assist Virtual Machine. This process can take up to an hour to complete.

**Note:** The Cisco Intersight Assist Virtual Machine will reboot during the software download process. You must refresh the web browser after the reboot is complete to follow the status of the download process.

26. When the software download is complete, a Cisco Intersight Assist login screen will appear.
27. Log into Cisco Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Cisco Intersight Assist status and **log out** of the software.

## Download and Install imm toolkit virtual machine

### Procedure 1. Install the imm-toolkit automation virtual machine

The imm Toolkit virtual machine provides a pre-configured environment with all the required automation tools, (Python, PowerShell + Ansible and Terraform), to provide the orchestration engines for the infrastructure deployment with Intersight Cloud Orchestrator (ICO). It will be added as an SSH Target in Cisco Intersight using the Cisco Intersight Assist appliance.

1. Download the imm toolkit ova from the following location: imm-toolkit
2. Log in to vCenter and select an existing cluster to deploy the "imm-toolkit" ova within.
3. Select **UPLOAD FILES**, locate the **imm-toolkitv0.1.5.ova** file which was previously downloaded, and click **Next**.

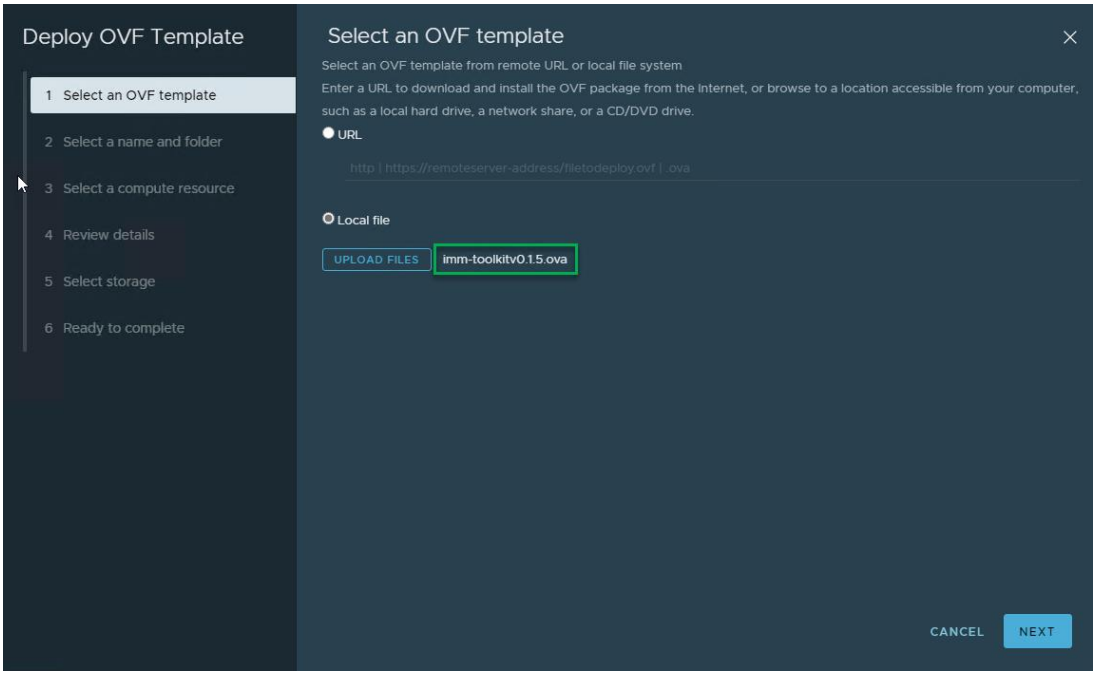


4. Assign the Virtual machine name and click Next.

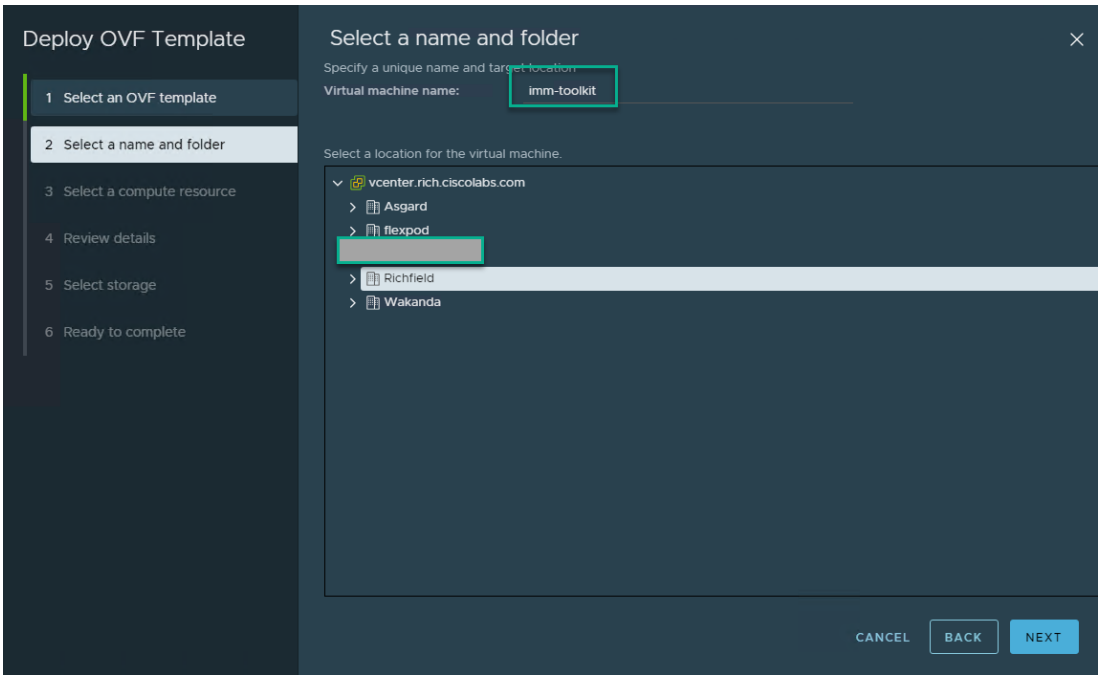


5. Select the **compute resource**. Click **Next**.

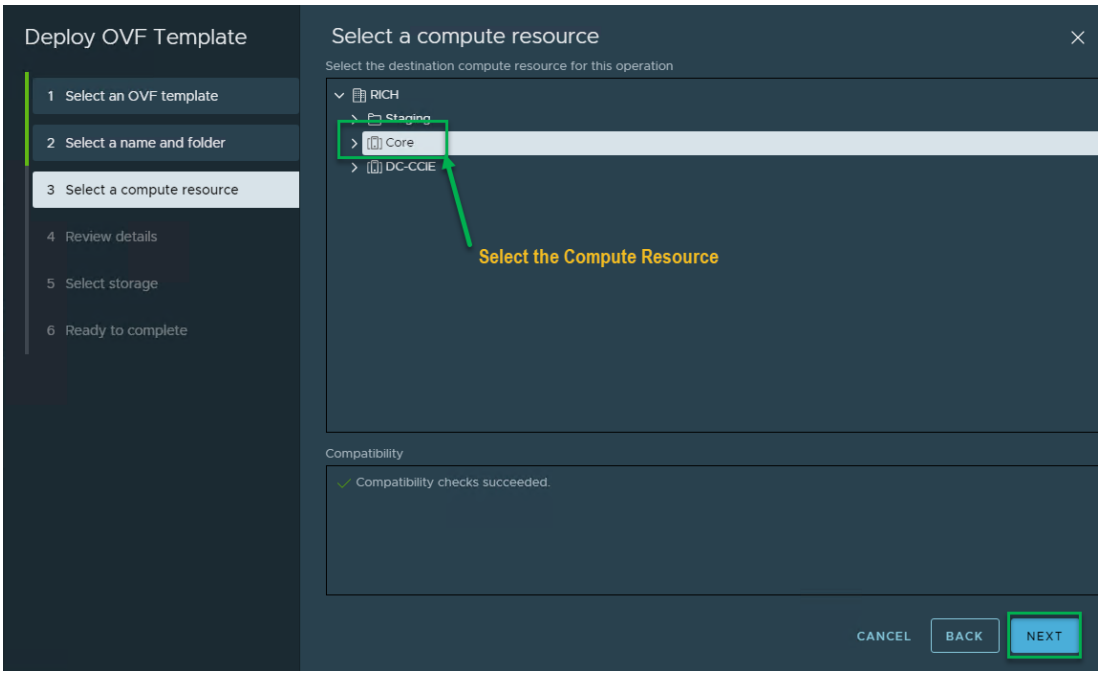


6. Assign the **storage** resource and Select the **virtual disk format**. Click **Next**.

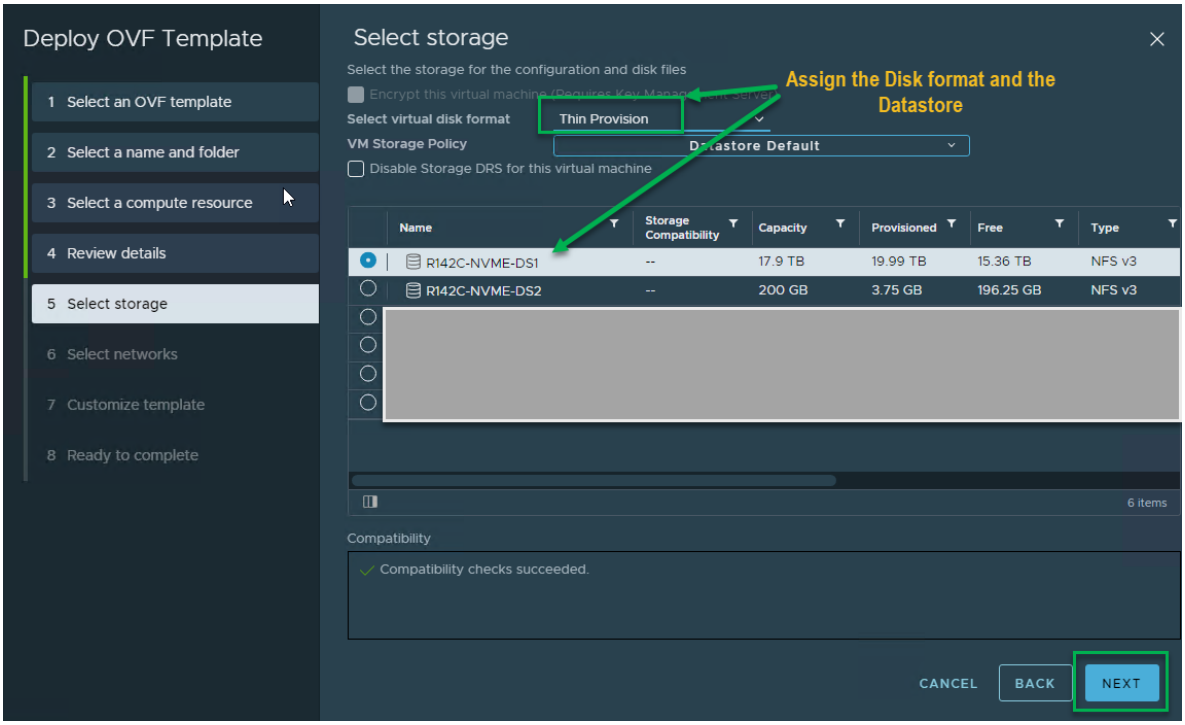


7. Select an existing network for the virtual appliance. In our environment VLAN 1001 (IB-MGMT) is already available and is the **VM Network** on this cluster. Click **Next**.

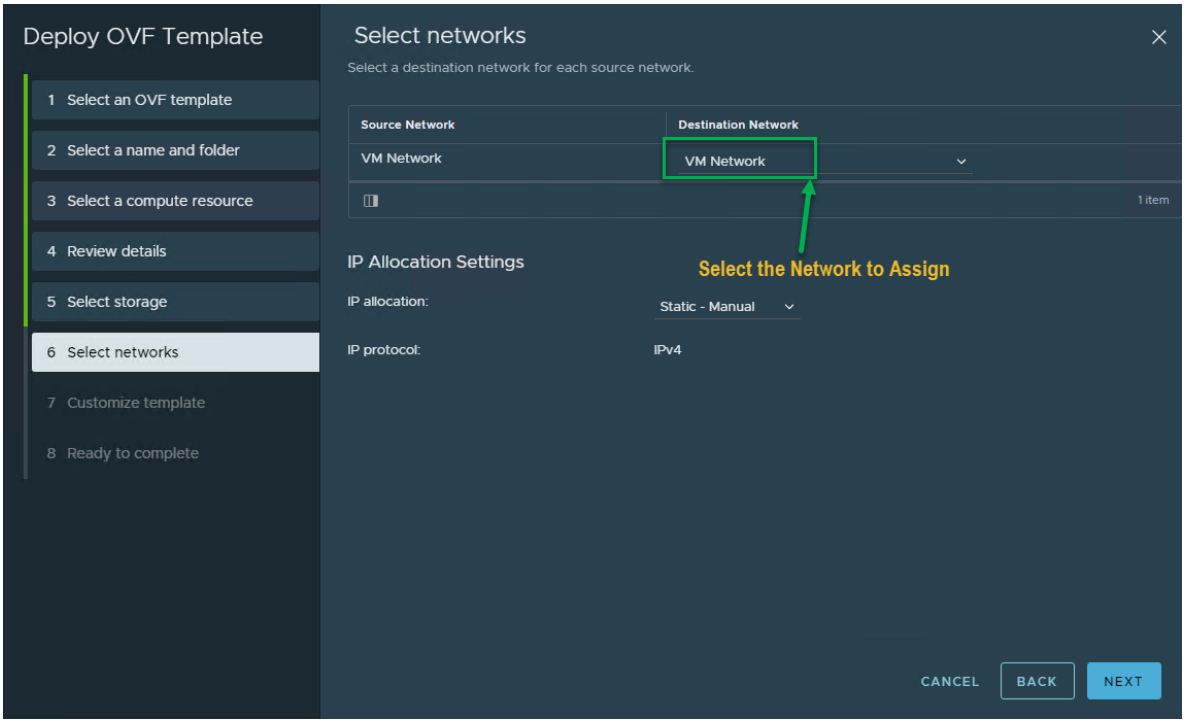


8. Customize the template with the settings for your environment.
   - IP Source: STATIC
   - Hostname: <imm-toolkit fqdn>
   - Network Prefix: <subnet-prefix>

- Gateway: <gateway>
- DNS Servers: <comma separated list>
- DNS Domains: <comma separated list of DNS domain suffixes>
- NTP Servers: <comma separated list of NTP servers>. Click **Next**.

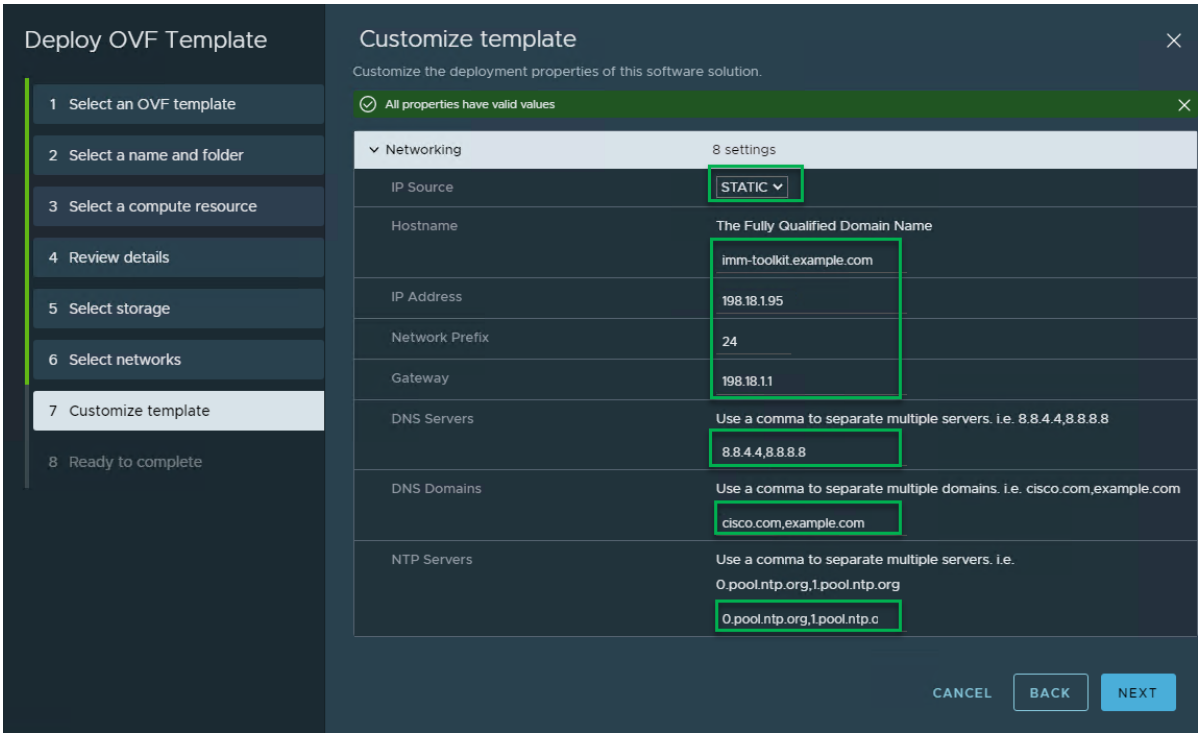


9. Log in to the IMM Toolkit Virtual Machine. The default password is “C1sc0123”. Change the password to a new value using the “`passwd`” Command.

```
imm-toolkit@imm-toolkit:~$ passwd
Changing password for imm-toolkit.
Current password:
New password:
Retype new password:
passwd: password updated successfully
imm-toolkit@imm-toolkit:~$
```

### Procedure 2. Add the imm-toolkit virtual machine as a Target to Intersight.

1. In Cisco Intersight go to System > Targets > Claim a New Target. Select Orchestrator in the left column and SSH Endpoint; and click Start.

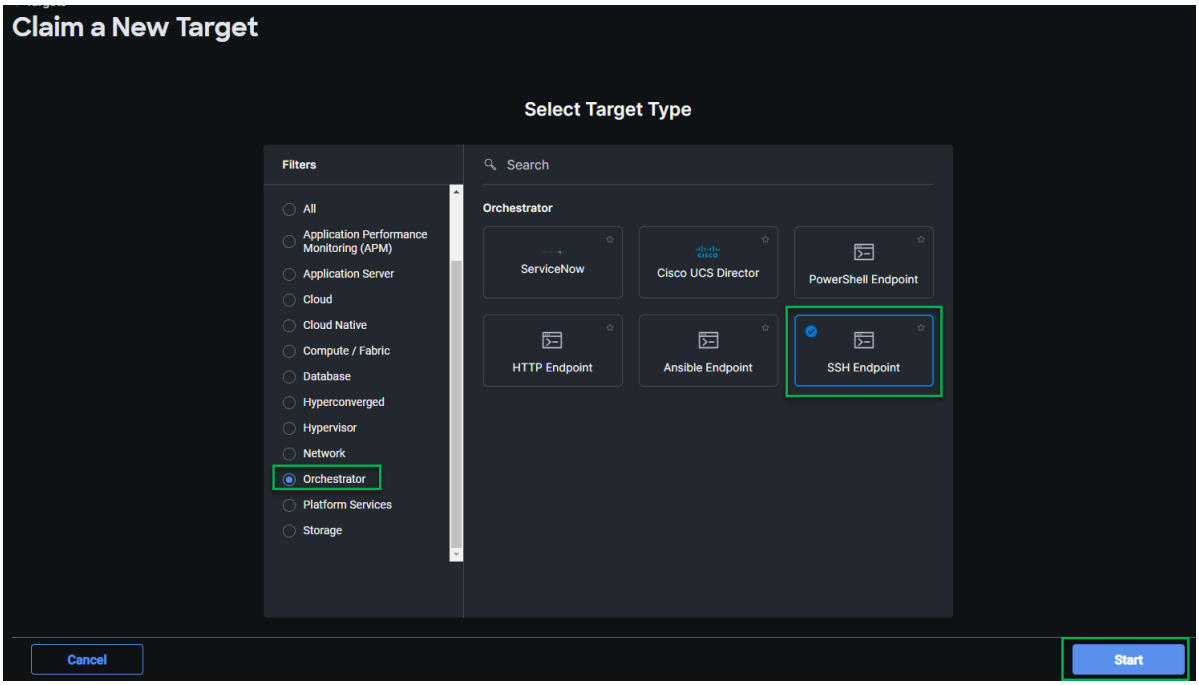


2. Fill in the information to connect to the host.
   - Select the Intersight Assist Appliance, previously deployed.
   - Assign a **Name** to the Target.
   - Fill in the Hostname or IP address For the “imm-toolkit” virtual machine.

   Add the Username and Password or configure key-based authentication on the system and add the keys here. Click **Claim**

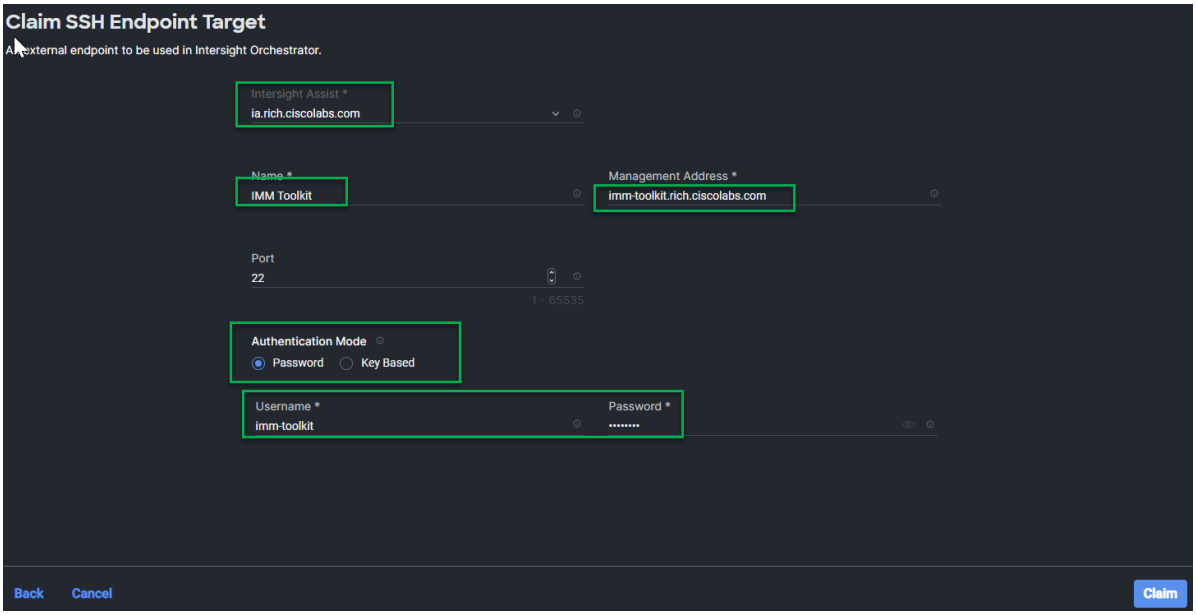

## Download and Install Cisco Intersight Managed Mode Transition Tool

### Procedure 1. Install the transition tool virtual machine

The Intersight Managed Mode Transition tool can be used to migrate existing UCS deployments from UCSM and UCS Central over to Cisco Intersight. It is extremely helpful in this transition as it will validate the health of the existing environment, pull in the configuration, and assigned identities.

Upon completion it will generate a health report and when ready import the configuration into Cisco Intersight. More information can be found here: Cisco Intersight Managed Mode Transition Tool

The purpose for this virtual machine for this white paper specifically is the new software repository added in v4.0.1. We will use this to host all the installation files and create the software repositories in Intersight to assign for OS installation, and later OS customization for VIB and package installations.

1. Download version 4.0.1 from the following location: UCS Tools page

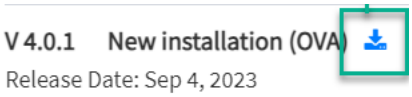


2. Log in to vCenter and select an existing cluster to deploy the OVA within.
3. Select **UPLOAD FILES**; locate the **IMM-Transition-4.0.1.ova** file which was previously downloaded, and click **Next**.

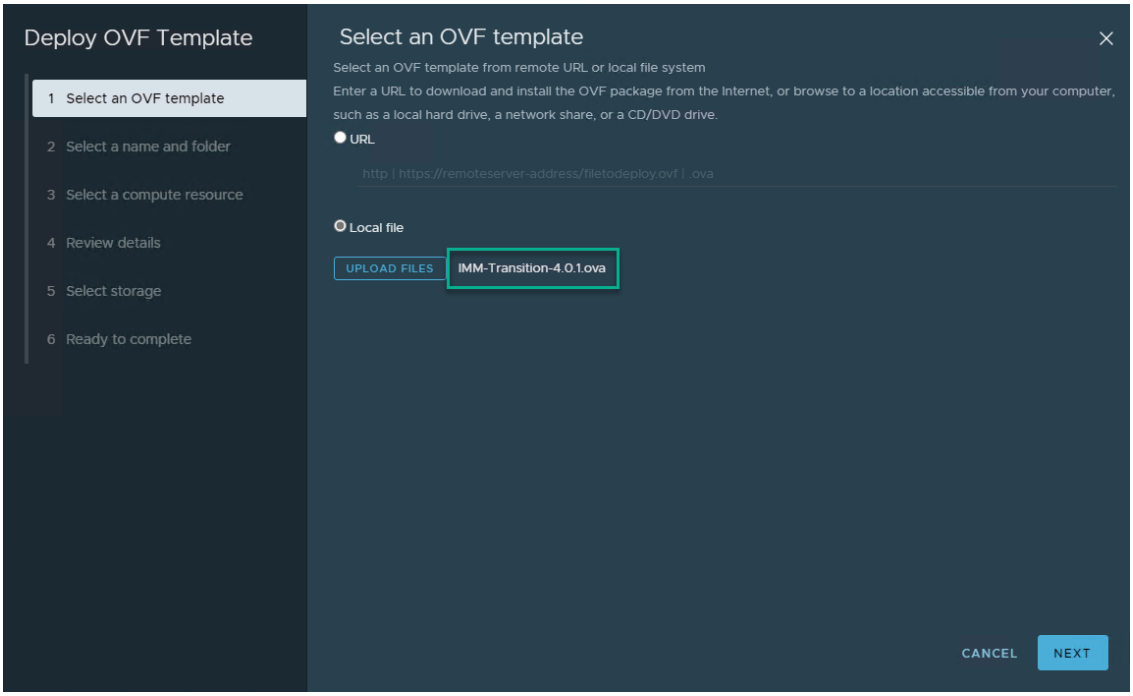


4. Assign the Virtual machine name and click Next.

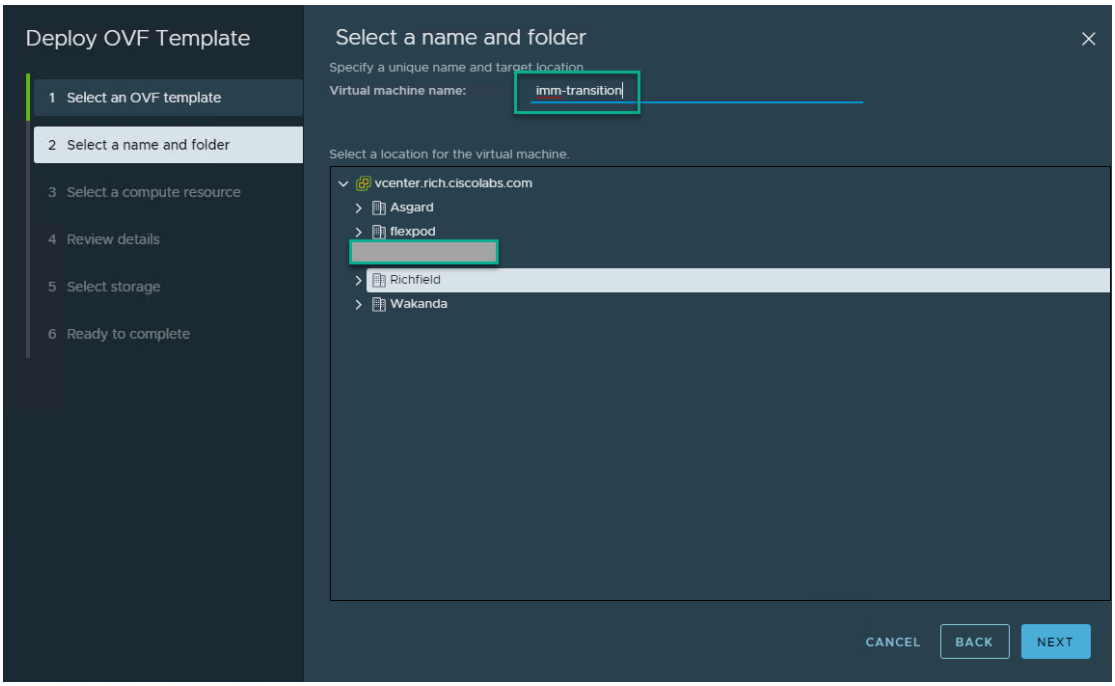


5. Select the **compute resource**. Click **Next**.

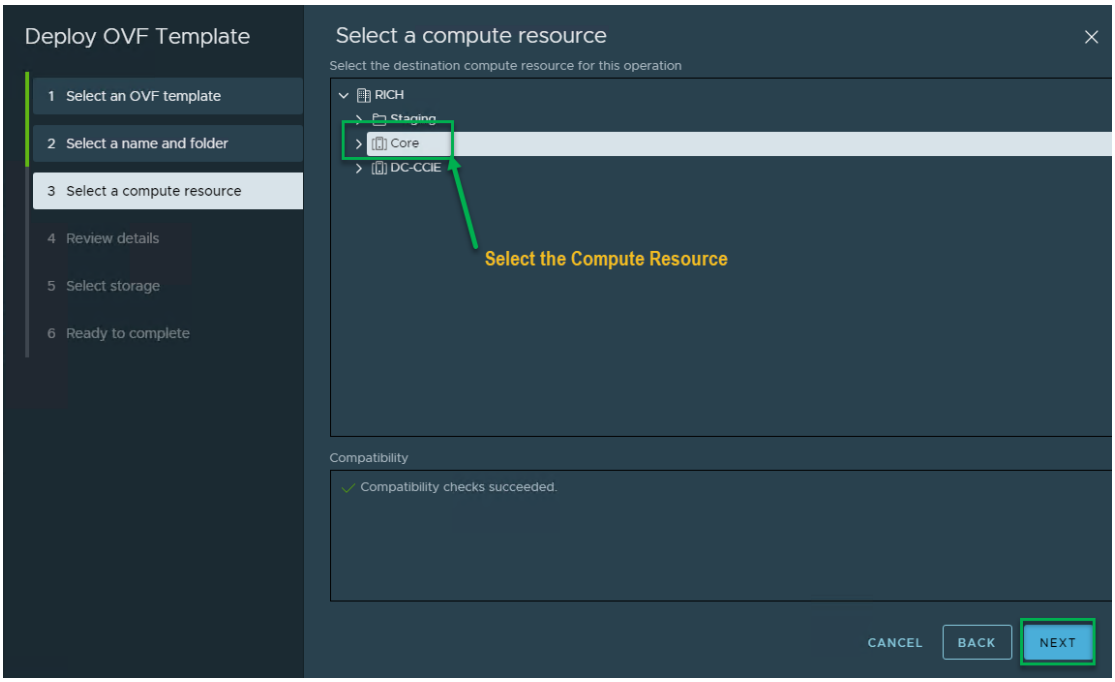


6. Assign the **storage** resource and Select the **virtual disk format**. Click **Next**.

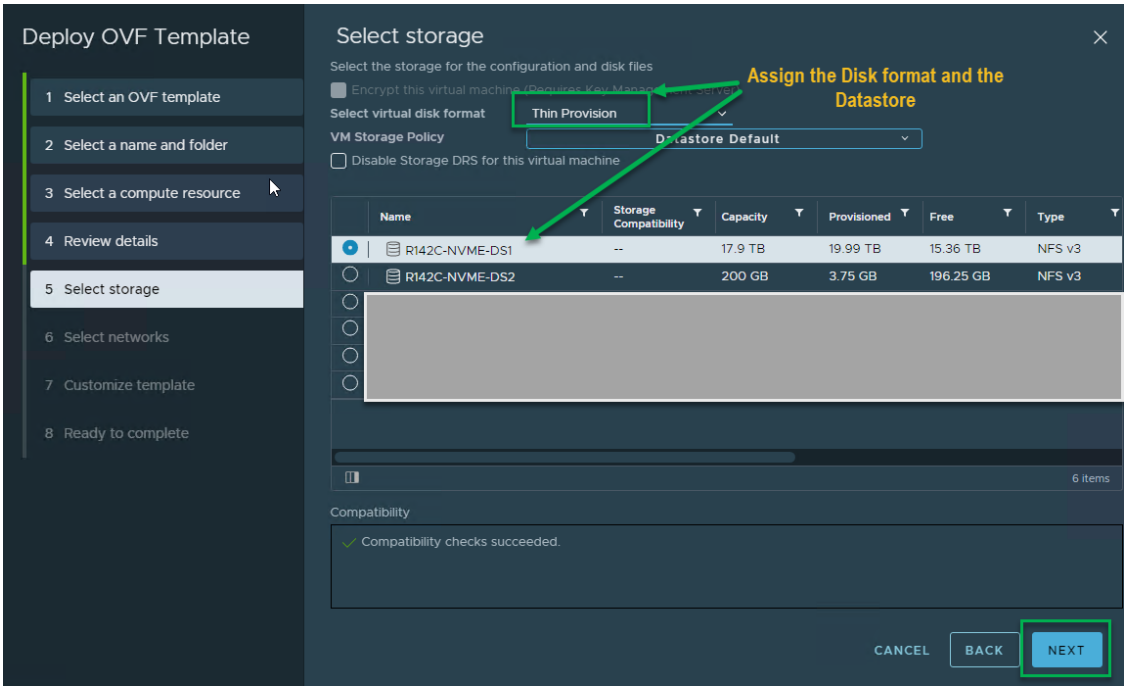


7. Select an existing network for the virtual appliance. In our environment VLAN 1001 (IB-MGMT) is already available and is the **VM Network** on this cluster. Click **Next**.

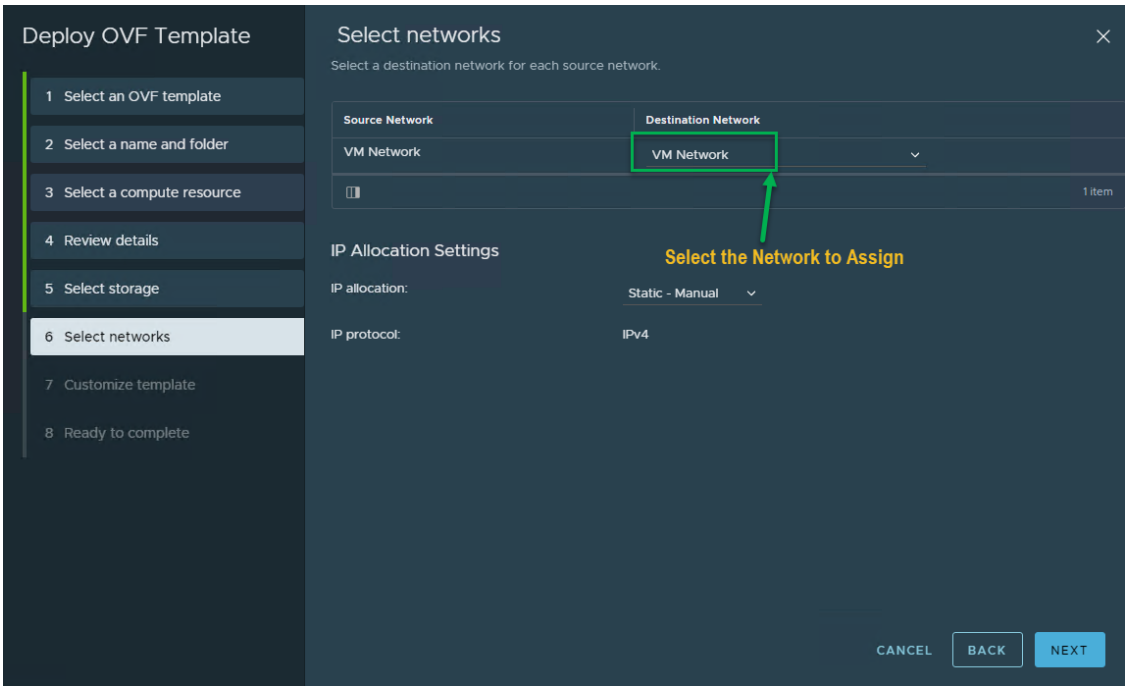


8. Customize the template with the settings for your environment.
   - IP Allocation: Static – Manual. Click **Next**.
   - Public Network Type: STATIC
   - Public Network IP: <imm-transition IP>
   - Public Network Netmask: <subnet-netmask>

- Public Network Gateway: <gateway>
- DNS: <comma separated list>
- NTP: <comma separated list of NTP servers>
- System Password: <admin password>
- Software Repository Disk Size: The default value of 100G is more than enough for our use case. You can increase this if you want to use it for other reasons.

Click **Next**.

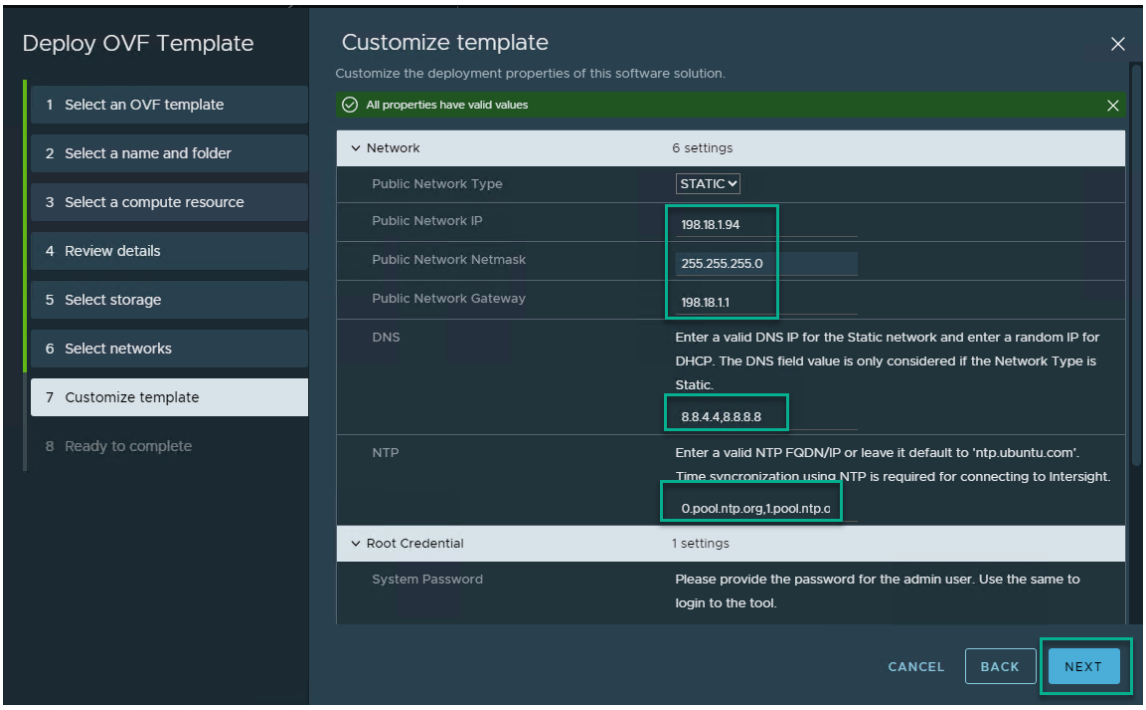


# Download Deployment images and add to software repository

This chapter contains the following:

- Download Deployment Images
- Add Images to the Transition tool Software Repository

## Download Deployment Images

### Procedure 1. Download Images and VIBs

**Note:** A valid Cisco VMware Customer Connect account is required to download images.

1. Download the Cisco Custom Image for ESXi 7.0 U3 Install CD
2. Select the VMware-ESXi-7.0.3i-20842708-Custom-Cisco-4.2.2-a.iso. Click Download.

**Note:** This image already has the 5.0.0.37 nfnic driver and the 1.0.45.0-1OEM nenic drivers that are needed for installation. If in the future the recommended driver's version changed and you needed to customize the installation media, go to section "Create a Custom ESXi ISO using VMware vCenter" for step-by-step instructions on how to build a custom ISO.

**Note:** A valid Cisco Connection Online (CCO) account is required to download the server configuration utility.

3. Download the Unified Computing System (UCS) Server Configuration Utility
4. Select version 6.2(3c) and Click Download

### Procedure 2. Obtain an API key and secret from Cisco Intersight

1. Login to Cisco Intersight (Cloud or On-Prem), i.e., http://intersight.com/
2. Login with your account credentials.
3. Go to System > **API Keys** and click **Generate API Key** on the top right.

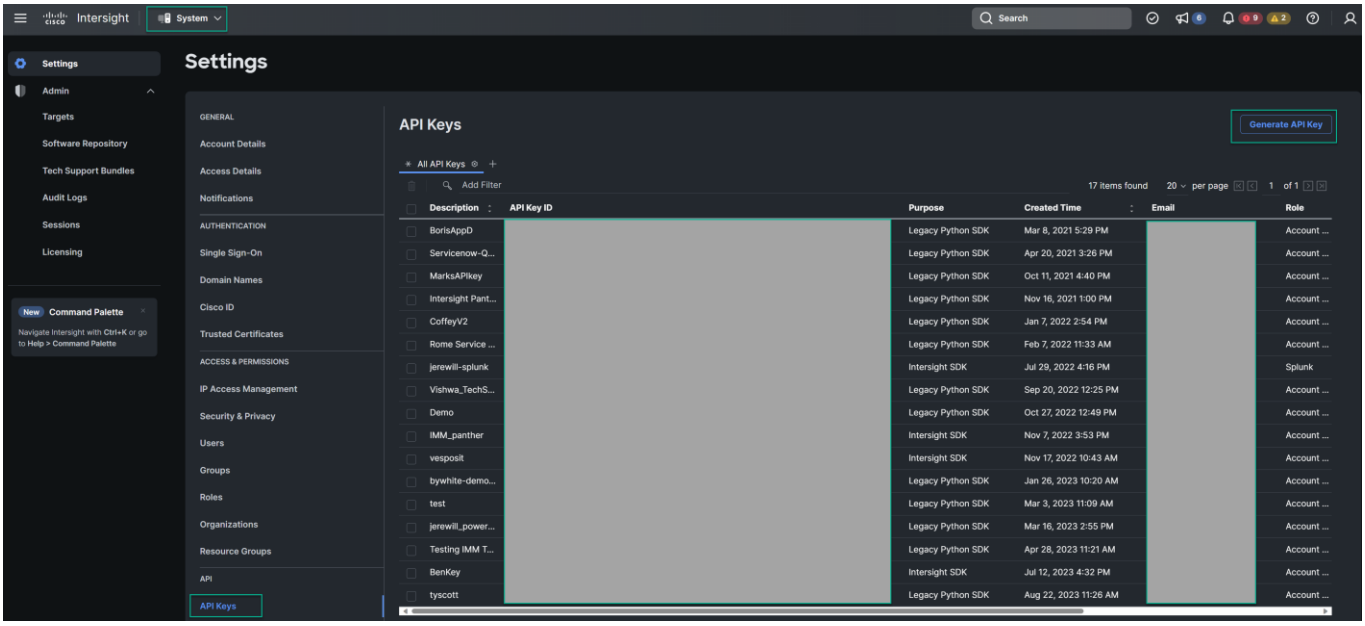


4. Provide a Description and **select API key for OpenAPI schema version 2**, click **Generate**.

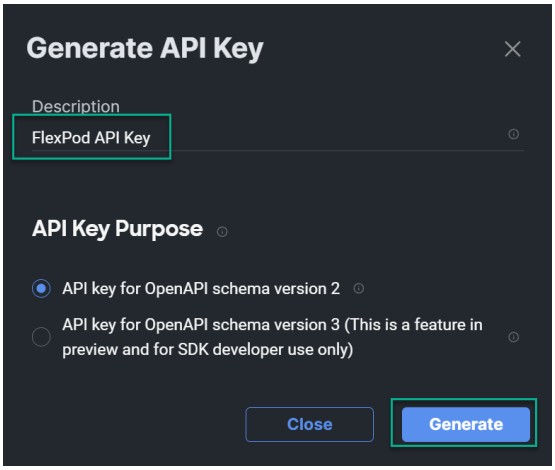


5. Copy the API Key ID and save the content to a secure location. Save the Secret Key text file in a secure location as well. These will be used by both the transition tool and Cisco Intersight Cloud Orchestrator (ICO) for the automation of the environment.

**Note:** This will be the only the time the Secret Key will be shown, as is detailed in the note at the top of the screenshot below. Make sure to save it to a secure location that you can access in the future. The API Key can be copied again from the portal if necessary.

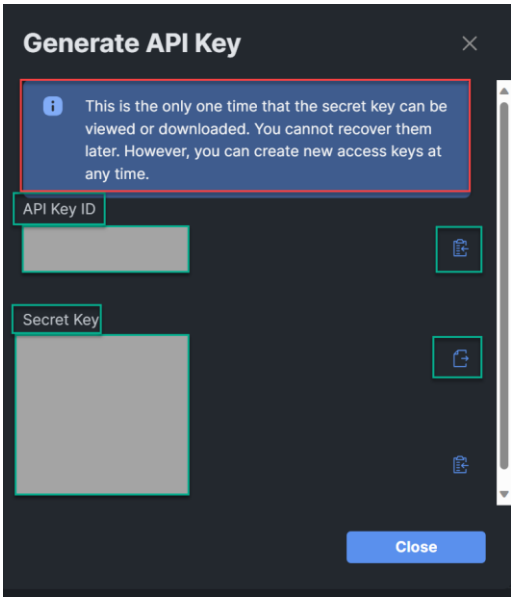


## Add Images to the Transition tool Software Repository

**Procedure 1. Add Images to transition tool repository**

1. Login to the Cisco Intersight Managed Mode Transition Tool: https://imm-transition-url
2. Enter the **Password** you assigned during the OVA deployment and click **Sign In**.
3. Our first step is to add an API key and secret to enable communication to the Cisco Intersight service. Click on Device Management on the left-hand column.

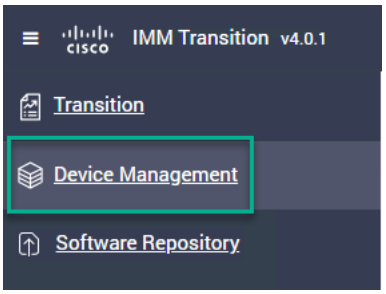


4. Click Add Device. Enter the Intersight instance details:
   - Device Type: Intersight
   - Intersight SaaS if using the cloud instance or Intersight Appliance VM for CVA/PVA
   - Enter the API Key and Secret Key obtained in the previous section.
   - Click Save.

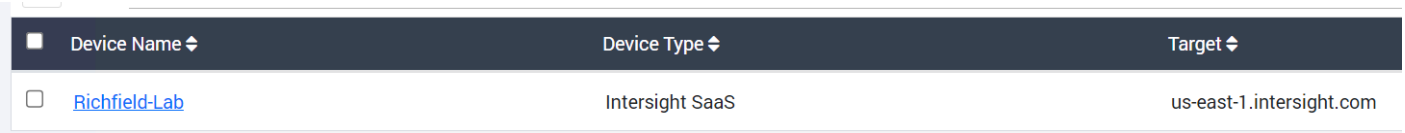

5. Click the **Software Repository** shortcut on the left-hand side.


or

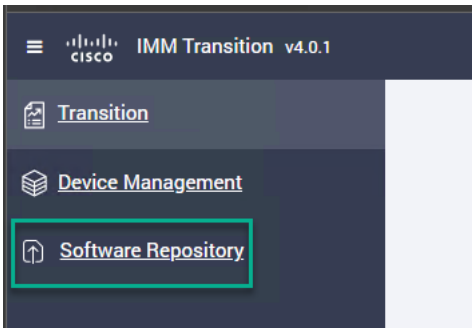


6. Click **New** to add the Images you have downloaded. Drag and drop the image to the browser and when complete you should see the following message in the lower right-hand corner.

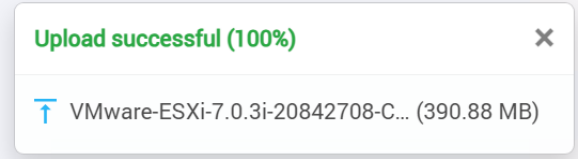


7. Repeat, for the Server Configuration Utility (SCU) image.

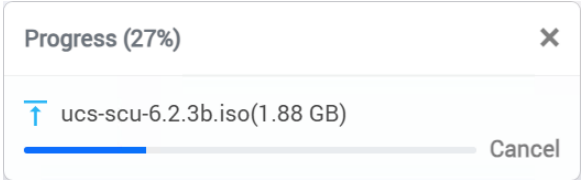


8. Sync the ESXi image with Cisco Intersight. Click the three ellipticals to the right of the image and select **Sync to Intersight**. Select the Intersight **Device** created earlier and the **Organization***. Click on **Fetch OS/Firmware data**. The Organization should be the Organization where the FlexPod solution will be deployed.

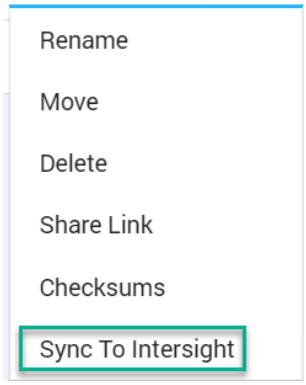


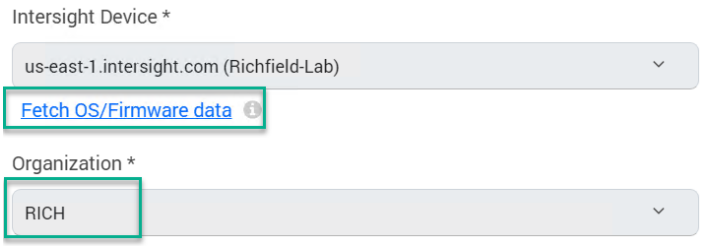


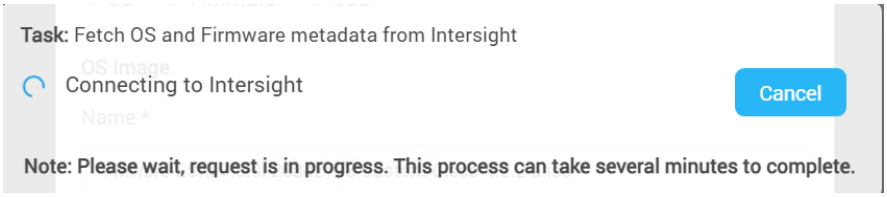


**Note:** This process will take a few minutes to complete.

9. You can Change the Name if you want, but make sure to set the **Vendor** and **Version** fields for the installation to succeed later. Click **Submit**.

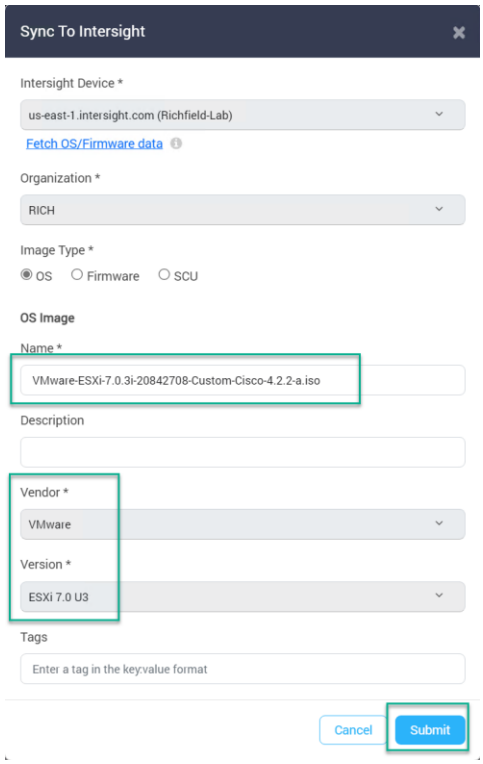


10. Repeat **Step** 13 and 14 for the SCU image. For the Supported Models field, make sure to include any server types that will be deployed with this workflow, i.e., UCSC-C245-M6 for an AMD based Rackmount, UCSX-210C-M6 for the X-9508 based nodes. Other examples could be UCSB-B200-M6 or UCSC-C220-M7, UCSX-420C-M7 for some other examples. When complete click **Submit**.

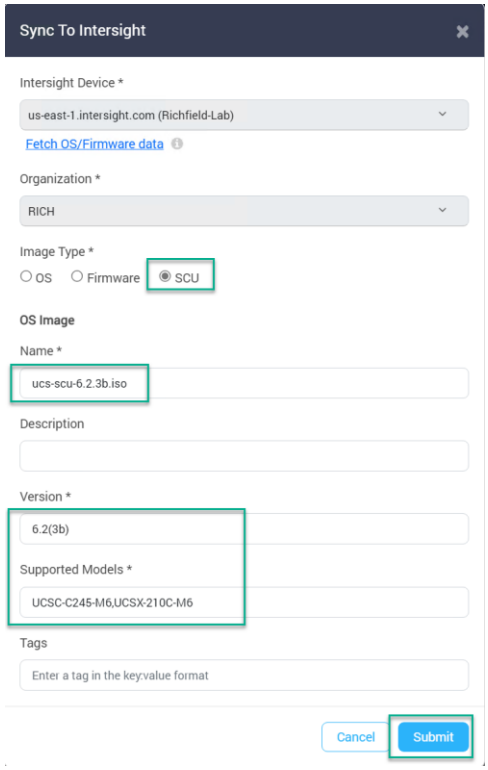

**Note:** If you are including an M7 series server, you will need **ucs-scu-6.3.1b.iso**, 6.2.3b is only for M6 and older.

11. When complete you should see both images in the Software Repository in the transition tool:

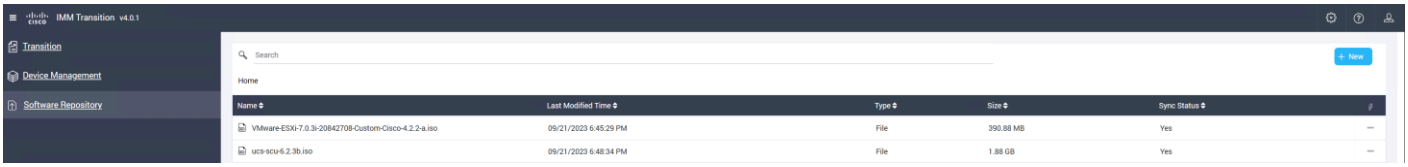


12. From Cisco Intersight you will find the software repositories that were created. Go to **System** > **Software Repository**. Click on the **OS Image Links** tab. By default, the Organization column is missing, let's add the Organization column to make sure the images were added where we need them to be. Click on the settings icon next to All OS Image Links > **Edit**. Add the **Organization** column and click **Submit**.

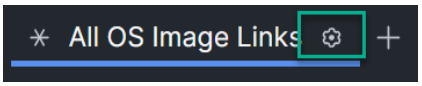


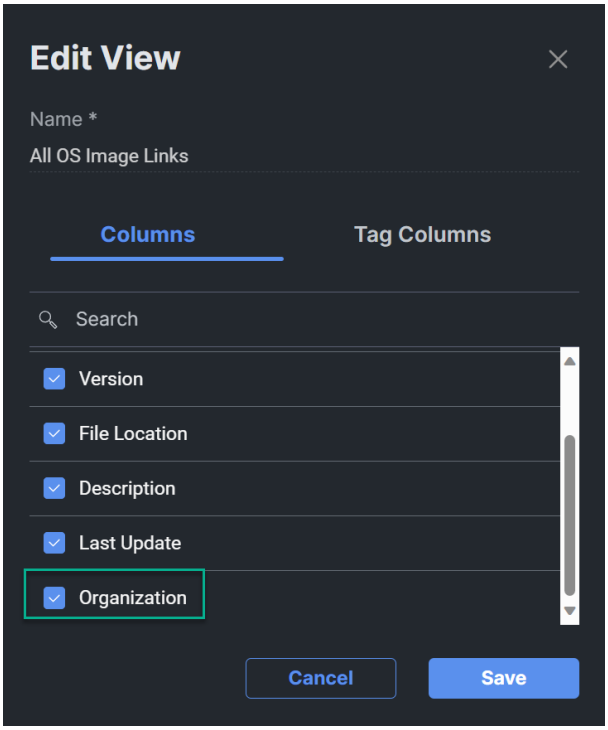


13. Validate that the Image and Organization are as expected. Notice that it automatically created the URL location for the image based on the hostname of the Cisco Intersight Managed Mode Transition Tool hostname.

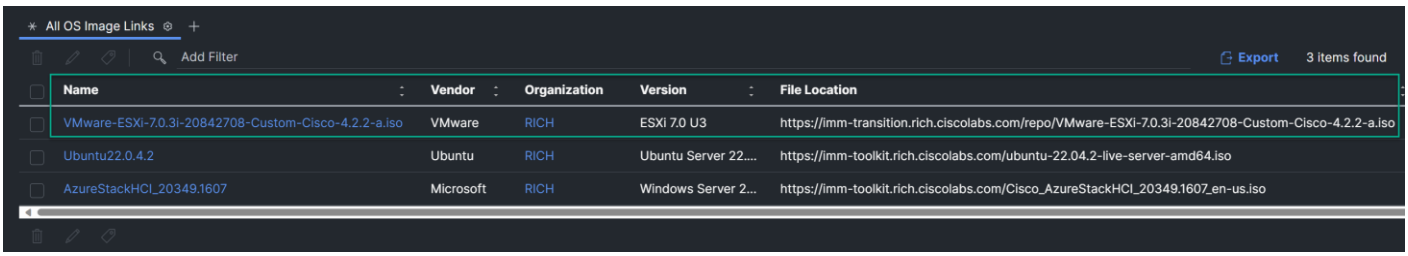


14. Validate the SCU image with the steps 16 and 17 as well.

All SCU Links

Add Filter

| Name | Vers... | Organization | Supported Models | File Location |
|---|---|---|---|---|
| ucs-scu-6.2.3b.iso | 6.2(3b) | RICH | UCSB-B200-M5,UCSC-C240-M5,UCSX-210C-M6 | https://imm-transition.rich.ciscolabs.com/repo/ucs-scu-6.2.3b.iso |

**Note:** The columns shown above are not in the default order. You can drag and re-arrange the columns according to the order you like, as I did above. This is supported in the other views as well.

# Cisco Intersight Cloud Orchestrator Workflow

This chapter contains the following:



## Add Workflow to Cisco Intersight Cloud Orchestrator

### Procedure 1. Import the DeployFlexPodDay0 Workflow to Cisco Intersight Cloud Orchestrator

1. Download the “DeployFlexPodDay0” Workflow from the GitHub Repository: Workflows
2. Go to **System** > **Cloud Orchestrator**:

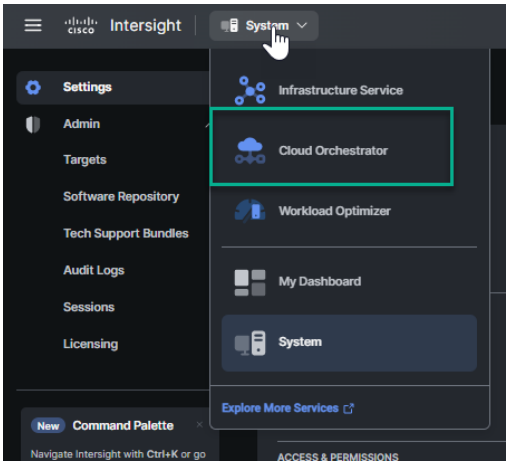

3. Select **Import**, to import the Workflow.


4. Keep the **Organization** as “default”. Click **Browse** to select the workflow. Once selected click **Next**.

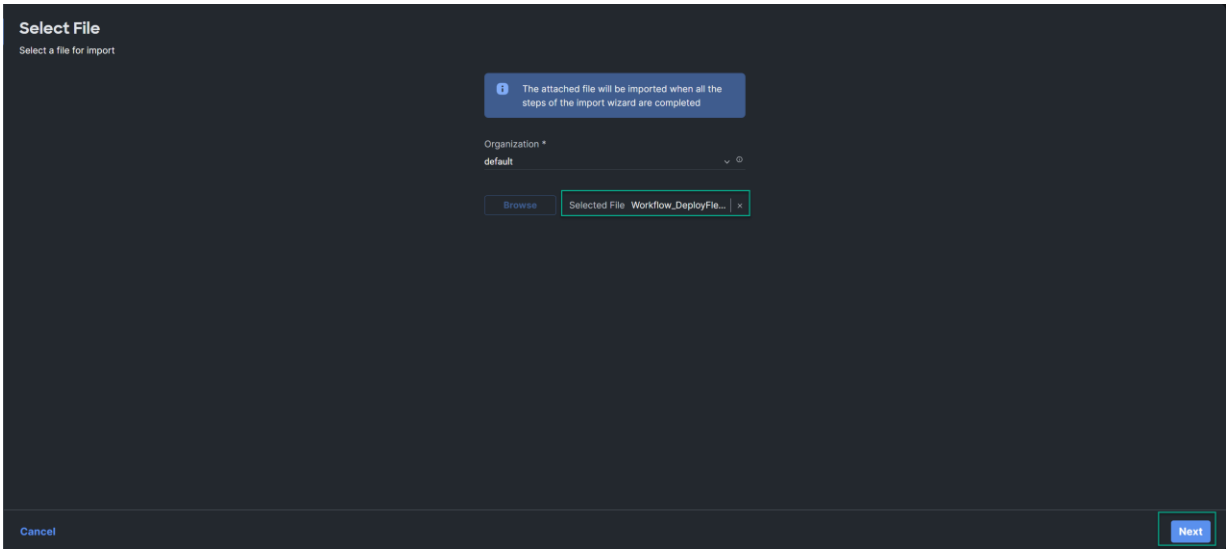


5. It will provide a summary of the **Data Types**, **Tasks**, and **Workflows** that will be imported. Click **Import**.

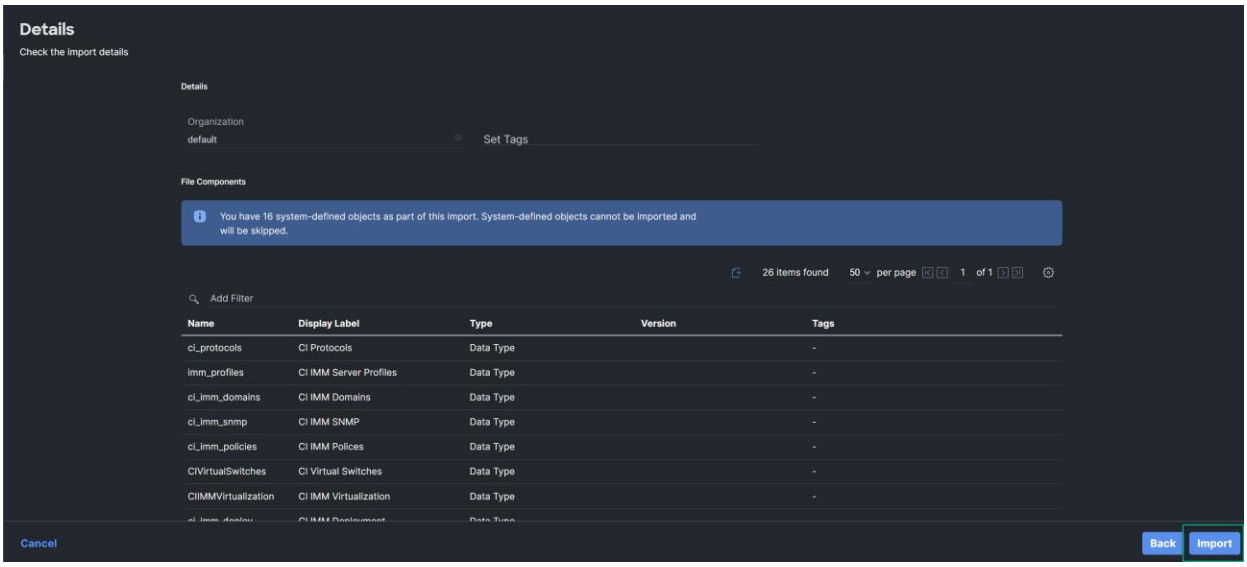


6. Click on the My Workflow tab and you should see the new Workflow at the top of the workflow list.

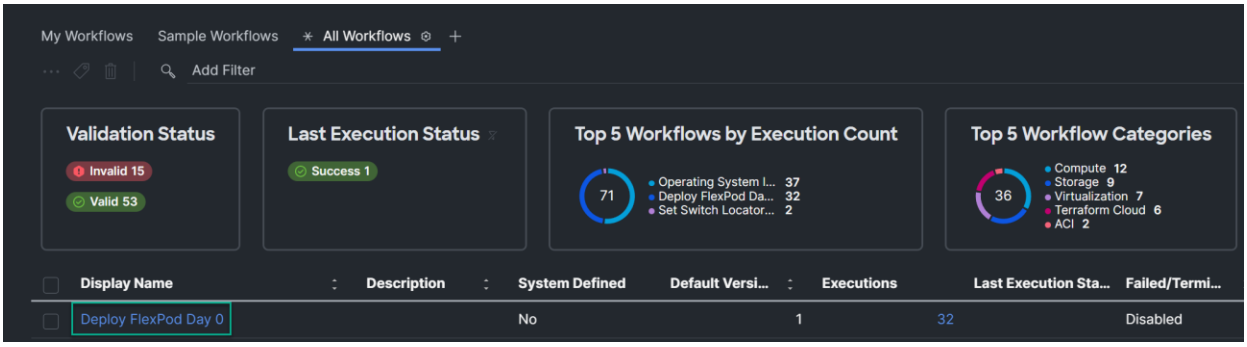

**Procedure 2. Edit API Key ID, API Secret Key, and Python Target.**

1. Click the link for Deploy FlexPod Day 0. This will take you into the designer view. We need to make a few modifications to prepare the workflow for your environment. Click on the General tab and under the Workflow Inputs section click edit (pencil icon) to select the API Key ID that was created earlier for the transition tool as well. Select it from the list of API Key's, click **Select**. Click **Save** on the **Edit Input** menu.

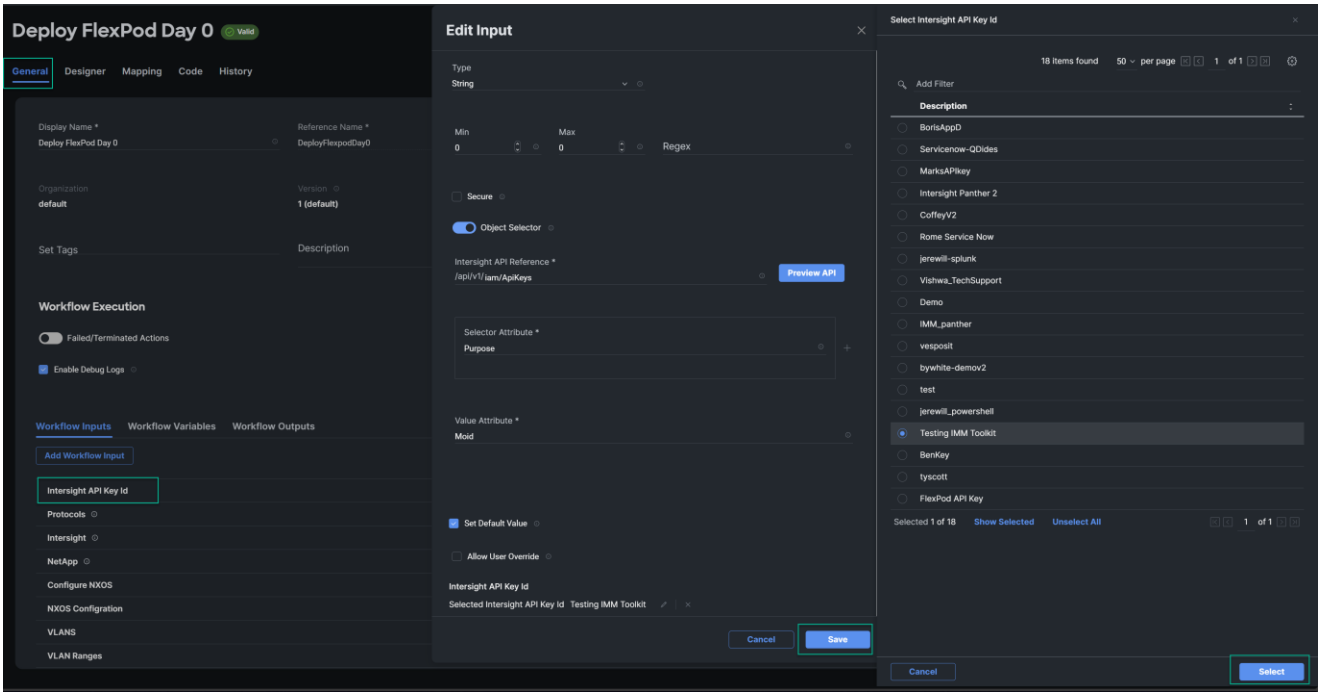


2. Scroll down in the Workflow Inputs menu and edit the **API Secret Key** input. Add the value of your API Secret Key (the text file you downloaded earlier) to this input value, and click Save.

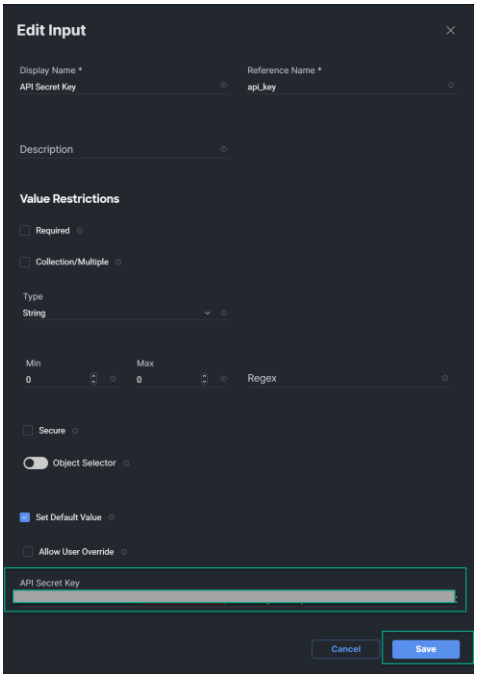


3. Scroll down in the Workflow Inputs menu and you will find **Python Target**. Click edit (pencil icon). At the bottom of the Edit Input screen click edit on the **Python Target**. Select the imm-toolkit target that was

added earlier (The virtual machine that is running the automation tasks), click **Select**, and **Save** on the Edit Input screen.

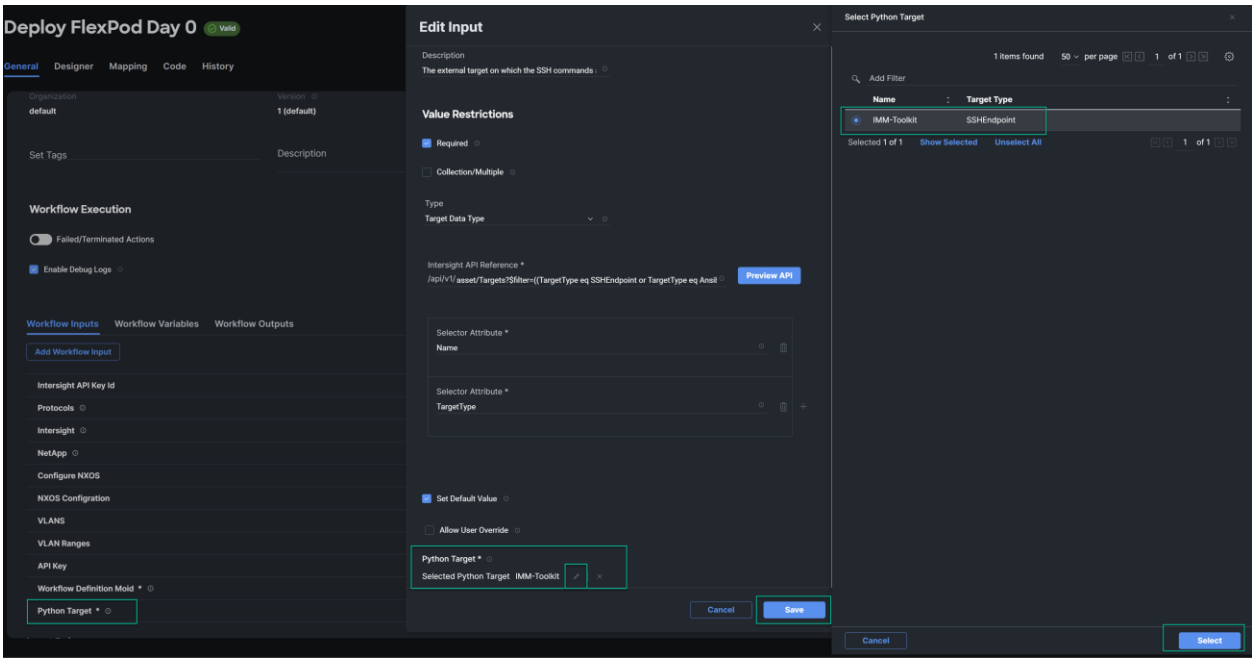


4. Click Save in the upper right, to save the changes to the workflow. If the target is reachable and the correct information was added to the API Key ID and API Secret Key, you should get a successful message.

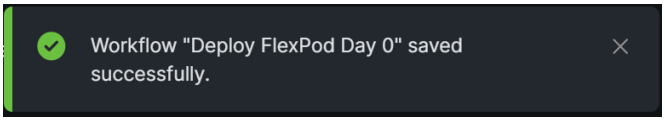


## Begin Workflow

### Procedure 1. Execute the Workflow.

1. Now we are ready to execute the workflow. Click **Execute** in the upper right corner next to **Save**. **Organization** means the organization you should run the workflow in. This organization most likely should remain as default, if the Cisco Intersight Assist Appliance was also added to the "default" **Organization**.

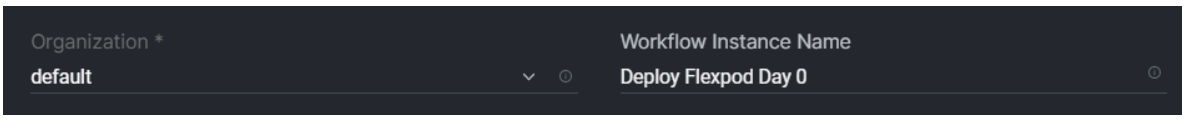


## Protocol Section

### Procedure 1. DHCP, DNS, NTP, and Time zone details.

1. Add the details for the DHCP, DNS, NTP, and Time zone configuration to be used for the FlexPod deployment.
   - **DHCP Servers**. You can add Multiple Servers using the + sign next to the input field. At least one server is required for this. DHCP servers is not today by the script. But please enter at least one for now.
   - **DNS Servers**. Below is a screenshot showing how you can add/remove with the + and trash icons. Add the number of DNS servers for the deployment.

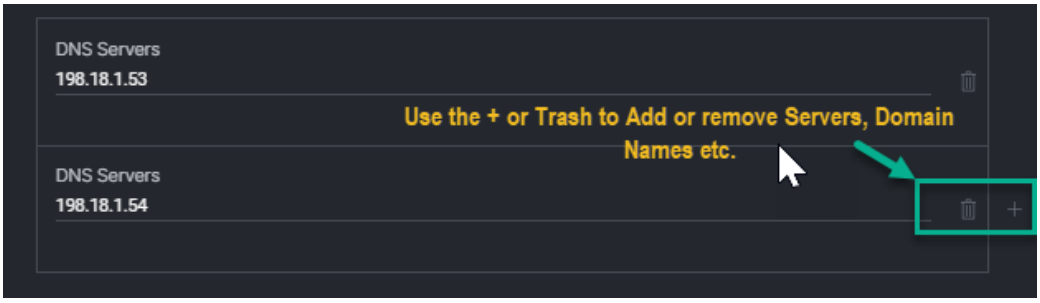


**Note:** For DNS and NTP servers in Cisco Intersight Managed Mode policies, you can use a maximum of two servers. Other appliances like NetApp and NX-OS can use more than two servers. In those cases, if you define more than two you will use them when the application/appliance supports more than two servers.

**Note:** When finished your input should look like this screenshot. For simplicity we reduced this down to a single item for each input field to simplify the view.

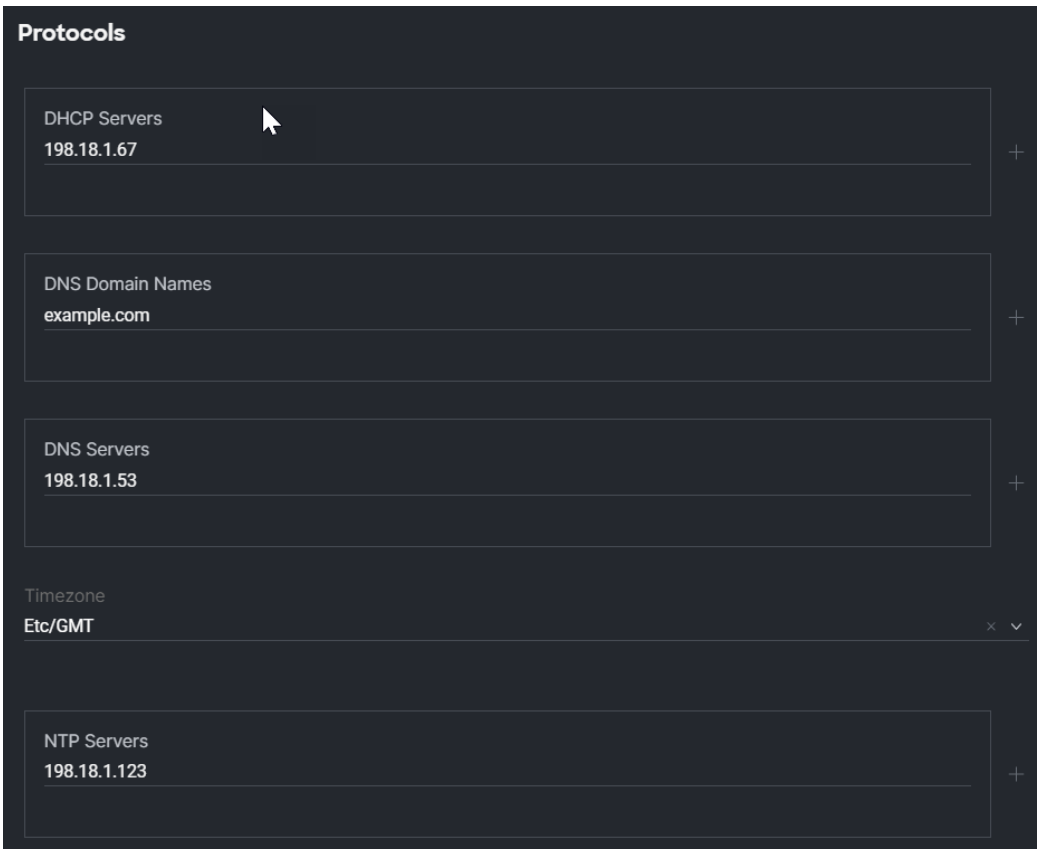

## UCS Domain and Server Profiles

### Procedure 1. Cisco Intersight – Domain Profile

1. Select the Cisco Intersight **Organization**. Here, the **Organization** will be where the pools, policies, and profiles will be created.
2. **Configure Quality of Service (QoS) Priorities**: Flag to determine if the automation should configure the domain profile with the classes/priorities to differentiate the traffic within the domain. By default, the domain will be configured with Jumbo Maximum Transmission Unit (MTU) and Best Effort priority for the Ethernet-based queues. If you want to configure the QoS classes, then select this checkbox. It is important to note that if you configure the QoS class on the domain, it is important that you ensure that QoS is configured across the environment that this domain will be sharing priority flows with.
3. **Discovery Protocol**: You can select cdp, lldp, or cdp-lldp to enable both protocols. This will be assigned to the virtual network interface cards (vNICs) and the virtual distributed virtual switches (vDS).

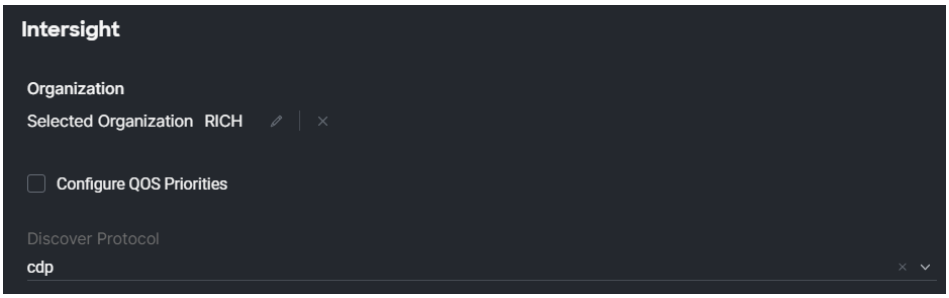


4. **Fabric Interconnects**: Select the Fabric Interconnect pair to apply this domain policy to. If you want to configure multiple domains as part of this deployment, click the plus sign down near the profiles section to add a second domain to the deployment. See example:

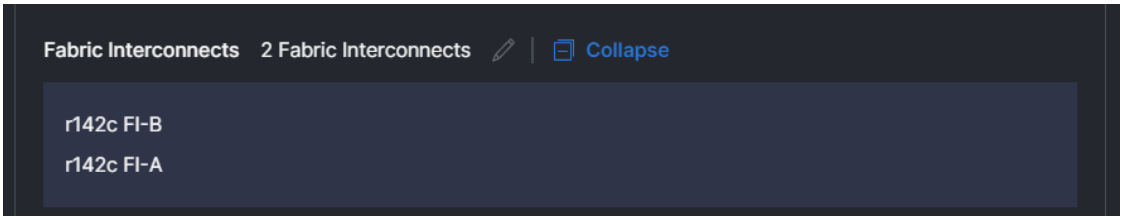


**Note:** The area to add a second domain is down near the VSANs section as shown here:

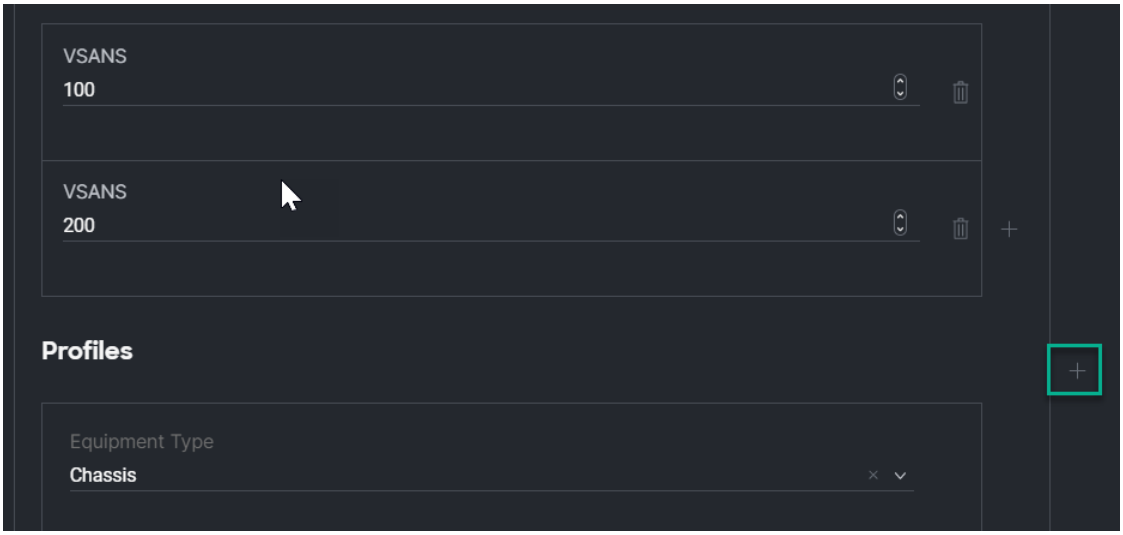

5. **Fibre Channel Switch Mode**: The Fibre Channel mode determines if the domain will connect to a Storage-Area Network (SAN; end-host) or if storage appliances will be directly connected to the domain. The design in this white paper uses directly attached storage (switch), but both modes are supported.

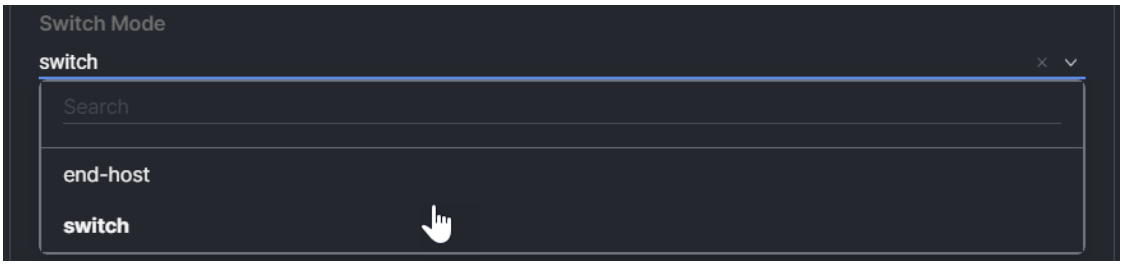


6. **Ethernet Breakout Speed**: This is required only if there are Ethernet ports within the domain that will be configured with breakout. Examples in this paper of the use case for the Cisco UCS C225 Rack Servers are shown using breakout ports. If this is not a part of your design, you can ignore this input.

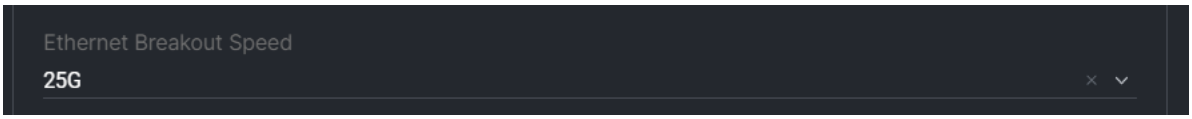


7. **Ethernet Uplink Ports**: The automation will configure these ports as port channels to the upstream network devices. The port definition must meet the format eth1/X or eth1/X/X for breakout.

**Note:** If you need a disjoint Layer 2, meaning that you will have a second path out of the FIs to a second network environment, such as a Demilitarized Zone (DMZ) or management network, the script will split the **Ethernet Uplink Ports** list in half, with half of the ports belonging to the upstream switch and the other half belonging to the disjoint network. Make sure you define the upstream uplinks first and the disjoint uplinks second. This screenshot shows an example with a disjoint network via the second group of ports. even though eth/29 and eth1/30 are lower in order, they were defined second because they go to the disjoint/secondary network uplinks, in this example.

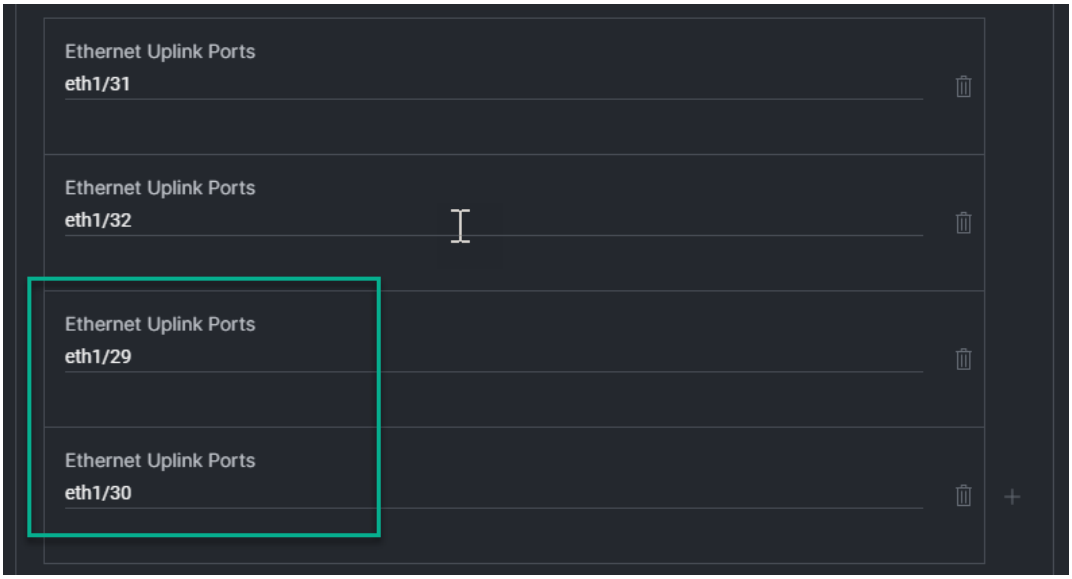


8. **Network Management Port**: This is the port on the Nexus OOB management switch that the Fis are plugged into.

9. **Network Data Ports**: These are the ports on the network switch which correlate to the **Ethernet Uplink Ports** defined in step 7. This will be used for both the port description and configuration of the switch ports if the Nexus switches are defined in the "**NX-OS**" section.

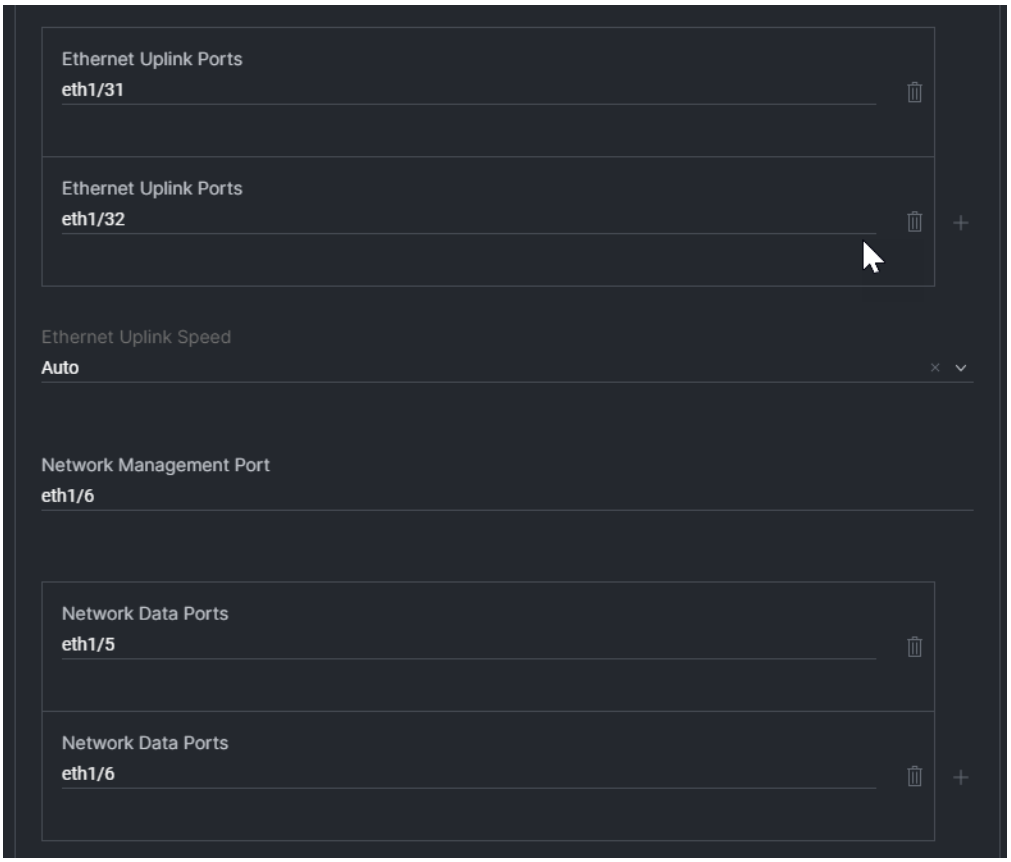


### Procedure 2. Fibre Channel

1. **FCP Uplink Speed**: In the Ethernet Uplinks we use auto-negotiation. Fibre Channel Uplinks do not support Auto with domain profiles, so the speed must be defined.
2. **FCP Ports**: This defines the ports to be configured, either as a Port Channel to a SAN environment or as storage ports if the storage appliances are directly connected to the domain. The format must be either fc1/X or fc1/X/X (case sensitive for the letters). In our example, the 6536 is utilizing port 35 and port 36, which will be configured for FC-Breakout, and we are just showing the first port in the breakout being used in this case. Use the + sign/trash bin on the right to add or remove ports as needed.
3. **VSANs**: This is a list of one or two VSANs based on your deployment. If your environment uses the same VSAN for both fabrics, this list has one VSAN. If your environment uses different VSANs for each fabric, you should define two VSANs, as shown in previous screenshot. Note that the VSAN ID will also be defined as the FCoE VLAN. Make sure that the VSAN doesn't overlap with a VLAN in the VLAN section because of the FCoE VLAN.


### Procedure 3. Profiles

You can create multiple profile blocks within the domain section for each chassis and rack-mount server you add. Use the plus sign on the far right to add multiple profile blocks to the domain. This section requires one profile block per chassis and rack server, but not for individual nodes within a chassis. Define only the profile section at the chassis or rack-server level.

1. **Equipment Type** and **Identifier**: The Identifier is used by the domain port policies to pre-assign the identities to the ports in the port policy. It ensures that the chassis and rack mounts are in the order you want them to be in the domain after discovery, with chassis 1 being chassis-1 and so forth. Note that you can use the same identifier for both chassis and rack mounts, meaning it can include chassis-1 and rack mount-1. The identifier needs to be unique to the equipment type per domain.

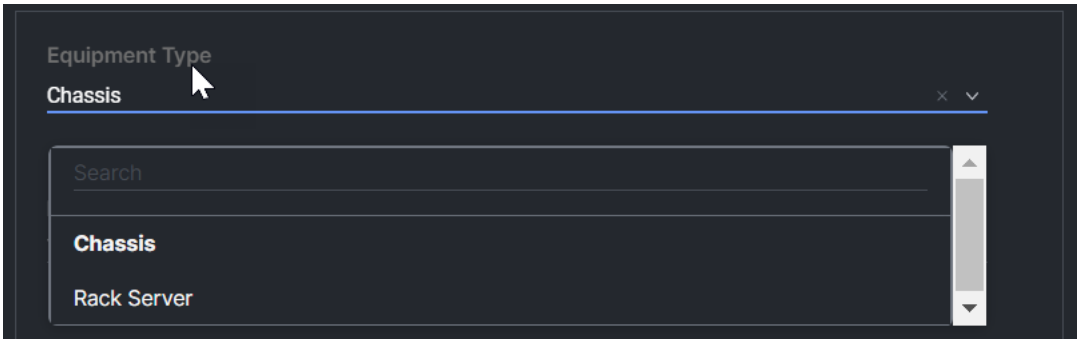


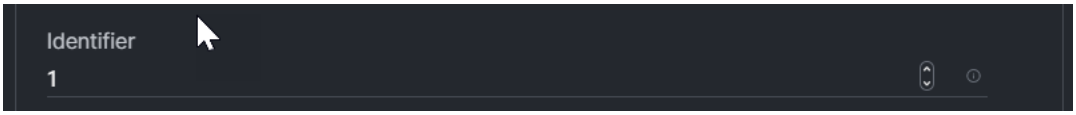


2. **Domain Ports**: The ports on the FI for the given profile block; in other words, these ports are the ones that the chassis or rack server are connected to. This correlates to the assigned identifier for configuring the device ID before discovery.

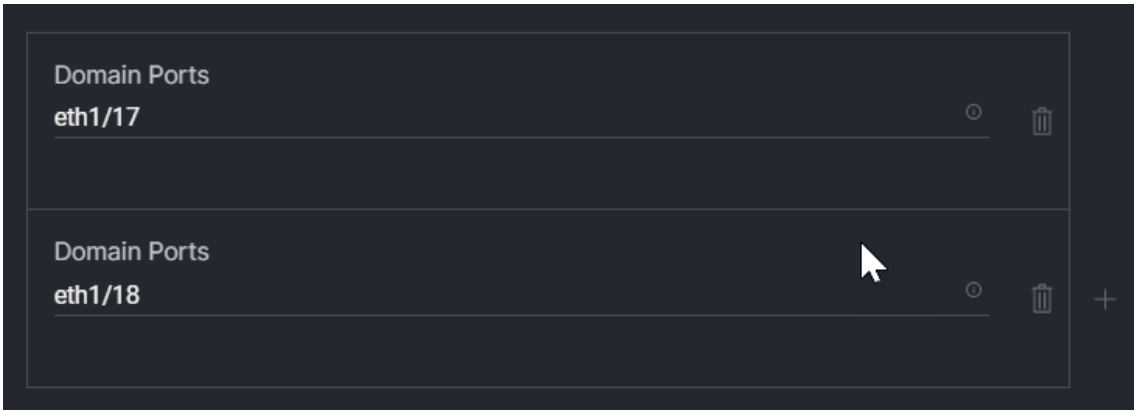


3. **Inband Start IP**, **Profile Starting Name** and **Suffix Digits**: We will focus on a chassis setup because a rack-server deployment is simply a one-to-one mapping of the IP and profile starting name with the suffix being ignored. In the case of a chassis, though; these options are used to determine the name of each node in the chassis. The profile starting name will be the hostname/profile name of slot 1. The script expects that you can use all 8 slots, so names and IPs should account for 8 nodes for each chassis even if there isn't a plan to use 8 nodes in each chassis. The script will also use the profile name as the hostname during the operating-system installation, so make sure to enter this information in DNS before beginning the deployment, or the addition of the host to vCenter will fail.

**Note:** The suffix digits are to allow for flexibility in the hostname configuration. For example, using the screenshot below, the last 2 digits are what is unique per host, that is, r142c-esx01, r142c-esx02, up to r142c-esx99. But if you wanted the hostnames to be r142c-esx-0001, then the suffix digits would be configured as 4 and the Profile Staring Name would be entered as r142c-esx-. Note that the automation will increment the last digit for each server so if you have more than 99 servers you are deploying the suffix should be more than two. But the profile start name will be restarted for each chassis that is defined as well.

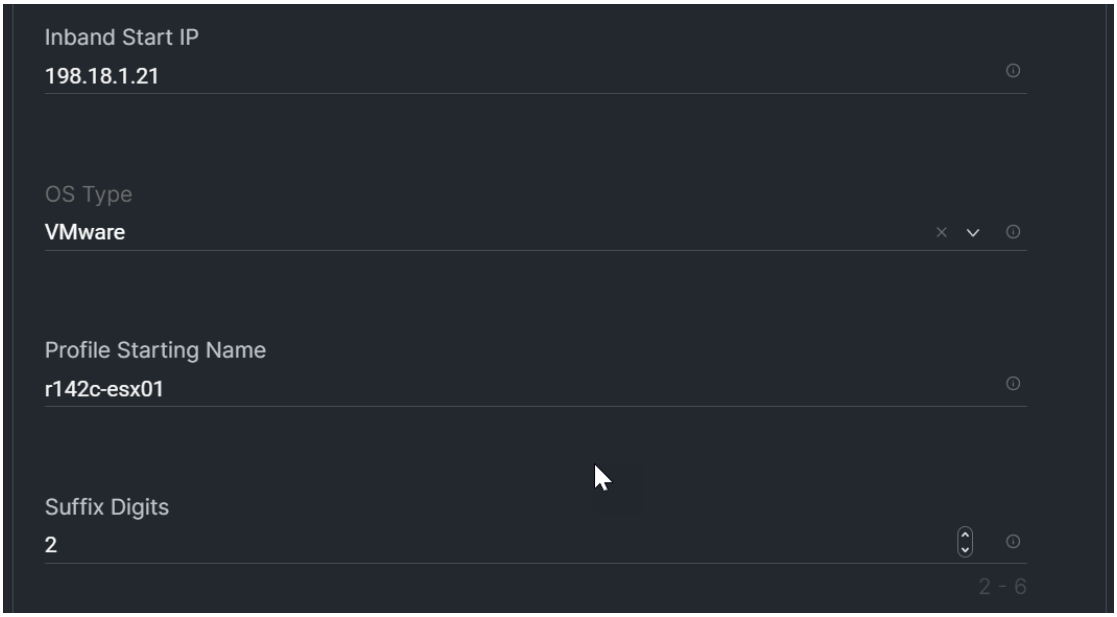


**Note:** In the case of a VMware deployment, the automation will configure vmk0 with the Inband Starting IP. For the Gateway/Subnet mask etc., that will be added later within the "**VLAN**" section.

4. **Server Firmware**: Select the firmware to be assigned to the Server Profile Template, which in turn will be applied to each server profile from the template. Select a version for the blades and a version for the rack-mount servers. At the time of the writing of this white paper the recommended version is 5.1(0.230054) for the blades and 4.2(3d) for the Rackmounts.


### Procedure 4. Pools

1. **Pools Prefix**: This will be the prefix added to the pools during creation. It is two digits and is a required parameter. Table 5 shows an example of how this prefix is applied to the various pools. The automation will create Media Access Control (MAC) pools, World Wide Node Name/World Wide Port Name (WWNN/WWPN) pools, and Universal Unique Identifier (UUID) pools using the defined prefix value.

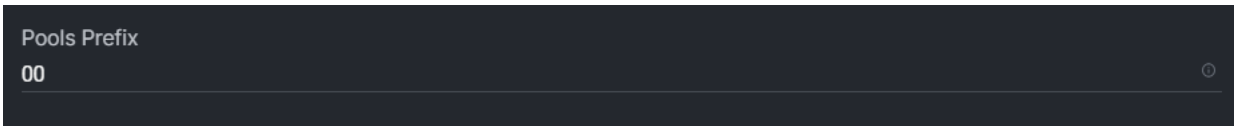


**Table 5.** Pools created by the wizard

| Pool type | Name | Prefix/suffix | Pool starting address | Size |
|---|---|---|---|---|
| IQN | iqn | Iqn.1984-12.com.cisco | 0 and ucs-host suffix | 1024 |
| MAC | mgmt-a | | 00:25:B5:{prefix}:10:00 | 1024 |
| MAC | mgmt-b | | 00:25:B5:{prefix}:20:00 | 1024 |
| UUID | uuid | 000025B5-0000-0000 | {prefix}00-000000000000 | 1024 |
| WWNN | wwnn | 198.18.1.96 | 20:00:00:25:B5:{prefix}:00:00 | 1024 |
| WWPN | wwpn-a | | 20:00:00:25:B5:{prefix}:A0:00 | 1024 |
| WWPN | wwpn-b | | 20:00:00:25:B5:{prefix}:B0:00 | 1024 |

**Note:** The number of MAC pools created is determined by the number of virtual switches that are defined in the "**Virtualization**" section. Two MAC Pools for each virtual switch, Fabric A and Fabric B respectively. The name of the MAC pools is determined based on the name of the virtual switch. In the special case of vSwitch0, an alternate name field is available, so pools and policies aren't named vSwitch0 that you have no control over the name.

**Note:** The WWNN/WWPN pools are created if any of the volumes are defined using "fcp" or "nvme-fc" for the mount "protocol".

**Note:** The IQN pool is configured only if the iSCSI VLANs are defined in the VLAN section and a volume with the "iscsi" mount "protocol" is defined.

### Procedure 5. Policies

1. **Boot Volume**: The options are iSCSI, M2, or SAN. This volume is assigned to install the operating system. If you are configuring SAN (previous screenshots), make sure to add a volume of type "boot" in the "Volumes" section.

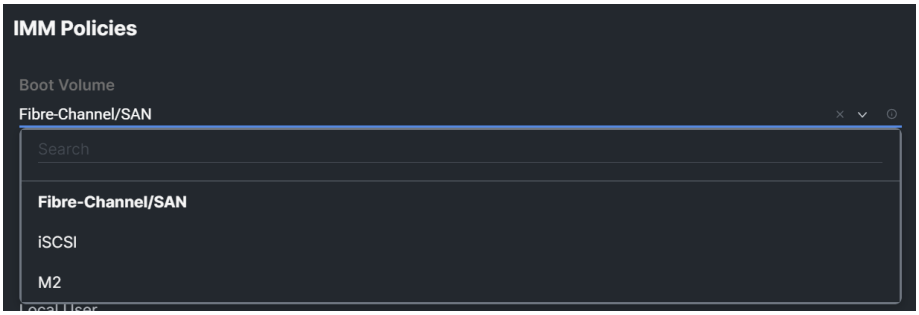


2. **Policies Prefix**: In this optional field you can assign a prefix to the name of all the policies that you want to create. The wizard does not ask for names of individual policies because that would defeat the purpose of a quick-start wizard. Optionally you can add a policy prefix to the default names that the wizard will create for each policy; that is, for the ntp policy if the policy prefix is *"MINE-"*; then the new name for the ntp policy would be *"MINE-ntp"*. The prefix convention is {prefix}<name>. Note that with Cisco Intersight you can always change the names of the policies if you want them to be different. It is an improvement over Cisco UCS Manager, which didn't allow for names to be changed.

3. **Local User** and **Local User Password**: The unique attributes to each customer environment are the only policy questions in this wizard. That is Usernames, passwords, and the like. So, Local User, Simple Network Management Protocol (SNMP), and Syslog are presented as questions to answer. The Local User Policy is used to log in to KVM directly, if Cisco Intersight SaaS was unavailable, as an example.

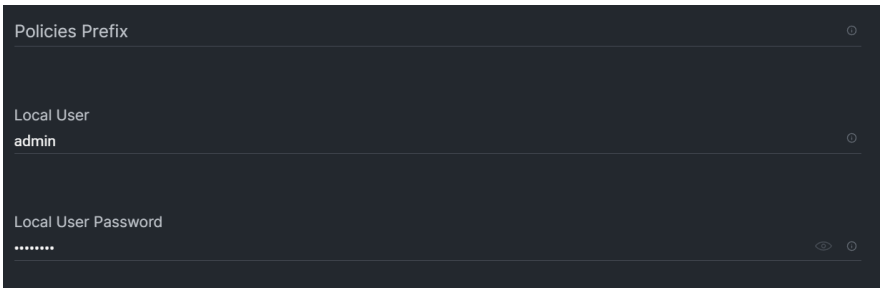


4. SNMP configuration.
   - **Contact**: Enter a value for the system contact that will be applied to the domain.
   - **Location**: Enter a value for the system location that will be applied to the domain.
   - **Username**: SNMPv3 user to be created with the SNMP Policy. The SNMP policy does not permit the use of admin as the username; the policy creation will fail if you use admin.
   - **Authentication and Privacy Password**: The automation will configuration a SNMP user with "AuthPriv" as the security level.
   - **SNMP Servers**: Add as many SNMP trap servers as you need.

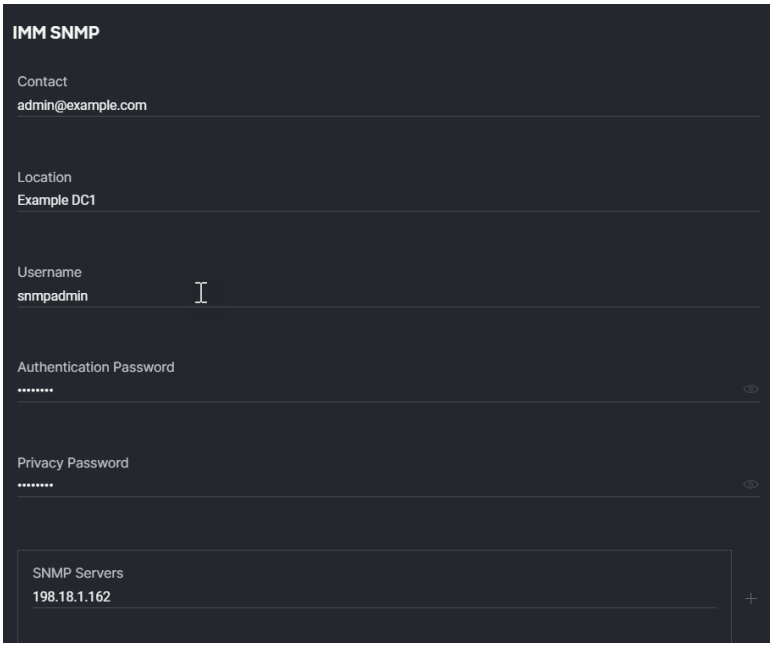


5. **Syslog Servers**: You can define two syslog servers. If you add more, the wizard will ignore them. The syslog policy only supports two servers.

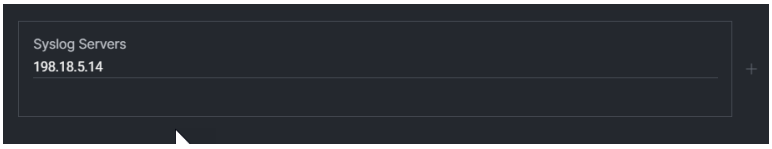


## Virtualization Environment

**Procedure 1. Define the configuration parameters for the virtualization environment**

1. **Datacenter**: The name of the datacenter defined/to define in vCenter.

**Note:** The wizard does not ask for clusters. It will add hosts to vCenter clusters based on the model of the server. Clusters are named based on the server model name. Please rename the cluster after deployment according to your deployment.

2. **License Type**: The script can deploy using either Standard or Enterprise licensing. If selecting Standard, make sure all the virtual switches are also Standard because Distributed Virtual Switches is an Enterprise feature.
3. **vCenter FQDN**: The hostname of the vCenter is called vCenter FQDN. With the ESXi hosts, you can use the profile name + the first domain in the DNS domains to register the hosts to the vCenter. Make sure to configure DNS before starting the OS deployment portion of this wizard.
4. **Type**: At the time of this writing only VMware has been tested, but OpenShift will be added when testing is completed.
5. **vCenter Username**: This username offers administrative privileges to configure hosts, clusters, datacenters, and virtual networking.
6. **vCenter and ESXi Passwords**: These passwords are required for the OS Installation and OS Configuration sections.

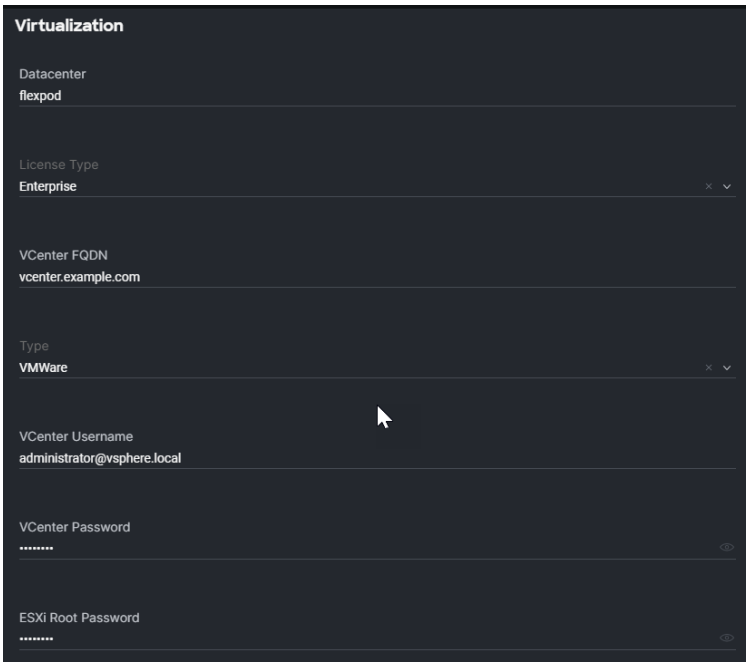


### Procedure 2. Virtual switches

1. **Data Types**: The type of traffic the virtual switch will carry. You can segment each data type on a per-virtual switch basis or combine multiple data types on the same virtual switch. In the case of VMware, management would be for the vmk0 interface. If you add the migration data type to the same virtual switch, then vMotion will also be added to the vmk0 interface. It is a recommended practice to separate these data types onto different virtual switches to allow the classification of these traffic types. Classification helps ensure that storage traffic is always preferred over guests' traffic, guests' traffic is always preferred over vMotion, and so forth. But you can customize this section based on your standards. Remember that two vNICs will be created for each virtual switch. So this section defines how many vNICs will be created.

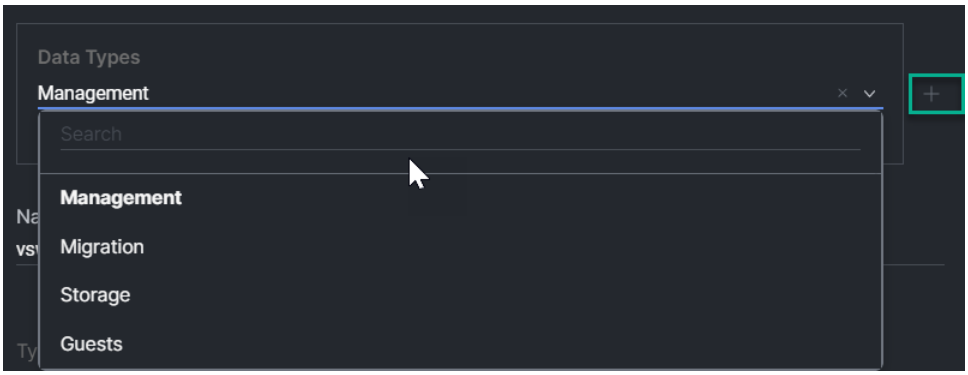


2. **Name**: We recommend that the name of the first virtual switch remain vSwitch0, with the type remaining as a standard vSwitch. The workflow can support changing this, but we don't recommend it. Any subsequent names can be according to your standards.

3. **Type**: If you choose Distributed vSwitch, make sure the license type is set to Enterprise for the hosts (refer to Table 6).

4. **Alternative Name**: This term is used only for vSwitch0. It provides an alternative name for the policies that you create for vSwitch0, because you cannot rename vSwitch0.

**Note:** Generally, we recommend that the Management traffic be left on vSwitch0 as a Standard Switch, and everything else can be added to one or more distributed switches.

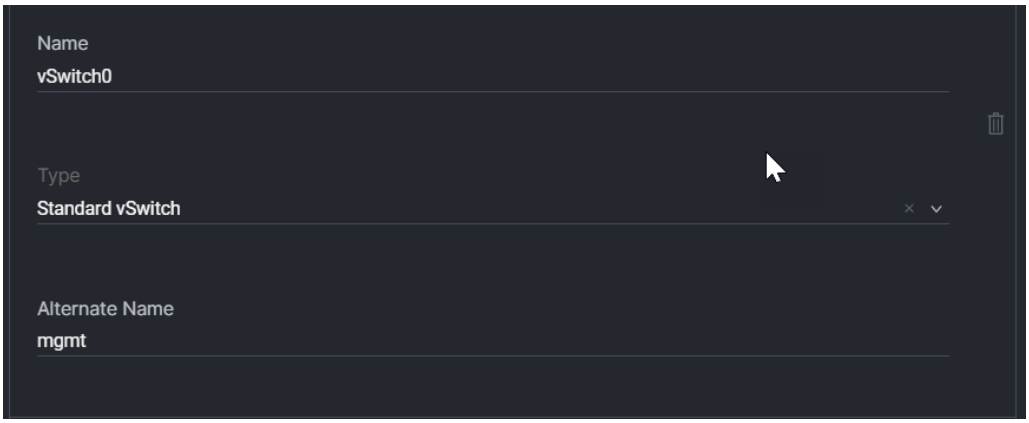


**Table 6.** Distributed virtual switches design uses

| Virtual switch | Type | Data type | Traffic |
|---|---|---|---|
| vSwitch0 | Standard vSwitch | Management | Management |
| migration | Distributed vSwitch | Migration | vMotion |
| Data | Distributed vSwitch | Storage | NFS, iSCSI, NVMe-TCP |
| Guests | Distributed vSwitch | Guests | Guest ddddddraffic |

## NetApp Storage Array Configuration

### Procedure 1. Base settings

1. **Username**: administrative user for API/Shell configuration. This must be an existing user account with administrative privileges.
2. **Password**: administrator password.
3. **Login Banner**: The banner presented to the user for Cluster Login

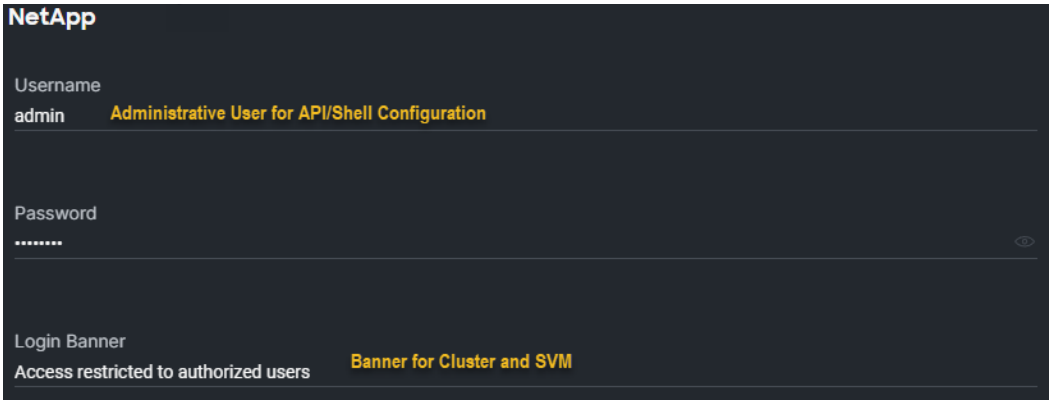

### Procedure 2. AutoSupport

1. **Mail Hosts**: The list of mail relay hosts for AutoSupport notifications.
2. **Proxy URL**: If you need a proxy, define the proxy URL here. When using authentication with the proxy, the format of the URL is http://<username>:<password>@<hostname>:<port>. i.e., http://proxy_user:password@proxy.example.com:80 or http://proxy.example.com:80 without authentication.
3. **From Address**: The email address AutoSupport will use as the source when sending notifications.
4. **To Addresses**: The list of email addresses to send AutoSupport email notifications to.

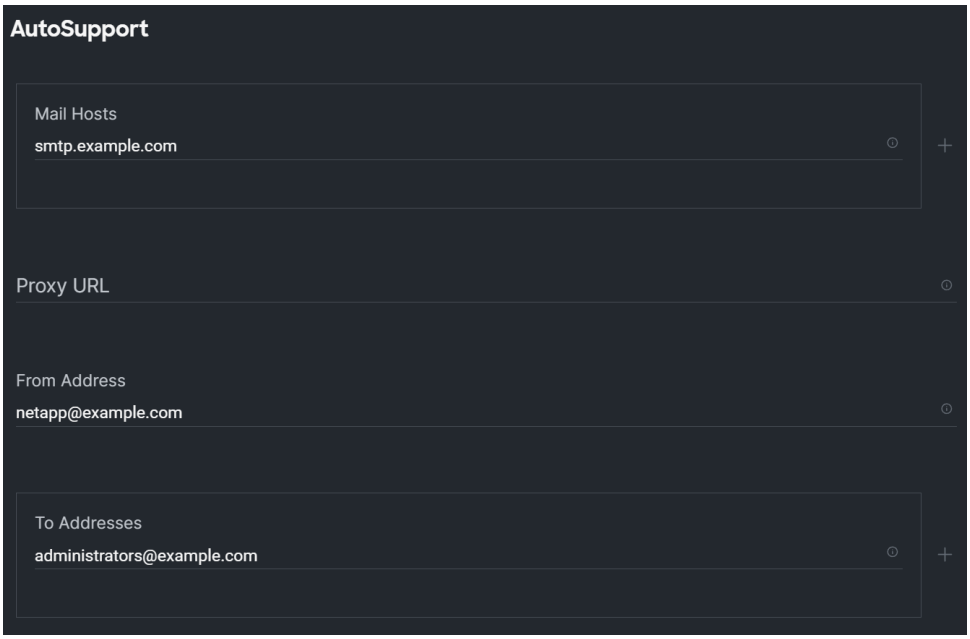


### Procedure 3. NetApp SNMP

1. NetApp SNMP configuration.
   - **Contact**: Enter a value for the system contact that will be applied to the array.
   - **Location**: Enter a value for the system location that will be applied to the array.
   - **Username**: SNMPv3 user to be created with the SNMP Policy. The SNMP policy does not permit the use of admin as the username; the policy creation will fail if you use admin.
   - **Authentication and Privacy Password**: The automation will configuration a SNMP user with “AuthPriv” as the security level.
   - **Trap Server**: If the server is in DNS, make sure to configure the FQDN of the trap server. When you add the trap server to the API, the appliance does a reverse lookup and adds the hostname from DNS if configured, so make sure it matches here. If it is not configured properly, if the script needs to run for a second time because of a failure all subsequent re-runs will fail.

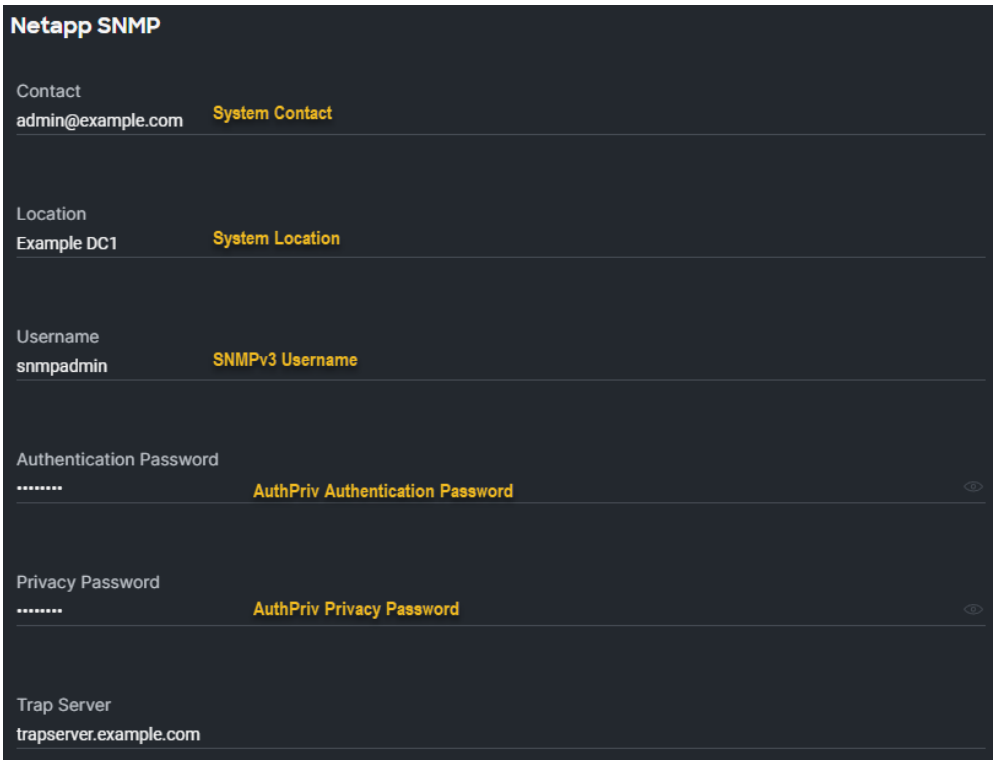


### Procedure 4. Storage Node Cluster(s)

1. **Cluster Name**: This name is for the NetApp Storage Array.
2. **Login Banner**: This banner is for the Cluster Login Screen.
3. **Node01 and Node02**: These names are the hostnames for the two cluster nodes.

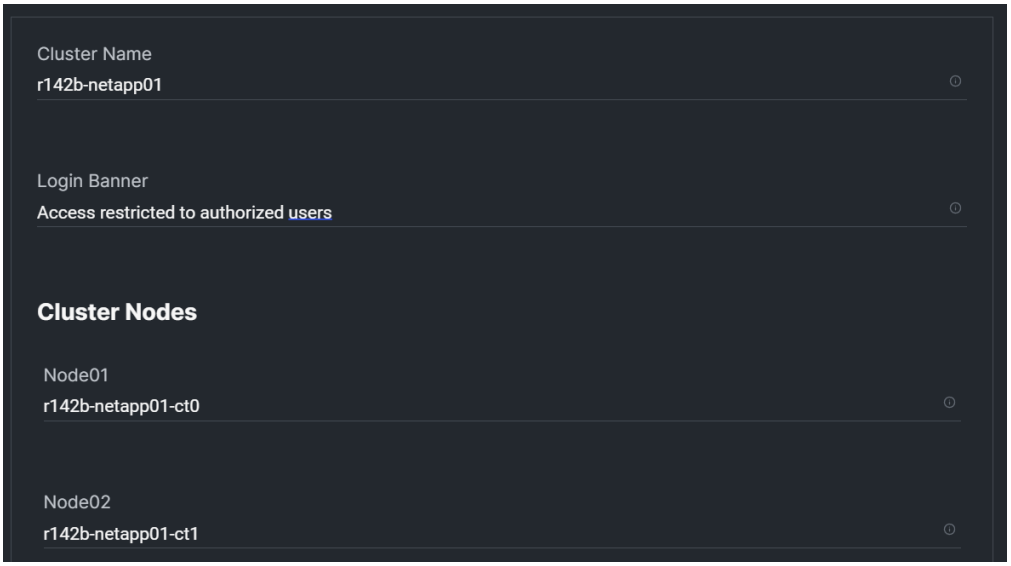


**Note:** VERY IMPORTANT. Node01 DNS Name is used for connecting to the appliance for the automation, because the cluster configuration will still not be complete. Make sure you enter this name into DNS before running the automation.

### Procedure 5. Storage Appliance Data Interfaces

1. **Data Speed**: Select the speed to apply to the ports.
2. **Data Ports**: This is the name of the data ports on the array. Because Node01 and Node02 should be a mirror image of each other. you need only to put in the names for one node.
3. **Network Data Ports**: These are the ports in the network switch that the array is connected to.
4. **Network Management Port**: This is the out-of-band management port for the array e0M interface.

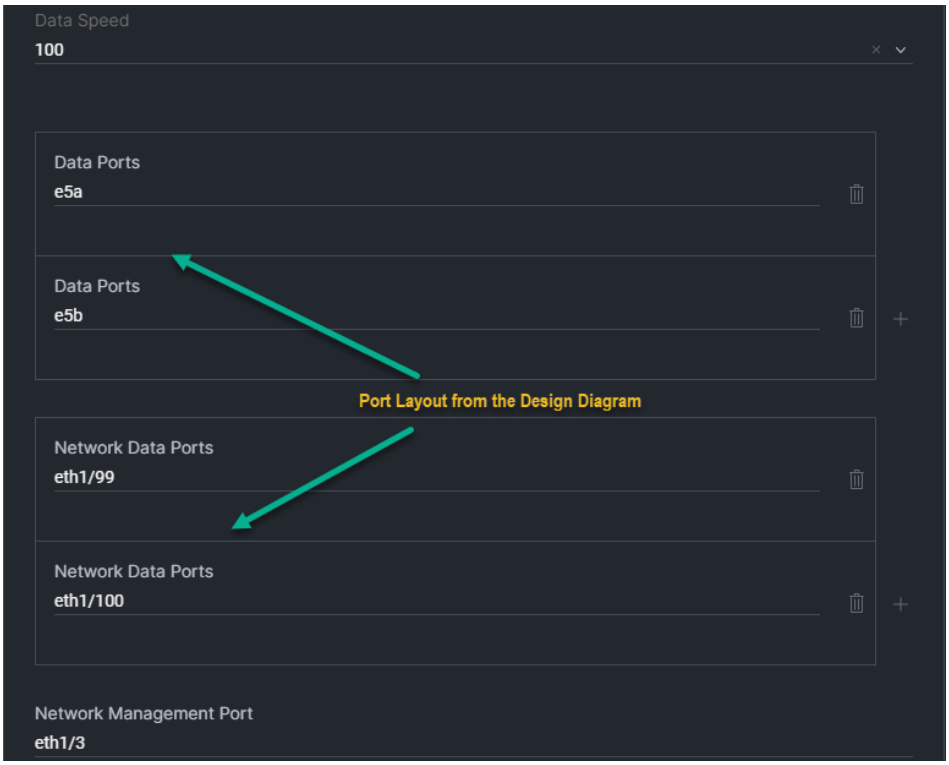


5. **FCP Ports**: This is the name of the Fibre Channel Ports on the array. Because Node01 and Node02 are a mirror image of each other, you need only to put in the names for one node.
6. **FCP Speed**: You can select the speed to apply to the Fibre Channel ports.

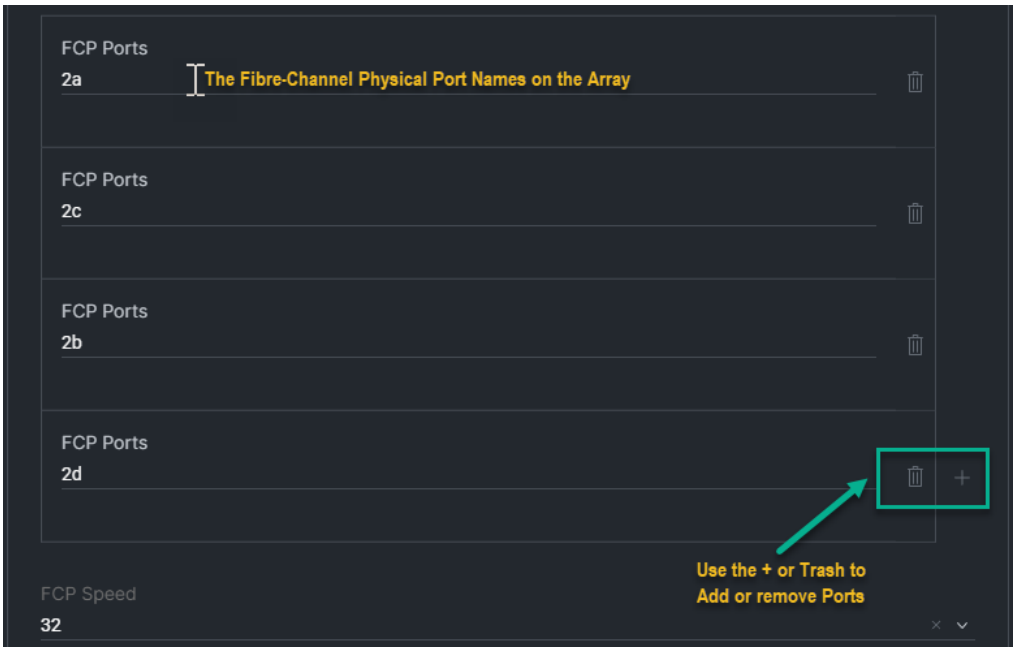


**Note:** In the example above four FCP Ports were defined. When NVMe over FC is in use, the list will be split in half with the first half of the list used for FCP and the second half of ports used for NVMe over FC. Make sure to list the ports in the order in which you want them to be consumed.

### Procedure 6. SVM

1. **SVM**: Storage Virtual Machine name.
1. **SVM Login Banner**: This is the login banner for the SVM login screen.

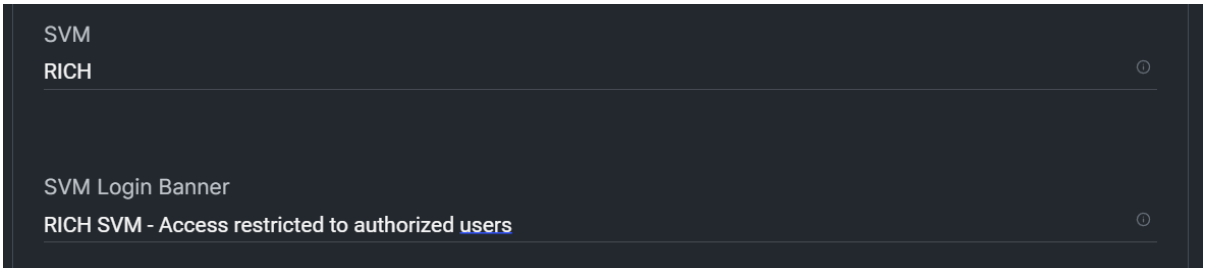


### Procedure 7. Volumes

1. **Name**: The name of the volume.
2. **OS Type**: The operating system of the device consuming the volume. Audit log uses the “netapp” type to signify that the local appliance uses it.
3. **Protocol**: This field defines the protocol used to mount the volume, logical unit number (LUN), and datastore. (Refer to Table 7) for explanations on the protocols.
4. **Size**: The size in GB for the volume.
5. **Volume Type**: The automation uses this setting to determine the configuration steps to apply to the volume such as vcls will be the volume used to assign vSphere Clustering Service (vCLS) virtual machines within a cluster, swap will be used to assign virtual machine swap file locations, etc. (Refer to Table 8.)

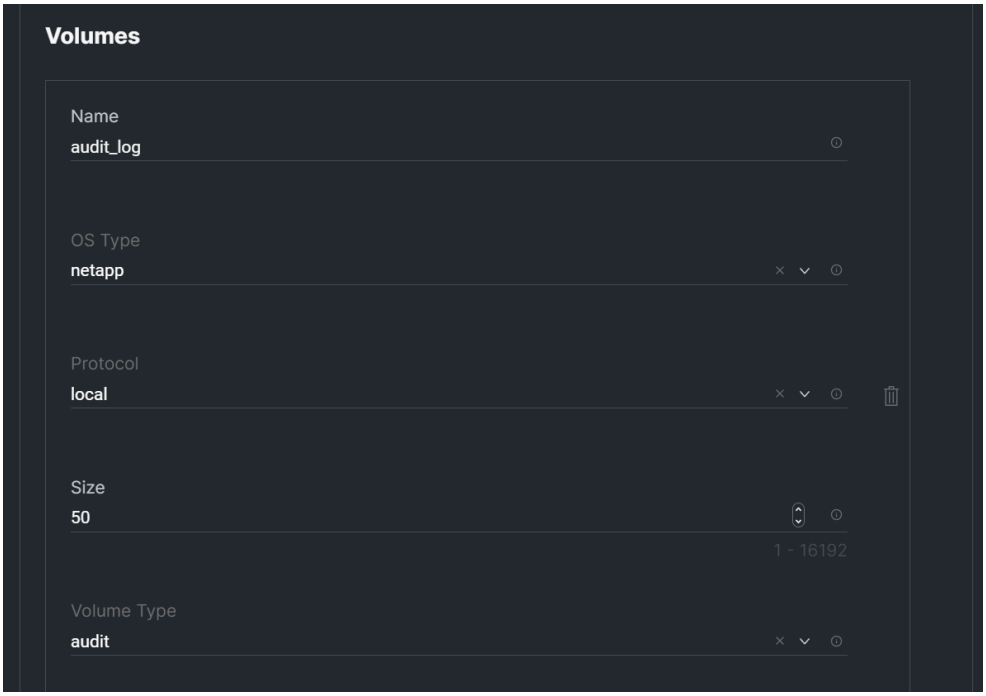


**Table 7.** Protocol options for mount the volume, lun, and datastore

| Protocol | Description |
|---|---|
| local | The volume will be for local data management on the storage array. |
| fcp | The datastore will be mounted with the Fibre Channel Protocol and a LUN will be created. |
| iscsi | The datastore will be mounted with iSCSI and a LUN will be created. |
| nfs | The datastore will be mounted using NFS. |
| nvme-fc | The datastore will be mounted using NVMe over Fibre Channel. |
| nvme-tcp | The datastore will be mounted using NVMe over TCP. |

**Table 8.** Recommended volumes the design uses

| Name | OS type | Protocol | Size | Volume type |
|---|---|---|---|---|
| audit_log | netapp | local | 50 | audit |
| esxi_boot* | vmware | fcp | 1024 | boot |
| infra_datastore | vmware | nfs | 1024 | data |
| infra_swap | vmware | nfs | 256 | swap |
| vCLS | vmware | nfs | 128 | vcls |
| nvme_datastore | vmware | nvme-tcp | 1024 | nvme |

**Notes:**

- The “audit_log” volume will be used to store SVM audit logs.
- The “esxi_boot” is a special datastore in that it hosts the boot volumes for all the ESXi hosts when the boot_volume is set to SAN in the “IMM Policies” section. Each boot volume is created as a 128GB drive and FC Zone policies will be defined for each ESXi host that will use a boot drive.
- The “infra_datastore” contains the virtual machines that manage the environment, including NetApp ONTAP tools and Active IQ Unified Manager. Note that vCenter, Cisco Intersight Assist, and the IMM Toolkit were all deployed prior to the FlexPod cluster deployment. You can migrate these applications to this datastore after the installation is complete. This datastore is mounted using NFS.
- The “infra_swap” volume is configured in vCenter as the virtual-machine swap file location to help improve the performance of the virtual machines. This datastore id mounted using NFS.
- The “vCLS” volume hostd VMware Cluster Lifecycle Services virtual machines. This datastore is mounted using NFS.
- The “nvme_datastore” is mounted using NVMe-over-TCP, and that is why it has the volume type of NVMe instead of data. The NVMe configuration is also different in that instead of configuring a LUN in the NetApp array with storage initiators, the NVMe is configured using the NVMe subsystem.

## Nexus Switch Configuration

**Procedure 1. Configuration of the Nexus management and network switches**

1. **Configure NXOS**: This flag determines whether you need to configure the network switches.
2. **Username**: The administrative user for switch configuration.
3. **Password**: The administrator password.

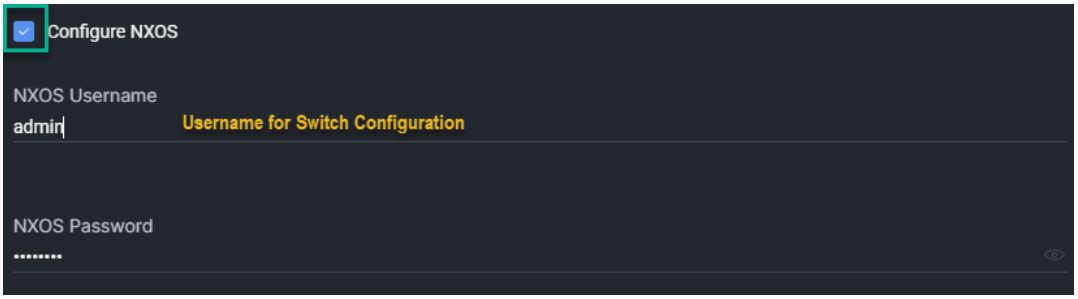


**Procedure 2. NXOS switches**

1. **Hostname**: Hostname is the short name for the switch. It will be combined with the first DNS domain for the FQDN of the switch, and the hostname is used for port descriptions, and switch configuration. The FQDN is used to connect to the switch. Make sure to add the FQDN to DNS before executing the workflow.
2. **Breakout Speed**: This entry is required only if you use breakout ports on the switch, much like the Intersight Managed Mode domain.
3. **Switch Type**: The type is either network or ooband. The script can configure the upstream switch pair as well as the ooband switch for the deployment. This entry tells the script what type of switch you are defining.

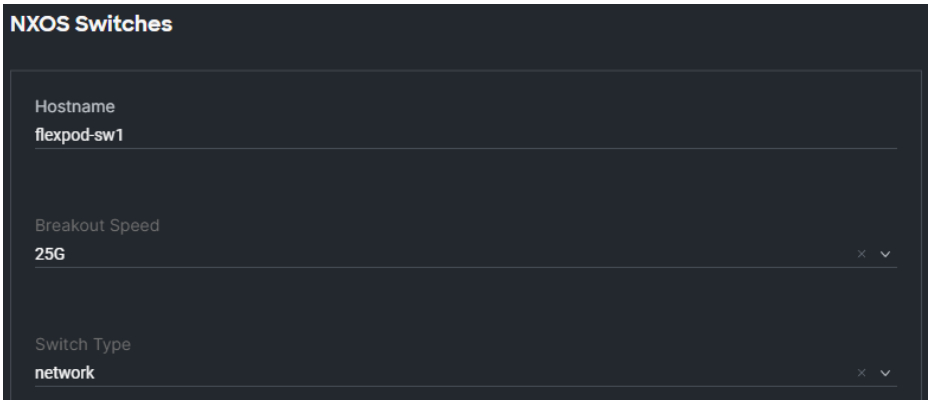


**Note:** When adding the switches, make sure that you add them as pairs, for Virtual Port Channel (VPC) and port configurations to match between the two. For Out-of-Band management switches, you can use a single switch.

4. **Configure VPC**: This flag determines if the script should configure VPC for a switch pair. Make sure you have defined both switches in the pair, in order.
5. **VPC domain ID**: This field is the domain ID for a pair of VPC switches. Make sure it matches for the switch pair.
6. **VPC Keepalive IP CIDR**: IP assigned to the switch for the keepalive port. The prefix is needed here in the event you do not use the mgmt0 for keepalive when it needs to configure the alternate ports.
7. **VPC Keepalive Ports**: This can be a single port or a list of ports for the keepalive ports.
8. **VPC Peer Ports**: The interfaces directly connected between the switches to form the VPC relationship.

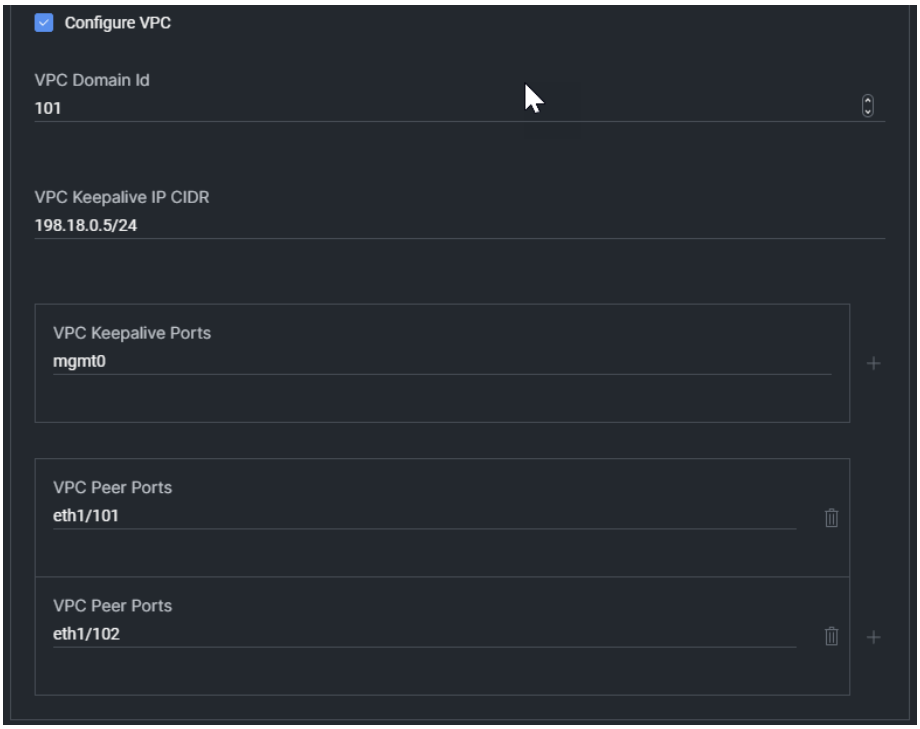

## VLANs

### Procedure 1. Assign Individual VLANs to be created

In the VLANs section we create individual VLANs correlating to the functions used by this design guide. You can also define more traditional VLANs with the Virtual Machine VLAN Type. But inband, ooband, iSCSI, NFS, Migration/vMotion, NVMe must be defined in this section if being used. This is where we will assign the subnet, IP ranges, gateway, subnet mask etc.

1. **VLAN Type**: This field shows the type of VLAN the configuration is used for. The VLAN type tells the automation what tasks to perform on a given VLAN. Table 9 gives details.


**Table 9.** VLAN types

| VLAN type | Usage |
|---|---|
| Inband | Configure Nexus VLAN if “Configure L2” is selected, SVI/HSRP if “Configure L3” is selected. This type of VLAN will configure inband VLAN pool with the pool range, NetApp inband MGMT interfaces in the SVM with the controller range, and Hypervisor vmk0 from the server range. |
| ooband | Configure Nexus SVI/HSRP if “Configure L3” is selected. This type of VLAN will configure ooband VLAN pool with the pool range and NetApp service-process interfaces in with the controller range |
| Migration/ vMotion | Configure Nexus SVI/HSRP if “Configure L3” is selected. Add vmk1 to the hypervisors with the server range. Only necessary if in the “Virtualization” section the migration Data Type was assigned to an individual virtual switch. |
| NFS | Configure Nexus VLAN if “Configure L2” is selected, SVI/HSRP if “Configure L3” is selected. Add vmk{num} to the hypervisor with the server range for mounting NFS volumes. |
| iSCSI | Configure Nexus VLAN if “Configure L2” is selected, SVI/HSRP if “Configure L3” is selected. Add a pair of vmk{num} on the hypervisors will be configured with these VLAN settings to connect to the storage array using the iSCSI protocol. Note that there should be a VLAN-A and a VLAN-B for redundancy of the fabrics. Additionally, a software storage controller interface for this VLAN type will be added to the hypervisor. |
| NVMe | Configure Nexus VLAN if “Configure L2” is selected, SVI/HSRP if “Configure L3” is selected. Add a pair of vmk{num} on the hypervisors will be configured with these VLAN settings to connect to the storage array using the NVMe protocol. Note that there should be a VLAN-A and a VLAN-B for redundancy of the fabrics. Additionally, a software storage controller interface for this VLAN type will be added to the hypervisor. |
| Virtual machine | Configure Nexus VLAN if “Configure L2” is selected, SVI/HSRP if “Configure L3” is selected. Use this VLAN type to configure additional VLANs that need to be added to the guest distributed virtual switch when you want to assign specific names to the VLANs. If you don’t care about the name of the VLAN, use the “VLAN ranges” section. |

2. VLAN ID: VLAN ID of the VLAN.

**Note:** Remember to make sure the VSAN IDs defined for the domain do not overlap with any consumed VLANs if you used the FCP deployment type.

3. **Name**: Name to Assign to the VLAN. The name will also be used by the automation to configure the interfaces of the NetApp appliance.
4. **Network CIDR**: This field contains the gateway address for the subnet including the network prefix/subnet mask.

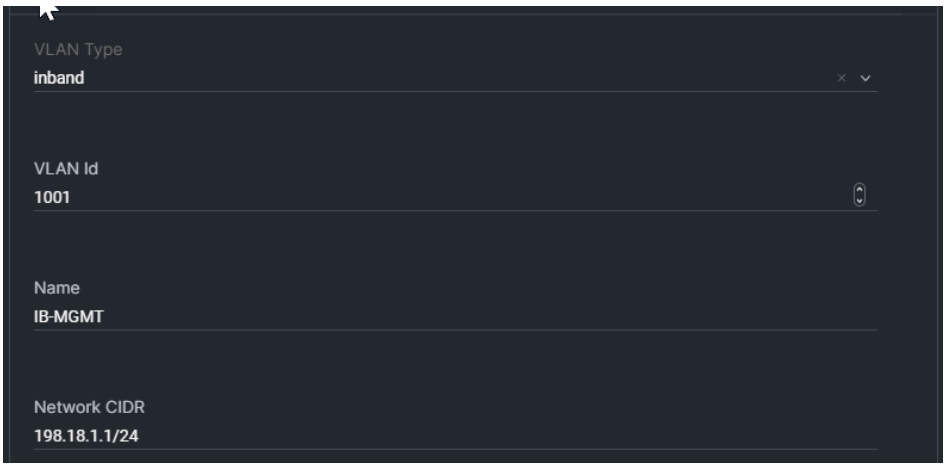


5. **Configure L2**: This flag determines if the process should configure the Layer 2 VLAN settings on the Cisco Nexus switches, including the VLAN and interface trunks for the NetApp Appliance and UCS Domain.
6. **Configure L3**: This flag determines if the process should configure the Layer 3 SVI/Hot Standby Router Protocol (HSRP) on the Cisco Nexus switches.
7. **Disjoint**: This flag is used to determines if the process should configure the VLAN on the UCS Domain as a disjoint/alternate uplink to the network. When you select Disjoint it determines the VLAN groups configured on the domain.
8. **Native VLAN**: Use this flag for uplink or disjoint uplink native VLAN configuration. For Server Profile Ethernet Network Groups, the VLAN type is assigned based on the virtual switch data types, with management vlan chosen first, then iSCSI if assigned, and finally migration.


9. **Controller Range**: This list is of IP addresses from the subnet to assign to the Storage controller.
10. **Pool Range**: This list is of IP addresses from the subnet to assign to the pool (inband/ooband).
11. **Server Range**: The list is of IP addresses from the subnet to Assign to the BareMetal server or hypervisor, depending on the VLAN type (refer to Table 10).

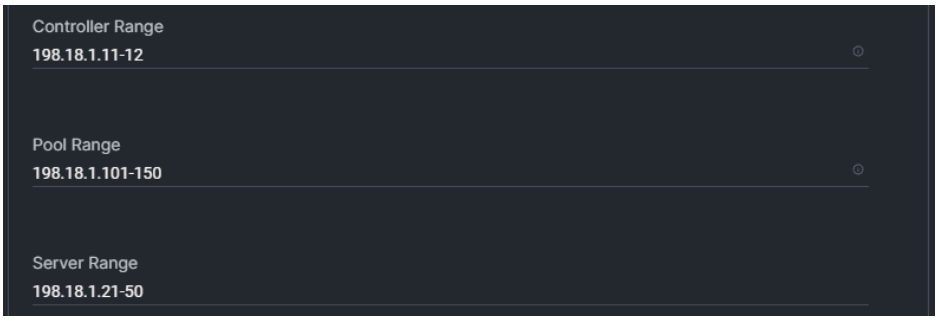


**Table 10.** VLANs for the validated design

| VLAN ID | VLAN name | VLAN type | Network CIDR | Controller range | Pool range | Server range |
|---|---|---|---|---|---|---|
| 1000 | OOB-MGMT | ooband | 198.18.0.1/24 | 198.18.0.11-15 | 198.18.0.101-150 | |
| 1001 | IB-MGMT | In-band | 198.18.1.1/24 | 198.18.1.11-12 | 198.18.1.101-150 | 198.18.1.21-50 |
| 1002 | vMotion | migration | 198.18.2.1/24 | | | 198.18.2.21-50 |
| 1011* | iSCSI-A | iscsi | 198.18.11.1/24 | 198.18.11.11-12 | | 198.18.11.21-50 |
| 1012* | iSCSI-B | iscsi | 198.18.12.1/24 | 198.18.12.11-12 | | 198.18.12.21-50 |
| 1013* | NVMe-TCP-A | nvme | 198.18.13.1/24 | 198.18.13.11-12 | | 198.18.13.21-50 |
| 1014* | NVMe-TCP-B | nvme | 198.18.14.1/24 | 198.18.14.11-12 | | 198.18.14.21-50 |
| 1015 | NFS | nfs | 198.18.15.1/24 | 198.18.15.11-12 | | 198.18.15.21-50 |

**Notes**: IP Range Requirements

- **OOB-MGMT.** The storage controller needs 5 IPs. one cluster IP, two IPs for the nodes, and 2 IPs for the NetApp appliance service processor interfaces. The pool range should be large enough to account for all the servers to be added to the domain(s).
- **IB-MGMT.** The storage controller needs two IPs, one for each node. The pool range should be large enough to account for all the servers to be added to the domain. The server range is configured as the vmk0 on the hypervisors; it should be a large enough range to account for all the servers added to the domain.
- **Migration.** The server range is configured as vmk1 on the hypervisors; it should be a large enough range to account for all the servers added to the domain. If a migration VLAN is not assigned, the "vmotion" flag is added to the management interface; in the case it is a VMware deployment.
- **iSCSI (A/B).** The storage controller needs 2 IPs, one for each node. The server range is configured as vmk(odd/even) on the hypervisors; it should be a large enough range to account for all the servers added to the domain(s). iSCSI VLANs are necessary only if you are planning to configure NVMe over TCP.
- **NVMe (A/B).** The storage controller needs 2 IPs, one for each node. The server range is configured as vmk(odd/even) on the hypervisors; it should be a large enough range to account for all the servers added to the domain(s). NVMe VLANs are necessary only if you are planning to configure NVMe over TCP.

## VLAN ranges

### Procedure 1. Define VLAN ranges to be created without Layer 3 definitions

**Note:** Unlike the previous "**VLANs**" section, the VLAN ranges are used to define VLANs that are deployed only with Layer2 configuration. These VLANs will not configure SVI and HSRP settings on the Cisco Nexus switches.

1. **VLAN Range**: List of VLANs, can include dash and comma, for instance 1-5,11-15.
2. **Name Prefix**: The automation will add -vl000X for 1-9, -vl00XX for 10-99, -vl0XXX for 100-999, and -vlXXXX for 1000-4094. This is to make sure the VLANs are in the correct sorting order when viewing them in the VLAN Policy.
3. **Configure L2**: This flag determines if the Cisco Nexus switch should be configured and add these VLANs to the trunk ports for the UCS Domain.
4. **Disjoint**: This flag identifies to the script that these VLANs belong to a disjoint network and to split the Ethernet Network Uplink ports. It also tells the automation to add these VLANs to the second uplink for disjoint traffic.

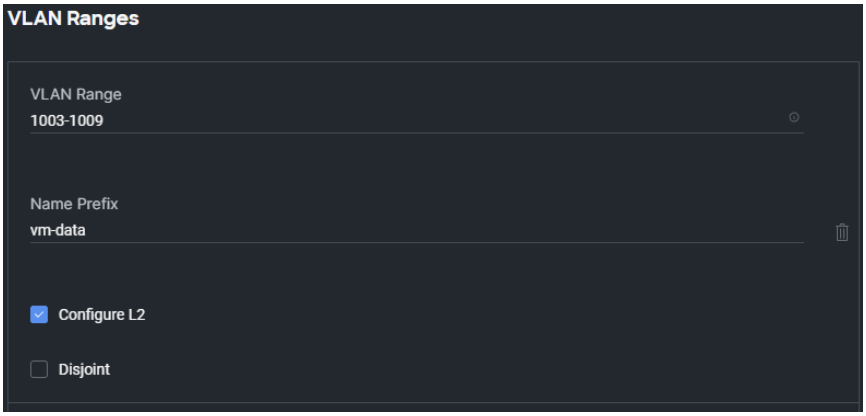


## Execute the Workflow

It would be good review the answers provided in the workflow. When you feel confident that the data entered matches what you would like to deploy, click the "**Execute**" button. If everything was entered correctly and the environment was cabled correctly, you should see a Status of "success" after approximately 4 hours:

- 10 minutes to configure the Cisco Nexus switches
- 15 minutes to configure the NetApp Storage Array
- 30-40 minutes to deploy the domain profile (assuming Fis must be rebooted to account for Fibre Channel connections. if not, subtract 20 minutes)
- 45 minutes to deploy server profiles with firmware upgrades: the testing environment consists of servers with 768 GB to 2 TB of RAM. Memory sizing may increase or decrease these times.
- 60 minutes to deploy the operating system; this value was found with 16 servers for the testing. You can deploy a maximum of 100 servers at one time. More than 30 servers will increase the time to install. Again, memory configuration can cause variances in these estimates.
- 60 minutes to configure the vCenter environment, due to reboots for driver installation for the servers.

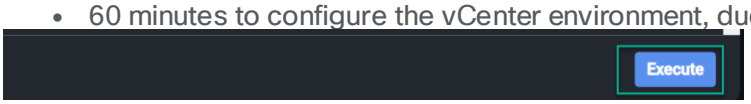

You can also view the status of many of the tasks from the Requests view in the Cisco Intersight.

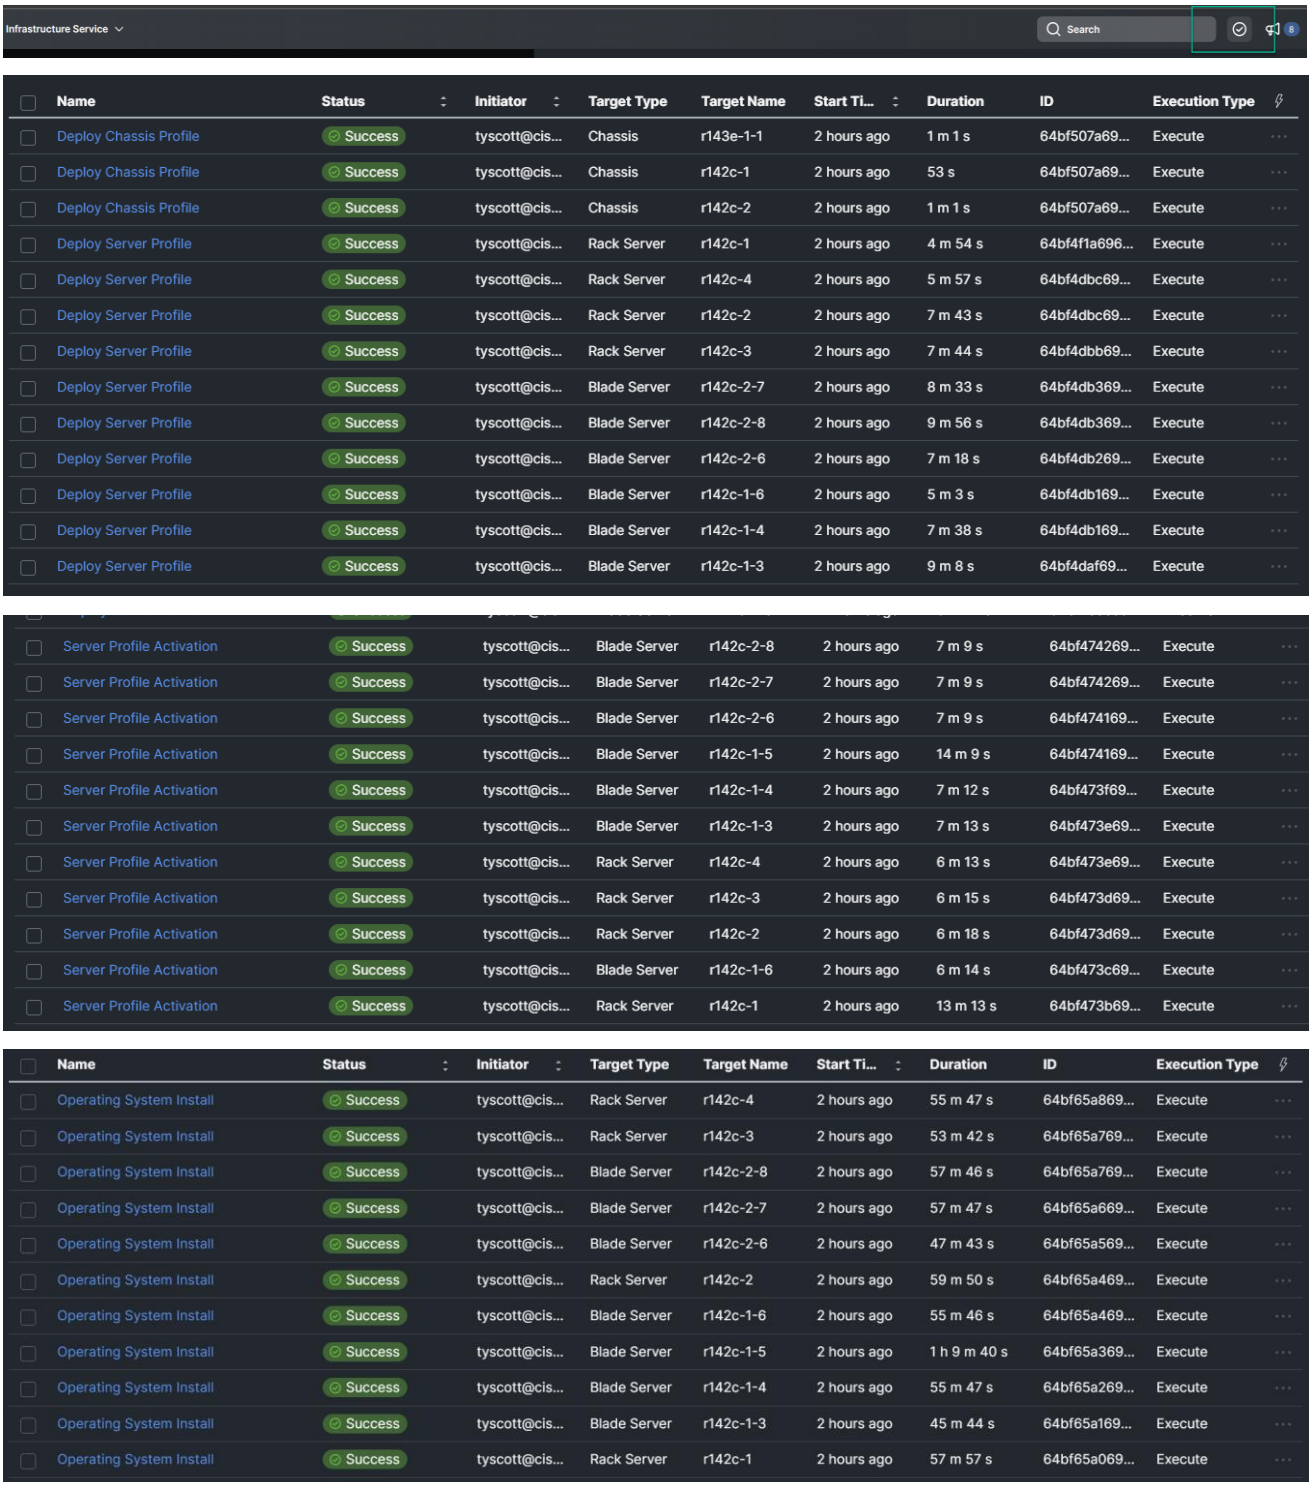


| Name | Status | Initiator | Target Type | Target Name | Start Ti... | Duration | ID | Execution Type |
|---|---|---|---|---|---|---|---|---|
| Deploy Chassis Profile | Success | tyscott@cis... | Chassis | r143e-1-1 | 2 hours ago | 1 m 1 s | 64bf507a69... | Execute |
| Deploy Chassis Profile | Success | tyscott@cis... | Chassis | r142c-1 | 2 hours ago | 53 s | 64bf507a69... | Execute |
| Deploy Chassis Profile | Success | tyscott@cis... | Chassis | r142c-2 | 2 hours ago | 1 m 1 s | 64bf507a69... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Rack Server | r142c-1 | 2 hours ago | 4 m 54 s | 64bf4f1a696... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Rack Server | r142c-4 | 2 hours ago | 5 m 57 s | 64bf4dbc69... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Rack Server | r142c-2 | 2 hours ago | 7 m 43 s | 64bf4dbc69... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Rack Server | r142c-3 | 2 hours ago | 7 m 44 s | 64bf4dbb69... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-2-7 | 2 hours ago | 8 m 33 s | 64bf4db369... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-2-8 | 2 hours ago | 9 m 56 s | 64bf4db369... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-2-6 | 2 hours ago | 7 m 18 s | 64bf4db269... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-1-6 | 2 hours ago | 5 m 3 s | 64bf4db169... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-1-4 | 2 hours ago | 7 m 38 s | 64bf4db169... | Execute |
| Deploy Server Profile | Success | tyscott@cis... | Blade Server | r142c-1-3 | 2 hours ago | 9 m 8 s | 64bf4daf69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-2-8 | 2 hours ago | 7 m 9 s | 64bf474269... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-2-7 | 2 hours ago | 7 m 9 s | 64bf474269... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-2-6 | 2 hours ago | 7 m 9 s | 64bf474169... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-1-5 | 2 hours ago | 14 m 9 s | 64bf474169... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-1-4 | 2 hours ago | 7 m 12 s | 64bf473f69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-1-3 | 2 hours ago | 7 m 13 s | 64bf473e69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Rack Server | r142c-4 | 2 hours ago | 6 m 13 s | 64bf473e69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Rack Server | r142c-3 | 2 hours ago | 6 m 15 s | 64bf473d69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Rack Server | r142c-2 | 2 hours ago | 6 m 18 s | 64bf473d69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Blade Server | r142c-1-6 | 2 hours ago | 6 m 14 s | 64bf473c69... | Execute |
| Server Profile Activation | Success | tyscott@cis... | Rack Server | r142c-1 | 2 hours ago | 13 m 13 s | 64bf473b69... | Execute |
| Operating System Install | Success | tyscott@cis... | Rack Server | r142c-4 | 2 hours ago | 55 m 47 s | 64bf65a869... | Execute |
| Operating System Install | Success | tyscott@cis... | Rack Server | r142c-3 | 2 hours ago | 53 m 42 s | 64bf65a769... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-2-8 | 2 hours ago | 57 m 46 s | 64bf65a769... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-2-7 | 2 hours ago | 57 m 47 s | 64bf65a669... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-2-6 | 2 hours ago | 47 m 43 s | 64bf65a569... | Execute |
| Operating System Install | Success | tyscott@cis... | Rack Server | r142c-2 | 2 hours ago | 59 m 50 s | 64bf65a469... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-1-6 | 2 hours ago | 55 m 46 s | 64bf65a469... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-1-5 | 2 hours ago | 1 h 9 m 40 s | 64bf65a369... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-1-4 | 2 hours ago | 55 m 47 s | 64bf65a269... | Execute |
| Operating System Install | Success | tyscott@cis... | Blade Server | r142c-1-3 | 2 hours ago | 45 m 44 s | 64bf65a169... | Execute |
| Operating System Install | Success | tyscott@cis... | Rack Server | r142c-1 | 2 hours ago | 57 m 57 s | 64bf65a069... | Execute |

When the deployment is finished, you will see a Success status within the Workflow Execution Pane.

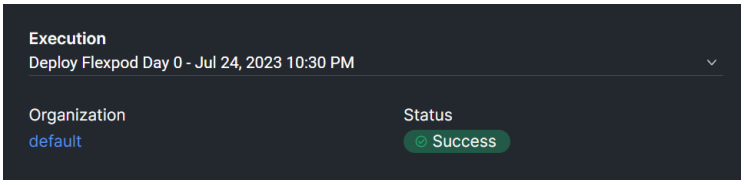

## Workflow Validation

Following are a few things to check in the virtualization environment to validate after the completion of the workflow.

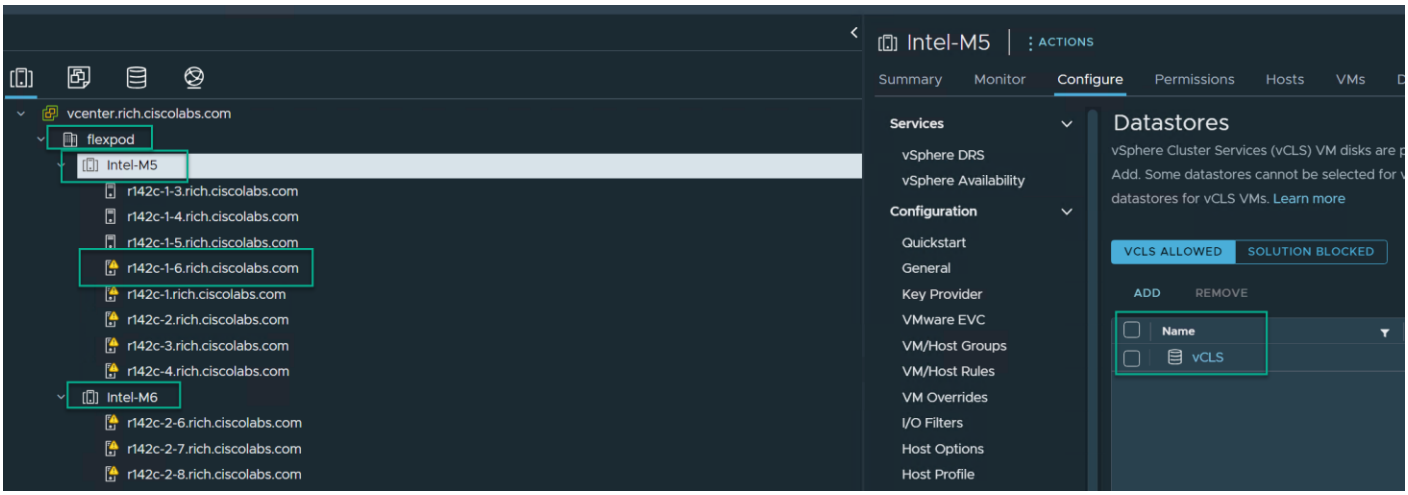


You will create the data center based on the name you gave in the workflow. We used the name “flexpod”. The cluster names are automatically derived based on the hardware profile of the servers in the domain. In our case, we had Intel M5 and M6 generation servers. It created a cluster for each model. You can change the names to suit your deployment, but this way was the easiest to group without asking the cluster name for individual servers, which then requires details for individual servers for profiles, location, IPs etc. If a new cluster is created, the configuration applied, if the license tier is Enterprise, a Disaster Recovery System (DRS) and high availability (HA) was enabled; If a standard license tier, then only HA was enabled. The vSphere Cluster Service (vCLS) datastore was assigned to the cluster based on the datastore created for this object.

**Note:** **Note**: Servers that have a TPM module installed will show up with a warning identifying that the TPM Encryption Recovery Key should be backed up.


This key should be backed up in Hardware needs replacement; if so, the key needs to be re-imported.

### Procedure 1. Avoiding boot failure when UEFI secure booted server profiles are moved

Typically, hosts are configured for boot from SAN. Cisco UCS supports stateless compute where you can move a server profile from one blade or compute node to another seamlessly.

When a server profile is moved from one blade to another and the blade server has the following conditions, the ESXi host runs into PSOD and ESXi will fail to boot:

- TPM is present in the node (Cisco UCS M5/M6/M7 family of servers).
- Host is installed with ESXi 7.0 U2 or later.
- Boot mode is Unified Extensible Hardware Interface (UEFI) secure.
- Error message: Unable to restore system configuration. A security violation was detected. [Boot time failures due to ESXi Configuration encryption]

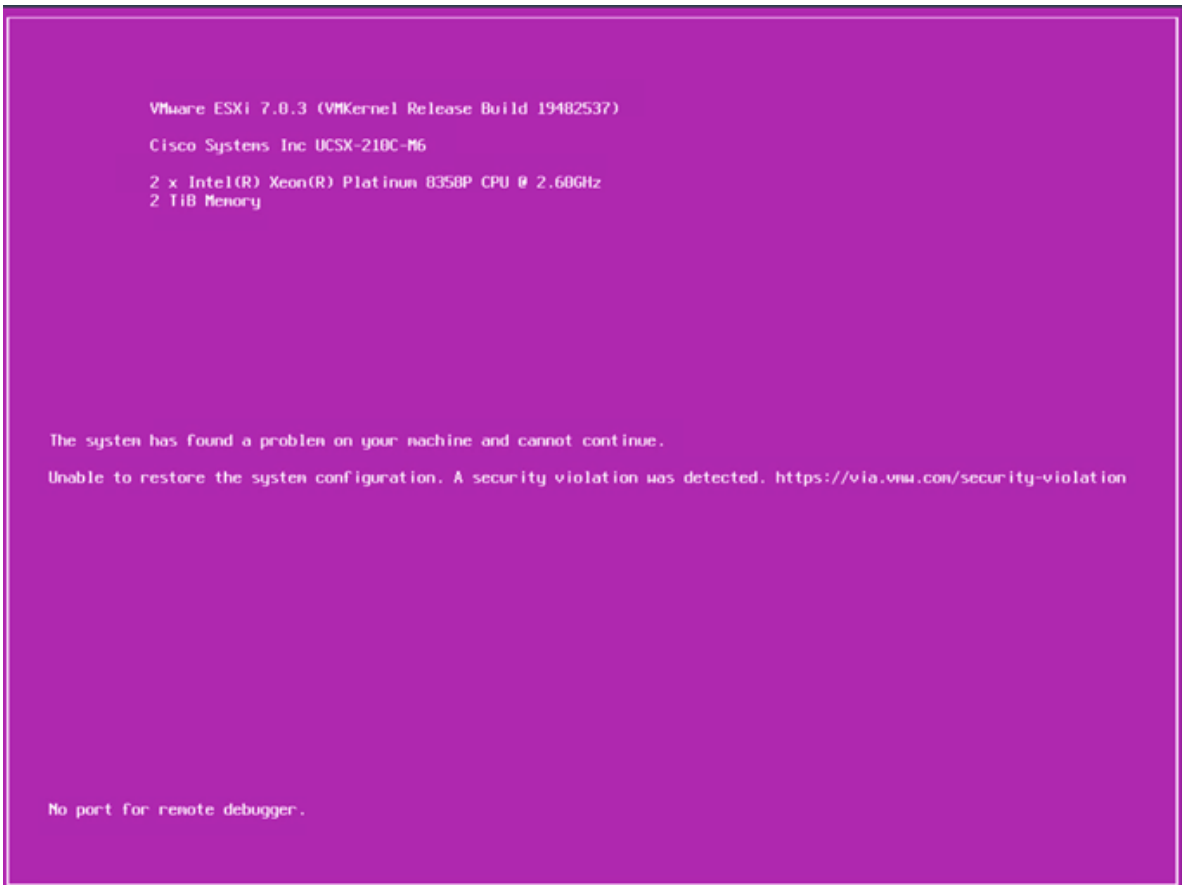


1. Log into the host using Secure Shell (SSH) Protocol.
2. Gather the recovery key using this command:

```
[root@aa02-esxi-1:~] esxcli system settings encryption recovery list
Recovery ID Key
------------------------------------ ---
{74AC4D68-FE47-491F-B529-6355D4AAF52C} 529012-402326-326163-088960-184364-097014-
312164-590080-407316-660658-634787-601062-601426-263837-330828-197047
```

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

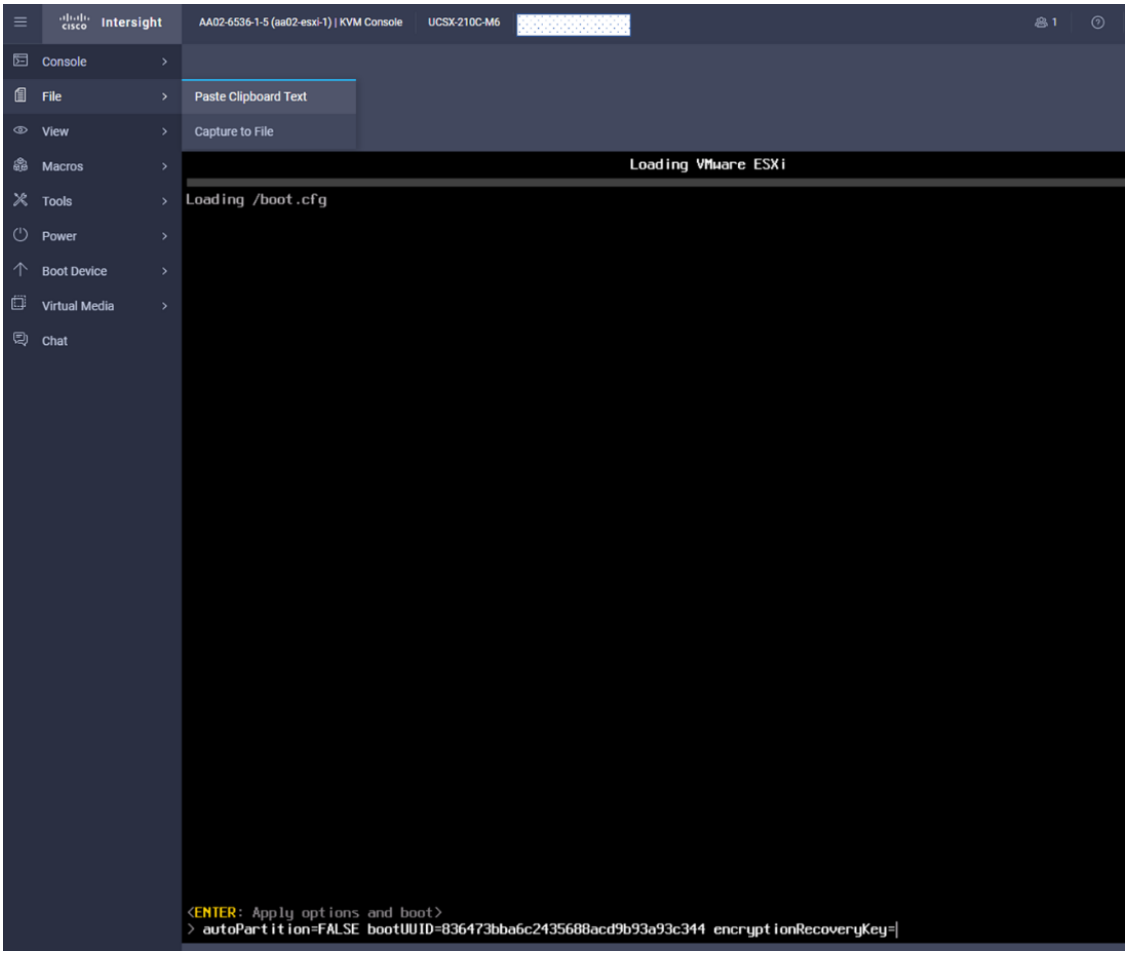


5. Add the recovery key using following boot option: encryptionRecoveryKey=*recovery_key*. Press Enter to continue the boot process.
6. To persist the change, enter the following command at the VMware ESXi SSH command prompt:

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

**Note:** For more information, refer to: https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-23FFB8BB-BD8B-46F1-BB59-D716418E889A.html.

## FlexPod Management Tools Setup

This chapter contains the following:

- Cisco Intersight Hardware Compatibility List (HCL) Status
- NetApp ONTAP Tools 9.12 Deployment
- Provision Datastores using NetApp ONTAP Tools (Optional)
- Virtual Volumes – vVol (Optional)
- NetApp SnapCenter 4.7 Configuration
- Active IQ Unified Manager 9.12P1 Installation
- Configure Active IQ Unified Manager

- Claim VMware vCenter Using Cisco Intersight Assist Appliance
- Claim NetApp Active IQ Manager Using Cisco Intersight Assist Appliance
- Claim Cisco Nexus Switches Using Cisco Intersight Assist Appliance

## Cisco Intersight HCL status

Cisco Intersight software evaluates the compatibility of a customer's Cisco UCS system to make sure that Cisco or Cisco partners have tested and validated the hardware and software. Cisco Intersight reports validation problems after checking the compatibility of the server model, processor, firmware, adapters, operating system, and drivers, and displays the compliance status with the Hardware Compatibility List (HCL).

To determine HCL compatibility for VMware ESXi, Cisco Intersight software uses Cisco UCS Tools, which is part of VMware ESXi Cisco custom ISO, and no additional configuration is required.

For more details about Cisco UCS Tools manual deployment and troubleshooting, refer to: Cisco UCS Tools

**Procedure 1. View compute node hardware compatibility**

1. To find detailed information about the hardware compatibility of a compute node, in Cisco Intersight software select **Infrastructure Service > Operate > Servers** in the left menu bar, click a server, select **HCL**.

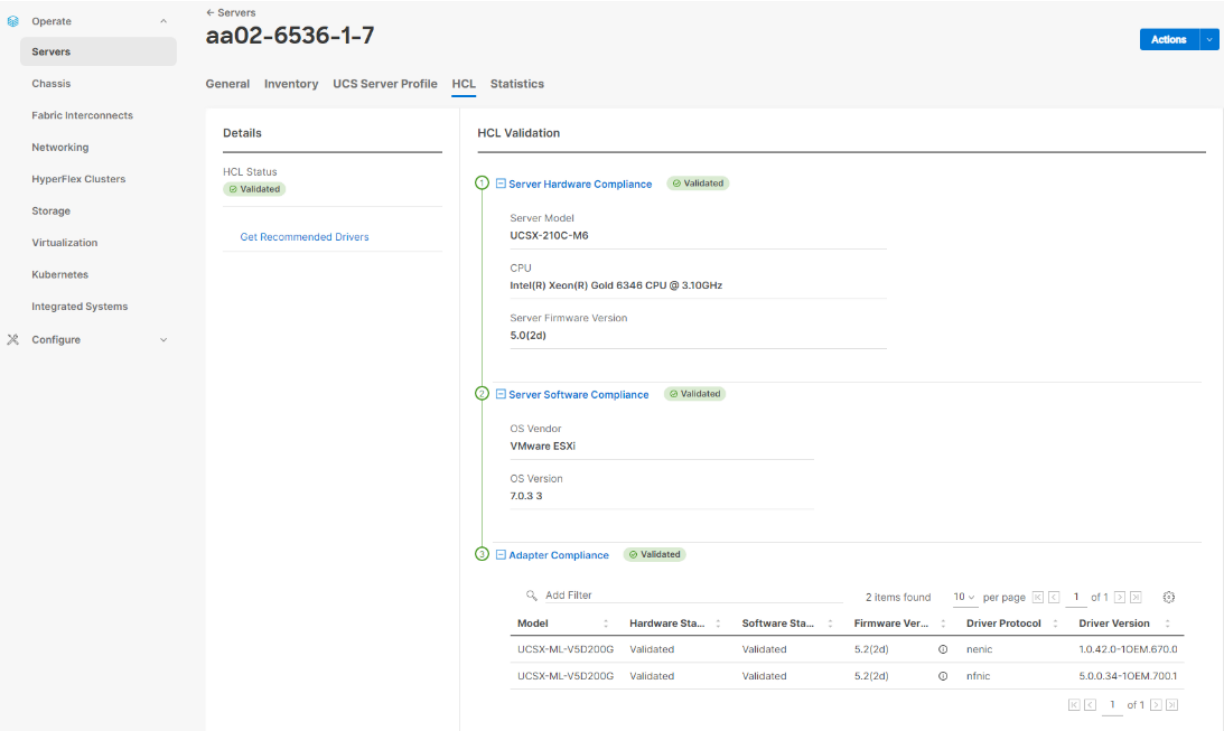


## NetApp ONTAP Tools 9.12 deployment

### 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 NetApp ONTAP Tools for VMware 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 later:

- Protocol licenses (NFS, FCP, and/or iSCSI)
- NetApp FlexClone is optional, but is required for performing test failover operations for Storage Replication Adapter SRA and for virtual volumes (vVols) operations of the vSphere API for Storage Awareness (VASA) provider.
- NetApp SnapRestore (for backup and recovery).
- The NetApp SnapManager Suite.
- NetApp SnapMirror or NetApp SnapVault (Optional - required for performing failover operations for Storage Replication Adapter (SRA) and VASA Provider when using vVols replication).

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

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

Table 11 lists the port requirements for NetApp ONTAP Tools.

**Table 11.** Port requirements for NetApp ONTAP Tools

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

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

### Procedure 1. Install NetApp ONTAP Tools manually

1. Download the NetApp ONTAP Tools 9.12 OVA (NETAPP-ONTAP-TOOLS-FOR-VMWARE-VSPHERE-9.12-9342.OVA) from NetApp support: NetApp Support Site - ONTAP tools for VMware vSphere (Downloads) - 9.12.
2. Launch the vSphere Web Client and navigate to **Hosts and Clusters**.
3. Select **ACTIONS** for the FlexPod-DC datacenter and select **Deploy OVF Template**.
4. Browse to the NetApp ONTAP tools OVA file and select the file.
5. Enter the VM name and select a datacenter or folder to deploy the VM and click **NEXT**.

6. Select a host cluster resource to deploy OVA and click **NEXT**.
7. Review the details and accept the license agreement.
8. Select the infra_datastore volume and Select the **Thin Provision** option for the virtual disk format.
9. From **Select Networks**, select a destination network (for example, IB-MGMT) and click **NEXT**.
10. From Customize Template, enter the NetApp ONTAP tools administrator password, vCenter name or IP address, and other configuration details and click **NEXT**.
11. Review the configuration details entered and click **FINISH** to complete the deployment of NetApp ONTAP-Tools VM.

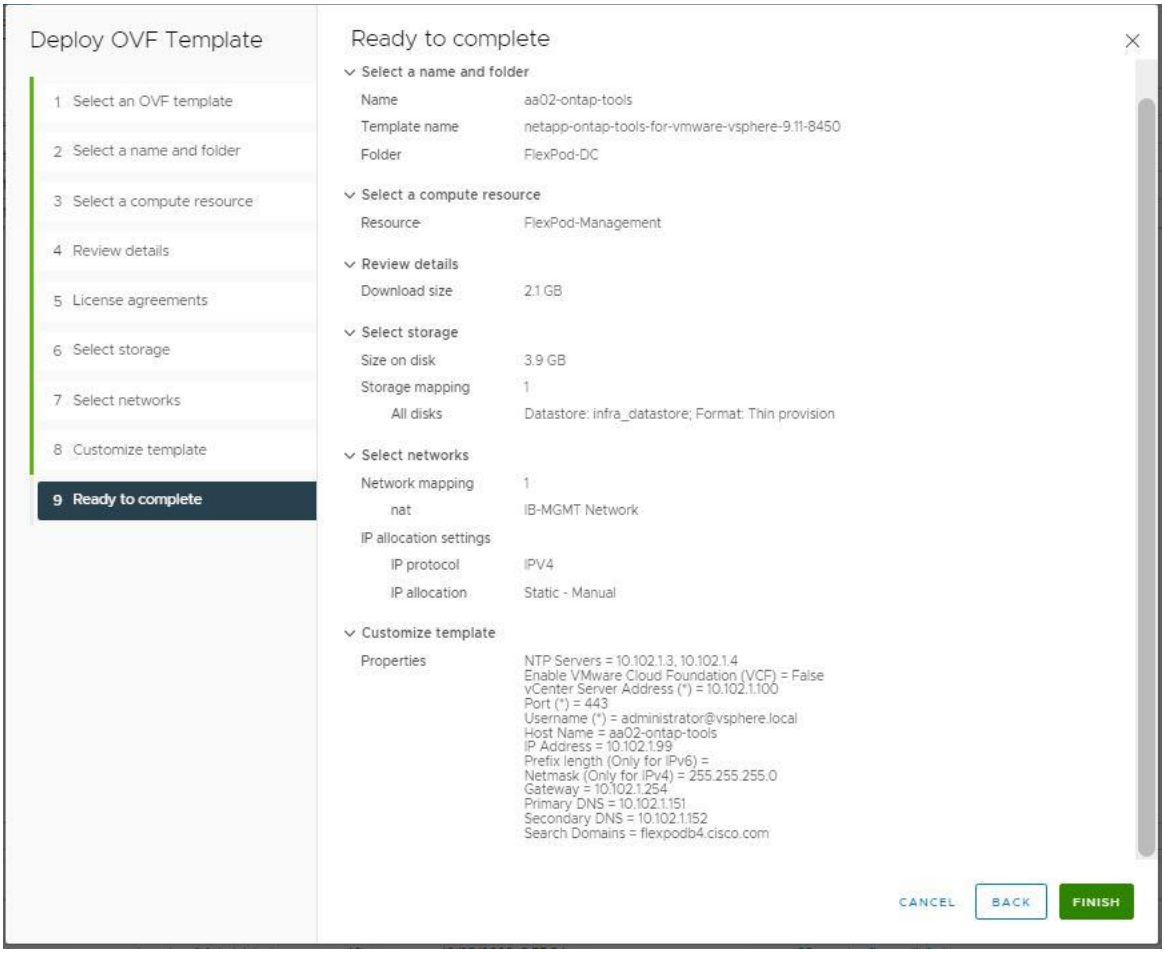


12. Power on the NetApp ONTAP-tools VM and open the VM console.
13. During the NetApp ONTAP-tools VM boot process, you should see a prompt to install VMware Tools. From vCenter, right-click the **ONTAP-tools VM > Guest OS > Install VMware Tools**.
14. Networking configuration and vCenter registration information was provided during the OVF template customization, so after the VM is up and running, NetApp ONTAP-Tools and VASA are registered with vCenter.
15. Refresh the vCenter Home Screen and confirm that the NetApp ONTAP tools is installed.

**Note:** The NetApp ONTAP tools vCenter plug-in is available only in the vSphere HTML5 Client and is not available in the vSphere Web Client.

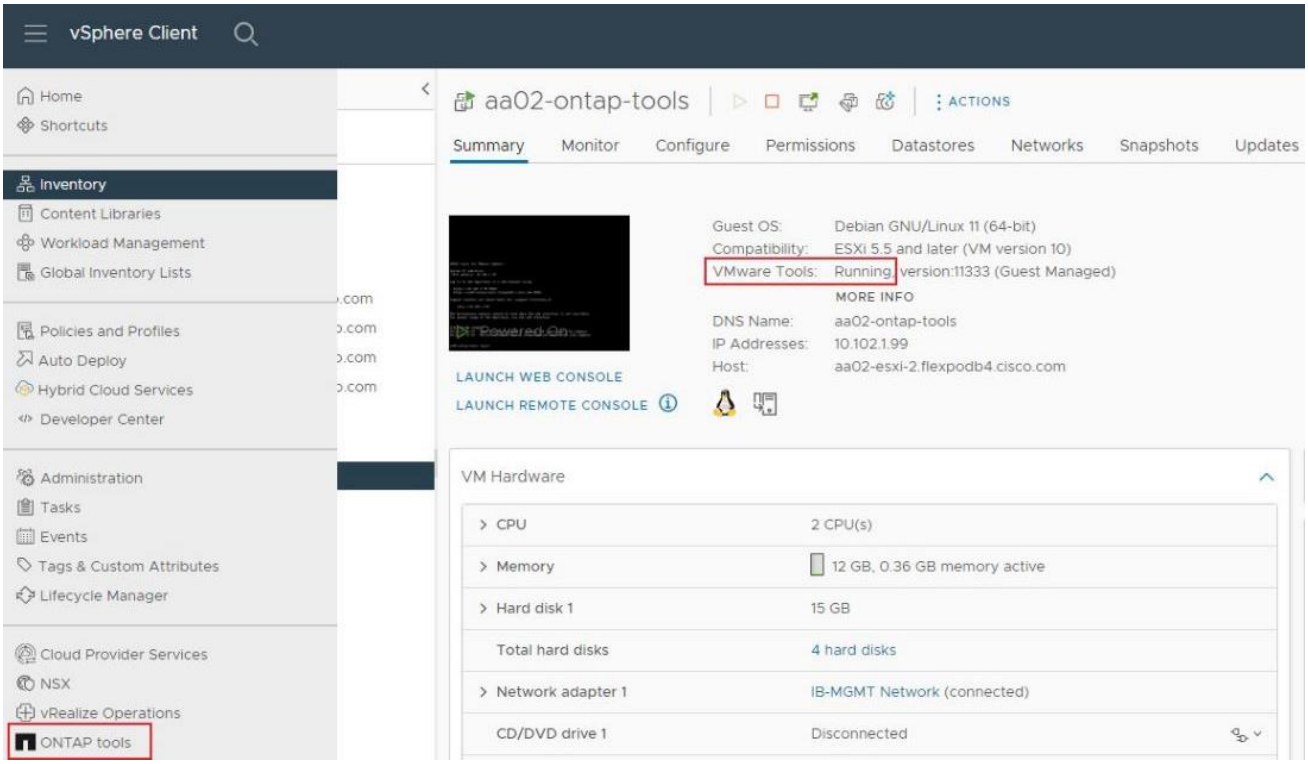


### Procedure 2. Download the NetApp NFS Plug-in for VAAI

**Note:** The NFS Plug-In for VAAI was previously installed on the ESXi hosts along with the Cisco UCS VIC drivers; it is not necessary to re-install the plug-in at this time. However, for any future additional ESXi host setup, instead of using esxcli commands, you can use NetApp ONTAP tools to install the NetApp NFS plug-in. The following steps upload the latest version of the plug-in to NetApp ONTAP Tools.

1. Download the NetApp NFS Plug-In 2.0 for VMware file from: NetApp NFS Plug-In for VMware VAAI
2. Unzip the file and extract NetApp_bootbank_NetAppNasPlugin_2.0-15.vib from **vib20 > NetAppNasPlugin.**
3. Rename the .vib file to NetAppNasPlugin.vib to match the predefined name that NetApp ONTAP Tools uses.
4. Click **Settings** in the NetApp ONTAP tool Getting Started page.
5. Click the NFS VAAI Tools tab.
6. Click **Change** in the Existing version section.
7. Browse and select the renamed .vib file, and then click **Upload** to upload the file to the virtual appliance.

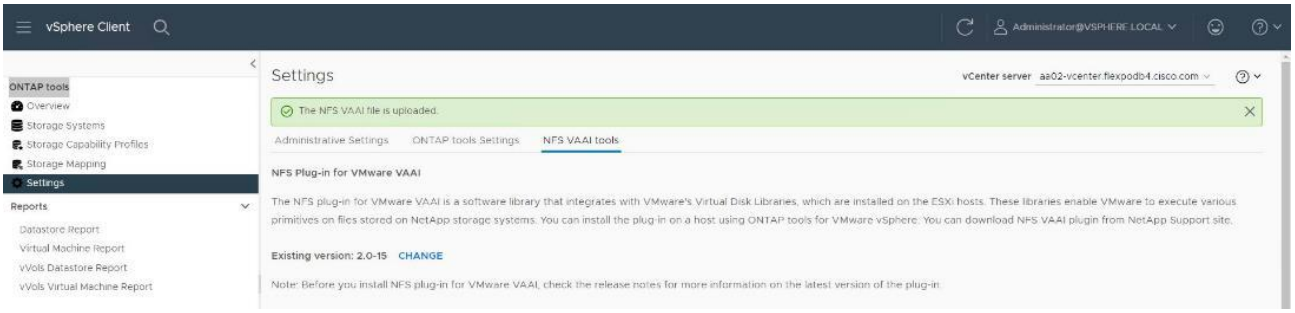

**Note:** The next step is required only on the hosts where the NetApp VAAI plug-in was not installed alongside Cisco VIC driver installation.

8. In the Install on “ESXi Hosts” section, select the ESXi host where the NFS Plug-In for VAAI is to be installed, and then click Install.
9. Reboot the ESXi host after the installation finishes.

### Procedure 3. Verify the VASA provider

**Note:** The VASA provider for NetApp ONTAP is enabled by default during the installation of the NetApp ONTAP Tools.

1. From the vSphere Client, click **Menu > ONTAP tools**.
2. Click Settings.
3. Click **Manage Capabilities** in the Administrative Settings tab.
4. In the Manage Capabilities dialog box, click **Enable VASA Provider** if it was not pre-enabled.
5. Enter the IP address of the virtual appliance for NetApp ONTAP tools, VASA Provider, and VMware Storage Replication Adapter (SRA), and the administrator password, and then click **Apply**.


### Procedure 4. Discover and add storage resources

1. Using the vSphere Web Client, log in to the vCenter. If the vSphere Web Client was previously opened, close the tab, and then reopen it.
2. In the Home screen, click the **Home** tab and click **ONTAP tools**.

**Note:** When using the cluster admin account, add storage from the cluster level.

**Note:** You can modify the storage credentials with the vsadmin account or another SVM-level account with role-based access control (RBAC) privileges. Refer to the [NetApp ONTAP 9 Administrator Authentication and RBAC Power Guide] for additional information.

3. Click **Storage Systems**, and then click **ADD** under Add Storage System.
4. Specify the vCenter Server where the storage will be located.
5. In the **Name or IP Address** field, enter the storage cluster management IP.
6. Enter admin for the username and the admin password for the cluster.
7. Confirm **Port** 443 to connect to this storage system.
8. Click **ADD** to add the storage configuration to NetApp ONTAP tools.

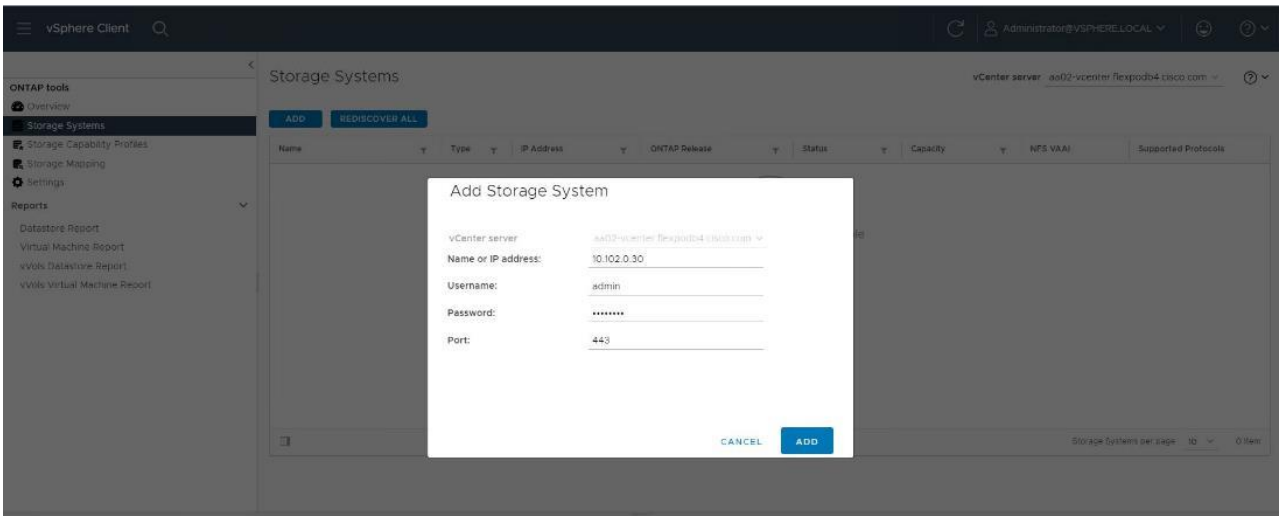


9. Wait for the Storage Systems to update. You might need to click **Refresh** to complete this update.

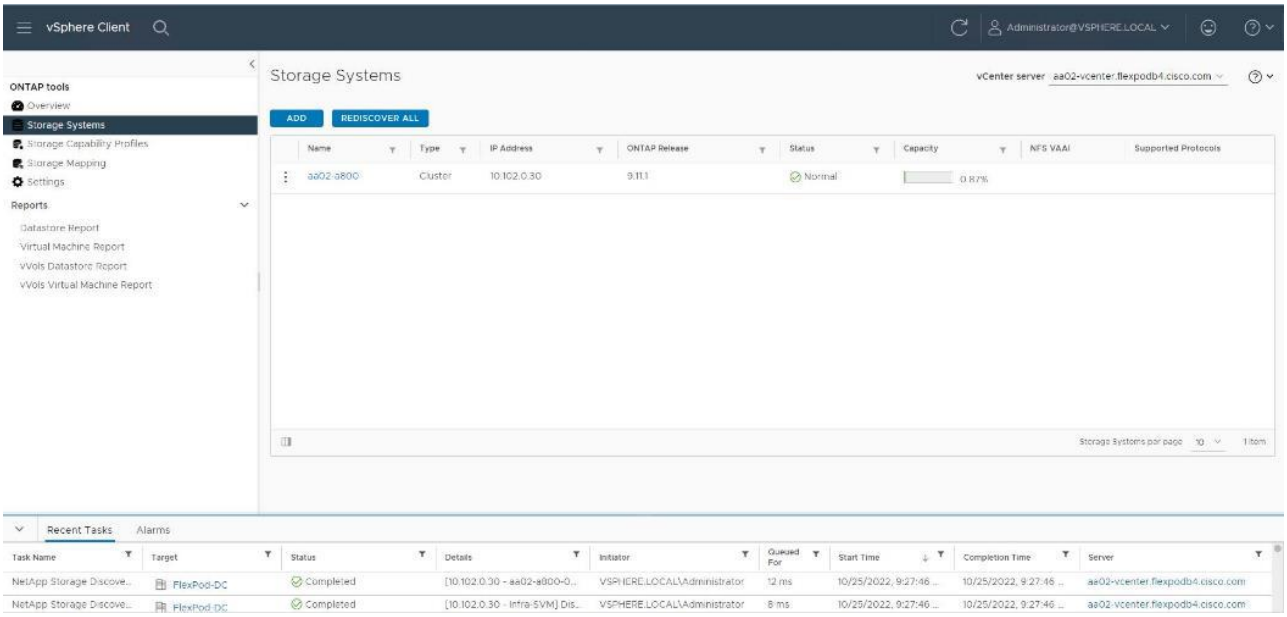


10. From the vSphere Client **Home** page, click **Hosts and Clusters**.

11. Right-click the FlexPod-DC datacenter; and click **NetApp ONTAP tools** > **Update Host and Storage Data**.

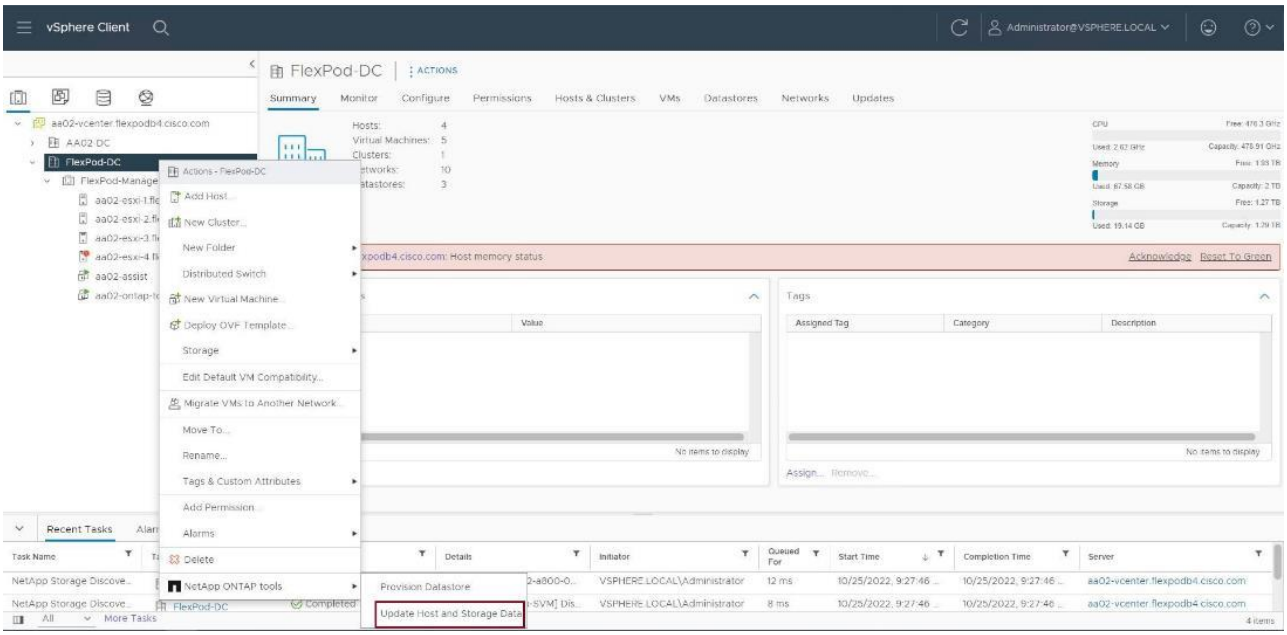


12. On the Confirmation dialog box, click **OK**. It might take a few minutes to update the data.

### Procedure 5. Optimal storage settings for ESXi hosts

**Note:** NetApp ONTAP tools enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers.

1. From the VMware vSphere Web Client Home page, click **vCenter > Hosts and Clusters**.
2. Select a host and then click Actions > NetApp ONTAP tools > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, select all the applicable values for the ESXi host.

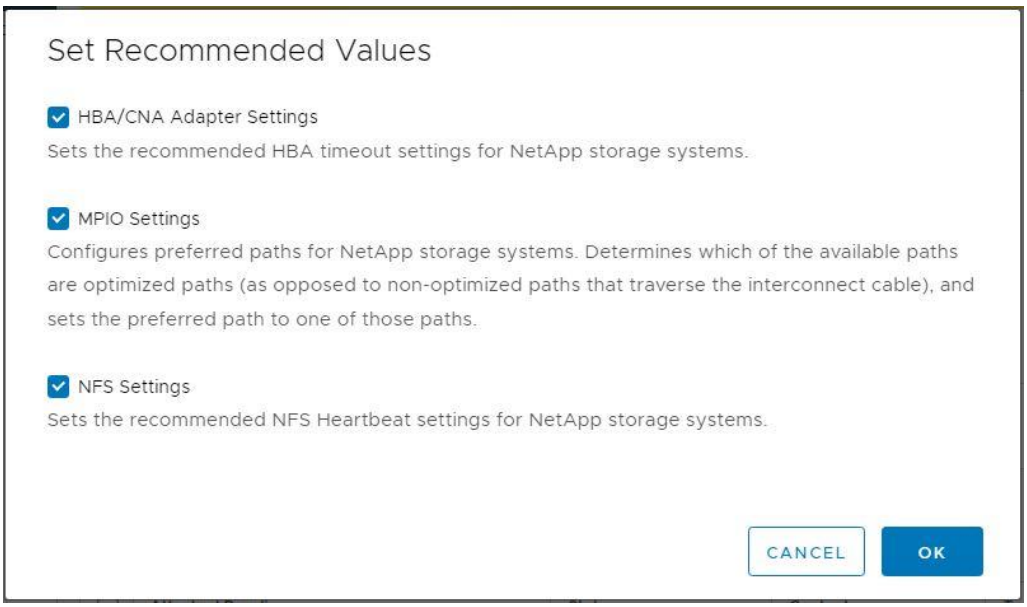

**Note:** This function sets values for host bus adapters (HBAs) and converged network adapters (CNAs), sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for NFS input/output (I/O). A vSphere host reboot may be required after you apply the settings.

4. Click **OK**.

## Provision datastores using NetApp ONTAP tools (optional)

Using NetApp ONTAP tools, the administrator can provision an NFS, FC, FC-NVMe or iSCSI datastore and attach it to a single or multiple hosts in the cluster. The following steps describe provisioning a datastore and attaching it to the cluster.

**Note:** It is a NetApp best practice to use NetApp ONTAP tools to provision any additional datastores for the FlexPod infrastructure. When using VSC to create vSphere datastores, all NetApp storage best practices are implemented during volume creation, and no additional configuration is needed to optimize performance of the datastore volumes.

### Storage capabilities

A storage capability is a set of storage system attributes that identifies a specific level of storage performance (storage service level), storage efficiency and other capabilities such as encryption for the storage object that is associated with the storage capability.

### Create the Storage Capability Profile

To use the automation features of VASA to advantage, you must first configure two primary components: the Storage Capability Profile (SCP) and the VM Storage Policy. The Storage Capability Profile expresses a specific set of storage characteristics into one or more profiles used to provision a virtual machine. The SCP is specified as part of the VM Storage Policy. NetApp ONTAP tools comes with several pre-configured SCPs such as Platinum, Bronze, and so on.

**Note:** The NetApp ONTAP tools for VMware vSphere plug-In also allow you to set QoS rules using a combination of maximum and/or minimum IOPs.

### Procedure 1. Review or edit the built-In profiles pre-configured with NetApp ONTAP tools

1. From the vCenter console, click **Menu > ONTAP tools**.
2. In the NetApp ONTAP tools click **Storage Capability Profiles**.
3. Select the **Platinum** Storage Capability Profile and select **Clone** from the toolbar.

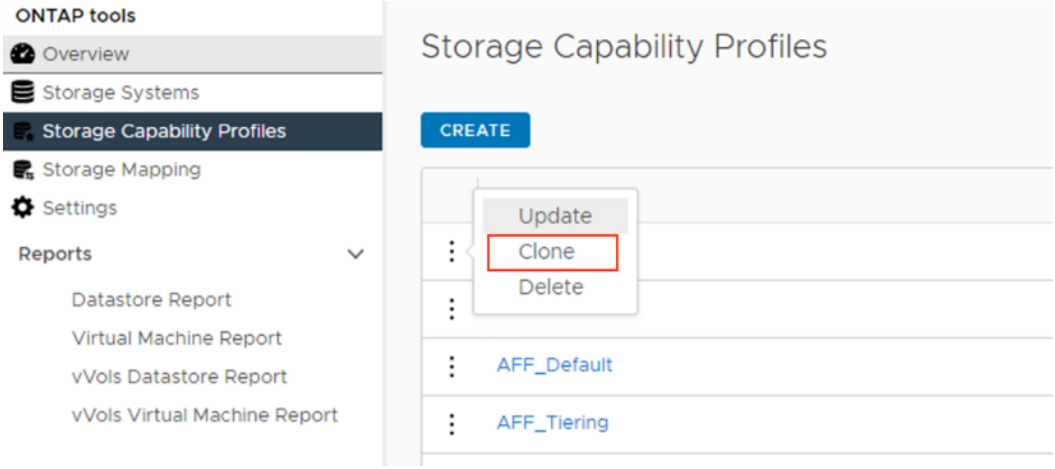

4. Enter a name for the cloned SCP (for example, AFF_Platinum_Encrypted) and add a description if desired. Click **NEXT**.

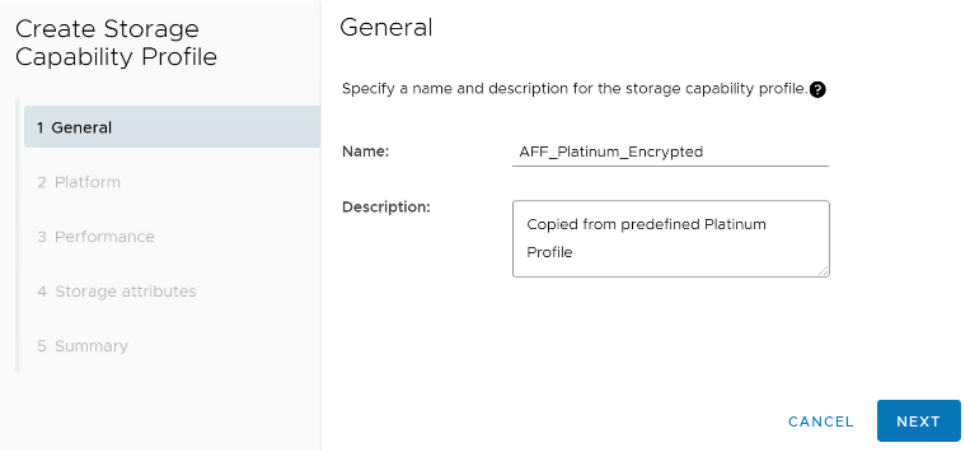


5. Select **All Flash FAS(AFF)** for the storage platform and then click **NEXT**.
6. Select **None** to allow unlimited performance or set the desired minimum and maximum IOPS for the QoS policy group. Click **NEXT**.
7. On the Storage attributes page**,** change the Encryption and Tiering policy to the desired settings and click **NEXT**. In the following example, Encryption was enabled.

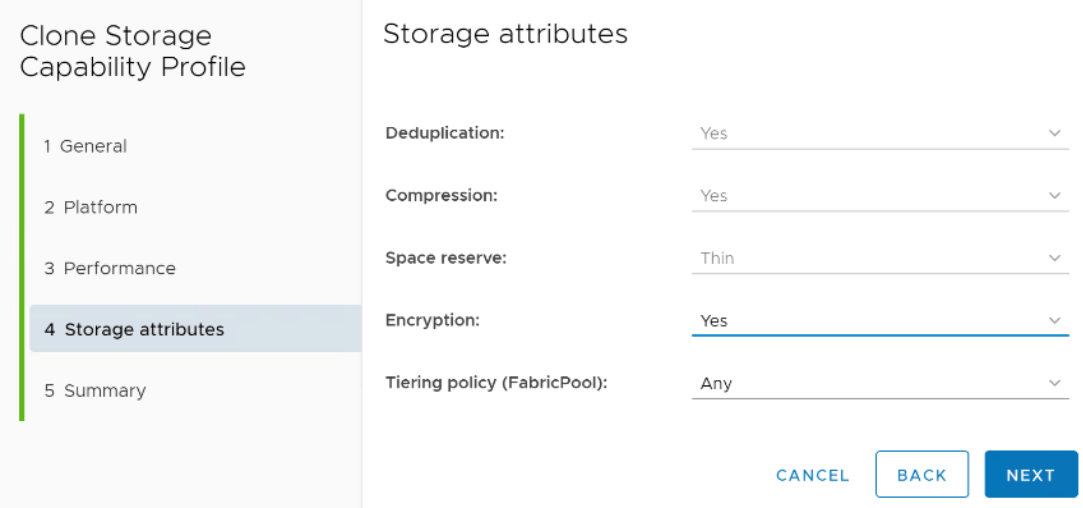


8. Review the summary page and click **FINISH** to create the storage capability profile.

**Note:** It is recommended to clone the Storage Capability Profile if you wish to make any changes to the predefined profiles rather than editing the built-in profile.

### Procedure 2. Create a VM storage policy

**Note:** You must create a VM storage policy and associate SCP to the datastore that meets the requirements defined in the SCP.

1. From the vCenter console, click **Menu > Policies and Profiles.**
2. Select VM Storage Policies and click CREATE.
3. Create a name for the VM storage policy and enter a description and click **NEXT**.

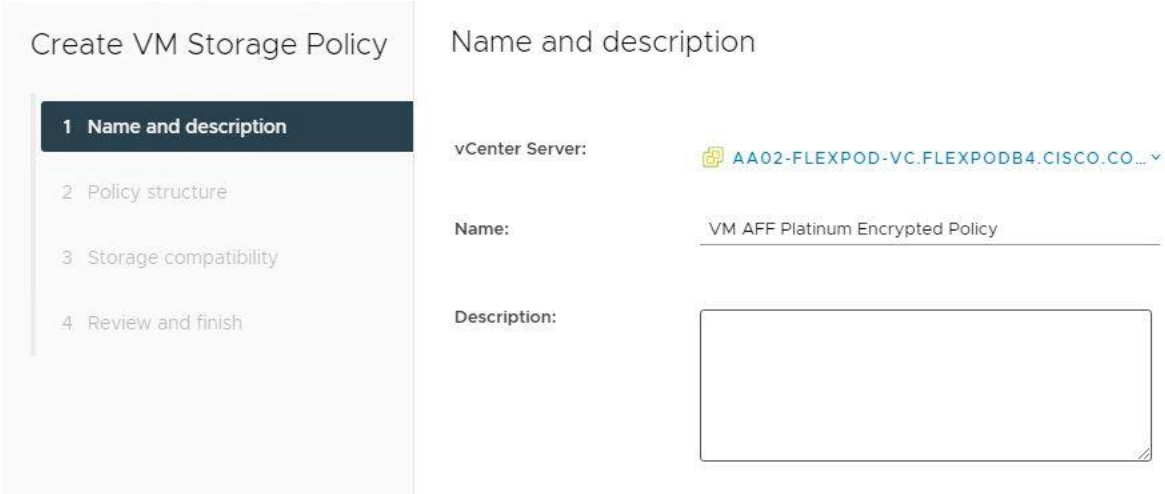


4. Select **Enable rules for NetApp.clustered.Data.ONTAP.VP.VASA10 storage** located under the Datastore specific rules section and click **NEXT**.

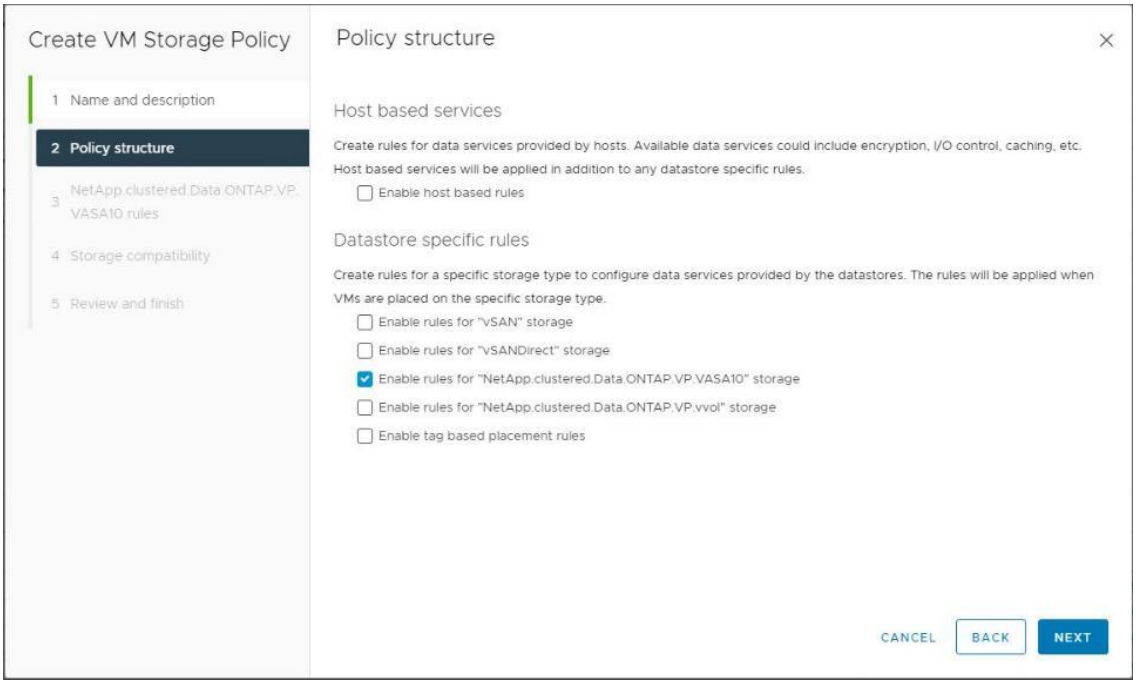


5. On the Placement tab select the SCP created in the previous step and click **NEXT**.

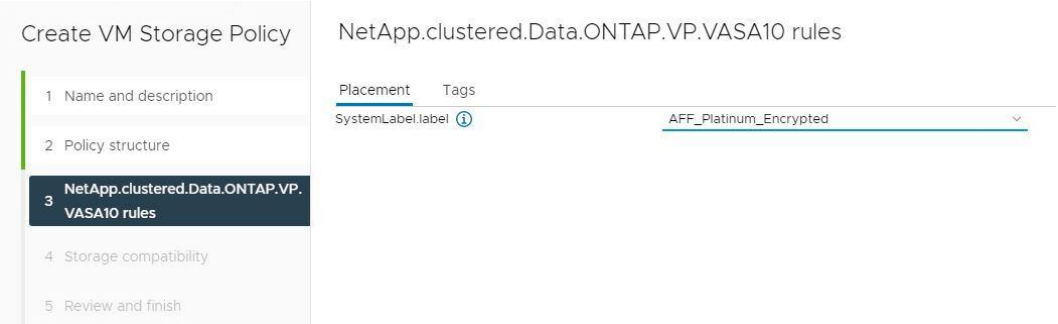


6. When all the datastores with matching capabilities are displayed, click **NEXT**.
7. Review the policy summary and click **FINISH**.

### Procedure 3. Provision NFS Datastore

1. From the vCenter console, click **Menu > ONTAP tools**.
2. From the NetApp ONTAP tools Home page, click **Overview**.
3. In the Getting Started tab, click **Provision**.
4. Click **Browse** to select the destination to provision the datastore.
5. Select the type as **NFS** and enter the datastore name (for example, NFS_DS_1).
6. Provide the size of the datastore and the NFS Protocol.
7. Check the storage capability profile and click **NEXT**.

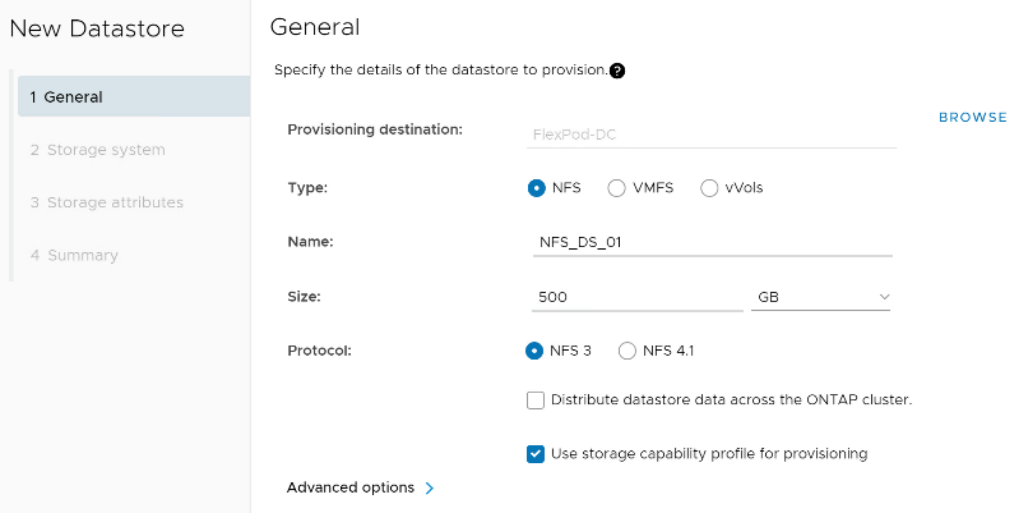


8. Select the desired Storage Capability Profile, cluster name, and the desired SVM to create the datastore. In this example, the Infra-SVM is selected.


9. Click **NEXT**.
10. Select the aggregate name and click **NEXT**.

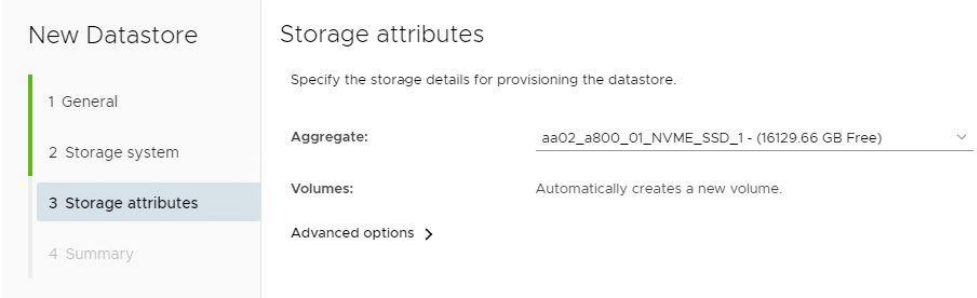


11. Review the Summary and click **FINISH**.


12. The datastore is created and mounted on the hosts in the cluster. Click Refresh from the vSphere Web Client to see the newly created datastore.
13. Distributed datastore is supported from NetApp ONTAP 9.8, which provides FlexGroup volume on NetApp ONTAP storage. To create a distributed datastore across the NetApp ONTAP Cluster, select NFS 4.1 and check the box for Distributed Datastore data across the NetApp ONTAP Cluster:

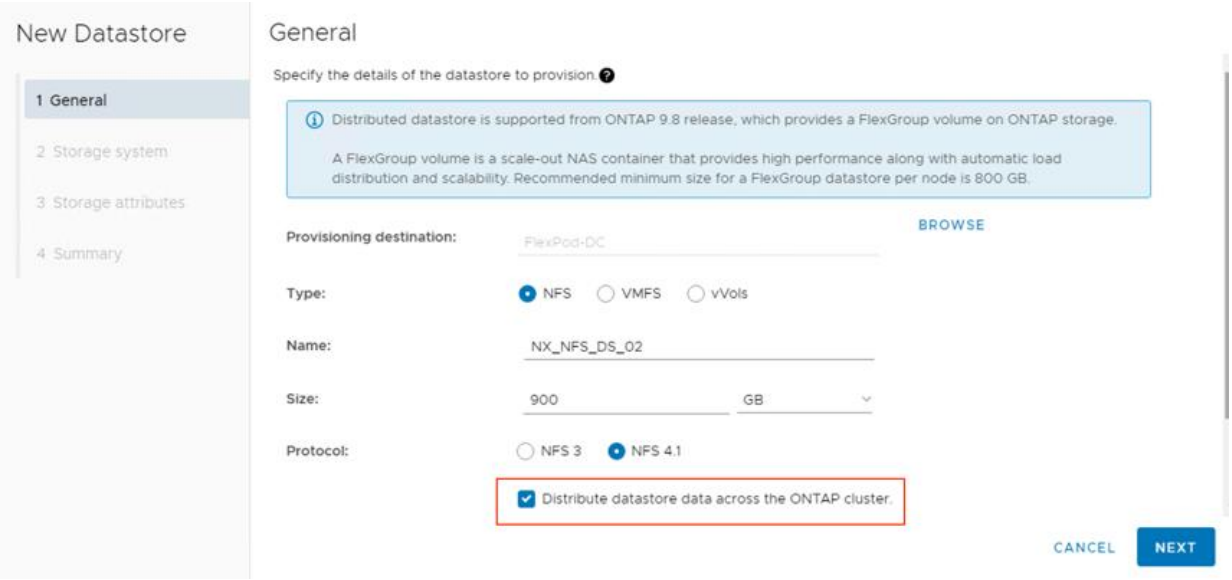


**Procedure 4. Provision FC datastore**

1. From the vCenter console, click **Menu > ONTAP tools**.
2. From the NetApp ONTAP tools Home page, click **Overview**.
3. In the Getting Started tab, click **Provision**.
4. Click **Browse** to select the destination to provision the datastore.
5. Select the type as **VMFS** and Enter the datastore name.
6. Provide the size of the datastore and the FC Protocol.

7. Check the Use storage capability profile and click **NEXT**.

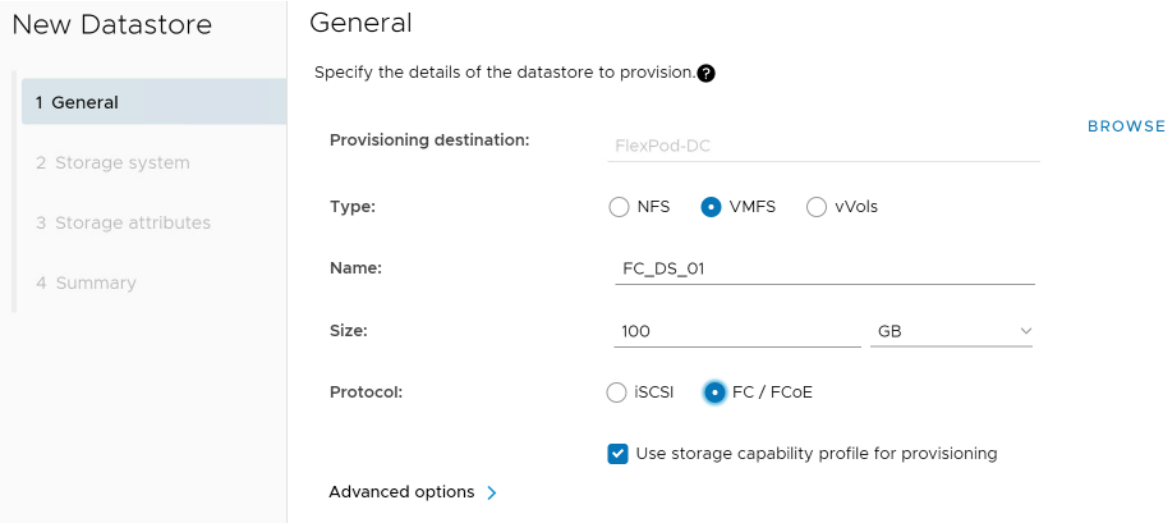


8. Select the **Storage Capability Profile**, **Storage System,** and the desired **Storage VM** to create the datastore.

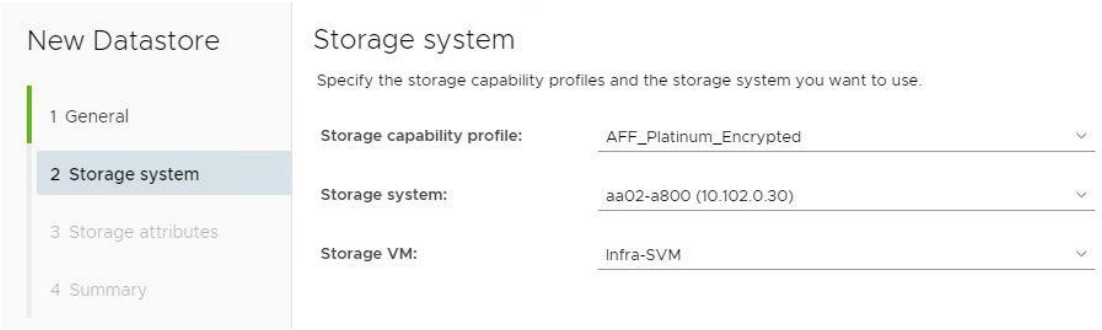


9. Click **NEXT**.
10. Select the aggregate name and click **NEXT**.

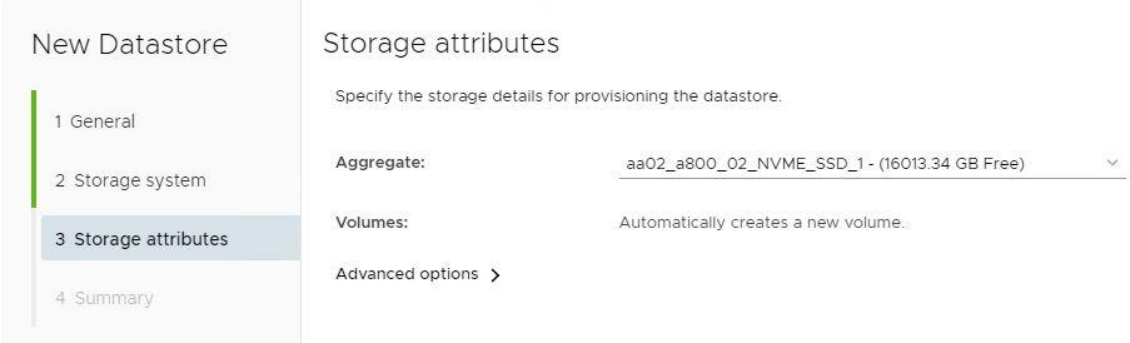


11. Review the Summary and click **FINISH**.

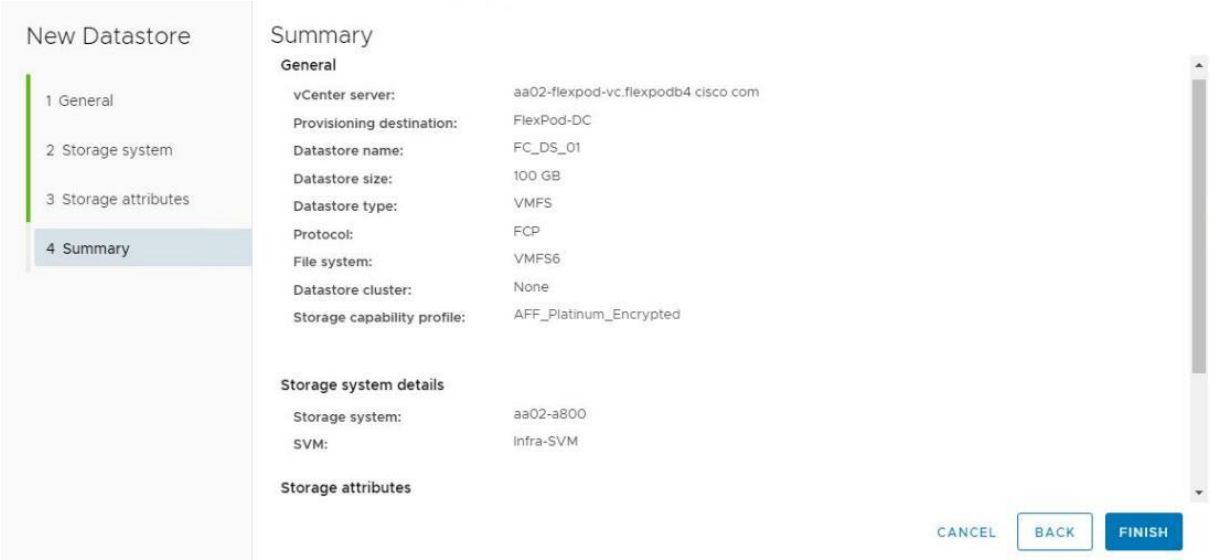


12. The datastore is created and mounted on all the hosts in the cluster. Click Refresh from the vSphere Web Client to see the newly created datastore.

**Procedure 5. Create virtual machine with Assigned VM Storage Policy**

1. Log into vCenter and navigate to the **VMs and Templates** tab and click to select the datacenter (for example, FlexPod-DC).
2. Click Actions and click New Virtual Machine.
3. Click Create a new virtual machine and click **NEXT**.
4. Enter a name for the VM and select the datacenter (for example, FlexPod-DC).
5. Select the cluster (for example, AA17-Cluster) and click **NEXT**.
6. Select the VM storage policy from the selections and select a compatible datastore. Click **NEXT**.

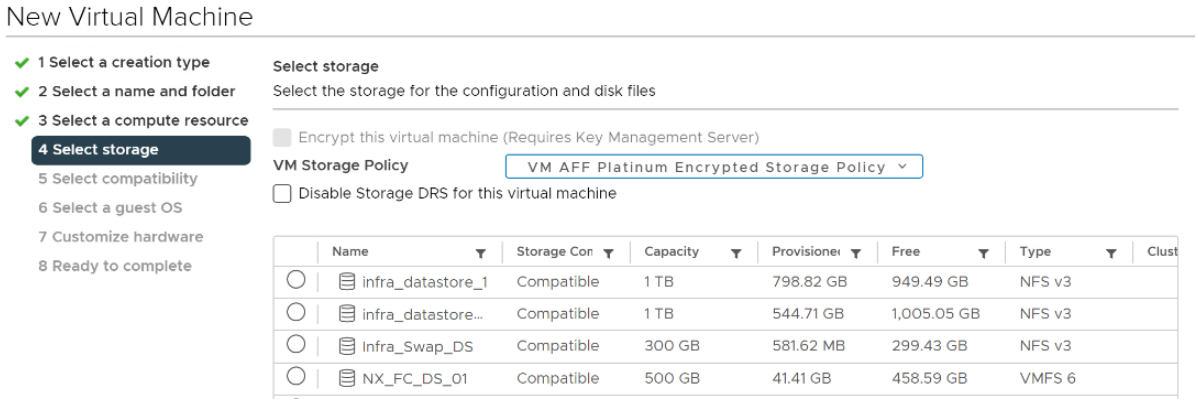


7. Select Compatibility (for example, ESXi 7.0 U2 or later) and click **NEXT**.
8. Select the Guest OS and click **NEXT**.
9. Customize the hardware for the VM and click **NEXT**.
10. Review the details and click **FINISH**.

**Note:** By selecting the VM storage policy in Step 6, the VM will be deployed on the compatible datastores.

## Virtual volumes (optional)

NetApp VASA Provider enables you to create and manage VMware virtual volumes (vVols). A vVols datastore consists of one or more FlexVol volumes within a storage container (also called "backing storage"). You can spread a virtual machine across one vVols datastore or multiple vVols datastores. All of the FlexVol volumes within the storage container must use the same protocol (NFS, iSCSI, or FCP) and the same SVMs.

For more information about vVOL datastore configuration, visit: FlexPod Datacenter with Cisco UCS X-Series VMware 7.0 U2, and NetApp ONTOP 9.9 – Virtual Volumes

### NetApp SnapCenter Plug-In 4.7 installation

SnapCenter Software is a centralized and scalable platform that provides application-consistent data protection for applications, databases, host file systems, and virtual machines running on NetApp ONTAP systems anywhere in the hybrid cloud.

#### NetApp SnapCenter architecture

The SnapCenter platform is based on a multitier architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent. The host agent that performs virtual machine and datastore backups for VMware vSphere is the SnapCenter Plug-In for VMware vSphere. It is packaged as a Linux appliance (Debian-based Open Virtual Appliance format) and is no longer part of the SnapCenter Plug-Ins Package for Windows. Additional information about deploying SnapCenter Server for application backups is available in the documentation listed in the Note that follows.

This guide focuses on deploying and configuring the SnapCenter Plug-In for VMware vSphere to protect virtual machines and their datastores.

**Note:** You must install SnapCenter Server and the necessary plug-ins to support application-consistent backups for Microsoft Sequenced Query Language (SQL), Microsoft Exchange, Oracle databases, and SAP HANA. Application-level protection is beyond the scope of this deployment guide.

**Note:** Refer to the SnapCenter documentation for more information or the application-specific white papers and technical reports for detailed information about how to deploy SnapCenter for a specific application configuration:

- SnapCenter documentation: https://docs.netapp.com/us-en/snapcenter/index.html
- Deploy FlexPod Datacenter for Microsoft SQL Server 2019 with VMware 7.0 on Cisco UCS B200 M6 and NetApp ONTAP 9.8: Deploy FlexPod Datacenter for Microsoft SQL Server 2019 with VMware 7.0 on Cisco UCS B200 M6 and NetApp ONTAP 9.8
- SnapCenter Plug-In for VMware vSphere documentation: SnapCenter Plug-In for VMware vSphere documentation (netapp.com)

#### Host and privilege requirements for the SnapCenter Plug-In for VMware vSphere

Review the following requirements before installing the SnapCenter Plug-In for VMware vSphere virtual appliance refer also to Table 12):

- SnapCenter Plug-In for VMware vSphere is deployed as a Linux-based virtual appliance.
- Virtual appliance must not be deployed in a folder name with special characters.
- A separate, unique instance of the virtual appliance must be deployed for each vCenter Server.

**Table 12.** Port requirements

| Port | Requirement |
|---|---|
| 8080(HTTPS) bidirectional | This port is used to manage the virtual appliance. |
| 8144(HTTPs) bidirectional | This port used for communication between SnapCenter Plug-In for VMware vSphere and vCenter |
| 443 (HTTPS) | This port used for communication between SnapCenter Plug-In for VMware vSphere and vCenter |

### License requirements for SnapCenter Plug-In for VMware vSphere

The licenses listed in Table 13 are required on the NetApp ONTAP storage system to back up and restore virtual machines in the virtual infrastructure:

**Table 13.** SnapCenter Plug-In for VMware vSphere License requirements

| Product | License requirements |
|---|---|
| NetApp ONTAP | **SnapManager Suite:** Used for backup operations<br>One of these: SnapMirror or SnapVault (for secondary data protection regardless of the type of relationship) |
| NetApp ONTAP primary destinations | To perform protection of VMware VMs and datastores, the following licenses should be installed:<br>**SnapRestore**: used for restoring operations<br>**FlexClone**: used for mount and attach operations |
| NetApp ONTAP secondary destinations | To perform protection of VMware VMs and datastores only:<br>**FlexClone**: used for mount and attach operations |
| VMware | **vSphere Standard, Enterprise, or Enterprise Plus**<br>A vSphere license is required to perform restore operations, which use Storage vMotion. vSphere Essentials or Essentials Plus licenses do not include Storage vMotion. |

**Note:** It is recommended (but not required) to add SnapCenter Standard licenses to secondary destinations. If SnapCenter Standard licenses are not enabled on secondary systems, you cannot use SnapCenter after a failover operation. A FlexClone license on secondary storage is required to perform mount and attach operations. A SnapRestore license is required to perform restore operations.

### Procedure 1. Manually deploy the SnapCenter Plug-In for VMware vSphere 4.7

1. Download the SnapCenter Plug-In for VMware vSphere OVA file from the NetApp support site (https://mysupport.netapp.com).
2. From VMware vCenter, navigate to the **VMs and Templates** tab, right-click the datacenter (for example, FlexPod-DC) and select **Deploy OVF Template**.
3. Specify the location of the OVF template and click **NEXT**.
4. On the Select a name and folder page, enter a unique name (for example, aa02-scv) and location (datacenter for example, FlexPod-DC) for the VM and click **NEXT** to continue.

5. On the Select a compute resource page, select the cluster, and click **NEXT**.
6. On the Review details page, verify the OVA template details and click **NEXT**.
7. On the License agreements page, read and check the box **I accept all license agreements**. Click **NEXT**.
8. On the Select storage page, select a datastore, change the datastore virtual disk format to **Thin Provision** and click **NEXT**.

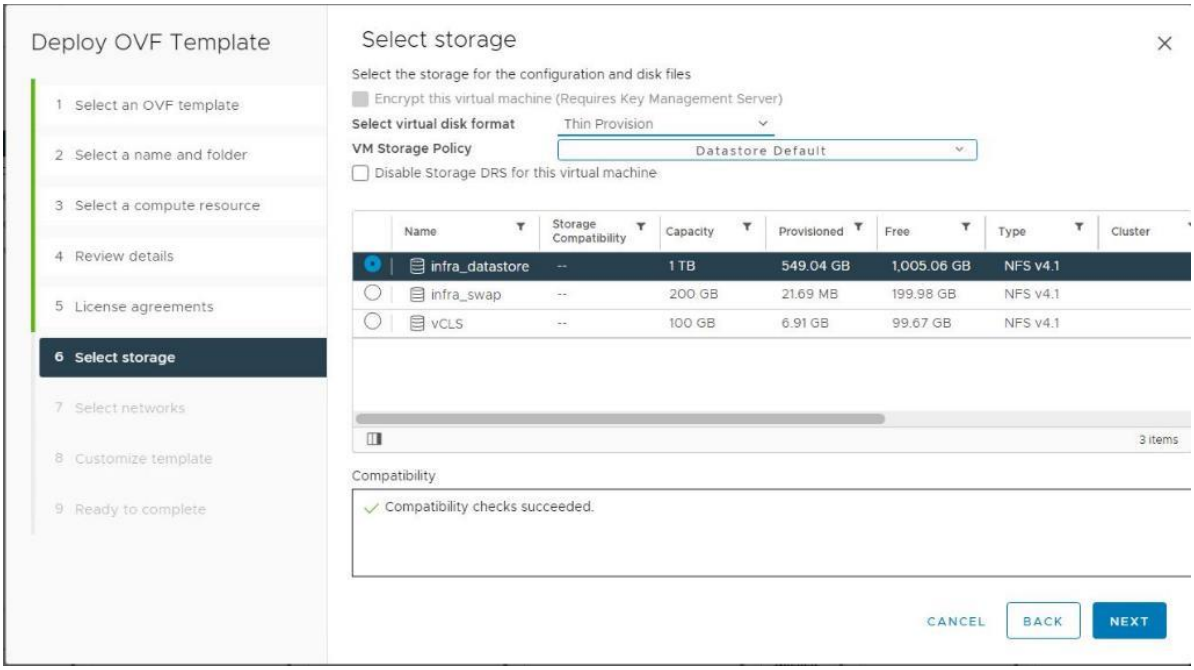


9. On the Select networks page, select a destination network, for example, IB-MGMT, and then click **NEXT**.
10. On the Customize template page, under Register to existing vCenter, enter the vCenter credentials.
11. In Create SCV credentials, create a username (for example, admin) and password.
12. In System Configuration, enter the maintenance user password.
13. In Setup Network Properties, enter the network information.
14. In Setup Date and Time, provide the NTP server address(es) and select the time zone where the vCenter is located.
15. Click **NEXT**.
16. On the Ready to complete page, review the page and click **FINISH**. The VM deployment will start. After the VM is deployed successfully, proceed to the next step.
17. Navigate to the SnapCenter VM, right-click, and select **Power > Power On** to start the virtual appliance.
18. While the virtual appliance is powering on, click **Install VMware tools**.
19. After the SnapCenter VM installation is complete and VM is ready to use, proceed to the next step.
20. Log into SnapCenter Plug-In for VMware vSphere using the IP address (https://<ip_address_of_SnapCenter>:8080 ) displayed on the appliance console screen with the credentials that you provided in the deployment wizard.

21. Verify on the Dashboard that the virtual appliance has successfully connected to vCenter and the SnapCenter Plug-In for VMware vSphere is successfully enabled and connected.

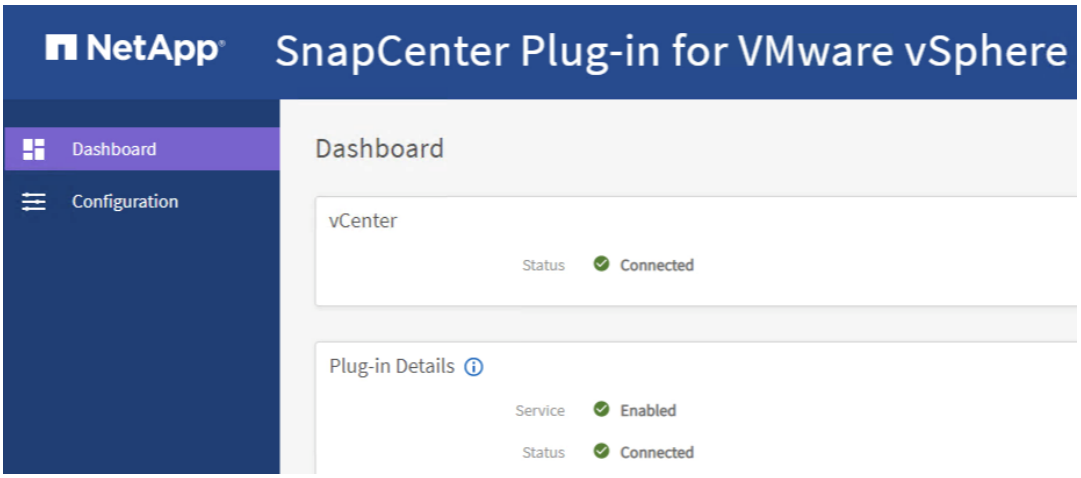


## NetApp SnapCenter Plug-In 4.7 configuration

### Procedure 1. SnapCenter Plug-In for VMware vSphere in vCenter Server

1. Navigate to VMware vSphere Web Client URL https://<vCenter Server>.
2. If you're currently logged into vCenter, log off, close the open tab, and sign on again to access the newly installed SnapCenter Plug-In for VMware vSphere.
3. After logging on, a blue banner will be displayed, indicating the SnapCenter plug-In was successfully deployed. Click **Refresh** to activate the plug-in.
4. On the VMware vSphere Web Client page, select **Menu > SnapCenter Plug-In for VMware vSphere** to launch the SnapCenter Plug-In for VMware GUI.

### Procedure 2. Add Storage System

1. Click Storage Systems.

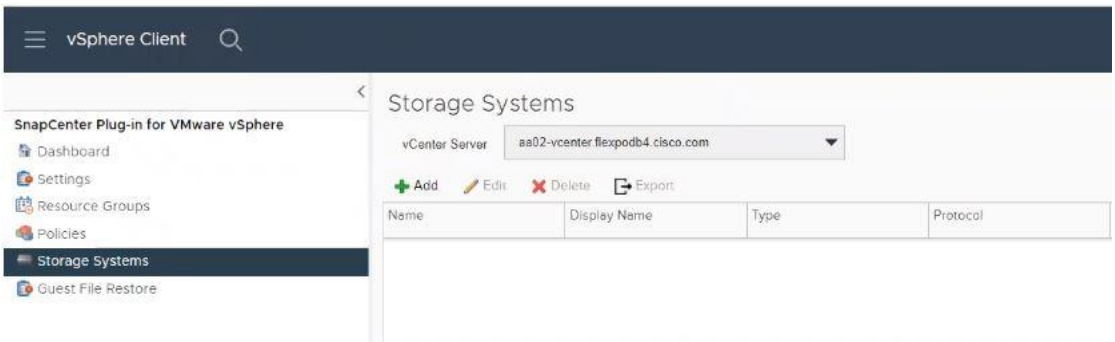


2. Click **+Add** to add a storage system (or SVM).
3. Enter Storage System, user credentials, and other required information in the following dialog box.
4. Check the box for Log SnapCenter server events to syslog and Send AutoSupport Notification for failed operation to storage system.

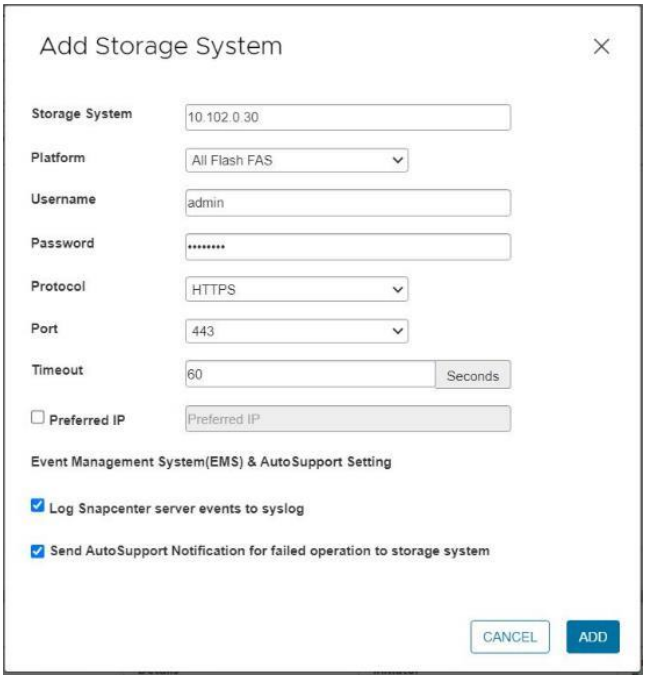


5. Click ADD.

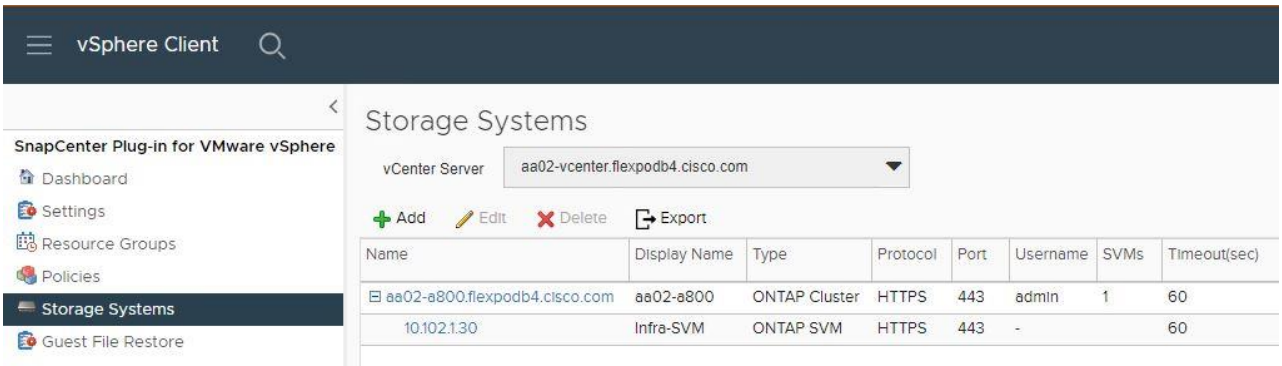


When the storage system is added, you can create backup policies and take scheduled backup of virtual machines and datastores. The SnapCenter Plug-In for VMware vSphere allows backup, restore, and on-demand backups.

For more information about backup policy configuration, refer to this CVD: [FlexPod Datacenter with Cisco UCS X-Series, VMware 7.0 U2, and NetApp ONTAP 9.9 – Management Tools Setup]

## Active IQ Unified Manager 9.11P1 installation

Active IQ Unified Manager enables you to monitor and manage the health and performance of NetApp ONTAP storage systems and virtual infrastructure from a single interface. Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems. Active IQ Unified Manager is required to integrate NetApp storage with the Cisco Intersight software.

This section describes the steps to deploy NetApp Active IQ Unified Manager 9.11P1 as a virtual appliance. Table 14 lists the recommended configuration for the virtual machine.

**Table 14.** Virtual machine configuration

| Hardware configuration | Recommended settings |
|---|---|
| RAM | 12 GB |
| Processors | 4 CPUs |
| CPU cycle capacity | 9572 MHz total |
| Free disk space/virtual disk size | 5 GB – Thin provisioned<br>152 GB – Thick provisioned |

**Note:** There is a limit to the number of nodes that a single instance of Active IQ Unified Manager can monitor before you need a second instance of Active IQ Unified Manager. Refer to the Unified Manager Best Practices Guide (TR-4621) for more details.

**Procedure 1. Install NetApp Active IQ Unified Manager 9.12 manually**

1. Download the NetApp Active IQ Unified Manager for VMware vSphere OVA file from: Active IQ Unified Manager – 9.12
2. In the VMware vCenter GUI, click **VMs and Templates** and then click **Actions> Deploy OVF Template**.
3. Specify the location of the OVF Template and click **NEXT**.
4. On the Select a name and folder page, enter a unique name for the VM, select a deployment location, and then click **NEXT**.
5. On the Select a compute resource screen, select the cluster where VM will be deployed and click **NEXT**.
6. On the Review details page, verify the OVA template details and click **NEXT**.

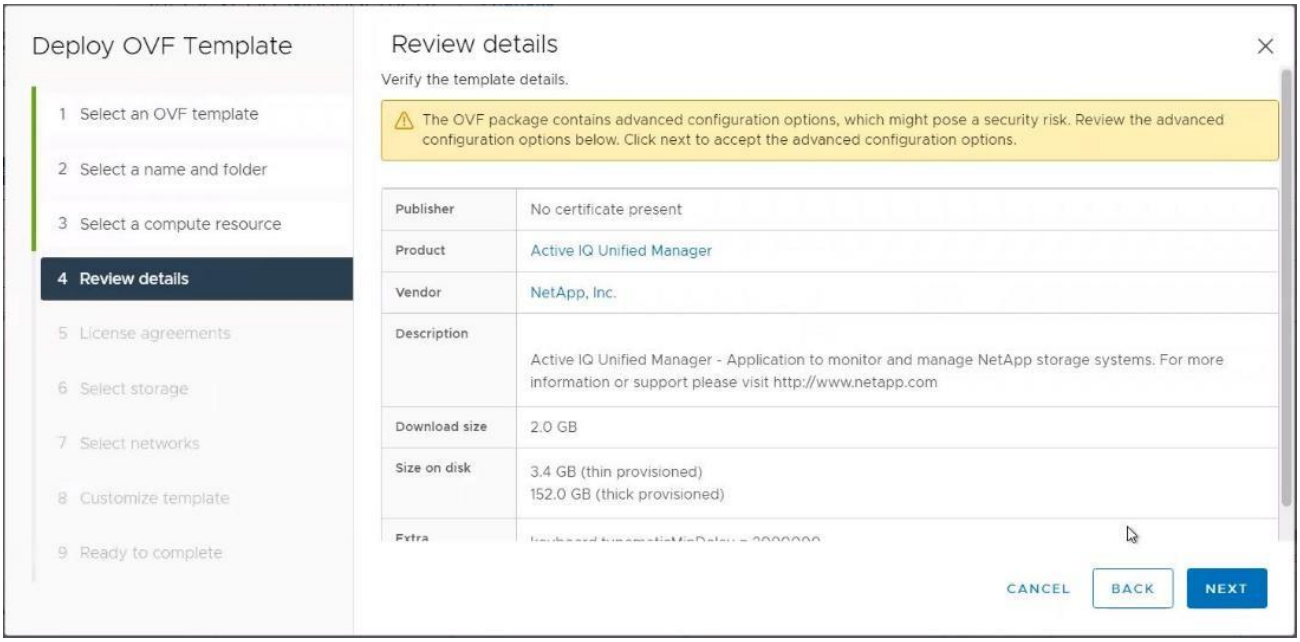


7. On the License agreements page, read and check the box for I accept all license agreements. Click **NEXT**.

8. On the Select storage page, select the following parameters for the VM deployment:
   - Select the disk format for the Virtual Machine Disk (VMDK) (for example, Thin Provisioning).
   - Select a VM Storage Policy (for example, Datastore Default).
   - Select a datastore to store the deployed OVA template.

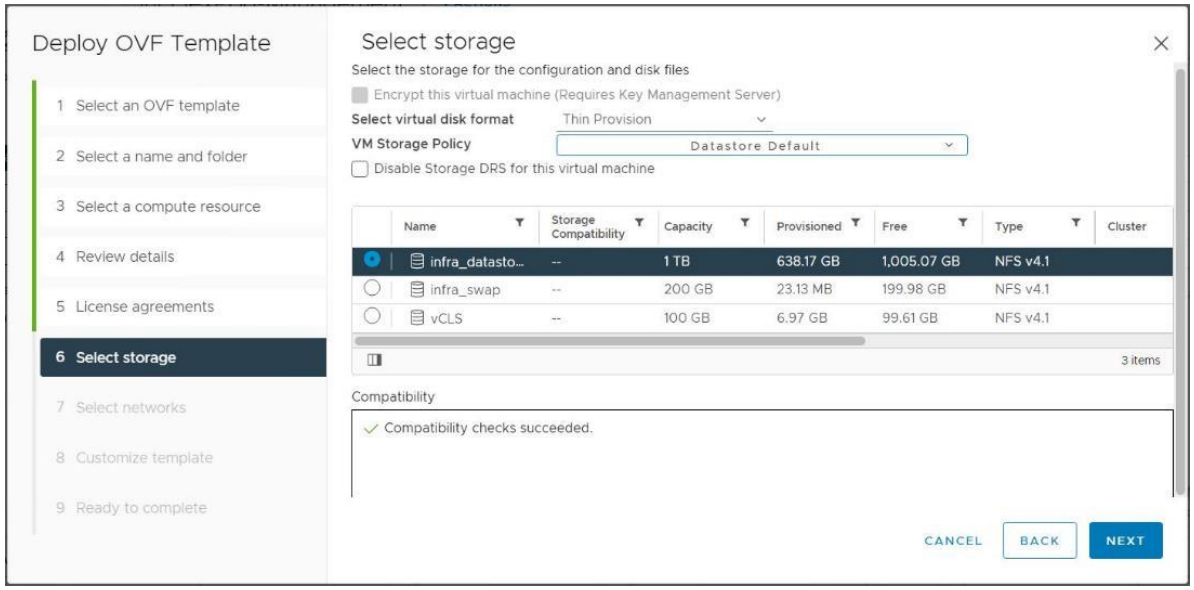


9. Click **NEXT**.
10. On the Select networks page, select the destination network (for example, IB-MGMT) and click **NEXT**.
11. On the Customize template page, provide network details such as hostname, IP address, gateway, and DNS.

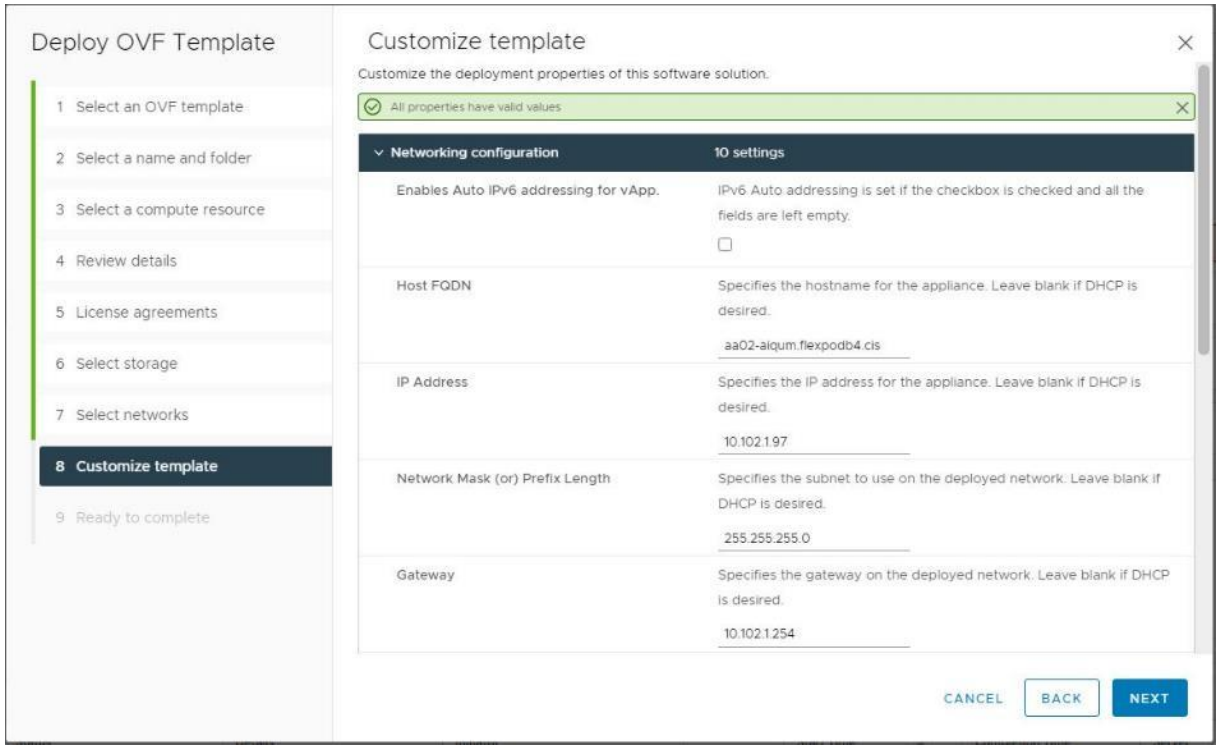

12. Leave the TimeZone value field blank but enter the Maintenance username and password.

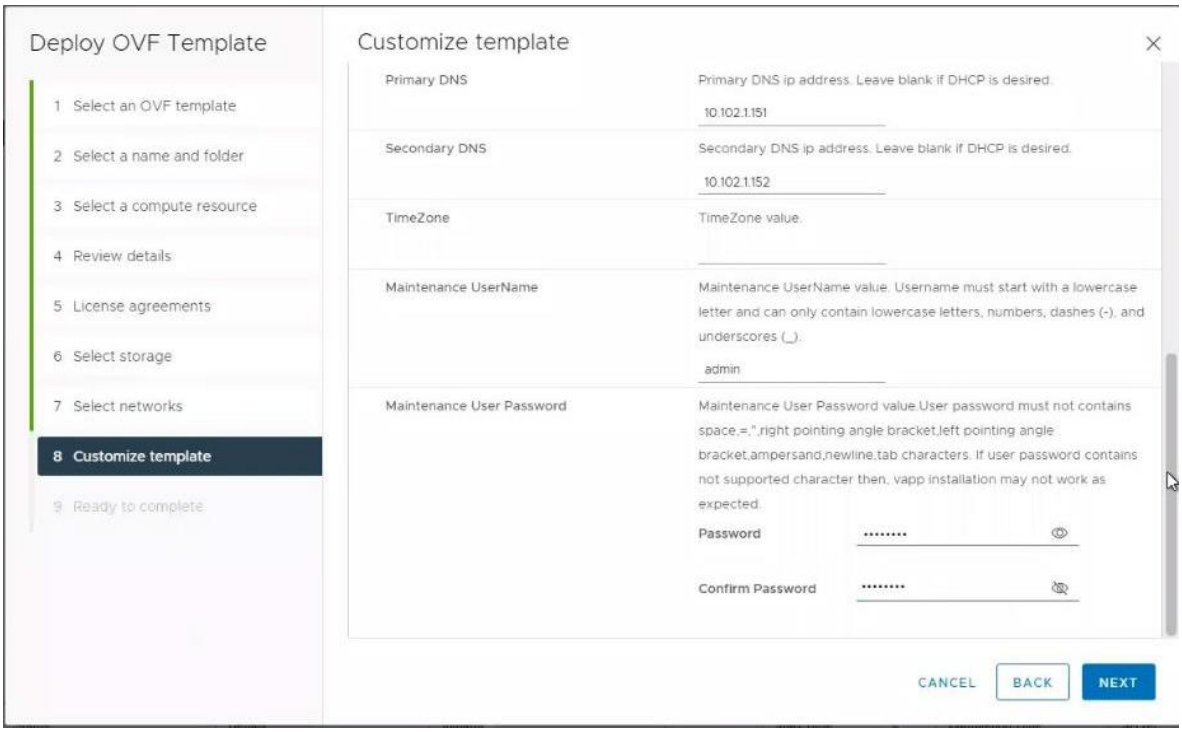


**Note:** Save the maintenance user account credentials in a secure location. You will use these credentials for the initial GUI login and to make any configuration changes to the appliance settings in the future.

13. Click **NEXT**.
14. On the Ready to complete page, review the settings and click **FINISH**. Wait for the VM deployment to complete before proceeding to the next step.
15. Select the newly created Active IQ Unified Manager VM, then right-click and select **Power > Power On**.
16. While the virtual machine is powering on, click the prompt in the yellow banner to **Install VMware tools**.

**Note:** Because of timing, VMware tools might not install correctly. In that case, you can manually install VMware tools after Active IQ Unified Manager VM is up and running.

17. Open the VM console for the Active IQ Unified Manager VM and configure the time zone information when displayed.

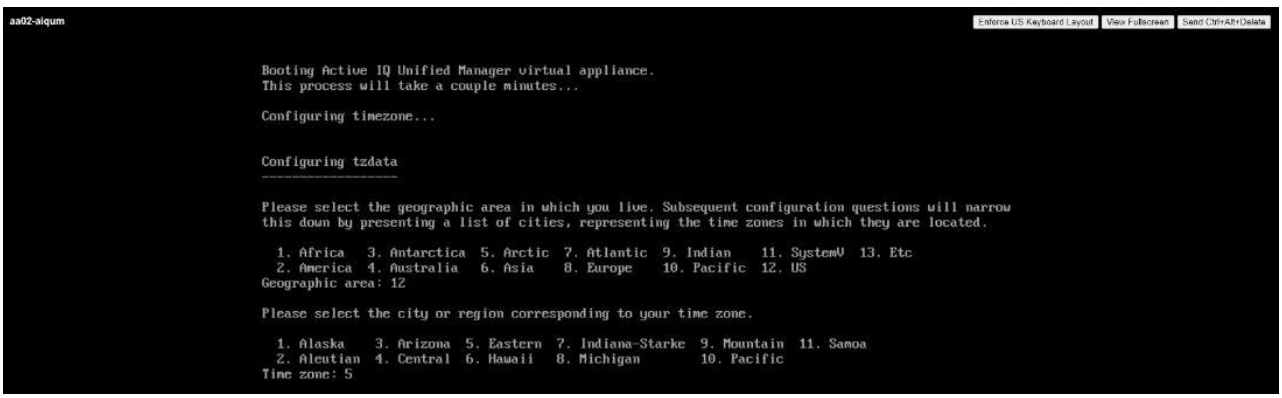


18. Wait for the AIQM web console to display the login prompt.

```
Active IQ Unified Manager



Log in to Active IQ Unified Manager in a web browser by using

    https://10.102.1.97/

or
    https://aa02-aiqum.flexpodb4.cisco.com/

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

aa02-aiqum login:
```

19. Log into NetApp Active IQ Unified Manager using the IP address or URL displayed on the web console.

## Configure Active IQ Unified Manager

**Procedure 1. Initial setup**

1. Launch a web browser and log into Active IQ Unified Manger using the URL shown in the VM console.
2. Enter the email address that Unified Manager will use to send alerts and the mail server configuration. Click **Continue**.
3. Select **Agree and Continue** on the Set up AutoSupport configuration.
4. Check the box for **Enable API Gateway** and click **Continue.**

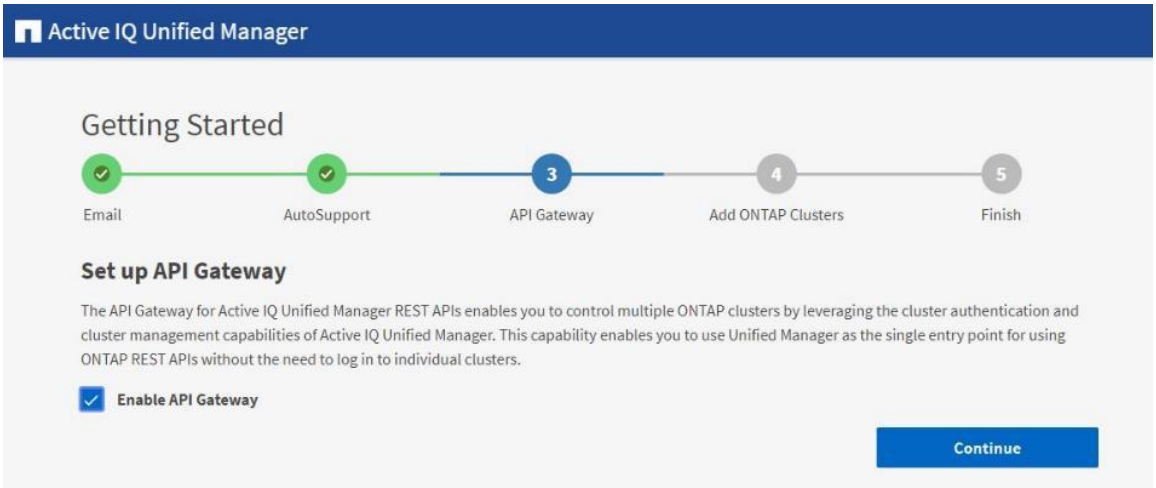


5. Enter the NetApp ONTAP cluster hostname or IP address and the admin login credentials.

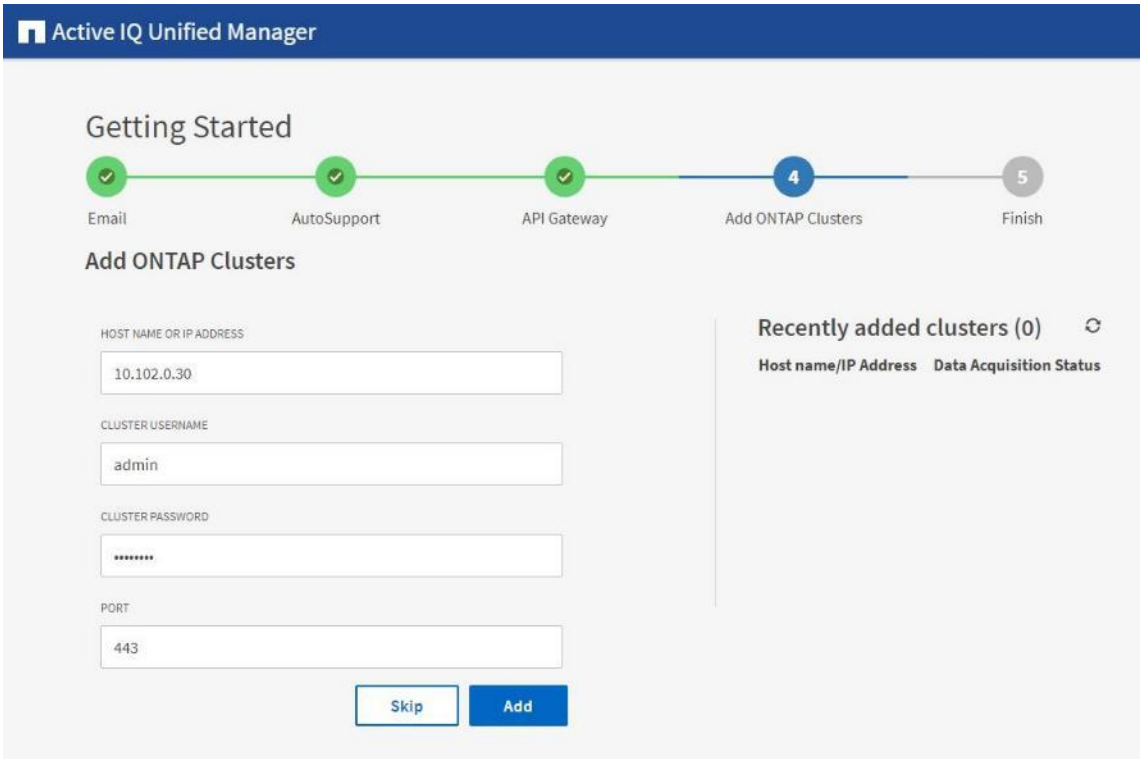


6. Click **Add**.
7. Click **Yes** to trust the self-signed cluster certificate and finish adding the storage system.

**Note:** The initial discovery process can take up to 15 minutes to complete.

### Procedure 2. Review security compliance with Active IQ Unified Manager

Active IQ Unified Manager identifies problems and makes recommendations to improve the security posture of NetApp ONTAP. Active IQ Unified Manager evaluates NetApp ONTAP storage based on recommendations made in the Security Hardening Guide for NetApp ONTAP 9. Items are identified according to their level of compliance with the recommendations. Review the Security Hardening Guide for NetApp ONTAP 9 (TR-4569) for additional information and recommendations for securing NetApp ONTAP 9.

**Note:** All events identified do not inherently apply to all environments, for example, Federal Information Processing Standard (FIPS) compliance.

The status icons in the security cards have the following meanings in relation to their compliance:

- - The parameter is configured as recommended.
- - The parameter is not configured as recommended.
- - Either the functionality is not enabled on the cluster, or the parameter is not configured as recommended, but this parameter does not contribute to the compliance of the object.

Note that volume encryption status does not contribute to whether the cluster or SVM are considered compliant.

8. Navigate to the URL of the Active IQ Unified Manager and login.
9. Select the **Dashboard** from the left menu bar in Active IQ Unified Manager.

10. Locate the **Security** card and note the compliance level of the cluster and SVM.

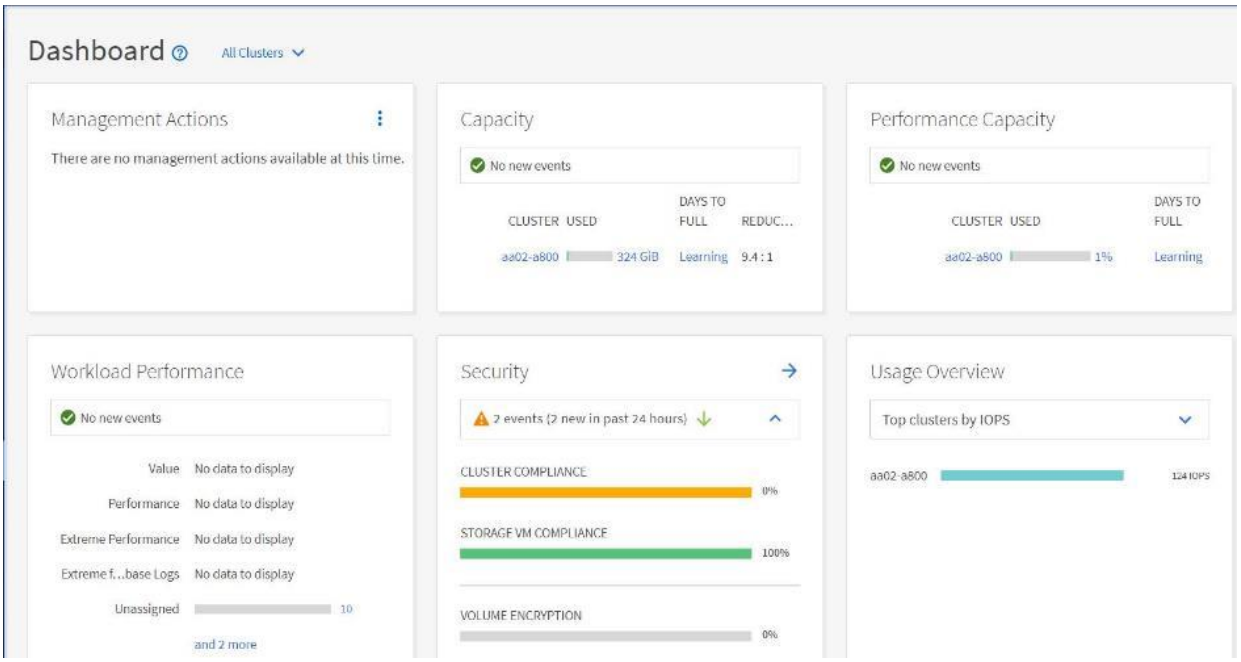


11. Click the blue arrow to expand the findings.

12. Locate the Individual Cluster section and the Cluster Compliance card. From the drop-down list select **View All**.

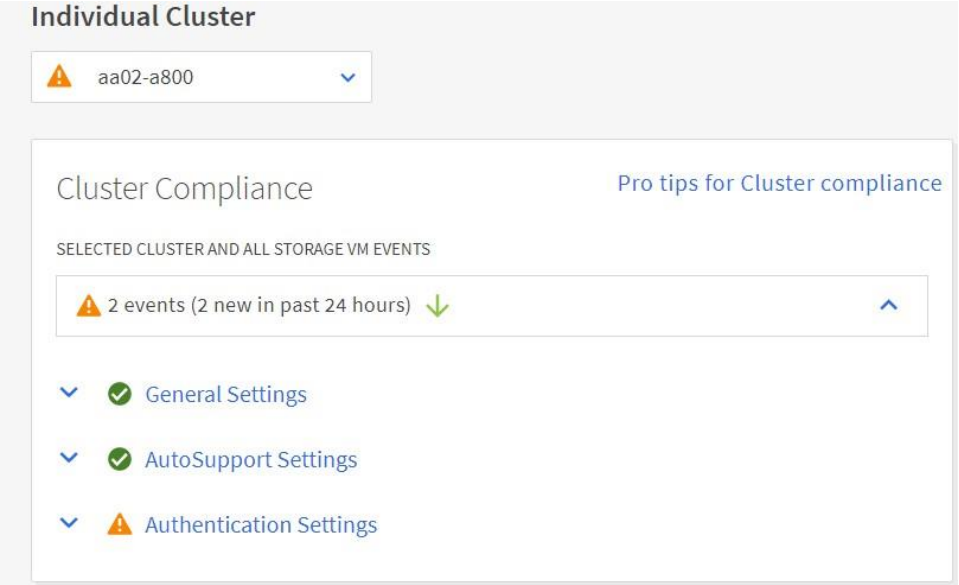


13. Select an event from the list and click the name of the event to view the remediation steps.

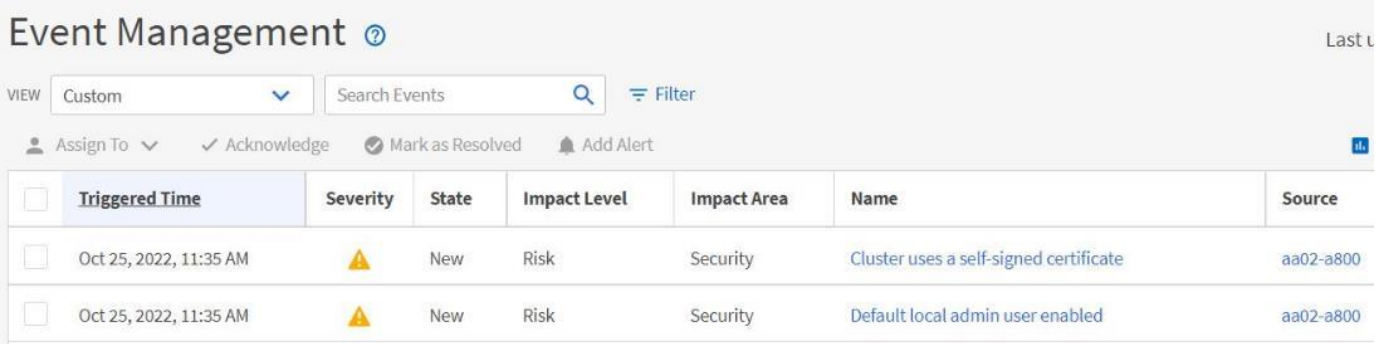


14. Remediate the risk if applicable to the current environment and perform the suggested actions to fix the problem.

### Remediate security compliance findings

**Note:** Active IQ identifies several security compliance risks after installation that you can correct immediately to improve the security posture of NetApp ONTAP. Click the event name to get more information and suggested actions to fix the problem.

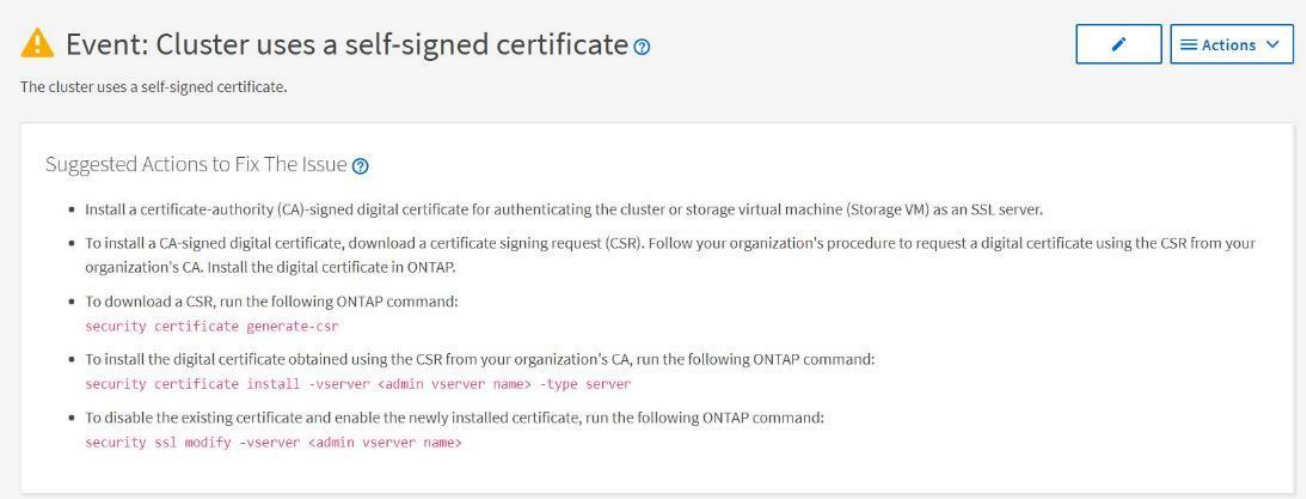


## Claim VMware vCenter using Cisco Intersight Assist Appliance

### Procedure 1. Claim the vCenter from the Cisco Intersight software

1. Log into **Cisco Intersight** and connect to the account for this FlexPod.
2. Select System > Administration > Targets and click Claim a New Target.
3. Under Select Target Type, select **VMware vCenter** under Hypervisor and click **Start**.
4. In the **VMware vCenter** window, verify that the correct Intersight Assist is selected.
5. Fill in the vCenter information. If you use Intersight Workflow Optimizer (IWO), turn on Datastore Browsing Enabled and Guest Metrics Enabled. If you want to use Hardware Support Manager (HSM) to be able to upgrade Intersight Management Mode server firmware from VMware Lifecycle Manager, turn on HSM. Click **Claim**.

**Note:** It is recommended to use an admin-level user other than administrator@vsphere.local to claim VMware vCenter to the Cisco Intersight application. The administrator@vsphere.local user has visibility to the vSphere Cluster Services (vCLS) virtual machines. These virtual machines would then be visible in the Cisco Intersight application and its operations could be executed on them. VMware does not recommend users executing operations on these virtual machines. With a user other than administrator@vsphere.local, the vCLS virtual machines would be inaccessible from the Cisco Intersight software.

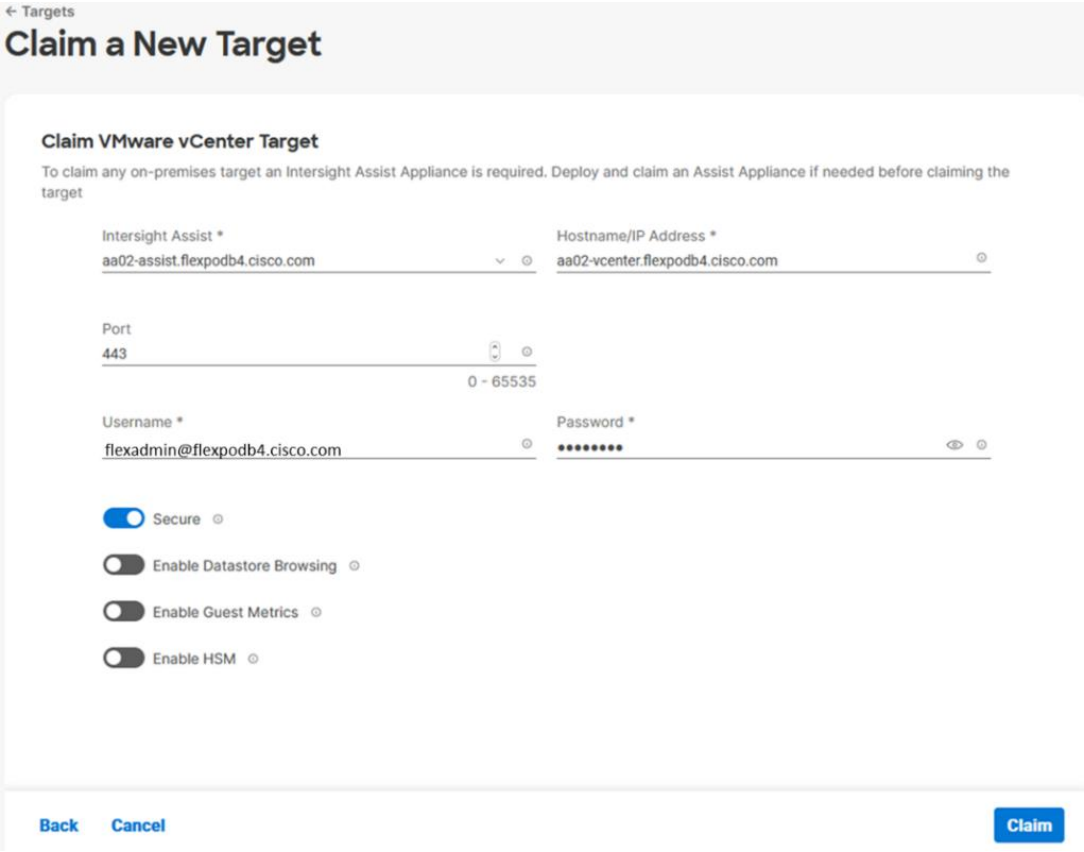


6. After a few minutes, the VMware vCenter will show Connected in the Targets list and will also appear under **Infrastructure Service > Operate > Virtualization**.
7. You now can view detailed information obtained from the vCenter by clicking **Infrastructure Service > Operate > Virtualization** and selecting the Datacenters tab. You can obtain other VMware vCenter information by navigating through the Virtualization tabs.

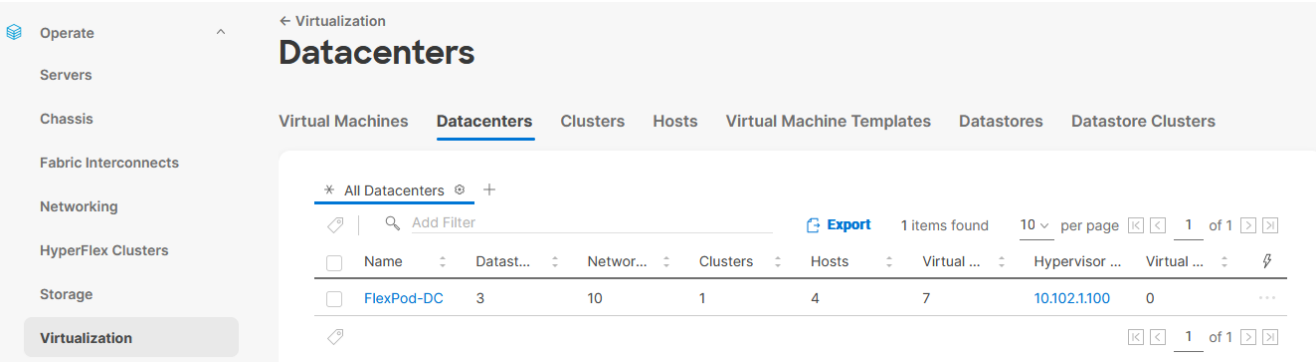

### Procedure 2. Interact with virtual machines

VMware vCenter integration with Cisco Intersight software allows you to interact directly with the virtual machines from the Cisco Intersight dashboard. In addition to obtaining in-depth information about a virtual machine, including the operating system, CPU, memory, host name, and IP addresses assigned to the virtual machines, you can use the application to perform the following actions on the virtual machines:

- **Start/Resume**
- **Stop**
- **Soft Stop**
- **Suspend**
- **Reset**
- **Launch VM Console**

1. Log into **Cisco Intersight** and connect to the account for this FlexPod.
2. Select Infrastructure Service > Operate > Virtualization.
3. Click the Virtual Machines tab.
4. Click "**...**" to the right of a VM and interact with various VM options.

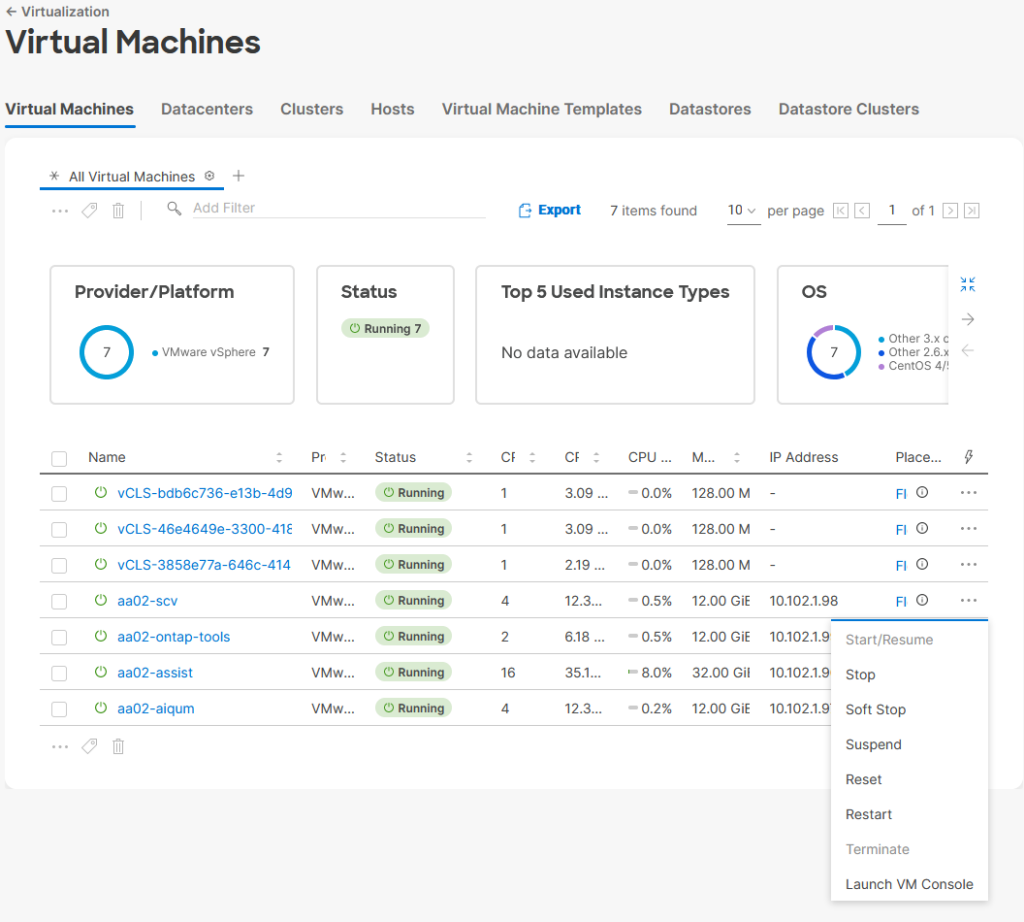

5. To gather more information about a VM, click a VM name. The same interactive options are available under **Actions**.

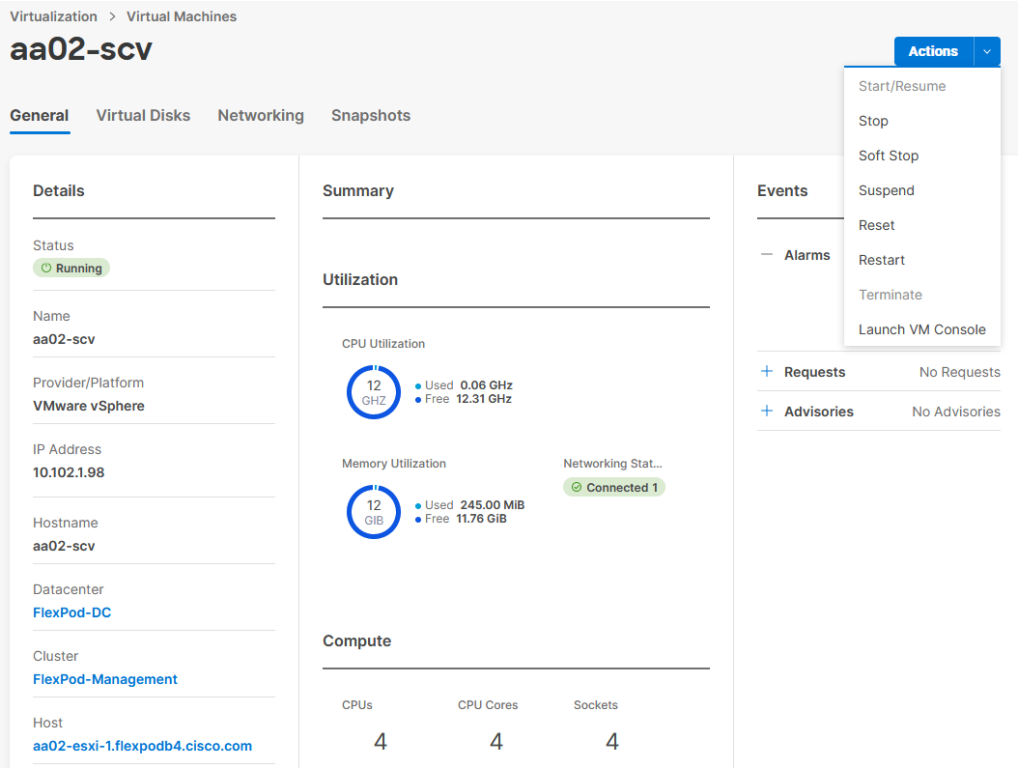


## Claim NetApp Active IQ Manager using Cisco Intersight Assist Appliance

### Procedure 1. Claim NetApp Active IQ Unified Manager into Cisco Intersight using Ansible

1. Clone the repository from https://github.com/NetApp-Automation/NetApp-AIQUM.
2. Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.
3. Update the variable files as mentioned in the README document in the repository.
4. To claim an existing AIQUM instance into the Intersight application, invoke the below ansible playbook:

```
ansible-playbook aiqum.yml -t intersight_claim
```

### Procedure 2. Manually claim the NetApp Active IQ Unified Manager into Cisco Intersight application

1. Log into **Cisco Intersight** and connect to the account for this FlexPod.
2. From Cisco Intersight, dashboard click **System > Administration > Targets**.
3. Click **Claim a New Target**. In the Select Target Type window, select NetApp Active IQ Unified Manager under Storage and click **Start**.
4. In the Claim NetApp Active IQ Unified Manager Target window, verify the correct Cisco Intersight Assist is selected.
5. Fill in the NetApp Active IQ Unified Manager information and click **Claim**.

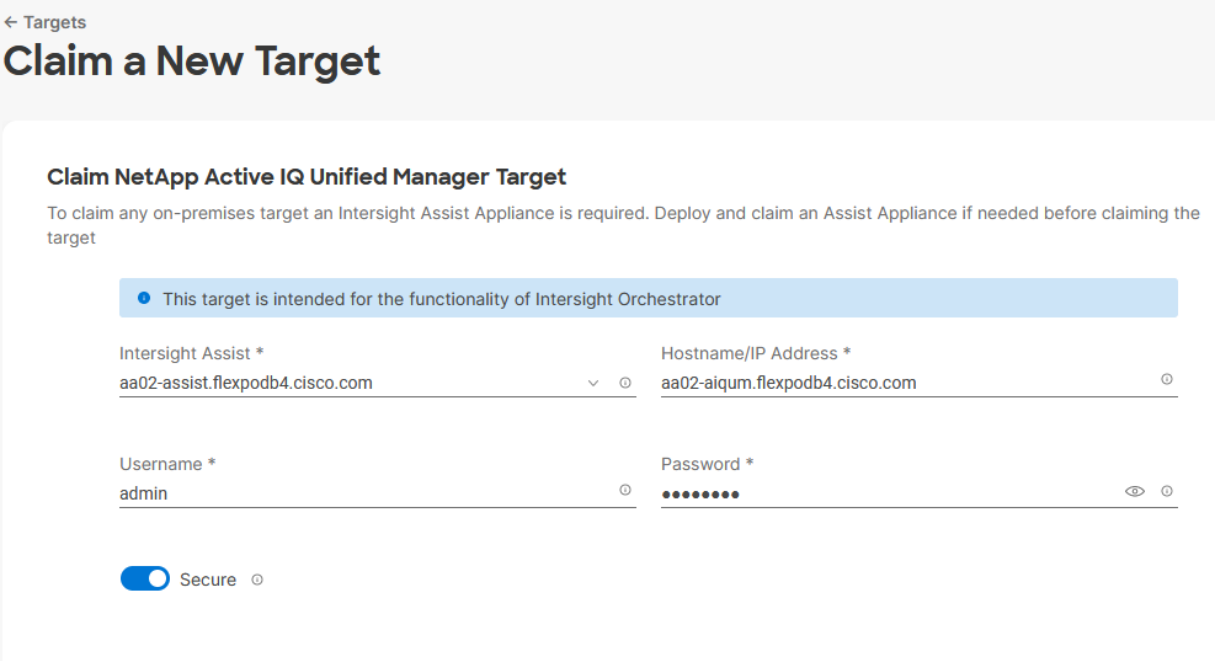


6. After a few minutes, the NetApp ONTAP Storage configured in the Active IQ Unified Manager will appear under the **Infrastructure Service > Operate > Storage** tab.

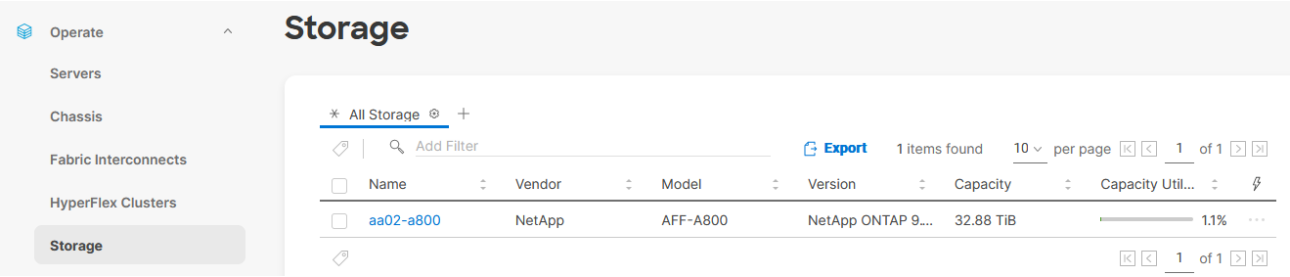


7. Click the storage cluster name to see detailed General, Inventory, and Checks information about the storage.

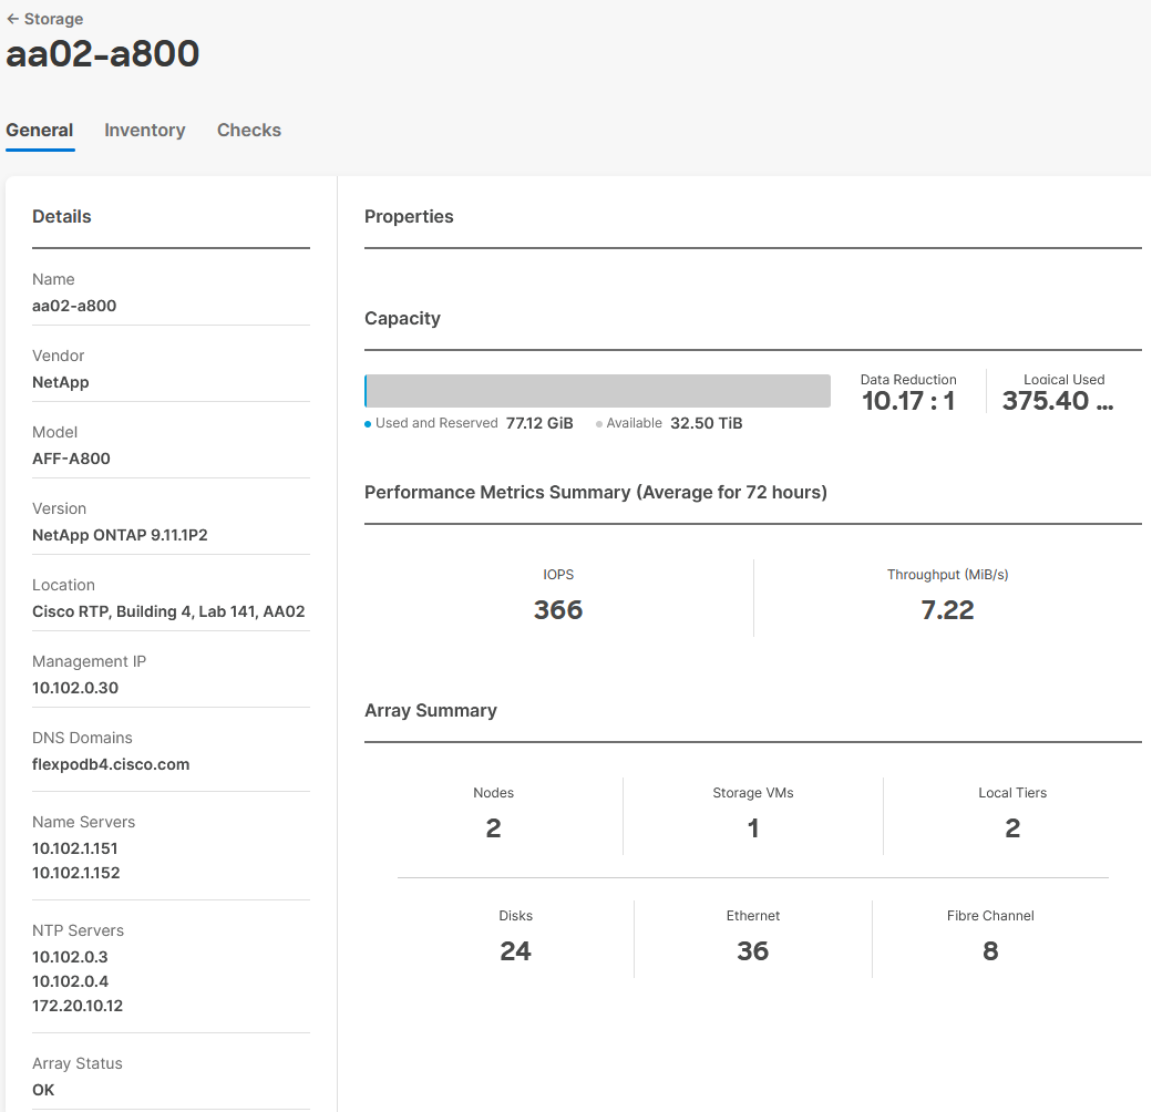


8. Click **My Dashboard > Storage** to see storage monitoring widgets.

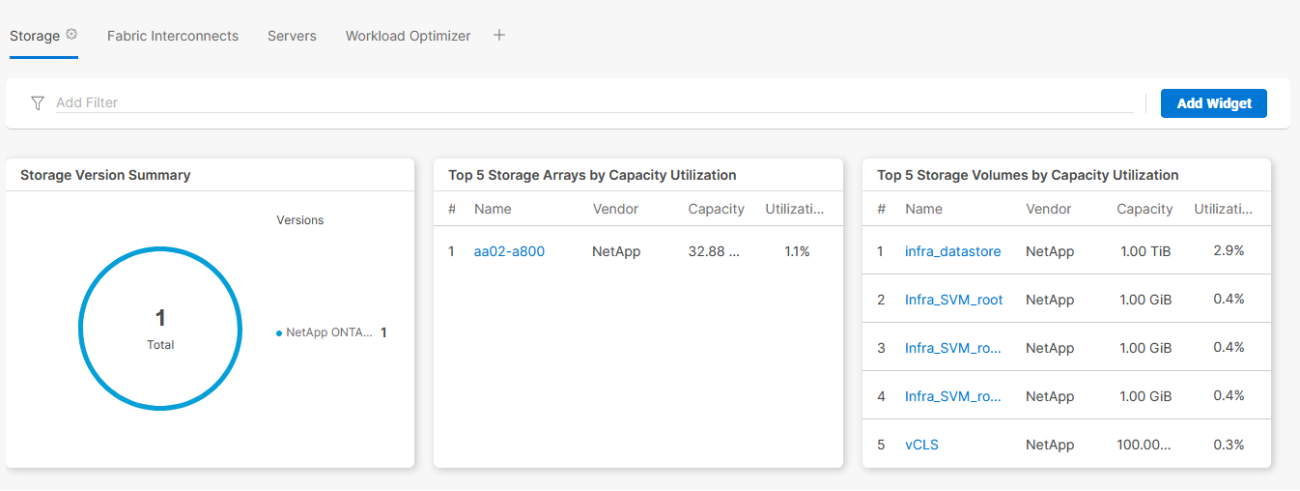

## Claim Cisco Nexus Switches using Cisco Intersight Assist Appliance

### Procedure 1. Claim Cisco Nexus Switches

1. Log into **Cisco Intersight** and connect to the account for this FlexPod.
2. From the Cisco Intersight application, click **System > Administration > Targets**.
3. Click **Claim a New Target**. In the Select Target Type window, select Cisco Nexus Switch under Network, and click **Start**.
4. In the Claim Cisco Nexus Switch Target window, verify the correct Intersight Assist is selected.
5. Fill in the Cisco Nexus Switch information and click **Claim**.

**Note:** You can use the admin user on the switch.

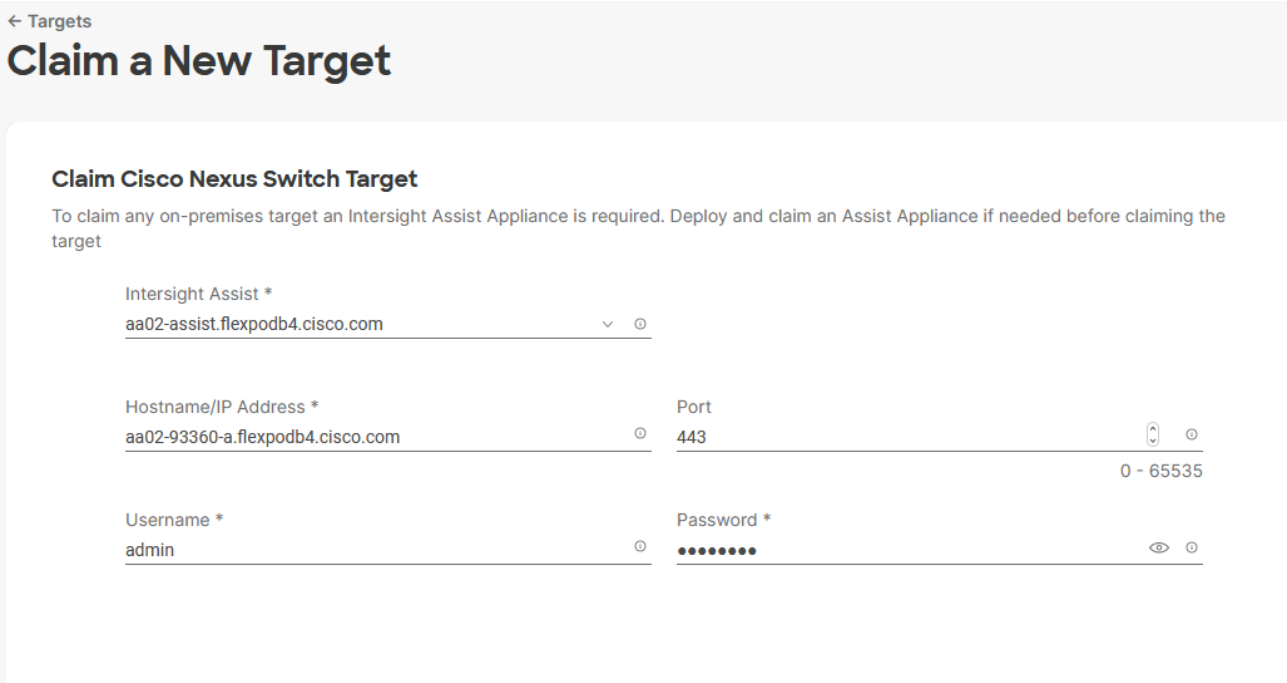


6. Follow the steps in this procedure to add the second Cisco Nexus switch.
7. After a few minutes, the two switches will appear under Infrastructure Service > Operate > Networking > Ethernet Switches.

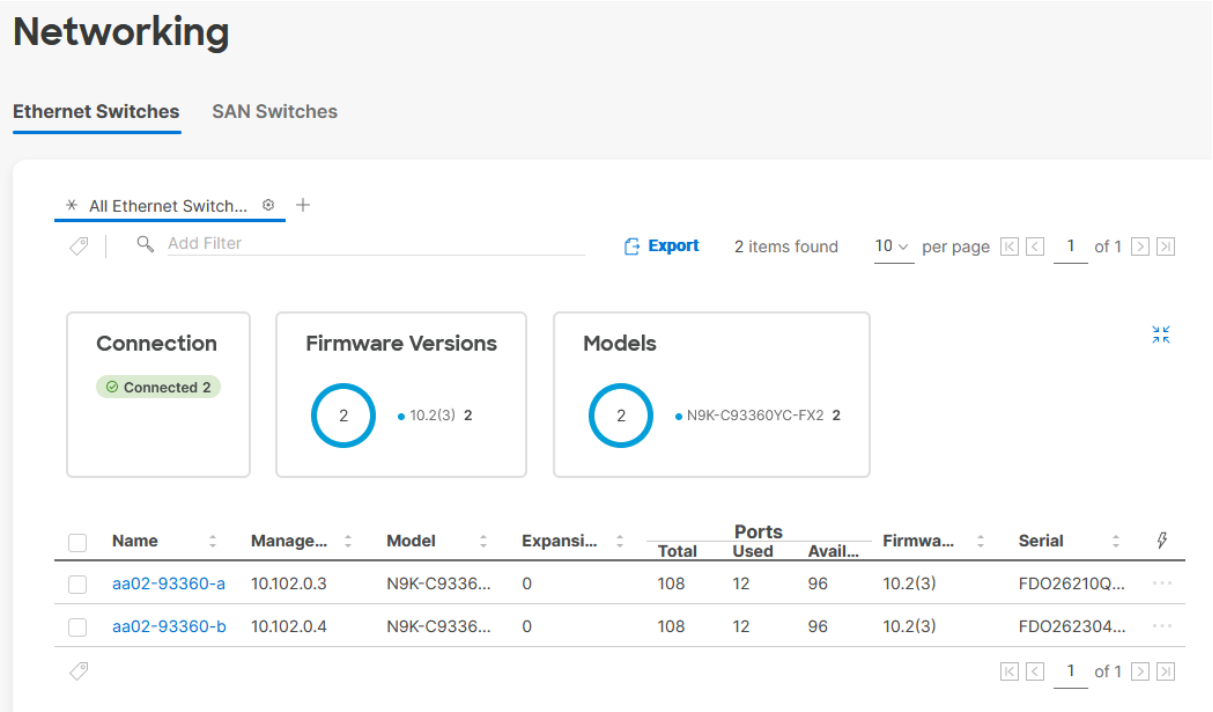


8. Click one of the switch names to get detailed General and Inventory information on the switch.

## Claim Cisco MDS Switches using Cisco Intersight Assist Appliance

### Procedure 1. Claim Cisco MDS Switches (if they are part of the FlexPod)

1. Log into the **Cisco Intersight application** and connect to the account for this FlexPod.
2. From the dashboard, click System > Administration > Targets.
3. Click **Claim a New Target**. In the Select Target Type window, select Cisco MDS Switch under Network and click **Start**.
4. In the Claim Cisco MDS Switch Target window, verify the correct Intersight Assist is selected.
5. Fill in the Cisco MDS Switch information including the use of port 8443 and click **Claim**.

**Note:** You can use the admin user on the switch.

← Targets

# Claim a New Target

**Claim Cisco MDS Switch Target**

To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

Intersight Assist *
aa02-assist.flexpodb4.cisco.com

Hostname/IP Address *
aa02-9132t-a.flexpodb4.cisco.com

Port
8443
0 - 65535

Username *
admin

Password *
••••••••

6. Follow the steps in this procedure to add the second Cisco MDS switch.
7. After a few minutes, the two switches will appear under Infrastructure Service > Operate > Networking > SAN Switches.

# Networking

Ethernet Switches | **SAN Switches**

All SAN Switches

Add Filter | Export | 2 items found | 10 per page | 1 of 1

**Connection**
Connected 2

**Firmware Versions**
2 • 9.2(2) 2

**Models**
2 • DS-C9132T-K9 2

| Name | Contract Status | Manag... | Model | Expans... | Ports Total | Ports Used | Ports Avail... | Firmw... |
|---|---|---|---|---|---|---|---|---|
| aa02-9132t-a | - | 10.102.0.7 | DS-C9132T... | 0 | 16 | 12 | 4 | 9.2(2) |
| aa02-9132t-b | - | 10.102.0.8 | DS-C9132T... | 0 | 16 | 12 | 4 | 9.2(2) |

1 of 1

8. Click one of the switch names to get detailed General and Inventory information about the switch.

## About the Authors

**John George, Technical Marketing Engineer, Cisco Systems, Inc.**

John has been involved in designing, developing, validating, and supporting the FlexPod Converged Infrastructure since it was developed almost 12 years ago. Before his roles with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a master's degree in computer engineering from Clemson University.

**Tyson Scott, Technical Solutions Architect, Cisco Systems, Inc.**

Tyson has been working in Cisco's sales organization for more than 12 years as a Technical Solutions Architect, in both Security and Data Center solutions. Prior to joining Cisco, Tyson was a CCIE instructor teaching Route Switch, Security, and Service Provider courses. Tyson developed content to provide self-paced learning in addition to teaching the courses (CCIE® certification— 13513 – Route Switch/Security/Service Provider).

**Roney Daniel, Technical Marketing Engineer, Hybrid Cloud Infra and OEM Solutions, NetApp Inc.**

Roney Daniel is a Technical Marketing Engineer (TME), at NetApp. He has more than 25 years of experience in the networking industry. Prior to NetApp, Roney worked at Cisco Systems in various roles with Cisco Technical Assistance Center (TAC), Financial Test Lab, Systems and solution engineering business units, and Cisco IT. He has a bachelor's degree in Electronics and Communication engineering and is a data center Cisco Certified Internetwork Expert (CCIE® certification— 42731).

**Kamini Singh, Technical Marketing Engineer, Hybrid Cloud Infra and OEM Solutions, NetApp**

Kamini Singh is a Technical Marketing Engineer (TME), at NetApp. She has 3 years of experience in data center infrastructure solutions. Kamini focuses on FlexPod hybrid cloud infrastructure solution design, implementation, validation, automation, and sales enablement. She holds a bachelor's degree in Electronics and Communication and a master's degree in communication systems.


### Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Haseeb Niazi, Principal Technical Marketing Engineer, Cisco Systems, Inc.
- Paniraja Koppa, Technical Marketing Engineer, Cisco Systems, Inc.
- Lisa DeRuyter-Wawrzynski, Information Developer, Cisco Systems, Inc.

# Appendix

This appendix is organized into the following:



**Note:** The features and functions explained in this appendix are optional configurations that can be helpful in configuring and managing the FlexPod deployment.

## FlexPod with Cisco Nexus SAN Switching configuration – Part 1

When using the Cisco Nexus switches for SAN switching, you should use the following alternate base switch setup. This configuration uses 100G Fibre Channel over Ethernet (FCoE) uplinks from the Cisco UCS fabric interconnects to the Cisco Nexus switches. You also can use 25G uplinks. Figure 5 shows the validation lab cabling for this setup.

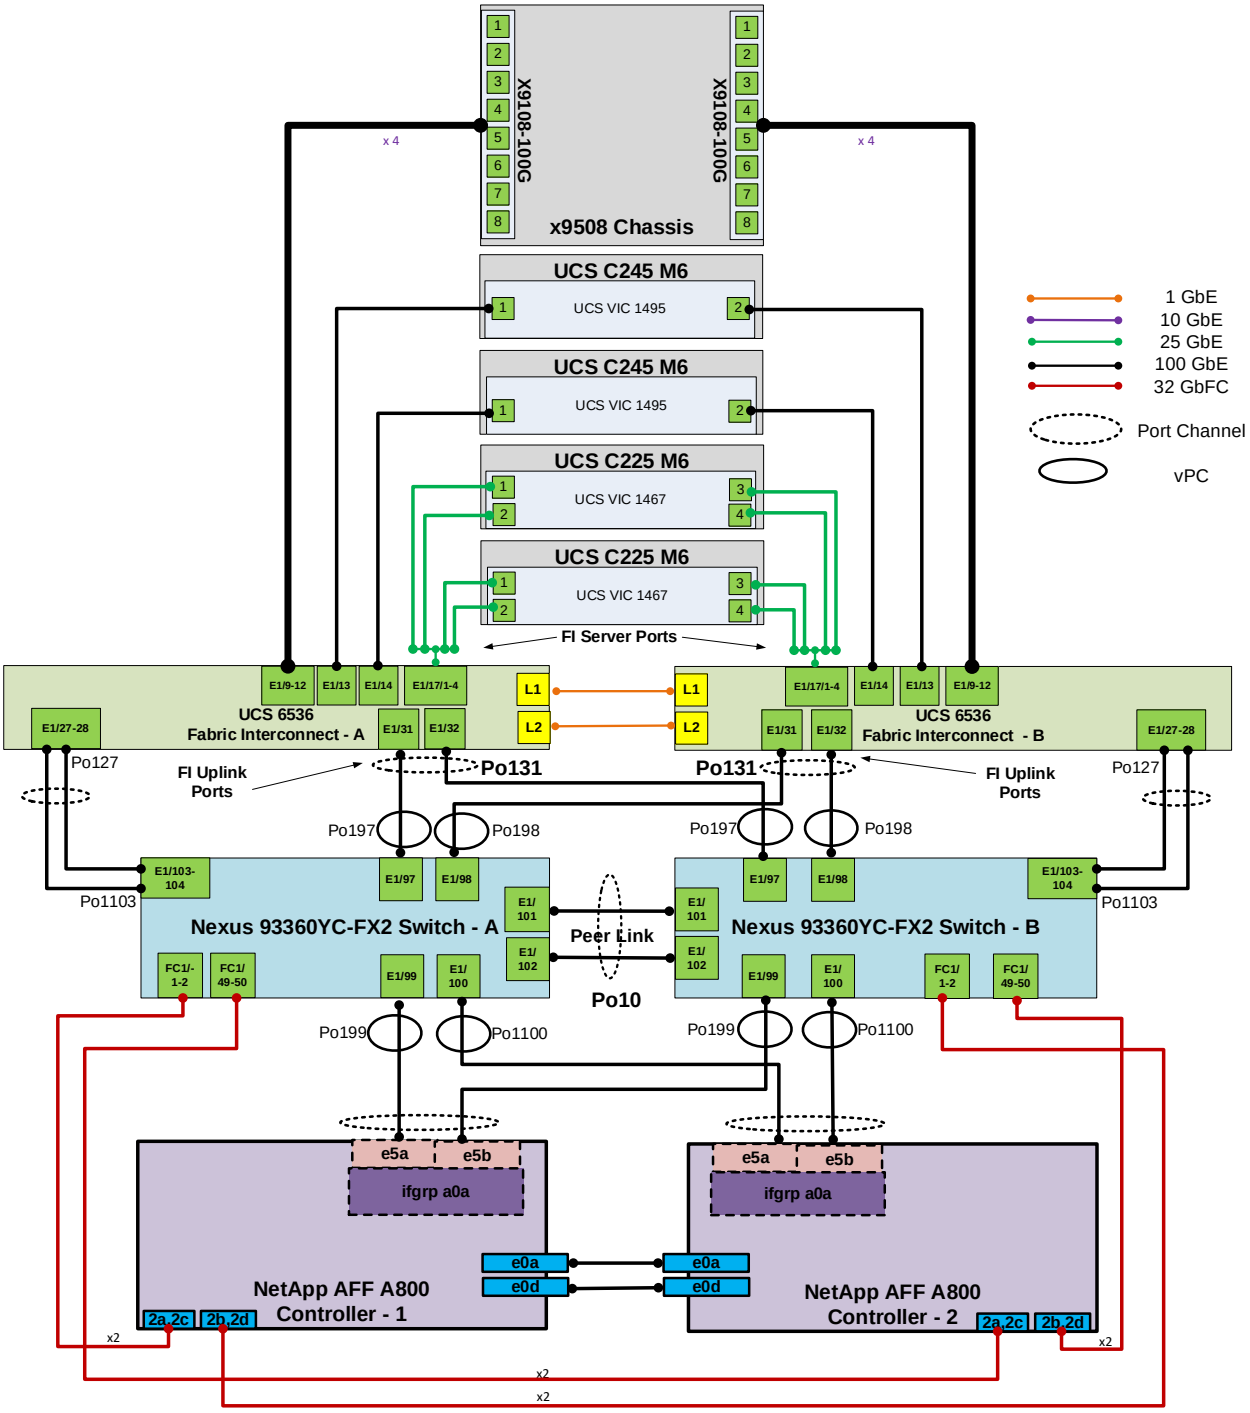


**Figure 5.**
Cisco Nexus SAN switching cabling with FCoE Fabric Interconnect uplinks

## FlexPod Cisco Nexus 93180YC-FX SAN Switching base configuration

The following procedures describe how to configure the Cisco Nexus 93180YC-FX switches for use in a base FlexPod environment that uses the switches for both LAN and SAN switching. This procedure assumes you are using Cisco Nexus 9000 10.2(4)M. It also assumes that you have created an FCoE uplink port-channel on the appropriate ports in the Cisco UCS Intersight Managed Mode Port Policies for each Cisco UCS fabric interconnect.

### Procedure 1. Set up initial configuration in Cisco Nexus 9300 Cloud Scale A

1. Configure the switch:

**Note:** On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power-on Auto Provisioning.

```
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-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-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]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: y
Configure default physical FC switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: y
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration:

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

### Procedure 2. Set up initial configuration in Cisco Nexus 9300 Cloud Scale B

1. Configure the switch:

**Note:** On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power-on Auto Provisioning.

```
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-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-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]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: y
Configure default physical FC switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: y
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration:

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

**Note:** SAN switching requires both the SAN_ENTERPRISE_PKG and FC_PORT_ACTIVATION_PKG licenses. Ensure these licenses are installed on each Cisco Nexus switch.

**Note:** This section is structured as a green-field switch setup. If you are setting up existing switches that are switching active traffic, execute this procedure down through Perform TCAM Carving and Configure Unified Ports in Cisco Nexus 9300 Cloud Scale A and B first on one switch and then, when that is completed, execute on the other switch.

### Procedure 3. Install feature-set fcoe in Cisco Nexus 9300 Cloud Scale A and B

1. Run the following commands to set global configurations:

```
config t
install feature-set fcoe
feature-set fcoe
system default switchport trunk mode auto
system default switchport mode F
```

**Note:** These steps are provided in case the basic Fibre Channel configurations were not configured in the switch setup script detailed in the previous section.

### Procedure 4. Set systemwide QoS configurations in Cisco Nexus 9300 Cloud Scale A and B

1. Run the following commands to set global configurations:

```
config t
system qos
service-policy type queuing input default-fcoe-in-que-policy
service-policy type queuing output default-fcoe-8q-out-policy
service-policy type network-qos default-fcoe-8q-nq-policy
copy run start
```

### Procedure 5. Perform TCAM carving and configure unified ports in Cisco Nexus 9300 Cloud Scale A and B

**Note:** SAN switching requires Ternary Content Addressable Memory (TCAM)carving for lossless Fibre Channel no-drop support. Also, you need to convert unified ports to Fibre Channel ports.

**Note:** On the Cisco Nexus 9300 Cloud Scale, unified ports are converted to Fibre Channel in groups of four in columns for example, 1,2,49,50.

1. Run the following commands:

```
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-ifacl 256
hardware access-list tcam region ing-redirect 256
slot 1
port 1,2,49,50,3,4,51,52 type fc
copy running-config startup-config
reload
This command will reboot the system. (y/n)? [n] y
```

2. After the switch reboots, log back in as admin. Then run the following commands:

```
show hardware access-list tcam region |i i ing-racl
show hardware access-list tcam region |i i ing-ifacl
show hardware access-list tcam region |i i ing-redirect
show int status
```

#### FlexPod Cisco Nexus 9300 Cloud Scale SAN and Ethernet Switching manual configuration

For the manual configuration of the Ethernet part of the Cisco Nexus 9300 Cloud Scale switches when using the switches for SAN switching, after the base configuration previously mentioned is set, return to FlexPod Cisco Nexus Switch manual configuration, and execute from there.

## FlexPod with Cisco Nexus 9300 Cloud Scale SAN Switching configuration – Part 2

**Note:** If the Cisco Nexus 9300 Cloud Scale switch is being used for SAN switching, you should complete this section in place of the "Cisco MDS" section of this document.

### FlexPod Cisco Nexus 9300 Cloud Scale SAN and Ethernet Switching manual configuration

This section details the manual configuration of the SAN part of the Cisco Nexus 9300 Cloud Scale switches when using the switches for SAN switching.

### Procedure 1. Enable features in Cisco Nexus 9300 Cloud Scale A and B

1. Log in as admin.

**Note:** SAN switching requires both the SAN_ENTERPRISE_PKG and FC_PORT_ACTIVATION_PKG licenses. Make sure these licenses are installed on each Cisco Nexus 9300 Cloud Scale switch.

2. Because basic Fibre Channel configurations were entered in the setup script, feature-set fcoe has been automatically in-stalled and enabled. Run the following commands:

```
config t
feature npiv
feature fport-channel-trunk
system default switchport trunk mode auto
system default switchport mode F
```

### Procedure 2. Configure FCoE VLAN and Fibre Channel ports in Cisco Nexus 9300 Cloud Scale A

1. From the global configuration mode, run the following commands:

```
vlan <vsan-a-id>
fcoe vsan <vsan-a-id>
name FCoE-VLAN-A

interface fc1/1
switchport description <st-clustername>-01:2a
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit
```

```
interface fc1/2
switchport description <st-clustername>-01:2c
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/49
switchport description <st-clustername>-02:2a
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/50
switchport description <st-clustername>-02:2c
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface Eth1/103
description <ucs-domainname>-a:FCoE:1/27
udld enable
channel-group 1103 mode active
no shutdown
exit

interface Eth1/104
description <ucs-domainname>-a:FCoE:1/28
udld enable
channel-group 1103 mode active
no shutdown
exit

interface port-channel1103
description <ucs-domainname>-a:FCoE
switchport mode trunk
switchport trunk allowed vlan <vsan-a-id>
spanning-tree port type edge trunk
```

```
mtu 9216

no negotiate auto
service-policy type qos input default-fcoe-in-policy
no shutdown
exit

interface vfc1103
switchport description <ucs-domainname>-a:FCoE
bind interface port-channel1103
switchport trunk allowed vsan <vsan-a-id>
switchport trunk mode on
no shutdown
exit
```

### Procedure 3. Configure FCoE VLAN and Fibre Channel ports in Cisco Nexus 9300 Cloud Scale B

1. From the global configuration mode, run the following commands:

```
vlan <vsan-b-id>
fcoe vsan <vsan-b-id>
name FCoE-VLAN-B

interface fc1/1
switchport description <st-clustername>-01:2b
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description <st-clustername>-01:2d
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/49
switchport description <st-clustername>-02:2b
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit
```

```
interface fc1/50
switchport description <st-clustername>-02:2d
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface Eth1/103
description <ucs-domainname>-b:FCoE:1/27
udld enable
channel-group 1103 mode active
no shutdown
exit

interface Eth1/104
description <ucs-domainname>-b:FCoE:1/28
udld enable
channel-group 1103 mode active
no shutdown
exit

interface port-channel1103
description <ucs-domainname>-b:FCoE
switchport mode trunk
switchport trunk allowed vlan <vsan-b-id>
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input default-fcoe-in-policy
no shutdown
exit

interface vfc1103
switchport description <ucs-domainname>-b:FCoE
bind interface port-channel1103
switchport trunk allowed vsan <vsan-b-id>
switchport trunk mode on
no shutdown
```

### Procedure 4. Create VSANs and add ports in Cisco Nexus 9300 Cloud Scale A

1. From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
vsan <vsan-a-id> interface fc1/1
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface fc1/2
Traffic on fc1/2 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface fc1/49
Traffic on fc1/49 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface fc1/50
Traffic on fc1/50 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface vfc1103
exit
zone smart-zoning enable vsan <vsan-a-id>
zoneset distribute full vsan <vsan-a-id>
copy run start
```

### Procedure 5. Create VSANs and add ports in Cisco Nexus 9300 Cloud Scale B

1. From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
vsan <vsan-b-id> interface fc1/1
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-b-id> interface fc1/2
Traffic on fc1/2 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-b-id> interface fc1/49
Traffic on fc1/49 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-b-id> interface fc1/50
Traffic on fc1/50 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-b-id> interface vfc1103
exit
zone smart-zoning enable vsan <vsan-b-id>
zoneset distribute full vsan <vsan-b-id>
copy run start
```

### Procedure 6. Create device-aliases in Cisco Nexus 93360YC-FX A to create zones

1. You can get the WWPN information required to create device-aliases and zones from NetApp using the following command:

```
network interface show -vserver Infra-SVM -data-protocol fcp
network interface show -vserver <svm-name> -data-protocol fc-nvme
```

2. To get the WWPN information for a server profile, by log into the Cisco Intersight dashboard and select each of the 3 server service profiles by going to **Infrastructure Service** > **Configure** > **Profiles** > **UCS Server Profiles** > **<Desired Server Profile>** > **Inventory** > **Network Adapters** > **<Adapter>** > **Interfaces**. The needed WWPNs are under HBA Interfaces.

3. Log in as admin and from the global configuration mode, run the following commands:

```
config t
device-alias mode enhanced
device-alias database
device-alias name <svm-name>-fcp-lif-01a pwwn <fcp-lif-01a-wwpn>
device-alias name <svm-name>-fcp-lif-02a pwwn <fcp-lif-02a-wwpn>
device-alias name FCP-<server1-hostname>-A pwwn <fcp-server1-wwpna>
device-alias name FCP-<server2-hostname>-A pwwn <fcp-server2-wwpna>
device-alias name FCP-<server3-hostname>-A pwwn <fcp-server3-wwpna>
device-alias name <svm-name>-fc-nvme-lif-01a pwwn <fc-nvme-lif-01a-wwpn>
device-alias name <svm-name>-fc-nvme-lif-02a pwwn <fc-nvme-lif-02a-wwpn>
device-alias name FC-NVMe-<server1-hostname>-A pwwn <fc-nvme-server1-wwpna>
device-alias name FC-NVMe-<server2-hostname>-A pwwn <fc-nvme-server2-wwpna>
device-alias name FC-NVMe-<server3-hostname>-A pwwn <fc-nvme-server3-wwpna>
device-alias commit
show device-alias database
```

### Procedure 7. Create device-aliases in Cisco Nexus 9300 Cloud Scale B to create zones

1. Log in as admin and, from the global configuration mode, run the following commands:

```
config t
device-alias mode enhanced
device-alias database
device-alias name <svm-name>-fcp-lif-01b pwwn <fcp-lif-01b-wwpn>
device-alias name <svm-name>-fcp-lif-02b pwwn <fcp-lif-02b-wwpn>
device-alias name FCP-<server1-hostname>-B pwwn <fcp-server1-wwpnb>
device-alias name FCP-<server2-hostname>-B pwwn <fcp-server2-wwpnb>
device-alias name FCP-<server3-hostname>-B pwwn <fcp-server3-wwpnb>
device-alias name <svm-name>-fc-nvme-lif-01b pwwn <fc-nvme-lif-01b-wwpn>
device-alias name <svm-name>-fc-nvme-lif-02b pwwn <fc-nvme-lif-02b-wwpn>
device-alias name FC-NVMe-<server1-hostname>-B pwwn <fc-nvme-server1-wwpnb>
device-alias name FC-NVMe-<server2-hostname>-B pwwn <fc-nvme-server2-wwpnb>
device-alias name FC-NVMe-<server3-hostname>-B pwwn <fc-nvme-server3-wwpnb>
```

```
device-alias commit
show device-alias database
```

### Procedure 8. Create zones and zonesets in Cisco Nexus 9300 Cloud Scale A

1. Run the following commands to create the required zones and zoneset on fabric A:

```
zone name FCP-<svm-name>-A vsan <vsan-a-id>
member device-alias FCP-<server1-hostname>-A init
member device-alias FCP-<server2-hostname>-A init
member device-alias FCP-<server3-hostname>-A init
member device-alias <svm-name>-fcp-lif-01a target
member device-alias <svm-name>-fcp-lif-02a target
exit
zone name FC-NVME-<svm-name>-A vsan <vsan-a-id>
member device-alias FC-NVME-<server1-hostname>-A init
member device-alias FC-NVME-<server2-hostname>-A init
member device-alias FC-NVME-<server3-hostname>-A init
member device-alias <svm-name>-fc-nvme-lif-01a target
member device-alias <svm-name>-fc-nvme-lif-02a target
exit
zoneset name FlexPod-Fabric-A vsan <vsan-a-id>
member FCP-<svm-name>-A
member FC-NVME-<svm-name>-A
exit
zoneset activate name FlexPod-Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```

### Procedure 9. Create zones and zonesets in Cisco Nexus 9300 Cloud Scale B

1. Run the following commands to create the required zones and zoneset on fabric B:

```
zone name FCP-<svm-name>-B vsan <vsan-b-id>
member device-alias FCP-<server1-hostname>-B init
member device-alias FCP-<server2-hostname>-B init
member device-alias FCP-<server3-hostname>-B init
member device-alias <svm-name>-fcp-lif-01b target
member device-alias <svm-name>-fcp-lif-02b target
exit
zone name FC-NVME-<svm-name>-B vsan <vsan-b-id>
member device-alias FC-NVME-<server1-hostname>-B init
member device-alias FC-NVME-<server2-hostname>-B init
member device-alias FC-NVME-<server3-hostname>-B init
member device-alias <svm-name>-fc-nvme-lif-01b target
member device-alias <svm-name>-fc-nvme-lif-02b target
exit
```

```
zoneset name FlexPod-Fabric-B vsan <vsan-b-id>
member FCP-<svm-name>-B
member FC-NVME-<svm-name>-B
exit
zoneset activate name FlexPod-Fabric-B vsan <vsan-b-id>
show zoneset active
copy r s
```

**Procedure 10. Switch testing commands**

You can use the following commands to check for correct switch configuration:

**Note:** To see complete results, some of these commands need to run after further configuration of the FlexPod components is complete.

```
show run
show run int
show int
show int status
show int brief
show flogi database
show device-alias database
show zone
show zoneset
show zoneset active
```

## Create a Custom ESXi ISO using VMware vCenter

In this white paper, the Cisco Custom Image for ESXi 7.0 U3 Install CD was used to install VMware ESXi. After this installation, the Cisco UCS VIC fnic driver, the lsi_mr3 driver, and the NetApp NFS Plug-In for VMware VAAI had to be installed or updated during the FlexPod deployment. You can use vCenter 7.0U3 or later to produce a FlexPod custom ISO containing the updated Cisco UCS VIC fnic driver, the lsi_mr3 driver, and the NetApp NFS Plug-In for VMware VAAI. You can use this ISO to install VMware ESXi 7.0U3 without having to do any additional driver updates.

**Procedure 1. Create a FlexPod ESXi Custom ISO using VMware vCenter**

1. Download the Cisco Custom Image for ESXi 7.0 U3 Offline Bundle. You can use this file (VMware-ESXi-7.0.3d-19482537-Custom-Cisco-4.2.2-a-depot.zip) to produce the FlexPod ESXi 7.0U3 CD ISO.
2. Download the following listed .zip files:
   - VMware ESXi 7.0 nfnic 5.0.0.34 Driver for Cisco VIC Adapters – Cisco-nfnic_5.0.0.34-1OEM.700.1.0.15843807_19966277.zip – extracted from the downloaded zip file
   - VMware ESXi 7.0 lsi_mr3 7.720.04.00-1OEM SAS Driver for Broadcom Megaraid 12Gbps – Broadcom-lsi-mr3_7.720.04.00-1OEM.700.1.0.15843807_19476191.zip – extracted from the downloaded zip file
   - NetApp NFS Plug-In for VMware VAAI 2.0 – NetAppNasPluginV2.0.zip
   - The Cisco VIC nenic driver is also normally be downloaded and added to the FlexPod Custom ISO, but the 1.0.42.0 nenic driver is already included in the Cisco Custom ISO.

3. Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.
4. Under the Menu at the top, select **Auto Deploy**.
5. If you see the following, select **ENABLE IMAGE BUILDER**.

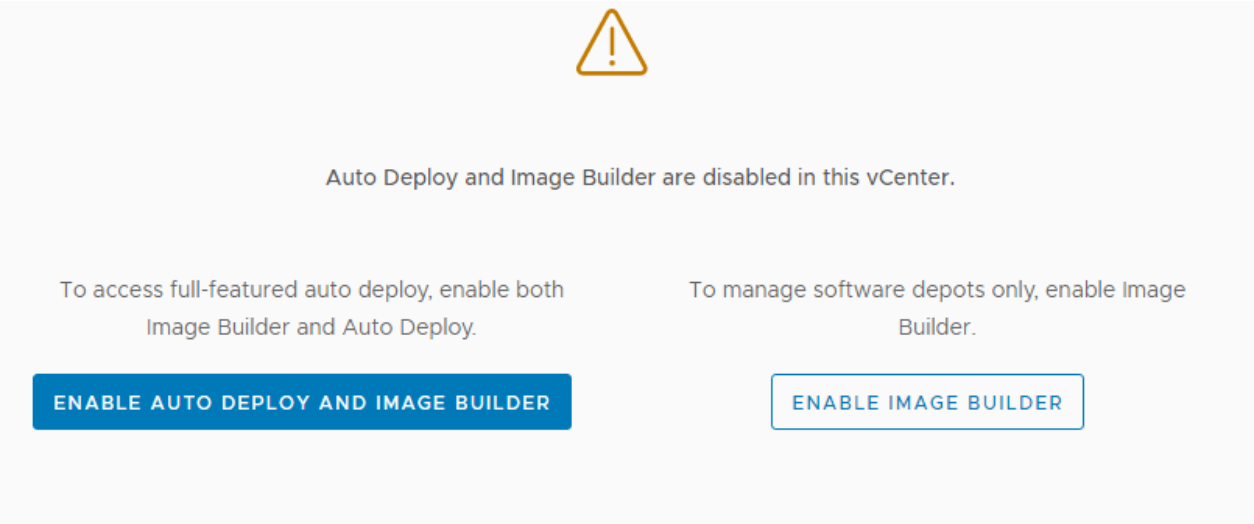


6. Click **IMPORT** to upload a software depot.
7. Name the depot “Cisco Custom ESXi 7.0U3.” Click **BROWSE**. Browse to the local location of the VMware-ESXi-7.0.3d-19482537-Custom-Cisco-4.2.2-a-depot.zip file downloaded previously, highlight it, and click **Open**.

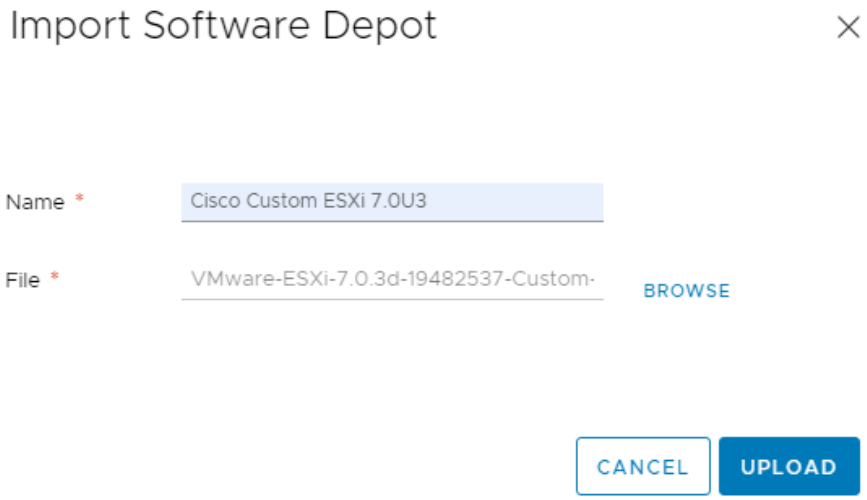


8. Click **UPLOAD** to upload the software depot.
9. Repeat steps 1-8 to add software depots for Cisco-nfnic_5.0.0.34-1OEM.700.1.0.15843807_19966277.zip, Broadcom-lsi-mr3_7.720.04.00-1OEM.700.1.0.15843807_19476191.zip, and NetAppNasPluginV2.0.zip.
10. Click **NEW** to add a custom software depot.
11. Select **Custom depot** and name the custom depot FlexPod-ESXi-7.0U3.

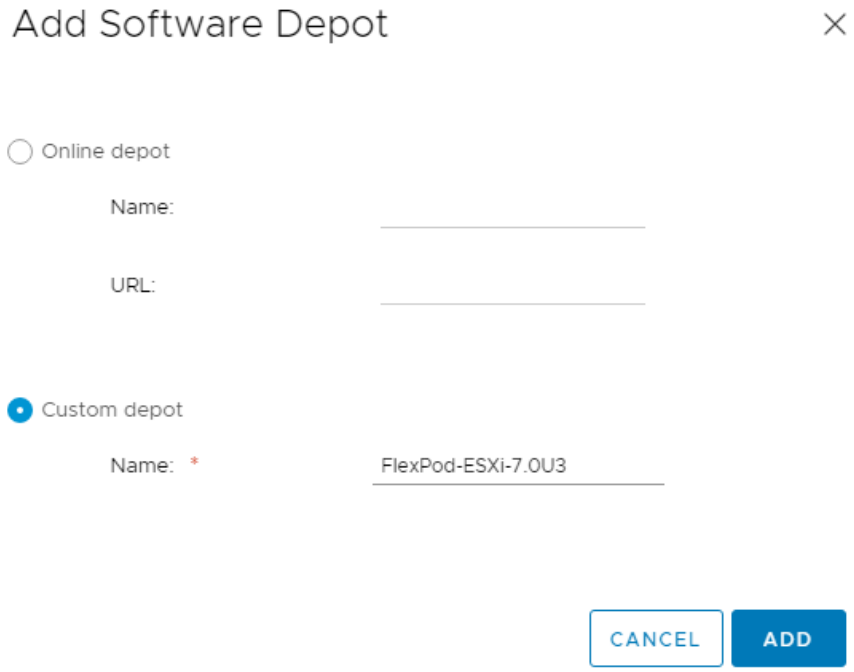


12. Click **ADD** to add the custom software depot.
13. From the drop-down list, select the Cisco Custom ESXi-7.0U3 (ZIP) software depot. Make sure the Image Profiles tab is selected and then click the radio button to select the Cisco-UCS-Addon-ESXi-7U3d-19482537_4.2.2-a image profile. Click **CLONE** to clone the image profile.
14. Name the clone FlexPod-ESXi-7.0U3. For Vendor, enter Cisco-NetApp. For Description, enter **Cisco Custom ISO ESXi 7.0U3 with Cisco VIC nfnic 5.0.0.34, LSI-MR3 7.720.04.0 and NetAppNasPluginv2.0**. Select **FlexPod-ESXi-7.0U3** for Software depot.

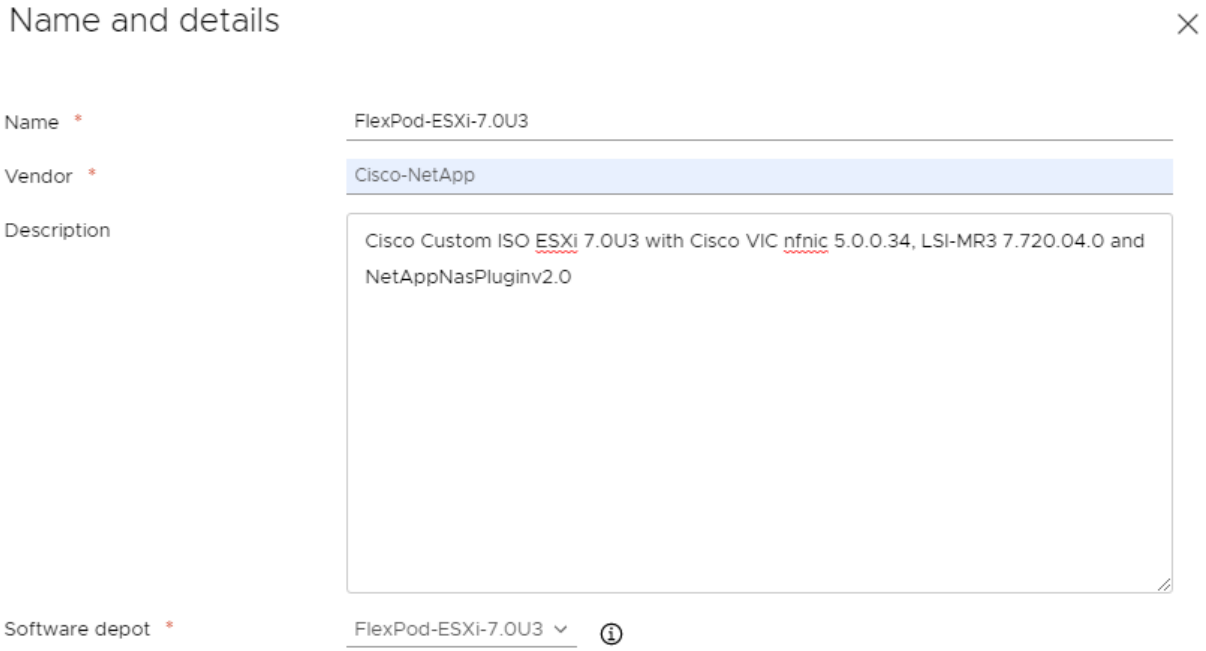

15. Click **NEXT**.
16. Under Available software packages, check lsi-mr3 7.720.04.00 and uncheck any other lsi-mr3 packages, check **NetAppNasPlugin 2.0-15**, and check **nfnic 5.0.0.34** and uncheck any other nfnic packages. Leave the remaining selections unchanged.

Select software packages

Acceptance level: Partner supported

| Name | Version | Acceptance Level | Vendor | Depot |
|---|---|---|---|---|
| lpnic | 11.4.62.0-1vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| lsi-mr3 | 7.720.04.00-1OEM.700... | VMware certified | BCM | LSI MR3 7.720.04.00 |
| lsi-mr3 | 7.718.02.00-1vmw.703.... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| lsi-msgpt2 | 20.00.06.00-4vmw.70... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| lsi-msgpt3 | 17.00.12.00-1vmw.703.... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| lsi-msgpt35 | 19.00.02.00-1vmw.703.... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| lsuv2-hpv2-hpsa-... | 1.0.0-3vmw.703.0.20.19... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-intelv2-nv... | 2.7.2173-1vmw.703.0.20... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-lsiv2-driver... | 1.0.0-10vmw.703.0.35.1... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-nvme-pcie-... | 1.0.0-1vmw.703.0.20.191... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-oem-dell-pl... | 1.0.0-1vmw.703.0.20.191... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-oem-hp-pl... | 1.0.0-1vmw.703.0.20.191... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-oem-lenov... | 1.0.0-1vmw.703.0.20.191... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| lsuv2-smartpqiv2... | 1.0.0-8vmw.703.0.20.19... | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| mtip32xx-native | 3.9.8-1vmw.703.0.20.19... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| native-misc-drive... | 7.0.3-0.35.19482537 | VMware certified | VMware | Cisco Custom ESXi 7.0... |
| ne1000 | 0.8.4-11vmw.703.0.20.1... | VMware certified | VMW | Cisco Custom ESXi 7.0... |

83 selected of 100 Items

## Select software packages

Acceptance level: Partner supported

| | Name | Version | Acceptance Level | Vendor | Depot |
|---|---|---|---|---|---|
| ☑ | ne1000 | 0.8.4-11vmw.703.0.20.1... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nenic | 1.0.42.0-1OEM.670.0.0.... | VMware certified | Cisco | Cisco Custom ESXi 7.0... |
| ☐ | nenic | 1.0.33.0-1vmw.703.0.20... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nenic-ens | 1.0.6.0-1OEM.700.1.0.15... | VMware certified | Cisco | Cisco Custom ESXi 7.0... |
| ☑ | NetAppNasPlugin | 2.0-15 | VMware accepted | NetApp | NetApp NAS Plugin v2.0 |
| ☑ | nfnic | 5.0.0.34-1OEM.700.1.0.1... | VMware certified | Cisco | Cisco nfnic 5.0.0.34 |
| ☐ | nfnic | 4.0.0.87-1OEM.670.0.0.... | VMware certified | Cisco | Cisco Custom ESXi 7.0... |
| ☐ | nfnic | 4.0.0.70-1vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nhpsa | 70.0051.0.100-4vmw.7... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nmlx4-core | 3.19.16.8-2vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nmlx4-en | 3.19.16.8-2vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nmlx4-rdma | 3.19.16.8-2vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nmlx5-core | 4.21.71.101-1OEM.702.0.... | VMware certified | MEL | Cisco Custom ESXi 7.0... |
| ☐ | nmlx5-core | 4.19.16.11-1vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☐ | nmlx5-rdma | 4.19.16.11-1vmw.703.0.2... | VMware certified | VMW | Cisco Custom ESXi 7.0... |
| ☑ | nmlx5-rdma | 4.21.71.101-1OEM.702.0.... | VMware certified | MEL | Cisco Custom ESXi 7.0... |
| ☑ | ntg3 | 4.1.7.0-0vmw.703.0.20.... | VMware certified | VMW | Cisco Custom ESXi 7.0... |

84 selected of 100 Items

17. Click **NEXT**.

## Ready to complete

| | |
|---|---|
| Name | FlexPod-ESXi-7.0U3 |
| Vendor | Cisco-NetApp |
| Acceptance level | Partner supported |
| Description | Cisco Custom ISO ESXi 7.0U3 with Cisco VIC nfnic 5.0.0.34, LSI-MR3 7.720.04.0 and NetAppNasPluginv2.0 |
| Software depot | FlexPod-ESXi-7.0U3 |
| Software packages | 84 |

18. Click **FINISH** to generate the depot.

19. Using the Software Depot pull-down menu, select the FlexPod-ESXi-7.0U3 (Custom) software depot. Under Image Profiles select the FlexPod-ESXi-7.0U3 image profile. Click **EXPORT** to export an image profile. ISO should be selected. Click **OK** to generate a bootable ESXi installable image.
20. When the Image profile export completes, click **DOWNLOAD** to download the ISO.
21. After it downloads, you can rename the ISO to a more descriptive name (for example, FlexPod-ESXi-7.0U3.iso).
22. Optionally, generate the ZIP archive to generate an offline bundle for the FlexPod image using **... > Export**.

## Active IQ Unified Manager user configuration

**Procedure 1. Add uocal Users to Active IQ Unified Manager**

1. Navigate to **Settings** > **General** section and click **Users**.

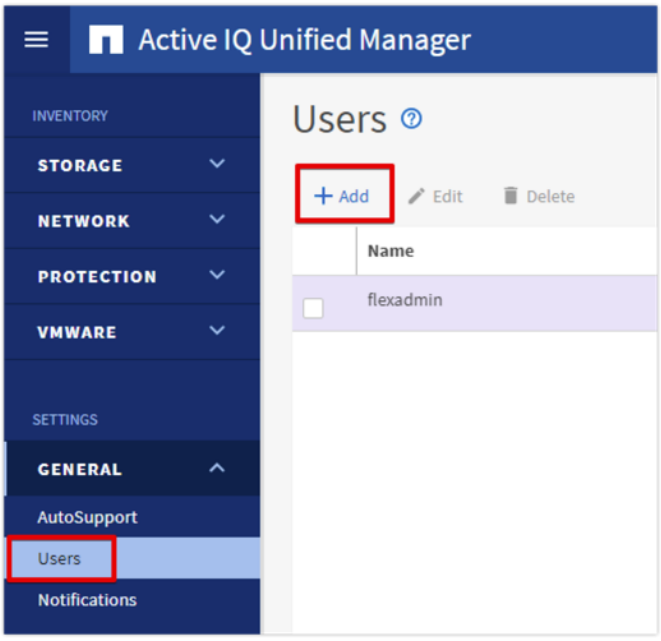


2. Click **+ Add** and complete the requested information:
   - Select Local User for the Type.
   - Enter a username and password.
   - Add the user's email address.
   - Select the appropriate role for the new user.

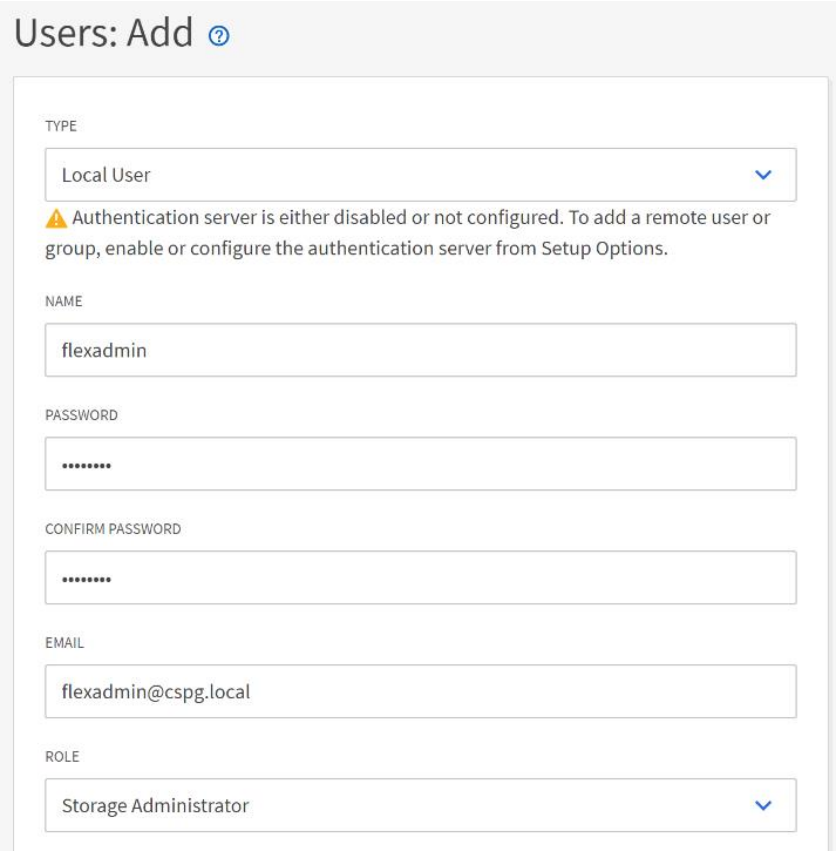


3. Click **SAVE** to finish adding the new user.

**Procedure 2. Configure remote authentication**

Simplify user management and authentication for Active IQ Unified Manager by integrating it with Microsoft Active Directory.

**Note:** You must be logged on as the maintenance user created during the installation or another user with application administrator privileges to configure remote authentication.

1. Navigate to the **General** and select **Remote Authentication**.
2. Select the option to enable remote authentication and define a remote user or remote group

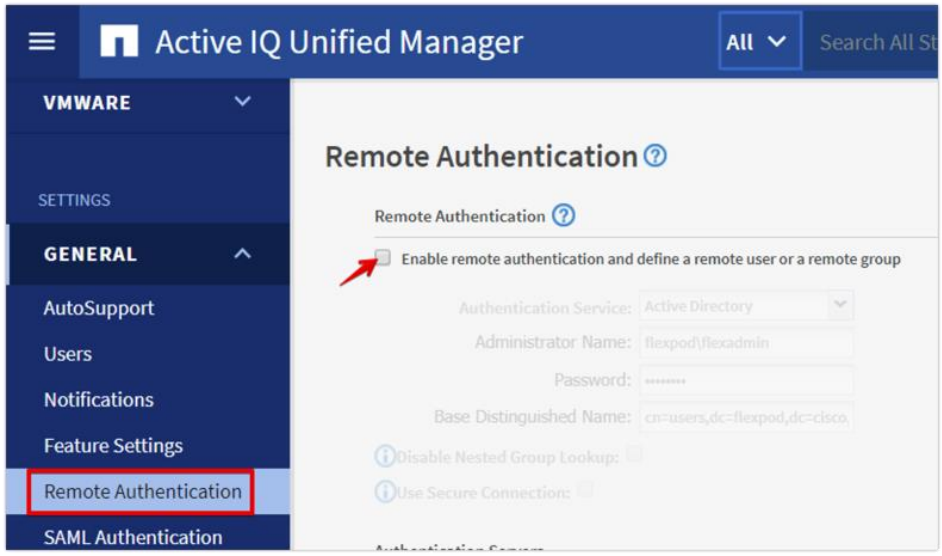


3. Select **Active Directory** from the authentication service list.
4. Enter the Active Directory service account name and password. The account name can be in the format of domain\user or user@domain.
5. Enter the base Distinguished Name (DN) where your Active Directory users reside.
6. If Active Directory Lightweight Directory Access Protocol (LDAP) communications are protected by Secure Shell (SSL) Protocol, enable the **Use Secure Connection** option.
7. Add one or more Active Directory domain controllers by clicking **Add** and entering the IP or FQDN of the domain controller.
8. Click **Save** to enable the configuration.

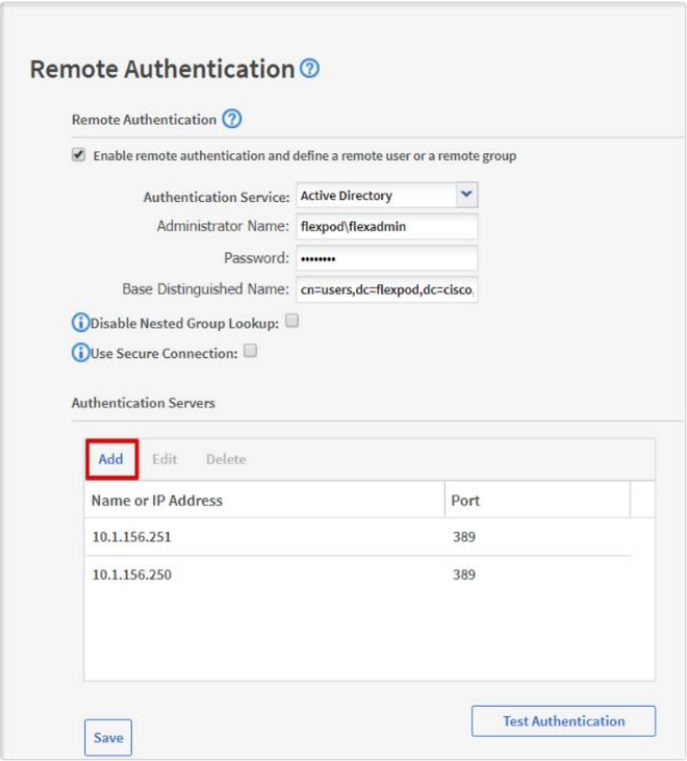

9. Click **Test Authentication** and enter an Active Directory username and password to test authentication with the Active Directory authentication servers. Click **Start**.


A result message indicating authentication was successful should display:

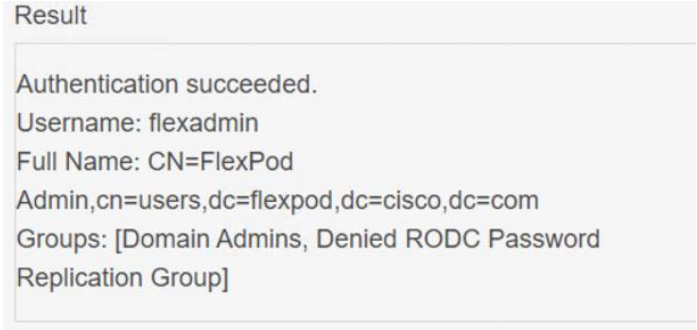


**Procedure 3. Add a remote user to Active IQ Unified Manager**

1. Navigate to the **General** section and select **Users**.
2. Click **Add** and select **Remote User** from the Type drop-down list.
3. Enter the following information on the form:
   - The **Name** of the Active Directory user
   - **Email** address of the user
   - The appropriate **Role** for the user

NAME

PASSWORD

CONFIRM PASSWORD

EMAIL

ROLE

Application Administrator

Save Cancel

4. Click **Save** to add the remote user to Active IQ Unified Manager.

## Active IQ Unified Manager vCenter configuration

Active IQ Unified Manager provides visibility into vCenter and the virtual machines running inside the datastores backed by NetApp ONTAP storage. Virtual machines and storage are monitored to enable quick identification of performance problems within the various components of the virtual infrastructure stack.

**Note:** Before adding vCenter into Active IQ Unified Manager, you must change the log level of the vCenter server .

### Procedure 1. Configure Active IQ Unified Manager vCenter

1. In the vSphere client navigate to **Menu > VMs and Templates** and select the vCenter instance from the top of the object tree.
2. Click the **Configure** tab, expand **Settings**, and select **General**.

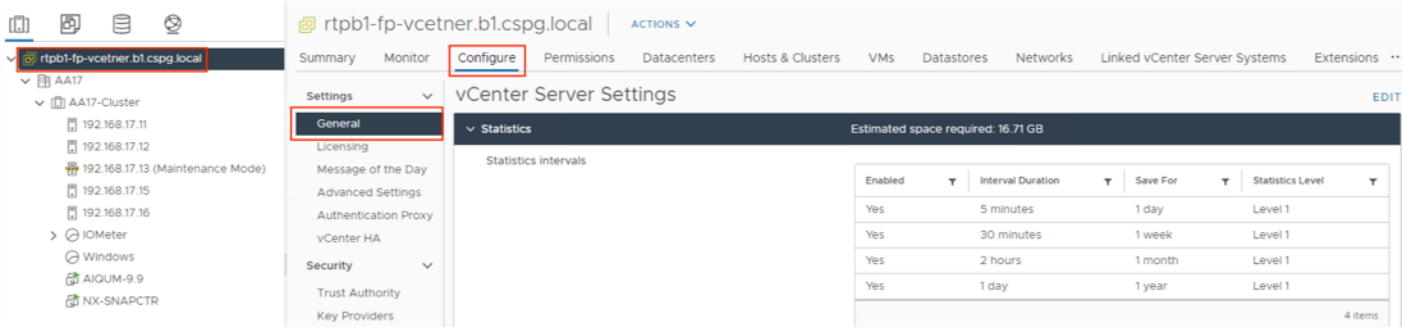


3. Click **EDIT**.
4. In the pop-up window under Statistics, locate the 5 minutes Interval Duration row and change the setting to **Level 3** under the Statistics Level column.
5. Click **SAVE**.

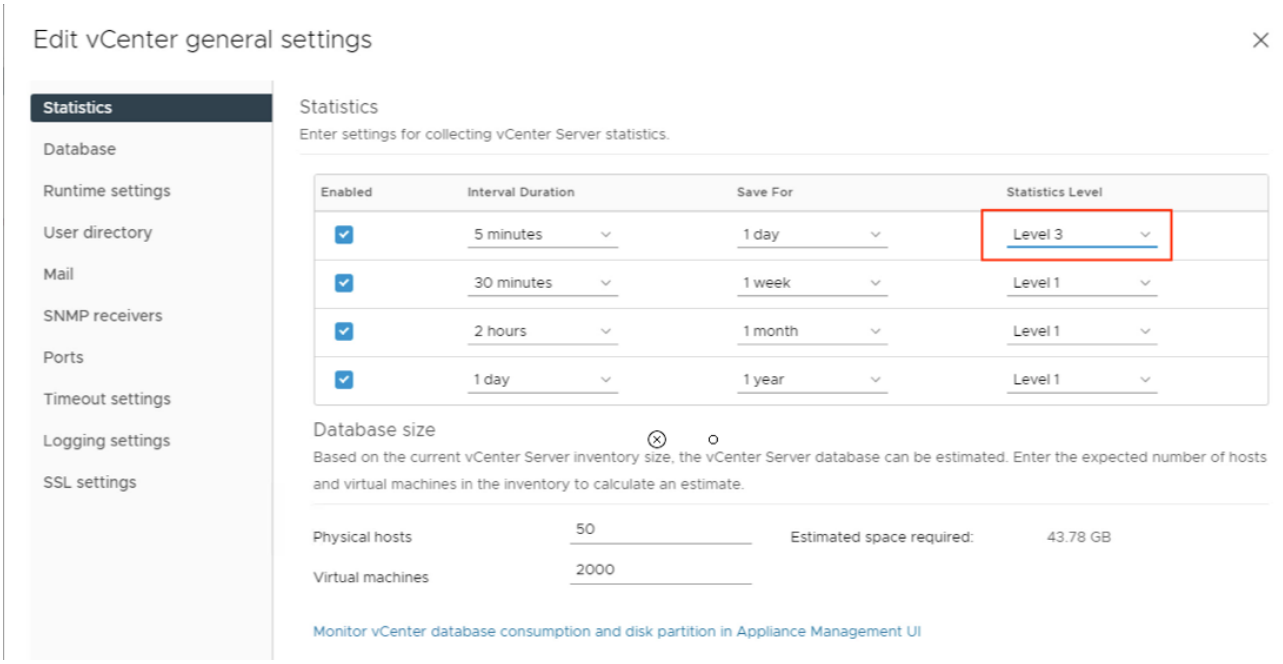


6. Switch to the Active IQ Unified Manager and navigate to the **VMware** section located under **Inventory**.
7. Expand VMware and select **vCenter**.

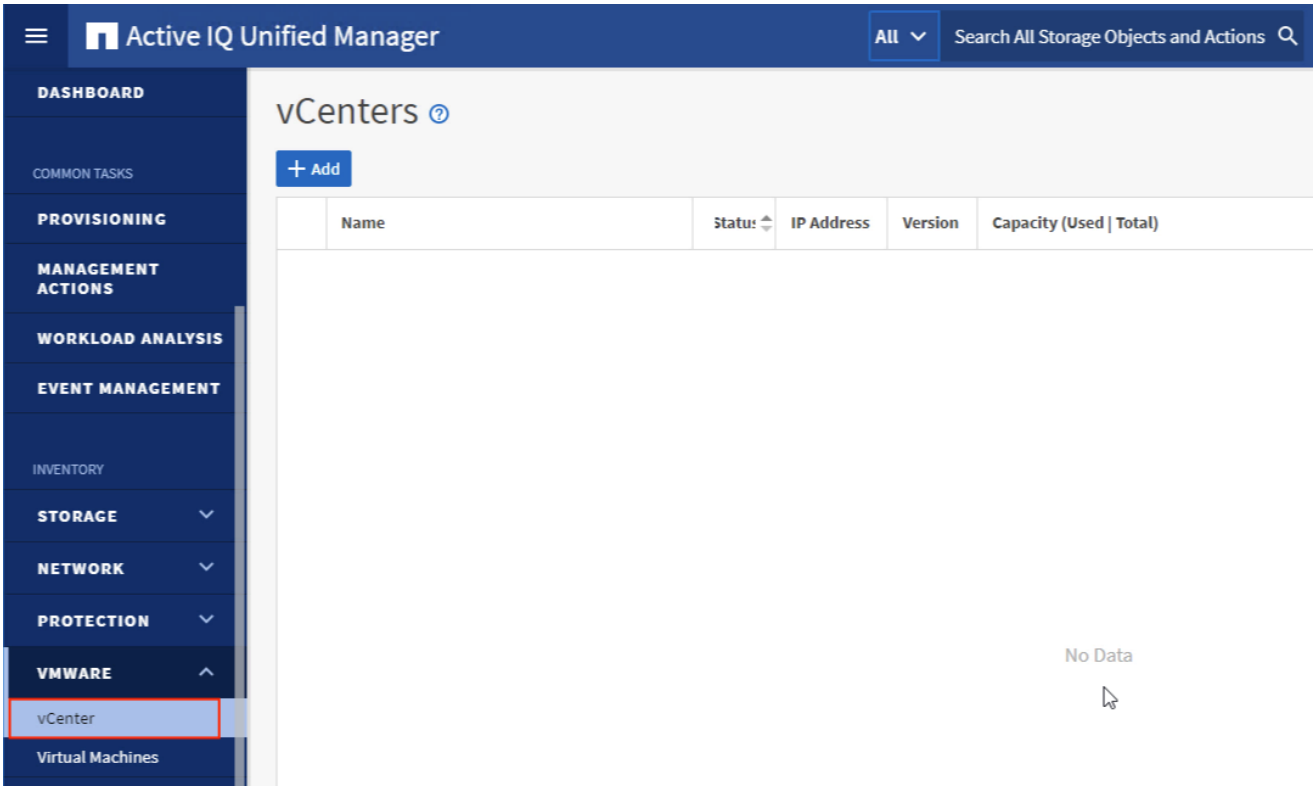


8. Click **Add**.
9. Enter the VMware vCenter server details and click **Save**.

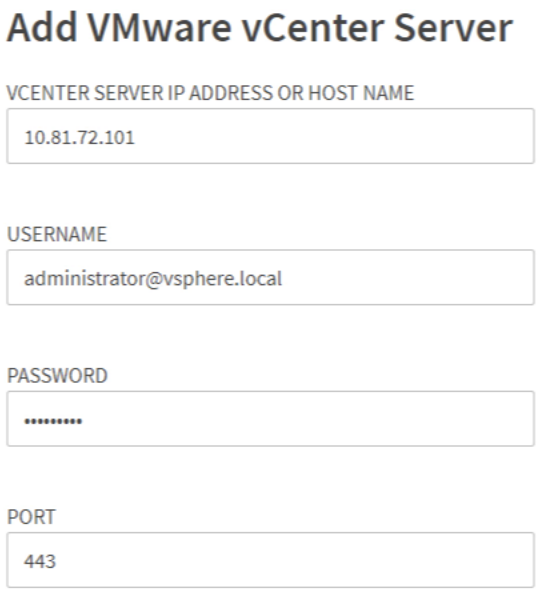


10. A dialog box asking to authorize the certificate will appear. Click **Yes** to accept the certificate and add the vCenter server.

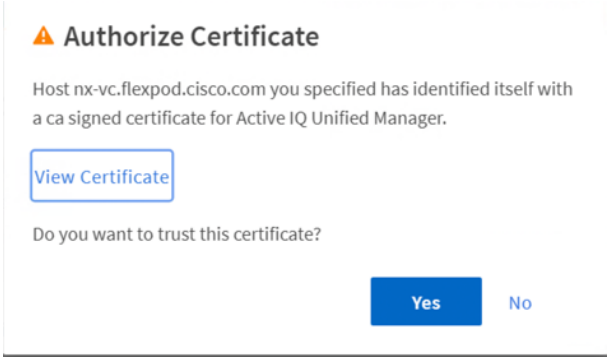


**Note:** It may take up to 15 minutes to discover vCenter. Performance data can take up to an hour to become available.

### Procedure 2. View virtual machine inventory

The virtual-machine inventory is automatically added to Active IQ Unified Manager during discovery of the vCenter server. You can view virtual machines in a hierarchical display detailing storage capacity, IOPS, and latency for each component in the virtual infrastructure to troubleshoot the source of any performance-related problems.

1. Log into NetApp Active IQ Unified Manager.
2. Navigate to the VMware section located under Inventory, expand the section, and click **Virtual Machines**.

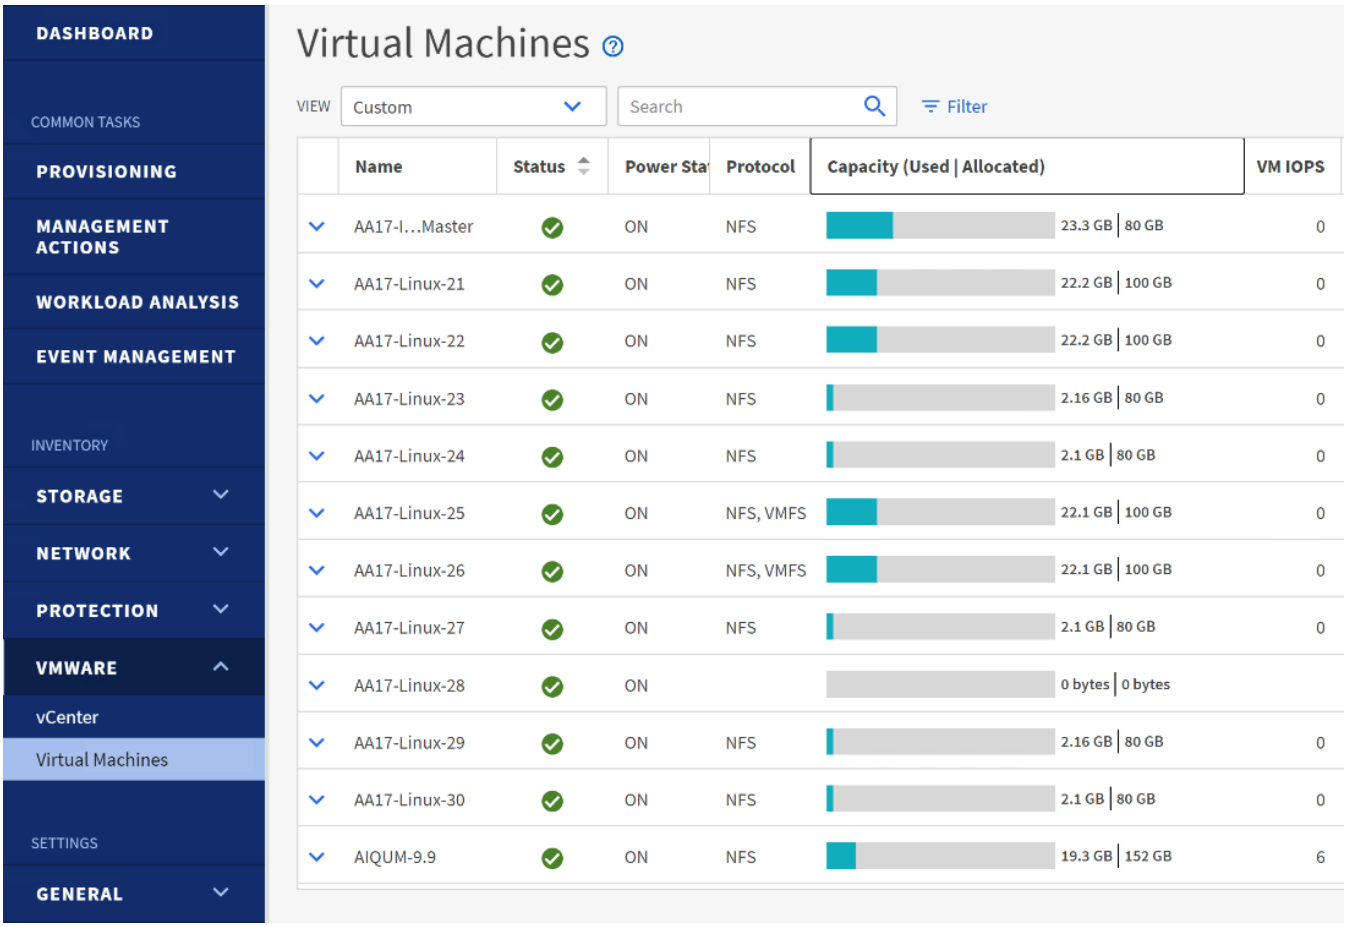


3. Select a virtual machine and click the blue caret to expose the topology view. Review the compute, network, and storage components and their associated IOPS and latency statistics.

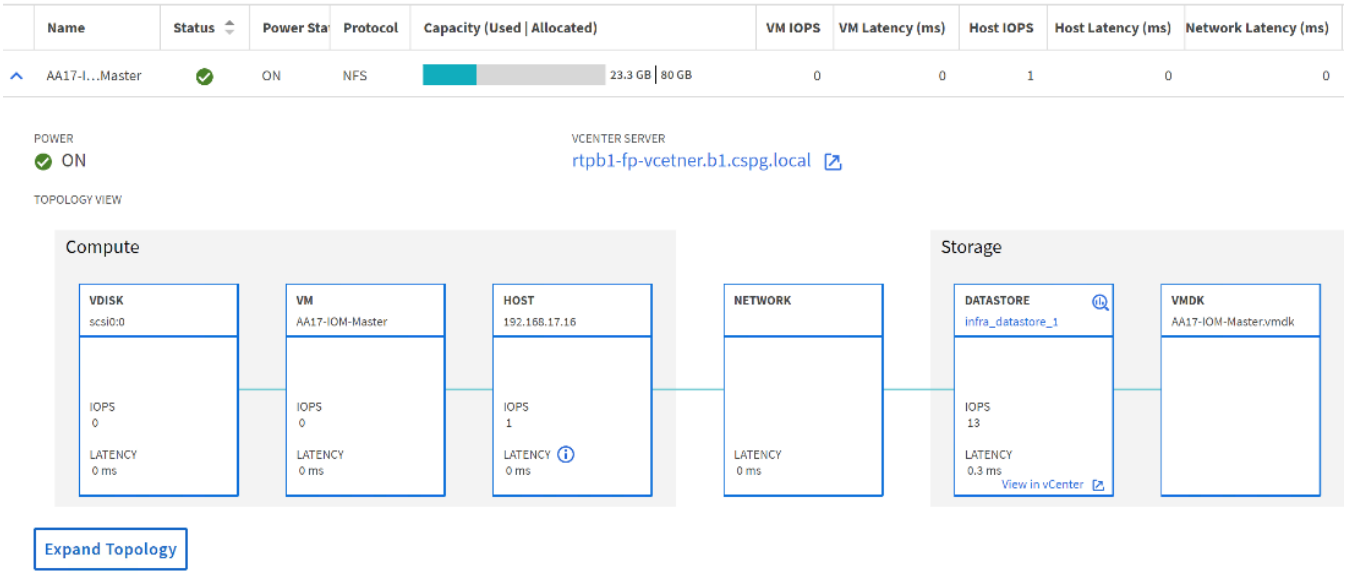


4. Click **Expand Topology** to see the entire hierarchy of the virtual machine and its virtual disks as it is connected through the virtual infrastructure stack. The virtual-machine components are mapped from vSphere and compute through the network to the storage.

## NetApp Active IQ

NetApp Active IQ is a data-driven service that uses artificial intelligence and machine learning to provide analytics and actionable intelligence for NetApp ONTAP storage systems. Active IQ uses AutoSupport data to deliver proactive guidance and best-practice recommendations to optimize storage performance and minimize risk. Additional Active IQ documentation is available on the Active IQ Documentation Resources webpage.

**Note:** Active IQ is automatically enabled when AutoSupport is configured on the NetApp ONTAP storage controllers.

### Procedure 1. Configure NetApp Active IQ

1. Navigate to the Active IQ portal at https://activeiq.netapp.com/.
2. Log in with NetApp support account ID.
3. At the Welcome screen enter the cluster name or one of the controller serial numbers in the search box. Active IQ will automatically begin searching for the cluster and display results:

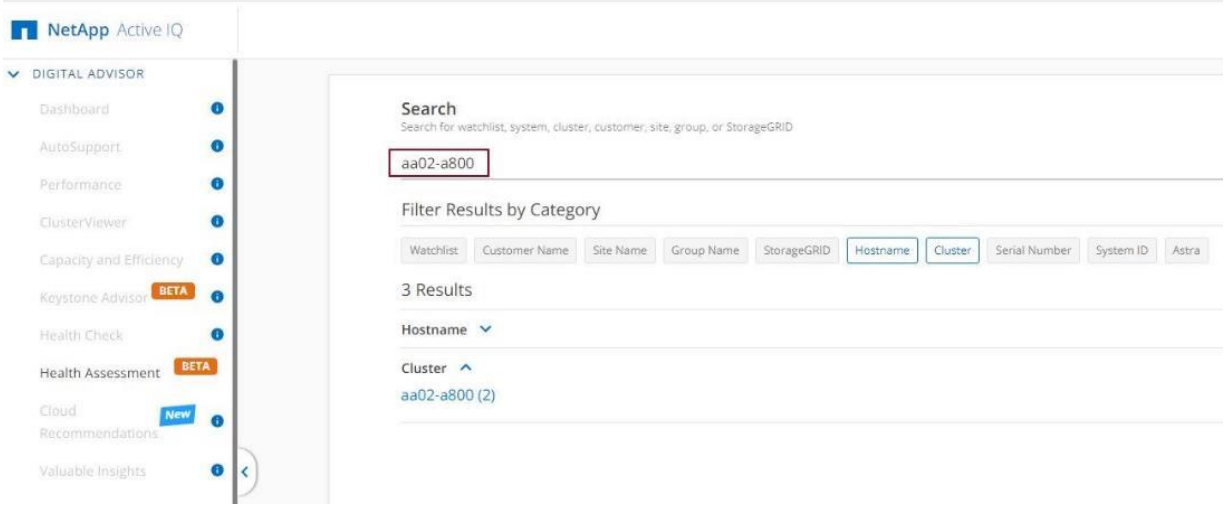


4. Click the <**cluster name**> (for example, aa02-a800) to launch the dashboard for this cluster.

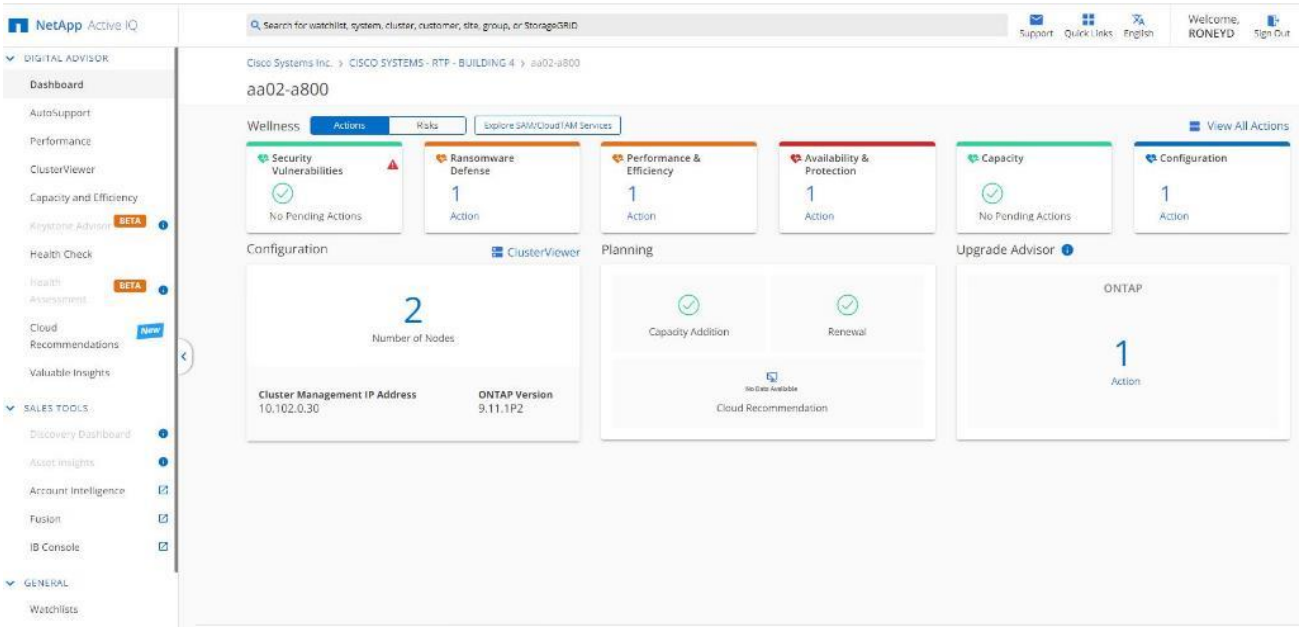

### Procedure 2. Add a WatchList to the Digital Advisor dashboard

The Active IQ Digital Advisor provides a summary dashboard and system wellness score based on the health and risks that Active IQ has identified. The dashboard provides a quick way to identify and get proactive recommendations on how to mitigate risks in the storage environment, including links to technical reports and mitigation plans. This procedure details the steps to create a watch list and launch the Digital Advisor dashboard for the WatchList.

1. Click **GENERAL > Watchlists** in the left menu bar.
2. Enter a name for the **WatchList**.
3. Select the radio button to add systems by serial number and enter the cluster serial numbers to the watch list.
4. Check the box for **Make this my default watchlist** if desired.

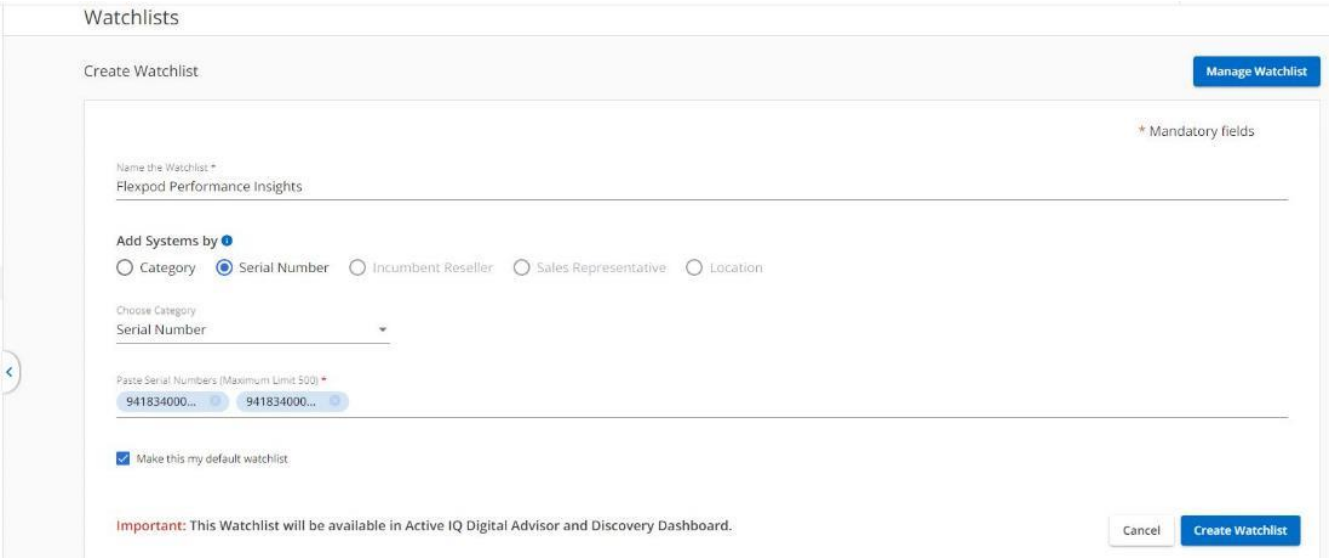


5. Click Create Watchlist.
6. Click **GENERAL** > **Watchlists** in the left menu bar again to list the watch list created.

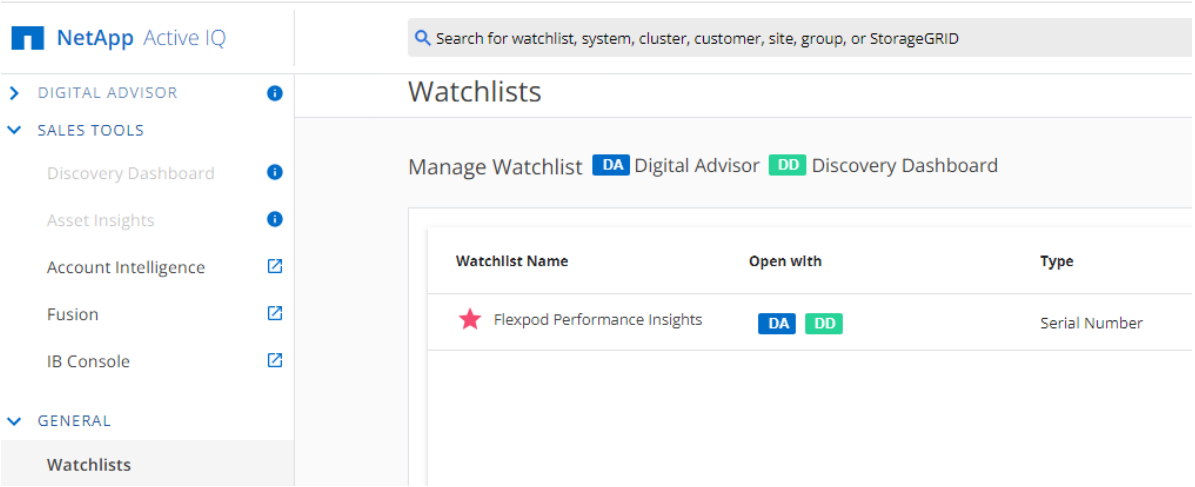


**Note:** The Discovery Dashboard functions have been moved to the IB Console (Install Based). Notice that Discovery Dashboard is greyed out under SALES TOOLS.

7. Click the blue box labelled DA to launch the specific WatchList in **Digital Advisor Dashboard**.
8. Review the enhanced dashboard to learn more about any recommended actions or risks.

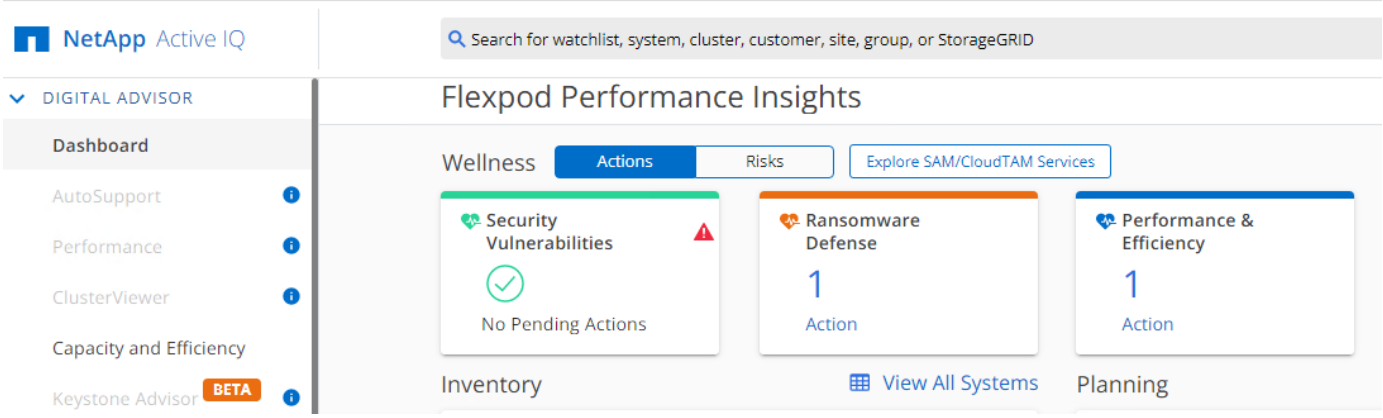


9. Switch between the **Actions** and **Risks** tabs to view the risks by category or a list of all risks with their impact and links to corrective actions.

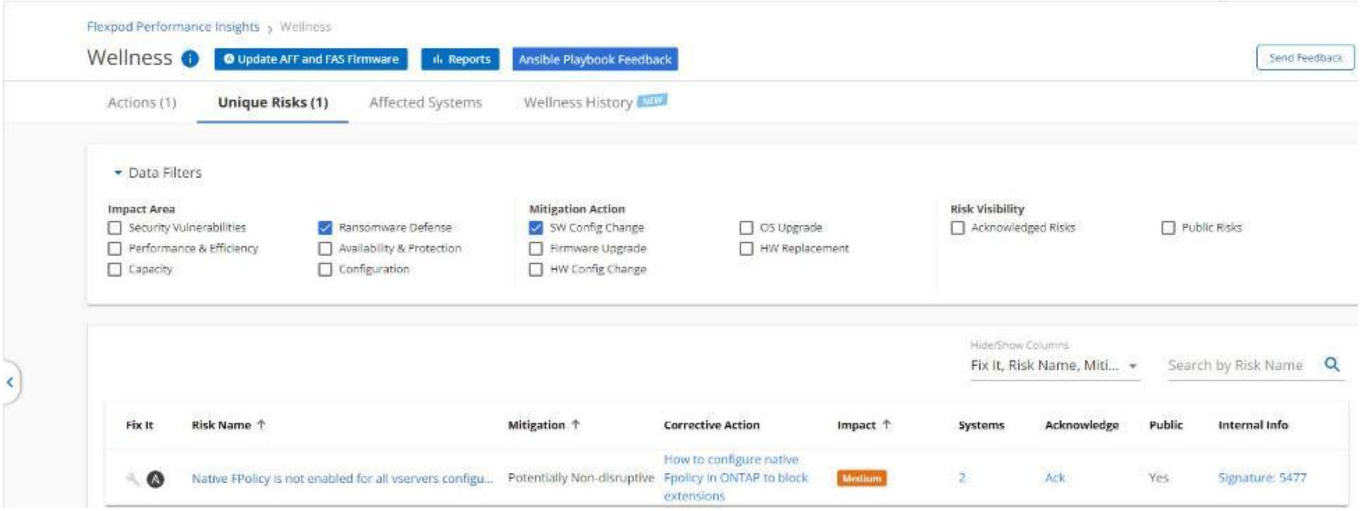


10. Click the links in the Corrective Action column to read the bug information or knowledge base article about how to remediate the risk.

**Note:** You can view additional tutorials and video walk-throughs of Active IQ features on the following page: https://docs.netapp.com/us-en/active-iq/.

## FlexPod backups

### Cisco Intersight SaaS platform

The Cisco Intersight SaaS platform maintains customer configurations online. No separate backup was created for the Cisco UCS configuration. If you are using a Cisco Intersight Private Virtual Appliance (PVA), ensure that the NetApp SnapCenter Plug-In for VMware vSphere is creating periodic backups of this appliance.

### Procedure 1. Cisco Nexus and MDS backups

You can manually back up the configuration of the Cisco Nexus 9000 and Cisco MDS 9132T switches at any time with the copy command, but you can enable automated backups using the NX-OS feature scheduler.

An example of setting up automated configuration backups of one of the NX-OS switches follows:

```
feature scheduler
scheduler logfile size 1024
scheduler job name backup-cfg
copy running-config tftp://<server-ip>/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
exit
scheduler schedule name daily
job name backup-cfg
time daily 2:00
end
```

**Note:** Using "vrf management" in the copy command is needed only when Mgmt0 interface is part of Virtual Route Forwarding (VRF) management.

1. Verify that the scheduler job has been correctly set up using the following command(s):

```
show scheduler job
Job Name: backup-cfg
--------------------
copy running-config tftp://10.1.156.150/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management

==============================================================================

show scheduler schedule
Schedule Name : daily
---------------------------
User Name : admin
Schedule Type : Run every day at 2 Hrs 0 Mins
Last Execution Time : Yet to be executed
----------------------------------------------
 Job Name Last Execution Status
----------------------------------------------
backup-cfg -NA-
==============================================================================
```

The documentation for the feature scheduler is available here: [Nexus 9000 – Configuring the Scheduler]

### Procedure 2. VMware VCSA backup

**Note:** Basic scheduled backup for the vCenter Server Appliance (VCSA) is available within the native capabilities of the VCSA.

1. Connect to the VCSA Console at **https://<VCSA IP>:5480**.
2. Log in as **root**.
3. Click **Backup** in the list to open the Backup Schedule Dialogue.

4. To the right of Backup Schedule, click **CONFIGURE**.
5. Specify the following:
   - The **Backup location** with the protocol to use (FTPS, HTTPS, SFTP, FTP,NFS, SMB, and HTTP)
   - The **User name** and **Password**; for the NFS (NFS3) example captured in the following illustration, the username is root; it uses a random password because NFSv3 sys security was configured
   - The Number of backups to retain

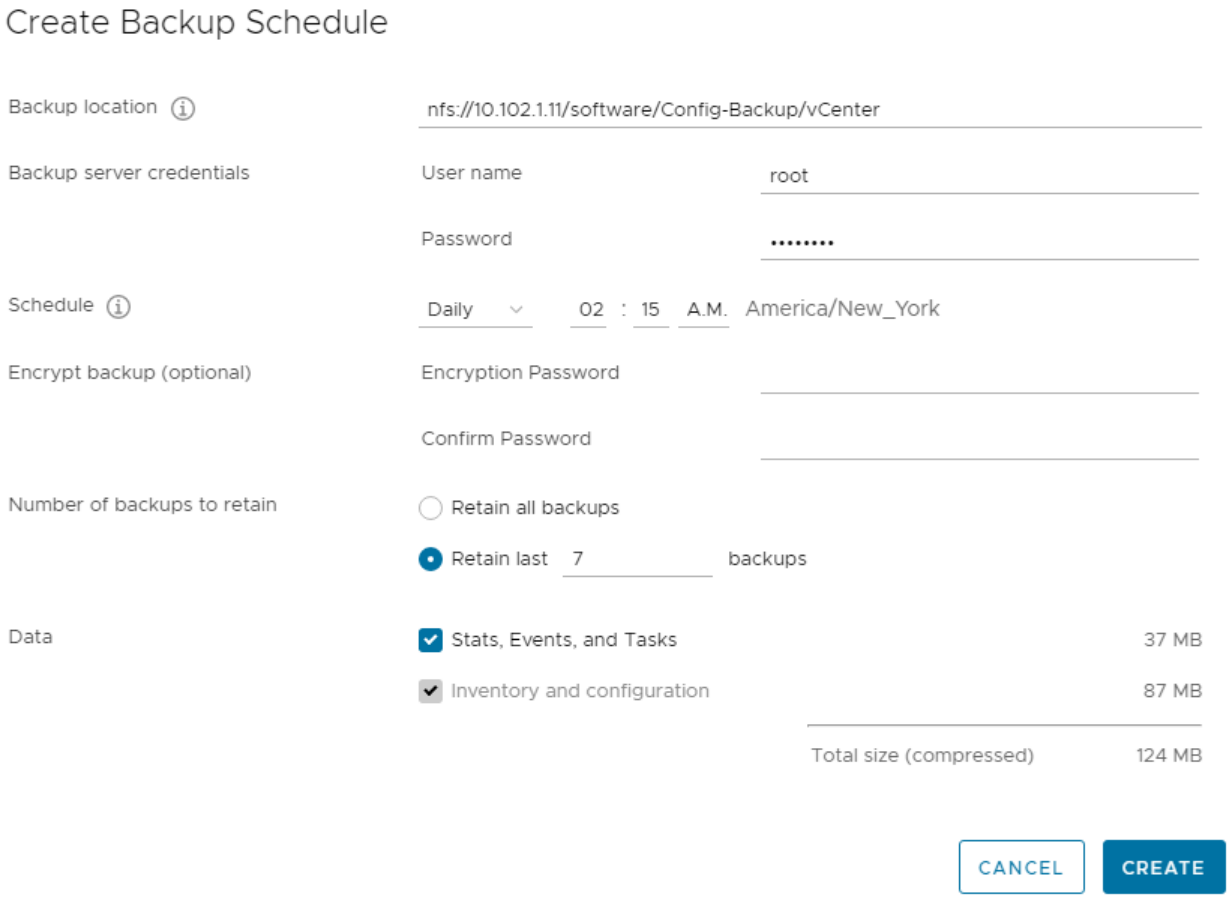


6. Click **CREATE**.

   The Backup Schedule Status should now show **Enabled**.
7. To test the backup setup, select **BACKUP NOW** and select "**Use backup location and user name from backup schedule**" to test the backup location.
8. You can initiate restoration with the backed-up files using the Restore function of the VCSA 7.0 Installer.

## Glossary of acronyms

**AAA**—Authentication, Authorization, and Accounting

**ACP**—Access-Control Policy

**ACI**—Cisco Application Centric Infrastructure

**ACK**—Acknowledge or Acknowledgement

**ACL**–Access-Control List

**AD**–Microsoft Active Directory

**AFI**–Address Family Identifier

**AMP**–Cisco Advanced Malware Protection

**AP**–Access Point

**API**–Application Programming Interface

**APIC**– Cisco Application Policy Infrastructure Controller (ACI)

**ASA**–Cisco Adaptive Security Appliance

**ASM**–Any-Source Multicast (PIM)

**ASR**–Aggregation Services Router

**Auto-RP**–Cisco Automatic Rendezvous Point Protocol (multicast)

**AVC**–Application Visibility and Control

**BFD**–Bidirectional Forwarding Detection

**BGP**–Border Gateway Protocol

**BMS**–Building Management System

**BSR**–Bootstrap Router (multicast)

**BYOD**–Bring Your Own Device

**CAPWAP**–Control and Provisioning of Wireless Access Points Protocol

**CDP**–Cisco Discovery Protocol

**CEF**–Cisco Express Forwarding

**CMD**–Cisco Meta Data

**CPU**–Central Processing Unit

**CSR**–Cloud Services Routers

**CTA**–Cognitive Threat Analytics

**CUWN**–Cisco Unified Wireless Network

**CVD**–Cisco Validated Design

**CYOD**–Choose Your Own Device

**DC**–Data Center

**DHCP**–Dynamic Host Configuration Protocol

**DM**–Dense-Mode (multicast)

**DMVPN**–Dynamic Multipoint Virtual Private Network

**DMZ**–Demilitarized Zone (firewall/networking construct)

**DNA**–Cisco Digital Network Architecture

**DNS**–Domain Name System

**DORA**–Discover, Offer, Request, ACK (DHCP process)

**DWDM**–Dense Wavelength Division Multiplexing

**ECMP**–Equal Cost Multi Path

**EID**–Endpoint Identifier

**EIGRP**–Enhanced Interior Gateway Routing Protocol

**EMI**–Electromagnetic Interference

**ETR**–Egress Tunnel Router (LISP)

**EVPN**–Ethernet Virtual Private Network (BGP EVPN with VXLAN data plane)

**FHR**–First-Hop Router (multicast)

**FHRP**–First-Hop Redundancy Protocol

**FMC**–Cisco Firepower Management Center

**FTD**–Cisco Firepower Threat Defense

**GBAC**–Group-Based Access Control

**Gbps**–Gigabits Per Second (interface/port speed reference)

**GRE**–Generic Routing Encapsulation

**GRT**–Global Routing Table

**HA**–High Availability

**HQ**–Headquarters

**HSRP**–Cisco Hot-Standby Routing Protocol

**HTDB**–Host-tracking Database (SD-Access control plane node construct)

**IBNS**–Identity-Based Networking Services (IBNS 2.0 is the current version)

**ICMP**– Internet Control Message Protocol

**IDF**–Intermediate Distribution Frame; essentially a wiring closet

**IEEE**–Institute of Electrical and Electronics Engineers

**IETF**–Internet Engineering Task Force

**IGP**–Interior Gateway Protocol

**IID**–Instance-ID (LISP)

**IoE**–Internet of Everything

**IoT**–Internet of Things

**IP**–Internet Protocol

**IPAM**–IP Address Management

**IPS**–Intrusion Prevention System

**IPSec**–Internet Protocol Security

**ISE**–Cisco Identity Services Engine

**ISR**–Integrated Services Router

**IS-IS**–Intermediate System-to-Intermediate System routing protocol

**ITR**–Ingress Tunnel Router (LISP)

**LACP**–Link Aggregation Control Protocol

**LAG**–Link Aggregation Group

**LAN**–Local Area Network

**L2 VNI**–Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN

**L3 VNI**– Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF

**LHR**–Last-Hop Router (multicast)

**LISP**–Location Identifier Separation Protocol

**MAC**–Media Access Control Address (OSI Layer 2 address)

**MAN**–Metro Area Network

**MEC**–Multichassis EtherChannel, sometimes referenced as ***MCEC***

**MDF**–Main Distribution Frame; essentially the central wiring point of the network.

**MnT**–Monitoring and Troubleshooting Node (Cisco ISE persona)

**MOH**–Music on Hold

**MPLS**–Multiprotocol Label Switching

**MR**–Map-resolver (LISP)

**MS**–Map-server (LISP)

**MSDP**–Multicast Source Discovery Protocol (multicast)

**MTU**–Maximum Transmission Unit

**NAC**–Network Access Control

**NAD**–Network Access Device

**NAT**–Network Address Translation

**NBAR**–Cisco Network-Based Application Recognition (NBAR2 is the current version)

**NFV**–Network Functions Virtualization

**NSF**–Non-Stop Forwarding

**OSI**–Open Systems Interconnection model

**OSPF**—Open Shortest Path First routing protocol

**OT**—Operational Technology

**PAgP**—Port Aggregation Protocol

**PAN**—Primary Administration Node (Cisco ISE persona)

**PCI DSS**—Payment Card Industry Data Security Standard

**PD**—Powered Devices (PoE)

**PETR**—Proxy-Egress Tunnel Router (LISP)

**PIM**—Protocol-Independent Multicast

**PITR**—Proxy-Ingress Tunnel Router (LISP)

**PoE**—Power over Ethernet (Generic term, may also refer to IEEE 802.3af, 15.4W at PSE)

**PoE+**—Power over Ethernet Plus (IEEE 802.3at, 30W at PSE)

**PSE**—Power Sourcing Equipment (PoE)

**PSN**—Policy Service Node (Cisco ISE persona)

**pxGrid**—Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)

**PxTR**—Proxy-Tunnel Router (LISP – device operating as both a PETR and PITR)

**QoS**—Quality of Service

**RADIUS**—Remote Authentication Dial-In User Service

**REST**—Representational State Transfer

**RFC**—Request for Comments Document (IETF)

**RIB**—Routing Information Base

**RLOC**—Routing Locator (LISP)

**RP**—Rendezvous Point (multicast)

**RP**—Redundancy Port (WLC)

**RP**—Route Processer

**RPF**—Reverse Path Forwarding

**RR**—Route Reflector (BGP)

**RTT**—Round-Trip Time

**SA**—Source Active (multicast)

**SAFI**—Subsequent Address Family Identifiers (BGP)

**SD**—Software-Defined

**SDA**—Cisco Software-Defined Access

**SDN**—Software-Defined Networking

**SFP**–Small Form-Factor Pluggable (1G transceiver)

**SFP+**– Small Form-Factor Pluggable (10G transceiver)

**SGACL**–Security-Group ACL

**SGT–**Scalable Group Tag, sometimes referenced as Security Group Tag

**SM**–Spare-mode (multicast)

**SNMP**–Simple Network Management Protocol

**SSID**–Service Set Identifier (wireless)

**SSM**–Source-Specific Multicast (PIM)

**SSO**–Stateful Switchover

**STP**–Spanning Tree Protocol

**SVI**–Switched Virtual Interface

**SVL**–Cisco StackWise® Virtual

**SWIM**–Software Image Management

**SXP**–Scalable Group Tag Exchange Protocol

**Syslog**–System Logging Protocol

**TACACS+**–Terminal Access Controller Access-Control System Plus

**TCP**–Transmission Control Protocol (OSI Layer 4)

**UCS**– Cisco Unified Computing System™ (Cisco UCS®)

**UDP**–User Datagram Protocol (OSI Layer 4)

**UPoE**–Cisco Universal Power over Ethernet (60W at PSE)

**UPoE+**– Cisco Universal Power over Ethernet Plus (90W at PSE)

**URL**–Uniform Resource Locator

**VLAN**–Virtual Local Area Network

**VM**—Virtual Machine

**VN**–Virtual Network, analogous to a VRF in SD-Access

**VNI**–Virtual Network Identifier (VXLAN)

**vPC**–virtual Port Channel (Cisco Nexus)

**VPLS**–Virtual Private LAN Service

**VPN**–Virtual Private Network

**VPNv4**–BGP address family that consists of a Route-Distinguisher (RD) prepended to an IPv4 prefix

**VPWS**–Virtual Private Wire Service

**VRF**–Virtual Routing and Forwarding

**VSL**–Virtual Switch Link (Cisco VSS component)

**VSS**–Cisco Virtual Switching System

**VXLAN**–Virtual Extensible LAN

**WAN**–Wide-Area Network

**WLAN**–Wireless Local Area Network (generally synonymous with IEEE 802.11-based networks)

**WoL**–Wake-on-LAN

**xTR**–Tunnel Router (LISP - device operating as both an ETR and ITR)

## Glossary of terms

This glossary addresses some terms used in this document to aid in understanding. It is not a complete list of all multicloud terminology. Some Cisco product links are supplied here also, where considered useful for the purposes of clarity, but it is by no means intended to be a complete list of all applicable Cisco products.

| | |
|---|---|
| **aaS/XaaS (IT capability provided as a Service)** | Some IT capability, X, provided as a service (XaaS). Some benefits follow:<br>• The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it.<br>• There are very low barriers to entry, so that services can be quickly adopted and dropped in response to business demand, without the penalty of inefficiently used CapEx.<br>• The service charge is an IT OpEx cost (pay-as-you-go), whereas the CapEx and the service infrastructure is the responsibility of the provider.<br>• Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes.<br>Such services are typically implemented as "microservices," which are accessed through Representational State Transfer (REST) APIs. This architectural style supports composition of service components into systems. Access to and management of aaS assets is through a web GUI and/or APIs, such that you can use Infrastructure-as-code (IaC) techniques d for automation, for example, Ansible and Terraform.<br>The provider can be any entity capable of implementing an aaS "cloud-native" architecture. The cloud-native architecture concept is well-documented and supported by open-source software and a rich ecosystem of services such as training and consultancy. The provider can be an internal IT department or any of many third-party companies using and supporting the same open-source platforms.<br>You can map service access control, integrated with corporate IAM, to specific users and business activities, enabling consistent policy controls across services, from wherever they are delivered. |
| **Ansible** | This infrastructure automation tool is used to implement processes for instantiating and configuring IT service components such as virtual machines on an IaaS platform. It supports the consistent execution of processes defined in YAML "playbooks" at scale, across multiple targets. Because the Ansible artefacts (playbooks) are text-based, you store them in a Source Code Management (SCM) system, such as GitHub. This setup allows for software development-like processes to be applied to infrastructure automation, such as, Infrastructure-as-code (IaC 0 (refer to this term later in this glossary).<br>https://www.ansible.com |
| **AWS (Amazon Web Services)** | Amazon Web Services, a provider of IaaS and PaaS<br>https://aws.amazon.com |
| **Azure** | Azure provides Microsoft IaaS and PaaS.<br>https://azure.microsoft.com/en-gb/ |

| | |
|---|---|
| **Co-located data center** | "A colocation center (CoLo)...is a type of data center where equipment, space, and bandwidth are available for rental to retail customers. Colocation facilities provide space, power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity."<br><br>https://en.wikipedia.org/wiki/Colocation_centre |
| **Containers**<br><br>**(Docker)** | A (Docker) container is a means to create a package of code for an application and its dependencies, such that the application can run on different platforms that support the Docker environment. In the context of aaS, microservices are typically packaged within Linux containers orchestrated by Kubernetes (K8s).<br><br>https://www.docker.com<br><br>https://www.cisco.com/c/en/us/products/cloud-systems-management/container-platform/index.html |
| **DevOps** | The underlying principle of DevOps is that the application development and operations teams should work closely together, ideally within the context of a toolchain that automates the stages of development, test, deployment, monitoring, and problem handling. DevOps is closely aligned with IaC, continuous integration and deployment (CI/CD), and Agile software development practices.<br><br>https://en.wikipedia.org/wiki/DevOps<br><br>https://en.wikipedia.org/wiki/CI/CD |
| **Edge compute** | Edge compute is the idea that it can be more efficient to process data at the edge of a network, close to the endpoints that originate that data, or to provide virtualized access services, such as at the network edge. Reasons could be related to low-latency response, reduction of the amount of unprocessed data being transported, efficiency of resource use, and so on. The generic label for this architecture is Multi-access Edge Computing (MEC), or Mobile Edge Computing for mobile networks specifically.<br><br>From an application experience perspective, it is important to be able to use, at the edge, the same operations model, processes, and tools used for any other compute node in the system.<br><br>https://en.wikipedia.org/wiki/Mobile_edge_computing |
| **IaaS**<br>**(Infrastructure as-a-Service)** | Infrastructure components provide aaS, located in data centers operated by a provider, typically accessed over the public Internet. IaaS provides a base platform for the deployment of workloads, typically with containers and Kubernetes (K8s). |
| **IaC**<br>**(Infrastructure as-Code)** | Given the ability to automate aaS through APIs, the implementation of the automation is typically with Python code, Ansible playbooks, and similar languages. These automation artefacts are programming code that define how the services are consumed. As such, they can be subject to the same code-management and software-development regimes as any other body of code, meaning that infrastructure automation can be subject to all of the quality and consistency benefits, CI/CD, traceability, automated testing, compliance checking, and so on, that could be applied to any coding project.<br><br>https://en.wikipedia.org/wiki/Infrastructure_as_code |
| **IAM**<br>**(Identity and Access Management)** | IAM is the means to control access to IT resources so that only those explicitly authorized to access given resources can do so. IAM is an essential foundation to a secure multicloud environment.<br><br>https://en.wikipedia.org/wiki/Identity_management |
| **IBM**<br>**(Cloud)** | IBM provided IaaS and PaaS.<br><br>https://www.ibm.com/cloud |

| | |
|---|---|
| **Cisco Intersight** | Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.<br>https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html |
| **GCP (Google Cloud Platform)** | Google provided IaaS and PaaS.<br>https://cloud.google.com/gcp |
| **Kubernetes (K8s)** | Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications.<br>https://kubernetes.io |
| **Microservices** | A microservices architecture is characterized by processes implementing fine-grained services, typically exposed by REST APIs; it can be composed into systems. The processes are often container-based, and the instantiation of the services are often managed with Kubernetes. Microservices managed in this way are intrinsically well-suited for deployment into IaaS environments, and as such, are the basis of a cloud native architecture.<br>https://en.wikipedia.org/wiki/Microservices |
| **PaaS (Platform-as-a-Service)** | PaaS is a layer of value-add services, typically for application development, deployment, monitoring, and general lifecycle management. The use of IaC with IaaS and PaaS is closely associated with DevOps practices. |
| **Private on-premises data center** | A private on-premises data center describes infrastructure housed within an environment owned by a given enterprise; it is distinguished from other forms of data centers, with the implication that the private data center is more secure, given that access is restricted to those that the enterprise authorizes. Thus, circumstances can arise where very sensitive IT assets are deployed only in a private data center, in contrast to using public IaaS. For many intents and purposes, the underlying technology can be identical, allowing for hybrid deployments where some IT assets are privately deployed but also accessible to other assets in public IaaS. IAM, VPNs, firewalls, and similar technologies are critical to underpin the security of such an arrangement. |
| **REST API** | Representational State Transfer (REST) APIs is a generic term for APIs accessed over HTTP(S), typically transporting data encoded in JSON or XML. REST APIs have the advantage that they support distributed systems, communicating over HTTP, which is a well-understood protocol from a security management perspective. REST APIs are another element of a cloud-native applications architecture, alongside microservices.<br>https://en.wikipedia.org/wiki/Representational_state_transfer |
| **SaaS (Software-as-a-Service)** | End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider. |
| **SAML (Security Assertion Markup Language)** | Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by the aaS for access control decisions.<br>https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language |
| **Terraform** | An IaC software tool for cloud services, based on declarative configuration files.<br>https://www.terraform.io |

## Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community at Cisco Developed UCS Integrations.

**Americas Headquarters**
Cisco Systems, Inc.
San Jose, CA

**Asia Pacific Headquarters**
Cisco Systems (USA) Pte. Ltd.
Singapore

**Europe Headquarters**
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.



Printed in USA 223053.4 10/23