FlexPod Datacenter with End-to-End 100G, Cisco Intersight Managed Mode, using Infrastructure as Code (IaC), VMware 7U3, and NetApp ONTAP 9.11

Deployment Guide for FlexPod Datacenter with End-to-End 100G, Cisco Intersight Managed Mode, VMware 7U3, and NetApp ONTAP 9.11

Published: February 2023

In partnership with:

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to: http://www.cisco.com/go/designzone.

Executive Summary

The FlexPod Datacenter 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 driven 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 UCS 5th Generation components into the FlexPod Datacenter and the ability to manage FlexPod components from the cloud using Cisco Intersight. Some of the key advantages of integrating Cisco UCS 5th generation components into the FlexPod infrastructure are:

- **Simpler and programmable infrastructure:** infrastructure as code delivered through a single partner integrable open API

- **End-to-End 100Gbps Ethernet:** utilizing the 5th Generation Cisco UCS VIC 15231, the 5th Generation Cisco UCS 6536 Fabric Interconnect, and the UCSX-I-9108-100G Intelligent Fabric Module to deliver 100Gbps Ethernet from the server through the network to the storage

- **End-to-End 32Gbps Fibre Channel:** utilizing the 5th Generation Cisco UCS VIC 15231, the 5th Generation Cisco UCS 6536 Fabric Interconnect, and the UCSX-I-9108-100G Intelligent Fabric Module to deliver 32Gbps Ethernet from the server (via 100Gbps FCoE) through the network to the storage

- **Innovative cloud operations:** continuous feature delivery and no need for maintaining on-premise virtual machines supporting management functions

- **Built for investment protections:** design ready for future technologies such as liquid cooling and high-Wattage CPUs; 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 is organized as follows:

- [Introduction](#)
- [Audience](#)
- [Purpose of this Document](#)
- [What's New in this Release?](#)

Introduction

The Cisco Unified Compute System (Cisco UCS) with Intersight Managed Mode (IMM) X-Series 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 Soft-ware-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 X-Series enables the next-generation cloud-operated FlexPod infrastructure that not only simplifies data-center management but also allows the infra-structure to adapt to the unpredictable needs of modern applications as well as traditional workloads. With the Cisco Intersight platform, customers get all the benefits of SaaS delivery and the full lifecycle management of Inter-sight-connected distributed servers and integrated NetApp storage systems across data centers, remote sites, branch offices, and edge environments.

Audience

The intended audience of this document includes but is not limited to 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 deployment guidance using Ansible playbooks around incorporating the Cisco Intersight managed Cisco UCS X-Series platform with end-to-end 100Gbps within FlexPod Datacenter infrastructure. The document explains both configurations and best practices for a successful deployment. This deployment guide also highlights integration of VMware vCenter and NetApp Active IQ Unified Manager to Cisco Intersight to deliver 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 100Gbps Ethernet and 32Gbps Fibre Channel in FlexPod Datacenter
- Integration of the 5[th] Generation Cisco UCS 6536 Fabric Interconnect into FlexPod Datacenter
- Integration of the 5[th] Generation Cisco UCS 15000-series VICs into FlexPod Datacenter
- Integration of the Cisco UCSX-I-9108-100G Intelligent Fabric Module into the X-Series 9508 Chassis
- Integration of the Cisco UCS C225 and C245 M6 Servers with AMD EPYC CPUs

- Addition of the Non-Volatile Memory Express over Transmission Control Protocol (NVMe-TCP) Storage Protocol with NetApp ONTAP 9.11.1

- An integrated, more complete end-to-end Infrastructure as Code (IaC) Day 0 configuration of the FlexPod Infrastructure utilizing Ansible Scripts

- VMware vSphere 7.0 Update 3

- Integration with the FlexPod XCS Integrated System in Cisco Intersight

Deployment Hardware and Software

This chapter is organized as follows:

Design Requirements

The FlexPod Datacenter with Cisco UCS and Cisco 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 can be replicated to expand and grow as the needs of the business grow

- Flexible design that can support different models of various components with ease

- Simplified design with ability to integrate and automate with external automation tools

- Cloud-enabled design which can be configured, managed, and orchestrated from the cloud using GUI or APIs

To deliver a solution which meets all these design requirements, various solution components are connected and configured as covered in the upcoming sections.

Physical Topology

The FlexPod Datacenter solution with end-to-end 100Gbps ethernet is built using the following hardware components:

- Cisco UCS X9508 Chassis with Cisco UCSX-I-9108-100G Intelligent Fabric Modules (IFMs) and up to eight Cisco UCS X210c M6 Compute Nodes with 3rd Generation Intel Xeon Scalable CPUs

- Fifth-generation Cisco UCS 6536 Fabric Interconnects to support 100GbE, 25GbE, and 32GFC connectivity from various components

- Cisco UCS C225 M6 and C245 M6 rack mount servers with AMD EPYC CPUs

- High-speed Cisco NX-OS-based Nexus 93360YC-FX2 switching design to support up to 100GE and 32GFC connectivity

- NetApp AFF A800/A400 end-to-end NVMe storage with 100G Ethernet and (optional) 32G Fibre Channel connectivity

- Cisco MDS 9132T* switches to support Fibre Channel storage configuration

Note: * Cisco MDS 9132T and FC connectivity is not needed when implementing IP-based connectivity design supporting iSCSI boot from SAN, NFS, and NVMe-TCP.

The software components of this solution consist of:

- Cisco Intersight SaaS platform to deploy, maintain and support the FlexPod components

- Cisco Intersight Assist Virtual Appliance to help connect NetApp ONTAP, VMware vCenter, and Cisco Nexus and MDS switches with Cisco Intersight

- NetApp Active IQ Unified Manager to monitor and manage the storage and for NetApp ONTAP integration with Cisco Intersight

- VMware vCenter to set up and manage the virtual infrastructure as well as Cisco Intersight integration

FlexPod Datacenter for IP-based Storage Access

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

Figure 1. FlexPod Datacenter Physical Topology for IP-based Storage Access

The reference hardware configuration includes:

- Two Cisco Nexus 93360YC-FX2 Switches in Cisco NX-OS mode provide the switching fabric.

- Two Cisco UCS 6536 Fabric Interconnects (FI) provide the chassis connectivity. One 100 Gigabit Ethernet port from each FI, configured as a Port-Channel, is connected to each Nexus 93360YC-FX2.

- One Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCSX 9108-100G Intelligent Fabric Modules (IFMs), where four 100 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI. If additional bandwidth is required, all eight 100G ports can be utilized.

- One NetApp AFF A800 HA pair connects to the Cisco Nexus 93360YC-FX2 Switches using two 100 GE ports from each controller configured as a Port-Channel.

- Two (one shown) UCS C245 rack mount servers connect to the Fabric Interconnects using two 100 GE ports per server

- Two (one shown) UCS C225 rack mount servers connect to the Fabric Interconnects via breakout using four 25 GE ports per server

FlexPod Datacenter for FC-based Storage Access

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

Figure 2.   FlexPod Datacenter Physical Topology for FC-based Storage Access

The reference hardware configuration includes:

- Two Cisco Nexus 93360YC-FX2 Switches in Cisco NX-OS mode provide the switching fabric.
- Two Cisco UCS 6536 Fabric Interconnects (FI) provide the chassis connectivity. One 100 Gigabit Ethernet port from each FI, configured as a Port-Channel, is connected to each Nexus 93360YC-FX2. Four FC ports are connected to the Cisco MDS 9132T switches via breakout using 32-Gbps Fibre Channel connections configured as a single port channel for SAN connectivity.
- One Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCSX 9108-100G Intelligent Fabric Modules (IFMs), where four 100 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI. If additional bandwidth is required, all eight 100G ports can be utilized.
- One NetApp AFF A800 HA pair connects to the Cisco Nexus 93360YC-FX2 Switches using two 100 GE ports from each controller configured as a Port-Channel. Two 32Gbps FC ports from each controller are connected to each Cisco MDS 9132T for SAN connectivity.
- Two (one shown) Cisco UCS C245 rack mount servers connect to the Fabric Interconnects using two 100 GE ports per server
- Two (one shown) Cisco UCS C225 rack mount servers connect to the Fabric Interconnects via breakout using four 25 GE ports per server

Note: 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: https://docs.netapp.com/us-en/ontap-systems/index.html

VLAN Configuration

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

Table 1.   VLAN Usage

VLAN ID	Name	Usage	IP Subnet used in this deployment
2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1).	
1020	OOB-MGMT-VLAN	Out-of-band management VLAN to connect management ports for various devices	10.102.0.0/24; GW: 10.102.0.254
1021	IB-MGMT-VLAN	In-band management VLAN utilized for all in-band management connectivity - for example, ESXi hosts, VM management, and so on.	10.102.1.0/24; GW: 10.102.1.254
1022	VM-Traffic	VM data traffic VLAN	10.102.2.0/24; GW: 10.102.2.254
3050	NFS-VLAN	NFS VLAN for mounting datastores in ESXi servers for VMs	192.168.50.0/24 **

VLAN ID	Name	Usage	IP Subnet used in this deployment
3010*	iSCSI-A	iSCSI-A path for storage traffic including boot-from-san traffic	192.168.10.0/24 **
3020*	iSCSI-B	iSCSI-B path for storage traffic including boot-from-san traffic	192.168.20.0/24 **
3030	NVMe-TCP-A	NVMe-TCP-A path when using NVMe-TCP	192.168.30.0/24 **
3040	NVMe-TCP-B	NVMe-TCP-B path when using NVMe-TCP	192.168.40.0/24 **
3000	vMotion	VMware vMotion traffic	192.168.0.0/24 **

* iSCSI VLANs are not required if using FC storage access.

** IP gateway is not needed since no routing is required for these subnets

Some of the key highlights of VLAN usage are as follows:

- VLAN 1020 allows customers to manage and access out-of-band management interfaces of various devices.
- VLAN 1021 is used for in-band management of VMs, ESXi hosts, and other infrastructure services
- VLAN 3050 provides ESXi hosts access to the NFS datastores hosted on the NetApp Controllers for deploying VMs.
- A pair of iSCSI VLANs (3010 and 3020) is configured to provide access to boot LUNs for ESXi hosts. These VLANs are not needed if customers are using FC-only connectivity.
- A pair of NVMe-TCP VLANs (3030 and 3040) is configured to provide access to NVMe datastores when NVMe-TCP is being used
- VLAN 3000 is used for VM vMotion

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

Table 2.   Virtual Machines

Virtual Machine Description	VLAN	IP Address	Comments
vCenter Server	1021	10.102.1.100	Hosted on either pre-existing management infrastructure or on FlexPod
NetApp ONTAP Tools for VMware	1021	10.102.1.99	Hosted on FlexPod

Virtual Machine Description	VLAN	IP Address	Comments
vSphere			
NetApp SnapCenter Plug-in for VMware vSphere	1021	10.102.1.98	Hosted on FlexPod
NetApp Active IQ Unified Manager	1021	10.102.1.97	Hosted on FlexPod
Cisco Intersight Assist	1021	10.102.1.96	Hosted on FlexPod

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(2c)	Cisco UCS GA release for infrastructure including FIs and IOM/IFM.
Network	Cisco Nexus 93360YC-FX2 NX-OS	10.2(3)F	
	Cisco MDS 9132T	9.2(2)	Requires SMART Licensing
Storage	NetApp AFF A800/A400	ONTAP 9.11.1P2	
Software	Cisco UCS X210c	5.0(2d)	Cisco UCS X-series GA release for compute nodes
	Cisco UCS C225/245 M6	4.2(2f)	
	Cisco Intersight Assist Appliance	1.0.9-456	1.0.9-342 initially installed and then automatically upgraded
	VMware vCenter	7.0 Update 3h	Build 20395099

Layer	Device	Image Bundle	Comments
	VMware ESXi	7.0 Update 3d	Build 19482537 included in Cisco Custom ISO
	VMware ESXi nfnic FC Driver	5.0.0.34	Supports FC-NVMe
	VMware ESXi nenic Ethernet Driver	1.0.42.0	
	NetApp ONTAP Tools for VMware vSphere	9.11	Formerly Virtual Storage Console (VSC)
	NetApp NFS Plug-in for VMware VAAI	2.0	
	NetApp SnapCenter for vSphere	4.7	Includes the vSphere plug-in for SnapCenter
	NetApp Active IQ Unified Manager	9.11P1	

FlexPod Cabling

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

The cabling diagram in this section contains the details for the prescribed and supported configuration of the NetApp AFF 800 running NetApp ONTAP 9.11.1.

Note: For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

Note: This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

Note: Be sure to use the cabling directions in this section as a guide.

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 fabric interconnect. Four 32Gb uplinks via breakout connect as port-channels from each Cisco UCS Fabric Interconnect to the MDS switches, and a total of eight 32Gb links connect the MDS switches to the NetApp AFF controllers. Also, 100Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the NetApp AFF controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (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**
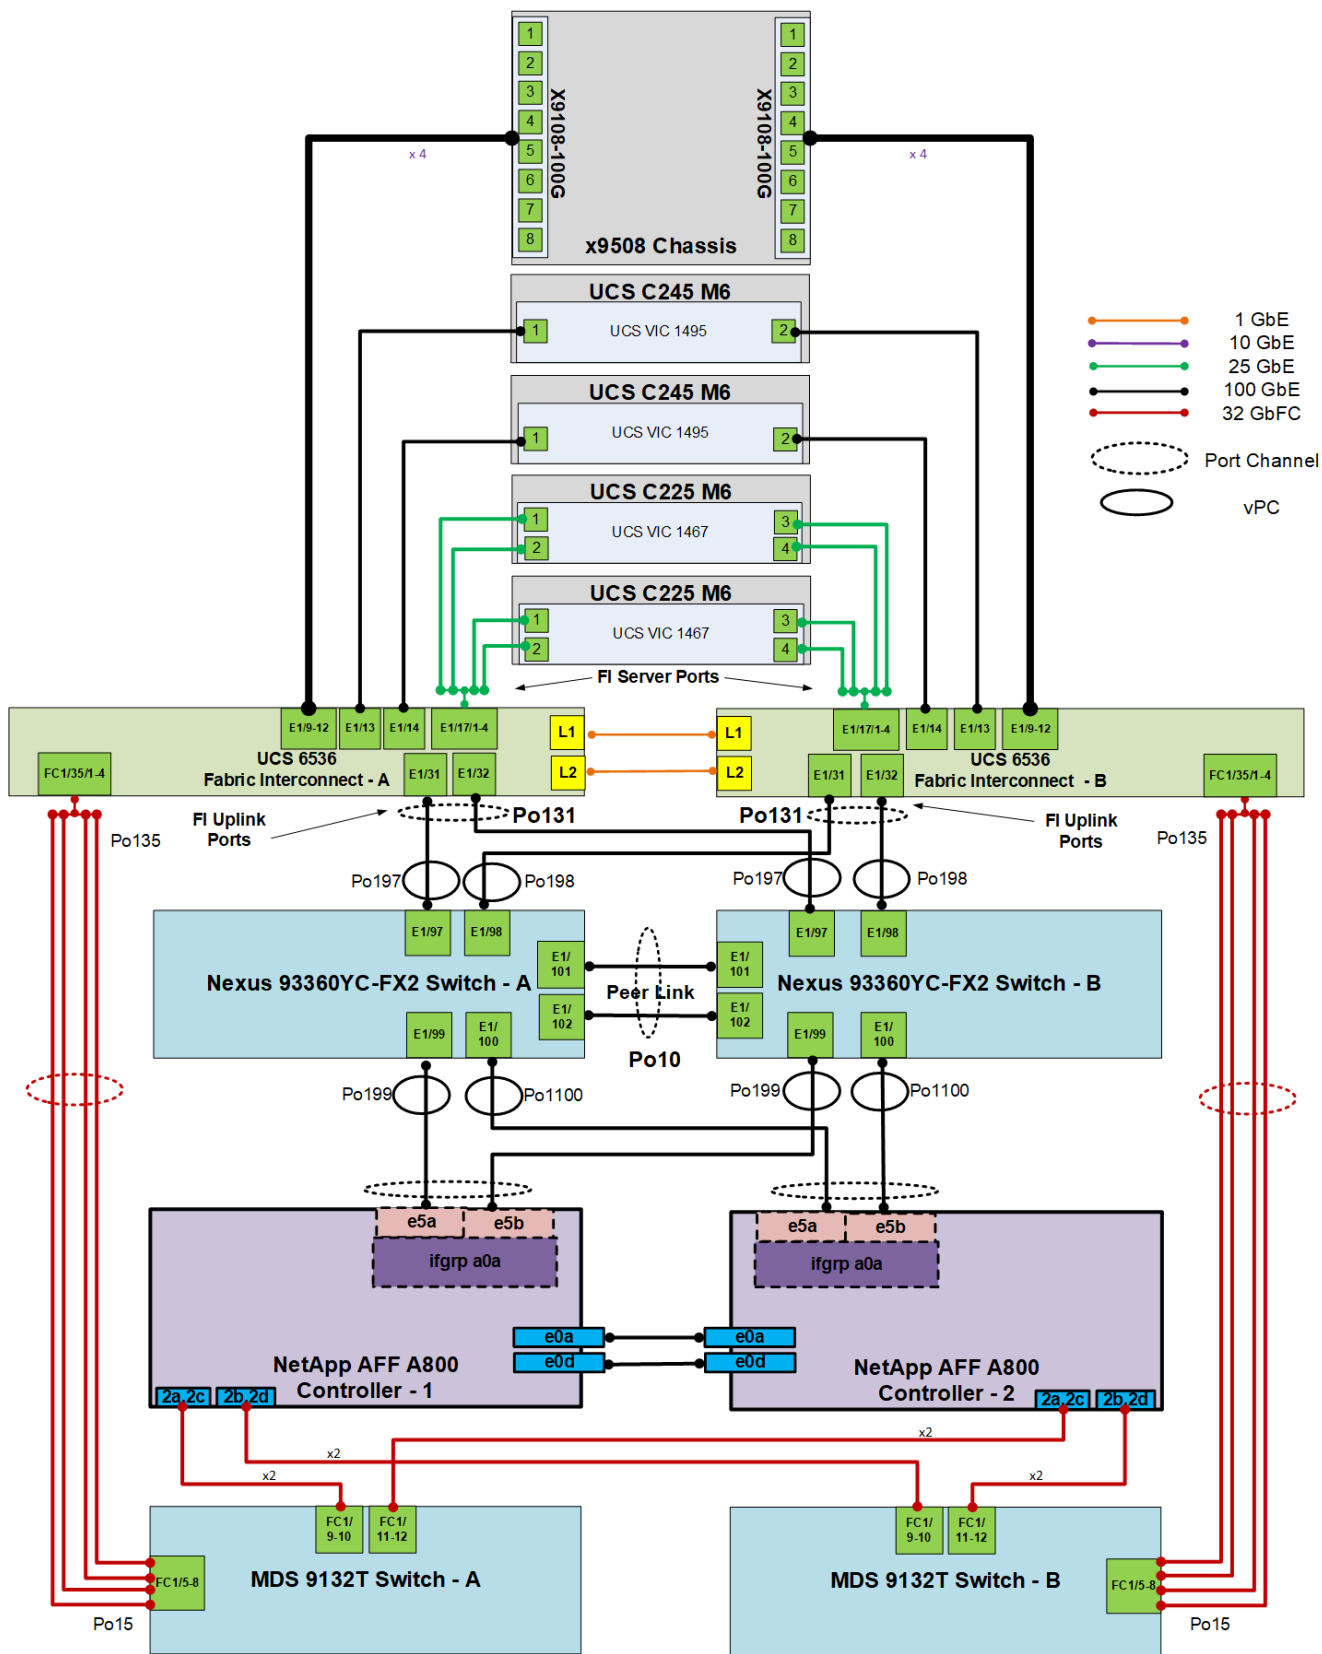
Ansible Automation Workflow and Solution Deployment

If using the published Ansible playbooks to configure the FlexPod infrastructure, complete this section of the document. If completing a manual configuration, skip to the next section of the document. The Ansible automated FlexPod solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, NetApp ONTAP Storage, Cisco UCS, Cisco MDS, and VMware ESXi.

Figure 4 illustrates the FlexPod solution implementation workflow which is explained in the following sections. The FlexPod infrastructure layers are first configured in the order illustrated.

Figure 4.   Ansible Automation Workflow
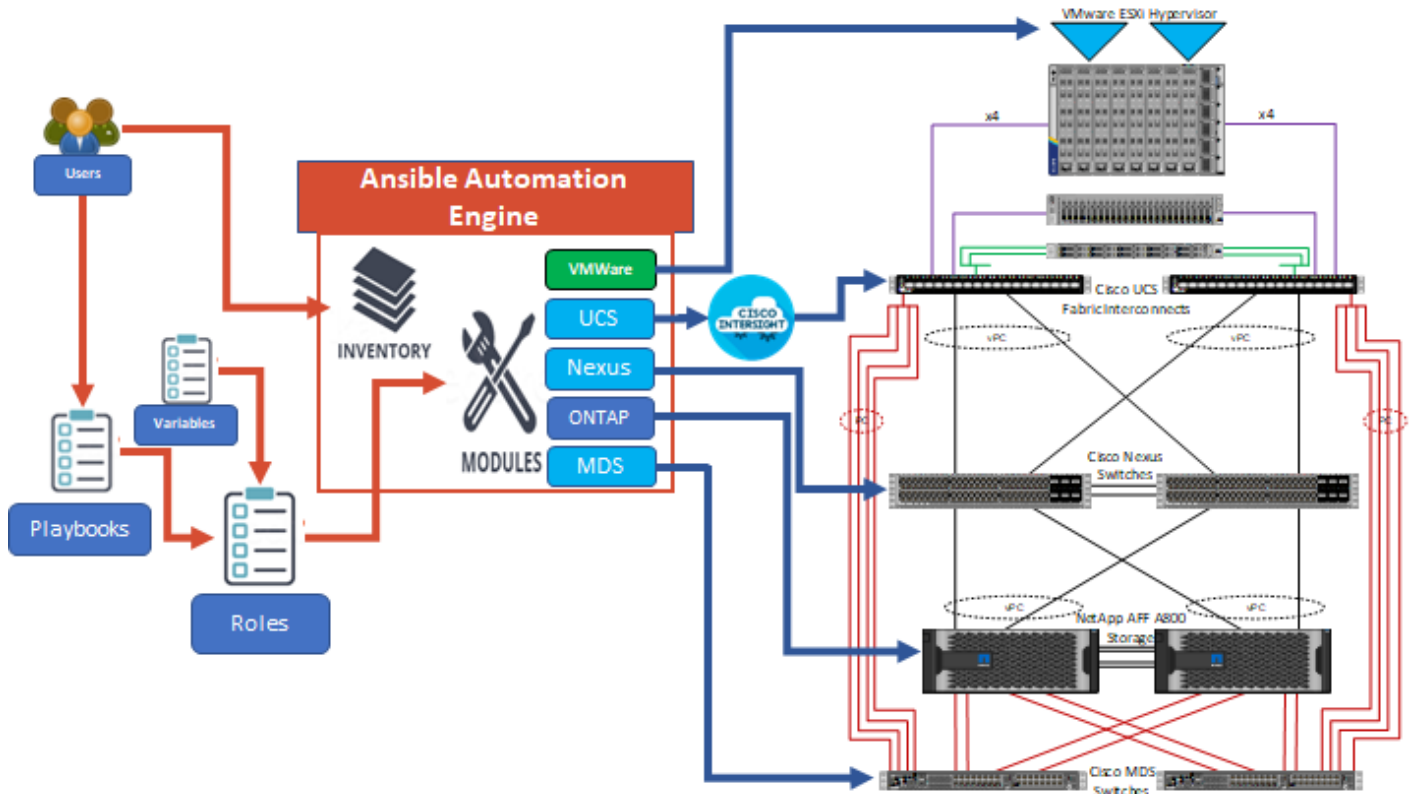
Prerequisites

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

- Ansible Community Documentation: https://docs.ansible.com/ansible_community.html
- To use the Ansible playbooks demonstrated in this document, the management workstation must also have a working installation of Git and access to the Cisco DevNet public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following links:

- ◦ Cisco DevNet (available in early 2023):
  https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/FlexPod-IMM-4.2.2
- ◦ GitHub repository: https://github.com/ucs-compute-solutions/FlexPod-IMM-4.2.2

- The Cisco Nexus and MDS Switches, NetApp Storage, and Cisco UCS must be physically racked, cabled, powered, and configured with management IP addresses before the Ansible-based installation procedure can begin as shown in the cabling diagram (Figure 3). If necessary, upgrade the Cisco Nexus Switches to release 10.2(3)F, and the Cisco MDS Switches to release 9.2(2).

- Before running each Ansible Playbook to setup the Network, Storage, Cisco UCS, and VMware ESXi various variables have to be updated based on the customers environment and specific implementation with values such as the VLANs, pools and ports on Cisco UCS, IP addresses for NFS, iSCSI, and NVMe-TCP interfaces and values needed for VMware ESXi.

- Day 2 Configuration tasks such as adding datastores or ESXi servers can be performed manually or with Cisco Intersight Cloud Orchestrator (ICO).

Procedure 1. Prepare Management Workstation (Control Machine)

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

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

Step 1.   Install the EPEL repository on the management host.

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

Step 2.   Install Ansible engine.

```
dnf install ansible
```

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

```
ansible --version
ansible [core 2.13.5]
  config file = /etc/ansible/ansible.cfg
  configured module search path = ['/root/.ansible/plugins/modules', '/usr/share/ansible/plugins/modules']
  ansible python module location = /usr/lib/python3.9/site-packages/ansible
  ansible collection location = /root/.ansible/collections:/usr/share/ansible/collections
  executable location = /usr/bin/ansible
  python version = 3.9.13 (main, Jun 24 2022, 15:32:51) [GCC 8.5.0 20210514 (Red Hat 8.5.0-13)]
  jinja version = 3.1.2
  libyaml = True
```

Step 4.   Install NetApp specific python modules.

```
pip3 install netapp-lib
```

Step 5.   Install ansible-galaxy collections and other dependencies for Cisco Nexus (and MDS), NetApp ONTAP, Cisco UCS, and VMware as follows:

```
ansible-galaxy collection install cisco.intersight
ansible-galaxy collection install cisco.nxos
pip3 install ansible-pylibssh
ansible-galaxy collection install netapp.ontap
ansible-galaxy collection install community.vmware
pip3 install wheel
```

```
pip3 install --upgrade pip setuptools
pip3 install -r ~/.ansible/collections/ansible_collections/community/vmware/requirements.txt
```

Note: The cisco.nxos collection is used for both Cisco Nexus and Cisco MDS configuration.

Procedure 2.   Clone GitHub Collection

Note: You need to use a GitHub repository from one public location; the first step in the process is to clone the GitHub collection named FlexPod-IMM-4.2.2 (https://github.com/ucs-compute-solutions/FlexPod-IMM-4.2.2.git) to a new empty folder on the management workstation. Cloning the repository creates a local copy, which is then used to run the playbooks that have been created for this solution.

Step 1.   From the management workstation, create a new folder for the project. The GitHub collection will be cloned in a new folder inside this one, named /root/FlexPod-IMM-4.2.2.

Step 2.   Open a command-line or console interface on the management workstation and change directories to the new folder just created.

Step 3.   Clone the GitHub collection using the following command:

```
git clone https://github.com/ucs-compute-solutions/FlexPod-IMM-4.2.2.git
```

Step 4.   Change directories to the new folder named FlexPod-IMM-4.2.2.

Network Switch Configuration

This chapter is organized as follows:

- [Physical Connectivity](#)
- [Initial Configuration](#)
- [Ansible Nexus Switch Configuration](#)

This chapter provides a detailed procedure for configuring the Cisco Nexus 93360YC-FX2 switches for use in a FlexPod environment. The Cisco Nexus 93360YC-FX2 will be used for LAN switching in this solution.

Note: The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 10.2(3)F.

- If using the Cisco Nexus 93360YC-FX2 switches or other Cisco Nexus switches for both LAN and SAN switching, please refer to section [FlexPod with Cisco Nexus 93360YC-FX2 SAN Switching Configuration](#) in the Appendix.

- The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

- This procedure sets up and uplink virtual port channel (vPC) with the IB-MGMT and OOB-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 as explained in section [FlexPod Cabling](#).

Initial Configuration

The following procedures describe this basic configuration of the Cisco Nexus switches for use in the FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 10.2(3)F, the Cisco suggested Nexus switch release at the time of this validation.

Procedure 1. Set Up Initial Configuration from a serial console

Set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>.

Step 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
```

Step 2.   Review the configuration summary before enabling the configuration.

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

Step 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.

Ansible Nexus Switch Configuration

Procedure 1.   Configure the Cisco Nexus switches from the management workstation

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

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

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

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/n9kA.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/n9kB.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/NEXUSconfig/defaults/main.yml

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

Step 3.   From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_Nexus.yml Ansible playbook.

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

Step 4. Once the Ansible playbook has been run on both switches, it is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the time-zone and daylight savings time or summertime, please see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 10.2(x)](#). Sample clock commands for the United States Eastern timezone are:

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

Step 5. ssh into each switch and execute the following commands:

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

NetApp ONTAP Storage Configuration

This chapter contains the following:

- NetApp AFF A400/A800 Controllers
- Disk Shelves
- NetApp ONTAP 9.11.1P2

NetApp AFF A400/A800 Controllers

See the following 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 ONTAP version. It also provides configuration information for all the NetApp storage appliances currently supported by 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 ONTAP that you plan to install, follow these steps at the NetApp Support site.

Procedure 1.	Confirm hardware and software components

Step 1. Access the HWU application to view the System Configuration guides. Click the Platforms menu to view the compatibility between different versions of the ONTAP software and the NetApp storage appliances with your desired specifications.

Step 2. Alternatively, to compare components by storage appliance, click **Compare Storage Systems**.

Controllers

Follow the physical installation procedures for the controllers found 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 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 9.11.1P2

Complete Configuration Worksheet

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

Ansible NetApp ONTAP Storage Configuration

End to End ONTAP Storage Configuration for a FlexPod is automated with Ansible. ONTAP Storage can be deployed via Ansible after the ONTAP Cluster setup is complete and the Cluster management network is configured.

A playbook by the name 'Setup_ONTAP.yml' is available at the root of this repository. It calls all the required roles to complete the setup of the ONTAP storage system.

The ONTAP setup is split into three sections, use the tags - ontap_config_part_1, ontap_config_part_2, and ontap_config_part_3 to execute parts of the playbook at the appropriate stage of setup.

Execute the playbook from the Ansible Control machine as an admin/ root user using the following commands:

- After setup of Cisco Nexus switches and bringing the NetApp storage cluster online with aggregates: ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_1
- After setup of Cisco UCS and deploying server profiles: ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_2
- After setup of VMware vSphere 7.0 U3 Setup: ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_3

If you would like to run a part of the deployment, you may use the appropriate tag that accompanies each task in the role and run the playbook by running the following command:

```
ansible-playbook -i inventory Setup_ONTAP.yml -t <tag_name>
```

Configure ONTAP Nodes

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

Table 4.   ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>

Cluster Detail	Cluster Detail Value
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.11.1P2 URL (http server hosting ONTAP software)	<url-boot-software>

Procedure 1.   Configure Node 01

Step 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…
```

Step 2.   Allow the system to boot up.

```
autoboot
```

Step 3.   Press Ctrl-C when prompted.

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

Step 4.   To install new software, select option 7 from the menu.

Step 5.   Enter `y` to continue the installation.

Step 6.   Select `e0M` for the network port for the download.

Step 7.   Enter `n` to skip the reboot.

Step 8.   Select option 7 from the menu: `Install new software first`

Step 9.   Enter `y` to continue the installation.

Step 10.   Enter the IP address, netmask, and default gateway for `e0M`.

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

Step 11.   Enter the 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>
```

Step 12.   Press Enter for the user name, indicating no user name.

Step 13.   Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

Step 14.   Enter `y` to reboot the node now.

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

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


Rebooting...
Files /cfcard/x86_64/freebsd/image2/VERSION and /var/VERSION differ
.
Setting default boot image to image2...
done.
Uptime: 37m44s
```

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

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

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

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

Step 16. Select option 4 for Clean Configuration and Initialize All Disks.

Step 17. Enter `y` to zero disks, reset config, and install a new file system.

Step 18. Enter `yes` to erase all the data on the disks.

Note: When initialization and creation of 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.
https://docs.netapp.com/us-en/ontap/concepts/root-data-partitioning-concept.html

---

| **Procedure 2.** | Configure Node 02 |

Step 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…
```

Step 2. Allow the system to boot up.

```
autoboot
```

Step 3. Press Ctrl-C when prompted.

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

Step 4. To install new software, select option 7.

Step 5. Enter `y` to continue the installation.

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

Step 7.   Enter `n` to skip the reboot.

Step 8.   Select option 7: `Install new software first`

Step 9.   Enter `y` to continue the installation.

Step 10.   Enter the IP address, netmask, and default gateway for e0M.

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

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

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

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

Step 12.   Press `Enter` for the username, indicating no user name.

Step 13.   Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

Step 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 ONTAP installation a prompt to reboot the node requests a Y/N response.

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

        `Press Ctrl-C for Boot Menu`

Step 16.   Select option 4 for Clean Configuration and Initialize All Disks.

Step 17.   Enter `y` to zero disks, reset config, and install a new file system.

Step 18.   Enter `yes` to erase all the data on the disks.

Note: When initialization and creation of root aggregate is complete, the storage system reboots.   For more information about root aggregate and disk partitioning, please refer to the following ONTAP documentation on root-data partitioning. https://docs.netapp.com/us-en/ontap/concepts/root-data-partitioning-concept.html

Procedure 3.   Set Up Node

Step 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 ONTAP 9.11.1P2 boots on the node for the first time.

Step 2. Follow the prompts to set up node 01.

```
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]: 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:
```

Step 3. To complete cluster setup, open a web browser and navigate to `https://<node01-mgmt-ip>`.

Table 5.   Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
Cluster Admin SVM	<cluster-adm-svm>
Infrastructure Data SVM	<infra-data-svm>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>

Cluster Detail	Cluster Detail Value
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>
SNMPv3 User	<snmp-v3-usr>
SNMPv3 Authentication Protocol	<snmp-v3-auth-proto>
SNMPv3 Privacy Protocol	<snmpv3-priv-proto>

Note: Cluster setup can also be performed using the CLI. This document describes the cluster setup using the NetApp ONTAP System Manager guided setup.

Step 4.   Complete the required information on the Initialize Storage System screen:

Step 5. In the Cluster screen:

    a. Enter the cluster name and administrator password.

    b. Complete the Networking information for the cluster and each node.

    c. Check the box for Use Domain Name Service (DNS) and enter the IP addresses of the DNS servers in a comma separated list.

    d. Check the box for Use time services (NTP) and enter the IP addresses of the time servers in a comma separated list.

Note: Here, the DNS and NTP server manual configuration for the cluster is optional. Ansible scripts will configure the same when ONTAP playbook with the tag "ontap_config_part_1" is executed.

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

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

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

Step 6. Click **Submit**.

Step 7. A few minutes will pass while the cluster is configured. When prompted, login to ONTAP System Manager to continue the cluster configuration.

Step 8. Click **Prepare Storage** to create data aggregates.

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

Procedure 4. Ansible ONTAP Storage Configuration - Part 1

Step 1. Edit the following variable files to ensure proper ONTAP Storage variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/ontap

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/vars/ontap_main.yml

Step 2. From /root/ FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_ONTAP.yml Ansible playbook with the associated tag for this section:

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

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

Cisco Intersight Managed Mode Configuration

This chapter contains the following:

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 (also referred to as Cisco IMM or Intersight managed MODE) is a new architecture that manages Cisco Unified Computing System (Cisco UCS) fabric inter-connect–attached systems through a Redfish-based standard model. Cisco Intersight managed mode standard-

izes both policy and operation management for Cisco UCS C-Series M6 and Cisco UCSX X210c M6 compute nodes used in this deployment guide.

Cisco UCS B-Series M6 servers, connected and managed through Cisco UCS FIs, are also supported by IMM. 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

Procedure 1.   Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

The Cisco UCS fabric interconnects need to be set up to support Cisco Intersight managed mode. When converting an existing pair of Cisco UCS fabric interconnects from Cisco UCS Manager mode to Intersight Mange Mode (IMM), first erase the configuration and reboot your system.

Note: Converting fabric interconnects to Cisco Intersight managed mode is a disruptive process, and configuration information will be lost. Customers are encouraged to make a backup of their existing configuration. If a software version that supports Intersight Managed Mode (4.1(3) or later) is already installed on Cisco UCS Fabric Interconnects, do not upgrade the software to a recommended recent release using Cisco UCS Manager. The software upgrade will be performed using Cisco Intersight to make sure Cisco UCS X-series firmware is part of the software upgrade.

Step 1.   Configure Fabric Interconnect A (FI-A). On the Basic System Configuration Dialog screen, set the management mode to Intersight. The remaining settings are similar to 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:
1.  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
```

Step 2.   After applying the settings, make sure you can ping the fabric interconnect management IP address. When Fabric Interconnect A is correctly set up and is available, Fabric Interconnect B will automatically discover Fabric Interconnect A during its setup process as shown in the next step.

Step 3.   Configure Fabric Interconnect B (FI-B). For the configuration method, select console. Fabric Interconnect B will detect the presence of Fabric Interconnect A and will prompt you to enter the admin password for Fabric Interconnect A. Provide the management IP address for Fabric Interconnect 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

Step 1.   Go to https://intersight.com and click Create an account.

Step 2.   Read and accept the license agreement. Click **Next**.

Step 3.   Provide an Account Name and click **Create**.

Step 4.   On successful creation of the Intersight account, the following page will be displayed:

Select a service

Select a service to start your Intersight
Journey

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

Procedure 3. Set up Cisco Intersight Licensing

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

Step 1. Log into the Cisco Smart Licensing portal:
https://software.cisco.com/software/smart-licensing/alerts.

Step 2. Verify that the correct virtual account is selected.

Step 3. Under **Inventory** > **General**, generate a new token for product registration.

Step 4. Copy this newly created token.

Create Registration Token

This will create a token that is used to register product instances, so that they can use licenses from this virtual account. Once it's created, go to the Smart Licensing configuration for your products and enter the token, to register them with this virtual account.

Virtual Account: Cisco ▮ ◦◦◦ Intersight

Description : RTP IMM

* Expire After: 30 Days

Between 1 - 365, 30 days recommended

Max. Number of Uses:

The token will be expired when either the expiration or the maximum uses is reached

☑ Allow export-controlled functionality on the products registered with this token ⓘ

Create Token   Cancel

Step 5. In Cisco Intersight click **Select Service > System**. Then on the left click **Administration > Licensing**.

Step 6. Under **Actions**, click **Register**.

Licensing

Actions ⌄

Set Products
Start Trial
Enable Get Subscription Information
Register

Subscription

Intersight provides the capability to have multiple active license tiers within a single intersight account. You can assign servers to a preferred tier. Learn more at Help Center.

Products

Intersight

DEFAULT
None
(Not Licensed Servers)

Essentials

Advantage

Step 7. Enter the copied token from the Cisco Smart Licensing portal. Click **Next**.

Step 8. Drop-down the pre-selected Default Tier * and select the license type (for example, Premier).

Step 9. Select **Move All Servers to Default Tier**.

Step 10. Click **Register**, then **Register** again.

Step 11. When the registration is successful (takes a few minutes), the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

Procedure 4. Set Up Cisco Intersight Resource Group

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

Step 1. Log into Cisco Intersight.

Step 2. At the top, select System. On the left, click **Settings** (the gear icon).

Step 3. Click **Resource Groups** in the middle panel.

Step 4. Click **+ Create Resource Group** in the top-right corner.

Step 5. Provide a name for the Resource Group (for example, AA02-rg).

← Resource Groups

Create Resource Group

Create Resource Group

Create a Resource Group to manage and access the targets.

General

Name *
AA02-rg

Description

Memberships

| Custom | All |

ⓘ The selected targets will be part of the Resource Group created.

0 items found   10 ∨ per page |< < 0 of 0 > >| ⚙

🔍 Add Filter

☐	Name	Status	Type	IP Address	Target ID

NO ITEMS AVAILABLE

|< < 0 of 0 > >|

Cancel   Create

Step 6. Under Memberships, select **Custom**.

Step 7. Click **Create**.

Procedure 5. Set Up Cisco Intersight Organization

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

Step 1. Log into the Cisco Intersight portal.

Step 2. At the top, select System. On the left, click **Settings** (the gear icon).

Step 3. Click **Organizations** in the middle panel.

Step 4. Click **+ Create Organization** in the top-right corner.

Step 5. Provide a name for the organization (for example, AA02).

Step 6. Select the Resource Group created in the last step (for example, AA02-rg).

Step 7. Click **Create**.

Procedure 6.   Claim Cisco UCS Fabric Interconnects in Cisco Intersight

Make sure the initial configuration for the fabric interconnects has been completed. Log into the Fabric Interconnect A Device Console using a web browser to capture the Cisco Intersight connectivity information.

Step 1. Use the management IP address of Fabric Interconnect A to access the device from a web browser and the previously configured admin password to log into the device.

Step 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 Cisco Intersight.
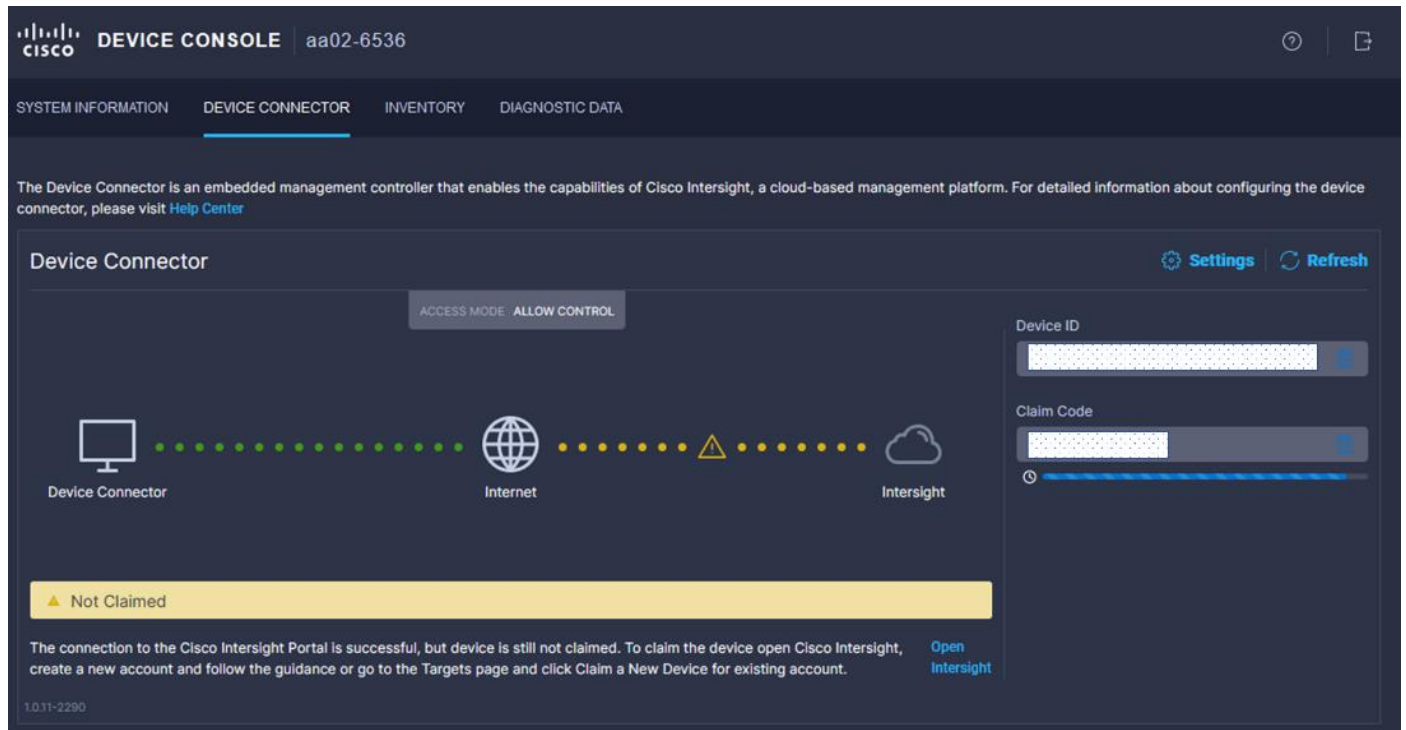
Step 3. Log into Cisco Intersight.

Step 4. At the top, select System. On the left, click **Administration > Targets.**

Step 5. Click **Claim a New Target**.

Step 6. Select **Cisco UCS Domain (Intersight Managed)** and click **Start**.

Step 7. Copy and paste the Device ID and Claim from the Cisco UCS FI to Intersight.

Step 8. Select the previously created Resource Group and click **Claim**.

Claim a New Target

Claim Cisco UCS Domain (Intersight Managed) Target

To claim your target, provide the Device ID, Claim Code and select the appropriate Resource Groups.

General

Device ID * ⓘ   Claim Code * ⓘ

Resource Groups

ⓘ Select the Resource Groups if required. However, this selection is not mandatory as one or more Resource Group type is 'All'. The claimed target will be part of all Organizations with the Resource Group type 'All'.

1 items found   10 ⌄ per page |< < 1 of 1 > >| ⚙

	Name	Usage	Description
☐	AA02-rg	AA02	

|< < 1 of 1 > >|

Back   **Cancel**   **Claim**

On a successful device claim, Cisco UCS FI should appear as a target in Cisco Intersight.

Targets

Procedure 7. Verify Addition of Cisco UCS Fabric Interconnects to Cisco Intersight

Step 1. Log into the web GUI of the Cisco UCS fabric interconnect and click the browser refresh button.

The fabric interconnect status should now be set to **Claimed**.
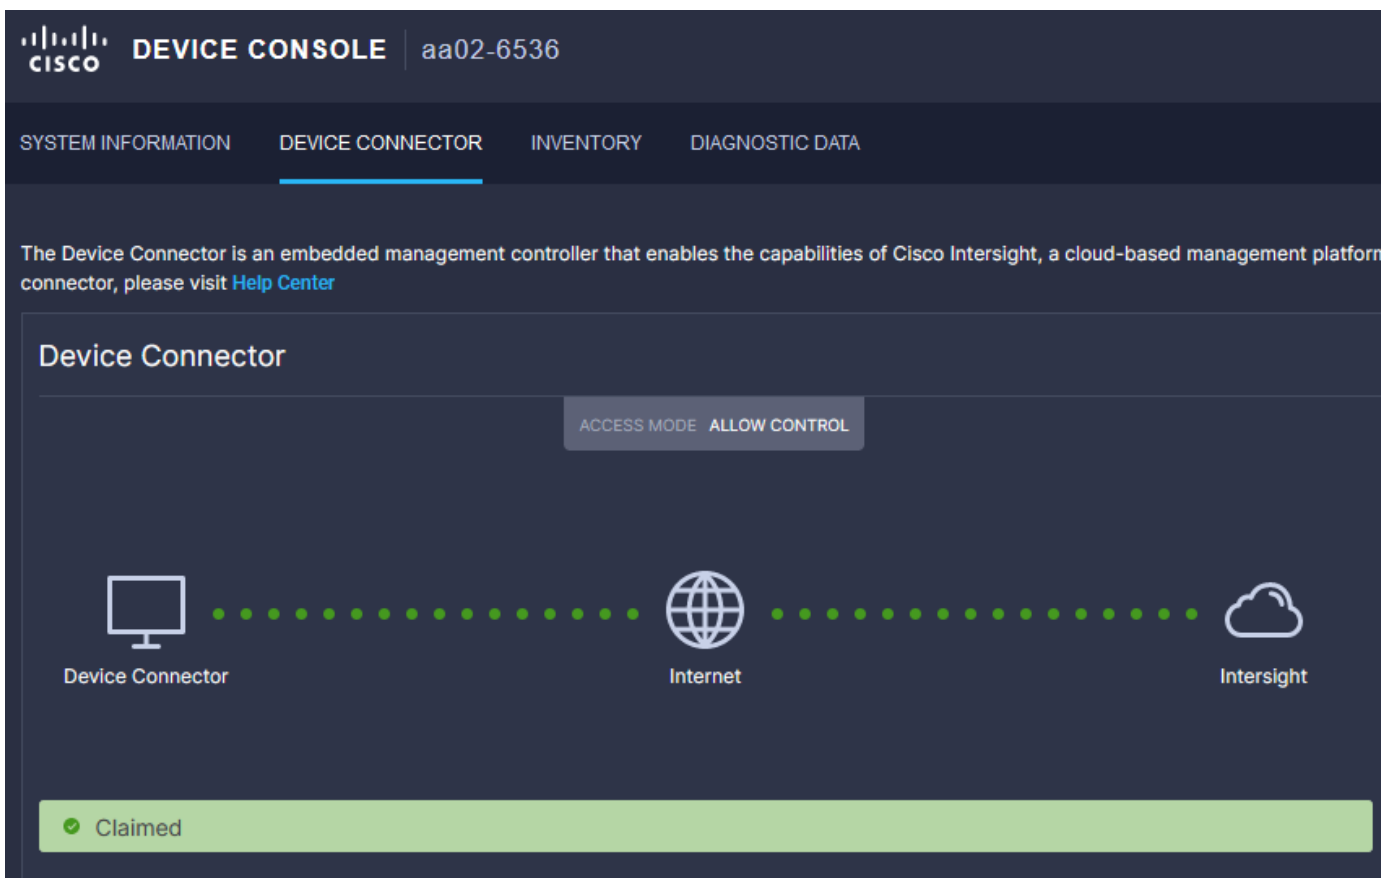
Procedure 8. Upgrade Fabric Interconnect Firmware using Cisco Intersight

If your Cisco UCS 6536 Fabric Interconnects are not already running firmware release 4.2(2c) (NX-OS version 9.3(5)I42(2c)), upgrade them to 4.2(2c).

Note: If Cisco UCS Fabric Interconnects were upgraded to the latest recommended software using Cisco UCS Manager, this upgrade process through Cisco Intersight will still work and will copy the Cisco X-Series firmware to the fabric interconnects.

Step 1. Log into the Cisco Intersight portal.

Step 2. At the top, using the pulldown select **Infrastructure Service** and then select **Fabric Interconnects** under Operate on the left.

Step 3. Click the three dots "**...**" at the end of the row for either of the Fabric Interconnects and select **Upgrade Firmware**.

Step 4. Click **Start**.

Step 5. Verify the Fabric Interconnect information and click **Next**.

Step 6. Enable **Advanced Mode** using the toggle switch and uncheck Fabric Interconnect Traffic Evacuation.

Step 7. Select 4.2(2c) release from the list and click **Next**.

Step 8. Verify the information and click **Upgrade** to start the upgrade process.

Step 9. Watch the Request panel of the main Intersight screen as the system will ask for user permission before upgrading each FI. Click on the Circle with Arrow and follow the prompts on screen to grant permission.

Step 10. Wait for both the FIs to successfully upgrade.

Procedure 9. Configure a Cisco UCS Domain Profile

Note: A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies, allows configuration of the ports and port channels, and configures the VLANs and VSANs in the network. It defines the characteristics of and configured ports on fabric interconnects. The domain-related policies can be attached to the profile either at the time of creation or later. One Cisco UCS domain profile can be assigned to one fabric interconnect domain.

Step 1. Log into the Cisco Intersight portal.

Step 2. At the top, use the pulldown to select **Infrastructure Service**. Then, under Configure select **Profiles**.

Step 3. In the main window, select **UCS Domain Profiles** and click **Create UCS Domain Profile**.

Step 4. On the Create UCS Domain Profile screen, click **Start**.

Create UCS Domain Profile

A UCS domain profile streamlines fabric interconnect assignment, port, and fabric interconnect configuration to eliminate failures caused by inconsistent configuration.

UCS Domain Assignment

Create a Fabric Interconnect pair and assign to a domain profile immediately or later.

‹ • · · ›

💡 **About UCS Domain Profile Creation**

☐ Do not show this page again

Cancel   **Start**

Procedure 10. General Configuration

Step 1.   Select the organization from the drop-down list (for example, AA02).

Step 2.   Provide a name for the domain profile (for example, AA02-6536-Domain-Profile).

Step 3.   Provide an optional Description.

General

Add a name, description and tag for the UCS domain profile.

Organization *

AA02

Name *

AA02-6536-Domain-Profile

Set Tags

Description

<= 1024

Close   **Back**   **Next**

Step 4.   Click **Next**.

Procedure 11. Cisco UCS Domain Assignment

Step 1.   Assign the Cisco UCS domain to this new domain profile by clicking **Assign Now** and selecting the previously added Cisco UCS domain (for example, AA02-6536).

Step 2. Click **Next**.

VLAN and VSAN Configuration

In this procedure, a single VLAN policy is created for both fabric interconnects and two individual VSAN policies are created because the VSAN IDs are unique for each fabric interconnect.

Procedure 1.	Create and Apply VLAN Policy

Step 1. Click **Select Policy** next to VLAN Configuration under Fabric Interconnect A.

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

▲ Fabric Interconnect A	0 of 2 Policies Configured
VLAN Configuration	**Select Policy** 📄
VSAN Configuration	**Select Policy** 📄

▲ Fabric Interconnect B	0 of 2 Policies Configured
VLAN Configuration	**Select Policy** 📄
VSAN Configuration	**Select Policy** 📄

Close   **Back** **Next**

Step 2. In the pane on the right, click **Create New**.

Step 3. Verify the correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-6536-VLAN).

General

Add a name, description and tag for the policy.

Organization *

AA02

Name *

AA02-6536-VLAN

Set Tags

Description

VLAN Policy for both FIs

<= 1024

1 General

2 Policy Details

Cancel   Next

Step 4.   Click **Next**.

Step 5.   Click **Add VLANs**.

Step 6.   Provide a name and VLAN ID for the native VLAN.

Add VLANs

Add VLANs to the policy

⚠ VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix *
Native-VLAN ⓘ

VLAN IDs *
2 ⓘ

🔵 Auto Allow On Uplinks ⓘ

⚪ Enable VLAN Sharing ⓘ

Multicast Policy *
Select Policy 🗎

Step 7.	Make sure **Auto Allow On** Uplinks is enabled.
Step 8.	To create the required Multicast policy, click **Select Policy** under Multicast*.
Step 9.	In the window on the right, Click **Create New** to create a new Multicast Policy.
Step 10.	Provide a Name for the Multicast Policy (for example, AA02-MCAST).
Step 11.	Provide optional Description and click **Next**.
Step 12.	Leave the Snooping State selected and click **Create**.

Create UCS Domain Profile > Create VLAN

Create Multicast Policy

✓ General

② Policy Details

Policy Details

Add policy details

Multicast Policy

⬤ Snooping State ⓘ

⬤ Querier State ⓘ

Step 13.   Click **Add** to add the VLAN.

Step 14.   Select **Set Native VLAN ID** and enter the VLAN number (for example, 2) under VLAN ID.

Policy Details

Add policy details

> ℹ This policy is applicable only for UCS Domains

VLANs

Add VLANs

⬤ Show VLAN Ranges

| | | | 2 items found | 50 ∨ per page |⟨ ⟨ 1 of 1 ⟩ ⟩| ⚙ |

🔍 Add Filter

☐	VLA... ↕	Name ↕	Shari... ↕	Prim... ↕	Multicast ...	Auto Allo...	⚡
☐	1	default	None			Yes	···
☐	2	Native-V...	None		AA02-M...	Yes	···

✏ 🗑 |⟨ ⟨ 1 of 1 ⟩ ⟩|

☑ Set Native VLAN ID

VLAN ID

2

Step 15.   Add the remaining VLANs for FlexPod by clicking Add VLANs and entering the VLANs one by one. Reuse the previously created multicast policy for all the VLANs.

The VLANs created during this validation are shown in the image below:

Create UCS Domain Profile

Create VLAN

Note: The iSCSI and NVMe-TCP VLANs shown in the screen image above are only needed when iSCSI and NVME-TCP are configured in the environment.

Step 16.   Click **Create** at bottom right to finish creating the VLAN policy and associated VLANs.

Step 17.   Click **Select Policy** next to VLAN Configuration for Fabric Interconnect B and select the same VLAN policy.

Procedure 2.   Create and Apply VSAN Policy (FC configuration only)

Step 1.   Click **Select Policy** next to VSAN Configuration under Fabric Interconnect A. Then, in the pane on the right, click **Create New**.

Step 2.   Verify the correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-6536-VSAN-Pol-A).

Note: A separate VSAN-Policy is created for each fabric interconnect.

Step 3.   Click **Next**.

Step 4. Optionally enable **Uplink Trunking**.

Policy Details

Add policy details

- This policy is applicable only for UCS Domains

🔵 Uplink Trunking ⓘ

Step 5. Click **Add VSAN** and provide a name (for example, VSAN-A), VSAN ID (for example, 101), and associated Fibre Channel over Ethernet (FCoE) VLAN ID (for example, 101) for SAN A.

Step 6. Set VLAN Scope as **Uplink**.

Add VSAN

Name *

VSAN-A   ⓘ

VSAN Scope ⓘ

◯ Storage & Uplink ⓘ   ◯ Storage ⓘ   ⦿ Uplink ⓘ

VSAN ID *

101   ⌄̂ ⓘ

1 - 4093

FCoE VLAN ID *

101   ⌄̂ ⓘ

Cancel   **Add**

Step 7. Click **Add**.

Step 8. Click **Create** to finish creating VSAN policy for fabric A.

Step 9.   Repeat these steps to create a new VSAN policy for SAN-B. Name the policy to identify the SAN-B configuration (for example, AA02-6536-VSAN-Pol-B) and use appropriate VSAN and FCoE VLAN (for example, 102).

Step 10.   Verify that a common VLAN policy and two unique VSAN policies are associated with the two fabric interconnects.

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

⌃ **Fabric Interconnect A**   2 of 2 Policies Configured

VLAN Configuration   ✕ | ✎ | AA02-6536-VLAN ▤

VSAN Configuration   ✕ | ✎ | AA02-6536-VSAN-Pol-A ▤

⌃ **Fabric Interconnect B**   2 of 2 Policies Configured

VLAN Configuration   ✕ | ✎ | AA02-6536-VLAN ▤

VSAN Configuration   ✕ | ✎ | AA02-6536-VSAN-Pol-B ▤

Step 11.   Click **Next**.

Procedure 3.   Ports Configuration

Step 1.   Click **Select Policy** for Fabric Interconnect A.

Step 2.   Click **Create New** in the pane on the right to define a new port configuration policy.

Note: Use two separate port policies for the fabric interconnects. Using separate policies provide flexibility when port configuration (port numbers or speed) differs between the two FIs. When configuring Fibre Channel, two port policies are required because each fabric interconnect uses a unique Fibre Channel VSAN ID.

Step 3.   Verify correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-6536-PortPol-A). Select the UCS-FI-6536 Switch Model.

Step 4.   Click **Next**.

Step 5.   Move the slider to set up unified ports. In this deployment, the last two ports were selected as Fibre Channel ports as 4x32G breakouts. Click **Next**.

Unified Port

Configure the port modes to carry FC or Ethernet traffic.

> **ⓘ** Move slider to configure unified ports and select port to set breakout.

Fibre Channel Ports

2 Fiber Channel Ports (Port 35,Port 36)

● FC   ● Ethernet   ☐ Port Modes

| FC | Ports 35-36 | Ethernet | Ports 1-34 |

Step 6. If any Ethernet ports need to be configured as breakouts, either 4x25G or 4x10G, for connecting Cisco UCS C-Series servers or a Cisco UCS 5108 chassis, configure them here. In the list, select the checkbox next to any ports that need to be configured as breakout or select the ports on the graphic. When all ports are selected, click **Configure** at the top of the window.

Breakout Options

Configure breakout ports on FC or Ethernet.

Ethernet   Fibre Channel

Configure   Selected Ports   Port 17 | Clear Selection

● FC   ● Ethernet   ☐ Port Modes

Step 7. In the Set Breakout popup, select either 4x10G or 4x25G and click **Set**.

Set Breakout

⚠ **Modifying the speed of an existing FC breakout port, will result in the deletion of previously configured port roles and port channel roles.**

Selected Ports   Port 17

○ **No Breakout**   ○ **4x10G**   ◉ **4x25G**

[Cancel]   [**Set**]

Step 8.   Under Breakout Options, select **Fibre Channel**. Select any ports that need the speed changed from 16G to 32G and click **Configure**.

Step 9.   In the Set Breakout popup, select 4x32G and click **Set**.

Set Breakout

⚠ **Modifying the speed of an existing FC breakout port, will result in the deletion of previously configured port roles and port channel roles.**

Selected Ports   Port 35, Port 36

○ **4x8G**   ○ **4x16G**   ◉ **4x32G**

[Cancel]   [**Set**]

Step 10.   Click **Next**.

Step 11.   In the list, select the checkbox next to any ports that need to be configured as server ports, including ports connected to chassis or C-Series servers. Ports can also be selected on the graphic. When all ports are selected, click **Configure**. Breakout and non-breakout ports cannot be configured together. If you need to configure breakout and non-breakout ports, do this configuration in two steps.

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

Port Roles   Port Channels   Pin Groups

Configure   Selected Ports   Port 9, Port 10, Port 11, Port 12, Port 13, Port 14   Clear Selection

- Unconfigured

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

Port Roles   Port Channels   Pin Groups

Configure   Selected Ports   Port 17/1, Port 17/2, Port 17/3, Port 17/4   Clear Selection

- Unconfigured   - Server

Step 12. From the drop-down list, select **Server** as the role. Also, unless you are using a Cisco Nexus 93180YC-FX3 as a FEX, leave Auto Negotiation enabled. If you need to do manual number of Chassis or C-Series Servers, enable Manual Chassis/Server Numbering.

Configure (6 Ports)

Configuration	

Selected Ports   **Port 9, Port 10, Port 11, Port 12, Port 13, Port 14**

Role

Server   ⌄

> ⓘ Auto Negotiation is not supported on N9K-C93180YC-FX3 for 100G speed ports. If the port is connected to N9K-C93180YC-FX3, the Auto Negotiation option should be disabled. Learn more at **Help Center**.

🔵◯   Auto Negotiation ⓘ

◯   Manual Chassis/Server Numbering ⓘ

Configure (4 Ports)

Configuration	

Selected Ports   **Port 17/1, Port 17/2, Port 17/3, Port 17/4**

Role

Server   ⌄

◯   Manual Chassis/Server Numbering ⓘ

Step 13.   Click **Save**.

Step 14.   Configure the Ethernet uplink port channel by selecting **Port Channel** in the main pane and then clicking **Create Port Channel**.

Step 15.   Select **Ethernet Uplink Port Channel** as the role, provide a port-channel ID (for example, 11), and select a value for Admin Speed from drop-down list (for example, Auto).

Note: You can create the Ethernet Network Group, Flow Control, Link Aggregation for defining disjoint Layer-2 domain or fine tune port-channel parameters. These policies were not used in this deployment and system default values were utilized.

Step 16. Under Link Control, click **Select Policy**. In the upper right, click **Create New**.

Step 17. Verify the correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-UDLD-Link-Control). Click **Next**.

Step 18. Leave the default values selected and click **Create**.

Step 19. Scroll down and select uplink ports from the list of available ports (for example, port 31 and 32)

Step 20. Click **Save**.

Procedure 4. Configure FC Port Channel (FC configuration only)

Note: An FC uplink port channels only needed when configuring FC SAN and can be skipped for IP-only (iSCSI) storage access.

Step 1. Configure a Fibre Channel Port Channel by selecting the **Port Channel** in the main pane again and clicking **Create Port Channel**.

Step 2. In the drop-down list under Role, select **FC Uplink Port Channel**.

Step 3. Provide a port-channel ID (for example, 135), select a value for Admin Speed (for example, 32Gbps), and provide a VSAN ID (for example, 101).

Create Port Channel

Configuration

> ❶ The combined maximum number of Ethernet Uplink, FCoE Uplink, and Appliance port channels permitted is 12 and the maximum number of FC port channels permitted is 4.

Role

FC Uplink Port Channel ⌄

Port Channel ID *	Admin Speed	VSAN ID *
135 ⌄ ⓘ	32Gbps ⌄ ⓘ	101 ⌄ ⓘ
1 - 256		1 - 4093

Select Member Ports

> ❶ FC or Ethernet ports with unconfigured role are available for port channel creation.

● Ethernet Uplink Port Channel

Step 4. Select ports (for example, 35/1,35/2,35/3,35/4).

Step 5. Click **Save**.

Step 6. Verify the port-channel IDs and ports after both the Ethernet uplink port channel and the Fibre Channel uplink port channel have been created.

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

Port Roles   **Port Channels**   Pin Groups

Create Port Channel

● Ethernet Uplink Port Channel   ● FC Uplink Port Channel

	ID	Role	Ports
☐	131	Ethernet Uplink Port C...	Port 31, Port 32
☐	135	FC Uplink Port Channel	Port 35/1, Port 35/2, P...

2 items found   10 ⌄ per page |< < 1 of 1 > >|

1 of 1

Step 7.   Click **Save** to create the port policy for Fabric Interconnect A.

Note: Use the summary screen to verify that the ports were selected and configured correctly.

Procedure 5.   Port Configuration for Fabric Interconnect B

Step 1.   Repeat the steps in <u>Ports Configuration</u> and <u>Configure FC Port Channel</u> to create the port policy for Fabric Interconnect B including the Ethernet port-channel and the FC port-channel (if configuring SAN). Use the following values for various parameters:

- Name of the port policy: AA02-PortPol-B

- Ethernet port-Channel ID: 132

- FC port-channel ID: 135

- FC VSAN ID: 102

Step 2. When the port configuration for both fabric interconnects is complete and looks good, click **Next**.

Procedure 6. UCS Domain Configuration

Under UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, four policies (NTP, Network Connectivity, SNMP, and System QoS) will be configured, as shown below:

Procedure 7. Configure NTP Policy

Step 1. Click **Select Policy** next to NTP and then, in the pane on the right, click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-NTP).

Step 3. Click **Next**.

Step 4. Enable NTP, provide the first NTP server IP address, and select the time zone from the drop-down list.

Step 5. Add a second NTP server by clicking **+** next to the first NTP server IP address.

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

Step 6. Click **Create**.

Procedure 8.	Configure Network Connectivity Policy

Step 1. Click **Select Policy** next to Network Connectivity and then, in the pane on the right, click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-NetConn).

Step 3. Click **Next**.

Step 4. Provide DNS server IP addresses for Cisco UCS (for example, 10.102.1.151 and 10.102.1.152).

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | **UCS Domain**

Common Properties

IPv4 Properties

Preferred IPv4 DNS Server
10.102.1.151

Alternate IPv4 DNS Server
10.102.1.152

Enable IPv6 ⓘ

Step 5.   Click **Create**.

Procedure 9.   Configure SNMP Policy

Step 1.   Click **Select Policy** next to SNMP and then, in the pane on the right, click **Create New**.

Step 2.   Verify correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-SNMP).

Step 3.   Click **Next**.

Step 4.   Provide a System Contact email address, a System Location, and optional Community Strings.

Step 5.   Under SNMP Users, click **Add SNMP User**.

Step 6.   This user id will be used for Cisco DCNM SAN to query the UCS Fabric Interconnects. Fill in a user name (for example, snmpadmin), Auth Type SHA, an Auth Password with confirmation, Privacy Type AES, and a Privacy Password with confirmation. Click **Add**.

Add SNMP User

×

Name *

snmpadmin

Security Level *

AuthPriv

Auth Type

SHA

Auth Password *

••••••••

Auth Password Confirmation *

••••••••

Privacy Type

AES

Privacy Password *

••••••••

Privacy Password Confirmation *

••••••••

Cancel   Add

Step 7. Optionally, add an SNMP Trap Destination (for example, the DCNM SAN IP Address). If the SNMP Trap Destination is V2, you must add Trap Community String.

Step 8. Click **Create**.

Procedure 10. Configure System QoS Policy

Step 1. Click **Select Policy** next to System QoS* and in the pane on the right, click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, AA02) and provide a name for the policy (for example, AA02-QoS).

Step 3. Click **Next**.

Step 4. Change the MTU for Best Effort class to **9216**.

Step 5. Keep the default selections or change the parameters if necessary.

Policy Details

Add policy details

ⓘ This policy is applicable only for UCS Domains

Configure Priorities

Platinum

Gold

Silver

Bronze

Best Effort	CoS	Weight	Allow Packet Drops	MTU
	Any ⓘ	5 ⓘ	☑ ⓘ	9216 ⓘ
		0 - 10		1500 - 9216
Fibre Channel	CoS	Weight	Allow Packet Drops	MTU
	3 ⓘ	5 ⓘ	☐ ⓘ	2240 ⓘ
	0 - 6	0 - 10		1500 - 9216

Step 6.   Click **Create**.

UCS Domain Configuration

Select the compute and management policies to be associated with the fabric interconnect.

Show Attached Policies (4)

∧ **Management** 3 of 4 Policies Configured

NTP	× | ✎ |	AA02-NTP
Syslog		**Select Policy**
Network Connectivity	× | ✎ |	AA02-NetConn
SNMP	× | ✎ |	AA02-SNMP

∧ **Network** 1 of 2 Policies Configured

System QoS *	× | ✎ |	AA02-QoS
Switch Control		**Select Policy**

Step 7.   Click **Next**.

Procedure 11. Summary

Step 1.   Verify all the settings including the fabric interconnect settings, by expanding the settings and make sure that the configuration is correct.

Procedure 12. Deploy the Cisco UCS Domain Profile

Step 1. From the UCS domain profile Summary view, click **Deploy**.

Step 2. Acknowledge any warnings and click **Deploy** again.

Note: The system will take some time to validate and configure the settings on the fabric interconnects. Log into the console servers to see when the Cisco UCS fabric interconnects have finished configuration and are successfully rebooted.

Procedure 13. Verify Cisco UCS Domain Profile Deployment

When the Cisco UCS domain profile has been successfully deployed, the Cisco UCS chassis and the blades should be successfully discovered.

Note: It takes a while to discover the blades for the first time. Watch the number of outstanding requests in Cisco Intersight:

Requests

	Name	Status	Initiator	Target Type	Target Name	Start Time	Duration
☐	Blade Discovery	⊘ Success	system@intersight	Blade Server	aa02-6536-1-1	10 minutes ago	9 m 2 s
☐	Blade Discovery	⊘ Success	system@intersight	Blade Server	aa02-6536-1-5	11 minutes ago	10 m 36 s
☐	Blade Discovery	⊘ Success	system@intersight	Blade Server	aa02-6536-1-3	11 minutes ago	10 m 36 s
☐	Blade Discovery	⊘ Success	system@intersight	Blade Server	aa02-6536-1-7	11 minutes ago	10 m 37 s
☐	Rack Server Disco...	⊘ Success	system@intersight	Rack Server	aa02-6536-2	10 minutes ago	10 m 21 s
☐	Rack Server Disco...	⊘ Success	system@intersight	Rack Server	aa02-6536-3	11 minutes ago	10 m 32 s
☐	Rack Server Disco...	⊘ Success	system@intersight	Rack Server	aa02-6536-1	11 minutes ago	11 m 18 s
☐	Chassis Inventory	⊘ Success	system@intersight	Chassis	aa02-6536-1	11 minutes ago	3 m 2 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	11 minutes ago	26 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	11 minutes ago	24 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	12 minutes ago	34 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	12 minutes ago	30 s
☐	Chassis Inventory	⊘ Success	system@intersight	Chassis	aa02-6536-1	12 minutes ago	2 m 17 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	12 minutes ago	25 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	12 minutes ago	24 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	13 minutes ago	36 s
☐	Chassis Discovery	⊘ Success	system@intersight	Chassis	aa02-6536-1	13 minutes ago	30 s
☐	Deploy Domain Pr...	⊘ Success	jogeorg2@cisco.c...	Fabric Interconnect	aa02-6536 FI-A	33 minutes ago	21 m 28 s
☐	Deploy Domain Pr...	⊘ Success	jogeorg2@cisco.c...	Fabric Interconnect	aa02-6536 FI-B	33 minutes ago	22 m 44 s

Step 1.   Log into Cisco Intersight. Under **Infrastructure Service > Configure > Profiles > UCS Domain Profiles**, verify that the domain profile has been successfully deployed.

Profiles

HyperFlex Cluster Profiles   UCS Chassis Profiles   **UCS Domain Profiles**   UCS Server Profiles

Create UCS Domain Profile

* All UCS Domain Pr... ⚙ +

··· 🖉 🏷 🗑 | 🔍 Add Filter   🔁 Export   1 items found   10 ∨ per page |◁ ◁   1   of 1 ▷ ▷|

	Name	Status	UCS Domain Fabric Interc...	Fabric Interc...	Last Update		𝄐
☐	AA02-6536-Domain-Profile	⊘ OK	aa02-6536 ...	aa02-6536 ...	7 minutes ago		···

··· 🖉 🏷 🗑 |◁ ◁   1   of 1 ▷ ▷|

Step 2.   Verify that the chassis (either UCSX-9508 or UCS 5108 chassis) has been discovered and is visible under **Infrastructure Service > Operate** > **Chassis**.

Chassis

* All Chassis ⚙ +

··· 🏷 | 🔍 Add Filter   🔁 Export   1 items found   10 ∨ per page |◁ ◁   1   of 1 ▷ ▷|

Health

(1)   ● Healthy **1**

	Name	Health	UCS Domain	Model	Chassis Profile	𝄐
☐	aa02-6536-1	⊘ Healthy	aa02-6536	UCSX-9508		···

··· 🏷 |◁ ◁   1   of 1 ▷ ▷|

Step 3.   Verify that the servers have been successfully discovered and are visible under **Infrastructure Service > Operate > Servers**.

Procedure 14. Ansible Cisco UCS IMM Configuration

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

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

Step 1. To execute the playbooks against your Intersight account, you need to create an API key and save a SecretKey.txt file from your Cisco Intersight account.

   a. In Cisco Intersight, select **System** > **Settings** > **API** > **API Keys**.

   b. Click Generate API Key.

   c. Under Generate API Key, enter a Description (for example, API Key for Ansible) and select API key for OpenAPI schema version 2. Click **Generate**.

Generate API Key ✕

Description

API Key for Ansible ⓘ

API Key Purpose ⓘ

⦿ API key for OpenAPI schema version 2 ⓘ

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

[Close] [**Generate**]

d. In the Generate API Key window, click the upper 📋 icon to copy the API Key ID to the clipboard. Paste this key into the api_key_id variable in the FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/ucs.yml variable file and save it.

e. Using an editor, open the FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/SecretKey.txt file and clear all text from the file. Then click the lower 📋 icon in the Generate API Key window and paste the Secret Key into the SecretKey.txt file and save it.

Step 2.   Edit the following variable files to ensure proper UCS variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/ucs.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/UCS-IMM/create_pools/defaults/main.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/UCS-IMM/create_server_policies/defaults/main.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/UCS-IMM/create_server_profile_template/defaults/main.yml

Note: It is critical when entering values in the variable files that either the FC and FC-NVMe NetApp LIF WWPNs or Infrastructure SVM iSCSI IQN be entered into the all.yml file so that UCS SAN boot and MDS device alias can be properly configured. LIF WWPNs can be queried by connecting to the NetApp cluster CLI interface and running "network interface show -vserver <svm-name>." If iSCSI SAN boot is being configured, the Infrastructure SVM's iSCSI IQN can be queried by running "vserver iscsi show -vserver <svm-name>."

Step 3. The /root/ FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2 directory contains three Ansible playbooks to set up Cisco UCS IMM server profile templates: Create_IMM_Pools.yml, Create_IMM_Server_Policies.yml, and Create_IMM_Server_Profile_Templates.yml. Run Create_IMM_Pools.yml only once. Then Create_IMM_Server_Policies.yml and Create_IMM_Server_Profile_Templates.yml are designed to be run more than once with different combinations of server_cpu_type andvic_type. It is important that when Create_IMM_Server_Policies.yml is run that Create_IMM_Server_Profiles_Templates.yml is run before changing the server_cpu_type and vic_type variables. Many of the policies and templates will be assigned unique names according to these variables. Also, since the UCS-IMM Ansible playbooks connect to the Cisco Intersight API website instead of hardware components, that the use of the inventory file is not needed. To set up the Cisco Intersight IMM pools, policies, and server profile templates, execute the following:

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

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

Cisco UCS IMM Setup Completion

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

Procedure 1. Derive Server Profiles

Step 1. From the Server profile template Summary screen, click **Derive Profiles**.

Note: This action can also be performed later by navigating to **Templates**, clicking **"..."** next to the template name and selecting **Derive Profiles**.

Step 2. Under the Server Assignment, select **Assign Now** and select Cisco UCS X210c M6 server(s). You can select one or more servers depending on the number of profiles to be deployed.

Server Assignment

Assign Now	Assign Server from a Resource Pool	Assign Later

🔍 Add Filter   ⤷   16 items found   10 ⌄ per page ⎢K⎢ ⎢<⎢   1   o

	Name	User Label	Health	Model	UCS Domain
☐	aa02-6536-7		⚠ Warning	UCSC-C225-...	aa02-6536
☐	aa02-6536-8		⚠ Warning	UCSC-C225-...	aa02-6536
☐	aa02-6536-5		⚠ Warning	UCSC-C225-...	aa02-6536
☐	aa02-6536-6		⚠ Warning	UCSC-C225-...	aa02-6536
☐	aa02-6536-1-1		✔ Healthy	UCSX-210C-M6	aa02-6536
☐	aa02-6536-1-3		✔ Healthy	UCSX-210C-M6	aa02-6536
☑	aa02-6536-1-5		✔ Healthy	UCSX-210C-M6	aa02-6536

Step 3.　　Click **Next**.

Note:　Cisco Intersight will fill in default information for the number of servers selected (1 in this case).

Step 4.　　Adjust the fields as needed. It is recommended to use the server hostname for the Server Profile name.

Details

Edit the description, tags, and auto-generated names of the profiles.

General

Organization *
AA02

Target Platform
UCS Server (FI-Attached)

Description
Supports iSCSI boot from SAN

<= 1024

Set Tags

Derive

1	Name *		Assigned Server
	aa02-esxi-2		aa02-6536-1-5

Step 5.   Click **Next**.

Step 6.   Verify the information and click **Derive** to create the Server Profile(s).

Step 7.   In the Infrastructure Service > Configure > Profiles > UCS Server Profiles list, select the profile(s) just created and click the **...** at the top of the column and select **Deploy**. Click Deploy to confirm.

Step 8.   Cisco Intersight will start deploying the server profile(s) and will take some time to apply all the policies. Use the Requests tab at the top right-hand corner of the window to see the progress.

When the Server Profile(s) are deployed successfully, they will appear under the Server Profiles with the status of OK.

Profiles

Create UCS Server Profile

✳ All UCS Server Prof... ⊗ +

··· ✎ ◌ 🗑 | 🔍 Add Filter   ⤴ Export   16 items found   27 ∨ per page |< < 1 of 1 > >|

	Name	Status	Target Platform	UCS Server Template	Server	Last Update	
☐	aa02-esxi-2	⊘ OK	UCS Server (FI-Attached)	Intel-5G-VIC-MLOM-iSCSI-...	aa02-6536-1-5	a few seconds ago	···
☐	aa02-esxi-5	⊘ OK	UCS Server (FI-Attached)	AA02-AMD-4G-VIC-2-FC-...	aa02-6536-1	2 minutes ago	···
☐	aa02-esxi-7	⊘ OK	UCS Server (FI-Attached)	AA02-AMD-4G-VIC-MLOM...	aa02-6536-5	2 minutes ago	···
☐	aa02-esxi-8	⊘ OK	UCS Server (FI-Attached)	AA02-AMD-4G-VIC-MLOM...	aa02-6536-6	2 minutes ago	···

Step 9.   Derive and Deploy all needed servers for your FlexPod environment.

SAN Switch Configuration

This chapter contains the following:

- [Physical Connectivity](#)

- [FlexPod Cisco MDS Base](#)

This chapter explains how to configure the Cisco MDS 9000s for use in a FlexPod environment. The configuration covered in this section is only needed when configuring Fibre Channel and FC-NVMe storage access.

Note: If FC connectivity is not required in the FlexPod deployment, this section can be skipped.

Note: If the Nexus 93360YC-FX2 switches are being used for SAN switching in this FlexPod Deployment, please refer to FlexPod with Cisco Nexus 93360YC-FX2 SAN Switching Configuration – Part 2 in the Appendix of this document.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in [Physical Topology](#) section.

FlexPod Cisco MDS Base

The following procedures describe how to configure the Cisco MDS switches for use in a base FlexPod environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.4(2c).

Procedure 1. Set up Cisco MDS 9132T A and 9132T B

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. Enter y to get to the System Admin Account Setup.

Step 1. Configure the switch using the command line:

```
         ---- 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 : <mds-A-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address : <mds-A-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter
```

```
IPv4 address of the default gateway : <mds-A-mgmt0-gw>

Configure advanced IP options? (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

Enable the telnet service? (yes/no) [n]: Enter

Configure congestion/no_credit drop for fc interfaces? (yes/no)     [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500.  [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

Configure the ntp server? (yes/no) [n]: Enter

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto

Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: y

Configure default zone mode (basic/enhanced) [basic]: Enter
```

Step 2.   Review the configuration.

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

Step 3.   To avoid possible timing errors with Ansible, perform a no shutdown on the FC interfaces con-
nected to the Cisco UCS fabric interconnects.

```
config t
int fc1/5-8
no shutdown
copy r s
exit
```

Step 4.   To set up the initial configuration of the Cisco MDS B switch, repeat steps 1-3 with appropriate
host and IP address information.

Procedure 2.   FlexPod Cisco MDS Switch Ansible Configuration

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

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

Step 2. Edit the following variable files to ensure proper MDS variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/mdsA.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/mdsB.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/MDSconfig/defaults/main.yml

Note: The FC and FC-NVMe NetApp LIF WWPNs should have already been entered into the all.yml file so that MDS device aliases can be properly configured. The Cisco UCS server initiator WWPNs for both FC and FC-NVMe should also be entered into all.yml. To query these WWPNs, log into Cisco Intersight 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 can be found under HBA Interfaces.

← Server Profiles

aa02-esxi-1

Actions ⌄

General   Server   **Inventory**

⊞ **Expand All**

Motherboard

Boot
Management
Controller
⊞ CPUs
⊞ Memory
⊟ Network Adapters
⎸ Adapter UCSX-ML-
V5D200G_FCH254474UN
⊞ Storage Controllers
TPM

Adapter UCSX-ML-V5D200G_FCH254474UN

General   **Interfaces**

DCE Interfaces

🔍 Add Filter

Name	IO Module Port	MAC Address
1	chassis-1-ioc-2-muxhostport-p...	4D:84:71:5B:10:01
2	chassis-1-ioc-2-muxhostport-p...	4D:84:71:5B:10:02
3	chassis-1-ioc-1-muxhostport-po...	4D:84:71:5B:10:03
4	chassis-1-ioc-1-muxhostport-po...	4D:84:71:5B:10:04

NIC Interfaces

Name	MAC Address	Fabric Interconnect A		Fabric Interconnect	
		Uplink Interface	Pin Group	Uplink Interface	Pin Grou
00-v...	00:25:B5:A2:0A:00	-	-	-	-
01-v...	00:25:B5:A2:0B:00	-	-	-	-
02-v...	00:25:B5:A2:0A:01	-	-	-	-
03-v...	00:25:B5:A2:0B:01	-	-	-	-

HBA Interfaces

Name	WWPN	Fabric Interconnect A	
		Uplink Interface	Pin Group
FC-NVMe-5G-MLOM-Fabric-A	20:00:00:25:B5:A2:0A:00	-	-
FC-NVMe-5G-MLOM-Fabric-B	20:00:00:25:B5:A2:0B:00	-	-
FCP-5G-MLOM-Fabric-A	20:00:00:25:B5:A2:0A:01	-	-
FCP-5G-MLOM-Fabric-B	20:00:00:25:B5:A2:0B:01	-	-

Step 3.   From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_MDS.yml Ansible playbook.

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

Step 4. Once the Ansible playbook has been run and configured both switches, it is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summertime, please see [Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 9.x](#). Sample clock commands for the United States Eastern timezone are:

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

Step 5. SSH into each switch and execute the following commands

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

Storage Configuration – ONTAP Boot Storage Setup

This chapter contains the following:

- [Ansible ONTAP Storage Configuration Part 2](#)

This configuration requires information from both the UCS server profiles and NetApp storage system. After creating the boot LUNs, initiator groups, and appropriate mappings between the two, UCS server profiles will be able to see the boot disks hosted on NetApp controllers.

Ansible ONTAP Storage Configuration Part 2

Procedure 1. Obtain the WWPNs for UCS Server Profiles (required only for FC configuration)

Step 1. From the Intersight GUI, follow: **CONFIGURE > Profiles.** Select **UCS Server Profile** and click on **[Server Profile Name].** Under **Inventory,** expand **Network Adapters** and click on the adapter. Select **Interfaces** sub tab and scroll down to find the WWPN information for various vHBAs.

← Server Profiles

aa02-esxi-1

Actions ⌄

General   Server   **Inventory**

⊞ **Expand All**

Motherboard
Boot
Management
Controller
⊞ CPUs
⊞ Memory
⊟ Network Adapters
   │ Adapter UCSX-ML-
      V5D200G_FCH254474UN
⊞ Storage Controllers
TPM

Adapter UCSX-ML-V5D200G_FCH254474UN

DCE Interfaces

🔍 Add Filter

Name	IO Module Port	MAC Address
1	chassis-1-ioc-2-muxhostp...	4D:84:71:5B:10:01
2	chassis-1-ioc-2-muxhostp...	4D:84:71:5B:10:02
3	chassis-1-ioc-1-muxhostpo...	4D:84:71:5B:10:03
4	chassis-1-ioc-1-muxhostpo...	4D:84:71:5B:10:04

NIC Interfaces

| Name | MAC Address | Fabric Interconnect A | | Fabric |
		Uplink Interface	Pin Group	Uplink Interface
00-v...	00:25:B5:A2:0A:00	-	-	-
01-v...	00:25:B5:A2:0B:00	-	-	-
02-v...	00:25:B5:A2:0A:01	-	-	-
03-v...	00:25:B5:A2:0B:01	-	-	-

HBA Interfaces

| Name | WWPN | Fabric Interconnec | |
		Uplink Interface	Pin Gro
FC-NVMe-5G-MLOM-Fabric-A	20:00:00:25:B5:A2:0A:00	-	-
FC-NVMe-5G-MLOM-Fabric-B	20:00:00:25:B5:A2:0B:00	-	-
FCP-5G-MLOM-Fabric-A	20:00:00:25:B5:A2:0A:01	-	-
FCP-5G-MLOM-Fabric-B	20:00:00:25:B5:A2:0B:01	-	-

Procedure 2.   Obtain the IQNs for UCS Server Profiles (required only for iSCSI configuration)

Step 1.   From Intersight GUI, go to: **CONFIGURE > Pools >** [**IQN Pool Name**] **> Usage** and find the IQN information for various ESXi servers:

← Pools

AA02-IQN-Pool

Actions ⌄

Details

Name
AA02-IQN-Pool

Type
IQN

Size
32

Used
4

Reserved
0

Available
28

Last Update
Oct 21, 2022 4:10 PM

Description
IQN Pool for iSCSI Configuration

Organization
AA02

Configuration & Usage

Configuration | **Usage**

✳ All Identifiers ⚙ +

4 items found | 10 ⌄ per page | 1 of 1

🔍 Add Filter

Status

4 • Used 4

Identifier	Status	Server Profile
iqn.2010-11.com.flexpod:AA02-ucshost:1	Used	aa02-esxi-4
iqn.2010-11.com.flexpod:AA02-ucshost:2	Used	aa02-esxi-2
iqn.2010-11.com.flexpod:AA02-ucshost:3	Used	aa02-esxi-6
iqn.2010-11.com.flexpod:AA02-ucshost:4	Used	aa02-esxi-8

1 of 1

Procedure 3. Configure ONTAP Boot Storage using Ansible

Step 1. Edit the following variable files to ensure proper ONTAP Boot Storage variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/vars/ontap_main.yml

Note: Update the "boot_luns_iscsi" and "boot_luns_fcp" variables under vars/ontap_main.yml file for iscsi and fcp boot storage configuration, respectively. Similarly, update the initiator IQNs and WWPNs related variables in group_vars/all.yml file. Initiator IQNs and WWPNs are for iscsi and fcp igroups, respectively.

Step 2. From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, invoke the ansible scripts for this section using the following command:

```
ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_2
```

VMware vSphere 7.0U3 Setup

This chapter contains the following:

VMware ESXi 7.0U3

This section provides detailed instructions for installing VMware ESXi 7.0U3 in a FlexPod environment. On successful completion of these steps, multiple ESXi hosts will be provisioned and ready to be added to VMware vCenter.

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

Download ESXi 7.0U3 from VMware

Procedure 1.   Download VMware ESXi ISO

Step 1.   Click the following link: [Cisco Custom Image for ESXi 7.0 U3 Install CD](#).

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

Step 2.   Download the .iso file.

Access Cisco Intersight and Launch KVM with vMedia

The Cisco Intersight KVM enables the administrators to begin the installation of the operating system (OS) through remote media. It is necessary to log into the Cisco Intersight to access KVM.

Procedure 1.   Log into Cisco Intersight and Access KVM

Step 1.   Log into Cisco Intersight.

Step 2.   From the main menu, select **Infrastructure Service** > **Servers**.

Step 3.   Find the Server with the desired Server Profile assigned and click "**...**" to see more options

Step 4.   Click **Launch vKVM**.
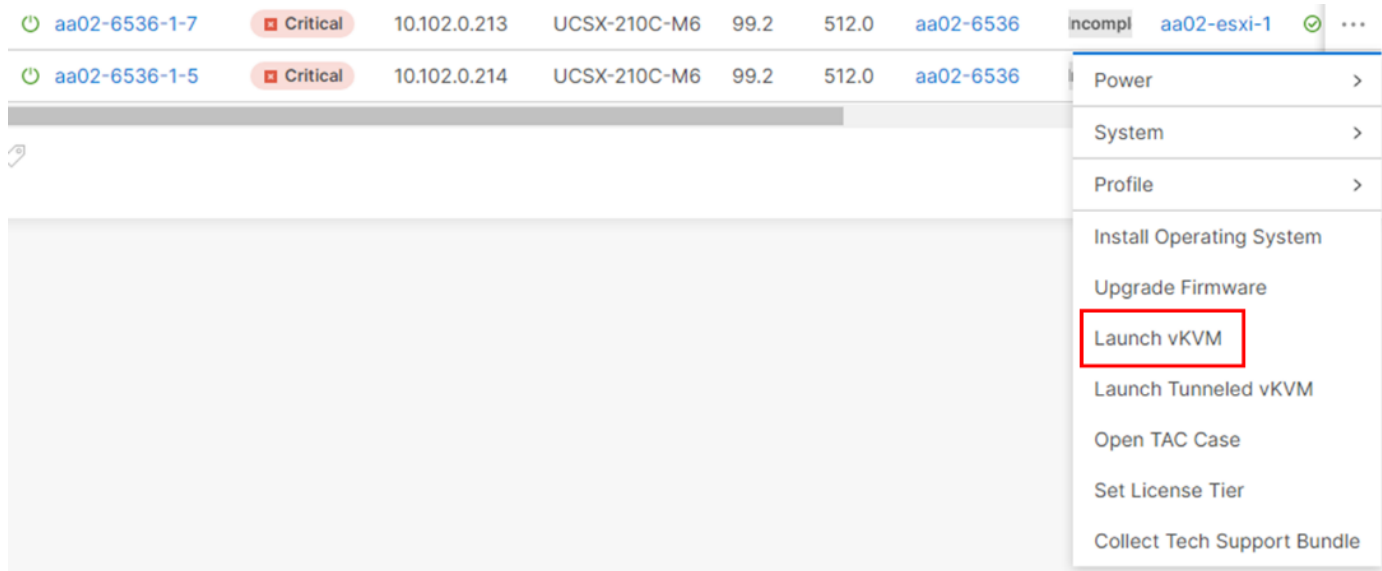
Note: Since the Cisco Custom ISO image will be mapped to the vKVM, it is important to use the standard vKVM and not the Tunneled vKVM and that the Cisco Intersight interface is being run from a subnet that has direct access to the subnet that the CIMC IPs (10.102.0.213 in this example) are provisioned on.

Step 5. Follow the prompts to ignore certificate workings (if any) and launch the HTML5 KVM console.

Step 6. Repeat steps 1 – 5 to launch the HTML5 KVM console for all the ESXi servers.

Set up VMware ESXi Installation

Procedure 1. Prepare the Server for the OS Installation

Note: Follow these steps on **each** ESXi host.

Step 1. In the KVM window, click **Virtual Media** > **vKVM-Mapped vDVD**.

Step 2. Browse and select the **ESXi installer ISO image** file downloaded in the last in Procedure 1 above (VMware-ESXi-7.0.3d-19482537-Custom-Cisco-4.2.2-a).

Step 3. Click **Map Drive**.

Step 4. Select **Power > Reset System** and **Confirm** to reboot the Server if the server is showing shell prompt. If the server is shutdown, select **Power > Power On System**.

Step 5. Monitor the server boot process in the KVM. The server should find the boot LUNs and begin to load the ESXi installer.

Note: If the ESXi installer fails to load because the software certificates cannot be validated, reset the server, and when prompted, press F2 to go into BIOS and set the system time and date to current. The ESXi installer should load properly.

Note: An alternative method to install VMware ESXi is to use the CIMC mounted vDVD, which is included in the Boot Order Policy as the last boot option. To use this option, edit the Intersight vMedia policy attached to the Server Profile Template and add an HTTPS mount to the Cisco Custom ISO mentioned above. For VMware ESXi OS boot install, an HTTPS mount is necessary, and the ISO must be placed on a server that can serve HTTPS. Since the server's SAN boot LUN is currently blank, the boot order will fall through to boot from this ISO. Once VMware ESXi is installed on the SAN boot LUN, VMware will boot from that LUN on subsequent reboots.

Install VMware ESXi

Procedure 1. Install VMware ESXi onto the bootable LUN of the UCS Servers

Note: Follow these steps on **each** host.

Step 1. After the ESXi installer is finished loading (from the last step), press **Enter** to continue with the installation.

Step 2. Read and accept the end-user license agreement (EULA). Press **F11** to accept and continue.

Note: It may be necessary to map function keys as User Defined Macros under the Macros menu in the KVM console.

Step 3. Select the NetApp boot LUN that was previously set up as the installation disk for ESXi and press **Enter** to continue with the installation.

Step 4. Select the appropriate keyboard layout and press **Enter**.

Step 5. Enter and confirm the root password and press **Enter**.

Step 6. The installer issues a warning that the selected disk will be repartitioned. Press **F11** to continue with the installation.

Step 7. After the installation is complete, press **Enter** to reboot the server. The ISO will be unmapped automatically.

Set up Management Networking for ESXi Hosts

Procedure 1. Add the Management Network for each VMware Host

Note: This is required for managing the host. To configure ESXi host with access to the management network, follow these steps on **each** ESXi host.

Step 1. After the server has finished rebooting, in the UCS KVM console, press **F2** to customize VMware ESXi.

Step 2. Log in as root, enter the password set during installation, and press **Enter** to log in.

Step 3. Use the down arrow key to select **Troubleshooting Options** and press **Enter**.

Step 4. Select Enable ESXi Shell and press Enter.

Step 5. Select **Enable SSH** and press **Enter**.

Step 6. Press **Esc** to exit the Troubleshooting Options menu.

Step 7. Select the **Configure Management Network** option and press **Enter**.

Step 8. Select Network Adapters and press **Enter**. Ensure the vmnic numbers align with the numbers under the Hardware Label (for example, vmnic0 and 00-vSwitch0-A). If these numbers do not align, note which vmnics are assigned to which vNICs (indicated under Hardware Label).

Note: In previous FlexPod CVDs, vmnic1 was selected at this stage as the second adapter in vSwitch0. It is important not to select vmnic1 at this stage. If using the Ansible configuration, if vmnic1 is selected here, the Ansible playbook will fail.

```
Network Adapters

Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.

    Device Name   Hardware Label (MAC Address)   Status
[X] vmnic0        00-vSwitch0... (...a2:0a:02)   Connected (...)
[ ] vmnic1        01-vSwitch0... (...a2:0b:02)   Connected
[ ] vmnic2        02-vDS0-i-5... (...a2:0a:03)   Connected
[ ] vmnic3        03-vDS0-i-5... (...a2:0b:03)   Connected
[ ] vmnic4        04-iSCSI-5G... (...a2:0a:04)   Connected (...)
[ ] vmnic5        05-iSCSI-5G... (...a2:0b:04)   Connected

<D> View Details  <Space> Toggle Selected        <Enter> OK  <Esc> Cancel
```

Step 9. Press **Enter**.

Note: In the UCS Configuration portion of this document, the IB-MGMT VLAN was set as the native VLAN on the 00-vSwitch0-A and 01-vSwitch0-B vNICs. Because of this, the IB-MGMT VLAN should not be set here and should remain **Not set**.

Step 10. Select **IPv4 Configuration** and press **Enter**.

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

Step 11. Select the **Set static IPv4 address and network configuration** option by using the arrow keys and space bar.

Step 12. Under **IPv4 Address**, enter the IP address for managing the ESXi host.

Step 13. Under **Subnet Mask**, enter the subnet mask.

Step 14. Under **Default Gateway**, enter the default gateway.

Step 15. Press **Enter** to accept the changes to the IP configuration.

Step 16. Select the I**Pv6 Configuration** option and press **Enter**.

Step 17. Using the spacebar, select **Disable IPv6 (restart required)** and press **Enter**.

Step 18. Select the **DNS Configuration** option and press **Enter**.

Note: If the IP address is configured manually, the DNS information must be provided.

Step 19. Using the spacebar, select Use the following DNS server addresses and hostname:

- Under **Primary DNS Server**, enter the IP address of the primary DNS server.

- Optional: Under **Alternate DNS Server,** enter the IP address of the secondary DNS server.

- Under **Hostname**, enter the fully qualified domain name (FQDN) for the ESXi host.

- Press **Enter** to accept the changes to the DNS configuration.

- Press **Esc** to exit the Configure Management Network submenu.

- Press **Y** to confirm the changes and reboot the ESXi host.

Procedure 2. (Optional) Reset VMware ESXi Host VMkernel Port MAC Address

Note: By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port it is placed on. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain the assigned MAC address unless the ESXi System Configuration is reset.

Step 1. From the ESXi console menu main screen, select **Macros > Static Macros > Ctrl + Alt + F's > Ctrl + Alt + F1** to access the VMware console command line interface.

Step 2. Log in as **root**.

Step 3. Type "`esxcfg-vmknic –l`" to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.

Step 4. To remove vmk0, type `esxcfg-vmknic –d "Management Network"`.

Step 5. To re-add vmk0 with a random MAC address, type `esxcfg-vmknic –a –i <vmk0-ip> -n <vmk0-netmask> "Management Network"`.

Step 6. Verify vmk0 has been re-added with a random MAC address by typing `esxcfg-vmknic –l`.

Step 7. Tag vmk0 as the management interface by typing `esxcli network ip interface tag add -i vmk0 -t Management`.

Step 8. When vmk0 was re-added, if a message pops up saying vmk1 was marked as the management interface, type `esxcli network ip interface tag remove -i vmk1 -t Management`.

Step 9. Press Ctrl-D to log out of the ESXi console.

Step 10. Select Macros > Static Macros > Ctrl + Alt + F's > Ctrl + Alt + F2 to return to the VMware ESXi menu.

FlexPod VMware ESXi Ansible Configuration

Procedure 1. Use Ansible to Configure All VMware ESXi Hosts from the Management Workstation

Step 1. Edit the following variable files to ensure proper VMwarevariables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/VMware/ESXIhosts/defaults/main.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/VMware/ESXIiscsi/defaults/main.yml (If using iSCSI boot)

Step 2. From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_ESXi.yml Ansible playbook:

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

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

Step 1. Click this link: https://customerconnect.vmware.com/downloads/details?downloadGroup=VC70U3H&productId=974&rPId=95488 and download the VMware-VCSA-all-7.0.3-20395099.iso.

Step 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 2 stages: installation and configuration.

Step 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.

Step 2. Mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).

Step 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 appears.

Step 4. Click **Install** to start the vCenter Server Appliance deployment wizard.

Step 5. Click **NEXT** in the Introduction section.

Step 6. Read and accept the license agreement and click **NEXT**.

Step 7. In the "vCenter Server deployment target" window, enter the 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, customers can choose the recently configured first ESXi host as an installation target. The recently configured ESXi host is shown in this deployment.

Step 8. Click **YES** to accept the certificate.

Step 9. Enter the Appliance VM name and password details shown in the "Set up vCenter Server VM" section. Click **NEXT**.

Step 10. In the "Select deployment size" section, select the Deployment size and Storage size. For example, select "Small" and "Default." Click **NEXT**.

Step 11. Select the datastore (for example, infra_datastore) for storage. Click **NEXT**.

Step 12. In the Network Settings section, configure the following settings:

  a. Select a Network: (for example, **IB-MGMT Network**)

Note: When the vCenter is running on the 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.

 b. IP version: **IPV4**

 c. IP assignment: **static**

 d. FQDN: <vcenter-fqdn>

 e. IP address: **<vcenter-ip**>

 f. Subnet mask or prefix length: **<vcenter-subnet-mask>**

 g. Default gateway: **<vcenter-gateway>**

 h. DNS Servers: <dns-server1>,<dns-server2>

Step 13. Click **NEXT**.

Step 14. Review all values and click **FINISH** to complete the installation.

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

Step 15. When Stage 1, Deploy vCenter Server, is complete, Click **CONTINUE** to proceed with stage 2.

Step 16. Click **NEXT**.

Step 17. In the vCenter Server configuration window, configure these settings:

 a. Time Synchronization Mode: Synchronize time with NTP servers.

 b. NTP Servers: NTP server IP addresses from IB-MGMT VLAN

 c. SSH access: **Enabled**.

Step 18. Click **NEXT**.

Step 19. Complete the SSO configuration as shown below (or according to your organization's security policies):
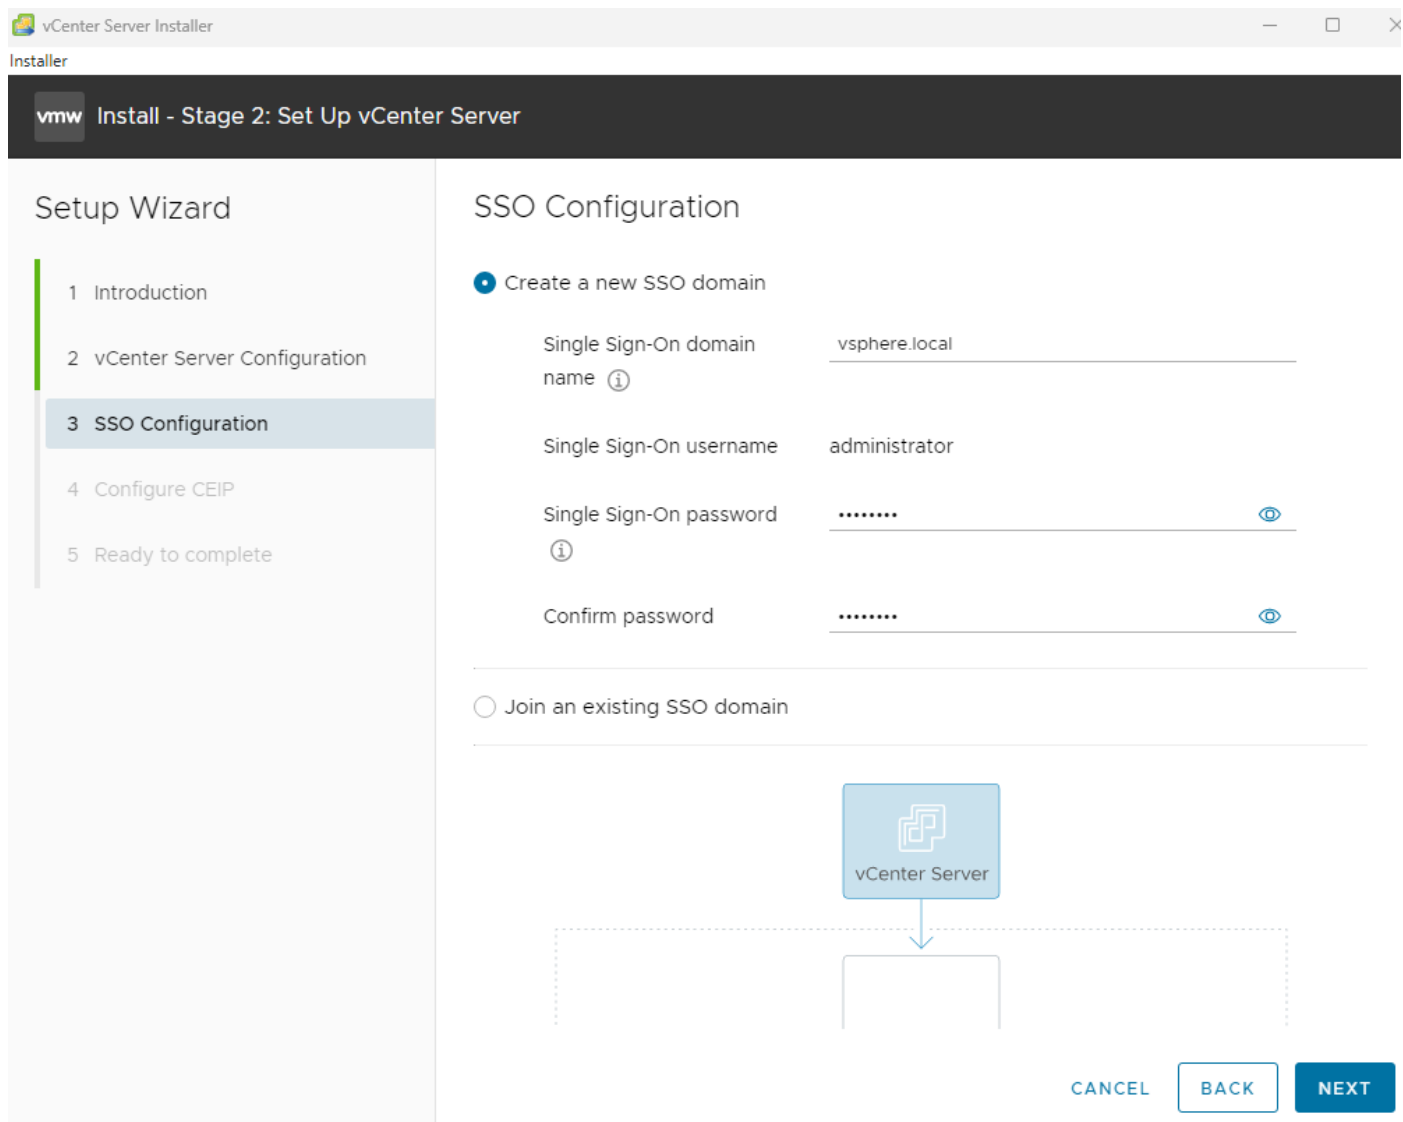
Step 20.   Click **NEXT**.

Step 21.   Decide whether to join VMware's Customer Experience Improvement Program (CEIP).

Step 22.   Click **NEXT**.

Step 23.   Review the configuration and click **FINISH**.

Step 24.   Click **OK**.

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

Step 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 setup for a 4-socket server. This setup can cause issues in the VMware ESXi cluster Admission Control.

Step 1. Open a web browser on the management workstation and navigate to the vCenter or ESXi server where the vCenter appliance was deployed and login.

Step 2. Click the **vCenter VM**, right-click and click **Edit settings**.

Step 3. In the **Edit settings** window, expand CPU and check the value of Sockets.

Step 4. If the number of Sockets match the server configuration, click **Cancel**.

Step 5. If the number of Sockets does not match the server configuration, it will need to be adjusted:

    a. Right-click the vCenter VM and click **Guest OS > Shut down**. Click **Yes** on the confirmation.

    b. When vCenter is shut down, right-click the vCenter VM and click **Edit settings**.

    c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to the server configuration.
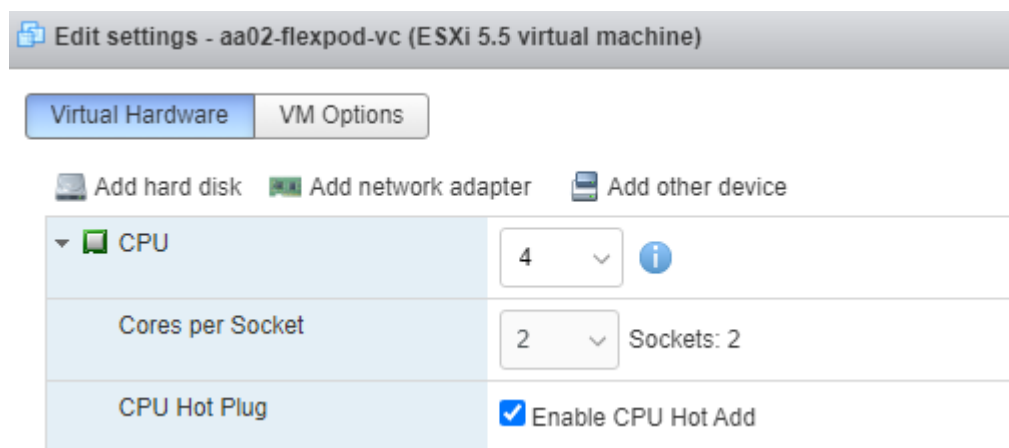
Step 6. Click **Save**.

Step 7. Right-click the vCenter VM and click **Power > Power on**. Wait approximately 10 minutes for vCenter to come up.

Procedure 4. Setup VMware vCenter Server

Step 1. Using a web browser, navigate to https://<vcenter-ip-address>:5480. Navigate the security screens.

Step 2. Log into the **VMware vCenter Server Management** interface as **root** with the root password set in the vCenter installation.

Step 3. In the menu on the left, click **Time**.

Step 4. Click **EDIT** to the right of Time zone.

Step 5. Select the appropriate Time zone and click **SAVE**.

Step 6. In the menu on the left select **Administration**.

Step 7. According to your Security Policy, adjust the settings for the root user and password.

Step 8. In the menu on the left click **Update**.

Step 9.   Follow the prompts to stage and install any available vCenter updates.

Step 10.   In the upper right-hand corner of the screen, click **root > Logout** to logout of the Appliance Management interface.

Step 11.   Using a web browser, navigate to https://<vcenter-fqdn> and navigate through security screens.

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

Step 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.

Step 13.   Log in using the Single Sign-On username (<u>administrator@vsphere.local</u>) and password created during the vCenter installation. Dismiss the Licensing warning.

Procedure 5.	Add AD User Authentication to vCenter (Optional)

Step 1.   In the **AD Infrastructure**, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).

Step 2.   Connect to https://<vcenter-fqdn> and select **LAUNCH VSPHERE CLIENT (HTML5)**.

Step 3.   Log in as **administrator@vsphere.local** (or the SSO user set up in vCenter installation) with the corresponding password.

Step 4.   Under the top-level menu, click **Administration**. In the list on the left, under **Single Sign On**, select **Configuration.**

Step 5.   In the center pane, under **Configuration**, select the **Identity Provider** tab.

Step 6.   In the list under **Type**, select **Active Directory Domain**.

Step 7.   Click **JOIN AD**.

Step 8.   Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click **JOIN**.

Step 9.   Click **Acknowledge**.

Step 10.   In the list on the left under **Deployment**, click **System Configuration**. Select the radio button to select the vCenter, then click **REBOOT NODE**.

Step 11.   Input a reboot reason and click **REBOOT**. The reboot will take approximately 10 minutes for full vCenter initialization.

Step 12.   Log back into the vCenter vSphere HTML5 Client as Administrator@vsphere.local.

Step 13.   Under the top-level menu, click **Administration**. In the list on the left, under **Single Sign On**, click **Configuration**.

Step 14.   In the center pane, under **Configuration**, click **the Identity Provider** tab. Under **Type**, select **Identity Sources**. Click **ADD**.

Step 15.   Make sure Active Directory (Integrated Windows Authentication) is selected, your Windows Domain name is listed and Use machine account is selected. Click **ADD**.

Step 16.   In the list select the **Active Directory (Integrated Windows Authentication)** Identity source type. If desired, select SET AS DEFAULT and click **OK**.

Step 17.   On the left under Access Control, select **Global Permissions**.

Step 18. In the center pane, click the **ADD** to add a Global Permission.

Step 19. In the **Add Permission** window, select your AD domain for the Domain.

Step 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 login to vCenter as an Administrator.

Step 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.

Step 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@domain.

vCenter and ESXi Ansible Setup

Procedure 1. Configure the VMware vCenter and the three management ESXi hosts

Step 1. Edit the following variable files to ensure proper VMware variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/VMware/ESXIpostvC/defaults/main.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/VMware/ESXIpostvCiscsi/defaults/main.yml

Step 2. From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_vCenter.yml Ansible playbook:

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

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

Step 3. In the center pane under **Virtual Machines**, click **Swap File location**.

Step 4. On the right, click **EDIT**.

Step 5. Select infra_swap and click **OK**.

Edit Swap File Location | aa02-esxi-2.flexpodb4.cisco.com ✕

Select a location to store the swap files.

◯ Virtual machine directory

    Store the swap files in the same directory as the virtual machine.

🔵 Use a specific datastore

    Store the swap files in the specified datastore. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

	Name ▼	Capacity ▼	Provisioned ▼	Free Space ▼	Type ▼	Thin Provisioned ▼
🔵	infra_swap	200 GB	19.32 MB	199.98 GB	NFS41	Supported
◯	infra_datasto…	1 TB	769.34 GB	951.53 GB	NFS41	Supported
◯	vCLS	100 GB	322.71 MB	99.68 GB	NFS41	Supported

3 items

CANCEL **OK**

Step 6. Repeat steps 1-6 to set the swap file location for each ESXi host.

Step 7. Right-click the cluster and select **Settings**. In the center pane under vSphere Cluster Services, select **Datastores**. In the center of the window, click **ADD**. Select the vCLS datastore and click **ADD**.

Add datastores | FlexPod-MGMT ✕

Select one or more datastores to add to the 'Allowed' list for vSphere Cluster Services (vCLS) VM disk placement. Solution blocked datastores are not visible in the table below. Order of allowed datastores list does not guarantee the order of placement of vCLS VM disks.

ⓘ Click of Add could result in storage migration of vCLS VM disks, which could impact the health of vCLS resulting in a downtime of DRS.   Learn more

Filter   Selected (1)

▼ Filter

☐	Name	↑	Type	Capacity	Free
☐	🗄 infra_datastore		NFS 4.1	1 TB	951.53 GB
☐	🗄 infra_swap		NFS 4.1	200 GB	199.98 GB
☐	🗄 nvme_datastore		VMFS 6	971.75 GB	970.33 GB
☑	🗄 vCLS		NFS 4.1	100 GB	99.68 GB

☑ 1 ▯▯   4 items

CANCEL   **ADD**

Step 8.   Select the first ESXi host. In the center pane under **Configure > Storage**, click **Storage Devices**. Make sure the NETAPP Fibre Channel Disk LUN 0 or NETAPP iSCSI Disk LUN 0 is selected.

Step 9.   Click the **Paths** tab.

Step 10.   Ensure that 4 paths appear, two of which should have the status Active (I/O). The output below shows the paths for an iSCSI LUN.

Storage Devices

REFRESH   ATTACH   DETACH   RENAME   TURN ON LED   TURN OFF LED   ERASE PARTITIONS   ...

	Name		LUN	
☐	Local ATA Disk (t10.ATA_____Micron_5300_MTFDDAV240TDS_____MSA24220AZL)		0	
☐	Local ATA Disk (t10.ATA_____Micron_5300_MTFDDAV240TDS_____MSA24220AZN)		0	
☑	NETAPP iSCSI Disk (naa.600a0980383135466224546943367858)		0	
☐	Local Marvell Processor (eui.0050430000000000)		0	

☑ 1 ☐   EXPORT ˅   4 item

Properties   **Paths**   Partition Details

ENABLE   DISABLE

	Runtime Name		Status		Target		Name		Preferred	
○	vmhba64:C0:T0:L0		◆ Active (I/O)		iqn.1992-08.com.netapp:sn...		vmhba64:C0:T0:L0			
○	vmhba64:C3:T0:L0		◆ Active (I/O)		iqn.1992-08.com.netapp:sn...		vmhba64:C3:T0:L0			
○	vmhba64:C2:T0:L0		◆ Active		iqn.1992-08.com.netapp:sn...		vmhba64:C2:T0:L0			
○	vmhba64:C1:T0:L0		◆ Active		iqn.1992-08.com.netapp:sn...		vmhba64:C1:T0:L0			

Step 11.   Repeat steps 9–11 for all ESXi hosts.

Procedure 2.   VMware ESXi 7.0 U3 TPM Attestation

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

Step 1.   For Cisco UCS servers that have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client.

Step 2.   In the vCenter HTML5 Interface, under **Hosts and Clusters** select the cluster.

Step 3.   In the center pane, click the **Monitor** tab.

Step 4. Click **Monitor > Security**. The Attestation status will show the status of the TPM:

Security

		Name	Attestation	Last verified	Attested by	TPM version	TXT	↑	Message
○	⠿	🖥 aa02-esxi-1.fl…	Passed	10/31/2022, 4:57:22 …	vCenter Server	2.0	false		
○	⠿	🖥 aa02-esxi-2.f…	Passed	10/31/2022, 9:05:02 …	vCenter Server	2.0	false		
○	⠿	🖥 aa02-esxi-3.f…	Passed	10/31/2022, 10:45:56 …	vCenter Server	2.0	false		
○	⠿	🖥 aa02-esxi-4.f…	Passed	10/31/2022, 10:45:59 …	vCenter Server	2.0	false		

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

Procedure 3. Avoiding Boot Failure When UEFI Secure Booted Server Profiles are Moved

Typically, hosts in FlexPod Datacenter are configured for boot from SAN. Cisco UCS supports stateless compute where a server profile can be moved from one blade or compute node to another seamlessly.

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

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

- Host installed with ESXi 7.0 U2 or above

- Boot mode is UEFI Secure

- Error message: Unable to restore system configuration. A security violation was detected. https://via.vmw.com/security-violation.

VMware ESXi 7.0.3 (VMKernel Release Build 19482537)

Cisco Systems Inc UCSX-210C-M6

2 x Intel(R) Xeon(R) Platinum 8358P CPU @ 2.60GHz
2 TiB Memory

The system has found a problem on your machine and cannot continue.

Unable to restore the system configuration. A security violation was detected. https://via.vmw.com/security-violation

No port for remote debugger.

Resolution:

Step 1. Log into the host using SSH.

Step 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}
529012-402326-326163-088960-184364-097014-312164-590080-407316-660658-634787-601062-601426-263837-330828-1970
47
```

Step 3. Store the keys from all hosts in a safe location.

Step 4. After associating the Server Profile to the new compute-node or blade, stop the ESXi boot sequence by pressing Shift + O when you see the ESXi boot screen.
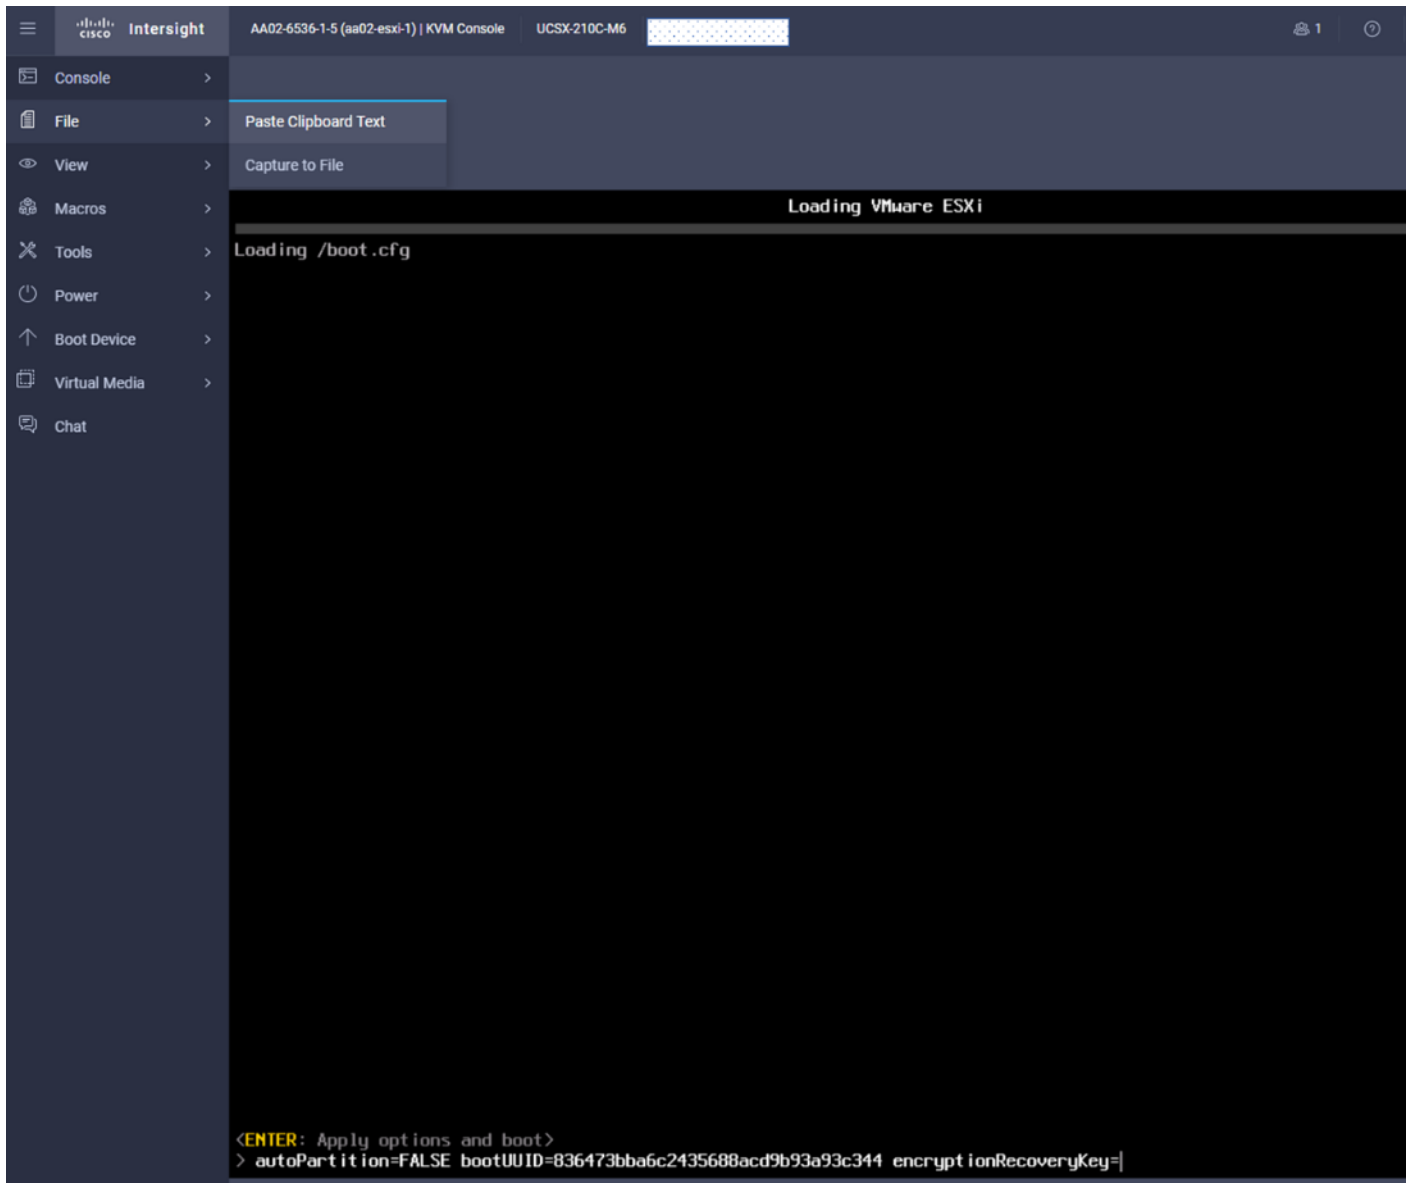
Step 5. Add the recovery key using following boot option: encryptionRecoveryKey=*recovery_key.* Press Enter to continue the boot process.

Step 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.

Storage Configuration – ONTAP NVMe Configuration and Finalizing ONTAP Storage

This chapter contains the following:

-

Ansible ONTAP Storage Configuration Part 3

Procedure 1. Configure the ONTAP NVMe setup and finalize ONTAP storage using Ansible

Step 1. Edit the following variable files to ensure proper variables are entered:

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/vars/ontap_main.yml

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

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

Step 2. From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, invoke the ansible scripts for this section using the following command:

```
ansible-playbook -i inventory Setup_ONTAP.yml -t ontap_config_part_3
```

Procedure 2. Configure ESXi Host NVMe over FC and NVMe over TCP Datastore

Step 1. To verify that the NVMe Fibre Channel Disk is mounted on each ESXi host, log into the VMware vCenter using a web-browser.

Step 2. Under **Hosts and Clusters** select an ESXi host running FC-NVMe. In the center pane, go to **Configure > Storage > Storage Devices**. The NVMe Fibre Channel Disk should be listed under Storage Devices.

Step 3. Select the NVMe Fibre Channel Disk, then select **Paths** underneath. Verify 2 paths have a status of Active (I/O) and 2 paths have a status of Active.

Step 4. Repeat [Step 3](#) for all the FC-NVMe hosts.

Step 5. Under **Hosts and Clusters** select an ESXi host running NVMe-TCP. In the center pane, go to **Configure > Storage > Storage Adapters**.

Step 6. Click **ADD SOFTWARE-ADAPTER > Add NVMe over TCP adapter**. Use the pulldown to select **vmnic4/nenic** and click **OK**. A new vmhba should appear under Storage Adapters.

Step 7. Click **ADD SOFTWARE-ADAPTER > Add NVMe over TCP adapter** to add a second vmhba. Use the pulldown to select **vmnic5/nenic** and click **OK**. A new vmhba should appear under Storage Adapters.

Step 8. Select the first VMware NVMe over TCP Storage Adapter added (for example, vmhba65). In the middle of the window, select the **Controllers** tab. Click **ADD CONTROLLER**.

Step 9.　Enter the IP address of nvme-tcp-lif-01a and click **DISCOVER CONTROLLERS**. Select the two controllers in the Infra-NVMe-TCP-A subnet and click **OK**. The two controllers should now appear under the Controllers tab.

Step 10.　Select the second VMware NVMe over TCP Storage Adapter added (for example, vmhba66). In the middle of the window, select the **Controllers** tab. Click **ADD CONTROLLER**.

Step 11.　Enter the IP address of nvme-tcp-lif-02b and click **DISCOVER CONTROLLERS**. Select the two controllers in the Infra-NVMe-TCP-B subnet and click **OK**. The two controllers should now appear under the Controllers tab.

Step 12.　Repeat steps 5-11 for all ESXi hosts running NVMe-TCP.

Step 13.　For any one of these hosts, right-click the host under **Hosts and Clusters** and select **Storage > New Datastore**. Leave VMFS selected and click **NEXT**.

Step 14. Name the datastore (for example, nvme_datastore) and select the **NVMe Disk**. Click **NEXT**.

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

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

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

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

Note: If any hosts are missing from the list, it may be necessary to put the host in Maintenance Mode and reboot the host. If you happen to have hosts with both FC-boot and iSCSI-boot and are running both FC-NVMe and NVMe-TCP, notice that the same datastore is mounted on both types of hosts and that the only difference in the storage configuration is what LIF the traffic is coming in on.

FlexPod Management Tools Setup

This chapter contains the following:

Cisco Intersight Hardware Compatibility List (HCL) Status

Cisco Intersight evaluates the compatibility of customer's UCS system to check if the hardware and software have been tested and validated by Cisco or Cisco partners. Intersight reports validation issues 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 uses Cisco UCS Tools. The Cisco UCS Tools is part of VMware ESXi Cisco custom ISO, and no additional configuration is required.

For more details on Cisco UCS Tools manual deployment and troubleshooting, refer to:
https://intersight.com/help/saas/resources/cisco_ucs_tools#about_cisco_ucs_tools

Procedure 1. View Compute Node Hardware Compatibility

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

NetApp ONTAP Tools 9.11 Deployment

The 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 topic describes the deployment procedures for the NetApp ONTAP Tools for VMware vSphere.

NetApp ONTAP Tools for VMware vSphere 9.11 Pre-installation Considerations

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

- Protocol licenses (NFS, FCP, and/or iSCSI)

- NetApp FlexClone ((optional) Required for performing test failover operations for SRA and for vVols operations of VASA Provider.

- NetApp SnapRestore (for backup and recovery).

- The NetApp SnapManager Suite.

- NetApp SnapMirror or NetApp SnapVault (Optional – required for performing failover operations for SRA and VASA Provider when using vVols replication).

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

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

Table 6.   Port Requirements for NetApp ONTAP Tools

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

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

Procedure 1.   Install NetApp ONTAP Tools via Ansible

Step 1.   Clone the repository from https://github.com/NetApp-Automation/ONTAP-Tools-for-VMware-vSphere.

Step 2.   Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.

Step 3.   Update the following variable files:

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

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

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

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

Procedure 2.   Download and Install 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, NetApp ONTAP-Tools can be utilized to install the NetApp NFS plug-in. The steps below upload the latest version of the plugin to ONTAP tools.

Step 1.   Download the NetApp NFS Plug-in 2.0 for VMware file from: https://mysupport.netapp.com/site/products/all/details/nfsplugin-vmware-vaai/downloads-tab.

Step 2.   Unzip the file and extract NetApp_bootbank_NetAppNasPlugin_2.0-15.vib from **vib20 > NetAppNasPlugin.**

Step 3. Rename the .vib file to NetAppNasPlugin.vib to match the predefined name that ONTAP tools uses.

Step 4. Click **Settings** in the ONTAP tool Getting Started page.

Step 5. Click **NFS VAAI Tools** tab.

Step 6. Click **Change** in the Existing version section.

Step 7. Browse and select the renamed .vib file, and then click **Upload** to upload the file to the virtual appliance.
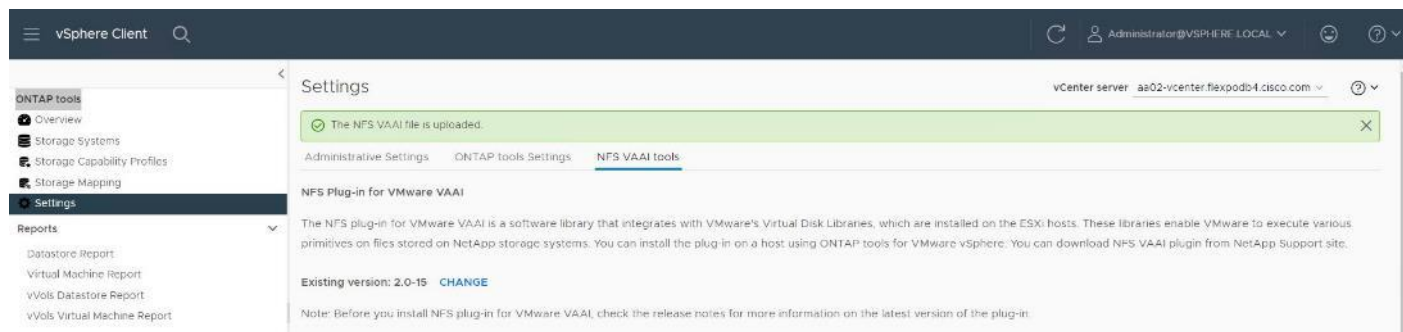
Note: The next step is only required on the hosts where NetApp VAAI plug-in was not installed alongside Cisco VIC driver installation.

Step 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.

Step 9. Reboot the ESXi host after the installation finishes.

Procedure 3. Verify the VASA Provider

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

Step 1. From the vSphere Client, click **Menu > ONTAP tools**.

Step 2. Click **Settings**.

Step 3. Click **Manage Capabilities** in the Administrative Settings tab.

Step 4. In the Manage Capabilities dialog box, click **Enable VASA Provider** if it was not pre-enabled.

Step 5. Enter the IP address of the virtual appliance for ONTAP tools, VASA Provider, and VMware Storage Replication Adapter (SRA) and the administrator password, and then click **Apply**.

Manage Capabilities

🟢 **Enable VASA Provider**

vStorage APIs for Storage Awareness (VASA) is a set of application program interfaces (APIs) that enables vSphere vCenter to recognize the capabilities of storage arrays.

　⚪ **Enable vVols replication**

　Enables replication of vVols when used with VMware Site Recovery Manager 8.3 or later.

⚪ **Enable Storage Replication Adapter (SRA)**

Storage Replication Adapter (SRA) allows VMware Site Recovery Manager (SRM) to integrate with third party storage array technology.

Enter authentication details for VASA Provider and SRA server:

IP address or hostname:　　　10.102.1.99

Username:　　　　　　　　　Administrator

Password:　　　　　　　　　••••••••

Procedure 4.　Update Host and Storage Data

Step 1.　　　From the vSphere Client Home Page, click **Hosts and Clusters**.

Step 2.　　　Right-click the FlexPod-DC datacenter, click **NetApp ONTAP tools > Update Host and Storage Data**.

Step 3.　　　On the Confirmation dialog box, click **OK**. It might take a few minutes to update the data.
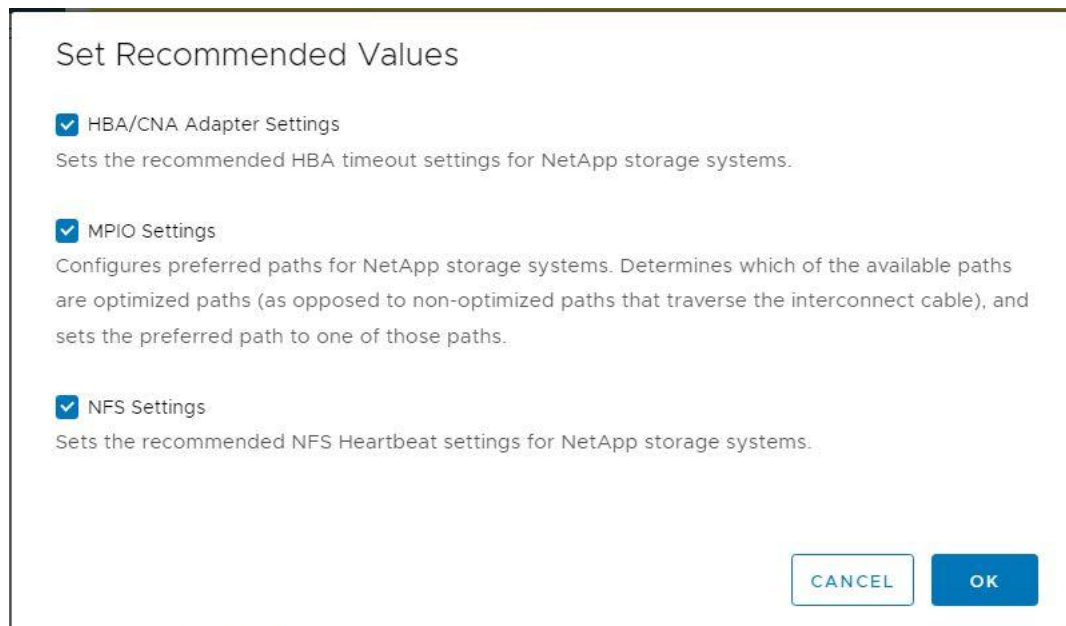
Note: ONTAP tools enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers.

Step 1.   From the VMware vSphere Web Client Home page, click **vCenter** > **Hosts and Clusters**.

Step 2.   Select a host and then click **Actions** > **NetApp ONTAP tools** > **Set Recommended Values**.

Step 3.   In the NetApp Recommended Settings dialog box, select all the applicable values for the ESXi host.

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

Step 4.   Click **OK**.

Provision Datastores using ONTAP Tools (Optional)

Using 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 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

In order to leverage the automation features of VASA two primary components must first be configured. 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 VM Storage Policy. NetApp ONTAP tools comes with several pre-configured SCPs such as Platinum, Bronze, and so on.

Note: The ONTAP tools for VMware vSphere plug-in also allows you to set Quality of Service (QoS) rule using a combination of maximum and/or minimum IOPs.

Procedure 1.	Review or Edit the Built-In Profiles Pre-Configured with ONTAP Tools

Step 1. From the vCenter console, click **Menu** > **ONTAP tools**.

Step 2. In the NetApp ONTAP tools click **Storage Capability Profiles**.

Step 3. Select the **Platinum** Storage Capability Profile and select **Clone** from the toolbar.
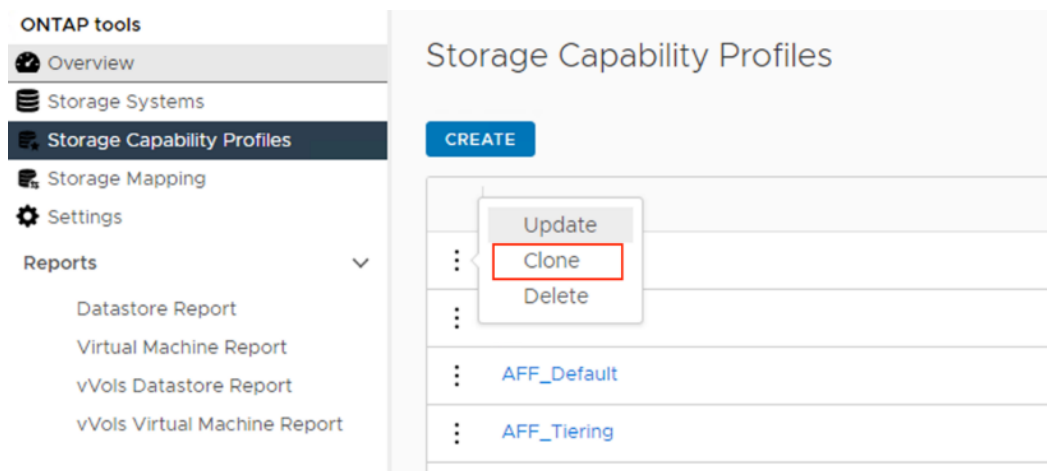
Step 4. Enter a name for the cloned SCP (for example, AFF_Platinum_Encrypted) and add a description if desired. Click **NEXT**.

Step 5. Select **All Flash FAS(AFF)** for the storage platform and click **NEXT**.

Step 6. Select **None** to allow unlimited performance or set a the desired minimum and maximum IOPS for the QoS policy group. Click **NEXT**.

Step 7. On the Storage attributes page**,** change the Encryption and Tiering policy to the desired settings and click **NEXT**. In the example below, Encryption was enabled.

Step 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 prede-fined 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.

Step 1. From the vCenter console, click **Menu** > **Policies and Profiles**.

Step 2. Select VM Storage Policies and click **CREATE**.

Step 3. Create a name for the VM storage policy and enter a description and click **NEXT**.

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

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

Step 6. All the datastores with matching capabilities are displayed, click **NEXT**.

Step 7. Review the policy summary and click **FINISH**.

Procedure 3. Provision NFS Datastore

Step 1. From the vCenter console, click **Menu > ONTAP tools**.

Step 2. From the ONTAP tools Home page, click **Overview**.

Step 3. In the Getting Started tab, click **Provision**.

Step 4. Click **Browse** to select the destination to provision the datastore.

Step 5. Select the type as **NFS** and Enter the datastore name (for example, NFS_DS_1).

Step 6. Provide the size of the datastore and the NFS Protocol.

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

Step 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.

Step 9. Click **NEXT**.

Step 10. Select the aggregate name and click **NEXT**.

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

Step 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.

Step 13. Distributed datastore is supported from ONTAP 9.8, which provides FlexGroup volume on ONTAP storage. To create a Distributed Datastore across the ONTAP Cluster select NFS 4.1 and check the box for Distributed Datastore data across the ONTAP Cluster as shown below.

Procedure 4.	Provision FC Datastore

Step 1. From the vCenter console, click **Menu > ONTAP tools**.

Step 2. From the ONTAP tools Home page, click **Overview**.

Step 3. In the Getting Started tab, click **Provision**.

Step 4. Click **Browse** to select the destination to provision the datastore.

Step 5.   Select the type as **VMFS** and Enter the datastore name.

Step 6.   Provide the size of the datastore and the FC Protocol.

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

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

Step 9.   Click **NEXT**.

Step 10.   Select the aggregate name and click **NEXT**.

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

Step 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

Step 1.   Log into vCenter and navigate to the **VMs and Templates** tab and click to select the datacenter (for example, FlexPod-DC).

Step 2.   Click **Actions** and click **New Virtual Machine**.

Step 3.   Click **Create a new virtual machine** and click **NEXT**.

Step 4.   Enter a name for the VM and select the datacenter (for example, FlexPod-DC).

Step 5.   Select the cluster (for example, AA17-Cluster) and click **NEXT**.

Step 6.   Select the VM storage policy from the selections and select a compatible datastore. Click **NEXT**.

New Virtual Machine

✔ 1 Select a creation type
✔ 2 Select a name and folder
✔ 3 Select a compute resource
   4 Select storage
   5 Select compatibility
   6 Select a guest OS
   7 Customize hardware
   8 Ready to complete

Select storage
Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy   VM AFF Platinum Encrypted Storage Policy ⌄

☐ Disable Storage DRS for this virtual machine

	Name	▼	Storage Con ▼	Capacity	▼	Provisione ▼	Free	▼	Type	▼	Clust
○	infra_datastore_1		Compatible	1 TB		798.82 GB	949.49 GB		NFS v3		
○	infra_datastore…		Compatible	1 TB		544.71 GB	1,005.05 GB		NFS v3		
○	Infra_Swap_DS		Compatible	300 GB		581.62 MB	299.43 GB		NFS v3		
○	NX_FC_DS_01		Compatible	500 GB		41.41 GB	458.59 GB		VMFS 6		

Step 7.   Select Compatibility (for example, ESXi 7.0 U2 or later) and click **NEXT**.

Step 8.   Select the Guest OS and click **NEXT**.

Step 9.   Customize the hardware for the VM and click **NEXT**.

Step 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 – vVol (Optional)

NetApp VASA Provider enables customers 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"). A virtual machine can be spread 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 on vVOL datastore configuration, see:
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_vmware_7u2.html#VirtualVolumesvVolOptional

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 VMs running on 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 on deploying SnapCenter server for application backups can be found in the documentation listed below.

This guide focuses on deploying and configuring the SnapCenter plug-in for VMware vSphere to protect virtual machines and VM datastores.

Note: You must install SnapCenter Server and the necessary plug-ins to support application-consistent backups for Microsoft 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 CVD's and technical reports for detailed information on 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: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/flexpod-sql-2019-vmware-on-ucs-netapp-ontap-wp.html
- 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:

- 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 7.   Port Requirements

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

License Requirements for SnapCenter Plug-In for VMware vSphere

The licenses listed in Table 8 are required on the ONTAP storage system to backup and restore VM's in the virtual infrastructure:

Table 8.   SnapCenter Plug-in for VMware vSphere License Requirements

Product	License Requirements
ONTAP	**SnapManager Suite:** Used for backup operations One of these: SnapMirror or SnapVault (for secondary data protection regardless of the type of relationship)

Product	License Requirements
ONTAP Primary Destinations	To perform protection of VMware VMs and datastores the following licenses should be installed: **SnapRestore**: used for restoring operations **FlexClone**: used for mount and attach operations
ONTAP Secondary Destinations	To perform protection of VMware VMs and datastores only: **FlexClone**: used for mount and attach operations
VMware	**vSphere Standard, Enterprise, or Enterprise Plus** 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, SnapCenter cannot be used 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. Deploy the SnapCenter Plug-In for VMware vSphere 4.7 using Ansible

Step 1. Clone the repository from https://github.com/NetApp-Automation/SnapCenter-Plug-in-for-VMware-vSphere.

Step 2. Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.

Step 3. Update the following variable files:

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

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

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

Note: The above ansible playbook will install SnapCenter Plug-in in vCenter and will also add ONTAP Storage System.

NetApp SnapCenter Plug-in 4.7 Configuration

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

Step 1. Navigate to VMware vSphere Web Client URL https://<vCenter Server>.

Note: If you're currently logged into vCenter, logoff, close the open tab and sign-on again to access the newly installed SnapCenter Plug-in for VMware vSphere.

Step 2. After logging on, a blue banner will be displayed indicating the SnapCenter plug-in was successfully deployed. Click **Refresh** to activate the plug-in.

Step 3. On the VMware vSphere Web Client page, select **Menu > SnapCenter Plug-in for VMware vSphere** to launch the SnapCenter Plug-in for VMware GUI.

Step 4. When the storage system is added, you can create backup policies and take scheduled backup of VMs and datastores. The SnapCenter plug-in for VMware vSphere allows backup, restore and on-demand backups.

For more information on backup policy configuration, refer to this CVD:
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_vmware_7u2.html#FlexPodManagementToolsSetup

Active IQ Unified Manager 9.11P1 Installation

Active IQ Unified Manager enables you to monitor and manage the health and performance of 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 Cisco Intersight.

This subject describes the procedure to deploy NetApp Active IQ Unified Manager 9.11P1 as a virtual appliance. Table 9 lists the recommended configuration for the VM.

Table 9.   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 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 a second instance of Active IQ Unified Manager is needed. See the Unified Manager Best Practices Guide (TR-4621) for more details.

Procedure 1.   Install NetApp Active IQ Unified Manager 9.11P1 using Ansible

Step 1. Clone the repository from https://github.com/NetApp-Automation/NetApp-AIQUM.

Step 2. Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.

Step 3. Update the variable files as mentioned in the README document in the repository.

Step 4. To install AIQUM and add an ONTAP cluster, invoke the below ansible playbook:

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

Configure Active IQ Unified Manager

Procedure 1. Initial Setup

Step 1. Launch a web browser and log into Active IQ Unified Manger using the URL shown in the VM console.

Step 2. Enter the email address that Unified Manager will use to send alerts and the mail server configuration. Click **Continue**.

Step 3. Select **Agree and Continue** on the Set up AutoSupport configuration.

Step 4. Check the box for **Enable API Gateway** and click **Continue.**

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

Step 6. Click **Add**.

Step 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.

Note: Adding ONTAP Cluster to AIQUM is also automated via Ansible. So, you can skip this step if already performed through Ansible.

Procedure 2. Review Security Compliance with Active IQ Unified Manager

Active IQ Unified Manager identifies issues and makes recommendations to improve the security posture of ONTAP. Active IQ Unified Manager evaluates ONTAP storage based on recommendations made in the Security Hardening Guide for 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 ONTAP 9.

Note: All events identified do not inherently apply to all environments, for example, FIPS compliance.

The status icons in the security cards have the following meanings in relation to their compliance:
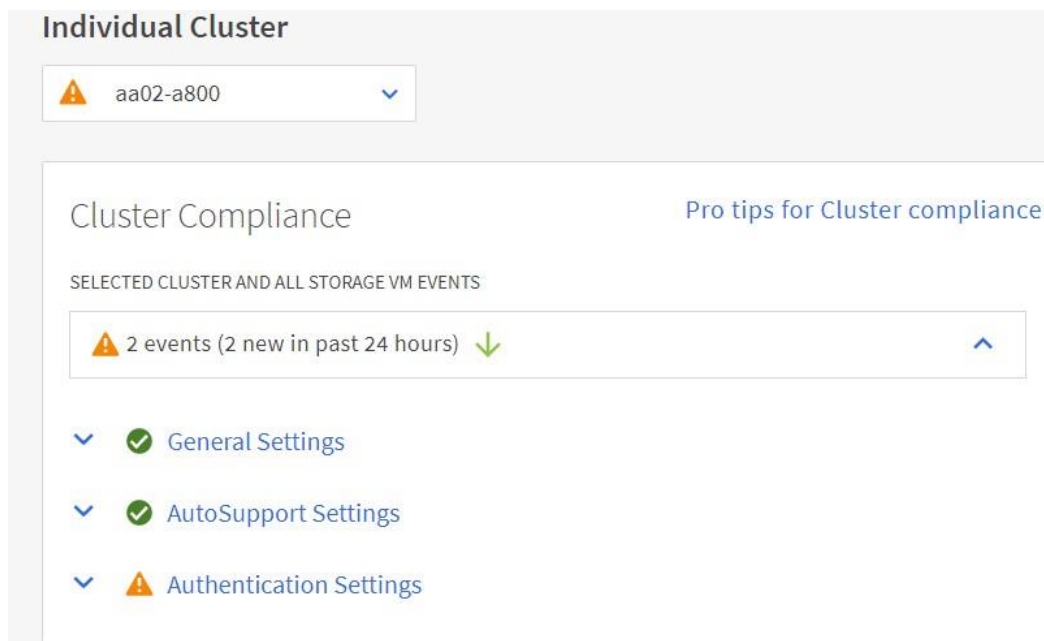
- ✅ - The parameter is configured as recommended.

- ⚠️ - The parameter is not configured as recommended.

- ℹ️ - 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.

Step 1. Navigate to the URL of the Active IQ Unified Manager and login.

Step 2. Select the **Dashboard** from the left menu bar in Active IQ Unified Manager.

Step 3. Locate the **Security** card and note the compliance level of the cluster and SVM.

Step 4. Click the blue arrow to expand the findings.

Step 5. Locate Individual Cluster section and the Cluster Compliance card. From the drop-down list select **View All**.

Step 6. Select an event from the list and click the name of the event to view the remediation steps.

Step 7. Remediate the risk if applicable to current environment and perform the suggested actions to fix the issue.

Remediate Security Compliance Findings

Note: Active IQ identifies several security compliance risks after installation that can be immediately corrected to improve the security posture of ONTAP. Click on the event name to get more information and suggested actions to fix the issue.

Deploy Cisco Intersight Assist Appliance

Cisco Intersight 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 Intersight device connector, Cisco Intersight Assist virtual appliance enables Cisco Intersight to communicate with these devices.

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

Figure 5. Managing NetApp and VMware vCenter through Cisco Intersight using Cisco Intersight Assist

---
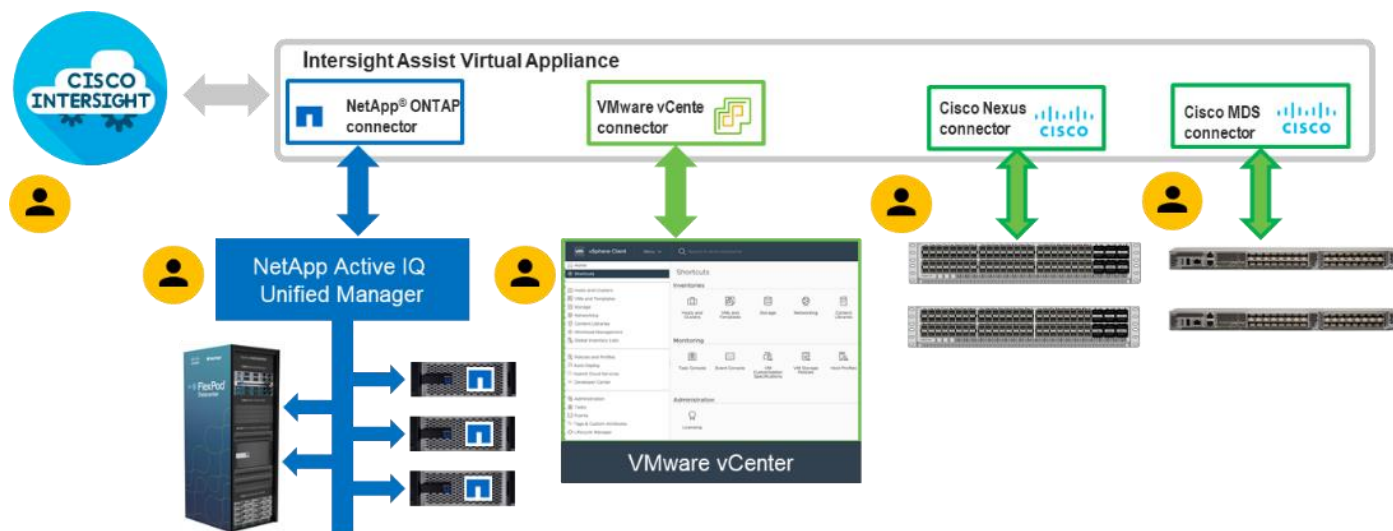
Procedure 1. Install Cisco Intersight Assist

Step 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
https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-499.

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

Procedure 2. Set up DNS entries

Step 1. Setting up Cisco Intersight Virtual Appliance requires an IP address and 2 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. This must be defined as an A record and 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. Hostnames in this format are used internally by the appliance to manage device connections.

Step 2. In this lab deployment the following information was used to deploy a Cisco Intersight Assist VM:

- **Hostname:** aa02-assist.flexpodb4.cisco.com

- **IP address**: 10.102.1.96

- **DNS Entries** (Windows AD/DNS):
  - A Record

aa02-assist	Host (A)	10.102.1.96	static

  - CNAME:

dc-aa02-assist	Alias (CNAME)	aa02-assist.flexpodb4.cisco.com.	static

  - PTR (reverse lookup):

10.102.1.96   Pointer (PTR)   aa02-assist.flexpodb4.cisco.com.   static

For more information, refer to:
https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Cisco_Intersight_Appliance_Getting_Started_Guide/b_Cisco_Intersight_Appliance_Install_and_Upgrade_Guide_chapter_00.html.

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.**

Step 1.   From Hosts and Clusters, right-click the cluster and click **Deploy OVF Template**.

Step 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**.

Step 3.   Name the Intersight Assist VM and select the location. Click **NEXT**.

Step 4.   Select the cluster and click **NEXT**.

Step 5.   Review details, click **Ignore All**, and click **NEXT**.

Step 6.   Select a deployment configuration. If only the Intersight Assist functionality is needed, a deployment size of **Tiny** can be used. If Intersight Workload Optimizer (IWO) is being used in this Intersight account, use the **Small** deployment size. Click **NEXT**.

Step 7.   Select the appropriate datastore (for example, infra_datastore) for storage and select the **Thin Provision** virtual disk format. Click **NEXT**.

Step 8.   Select appropriate management network (for example, IB-MGMT Network) for the OVA. Click **NEXT**.

Note: The Cisco Intersight Assist VM must be able to access both the IB-MGMT network on FlexPod and Intersight.com. Select and configure the management network appropriately. If selecting IB-MGMT network on FlexPod, make sure the routing and firewall is setup correctly to access the Internet.

Step 9.   Fill in all values to customize the template. Click **NEXT**.

Step 10.   Review the deployment information and click **FINISH** to deploy the appliance.

Step 11.   When the OVA deployment is complete, right-click the Intersight Assist VM and click **Edit Settings**.

Step 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:

Edit Settings | aa02-assist

Virtual Hardware VM Options

∨ CPU	16 ∨
Cores per Socket	1 ∨
	Sockets: 16

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

Edit Settings | aa02-assist

Virtual Hardware VM Options

∨ CPU *	16 ∨
Cores per Socket	8 ∨
	Sockets: 2

Step 14. Click **OK**.

Step 15. Right-click the Intersight Assist VM and select **Power > Power On**.

Step 16. When the VM powers on and 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.

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

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

Step 19.   Log into Cisco Intersight and connect to the appropriate account.

Step 20.   From Cisco Intersight, at the top select **System**, then click **Administration > Targets**.

Step 21.   Click **Claim a New Target**. Select Cisco Intersight Assist and click **Start**.

Step 22.   Copy and paste the Device ID and Claim Code shown in the Intersight Assist web interface to the Cisco Intersight Device Claim window.

Step 23.   Select the Resource Group and click **Claim**.

Claim Cisco Intersight Assist Target

To claim your target, provide the Device ID, Claim Code and select the appropriate Resource Groups.

General

Device ID *   ⓘ   Claim Code *   ⓘ

Resource Groups

ⓘ Select the Resource Groups if required. However, this selection is not mandatory as one or more Resource Group type is 'All'.
    The claimed target will be part of all Organizations with the Resource Group type 'All'.

1 items found   10 ∨   per page |◁ ◁   1   of 1 ▷ ▷| ⚙

	Name	Usage	Description
☐	AA02-rg	AA02	

|◁ ◁   1   of 1 ▷ ▷|

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

Step 25. In the Intersight Assist web interface, verify that Intersight Assist is Connected Successfully, and click **Continue**.

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

Note: The Cisco Intersight Assist VM will reboot during the software download process. It will be necessary to refresh the Web Browser after the reboot is complete to follow the status of the download process.

Step 26. When the software download is complete, an Intersight Assist login screen will appear.

Step 27. Log into Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Intersight Assist status and **log out** of Intersight Assist.

Claim VMware vCenter using Cisco Intersight Assist Appliance

Procedure 1.   Claim the vCenter from Cisco Intersight

Step 1. Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2. Select **System** > **Administration** > **Targets** and click **Claim a New Target**.

Step 3. Under Select Target Type, select **VMware vCenter** under Hypervisor and click **Start**.

Step 4. In the **VMware vCenter** window, verify the correct Intersight Assist is selected.

Step 5. Fill in the vCenter information. If Intersight Workflow Optimizer (IWO) will be used, turn on Datastore Browsing Enabled and Guest Metrics Enabled. If it is desired to use Hardware Support Manager (HSM) to be able to upgrade IMM server firmware from VMware Lifecycle Manager, turn on HSM. Click **Claim**.

← Targets

Claim a New Target

Claim VMware vCenter 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-vcenter.flexpodb4.cisco.com

Port
443

0 - 65535

Username *
administrator@vsphere.local

Password *
••••••••

🔵 Secure

⚫ Enable Datastore Browsing

⚫ Enable Guest Metrics

⚫ Enable HSM

Back   **Cancel**

Claim

Step 6. After a few minutes, the VMware vCenter will show Connected in the Targets list and will also appear under **Infrastructure Service > Operate > Virtualization**.

Step 7. Detailed information obtained from the vCenter can now be viewed by clicking **Infrastructure Service > Operate > Virtualization** and selecting the Datacenters tab. Other VMware vCenter information can be obtained by navigating through the Virtualization tabs.

Procedure 2. Interact with Virtual Machines

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

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

Step 1. Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2. Select **Infrastructure Service** > **Operate** > **Virtualization**.

Step 3. Click the **Virtual Machines** tab.

Step 4. Click "**...**" to the right of a VM and interact with various VM options.

← Virtualization

Virtual Machines

Virtual Machines   Datacenters   Clusters   Hosts   Virtual Machine Templates   Datastores   Datastore Clusters

✳ All Virtual Machines ⚙ +

··· 🏷 🗑 | 🔍 Add Filter   ↪ **Export**   7 items found   10 ⌄ per page ⎁⎀ 1 of 1 ▷ ▷|

	Provider/Platform	Status	Top 5 Used Instance Types	OS
	7 ● VMware vSphere 7	⏻ Running 7	No data available	7 ● Other 3.x ○ ● Other 2.6.x ● CentOS 4/5

	Name	Pr	Status	CF	CF	CPU ...	M...	IP Address	Place...	⚡
☐	⏻ vCLS-bdb6c736-e13b-4d9	VMw...	⏻ Running	1	3.09 ...	▬ 0.0%	128.00 M	-	FI ⓘ	···
☐	⏻ vCLS-46e4649e-3300-418	VMw...	⏻ Running	1	3.09 ...	▬ 0.0%	128.00 M	-	FI ⓘ	···
☐	⏻ vCLS-3858e77a-646c-414	VMw...	⏻ Running	1	2.19 ...	▬ 0.0%	128.00 M	-	FI ⓘ	···
☐	⏻ aa02-scv	VMw...	⏻ Running	4	12.3...	▬ 0.5%	12.00 GiE	10.102.1.98	FI ⓘ	···
☐	⏻ aa02-ontap-tools	VMw...	⏻ Running	2	6.18 ...	▬ 0.5%	12.00 GiE	10.102.1.9		

Start/Resume

Stop

| ☐ | ⏻ aa02-assist | VMw... | ⏻ Running | 16 | 35.1... | ▬ 8.0% | 32.00 GiE | 10.102.1.9 | Soft Stop |
| ☐ | ⏻ aa02-aiqum | VMw... | ⏻ Running | 4 | 12.3... | ▬ 0.2% | 12.00 GiE | 10.102.1.9 | Suspend |

Reset

Restart

Terminate

Launch VM Console

··· 🏷 🗑

Step 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

Step 1. Clone the repository from https://github.com/NetApp-Automation/NetApp-AIQUM.

Step 2. Follow the instructions in the README file in the repository to ensure the Ansible environment is configured properly.

Step 3. Update the variable files as mentioned in the README document in the repository.

Step 4. To claim an existing AIQUM instance into Intersight, 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

Step 1. Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2. From Cisco Intersight, click **System > Administration > Targets**.

Step 3. Click **Claim a New Target**. In the Select Target Type window, select NetApp Active IQ Unified Manager under Storage and click **Start**.

Step 4. In the Claim NetApp Active IQ Unified Manager Target window, verify the correct Intersight Assist is selected.

Step 5. Fill in the NetApp Active IQ Unified Manager information and click **Claim**.

← Targets

Claim a New Target

Claim NetApp Active IQ Unified Manager 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

ⓘ This target is intended for the functionality of Intersight Orchestrator

Intersight Assist *	Hostname/IP Address *
aa02-assist.flexpodb4.cisco.com	aa02-aiqum.flexpodb4.cisco.com

Username *	Password *
admin	••••••••

🔵 Secure ⓘ

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

Storage

Operate
- Servers
- Chassis
- Fabric Interconnects
- HyperFlex Clusters
- **Storage**

✷ All Storage ⊙ +

🔍 Add Filter   ⤓ **Export**   1 items found   10 ˅ per page   1 of 1

	Name	Vendor	Model	Version	Capacity	Capacity Util...	
☐	aa02-a800	NetApp	AFF-A800	NetApp ONTAP 9...	32.88 TiB	1.1%	...

1 of 1

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

← Storage

aa02-a800

General   Inventory   Checks

Details

Name
aa02-a800

Vendor
NetApp

Model
AFF-A800

Version
NetApp ONTAP 9.11.1P2

Location
Cisco RTP, Building 4, Lab 141, AA02

Management IP
10.102.0.30

DNS Domains
flexpodb4.cisco.com

Name Servers
10.102.1.151
10.102.1.152

NTP Servers
10.102.0.3
10.102.0.4
172.20.10.12

Array Status
OK

Properties

Capacity

Data Reduction	Logical Used
10.17 : 1	**375.40 …**

● Used and Reserved **77.12 GiB**   ● Available **32.50 TiB**

Performance Metrics Summary (Average for 72 hours)

IOPS	Throughput (MiB/s)
366	**7.22**

Array Summary

Nodes	Storage VMs	Local Tiers
2	**1**	**2**

Disks	Ethernet	Fibre Channel
24	**36**	**8**

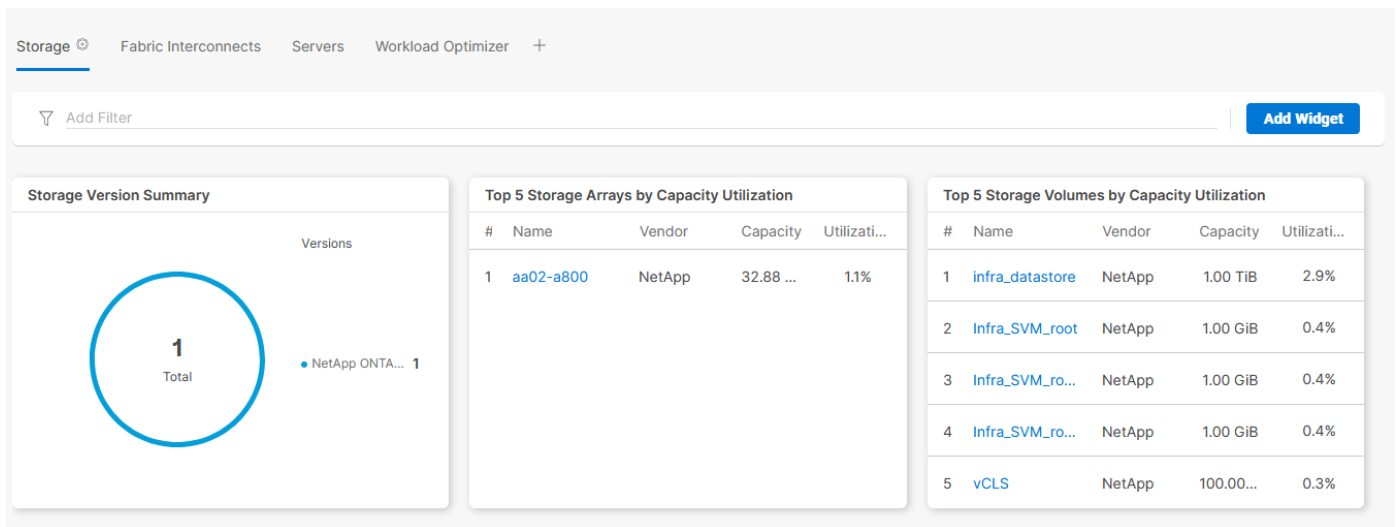**Step 8.**   Click **My Dashboard > Storage** to see storage monitoring widgets.

Claim Cisco Nexus Switches using Cisco Intersight Assist Appliance

Procedure 1. Claim Cisco Nexus Switches

Step 1. Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2. From Cisco Intersight, click **System > Administration > Targets**.

Step 3. Click **Claim a New Target**. In the Select Target Type window, select Cisco Nexus Switch under Network and click **Start**.

Step 4. In the Claim Cisco Nexus Switch Target window, verify the correct Intersight Assist is selected.

Step 5. Fill in the Cisco Nexus Switch information and click **Claim**.

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

Claim a New Target

← Targets

Claim Cisco Nexus 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 *

aa02-93360-a.flexpodb4.cisco.com

Port

443

0 - 65535

Username *

admin

Password *

••••••••

Step 6. Repeat the steps in this procedure to add the second Cisco Nexus Switch.

Step 7. After a few minutes, the two switches will appear under **Infrastructure Service** > **Operate** > **Networking** > **Ethernet Switches**.

Networking

Ethernet Switches   SAN Switches

Step 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)

Step 1. Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2. From Cisco Intersight, click **System > Administration > Targets**.

Step 3. Click **Claim a New Target**. In the Select Target Type window, select Cisco MDS Switch under Network and click **Start**.

Step 4. In the Claim Cisco MDS Switch Target window, verify the correct Intersight Assist is selected.

Step 5. Fill in the Cisco MDS Switch information including use of Port 8443 and click **Claim**.

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

Step 6. Repeat the steps in this procedure to add the second Cisco MDS Switch.

Step 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	Firmware Versions	Models
⊘ Connected 2	2 • 9.2(2) **2**	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 ▷ ▷|

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

Create a FlexPod XCS Integrated System

Procedure 1.   Creating a FlexPod XCS Integrated System

Step 1.   Log into **Cisco Intersight** and connect to the account for this FlexPod.

Step 2.   From Cisco Intersight, click **Infrastructure Service** > **Operate** > **Integrated Systems**.

Step 3.   Click **Create Integrated System**. In the center pane, select **FlexPod** and click **Start**.

Step 4.   Select the correct Organization (for example, AA02), provide a suitable name, and optionally any Tags or a Description and click **Next**.

Step 5. Select the UCS Domain used in this FlexPod and click **Next**.

Step 6. Select the two Cisco Nexus switches used in this FlexPod and click **Next**.

Create Integrated System

Network Switch Selection

Select HA pair of Nexus Switches

		^ Ethernet Switches

2 items found   10 ∨ per page |◁ ◁ 1 of 1 ▷ ▷| ⚙

🔍 Add Filter

☑	Name	⇅	Manag...	⇅	Model	⇅	Firmwa...	⇅	⚡
☑	aa02-93360-a		10.102.0.3		N9K-C9336...		10.2(3)		...
☑	aa02-93360-b		10.102.0.4		N9K-C9336...		10.2(3)		...

Selected 2 of 2   **Show Selected**   **Unselect All** |◁ ◁ 1 of 1 ▷ ▷|

Step 7.   Select all NetApp storage used in this FlexPod and click **Next**.

Create Integrated System

Storage Array Selection

Select one or more Storage Arrays

1 items found   10 ∨ per page |◁ ◁ 1 of 1 ▷ ▷| ⚙

🔍 Add Filter

☑	Name	⇅	Vendor	⇅	Version	⇅	Capacity	⇅
☑	aa02-a800		NetApp		NetApp ONTA...		32.88 TiB	

Selected 1 of 1   **Show Selected**   **Unselect All** |◁ ◁ 1 of 1 ▷ ▷|

Step 8.   Review the Summary information and click **Create**. After a few minutes, the FlexPod Integrated System will appear under Integrated Systems.

Integrated Systems

Note: You can click the "**...**" to the right of the FlexPod name and run an Interoperability check on the FlexPod. This check will take information on the FlexPod already checked against the Cisco UCS Hardware Compatibility List (HCL) and also check this information against the NetApp Interoperability Matrix Tool (IMT).

Step 9.   Click on the FlexPod name to see detailed General, Inventory, and Interoperability data on the FlexPod XCS Integrated System.

Step 10.   Select **My Dashboard** > **FlexPod** to see several informational widgets on FlexPod Integrated Systems.
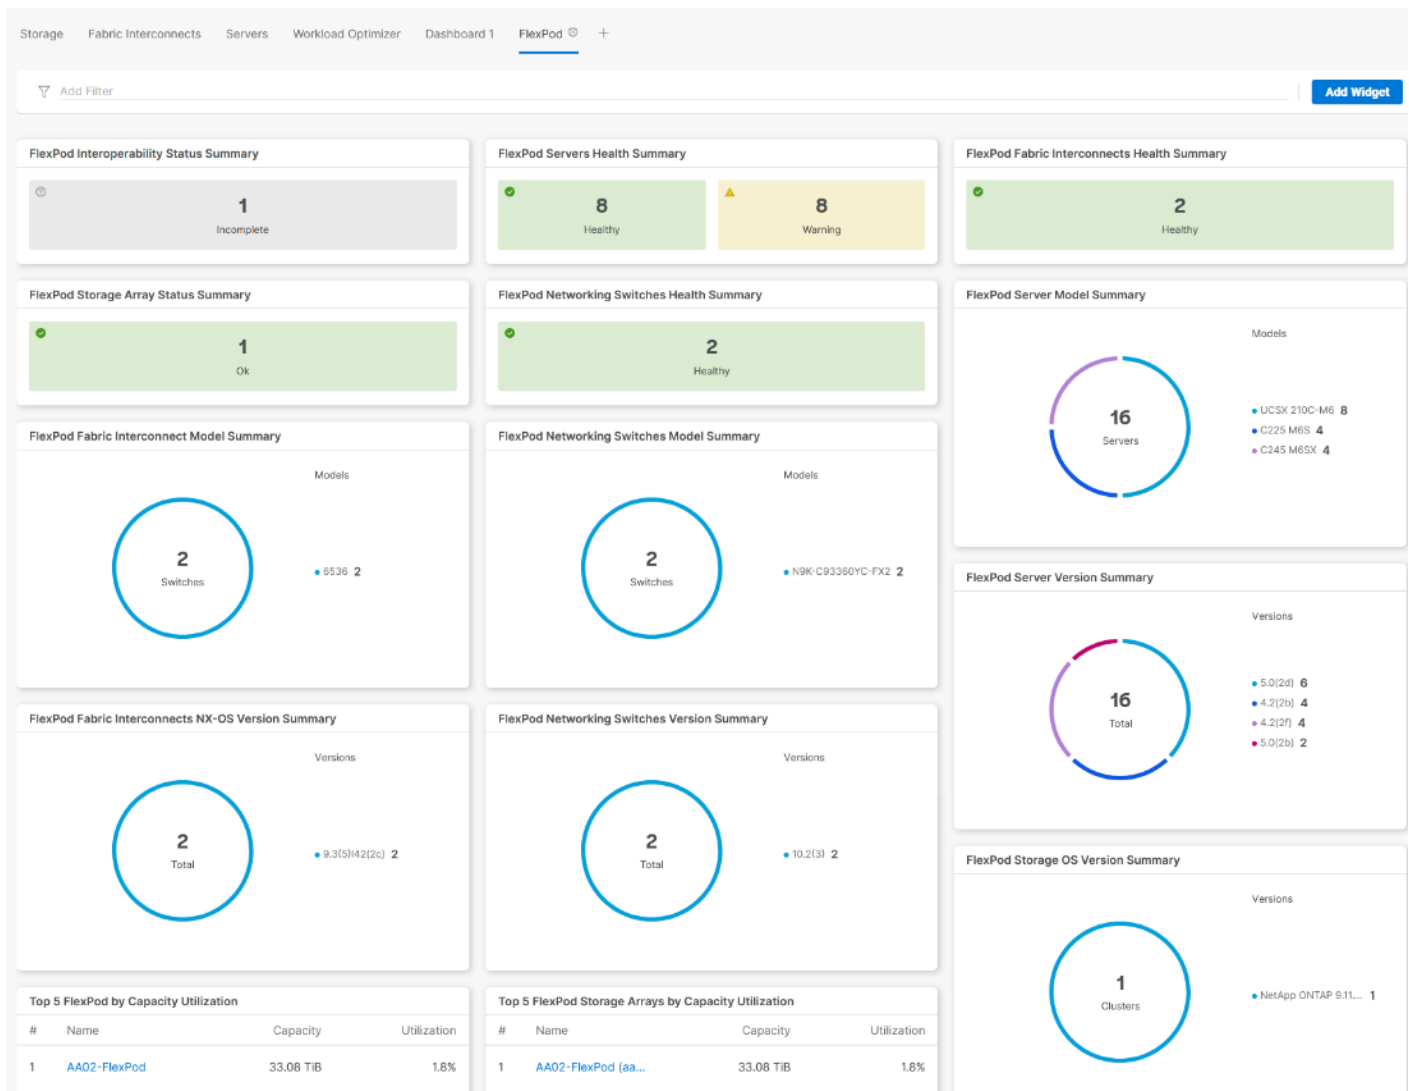
Cisco Data Center Network Manager (DCNM)–SAN

If you have fibre-channel SAN in your FlexPod, Cisco DCNM-SAN can be used to monitor, configure, and analyze Cisco fibre channel fabrics. Cisco DCNM-SAN is deployed as a virtual appliance from an OVA and is managed through a web browser. SAN Analytics can be added to provide insights into your fabric by allowing you to monitor, analyze, identify, and troubleshoot performance issues.

Prerequisites

The following prerequisites need to be configured:

- Licensing. Cisco DCNM-SAN includes a 60-day server-based trial license that can be used to monitor and configure Cisco MDS Fibre Channel switches and monitor Cisco Nexus switches. Both DCNM server-based and switch-based licenses can be purchased. Additionally, SAN Insights and SAN Analytics requires an additional switch-based license on each switch. Cisco MDS 32Gbps Fibre Channel switches provide a 120-day grace period to trial SAN Analytics.

Note: If using Cisco Nexus 93180YC-FX, 93360YC-FX2, or 9336C-FX2-E for SAN switching, the Nexus switch does not support SAN Analytics.

- Passwords. Cisco DCNM-SAN passwords should adhere to the following password requirements:
  - It must be at least eight characters long and contain at least one alphabet and one numeral.
  - It can contain a combination of alphabets, numerals, and special characters.
  - Do not use any of these special characters in the DCNM password for all platforms: <SPACE> " & $ % ' ^ = < > ; : ` \ | / , .*

- DCNM SNMPv3 user on switches. Each switch (both Cisco MDS and Nexus) needs an SNMPv3 user added for DCNM to use to query and configure the switch. On each switch, enter the following command in configure terminal mode (in the example, the userid is snmpuser):

```
snmp-server user snmpadmin network-admin auth sha <password> priv aes-128 <privacy-password>
```

- On Cisco MDS switches, type show run. If snmpadmin passphrase lifetime 0 is present, enter username snm-padmin passphrase lifetime 99999 warntime 14 gracetime 3.

Note: It is important to use auth type sha and privacy auth aes-128 for both the switch and UCS snmpadmin users.

- Type "copy run start" on all switches to save the running configuration to the startup configuration.

- An SNMP Policy was added to the UCS Domain Profile in IMM to create the snmpadmin user there.

Procedure 1. Deploy the Cisco DCNM-SAN OVA

Step 1. Download the Cisco DCNM 11.5(4). Open Virtual Appliance for VMware from https://software.cisco.com/download/home/281722751/type/282088134/release/11.5(4). Extract dcnm-va.11.5.4.ova from the ZIP file.

Step 2. In the VMware vCenter HTML5 interface, select **Inventory** > **Hosts and Clusters**.

Step 3. Right-click the FlexPod-Management cluster and select **Deploy OVF Template**.

Step 4. Select Local file then click **UPLOAD FILES**. Navigate to select dcnm-va.11.5.4.ova and click **Open**. Click **NEXT**.

Deploy OVF Template

1 **Select an OVF template**

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template   ✕

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

    http | https://remoteserver-address/filetodeploy.ovf | .ova

● Local file

    [UPLOAD FILES]   dcnm-va.11.5.4.ova

                                                              CANCEL   **NEXT**

Step 5.   Name the virtual machine and select the FlexPod-DC datacenter. Click **NEXT**.

Step 6.   Select the FlexPod-Management cluster and click **NEXT**.

Step 7.   Review the details and click **NEXT**.

Step 8.   Scroll through and accept the license agreements. Click **NEXT**.

Step 9.   Select the appropriate deployment configuration size and click **NEXT**.

Note: If using the SAN Insights and SAN Analytics feature, it is recommended to use the Huge size.

Step 10. Select infra_datastore and the Thin Provision virtual disk format. Click **NEXT**.

Step 11. Select IB-MGMT Network for the first and third Source Networks. Select OOB-MGMT Network for the second enhanced-fabric-mgmt Source Network. Click **NEXT**.

Deploy OVF Template

Select networks   ✕

Select a destination network for each source network.

1. Select an OVF template
2. Select a name and folder
3. Select a compute resource
4. Review details
5. License agreements
6. Configuration
7. Select storage
8. **Select networks**
9. Customize template
10. Ready to complete

Source Network	Destination Network
dcnm-mgmt	IB-MGMT Network ⌄
enhanced-fabric-mgmt	OOB-MGMT Network ⌄
enhanced-fabric-inband	IB-MGMT Network ⌄

3 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL   BACK   NEXT

Step 12. Fill in the management IP address, subnet mask, and gateway. Set the Extra Disk Size according to how many Cisco MDS switches you will be monitoring with this DCNM. If you are only monitoring the two Cisco MDS switches in this FlexPod deployment, set this field to 32. Click **NEXT**.

Step 13. Review the settings and click **FINISH** to deploy the OVA.

Deploy OVF Template

1. Select an OVF template
2. Select a name and folder
3. Select a compute resource
4. Review details
5. License agreements
6. Configuration
7. Select storage
8. Select networks
9. Customize template
10. **Ready to complete**

Ready to complete

Review your selections before finishing the wizard

∨ **Select a name and folder**

Name	aa02-dcnm
Template name	dcnm
Folder	FlexPod-DC

∨ **Select a compute resource**

Resource	FlexPod-Management

∨ **Review details**

Download size	5.4 GB

∨ **Select storage**

Size on disk	9.4 GB
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thin provision

∨ **Select networks**

Network mapping	3
dcnm-mgmt	IB-MGMT Network
enhanced-fabric-mgmt	OOB-MGMT Network
enhanced-fabric-inband	IB-MGMT Network

CANCEL   BACK   FINISH

Step 14.　　　After deployment is complete, right-click the newly deployed DCNM VM and click **Edit Settings**. Expand CPU and adjust the Cores per Socket setting until the number of Sockets is set to match the number of CPUs in the UCS servers used in this deployment. The following example shows 2 sockets. Click **OK**.

Edit Settings | aa02-dcnm

Virtual Hardware   VM Options

ADD NEW DEVICE ˅

CPU *	32 ˅	ⓘ
Cores per Socket	16 ˅	
	Sockets: 2	
CPU Hot Plug	☐ Enable CPU Hot Add	
Reservation	0 ˅	MHz ˅
Limit	Unlimited ˅	MHz ˅
Shares	Normal ˅ 32000 ˅	
Hardware virtualization	☐ Expose hardware assisted virtualization to the guest OS	
Performance Counters	☐ Enable virtualized CPU performance counters	
CPU/MMU Virtualization	Automatic ˅	ⓘ
> Memory	128 ˅	GB ˅

Step 15. Right-click the newly deployed DCNM VM and click **Open Remote Console**. Once the console is up, click the green arrow to power on the VM. Once the VM has powered up, point a web browser to the URL displayed on the console.

Step 16. Navigate the security prompts and click **Get started**.

Step 17. Make sure Fresh installation – Standalone is selected and click **Continue**.

Step 18. Select SAN only for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click **Next**.

Step 19. Enter and repeat the administrator, database, and root passwords and click **Next**.

Step 20. Enter the DCNM FQDN, a comma-separated list of DNS servers, a comma-separated list of NTP servers, and select the appropriate time zone. Click **Next**.

Cisco DCNM Installer

Install Mode | Administration | **System Settings** | Network Settings | Applications | HA Settings | Summary

Please enter the following system settings

Fully Qualified Host Name *
Fully Qualified Host Name as per RFC1123, section 2.1, for example:
myhost.mydomain.com. Digit-only host names are not allowed.

```
aa02-dcnm.flexpodb4.cisco.com
```

DNS Server Address List *
Comma-separated list of DNS Server addresses (IPv4 or IPv6)

```
10.102.1.151,10.102.1.152
```

NTP Server Address List *
Comma-separated list of NTP Server addresses (RFC1123-compliant name, IPv4 or IPv6)

```
10.102.1.3,10.102.1.4
```

Timezone *

```
America/New_York
```

Previous   Next

Step 21.   The Management Network settings should be filled in. For Out-of-Band Network, enter an IP address in the Out-of-Band management subnet. For the Out-of-Band Network, only input the IPV4 address with prefix. Do not put in the Gateway IPv4 Address. Do not enter any information for the In-Band Network. Scroll down and click **Next**.

Step 22.   If necessary, enter data for the Device connector configuration. Leave Internal Application Services Network set at the default setting and click **Next**.

Step 23.   Review the Summary details and click **Start installation**.

Step 24.   When the Installation status is complete, click **Continue**.

Step 25.   In the vCenter HTML5 client under Hosts and Clusters, select the DCNM VM and click the Summary tab. If an alert is present that states "A newer version of VMware Tools is available for this virtual machine.," click **Upgrade VMware Tools**. Select Automatic Upgrade and click **UPGRADE**. Wait for the VMware Tools upgrade to complete.

Procedure 2.   Configure DCNM-SAN

Step 1.   When the DCNM installation is complete, the browser should redirect to the DCNM management URL.

Step 2.   Log in as admin with the password previously entered.

Step 3.   On the message that appears, select **Do not show this message again** and click **No**.

Step 4.   If you have purchased DCNM server-based or switch-based licenses, follow the instructions that came with the licenses to install them. A new DCNM installation also has a 60-day trial license.

Step 5.   In the menu on the left, click **Inventory** > **Discovery** > **LAN Switches**.

Step 6.   Click ➕ to add LAN switches. In the Add LAN Devices window, enter the mgmt0 IP address of the Nexus switch A in the Seed Switch box. Enter the snmpadmin user name and password set up in the Pre-requisites section above. Set Auth-Privacy to SHA_AES. Click **Next**.

Add LAN Devices

Discovery Type:	◉ Hops from seed switch ○ Switch list
Seed Switch:	10.102.0.3
Max Hops from Seed:	0 1 2 3 4 5 6 7
User Name:	snmpadmin
Password:	••••••••
Auth-Privacy:	SHA_AES ▼
Add Switches To Group:	Default_LAN ▼
Scan Time:	3 secs ▼

Next   Cancel

Step 7.   LAN switch discovery will take a few minutes. In the LAN Discovery list that appears, the two Nexus switches and two Fabric Interconnects that are part of this FlexPod should appear with a status of "man-ageable." Using the checkboxes on the left, select the two Nexus switches and two Fabric Interconnects that are part of this FlexPod. Click **Add**.

Step 8.   After a few minutes, click the **Refresh** icon in the upper right-hand corner, and detailed infor-mation abut the two Nexus switches and two Fabric Interconnects that are part of this FlexPod will display.

		Switch	IP Address	Serial No	Managed	SNMP Status	Role
1	☐	aa02-6536-A	10.102.0.18		true	SSH: Failed to…	
2	☐	aa02-6536-B	10.102.0.19		true	SSH: Failed to…	
3	☐	aa02-93360-a	10.102.0.3		true	ok	
4	☐	aa02-93360-b	10.102.0.4		true	ok	

Step 9. In the menu on the left, click **Inventory** > **Discovery** > **SAN Switches**.

Step 10. Click ⊞ to add a switching fabric.

Step 11. Enter either the IP address or hostname of the first Cisco MDS 9132T switch. Leave Use SNMPv3/SSH selected. Set Auth-Privacy to **SHA_AES**. Enter the snmpadmin user name and password set up in the Prerequisites section. Click **Options>>**. Enter the UCS admin user name and password. Click **Add**.

Note: If Cisco Nexus 93180YC-FX, 93360YC-FX2, or 9336C-FX2-E switches are being used for SAN switching, substitute them for MDS 9132Ts. They will need to be added again under SAN switches since LAN and SAN switching are handled separately in DCNM.

Step 12. Repeat steps 9-11 to add the second Cisco MDS 9132T and Fabric Interconnect.

The two SAN fabrics now appear in the Inventory.

Selected 0 / Total 2

+	✕	✎	⤴ Move	↻ Rediscover	☒ Purge			

	Name	SeedSwitch	Status	SNMPv3/SSH	User/Cmnty
☐	Fabric_aa02-9132t-a	10.102.0.7	managedContinuously	true	snmpadmin
☐	Fabric_aa02-9132t-b	10.102.0.8	managedContinuously	true	snmpadmin

Step 13. Select **Inventory** > **Discovery** > **Virtual Machine Manager**.

Step 14. Click ➕ to add the vCenter.

Step 15. In the Add VCenter window, enter the IP address of the vCenter VCSA. Enter the administrator@vsphere.local user name and password. Click **Add**. The vCenter should now appear in the inventory.

Step 16. Select **Inventory** > **Switches**. All LAN and SAN switches should now appear in the inventory.

Step 17. Select **Administration** > **Performance Setup** > **LAN Collections**.

Step 18. Select the Default_LAN group and all information you would like to collect. Click **Apply**. Click **Yes** to restart the Performance Collector.

🏠 | **Administration / Performance Setup / LAN Collections**

For all selected licensed LAN Switches collect: ☑ Trunks ☑ Access ☑ Errors & Discards ☑ Temperature Sensor   **Apply**

▼ ☑ 🔅 Default_LAN
   ☑ 🖥 aa02-93360-a
   ☑ 🖥 aa02-93360-b

Step 19. Select **Administration** > **Performance Setup** > **SAN Collections**.

Step 20. Select both fabrics. Select all information you would like to collect and click **Apply**. Click **Yes** to restart the Performance Collector.

🏠 | Administration / Performance Setup / SAN Collections

Total 2 ↻

		Name	ISL/NPV Links	Hosts	Storage	FC Flows	FC Ethernet
1	☑	Fabric_aa02-9132t-a	☑	☑	☑	☐	☐
2	☑	Fabric_aa02-9132t-b	☑	☑	☑	☐	☐

Apply

Step 21. Select **Configure** > **SAN** > **Device Alias**. Since device-alias mode enhanced was configured in the Cisco MDS 9132T switches, Device Aliases can be created and deleted from DCNM and pushed to the MDS switches.

Step 22. Select **Configure** > **SAN** > **Zoning**. Just as Device Aliases can be created and deleted from DCNM, zones can be created, deleted, and modified in DCNM and pushed to the MDS switches. Make sure to enable Smart Zoning and to Zone by Device Alias.

You can now explore all of the different options and information provided by DCNM SAN. See [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5(x)](#).

Configure SAN Insights in DCNM SAN

The SAN Insights feature enables you to configure, monitor, and view the flow analytics in fabrics. Cisco DCNM enables you to visually see health-related indicators in the interface so that you can quickly identify issues in fabrics. Also, the health indicators enable you to understand the problems in fabrics. The SAN Insights feature also pro-vides more comprehensive end-to-end flow-based data from host to LUN.

- Ensure that the time configurations set above, including daylight savings settings are consistent across the MDS switches and Cisco DCNM.

- SAN Insights requires installation of a SMART SAN Analytics license on each switch. To trial the feature, each switch includes a one-time 120-day grace period for SAN Analytics from the time the feature is first enabled.

- SAN Insights supports current Fibre Channel Protocol (SCSI) and NVMe over Fibre Channel (NVMe).

- SAN Insights works by enabling SAN Analytics and Telemetry Streaming on each switch. The switches then stream the SAN Analytics data to DCNM, which collects, correlates, and displays statistics. All configurations can be done from DCNM.

- Only Cisco MDS switches support SAN Analytics. Cisco Nexus switches do not support SAN Analytics.

- For more information on SAN Insights, see the [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5(x)](#).

- For more information on SAN Analytics, see https://www.cisco.com/c/en/us/td/docs/dcn/mds9000/sw/9x/configuration/san-analytics/cisco-mds-9000-san-analytics-telemetry-streaming-configuration-guide-9x.html.

Procedure 1.	Configure SAN Insights in DCNM SAN

Step 1. Click **Configure** > **SAN** > **SAN Insights**. Click **Continue**.

Step 2. Select **Fabric A**. Click **Continue**.

Step 3. Select the **Fabric A Cisco MDS switch**. Under Install Query click **None** and from the drop-down list click **Storage**. Under Subscriptions, select **SCSI & NVMe** or whatever you have currently installed. Optionally, under Receiver, select the IP address in the Out-of-Band Management subnet configured for DCNM. Click **Save**, then click **Continue**.

2. Select Switches

Choose the switch(es) on which SAN Insights is to be configured in Fabric_aa02-9132t-a

DCNM server time: 12:47:14.026 EDT Thursday October 27 2022

Selected 1 / Total 1

	Switch		Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Interval	Receiver
☑	aa02-9132t-a	ⓘ	DS-C9132T-K9	9.2(2)	Yes	12:47:15.568 EDT Thu Oct 27 2022	SCSI & NVMe	Storage	30	10.102.0.9

Disable Analytics...   Show Quick Filter ▼ 🔽

Step 4.   Review the information and click **Continue**.

Step 5.   Expand the switch and then the module. Under Enable / Disable SCSI Telemetry, click the left icon to enable telemetry on the ports connected to the NetApp AFF A800. Under Enable / Disable NVMe Telemetry, click the left icon to enable telemetry on the ports connected to the NetApp AFF A800. Click **Continue**.

4. Select Interfaces

Choose the switch interfaces that will generate analytics data within Fabric_aa02-9132t-a

Total Top Level Rows 1

Switch	Module	Interface	Connected To	Type	Analytics Status	Enable / Disable SCSI Telemetry	Enable / Disable NVMe Telemetry
▼ aa02-9132t-a	1 module(s)	6 interface(s)		Storage			
▼	DS-C9132T-K9-S...	6 interface(s)					
		fc1/1	AA02-A400-Infra-SVM...	Storage	disabled	■ ■	■ ■
		fc1/2	AA02-A400-Infra-SVM...	Storage	disabled	■ ■	■ ■
		fc1/9	50:0a:09:81:80:71:50:9c	Storage	disabled	■ ■ pending enable	■ ■ pending enable
		fc1/10	50:0a:09:83:80:71:50:9c	Storage	disabled	■ ■ pending enable	■ ■ pending enable
		fc1/11	50:0a:09:81:80:41:50:95	Storage	disabled	■ ■ pending enable	■ ■ pending enable
		fc1/12	50:0a:09:83:80:41:50:95	Storage	disabled	■ ■ pending enable	■ ■ pending enable

Step 6.   Review the information and click **Commit** to push the configuration to the Cisco MDS switch.

Step 7.   Ensure that the two operations were successful and click **Close**.

Step 8.   Repeat steps 1 – 7 to install SAN Analytics and Telemetry on the Fabric B switch.

Note: After approximately two hours, you can view SAN Analytics data under the Dashboard and Monitor.

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.

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

Roney Daniel is a Technical Marketing engineer at NetApp. He has over 25 years of experience in the networking industry. Prior to NetApp, Roney worked at Cisco Systems in various roles with Cisco TAC, Financial Test Lab, Systems and solution engineering BUs and Cisco IT. He has a bachelor's degree in Electronics and Communication engineering and is a data center Cisco Certified Internetwork Expert (CCIE 42731).

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

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

Acknowledgements

Appendix

This appendix is organized into the following:

- [FlexPod with Cisco Nexus SAN Switching Configuration – Part 1](#)
- [FlexPod with Cisco Nexus 93360YC-FX2 SAN Switching Configuration – Part 2](#)
- [Create a FlexPod ESXi Custom ISO using VMware vCenter](#)
- [Active IQ Unified Manager User Configuration](#)
- [Active IQ Unified Manager vCenter Configuration](#)
- [NetApp Active IQ](#)
- [FlexPod Backups](#)
- [Glossary of Acronyms](#)
- [Glossary of Terms](#)

Note: The features and functionality explained in this Appendix are optional configurations which can be helpful in configuring and managing the FlexPod deployment.

FlexPod with Cisco Nexus SAN Switching Configuration – Part 1

If the Cisco Nexus switches are to be used for both LAN and SAN switching in the FlexPod configuration, either an automated configuration with Ansible or a manual configuration can be done. For either configuration method, the following base switch setup must be done manually. Figure 6 shows the validation lab cabling for this setup.

Figure 6. Cisco Nexus SAN Switching Cabling with FCoE Fabric Interconnect Uplinks
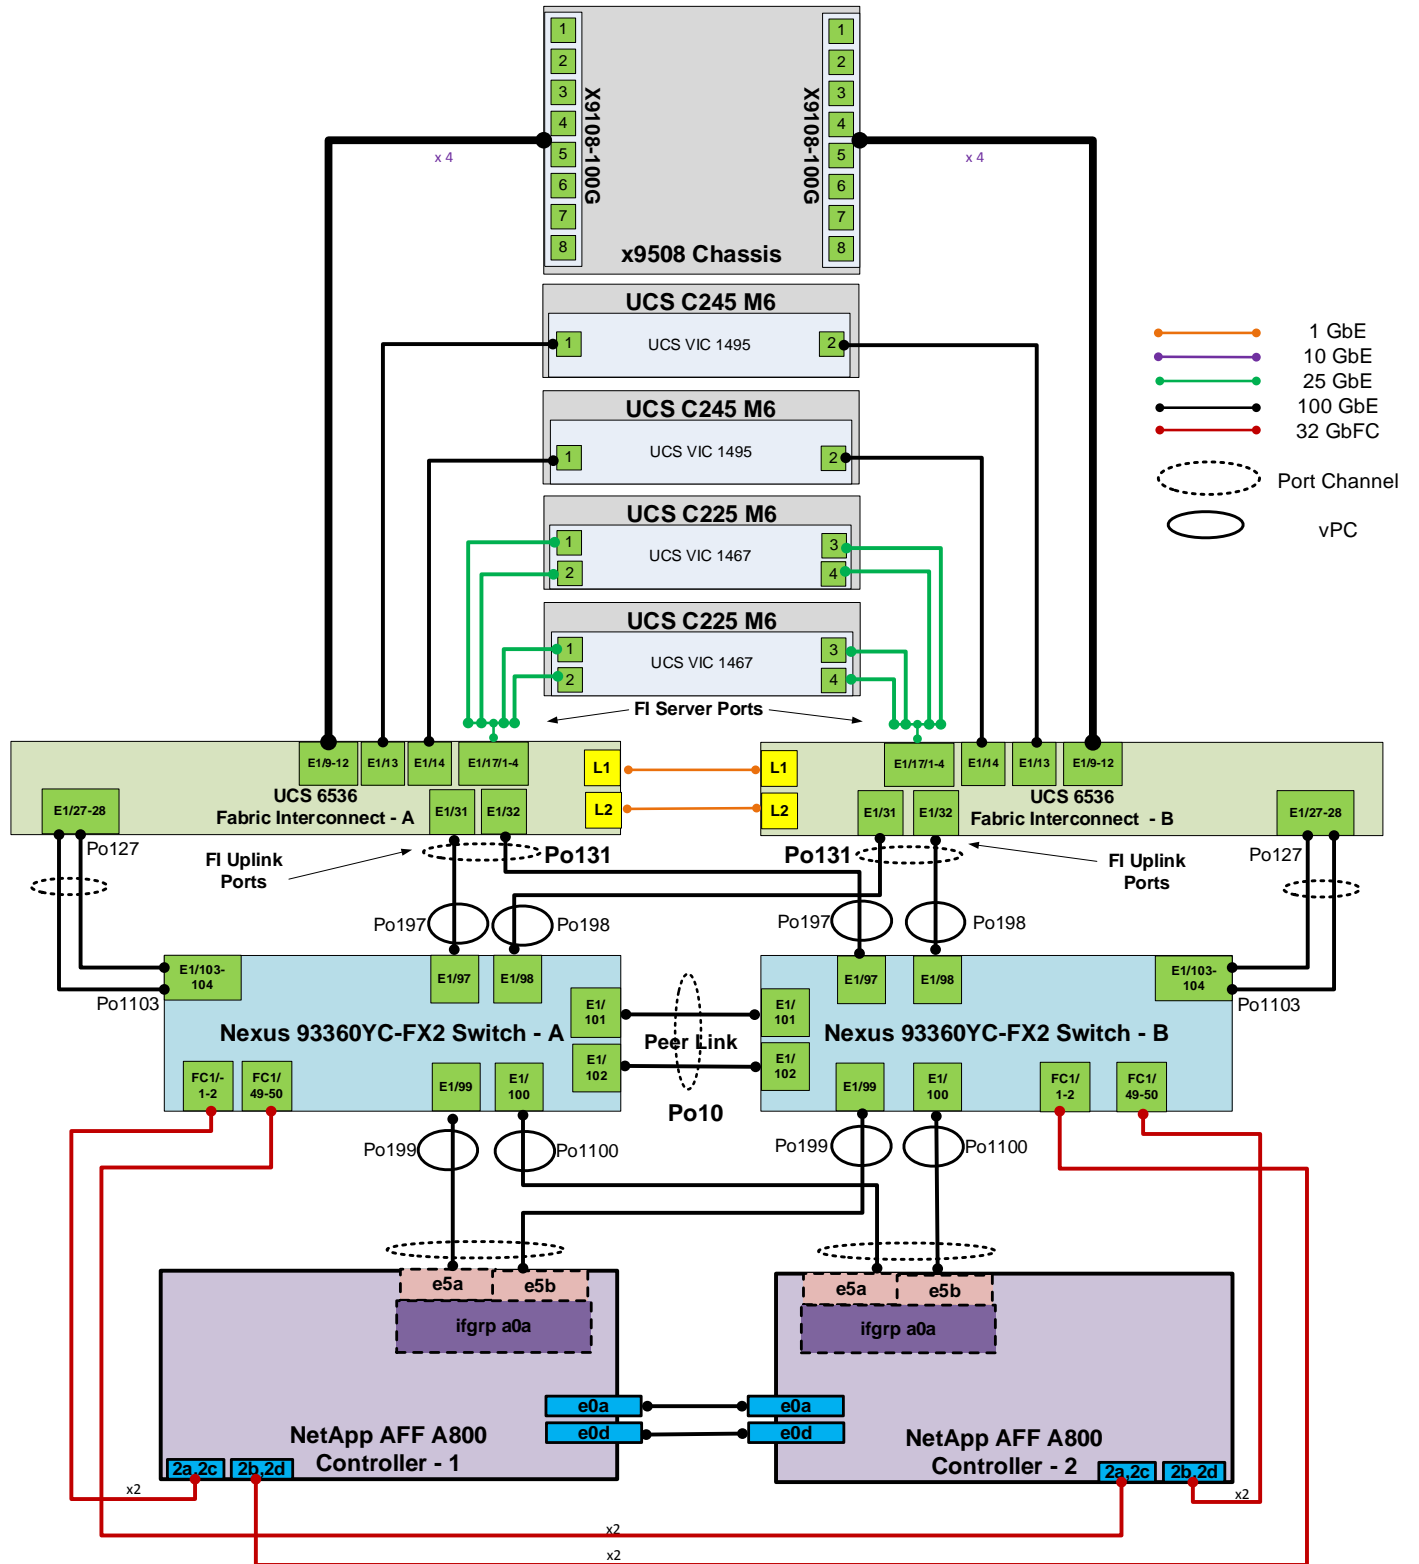
FlexPod Cisco Nexus 93180YC-FX SAN Switching Base Configuration

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're using Cisco Nexus 9000 10.2(3)F. This procedure also assumes that you have created an FCoE Uplink Port Channel on the appropriate ports in the Cisco UCS IMM Port Policies for each UCS fabric interconnect.

Procedure 1.   Set Up Initial Configuration in Cisco Nexus 93360YC-FX2 A

Step 1.   Configure the switch:

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should auto-matically 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
```

Step 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 93360YC-FX2 B

Step 1.   Configure the switch:

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should auto-matically 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
```

Step 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 Nexus switch.

Note: This section is structured as a green field switch setup. If existing switches that are switching active traffic are being setup, execute this procedure down through Perform TCAM Carving and Configure Unified Ports in Cisco Nexus 93360YC-FX2 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 93360YC-FX2 A and B |

Step 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 FC configurations were not configured in the switch setup script de-tailed in the previous section.

| **Procedure 4.**   Set System-Wide QoS Configurations in Cisco Nexus 93360YC-FX2 A and B |

Step 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
```

Note: SAN switching requires TCAM carving for lossless fibre channel no-drop support. Also, unified ports need to be converted to fc ports.

Note: On the Cisco Nexus 93360YC-FX2, UP ports are converted to FC in groups of 4 in columns, for example, 1,2,49,50.

Step 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
```

Step 2. After the switch reboots, log back in as admin. 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 93360YC-FX2 SAN Switching Ethernet Switching Automated Configuration

For the automated configuration of the Ethernet part of the Cisco Nexus 93360YC-FX2 switches when using the switches for SAN switching, once the base configuration is set, return to Ansible Nexus Switch Configuration, and execute from there.

FlexPod with Cisco Nexus 93360YC-FX2 SAN Switching Configuration - Part 2

Note: If the Cisco Nexus 93360YC-FX2 switch is being used for SAN Switching, this section should be completed in place of the Cisco MDS section of this document.

Automate the configuration of the SAN part of the Cisco Nexus 93180YC-FX switches when using the switches for SAN switching.

Step 1. Verify Nexus switch ssh keys are in /root/.ssh/known_hosts. Adjust known_hosts as necessary if errors occur.

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

Step 2. Edit the /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/inventory file putting the Cisco Nexus A information in for MDS A and the Cisco Nexus B information in for MDS B.

Step 3. Edit the following variable files to ensure proper Cisco Nexus SAN variables are entered:

- /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/group_vars/all.yml

- /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/mdsA.yml

- /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/host_vars/mdsB.yml

- /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2/roles/NEXUSSANconfig/defaults/main.yml

Note: The SAN variables and port descriptions from the mdsA.yml and mdsB.yml files will be used for the SAN configuration in the Cisco Nexus 93360YC-FX2 switches.

Step 4.   From /root/FlexPod-IMM-4.2.2/FlexPod-IMM-4.2.2, run the Setup_NexusSAN.yml Ansible playbook.

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

Procedure 2.   Switch Testing Commands

The following commands can be used to check for correct switch configuration:

Note: Some of these commands need to run after further configuration of the FlexPod components are complete to see complete results.

```
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 FlexPod ESXi Custom ISO using VMware vCenter

In this Cisco Validated Design (CVD), 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. vCenter 7.0U3 or later can be used to produce a FlexPod custom ISO containing the updated UCS VIC fnic driver, the lsi_mr3 driver, and the NetApp NFS Plug-in for VMware VAAI. This ISO can be used 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

Step 1.   Download the [Cisco Custom Image for ESXi 7.0 U3 Offline Bundle](#). This file (VMware-ESXi-7.0.3d-19482537-Custom-Cisco-4.2.2-a-depot.zip) can be used to produce the FlexPod ESXi 7.0U3 CD ISO.

Step 2.   Download the following listed .zip files:

- [VMware ESXi 7.0 nfnic 5.0.0.34 Driver for Cisco VIC Adapters](#) – Cis-co-nfnic_5.0.0.34-1OEM.700.1.0.15843807_19966277.zip – extracted from the downloaded zip

- [VMware ESXi 7.0 lsi_mr3 7.720.04.00-1OEM SAS Driver for Broadcom Megaraid 12Gbps](#) – Broad-com-lsi-mr3_7.720.04.00-1OEM.700.1.0.15843807_19476191.zip – extracted from the downloaded zip

- [NetApp NFS Plug-in for VMware VAAI 2.0](#) – NetAppNasPluginV2.0.zip

- The Cisco VIC nenic driver would 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.

Step 3.   Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.

Step 4.   Under the Menu at the top, select **Auto Deploy**.

Step 5.   If you see the following, select **ENABLE IMAGE BUILDER**.

⚠️

Auto Deploy and Image Builder are disabled in this vCenter.

To access full-featured auto deploy, enable both Image Builder and Auto Deploy.

To manage software depots only, enable Image Builder.

ENABLE AUTO DEPLOY AND IMAGE BUILDER

ENABLE IMAGE BUILDER

Step 6.   Click **IMPORT** to upload a software depot.

Step 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 above, highlight it, and click **Open**.

Import Software Depot   ✕

| Name * | Cisco Custom ESXi 7.0U3 |
| File * | VMware-ESXi-7.0.3d-19482537-Custom-   BROWSE |

CANCEL   UPLOAD

Step 8. Click **UPLOAD** to upload the software depot.

Step 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.

Step 10. Click **NEW** to add a custom software depot.

Step 11. Select **Custom depot** and name the custom depot FlexPod-ESXi-7.0U3.

Add Software Depot ✕

○ Online depot

    Name: _____

    URL: _____

● Custom depot

    Name: *   FlexPod-ESXi-7.0U3

[CANCEL] [**ADD**]

Step 12. Click **ADD** to add the custom software depot.

Step 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.

Step 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.

Name and details

Name *	FlexPod-ESXi-7.0U3
Vendor *	Cisco-NetApp
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 ⌄ ⓘ

Step 15. Click **NEXT**.

Step 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

Step 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

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

Step 19. Using the Software Depot pulldown, 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.

Step 20. Once the Image profile export completes, click **DOWNLOAD** to download the ISO.

Step 21. Once downloaded, you can rename the ISO to a more descriptive name (for example, FlexPod-ESXi-7.0U3.iso).

Step 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 Local Users to Active IQ Unified Manager

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

Step 2. Click **+ Add** and complete the requested information:

    a. Select Local User for the Type.

    b. Enter a username and password.

    c. Add the user's email address.

    d. Select the appropriate role for the new user.

Users: Add ⓘ

TYPE

Local User   ⌄

⚠ Authentication server is either disabled or not configured. To add a remote user or
group, enable or configure the authentication server from Setup Options.

NAME

flexadmin

PASSWORD

••••••••

CONFIRM PASSWORD

••••••••

EMAIL

flexadmin@cspg.local

ROLE

Storage Administrator   ⌄

Step 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 Ap-
plication Administrator privileges to configure remote authentication.

Step 1.   Navigate to the **General** and select **Remote Authentication**.

Step 2.   Select the option to enable Remote Authentication and define a remote user or remote group.
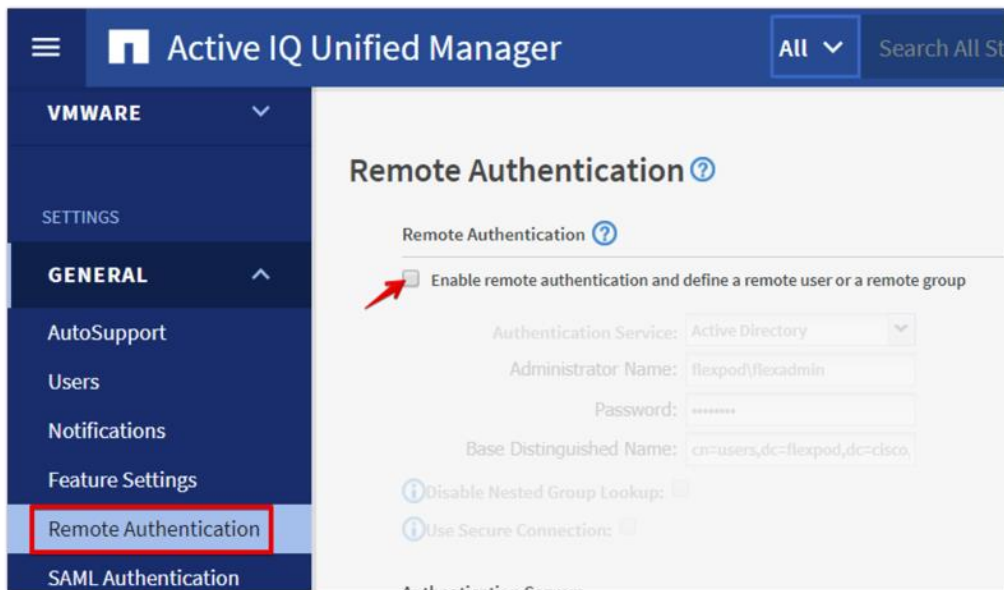
Step 3. Select **Active Directory** from the authentication service list.

Step 4. Enter the Active Directory service account name and password. The account name can be in the format of domain\user or user@domain.

Step 5. Enter the base DN where your Active Directory users reside.

Step 6. If Active Directory LDAP communications are protected via SSL enable the **Use Secure Connection** option.

Step 7. Add one or more Active Directory domain controllers by clicking **Add** and entering the IP or FQDN of the domain controller.

Step 8. Click **Save** to enable the configuration.

Step 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 displays indicating authentication was successful:

Result

Authentication succeeded.
Username: flexadmin
Full Name: CN=FlexPod
Admin,cn=users,dc=flexpod,dc=cisco,dc=com
Groups: [Domain Admins, Denied RODC Password
Replication Group]

Procedure 3. Add a Remote User to Active IQ Unified Manager

Step 1. Navigate to the **General** section and select **Users**.

Step 2. Click **Add** and select **Remote User** from the Type drop-down list.

Step 3. Enter the following information into the form:

a. The username of the Active Directory user.

b. Email address of the user.

c. Select the appropriate role for the user.

NAME

PASSWORD

CONFIRM PASSWORD

EMAIL

ROLE

Application Administrator

Save   Cancel

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

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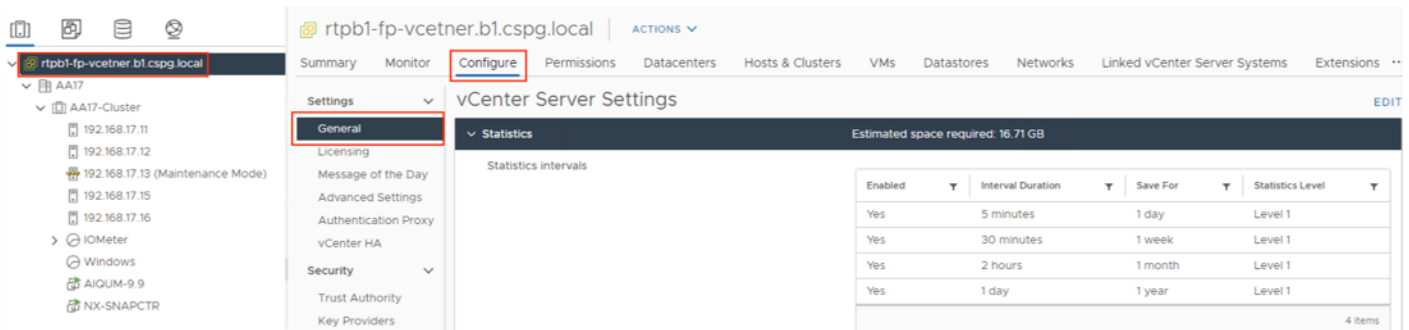
Active IQ Unified Manager provides visibility into vCenter and the virtual machines running inside the datastores backed by ONTAP storage. Virtual machines and storage are monitored to enable quick identification of performance issues within the various components of the virtual infrastructure stack.

Note: Before adding vCenter into Active IQ Unified Manager, the log level of the vCenter server must be changed.

Procedure 1. Configure Active IQ Unified Manager vCenter

Step 1. In the vSphere client navigate to **Menu > VMs and Templates** and select the vCenter instance from the top of the object tree.

Step 2. Click the **Configure** tab, expand **Settings**, and select **General**.

Step 3. Click **EDIT**.

Step 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.

Step 5. Click **SAVE**.

Step 6. Switch to the Active IQ Unified Manager and navigate to the **VMware** section located under **In-ventory**.

Step 7. Expand VMware and select **vCenter**.

Step 8.   Click **Add**.

Step 9.   Enter the VMware vCenter server details and click **Save**.

Add VMware vCenter Server

VCENTER SERVER IP ADDRESS OR HOST NAME

> 10.81.72.101

USERNAME

> administrator@vsphere.local

PASSWORD

> ••••••••

PORT

> 443

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

⚠ **Authorize Certificate**

Host nx-vc.flexpod.cisco.com you specified has identified itself with
a ca signed certificate for Active IQ Unified Manager.

View Certificate

Do you want to trust this certificate?

Yes   No
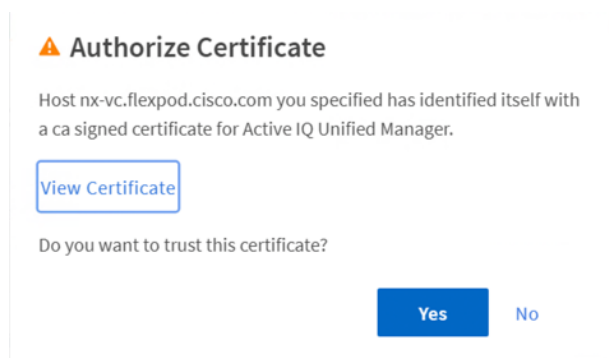
Note: It may take up to 15 minutes to discover vCenter. Performance data can take up to an hour to become available.

Procedure 2.   View Virtual Machine Inventory

The virtual machine inventory is automatically added to Active IQ Unified Manager during discovery of the vCenter server. Virtual machines can be viewed 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 issues.

Step 1.   Log into NetApp Active IQ Unified Manager.

Step 2.   Navigate to the VMware section located under Inventory, expand the section, and click **Virtual Machines**.

Virtual Machines ⦾

	Name	Status	Power Sta	Protocol	Capacity (Used | Allocated)	VM IOPS
⌄	AA17-I…Master	✓	ON	NFS	23.3 GB | 80 GB	0
⌄	AA17-Linux-21	✓	ON	NFS	22.2 GB | 100 GB	0
⌄	AA17-Linux-22	✓	ON	NFS	22.2 GB | 100 GB	0
⌄	AA17-Linux-23	✓	ON	NFS	2.16 GB | 80 GB	0
⌄	AA17-Linux-24	✓	ON	NFS	2.1 GB | 80 GB	0
⌄	AA17-Linux-25	✓	ON	NFS, VMFS	22.1 GB | 100 GB	0
⌄	AA17-Linux-26	✓	ON	NFS, VMFS	22.1 GB | 100 GB	0
⌄	AA17-Linux-27	✓	ON	NFS	2.1 GB | 80 GB	0
⌄	AA17-Linux-28	✓	ON		0 bytes | 0 bytes	
⌄	AA17-Linux-29	✓	ON	NFS	2.16 GB | 80 GB	0
⌄	AA17-Linux-30	✓	ON	NFS	2.1 GB | 80 GB	0
⌄	AIQUM-9.9	✓	ON	NFS	19.3 GB | 152 GB	6

Sidebar navigation:

DASHBOARD

COMMON TASKS
- PROVISIONING
- MANAGEMENT ACTIONS
- WORKLOAD ANALYSIS
- EVENT MANAGEMENT

INVENTORY
- STORAGE
- NETWORK
- PROTECTION
- VMWARE
  - vCenter
  - Virtual Machines

SETTINGS
- GENERAL

VIEW: Custom | Search | ☰ Filter

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

	Name	Status	Power Sta	Protocol	Capacity (Used | Allocated)	VM IOPS	VM Latency (ms)	Host IOPS	Host Latency (ms)	Network Latency (ms)
⌃	AA17-I…Master	✓	ON	NFS	23.3 GB | 80 GB	0	0	1	0	0

POWER
✓ ON

VCENTER SERVER
rtpb1-fp-vcetner.b1.cspg.local ↗

TOPOLOGY VIEW

Compute

VDISK	VM	HOST
scsi0:0	AA17-IOM-Master	192.168.17.16
IOPS 0	IOPS 0	IOPS 1
LATENCY 0 ms	LATENCY 0 ms	LATENCY ⓘ 0 ms

NETWORK

LATENCY
0 ms

Storage

DATASTORE	VMDK
infra_datastore_1	AA17-IOM-Master.vmdk
IOPS 13	
LATENCY 0.3 ms	
View in vCenter ↗	
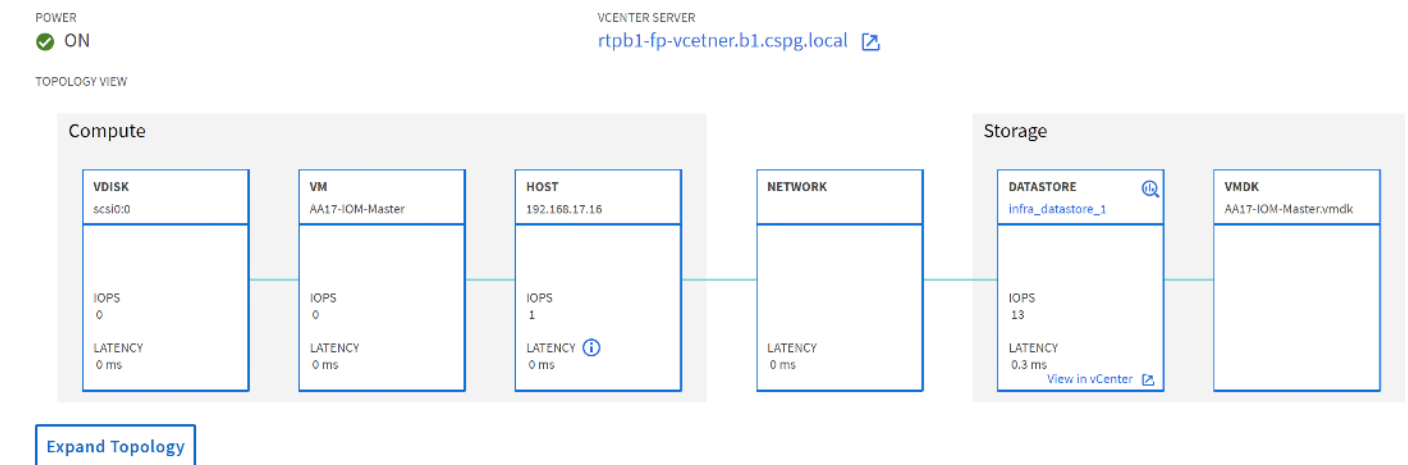
Expand Topology

Step 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 VM 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 leverages artificial intelligence and machine learning to provide analytics and actionable intelligence for ONTAP storage systems. Active IQ uses AutoSupport data to deliver proactive guidance and best practices recommendations to optimize storage performance and minimize risk. Additional Active IQ documentation is available on the [Active IQ Documentation Resources](#) web page.
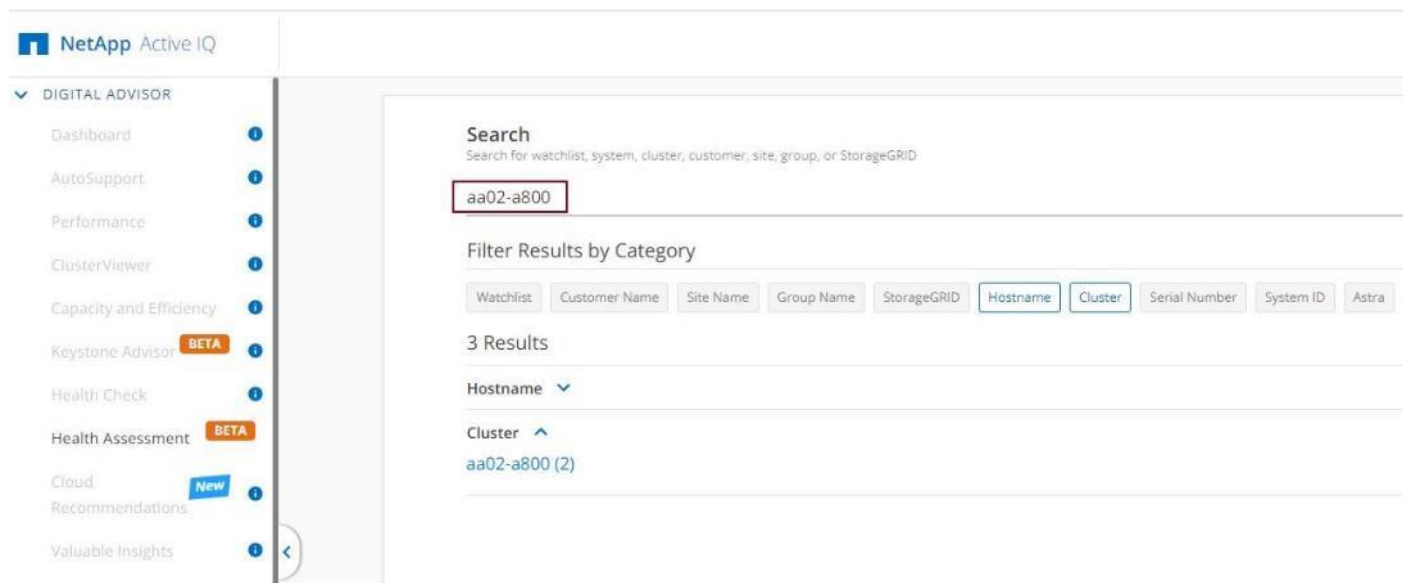
Note: Active IQ is automatically enabled when AutoSupport is configured on the NetApp ONTAP storage controllers.

Procedure 1.   Configure NetApp Active IQ

Step 1.   Navigate to the Active IQ portal at https://activeiq.netapp.com/.

Step 2.   Login with NetApp support account ID.

Step 3.   At the Welcome screen enter the cluster name or one of controller serial numbers in the search box. Active IQ will automatically begin searching for the cluster and display results below:

Step 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 recommen-dations 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 watchlist and launch Digital advisor dashboard for the watchlist.

Step 1.   Click **GENERAL > Watchlists** in the left menu bar.

Step 2.   Enter a name for the watchlist.

Step 3.   Select the radio button to add systems by serial number and enter the cluster serial numbers to the watchlist.

Step 4.   Check the box for **Make this my default watchlist** if desired.

Step 5. Click **Create Watchlist**.

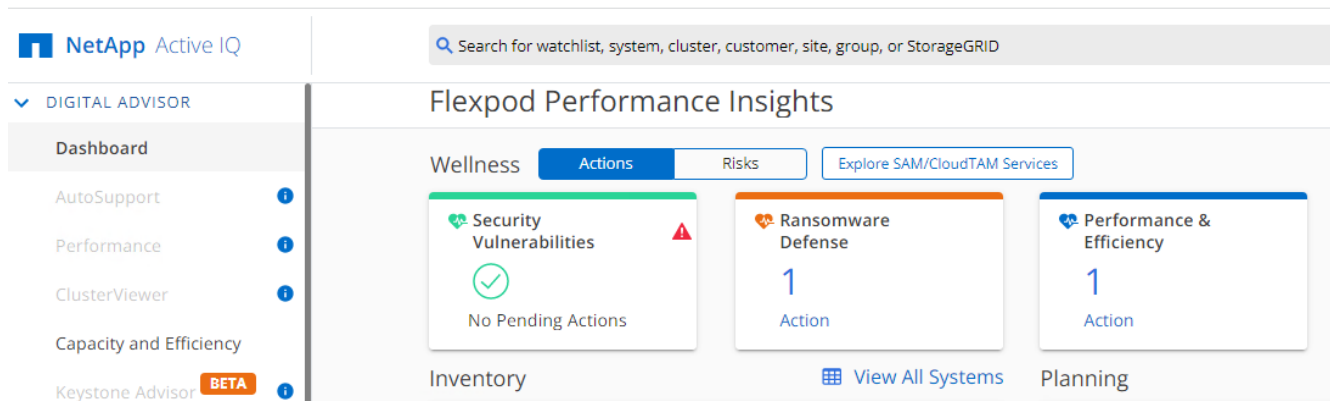**Step 6.** Click **GENERAL** > **Watchlists** in the left menu bar again to list the watchlist created.

Note: The Discovery Dashboard functionality has been moved to IB (Installed Base) console. Notice that Discovery Dashboard is greyed out under SALES TOOLS.

Step 7. Click the blue box labelled DA to launch the specific watchlist in **Digital Advisor Dashboard**.

Step 8. Review the enhanced dashboard to learn more about any recommended actions or risks.

Step 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.

Step 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: Additional tutorials and video walk-throughs of Active IQ features can be viewed on the following page: https://docs.netapp.com/us-en/active-iq/

FlexPod Backups

Cisco Intersight SaaS Platform

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

Procedure 1. Cisco Nexus and MDS Backups

The configuration of the Cisco Nexus 9000 and Cisco MDS 9132T switches can be backed up manually at any time with the copy command, but automated backups can be enabled using the NX-OS feature scheduler.

An example of setting up automated configuration backups of one of the NX-OS switches is shown below:

```
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 only needed when Mgmt0 interface is part of VRF management.

Verify the scheduler job has been correctly setup using 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 can be found here:
https://www.cisco.com/c/en/us/td/docs/dcn/nx-os/nexus9000/102x/configuration/system-management/cisco-nexus-9000-series-nx-os-system-management-configuration-guide-102x/m-configuring-the-scheduler-10x.html

Procedure 2. VMware VCSA Backup

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

Step 1. Connect to the VCSA Console at **https://<VCSA IP>:5480**.

Step 2. Log in as **root**.

Step 3. Click **Backup** in the list to open the Backup Schedule Dialogue.

Step 4. To the right of Backup Schedule, click **CONFIGURE**.

Step 5. Specify the following:

   a. The Backup location with the protocol to use (FTPS,HTTPS,SFTP,FTP,NFS,SMB, and HTTP)

   b. The Username and Password. For the NFS (NFS3) example captured below, the username is root and use a random password because NFSv3 sys security was configured.

   c. The Number of backups to retain.

Create Backup Schedule

Backup location ⓘ	nfs://10.102.1.11/software/Config-Backup/vCenter	
Backup server credentials	User name	root
	Password	••••••••
Schedule ⓘ	Daily ∨ 02 : 15 A.M. America/New_York	
Encrypt backup (optional)	Encryption Password	
	Confirm Password	
Number of backups to retain	○ Retain all backups	
	● Retain last 7 backups	
Data	☑ Stats, Events, and Tasks	37 MB
	☑ Inventory and configuration	87 MB
	Total size (compressed)	124 MB

CANCEL   **CREATE**

Step 6.   Click **CREATE**.

The Backup Schedule Status should now show **Enabled**.

Step 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.

Step 8.   Restoration can be initiated 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 Adaptative 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

GbE—Gigabit Ethernet

Gbit/s—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)

PnP–Plug-n-Play

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 (1 GbE transceiver)

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

SGACL–Security-Group ACL

SGT–Scalable Group Tag, sometimes reference 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

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, for the purposes of aiding understanding. This 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 this 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 are: • The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it. • 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 utilized CapEx. • 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. • Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes. Such services are typically implemented as "microservices," which are accessed via REST APIs. This architectural style supports composition of service components into systems. Access to and management of aaS assets is via a web GUI and/or APIs, such that Infrastructure-as-code (IaC) techniques can be used for automation, for example, Ansible and Terraform. 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. Service access control, integrated with corporate IAM, can be mapped to specific users and business activities, enabling consistent policy controls across services, wherever they are delivered from.
Ansible	An infrastructure automation tool, used to implement processes for instantiating and configuring IT service components, such as VMs on an IaaS platform. Supports the consistent execution of processes defined in YAML "playbooks" at scale, across multiple targets. Because the Ansible artefacts (playbooks) are text-based, they can be stored in a Source Code Management (SCM) system, such as GitHub. This allows for software development like processes to be applied to infrastructure automation, such as, Infrastructure-as-code (see IaC below). https://www.ansible.com
AWS **(Amazon Web Services)**	Provider of IaaS and PaaS. https://aws.amazon.com
Azure	Microsoft IaaS and PaaS. 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." https://en.wikipedia.org/wiki/Colocation_centre

Containers (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 which support the Docker environment. In the context of aaS, microservices are typically packaged within Linux containers orchestrated by Kubernetes (K8s). https://www.docker.com https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/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 issue handling. DevOps is closely aligned with IaC, continuous integration and deployment (CI/CD), and Agile software development practices. https://en.wikipedia.org/wiki/DevOps 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. This could be for reasons related to low latency response, reduction of the amount of unprocessed data being transported, efficiency of resource utilization, and so on. The generic label for this is Multi-access Edge Computing (MEC), or Mobile Edge Computing for mobile networks specifically. From an application experience perspective, it is important to be able to utilize, at the edge, the same operations model, processes, and tools used for any other compute node in the system. https://en.wikipedia.org/wiki/Mobile_edge_computing
IaaS (Infrastructure as-a-Service)	Infrastructure components provided 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 (Infrastructure as-Code)	Given the ability to automate aaS via APIs, the implementation of the automation is typically via Python code, Ansible playbooks, and similar. 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. This means 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. https://en.wikipedia.org/wiki/Infrastructure_as_code
IAM (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. https://en.wikipedia.org/wiki/Identity_management
IBM (Cloud)	IBM IaaS and PaaS. https://www.ibm.com/cloud

Intersight	Cisco Intersight™ is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html
GCP **(Google Cloud Platform)**	Google IaaS and PaaS. https://cloud.google.com/gcp
Kubernetes **(K8s)**	Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. https://kubernetes.io
Microservices	A microservices architecture is characterized by processes implementing fine-grained services, typically exposed via REST APIs and which can be composed into systems. The processes are often container-based, and the instantiation of the services 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. 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 very closely associated with DevOps practices.
Private on-premises data center	A data center infrastructure housed within an environment owned by a given enterprise is distinguished from other forms of data center, with the implication that the private data center is more secure, given that access is restricted to those authorized by the enterprise. Thus, circumstances can arise where very sensitive IT assets are only deployed 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 are key technologies needed 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. https://en.wikipedia.org/wiki/Representational_state_transfer
SaaS **(Software-as-a-Servic e)**	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**	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

Markup Language)	the aaS for access control decisions. https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language
Terraform	An open-source IaC software tool for cloud services, based on declarative configuration files. https://www.terraform.io

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://cs.co/en-cvds) at https://cs.co/en-cvds.

CVD Program