



The bridge to possible

Design and Deployment Guide

Cisco Public

FlashStack as a Workload Domain for VMware Cloud Foundation Design and Deployment Guide

Published: May 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 FlashStack Datacenter solution is a validated approach for deploying Cisco and Pure Storage technologies and products to build a shared private and public cloud infrastructure. Cisco and Pure Storage have partnered to deliver a series of FlashStack solutions that enable strategic data-center platforms. The success of the FlashStack 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 the deployment details of incorporating FlashStack Datacenter as a workload domain for VMware Cloud Foundation. For an in depth design description, refer to the following guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_vsi_ucs_xseries_5gen_design.html

VMware Cloud Foundation provides a complete set of software defined services to run enterprise apps, both traditional and containerized, in private or public cloud environments. VMware Cloud Foundation simplifies the private cloud deployment and provides a streamlined path to the hybrid cloud by delivering a single integrated solution that is easy to deploy, operate and manage.

VMware Cloud Foundation (VCF) provides following benefits in a data center environment:

- **Integrated Stack:** VCF is an engineered solution that integrates the entire VMware software-defined stack with guaranteed interoperability.
- **Standardized Architecture:** VCF is built upon standard VMware Validated Design architecture and therefore ensures quick, repeatable deployments while eliminating risk of misconfigurations.
- **Lifecycle Management:** VCF includes lifecycle management services that automate day 0 to day 2 operations, resource provisioning plus patching and upgrades.

Some of the key advantages of integrating FlashStack Datacenter as a workload domain for VMware Cloud Foundation are:

- **Simpler and programmable infrastructure:** FlashStack infrastructure delivered as infrastructure-as-a-code through a single partner integrable open API.
- **Latest hardware and software compute innovations:** Policy-based configurations, delivered using Cisco Intersight, to deploy and manage the latest processor, memory, network, and power/cooling improvements.
- **Storage Modernization:** Deliver high-speed, consistent, low latency, multi-tenant storage using a range of Pure Storage all-flash arrays.
- **Innovative cloud operations:** Continuous feature delivery with no need for maintaining on-premises virtual machines supporting management functions.
- **Built for investment protection:** Design ready for future technologies such as liquid cooling and high-Wattage CPUs; CXL-ready.

The FlashStack workload domain includes integration of Cisco Intersight with Cisco UCS, Cisco Nexus, and MDS, an Intersight connector for Pure Storage arrays, plus VMware vCenter and Pure Storage vSphere Plugins to deliver monitoring, orchestration, and workload optimization capabilities for different layers (virtualization and storage) of the FlashStack infrastructure. The modular nature of the Cisco Intersight platform also provides an easy upgrade path to additional services, such as Intersight Workload Optimization and Intersight Cloud Orchestrator.

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

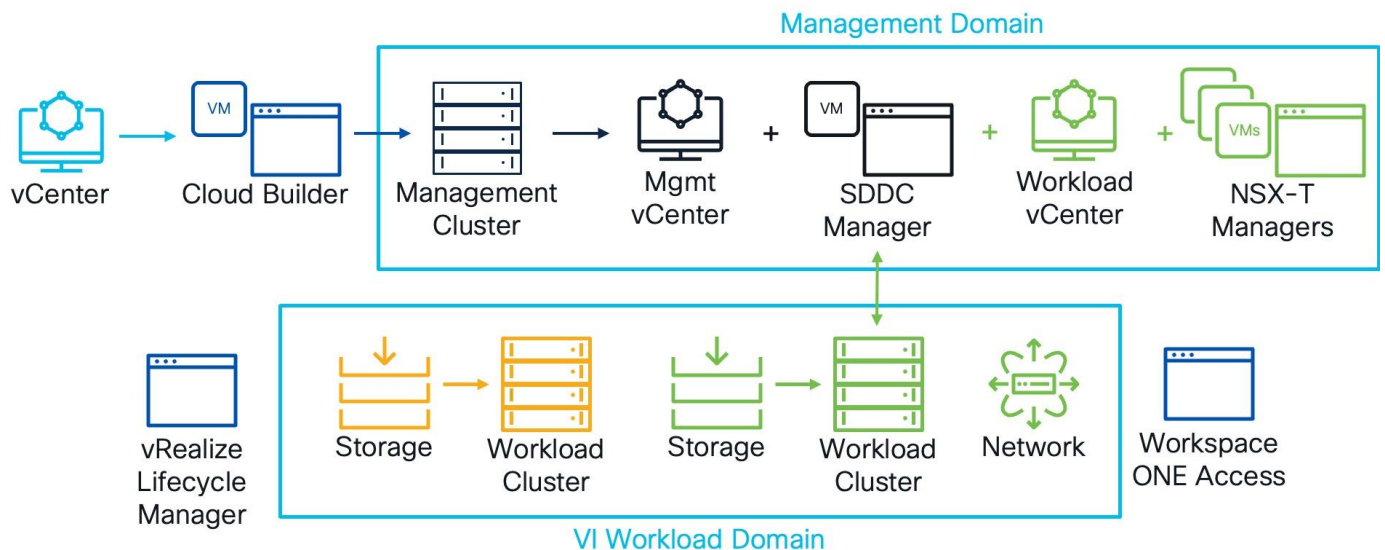
Solution Overview

This chapter contains the following:

- [Audience](#)
- [Purpose of this Document](#)
- [What's New in this Release?](#)
- [Infrastructure as Code with Ansible to setup FlashStack and VCF Management Domain](#)

VMware Cloud Foundation enables data center administrators to provision an application environment in a quick, repeatable, and automated manner. VMware Cloud Foundation consists of workload domains which represent an application-ready infrastructure. A workload domain represents a logical unit that groups ESXi hosts managed by a vCenter Server instance with specific characteristics according to VMware best practices.

To deploy and manage the workload domains, VMware Cloud Foundation introduces VMware Cloud Builder and VMware Cloud Foundation Software Defined Data Center (SDDC) Manager. VMware Cloud Builder automates the deployment of the software defined stack, creating the first software defined unit known as the management domain. After the management domain is successfully setup, using the newly deployed SDDC Manager, a virtual infrastructure administrator or cloud administrator provisions the FlashStack Datacenter blades as a new workload domain.



Workload domain installation requires administrators to configure network, compute, and storage as well as install VMware vSphere ESXi software on the hosts that become part of workload domains (including the management domain). To automate the infrastructure setup, Cisco Intersight configurations are (optionally) configured using RedHat Ansible playbooks for an easy on-boarding experience.

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 on following two key areas:

-
- Deploying VMware Cloud Foundation management domain on Cisco UCS C240 M5 “vSAN-ready” servers managed using Cisco Intersight.
 - Configuring Cisco UCS X210c compute nodes in the FlashStack configuration and adding these FlashStack ESXi hosts to VMware Cloud Foundation as a Virtual Infrastructure (VI) workload domain.

While VMware Cloud Foundation can be utilized in public cloud such as VMware Cloud on AWS as well as hybrid cloud solutions, the discussion in this document focuses solely on the on-prem data center design and deployment. This document augments the FlashStack Datacenter with Cisco UCS X-Series Cisco Validated Design (CVD):

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/FlashStack_xseries_esxi7u2_design.html and explains new and changed information around VMware Cloud Foundation deployment. For a complete FlashStack configuration including various management components, refer to: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/FlashStack_xseries_vmware_7u2.html.

What’s New in this Release?

The following elements distinguish this FlashStack Datacenter Cisco Validated Design from previous designs:

- VMware Cloud Foundation management domain deployment on vSAN ready nodes.
- Integration of FlashStack Datacenter as a workload domain in VMware Cloud Foundation.
- Automated configuration of the ESXi hosts for both the VMware Cloud Foundation management and workload domains using Cisco Intersight.

Like all other FlashStack solution designs, FlashStack as a workload domain for VMware Cloud Foundation solution is configurable according to demand and usage. Customers can purchase exactly the infrastructure they need for their current application requirements and can then scale up by adding more resources to the FlashStack system or scale out by adding more FlashStack instances. Because the workload domain management of VMware Cloud Foundation is hosted in a separate domain, and infrastructure management is hosted by Cisco Intersight in the cloud, the solution can respond to the speed and scale of customer deployments swiftly at cloud-scale.

Deployment Hardware and Software

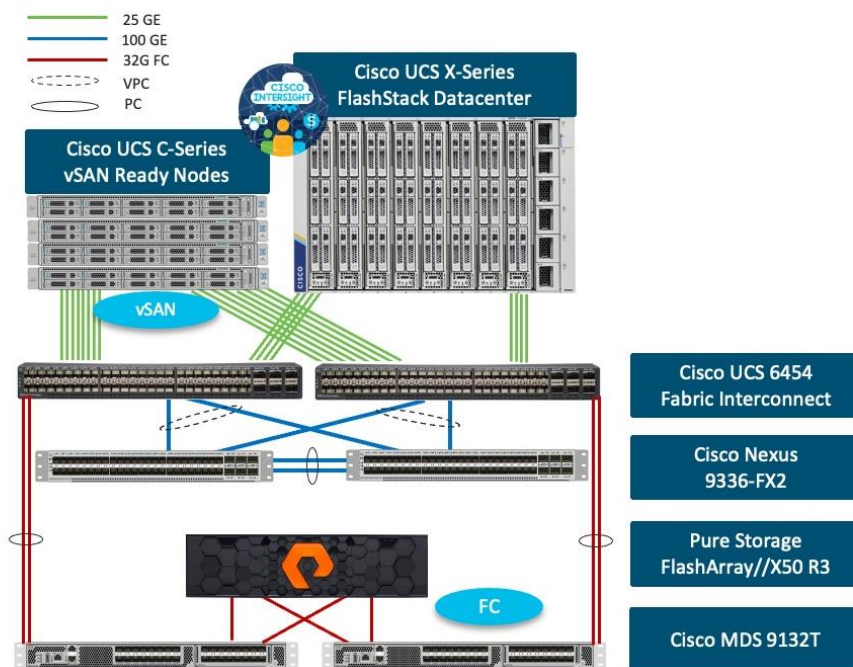
This chapter contains the following:

- [Installation Overview](#)
- [Design Requirements](#)
- [Physical Topology](#)
- [Ansible Automation for Solution Deployment](#)

The FlashStack as a workload domain for VMware Cloud Foundation solution delivers a VMware Cloud Foundation VI workload domain built on Cisco UCS X-Series based FlashStack infrastructure. For setting up the VMware Cloud Foundation management domain, 4 Cisco UCS C240 M5 servers with vSAN certified components are utilized. The VMware vSphere 7.0 U3 hypervisor is installed on M.2 boot optimized Solid State Drives (SSD) and vSAN is configured by VMware Cloud Builder as principal storage. For setting up the VMware Cloud Foundation VI workload domain, 3 UCS X210c compute nodes are utilized. VMware vSphere 7.0 U3 hypervisor is installed on the Fibre Channel (FC) LUNs hosted on Pure Storage FlashArray//X50 R3 storage array. The Pure Storage FlashArray//X50 R3 also provides Fibre Channel based principal storage for setting up the VMware infrastructure. Additional principal and/or secondary storage can be provisioned from the Pure Storage array as FC LUNs using VMFS or as vVols.

The Cisco UCS X-Series chassis and all the management rack servers are connected to a single pair of Cisco UCS 6454 Fabric Interconnects configured for Cisco Intersight Managed Mode (IMM).

Figure 1. FlashStack Solution for VMware Cloud Foundation



Note: Some customers might own Cisco UCS C-Series systems that are not supported in Intersight Managed Mode (IMM) because of unsupported components. These C-Series servers cannot be connected

to the same Cisco UCS FIs where the FlashStack Cisco UCS X-Series chassis is connected and would need to be connected to a separate pair of FIs and be configured in Cisco UCSM mode.

Installation Overview

Installation of the VCF on FlashStack solution generally follows the normal process for a standard FlashStack for Virtual Server Infrastructure (VSI) deployment. Because of this, many steps of the installation in this document will refer to the existing FlashStack for VSI deployment Cisco Validated Design (CVD) document. Design changes or deviations from the base CVD will be noted in each step. As such, the document assumes that this deployment is a new environment, although it is possible to use an existing FlashStack installation as a VCF workload domain as well. An existing VMware vCenter server must be deployed and available for the installation, to be used for building an ESXi installation ISO image, and for hosting the VCF Cloud Builder virtual machine.

The solution installation has the following steps:

1. Mount, install, cable and power all equipment into the cabinets or racks.
2. Configure prerequisites such as NTP and DNS A records.
3. Configure the Cisco Nexus network switches with the appropriate interfaces, port-channels, vPC, VLANs and gateways.
4. Configure the Pure Storage FlashArray with the recommended settings, creating the volumes.
5. Configure Cisco UCS via Cisco Intersight, which creates all the pools, policies and profiles for the C-series and X-series servers. This process can be automated using the supplied RedHat Ansible playbooks.
6. Configure the Cisco MDS storage switches, creating the aliases, zones and zonesets with the worldwide port names (WWPNs) from the Cisco UCS configuration created earlier.
7. Create hosts and host groups on the Pure Storage FlashArray and connect the hosts and groups to their volumes.
8. Create the ESXi custom installation ISO image using vCenter Image Builder on an existing VMware vCenter server.
9. Install ESXi on the management and workload hosts using the custom image.
10. Configure ESXi management interfaces on the hosts.
11. Configure ESXi settings to prepare them for VCF. This process can be automated using the supplied RedHat Ansible playbooks.
12. Deploy the VCF Cloud Builder VM from the OVA file on an existing VMware vCenter server.
13. Use Cloud Builder to deploy the VCF management domain on the 4 vSAN ready hosts, which also deploys the VMware Software Defined Data Center (SDDC) virtual machine in the management domain.
14. Use SDDC to commission the workload domain hosts, and to deploy the workload domain.

Design Requirements

The FlashStack as a workload domain for VMware Cloud Foundation design 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 VMware Cloud Foundation and other external automation tools
- Cloud-enabled design which can be configured, managed, and orchestrated from the cloud using GUI or APIs

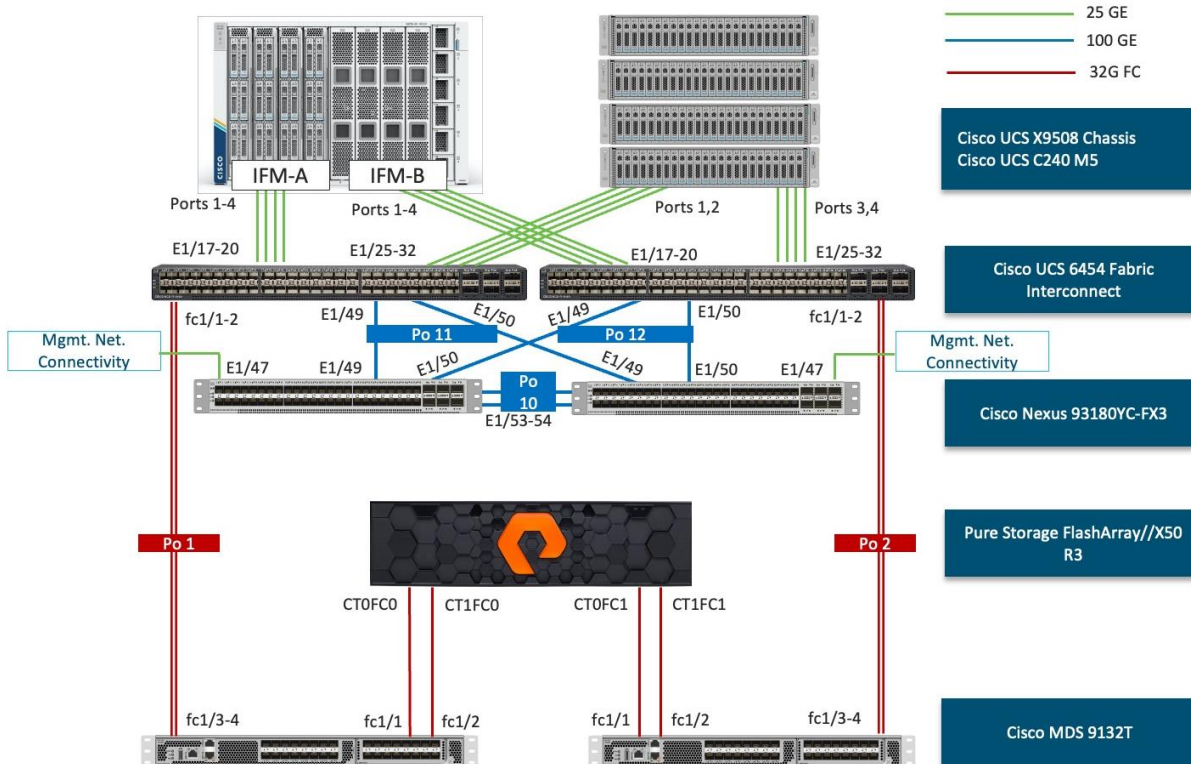
Physical Topology

FlashStack as a workload domain for VMware Cloud Foundation was validated using a Fibre Channel (FC) boot from SAN configuration.

FlashStack Datacenter with Fibre Channel Design

For the FC design, the Pure Storage FlashArray//X50 R3 and the Cisco UCS X-Series blades are connected through Cisco MDS 9132T Fibre Channel Switches and boot from SAN for stateless computing via the FC network. When adding FlashStack as a VI workload domain, additional FC LUNs are used as VMFS datastores for VCF principal storage. The physical topology is shown in [Figure 2](#).

Figure 2. Physical Topology



The components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the rack server and chassis connectivity.

- 4 Cisco UCS C-Series* vSAN ready nodes are connected to the Fabric Interconnects (FI) and are managed using Cisco Intersight. Two 25 Gigabit Ethernet ports from each Cisco UCS C-Series server are connected, one to each FI.
- The Cisco UCS X9508 Chassis connects to the FIs using Cisco UCSX 9108-25G Intelligent Fabric Modules (IFMs), where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI. The remaining 4 ports from each IFM can be connected to the same FIs if additional bandwidth is required.
- Cisco Nexus 93180YC-FX3 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100 Gigabit Ethernet uplink ports connect to the Cisco Nexus 93180YC-FX3 Switches in a vPC configuration.
- For Cisco UCS to SAN connectivity, Cisco UCS 6454 Fabric Interconnects connect to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel uplinks, each configured as a single port channel.
- For Pure Storage FlashArray//X50 R3 connectivity, each controller connects to both Cisco MDS 9132T switches using 32-Gbps Fibre Channel.

Note: * Since Cisco UCS C-series is being managed and configured by Cisco Intersight Managed Mode, the vSAN ready nodes must satisfy the software and hardware requirements outlined here:

https://intersight.com/help/saas/supported_systems

VLAN Configuration

[Table 1](#) lists VLANs configured for setting up the FlashStack environment along with their usage.

Table 1. VLAN Usage

VLAN ID	Name	Description	Subnet
2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1)	
1010	OOB-Mgmt	Existing management VLAN where all the management interfaces for various devices will be connected	10.101.0.0/24
1011	IB-Mgmt	FlashStack In-band management VLAN utilized for all in-band management connectivity such as ESXi hosts, VM management, and VCF components (Cloud Builder, SDDC Manager, all NSX managers, all vCenters)	10.101.1.0/24
1012	VM-Traffic	Application VLAN (one of many) where application VMs will be deployed. Adjust the name and add more VLANs as needed.	10.101.2.0/24
3001	Mgmt-vSAN	vSAN VLAN for the management domain	192.168.1.0/24
3002	Mgmt-Host-Overlay	NSX-T Host Overlay Network VLAN for the management domain	192.168.2.0/24
3003	WD-Host-Overlay	NSX-T Host Overlay Network VLAN for the FlashStack VI workload domain	192.168.3.0/24
3030	vMotion	Common vMotion VLAN for both management and VI workload domains	192.168.31.0/24

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

- VLAN 1010 is the management VLAN where out of band management interfaces of all the physical devices are connected.
- VLAN 1011 is used for in-band management of VMs, ESXi hosts, and other infrastructure services in the FlashStack environment. This VLAN is also used for deploying VMware Cloud Foundation components.
- VLAN 3001 is used for VMware Cloud Foundation management domain vSAN configuration.
- VLANs 3002 and 3003 are separate NSX-T host overlay VLANs for VMware Cloud Foundation management and FlashStack VI workload domains. Depending on the customer requirements, a single VLAN can be used.
- VLAN 3030 is a common vMotion VLAN for VMware Cloud Foundation management and FlashStack VI workload domains. Depending on the customer requirements, separate VLANs can be configured to isolate vMotion traffic.

Physical Components

[Table 2](#) lists the required hardware components used to build the validated solution. Customers are encouraged to review their requirements and adjust the size or quantity of various components as needed.

Table 2. FlashStack as a workload domain for VMware Cloud Foundation hardware components

Component	Hardware	Comments
Cisco Nexus Switches	Two Cisco Nexus 93180YC-FX3 switches	
Cisco MDS Switches	Two Cisco MDS 9132T switches	
Pure Storage FlashArray	A Pure Storage FlashArray//X50 R3 with appropriate storage and network connectivity	Customer requirements will determine the amount and type of storage. The Pure Storage Array should support 32Gbps (or 16 Gbps) FC connectivity, and optional 25Gbps (or 100 Gbps) ethernet.
Fabric Interconnects	Two Cisco UCS 6454 Fabric Interconnects	These fabric interconnects will be shared between the management and the workload domain
Management Domain Compute		
Cisco UCS Servers	A minimum of four Cisco UCS C-Series vSAN ready (or vSAN compatible) nodes	vSAN ready nodes are recommended for ease of deployment however, customers can also utilize existing Cisco UCS C-Series servers with vSAN supported components
FlashStack VI Workload Domain Compute		
Cisco UCS Chassis	A minimum of one UCS X9508 chassis.	Single chassis can host up to 8 Cisco UCS X210c compute nodes
Cisco UCS Compute Nodes	A minimum of three Cisco UCS X210c compute nodes	Four compute nodes are recommended but three compute nodes will work.

Software Components

[Table 3](#) lists various software releases used in the solution.

Table 3. Software components and versions

Component	Version
Cisco Nexus 93180YC-FX3	9.3(10)
Cisco MDS 9132T	9.2(2)
Cisco UCS Fabric Interconnects	4.2(3b)
Cisco UCS C-Series vSAN ready nodes	4.2(3b)
Cisco UCS X210c compute nodes	5.0(4a)
Cisco Intersight Assist Appliance	1.0.9-499 (will automatically upgrade to latest version when claimed in Cisco Intersight)
Pure Storage FlashArray//X50 R3 - Purity//FA	6.4.3
Pure Storage VMware Appliance	4.0.0
Pure Storage vSphere Client Plugin	5.2.1
Pure Storage FlashArray VASA Provider	1.2
VMware Cloud Foundation	
Cloud Builder VM	4.5
SDDC Manager	4.5
VMware NSX-T	3.2.1.2.0
VMware vCenter	7.0 Update 3h
VMware ESXi	7.0 Update 3g
Cisco VIC FC Driver (nfnic)	5.0.0.37
Cisco VIC Ethernet Driver (nenic)	1.0.45.0

Ansible Automation for Solution Deployment

This section provides information about setting up and running Ansible playbooks to configure the infrastructure for VMware Cloud Foundation. Skip this section if the infrastructure configuration for VMware Cloud Foundation hosts is being implemented manually.

Ansible automation requires a management workstation (control machine) to run Ansible playbooks for configuring Cisco Nexus, Pure Storage, Cisco UCS, Cisco MDS, and VMware ESXi.

Management Workstation

A management workstation is a VM where Ansible is installed and has access to the Internet to download various packages and clone the playbook repository. Instructions for installing the workstation Operating System (OS) or complete setup of Ansible are not included in this document, however, basic installation and configuration of Ansible is provided as a reference. A guide for installing and getting started with Ansible can be found at: https://docs.ansible.com/ansible_community.html.

Prepare Management Workstation (Control Node)

In this section, the installation steps are performed on the CentOS Stream 8 management host to prepare the host for solution deployment to support the automation of Cisco UCS, Cisco Nexus, Pure Storage, Cisco MDS and VMware ESXi using Ansible Playbooks. The following steps were performed on a CentOS Stream 8 Virtual Machine* as the root user.

Note: * CentOS Stream 8 “Server with GUI” option was selected when installing the operating system.

Procedure 1. Prepare the Management Workstation

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

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

Step 2. Install Ansible.

```
dnf install ansible
```

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

```
ansible --version
ansible [core 2.13.3]
  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. Update pip and setuptools.

```
pip3 install --upgrade pip
pip3 install --upgrade setuptools
```

Step 5. Install ansible-galaxy collections for Cisco Intersight, NX-OS, and VMware as follows:

```
ansible-galaxy collection install cisco.intersight
ansible-galaxy collection install cisco.nxos
pip3 install ansible-pylibssh
ansible-galaxy collection install community.vmware
pip3 install -r ~/.ansible/collections/ansible_collections/community/vmware/requirements.txt
```

Troubleshooting Tip

In some instances, the following error messages might be seen when executing VMware-specific ansible playbooks:

```
An exception occurred during task execution. To see the full traceback, use -vvv. The error was:
ModuleNotFoundError: No module named 'requests'
fatal: [10.101.1.101 -> localhost]: FAILED! => {"changed": false, "msg": "Failed to import the required Python library (requests) on aa01-linux8.vm.vcf.local's Python /usr/bin/python3.8. Please read the module documentation and install it in the appropriate location. If the required library is installed, but Ansible is using the wrong Python interpreter, please consult the documentation on ansible_python_interpreter"}
```

```
An exception occurred during task execution. To see the full traceback, use -vvv. The error was:
ModuleNotFoundError: No module named 'pyVim'
fatal: [10.101.1.101 -> localhost]: FAILED! => {"changed": false, "msg": "Failed to import the required Python library (PyVmomi) on aa01-linux8.vm.vcf.local's Python /usr/bin/python3.8. Please read the module documentation and install it in the appropriate location. If the required library is installed, but Ansible is using the wrong Python interpreter, please consult the documentation on ansible_python_interpreter"}
```

To fix this issue, use the appropriate version of PIP to install “requests” and “pyvmomi”:

```
pip3.8 install requests
```

```
pip3.8 install pyVmomi
```

Ansible Playbooks

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

- Cisco DevNet: <https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/CVD-FlashStack-IMM-VCF>
- GitHub repository: <https://github.com/ucs-compute-solutions/CVD-FlashStack-IMM-VCF>

The provided Ansible playbooks in the GitHub repository can be used to configure Cisco Intersight, creating the pools, policies and templates for the service profiles used by the four VCF Management servers, and the VI Workload Domain hosts. In addition, extra playbooks are provided which perform some post-installation configurations against the ESXi hosts, which are necessary prior to using the hosts in a VCF domain. Instructions for configuring the variables required, and the use of the playbooks are provided in the README document of the GitHub repository. Additional playbooks are available from the base FlashStack GitHub repository, which can be used to automate the configuration of MDS, Nexus and Pure Storage, which can be found here: https://github.com/ucs-compute-solutions/FlashStack_IMM_Ansible

Switch Configuration

This chapter contains the following:

- [Physical Connectivity](#)
- [Initial Configuration](#)
- [Enable Cisco Nexus Features and Global Configuration](#)
- [Create VLANs](#)
- [Create Port Channels](#)
- [Create Port Channel Parameters](#)
- [Configure Virtual Port Channels](#)
- [Configure IP Gateways](#)

This chapter provides the procedure to configure the Cisco Nexus 93180YC-FX3 switches used for ethernet LAN switching in this solution. The switch configuration for this validated design is based on the switching configuration explained in FlashStack Datacenter with Cisco UCS X-Series Cisco Validated Design (CVD): https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#NetworkSwitchConfiguration therefore this section only explains the changes to switching configuration from the base CVD.

Physical Connectivity

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

Initial Configuration

For setting up the initial switch configuration, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashStackCiscoNexusBase

Enable Cisco Nexus Features and Global Configuration

To enable the required Cisco Nexus features, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashStackCiscoNexusSwitchConfiguration

Create VLANs

Procedure 1. Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), follow this step on both switches. Refer to the VLAN information in [Table 1](#) for setting up all required VLANs.

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

```
vlan <native-vlan-id for example 2>
name Native-Vlan
vlan <oob-mgmt-vlan-id for example 1010>
name OOB-Mgmt
vlan <ib-mgmt-vlan-id for example 1011>
name IB-Mgmt
vlan <application-vm-vlan-id for example 1012>
name VM-Traffic
vlan <vsan-vlan-id for example 3001>
```



```
name Mgmt-vSAN
vlan <nsx-mgmt-host-overlay-vlan-id for example 3002>
name Mgmt-Host-Overlay
vlan <nsx-WorkloadDomain-host-overlay-vlan-id for example 3003>
name WD-Host-Overlay
vlan <vmotion-vlan-id for example 3030>
name vMotion
```

Note: Separate vMotion VLANs for management and VI workload domain can be configured for traffic isolation.

Create Port Channels

To set up Port Channels on both Nexus switches, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashStackCiscoNexusSwitchConfiguration

Create Port Channel Parameters

Procedure 1. Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow these steps on both Cisco Nexus switches.

Step 1. From the global configuration mode, run the following commands to setup VPC Peer-Link port-channel:

```
interface Po10
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <oob-mgmt-vlan-id>, <ib-mgmt-vlan-id>, <application-vlan-id>, <vsan-vlan-id>, <
nsx-mgmt-host-overlay-vlan-id>, < nsx-WorkloadDomain-host-overlay-vlan-id>, <vmotion-vlan-id>
spanning-tree port type network
```

Step 2. From the global configuration mode, run the following commands to setup port-channels for UCS FI 6454 connectivity:

```
interface Po11
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <oob-mgmt-vlan-id>, <ib-mgmt-vlan-id>, <vm-traffic-vlan-id>, <vsan-vlan-id>, <
nsx-mgmt-host-overlay-vlan-id>, < nsx-WorkloadDomain-host-overlay-vlan-id>, <vmotion-vlan-id>
spanning-tree port type edge trunk
mtu 9216
!
interface Po12
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <oob-mgmt-vlan-id>, <ib-mgmt-vlan-id>, <vm-traffic-vlan-id>, <vsan-vlan-id>, <
nsx-mgmt-host-overlay-vlan-id>, < nsx-WorkloadDomain-host-overlay-vlan-id>, <vmotion-vlan-id>
spanning-tree port type edge trunk
mtu 9216
```

Step 3. From the global configuration mode, run the following commands to setup port-channels for connectivity to existing management switch:

```
interface Po101
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <oob-mgmt-vlan-id>, <ib-mgmt-vlan-id>
spanning-tree port type network
mtu 9216
!
exit
copy run start
```

UDLD for Cisco UCS Interfaces

For fibre-optic connections between Cisco UCS Fabric Interconnects and Cisco Nexus 93180YC-FX3 switches, UDLD configuration is automatically enabled, and no additional configuration is required on either device.

Configure Virtual Port Channels

For setting up Virtual Port Channel configuration on both Cisco Nexus switches, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashStackCiscoNexusSwitchConfiguration

Configure IP Gateways

VMware Cloud Foundation installation checks for gateways when configuring various VM Kernel ports on the ESXi hosts. If IP gateways for the VLANs explained below are present on the upstream switches, the configuration in this step can be skipped. If some or all the gateways are not configured, use Hot Standby Router Protocol (HSRP) and Switched Virtual Interface (SVI) on the Nexus switches to setup gateways for:

- Out-of-band management*
- in-band management*
- Application VM
- vSAN
- NSX host-overlay networks

Note: * Gateways for management networks will most likely be pre-configured in existing customer environments therefore exercise extreme caution when configuring new management IP gateways.

Procedure 1. Configure Nexus-A Switch

Step 1. From the global configuration mode, run the following commands to setup VPC Peer-Link port-channel:

```
feature interface-vlan
feature hsrp

interface Vlan1010
  description GW for Out-of-Band Mgmt 10.101.0.0/24 Network
  no shutdown
  no ip redirects
  ip address 10.101.0.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1010
  preempt delay minimum 300
  priority 105
  ip 10.101.0.254

interface Vlan1011
  description GW for In-band Management 10.101.1.0/24 Network
  no shutdown
  no ip redirects
  ip address 10.101.1.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1011
  preempt delay minimum 300
  priority 105
  ip 10.101.1.254
```

```

interface Vlan1012
  description GW for Application VM Traffic 10.101.2.0/24 Network
  no shutdown
  ! MTU should be adjusted based on application requirements
  mtu 1500
  no ip redirects
  ip address 10.101.2.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1012
    preempt delay minimum 300
    priority 105
    ip 10.101.2.254

interface Vlan3001
  description Gateway for Management Domain vSAN Network
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.1.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 3001
    preempt delay minimum 300
    priority 105
    ip 192.168.1.254

interface Vlan3002
  description Gateway for NSX Management Domain Host Overlay VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.2.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 3002
    preempt delay minimum 300
    priority 105
    ip 192.168.2.254

interface Vlan3003
  description Gateway for NSX Worload Domain Host Overlay VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.3.251/24
  hsrp version 2
  hsrp 3003
    preempt delay minimum 300
    priority 105
    ip 192.168.3.254

interface Vlan3030
  description Gateway for vMotion VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.30.251/24
  no ipv6 redirects
  hsrp version 2
  hsrp 3030
    preempt delay minimum 300
    priority 105
    ip 192.168.30.254

```

Procedure 2. Configure Nexus-B Switch

Step 1. From the global configuration mode, run the following commands to setup VPC Peer-Link port-channel:

```
feature interface-vlan
feature hsrp

interface Vlan1010
  description GW for Out-of-Band Mgmt 10.101.0.0/24 Network
  no shutdown
  no ip redirects
  ip address 10.101.0.252/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1010
    ip 10.101.0.254

interface Vlan1011
  description GW for In-band Management 10.101.1.0/24 Network
  no shutdown
  no ip redirects
  ip address 10.101.1.252/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1011
    ip 10.101.1.254

interface Vlan1012
  description GW for Application VM Traffic 10.101.2.0/24 Network
  no shutdown
  ! MTU should be adjusted based on application requirements
  mtu 1500
  no ip redirects
  ip address 10.101.2.252/24
  no ipv6 redirects
  hsrp version 2
  hsrp 1012
    ip 10.101.2.254

interface Vlan3001
  description Gateway for Management Domain vSAN Network
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.1.252/24
  no ipv6 redirects
  hsrp version 2
  hsrp 3001
    ip 192.168.1.254

interface Vlan3002
  description Gateway for NSX Management Domain Host Overlay VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.2.252/24
  no ipv6 redirects
  hsrp version 2
  hsrp 3002
    ip 192.168.2.254

interface Vlan3003
  description Gateway for NSX Worload Domain Host Overlay VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.3.252/24
  hsrp version 2
  hsrp 3003
    ip 192.168.3.254

interface Vlan3030
  description Gateway for vMotion VLAN
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.30.252/24
```

```
no ipv6 redirects
hsrp version 2
hsrp 3030
 ip 192.168.30.254
```

Storage Configuration

This chapter contains the following:

- [Pure Storage Purity//FA Configuration](#)

Pure Storage Purity//FA Configuration

Complete the Pure Storage FlashArray setup for Fibre Channel based storage access explained in the following section:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashArrayConfiguration

Note: Any iSCSI or FC-NVMe configuration sections can be skipped since this deployment only covers Fibre Channel based storage design for FlashStack.

At the completion of this step, the Pure Storage FlashArray management connectivity, volume configuration, FC and management interfaces, and boot LUNs for three ESXi workload domain hosts that support boot from SAN using FC are ready.

Note: Configuration of the hosts and the host initiators can be skipped, as they will be configured during a later step, after the Cisco UCS X-series blades are configured.

Cisco Intersight Managed Mode – Initial Setup

This chapter contains the following:

- [Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects](#)
- [Set up Cisco Intersight Account](#)
- [Set up Cisco Intersight Licensing](#)
- [Set Up Cisco Intersight Resource Group](#)
- [Set Up Cisco Intersight Organization](#)
- [Claim Cisco UCS Fabric Interconnects in Cisco Intersight](#)
- [Upgrade Fabric Interconnect Firmware using Cisco Intersight](#)

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 interconnect-attached systems. Cisco Intersight managed mode standardizes both policy and operation management for Cisco UCS C-series M5 and Cisco UCSX X210c M6 compute nodes used in this deployment guide. 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

During the initial setup, Cisco UCS FIs are configured in Intersight Managed Mode and added to a newly created Intersight account. Intersight organization creation, resource group definition and license setup are also part of the initial setup. At the end of this section, customers can start creating various chassis and server level policies and profiles to deploy UCS compute nodes.

Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

For setting up Cisco UCS 6454 Fabric Interconnects in Intersight Managed Mode, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscoIntersightManagedModeConfiguration

Note: 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.

Set up Cisco Intersight Account

For setting up a new Cisco Intersight Account, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscoIntersightAccount

Note: Setting up a new Intersight account is not necessary if customers plan to add the Cisco UCS FIs to an existing account.

Set up Cisco Intersight Licensing

All new Cisco Intersight accounts need to be enabled for Cisco Smart Software Licensing. For setting up Cisco Intersight licensing, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscolntersightAccount

Set Up Cisco Intersight Resource Group

A Cisco Intersight resource group is created where resources such as various targets will be logically grouped. A single resource group is created to host all the resources in this deployment. To configure a resource group, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscolntersightAccount

Set Up Cisco Intersight Organization

All Cisco Intersight managed mode configurations including policies and profiles are defined under an organization. To define a new organization, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscolntersightAccount

This deployment guide uses an example organization “AA01” throughout the document.

Claim Cisco UCS Fabric Interconnects in Cisco Intersight

Before claiming the UCS Fabric Interconnects in Intersight, make sure the initial configuration for the fabric interconnects has been completed. To claim the UCS Fabric Interconnects, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CiscolntersightAccount

Upgrade Fabric Interconnect Firmware using Cisco Intersight

Cisco UCS Manager does not support Cisco UCS X-Series therefore Fabric Interconnect software upgrade performed using UCS Manager does not contain the firmware for Cisco UCS X-series. If Cisco UCS Fabric Interconnects are being converted from UCSM to Intersight Managed Mode, before setting up UCS domain profile and discovering the chassis, upgrade the Fabric Interconnect firmware to release 4.2(3b) (outlined in Table 3) using Cisco Intersight by completing the steps in:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#UpgradeCiscoUCSFabricInterconnectFirmwareusingCiscoIntersight

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

Cisco Intersight Managed Mode – Domain Profile Setup

This chapter contains the following:

- [General Configuration](#)
- [UCS Domain Assignment](#)
- [VLAN and VSAN Configuration](#)
- [Port Configuration](#)
- [UCS Domain Configuration](#)
- [Review and Deploy the Domain Profile](#)
- [Configure Cisco UCS Chassis Profile \(optional\)](#)

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. 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.

Domain profile setup has the following steps:

- General configuration – name and organization assignment
- UCS Domain Assignment – assign previously claimed Cisco UCS Fabric Interconnects to the domain profile
- VLAN and VSAN configuration – define required VLANs and VSANs
- Port configuration – configure server and uplink ports and port-channels for Ethernet and FC traffic
- UCS domain configuration – policies such as NTP, DNS and QoS
- Review and deploy – review the configuration and deploy the UCS domain profile

General Configuration

To configure the name, description, and organization for the UCS domain profile, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#ConfigureaCiscoUCSDomainProfile

UCS Domain Assignment

To assign the Cisco UCS Fabric Interconnects to the UCS domain profile, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#ConfigureaCiscoUCSDomainProfile

VLAN and VSAN Configuration

To define the VLANs and VSANs, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#VLANandVSANConfiguration

The VLAN are explained in [Table 1](#). When the VLANs are successfully configured, Cisco Intersight displays a screen like [Figure 3](#).

Figure 3. VLANs used in UCS Domain Profile

VLANs

[Add VLANs](#)

Show VLAN Ranges

Add Filter [Export](#) 11 items found 50 per page 1 of 1

<input type="checkbox"/>	VLAN ...	Name	Sharing...	Multicast Policy	Auto All...	Primary VLAN ID	
<input type="checkbox"/>	2	Native-VLAN	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	1010	OOB-Mgmt_1010	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	1011	IB-Mgmt_1011	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	1012	VM-Traffic_1012	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	3001	vSAN-VLAN_3001	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	3002	Host_Overlay_Mgmt_3002	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	3003	Host_Overlay_WD_3003	None	AA01-Multicast-Policy	Yes		...
<input type="checkbox"/>	3030	vMotion-VLAN_3030	None	AA01-Multicast-Policy	Yes		...

Define two separate VSANs for the SAN-A and SAN-B paths as explained in the link above. In this document, VSAN 101 and 102 were defined for SAN-A and SAN-B, respectively. The VSANs are not required for the VMware Cloud Foundation management domain deployment but are used in FlashStack VI workload domain for boot from SAN configuration and for VMware Cloud Foundation principal storage.

Note: In this deployment, a single VLAN policy is shared by both Fabric Interconnects, but separate VSAN policies are defined for each Fabric Interconnect as shown in [Figure 4](#).

Figure 4. UCS Domain Profile VLAN and VSAN policy mapping

Policies

Port Configuration **VLAN & VSAN Configuration** UCS Domain Configuration

^ Fabric Interconnect A Configured

General Identifiers Connectivity

VLAN Configuration AA01-VLAN-Policy

VSAN Configuration AA01-VSAN-Policy-FI-A

^ Fabric Interconnect B Configured

General Identifiers Connectivity

VLAN Configuration AA01-VLAN-Policy

VSAN Configuration AA01-VSAN-Policy-FI-B

Port Configuration

To define the port roles and port-channels, complete the steps explained in Procedures 2 through 5 here: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#VLANandVSANConfiguration

In this deployment, various port roles and associated port-channels used to connect to different devices are shown in Figure 2, Figure 5, and Figure 6 show various port roles and associated port-channel numbers as defined in Cisco Intersight.

Figure 5. Cisco UCS Fabric Interconnect A port configuration



Figure 6. Cisco UCS Fabric Interconnect B port configuration



UCS Domain Configuration

To define the NTP server(s), DNS server(s), and to set the jumbo MTU for the best effort queue in QoS, complete the steps explained in Procedures 5 through 9 here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#VLANandVSANConfiguration

Review and Deploy the Domain Profile

To verify the configuration and to deploy the domain profile, complete the steps explained in Procedure 10 and 11 here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#VLANandVSANConfiguration

On successful deployment of the UCS domain profile, the ethernet port channels are enabled and the Cisco UCS rack servers and compute nodes are successfully discovered.

Figure 7. Discovered compute nodes and rack server example

<input type="checkbox"/>	Name	Health	Model
<input type="checkbox"/>	AA01-6454-1-2	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-1-3	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-1-4	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-1-6	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-1-7	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-1-8	Healthy	UCSX-210C-M6
<input type="checkbox"/>	AA01-6454-5	Healthy	UCSC-C240-M5L
<input type="checkbox"/>	AA01-6454-6	Healthy	UCSC-C240-M5L

Configure Cisco UCS Chassis Profile (optional)

Cisco UCS Chassis profile in Cisco Intersight allow customers to configure various parameters for chassis, including:

- IMC Access Policy: IP configuration for the in-band chassis connectivity. This setting is independent of Server IP connectivity and only applies to communication to and from chassis.
- SNMP Policy, and SNMP trap settings.

-
- Power Policy to enable power management and power supply redundancy mode.
 - Thermal Policy to control the speed of FANs.

A chassis policy can be assigned to any number of chassis profiles to provide a configuration baseline for a chassis. In this deployment, no chassis profile was created or attached but customers can configure some or all the policies and attach them to the chassis as needed. For more details on configuring UCS chassis policies, refer to:

https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide/b_intersight_managed_mode_guide_chapter_01100.html

Cisco Intersight Managed Mode – Server Profile Template

This chapter contains the following:

- [vNIC and vHBA Placement for Server Profile Templates](#)
- [Server Profile Template Creation](#)
- [Derive Management Domain Server Profile](#)
- [Derive VI Workload Domain Server Profile](#)

In Cisco Intersight Managed Mode, a server profile enables resource management by simplifying policy alignment and server configuration. The server profiles are derived from a server profile template. Server profile template and its associated policies can be created using the server profile template wizard.

In this deployment, two separate server profile templates are created for VMware Cloud Foundation management hosts and FlashStack VI workload domain hosts because of several differences in the two types of hosts. The two server profile templates both share certain elements such as UUID pools, management access policies, adapter policies etc. but have some unique configurations such as boot policy, BIOS policy and LAN/SAN connectivity policy.

Note: This chapter explains the configuration of both types of server profile templates. Customers can deploy one or both templates depending on their environment.

vNIC and vHBA Placement for Server Profile Templates

This section explains the vNIC and vHBA definitions and placement for both types of server profile templates.

Management Domain Host vNIC Placement

Four vNICs are configured and manually placed as listed in [Table 4](#).

Table 4. vNIC placement for Management Domain hosts

vNIC/vHBA Name	Slot	Switch ID	PCI Order
00-VDS01-A	MLOM	A	0
01-VDS01-B	MLOM	B	1
02-VDS02-A	MLOM	A	2
03-VDS02-B	MLOM	B	3

FlashStack VI Workload Domain Host vNIC and vHBA Placement

Four vNICs and two vHBAs are configured and manually placed as listed in [Table 5](#).

Table 5. vHBA and vNIC placement for FlashStack VI workload domain FC connected storage

vNIC/vHBA Name	Slot	Switch ID	PCI Order
00-VDS01-A	MLOM	A	0
01-VDS01-B	MLOM	B	1
02-VDS02-A	MLOM	A	2

vNIC/vHBA Name	Slot	Switch ID	PCI Order
03-VDS02-B	MLOM	B	3
vHBA-A	MLOM	A	4
vHBA-B	MLOM	B	5

Server Profile Template Creation

Following two server profiles templates will be configured for this deployment:

- Management Domain host template
- FlashStack VI workload domain template

Procedure 1. Configure a Server Profile Template

Step 1. Log in to the Cisco Intersight.

Step 2. Go to **Infrastructure Service > Configure > Templates** and in the main window click **Create UCS Server Profile Template**.

Procedure 2. General Configuration

Step 1. Select the organization from the drop-down list (for example, AA01).

Step 2. Provide a name for the server profile template. The names used in this deployment are:

- VCF-MgmtHost-Template (UCS C240 M5 management hosts)
- AA01-WD-FC-Boot-Template (FlashStack FC boot from SAN)

Step 3. Select **UCS Server (FI-Attached)**.

Step 4. Provide an optional description.

General

Enter a name, description, tag and select a platform for the server profile template.

Organization *
AA01

Name *
VCF-MgmtHost-Template

Target Platform UCS Server (Standalone) UCS Server (FI-Attached)

Set Tags

Description
VCF Mangement Hosts

<= 1024

Step 5. Click **Next**.

Procedure 3. Compute Configuration - UUID Pool

- Step 1.** Click **Select Pool** under UUID Pool 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, AA01) and provide a name for the UUID Pool (for example, AA01-UUID-Pool).
- Step 3.** Provide an optional Description and click **Next**.
- Step 4.** Provide a UUID Prefix (for example, a random prefix of AA010000-0000-0001 was used).
- Step 5.** Add a UUID block.

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *

AA010000-0000-0001 ⓘ

UUID Blocks

From	Size	
AA01-000000000001 ⓘ	50 ⓘ	+ 1 - 1024

- Step 6.** Click **Create**.

Procedure 4. Compute Configuration – BIOS policy

Note: Since the management hosts in this deployment are Cisco UCS C240 M5 servers while the VI workload domain servers are Cisco UCS X210c M6 servers, different BIOS policies will be created for each of the server profile templates.

- Step 1.** Click **Select Policy** next to BIOS and in the pane on the right, click **Create New**.
- Step 2.** Verify correct organization is selected from the drop-down list (for example, AA01) and provide a name for the policy (for example, AA01-M5-BIOS-Policy or AA01-M6-BIOS-Policy).
- Step 3.** Click **Next**.
- Step 4.** On the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on “Virtualization” workload recommendations in the performance tuning guide for Cisco UCS servers. Use the settings listed below:

Procedure 5. Configure UCS M6 Server BIOS Policy

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

- Step 1.** Set the parameters below and leave all other parameters set to “platform-default.”
 - Memory > NVM Performance Setting: Balanced Profile
 - Power and Performance > Enhanced CPU Performance: Auto
 - Processor > Energy Efficient Turbo: enabled

- Processor > Processor C1E: enabled
- Processor > Processor C6 Report: enabled
- Server Management > Consistent Device Naming: enabled

Procedure 6. Configure UCS M5 Server BIOS Policy

For detailed information, see: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/white-paper-c11-744678.html>

Step 1. Set the parameters below and leave all other parameters set to “platform-default.”

- Memory > NVM Performance Setting: Balanced Profile
- Processor > Power Technology: custom
- Processor > Processor C1E: disabled
- Processor > Processor C3 Report: disabled
- Processor > Processor C6 Report: disabled
- Processor > CPU C State: disabled
- Server Management > Consistent Device Naming: enabled

Step 2. Click **Create**.

Procedure 7. Compute Configuration - Boot Order policy for Management Domain hosts

Note: Management hosts are equipped with Cisco UCS Boot Optimized M.2 drive where ESXi will be installed for local boot. The policy explained below may need to be adjusted if customers have a different hard disk configuration or boot drive. The FC boot order policy for the VI workload domain hosts is different and is explained in the next procedure.

Step 1. Click **Select Policy** next to BIOS Configuration 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, AA01) and provide a name for the policy (for example, Local-BootOrder-Pol).

Step 3. Click **Next**.

Step 4. For Configured Boot Mode option, select **Unified Extensible Firmware Interface (UEFI)**.

Step 5. Turn on **Enable Secure Boot**.

Policy Details

Add policy details


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

Configured Boot Mode 

Unified Extensible Firmware Interface (UEFI) Legacy

Enable Secure Boot 

Add Boot Device 

Step 6. Click **Add Boot Device** drop-down list and select **Virtual Media**.

Step 7. Provide a device name (for example, KVM-Mapped-ISO) and then, for the subtype, select **KVM Mapped DVD**.

— Virtual Media (KVM-Mapped-ISO) Enabled 🗑️ ^ v

Device Name *
KVM-Mapped-ISO ⊙

Sub-Type
KVM MAPPED DVD v ⊙

Step 8. From the **Add Boot Device** drop-down list, select **Local Disk**.

Step 9. Provide the Device Name (for example Local-Boot).

— Local Disk (Local-Boot) Enabled 🗑️ ^ v

Device Name *
Local-Boot ⊙ Slot ⊙

Bootloader Name ⊙ Bootloader Description ⊙

Bootloader Path ⊙

Step 10. Verify the order of the boot policies and adjust the boot order, as necessary.

Add Boot Device v

+ Virtual Media (KVM-Mapped-ISO) Enabled 🗑️ ^ v

+ Local Disk (Local-Boot) Enabled 🗑️ ^ v

Step 11. Click **Create**.

Procedure 8. Compute Configuration - Boot Order policy for VI Workload Domain hosts

Step 1. Click **Select Policy** next to BIOS Configuration 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, AA01) and provide a name for the policy (for example, FC-BootOrder-Pol).

Step 3. Click **Next**.


Step 4. For Configured Boot Mode option, select **Unified Extensible Firmware Interface (UEFI)**.

Step 5. Turn on Enable **Secure Boot**.

Policy Details

Add policy details

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

Configured Boot Mode 

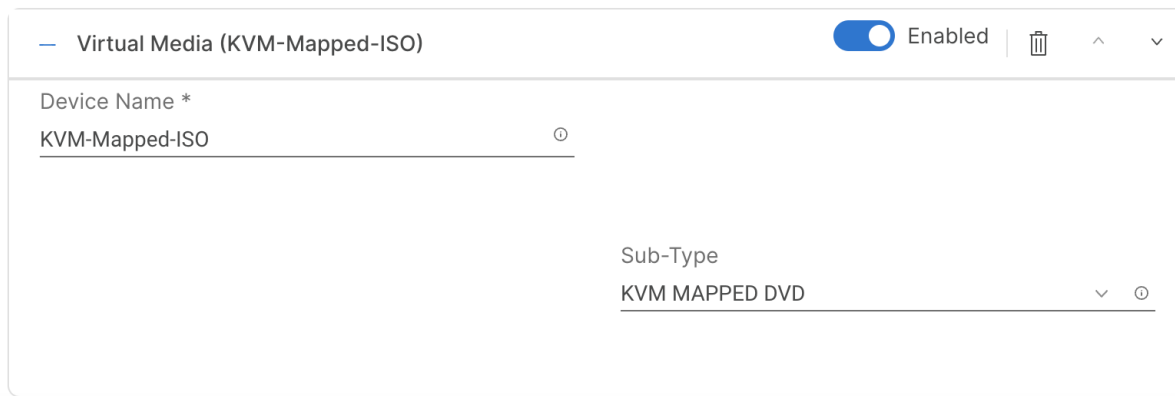
Unified Extensible Firmware Interface (UEFI) Legacy

Enable Secure Boot 

[Add Boot Device](#) 

Step 6. Click **Add Boot Device** drop-down list and select **Virtual Media**.

Step 7. Provide a device name (for example, KVM-Mapped-ISO) and then, for the subtype, select **KVM Mapped DVD**.



The screenshot shows a configuration form for a boot device. At the top, it says "Virtual Media (KVM-Mapped-ISO)" with an "Enabled" toggle switch and a trash icon. Below this, there are two input fields: "Device Name *" with the value "KVM-Mapped-ISO" and "Sub-Type" with the value "KVM MAPPED DVD".

For Fibre Channel SAN boot, all four Pure Storage controller FC interfaces will be added as boot options. The four targets are named as follows:

- **FCP-X50R3-CT0-FC0:** Pure Storage Controller 1, FC port 0 for Fibre Channel SAN A
- **FCP-X50R3-CT1-FC0:** Pure Storage Controller 2, FC port 0 for Fibre Channel SAN A
- **FCP-X50R3-CT0-FC1:** Pure Storage Controller 1, FC port 1 for Fibre Channel SAN B
- **FCP-X50R3-CT1-FC1:** Pure Storage Controller 2, FC port 1 for Fibre Channel SAN B

Step 8. From the Add Boot Device drop-down list, select **SAN Boot**.

Step 9. Provide the Device Name: FCP-X50R3-CT0-FC0 and the Logical Unit Number (LUN) value (for example, 1).

Step 10. Provide an interface name (e.g., FCP-Fabric-A or vHBA-B). This value is important and should match the appropriate vHBA name for SAN-A or SAN-B.

Note: FCP-Fabric-A is used to access FCP-X50R3-CT0-FC0 and FCP-X50R3-CT1-FC0 and FCP-Fabric-B is used to access FCP-X50R3-CT0-FC1 and FCP-X50R3-CT1-FC1.

Step 11. Add the appropriate World Wide Port Name (WWPN) of Pure Storage FC port as the Target WWPN.

— SAN Boot (FCP-X50R3-CT0-FC0)

 Enabled
 🗑️
^
v

Device Name *	LUN	
FCP-X50R3-CT0-FC0	1	0 - 255
Slot	Interface Name *	
	FCP-Fabric-A	
Target WWPN *		
52:4a:93:7f:8b:d2:75:00		
Bootloader Name	Bootloader Description	
Bootloader Path		

Step 12. Repeat steps 8-11 three more times to add all the remaining Pure Storage targets.

Step 13. Verify the order of the boot policies and adjust the boot order as necessary using arrows next to delete button.

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

Configured Boot Mode ⊙

Unified Extensible Firmware Interface (UEFI)
 Legacy

Enable Secure Boot ⊙

Add Boot Device v

+ Virtual Media (KVM-Mapped-ISO)

 Enabled
 🗑️
^
v

+ SAN Boot (FCP-X50R3-CT0-FC0)

 Enabled
 🗑️
^
v

+ SAN Boot (FCP-X50R3-CT1-FC0)

 Enabled
 🗑️
^
v

+ SAN Boot (FCP-X50R3-CT0-FC1)

 Enabled
 🗑️
^
v

+ SAN Boot (FCP-X50R3-CT1-FC1)

 Enabled
 🗑️
^
v

Step 14. Click **Create**.

Procedure 9. Compute Configuration – Configure Virtual Media Policy

This procedure enables you to configure the Virtual Media Policy to allow mapping an ISO file as installation source for operating system.

Step 1. Click **Select Policy** next to Virtual Media 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, AA01) and provide a name for the policy (for example, AA01-vMedia-Policy).

Step 3. Turn on **Enable Virtual Media**, **Enable Virtual Media Encryption**, and **Enable Low Power USB**.

Step 4. Do not Add Virtual Media at this time.


Policy Details


Add policy details

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

Configuration

Enable Virtual Media 

Enable Virtual Media Encryption 

Enable Low Power USB 

[Add Virtual Media](#)



0 items found

26  per page   0 of 0  



Name


Type

Protocol

File Location

NO ITEMS AVAILABLE



  0 of 0  

Step 5. Click **Create**.

Step 6. Click **Next** to move to Management Configuration.

Management Configuration

The following four policies will be added to the management configuration:

- IMC Access to define the pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM
- Virtual KVM to allow the Tunneled KVM

Procedure 1. Management Configuration - Cisco IMC Access Policy

Step 1. Click **Select Policy** next to IMC Access 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, AA01) and provide a name for the policy (for example, AA01-IMC-Access-Policy).

Step 3. Click **Next**.


Note: Customers can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 1011) or out-of-band management access via the Mgmt0 interfaces of the FIs. In-band management access was configured in this deployment guide.

Step 4. Enable **In-Band Configuration** and provide the in-band management VLAN (for example, 1011).

Step 5. Make sure **IPv4 address configuration** is selected.

Policy Details

Add policy details

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

! A minimum of one configuration must be enabled. Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Check here for more info, [Help Centre](#)

In-Band Configuration 

Enabled

VLAN ID *

1011   
4 - 4093

IPv4 address configuration 

IPv6 address configuration 

IP Pool *

[Select IP Pool](#) 

Out-Of-Band Configuration 

Enabled

Step 6. Under IP Pool, click **Select IP Pool** and then, in the pane on the right, click **Create New**.

Step 7. Verify correct organization is selected from the drop-down list (for example, AA01) and provide a name for the pool (for example, AA01-Mgmt-IP-Pool).

Step 8. Select **Configure IPv4 Pool** and provide the information to define a pool for KVM IP address assignment including an IP Block.

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Previously saved parameters cannot be changed. You can find Cisco recommendations at [Help Center](#).

Configuration

Netmask *	Gateway
255.255.255.0	10.101.1.254
Primary DNS	Secondary DNS
172.20.4.53	172.20.4.54

IP Blocks

From	Size	
10.101.1.201	10	1 - 1024

Note: The management IP pool subnet should be routable from the host that is trying to access the KVM session. In the example shown here, the hosts trying to establish an KVM connection would need to be able to route to 10.101.1.0/24 subnet.

Step 9. Click **Next**.

Step 10. Unselect **Configure IPv6 Pool**.

Step 11. Click **Create** to finish configuring the IP address pool.

Step 12. Click **Create** to finish configuring the IMC access policy.

Procedure 2. Management Configuration - IPMI Over LAN policy

Step 1. Click **Select Policy** next to IPMI Over LAN 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, AA01) and provide a name for the policy (for example, AA01-Enable-IPMIoLAN-Policy).

Step 3. Turn on **Enable IPMI Over LAN**.

Step 4. Click **Create**.

Policy Details

Add policy details

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

Enable IPMI Over LAN

Procedure 3. Management Configuration - Local User policy

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

Step 2. Verify correct organization is selected from the drop-down list (for example, AA01) and provide a name for the policy (for example, AA01-LocalUser-Pol).

Step 3. Verify that **UCS Server (FI-Attached)** is selected.

Step 4. Verify that **Enforce Strong Password** is selected.


Policy Details

Add policy details

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

Password Properties

Enforce Strong Password 

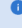
Enable Password Expiry 

Password History

5   
0 - 5

Always Send User Password 

Local Users

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

[Add New User](#)

Step 5. Click **Add New User** and then click **+** next to the New User

Step 6. Provide the username (for example, fsadmin), select a role (for example, admin), and provide a password.

Note: The username and password combination defined here can be used to log into KVMs as well as for IPMI access. The default admin user and password also allow customers to log into KVM.

Step 7. Click **Create** to finish configuring the user.

Step 8. Click **Create** to finish configuring local user policy.

Step 9. Click **Next** to move to Storage Configuration.

Procedure 4. Management Configuration - Virtual KVM Policy

Step 1. Click **Select Policy** next to Virtual KVM 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, AA01) and provide a name for the policy (for example, AA01-KVM-Policy).

Step 3. Verify that **UCS Server (FI-Attached)** is selected.

Step 4. Turn on **Allow Tunneled vKVM** and leave the other two options on as well.

Policy Details

Add policy details



All Platforms

UCS Server (Standalone)

UCS Server (FI-Attached)

Enable Virtual KVM

Max Sessions *

4
1 - 4

Enable Video Encryption

Allow Tunneled vKVM

Step 5. Click **Create**.

Note: To enable Tunneled KVM, make sure under **System > Settings > Security and Privacy>Configure**, “Allow Tunneled vKVM Launch” and “Allow Tunneled vKVM Configuration” is turned on.

Configure Security & Privacy Settings

^ Data Collection

Allow Tech Support Bundle Collection

• If Tech Support Bundle Collection is disallowed, the tech support bundle collection is not possible and Support Case Manager and Proactive RMA cannot perform properly. Learn more at [Help Center](#).

^ Connection to Intersight

Allow Tunneled vKVM Launch

• Allows Tunneled vKVM launch for all the setups claimed to the account. Learn more at [Help Center](#).

Allow Tunneled vKVM Configuration

• Allows configuration of Tunneled vKVM for all the setups claimed to the account. Learn more at [Help Center](#).

Step 6. Click **Next** to move to Storage Configuration.

Procedure 5. Storage Configuration

The Cisco UCS C240 M5 management hosts used in this deployment contain:

- A single M.2 drive for ESXi installation
- An SSD drive for caching tier
- Multiple HDDs for capacity tier

No special configuration (such as RAID) is needed for the M.2 drive and all the SSDs and HDDs are presented to operating system in JBOD configuration. VMware vSAN configures the caching and capacity disks as needed for vSAN setup. [Figure 8](#) shows a sample SSD/HDD configuration used in the validation environment. The RAID controller, SSD and HDD models are all certified by VMware for vSAN configuration.

Figure 8. Example SSD/HDD layout (lab host)

General		Physical Drives	Virtual Drives						
...									
<input type="checkbox"/>	Name	Disk Firmw...	Size (MiB)	Model	Vendor	Protocol	Type	Drive State	
<input type="checkbox"/>	Disk 1	A3Z4	7630328	UCS-HD8T7KL4KN	HGST	SAS	HDD	Jbod	...
<input type="checkbox"/>	Disk 2	A3Z4	7630328	UCS-HD8T7KL4KN	HGST	SAS	HDD	Jbod	...
<input type="checkbox"/>	Disk 3	A3Z4	7630328	UCS-HD8T7KL4KN	HGST	SAS	HDD	Jbod	...
<input type="checkbox"/>	Disk 4	A3Z4	7630328	UCS-HD8T7KL4KN	HGST	SAS	HDD	Jbod	...
<input type="checkbox"/>	Disk 14	0104	3051757	UCS-SD32T123X-EP	TOSHIBA	SAS	SSD	Jbod	...

Step 1. Click **Next** on the Storage Configuration screen to proceed to Network Configuration. No configuration is needed for the local storage system.

Network Configuration

Network configuration explains both LAN and SAN connectivity policies.

Procedure 1. Network Configuration - LAN Connectivity

LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For consistent vNIC and vHBA placement, manual vHBA/vNIC placement is utilized.

Note: Two separate LAN connectivity policies should be configured: one for management domain hosts and one for VI workload domain hosts.

The Management Domain hosts, and FlashStack VI workload domain hosts each use 4 vNICs configured as shown in [Table 6](#).

Table 6. vNICs for setting up LAN Connectivity Policy

vNIC/vHBA Name	Slot	Switch ID	PCI Order
00-VDS01-A	MLOM	A	0
01-VDS01-B	MLOM	B	1
02-VDS02-A	MLOM	A	2
03-VDS02-B	MLOM	B	3

Step 1. Click **Select Policy** next to LAN 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, AA01) and provide a name for the policy (for example, AA01-MgmtHost-LanConn-Pol or AA01-VI-FC-LanConn-Pol). Click **Next**.
- Step 3.** Under vNIC Configuration, select **Manual vNICs Placement**.
- Step 4.** Click Add vNIC.

vNIC Configuration

Manual vNICs Placement

Auto vNICs Placement

i For manual placement option you need to specify placement for each vNIC. Learn more at

Help Center

Add vNIC

Graphic vNICs Editor

Procedure 2. Network Configuration – LAN Connectivity – Define MAC Pool for Fabric Interconnects A and B

Note: If the MAC address pool has not been defined yet, when creating the first vNIC new MAC address pools will need to be created. Two separate MAC address pools are configured: MAC-Pool-A will be used for all Fabric-A vNICs, and MAC-Pool-B will be used for all Fabric-B vNICs.

Table 7. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs
MAC-Pool-A	00:25:B5:A1:0A:00	256*	00-VDS01-A, 02-VDS02-A
MAC-Pool-B	00:25:B5:A1:0B:00	256*	01-VDS01-B, 03-VDS02-B

Note: Each server requires 2 MAC addresses from each pool. Adjust the size of the pool according to your requirements. “A1” in the MAC address pool above is a unique identifier representing the rack ID while 0A/0B identifies the Fabric A or Fabric B. Adding a unique identifier help with troubleshooting of switching issues.

- Step 1.** Click **Select Pool** under MAC Address Pool 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, AA01) and provide a name for the pool from Table 7 depending on the vNIC being created (for example, AA01-MAC-Pool-A for Fabric A vNICs and AA01-MAC-Pool-B for Fabric B vNICs).
- Step 3.** Click **Next**.
- Step 4.** Provide the starting MAC address from [Table 7](#) (for example, 00:25:B5:A1:0A:00).
- Step 5.** Provide the size of the MAC address pool from [Table 7](#) (for example, 256).

Pool Details

Collection of MAC Blocks.

MAC Blocks

From	Size	
00:25:B5:A1:0A:00	256	1 - 1024

Step 6. Click **Create** to finish creating the MAC address pool.

Step 7. From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information from [Table 6](#).

General

Name *

00-VDS01-A

Pin Group Name

MAC

Pool

Static

MAC Pool *

Selected Pool AA01-Mac-Pool-A | | |

Placement

Simple

Advanced

Slot ID *

MLOM

PCI Link

0

0 - 1

Switch ID *

A

PCI Order

0

Step 8. For Consistent Device Naming (CDN), from the drop-down list, select **vNIC Name**.

Step 9. Verify that **Failover** is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.

Consistent Device Naming (CDN)

Source

vNIC Name ▼ ⓘ

Failover

Enabled ⓘ

Procedure 3. Network Configuration – LAN Connectivity – Define Ethernet Network Group Policy for a vNIC

Ethernet Network Group policies are created and reused on applicable vNICs as explained below. Ethernet network group policy defines the VLANs allowed for a particular vNIC therefore multiple network group policies will be defined as follows:

Table 8. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs
Mgmt-VDS01-NetGrp	Native-VLAN (2)	00-VDS01-A, 01-VDS01-B	OOB-MGMT*, IB-MGMT, vSAN, vMotion
Mgmt-VDS02-NetGrp	Native-VLAN (2)	02-VDS02-A, 03-VDS02-B	Mgmt-Host-Overlay
WD-VDS01-NetGrp	Native-VLAN (2)	00-VDS01-A, 01-VDS01-B	OOB-MGMT*, IB-MGMT
WD-VDS02-NetGrp	Native-VLAN (2)	02-VDS02-A, 03-VDS02-B	WD-Host-Overlay, vMotion, VM-Traffic

Note: * Adding Out of Band Management VLAN is optional and depends on customer networking requirements.

Step 1. Click **Select Policy** under Ethernet Network Group Policy 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, AA01) and provide a name for the policy from the Table 8 (for example, Mgmt-VDS01-NetGrp).

Step 3. Click **Next**.

Step 4. Enter the **Allowed VLANs** and **Native VLAN** from the [Table 8](#).

Policy Details

Add policy details

VLAN Settings

Allowed VLANs

1010,1011,3001,3030 ⓘ

Native VLAN

2 ⌵ ⓘ

1 - 4093

Step 5. Click **Create** to finish configuring the Ethernet network group policy.

Note: When ethernet group policies are shared between two vNICs, the ethernet group policy only needs to be defined for the first vNIC. For subsequent vNIC policy mapping, just click **Select Policy** and pick the previously defined ethernet group policy from the list.

Procedure 4. Network Configuration – LAN Connectivity – Create Ethernet Network Control Policy

Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created here and reused for all the vNICs.

Step 1. Click **Select Policy** under Ethernet Network Control Policy 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, AA01) and provide a name for the policy (for example, AA01-Enable-CDP-LLDP).

Step 3. Click **Next**.

Step 4. Enable **Cisco Discovery Protocol** and both **Enable Transmit** and **Enable Receive** under LLDP.

Policy Details

Add policy details

 This policy is applicable only for UCS Servers (FI-Attached)


Enable CDP 

Mac Register Mode 

Only Native VLAN All Host VLANs

Action on Uplink Fail 

Link Down Warning

 Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.

MAC Security

Forge 

Allow Deny

LLDP

Enable Transmit 

Enable Receive 

Step 5. Click **Create** to finish creating Ethernet network control policy.

Procedure 5. Network Configuration – LAN Connectivity – Create Ethernet QoS Policy

Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for the vNICs. A single policy will be created and reused for all the vNICs.

Step 1. Click **Select Policy** under Ethernet 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, AA01) and provide a name for the policy (for example, AA01-EthernetQos-Pol).

Step 3. Click **Next**.











Step 4. Change the MTU, Bytes value to 9000.

Policy Details

Add policy details


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

QoS Settings

<p>MTU, Bytes</p> <p>9000  </p> <hr/> <p>1500 - 9000</p>	<p>Rate Limit, Mbps</p> <p>0  </p> <hr/> <p>0 - 100000</p>
<p>Class of Service</p> <p>0  </p> <hr/> <p>0 - 6</p>	<p>Burst</p> <p>10240  </p> <hr/> <p>1 - 1000000</p>
<p>Priority</p> <p>Best-effort  </p>	

Enable Trust Host CoS 

Step 5. Click **Create** to finish setting up the Ethernet QoS policy.

Procedure 6. Network Configuration - LAN Connectivity - Create Ethernet Adapter Policy

An Ethernet adapter policy is used to set the interrupts and configure the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

Customers can also configure a modified ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, AA01-EthAdapter-HighTraffic-Policy, is created and attached to the 00-VDS01-A and 01-VDS01-B interfaces on management domain hosts and 02-VDS02-A and 03-VDS02-B interfaces on VI workload domain hosts which handle vMotion.

Table 9. Ethernet Adapter Policy association to vNICs

Host Type	Policy Name	Apply to vNICs	Description
Management Domain Host	AA01-EthAdapter-HighTraffic-Policy	00-VDS01-A, 01-VDS01-B	Support vMotion
Management Domain Host	AA01-EthAdapter-VMware-Policy	02-VDS02-A, 03-VDS02-B	Application Traffic

Host Type	Policy Name	Apply to vNICs	Description
VI Workload Domain	AA01-EthAdapter-VMware-Policy	00-VDS01-A, 01-VDS01-B	Management
VI Workload Domain	AA01-EthAdapter-HighTraffic-Policy	02-VDS02-A, 03-VDS02-B	Support vMotion

Step 1. Click **Select Policy** under Ethernet Adapter 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, AA01) and provide a name for the policy (for example, AA01-EthAdapter-VMware-Policy).

Step 3. Click **Select Default Configuration** under Ethernet Adapter Default Configuration.

General

Add a name, description and tag for the policy.

Organization *

AA01 ▼

Name *

AA01-EthAdapter-VMware-Policy

Set Tags

Description ⌵

<= 1024

Ethernet Adapter Default Configuration ⊙

Select Default Configuration 

Step 4. From the list, select **VMware**.

Step 5. Click **Next**.

Step 6. For the AA01-EthAdapter-VMware policy, click **Create** and skip the rest of the steps in this “Create Ethernet Adapter Policy” section.

Step 7. For the AA01-VMware-High-Traffic policy, make the following modifications to the policy:

- Increase Interrupts to 11
- Increase Receive Queue Count to 8
- Increase Completion Queue Count to 9
- Enable Receive Side Scaling
- Set Receive Ring Size and Transmit Ring Size to 4096

Interrupt Settings

Interrupts: 1 - 1024
 Interrupt Mode: MSIx
 Interrupt Timer, us: 0 - 65535

Interrupt Coalescing Type: Min

Receive

Receive Queue Count: 1 - 1000
 Receive Ring Size: 64 - 16384

Transmit

Transmit Queue Count: 1 - 1000
 Transmit Ring Size: 64 - 16384

Completion

Completion Queue Count: 1 - 2000
 Completion Ring Size: 1 - 256

Uplink Failback Timeout (seconds):

Receive Side Scaling

Enable Receive Side Scaling

Step 8. Click **Create**.

Step 9. Click **Add** to add the vNIC to the LAN connectivity policy.

Step 10. Go back to step 4 [Add vNIC](#) and repeat vNIC creation for all four vNICs.

Step 11. Verify all four vNICs were successfully created for appropriate LAN connectivity Policy.

<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="text" value="Add Filter"/> <input type="button" value="4 items found"/> <input type="text" value="50"/> per page <input type="button" value="1"/> of 1 <input type="button" value="⚙️"/>							
<input type="checkbox"/>	Name	Slo...	Switch ID	PCI Order	Failover	MAC Pool	<input type="button" value="⚡"/>
<input type="checkbox"/>	00-VDS01-A	MLOM	A	0	Disabled	AA01-Mac-Pool-A	<input type="button" value="⋮"/>
<input type="checkbox"/>	02-VDS02-A	MLOM	A	2	Disabled	AA01-Mac-Pool-A	<input type="button" value="⋮"/>
<input type="checkbox"/>	01-VDS01-B	MLOM	B	1	Disabled	AA01-Mac-Pool-B	<input type="button" value="⋮"/>
<input type="checkbox"/>	03-VDS02-B	MLOM	B	3	Disabled	AA01-Mac-Pool-B	<input type="button" value="⋮"/>

of 1

Step 12. Click **Create** to finish creating the LAN Connectivity policy.

Procedure 7. Network Connectivity – Create SAN Connectivity Policy (only for VI workload domain)

A SAN connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables customers to configure the vHBAs that the servers use to communicate with the SAN. [Table 10](#) lists the details of two vHBAs that are used to provide FC connectivity and boot from SAN functionality.

Note: SAN Connectivity policy is not needed for management domain hosts and can be skipped when configuring the Server Profile Template for the management hosts.

Table 10. vHBAs for FlashStack VI workload domain (boot from FC)

vNIC/vHBA Name	Slot	Switch ID	PCI Order
FCP-Fabric-A	MLOM	A	4
FCP-Fabric-B	MLOM	B	5

Step 1. Click **Select Policy** next to SAN 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, AA01) and provide a name for the policy (for example, VCF-VIC-SAN-Connectivity-Policy).

Step 3. Select **Manual vHBAs Placement**.

Step 4. Select **Pool** under WWNN Address.

Policy Details

Add policy details

Manual vHBAs Placement

Auto vHBAs Placement

WWNN

Pool

Static

Procedure 8. Network Connectivity – SAN Connectivity – WWNN Pool

If the WWNN address pools have not been previously defined, a new WWNN address pool must be defined when adding the SAN connectivity policy.

Step 1. Click **Select Pool** under WWNN Address Pool 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, AA01) and provide a name for the policy (for example, AA01-WWNN-Pool).

Step 3. Click **Next**.

Step 4. Provide the starting WWNN block address and the size of the pool.

Pool Details

Block of WWNN Identifiers.

WWNN Blocks	
From	Size
20:00:00:25:B5:A1:00:00	32
1 - 1024	

Note: As a best practice, some additional information is always encoded into the WWNN address pool for ease of troubleshooting. For example, in the address 20:00:00:25:B5:A1:00:00, A1 is the rack ID.

Step 5. Click **Create** to finish creating the WWNN address pool.

Procedure 9. Network Connectivity - SAN Connectivity - Create vHBA for SAN A

Step 1. Click **Add vHBA**.

Step 2. For vHBA Type, select **fc-initiator** from the drop-down list.

Procedure 10. Network Connectivity - SAN Connectivity - WWPN Pool for SAN A

If the WWPN address pool has not been previously defined, a new WWPN address pool for Fabric A must be defined when adding a vHBA.

Step 1. Click **Select Pool** under WWPN Address Pool 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, AA01) and provide a name for the policy (for example, VCF-WWPN-Pool-A).

Step 3. Provide the starting WWPN block address for SAN A and the size of the pool.

Note: As a best practice, in FlashStack some additional information is always encoded into the WWPN address pool for ease of troubleshooting. For example, in the address 20:00:00:25:B5:A1:0A:00, A1 is the rack ID and 0A signifies SAN A.

Pool Details

Block of WWPN Identifiers.

WWPN Blocks	
From	Size
20:00:00:25:B5:A1:0A:00	32
1 - 1024	

Step 4. Click **Create** to finish creating the WWPN pool.

Step 5. Back in the Create vHBA window, provide the Name (for example, FCP-Fabric-A), select **Advanced** under placement option, and add Slot ID (for example, MLOM), Switch ID (for example, A) and PCI Order from [Table 10](#).

General

Name *	FCP-Fabric-A	vHBA Type	fc-initiator
Pin Group Name			

WWPN

Pool Static

WWPN Pool *
 Selected Pool VCF-WWPN-Pool-A

Placement

Simple Advanced

When Simple Placement is selected, the Slot ID and PCI Link are automatically determined by the system. vHBAs are deployed on the first VIC. The Slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. Simple assignment is not applicable for 13xx series VICs that support dual-link.

Switch ID *
 A

PCI Order
 4

Procedure 11. Network Connectivity - SAN Connectivity - Fibre Channel Network for SAN A

A Fibre Channel network policy governs the VSAN configuration for the virtual interfaces. In this deployment, VSAN 101 is used for FCP-Fabric-A.

Step 1. Click **Select Policy** under Fibre Channel Network 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, AA01) and provide a name for the policy (for example, AA01-FC-Network-SAN-A).

Step 3. For the scope, make sure **UCS Server (FI-Attached)** is selected.

Step 4. Under VSAN ID, provide the VSAN information (for example, 101).

Policy Details

Add policy details

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

Fibre Channel Network

VSAN ID
 101

1 - 4094

Step 5. Click **Create** to finish creating the Fibre Channel network policy.

Procedure 12. Network Connectivity – SAN Connectivity – Fibre Channel QoS

The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. The Fibre Channel QoS policy used in this deployment uses default values and will be shared by all vHBAs.

Step 1. Click **Select Policy** under Fibre Channel QoS 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, AA01) and provide a name for the policy (for example, AA01-FC-QoS-Policy).

Step 3. For the scope, select **UCS Server (FI-Attached)**.

Step 4. Do not change the default values on the Policy Details screen.

Step 5. Click **Create** to finish creating the Fibre Channel QoS policy.

Procedure 13. Network Connectivity – SAN Connectivity – Fibre Channel Adapter

A Fibre Channel adapter policy governs the host-side behavior of the adapter, including the way that the adapter handles traffic. This validation uses the default values for the adapter policy, and the policy will be shared by all the vHBAs.

Step 1. Click **Select Policy** under Fibre Channel Adapter 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, AA01) and provide a name for the policy (for example, AA01-FC-Adapter-Policy).

Step 3. For the scope, select **UCS Server (FI-Attached)**.

Step 4. Do not change the default values on the Policy Details screen.

Step 5. Click **Create** to finish creating the Fibre Channel adapter policy.

Step 6. Click **Add** to create FCP-Fabric-A.

Procedure 14. Network Connectivity – SAN Connectivity – Add vHBA-B for SAN B

Step 1. Click **Add vHBA**.

Step 2. For vHBA Type, select **fc-initiator** from the drop-down list.

Procedure 15. Network Connectivity – SAN Connectivity – WWPN Pool for SAN B

If the WWPN address pool has not been previously defined, a WWPN address pool for Fabric B must be defined for vHBA-B.

Step 1. Click **Select Pool** under WWPN Address Pool 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, AA01) and provide a name for the policy (for example, VCF-WWPN-Pool-B).

Step 3. Provide the starting WWPN block address for SAN B and the size of the pool.

Note: As a best practice, in FlashStack some additional information is always encoded into the WWPN address pool for ease of troubleshooting. For example, in the address 20:00:00:25:B5:A1:0B:00, A1 is the rack ID and 0B signifies SAN B.

Pool Details

Block of WWPN Identifiers.

WWPN Blocks

From	Size	
20:00:00:25:B5:A1:0B:00	32	1 - 1024

Step 4. Click **Create** to finish creating the WWPN pool.

Step 5. Back in the Create vHBA window, provide the Name (for example, FCP-Fabric-B), select **Advanced** under placement option, and add Slot ID (for example, MLOM), Switch ID (for example, B) and PCI Order from [Table 10](#).

General

Name *	FCP-Fabric-B	vHBA Type	fc-initiator
Pin Group Name			

WWPN

Pool Static

WWPN Pool *

Selected Pool VCF-WWPN-Pool-B

Placement

Simple Advanced

- When Simple Placement is selected, the Slot ID and PCI Link are automatically determined by the system. vHBAs are deployed on the first VIC. The Slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. Simple assignment is not applicable for 13xx series VICs that support dual-link.

Switch ID *
B

PCI Order
5

Procedure 16. Network Connectivity - SAN Connectivity - Fibre Channel Network for SAN B

In this deployment, VSAN 102 will be used for FCP-Fabric-B.

Step 1. Click **Select Policy** under Fibre Channel Network 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, AA01) and provide a name for the policy (for example, AA01-FC-Network-SAN-B).

Step 3. For the scope, select **UCS Server (FI-Attached)**.

Step 4. Under VSAN ID, provide the VSAN information (for example, 102).

Policy Details

Add policy details

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

Fibre Channel Network

VSAN ID

102

1 - 4094

Step 5. Click **Create**.

Procedure 17. Network Connectivity – SAN Connectivity – Fibre Channel QoS

Step 1. Click **Select Policy** under Fibre Channel QoS and then, in the pane on the right, select the previously created QoS policy AA01-FC-QoS.

Procedure 18. Network Connectivity – SAN Connectivity – Fibre Channel Adapter

Step 1. Click **Select Policy** under Fibre Channel Adapter and then, in the pane on the right, select the previously created Adapter policy AA01-FC-Adapter.

Step 2. Verify all the vHBA policies are mapped.

Persistent LUN Bindings

Persistent LUN Bindings 

Fibre Channel Network * 

Selected Policy AA01-FC-Network-SAN-B |  |  | 

Fibre Channel QoS * 

Selected Policy AA01-FC-QoS-Policy |  |  | 

Fibre Channel Adapter * 

Selected Policy AA01-FC-Adapter-Policy |  |  | 

FC Zone 

[Select Policy\(s\)](#) 

Step 3. Click **Add** to add the FCP-Fabric-B.

Step 4. Verify both the vHBAs are added to the SAN connectivity policy.

Add vHBA

Graphic vHBAs Editor

🗑️ ✎ 📄 | 🔍 Add Filter 📄 2 items found 21 per page ⏪ ⏩ 1 of 1 ⏪ ⏩ ⚙️

<input type="checkbox"/>	Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN P...	⚡
<input type="checkbox"/>	FCP-Fabric-A	Auto	A	4	-	VCF-WWPN-P...	...
<input type="checkbox"/>	FCP-Fabric-B	Auto	B	5	-	VCF-WWPN-P...	...

🗑️ ✎ 📄 ⏪ ⏩ 1 of 1 ⏪ ⏩

Step 5. Click **Create** to finish creating SAN connectivity policy.

Step 6. When the LAN connectivity policy and SAN connectivity policy (for FC) is created, click **Next** to move to the Summary screen.

Procedure 1. Summary verification

Step 1. On the summary screen, verify policies mapped to various settings.

Step 2. Click **Close** to finish Server Profile Template creation.

Note: Remember to create both management domain host and VI workload domain host Server Profile Templates using the [Server Profile Templates Creation](#) procedure.

Derive Management Domain Server Profile

Procedure 1. Derive One or more Server Profiles

Step 1. From the **Infrastructure Services > Configure > Templates**, click “...” next to the management host template name and select **Derive Profiles**.

Step 2. Under the Server Assignment, select **Assign Now** and pick four Cisco UCS C240 M5 racks servers. Customers can adjust the number of servers depending on the number of profiles to be deployed.

Note: In this deployment **four** (minimum) management domain hosts will be derived. Only two out of four servers shown in screen capture below.

Server Assignment

Assign Now Assign Server from a Resource Pool Assign Later

🔍 Add Filter 📄 8 items found

<input type="checkbox"/>	Name	Model	UCS Domain
<input checked="" type="checkbox"/>	AA01-6454-5	UCSC-C240-M5L	AA01-6454
<input checked="" type="checkbox"/>	AA01-6454-6	UCSC-C240-M5L	AA01-6454

Step 3. Click **Next**.

Step 4. Cisco Intersight will fill in “default” information for the selected servers (only two out of four servers shown):

Derive

Profile Name Prefix	Digits Count	Start Index for Suffix
VCF-MgmtDomHost-Template_DERIVED-	1	1
	>= 1	>= 0
<hr/>		
1 Name *	Assigned Server	
VCF-MgmtDomHost-Template_DERIVED-1	AA01-6454-1	
<hr/>		
2 Name *	Assigned Server	
VCF-MgmtDomHost-Template_DERIVED-2	AA01-6454-2	

Step 5. Adjust the Prefix name and number (if needed).

Step 6. Click **Next**.

Step 7. Verify the information and click **Derive** to create the Server Profiles.

Step 8. Intersight will start configuring the server profiles and will take some time to apply all the policies. Use the Requests tab to see the progress.



Step 9. When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK (only two out of four servers shown in the screen capture):

* All UCS Server Prof... +			
...	🗑️	🔍 Add Filter	
<input type="checkbox"/>	Name	Status	UCS Server Template
<input type="checkbox"/>	VCF-MgmtDomHost-01	OK	VCF-MgmtDomHost-Template
<input type="checkbox"/>	VCF-MgmtDomHost-02	OK	VCF-MgmtDomHost-Template

Derive VI Workload Domain Server Profile

Procedure 1. Derive One or more Server Profiles

Step 1. From the **Infrastructure Services > Configure > Templates**, click “...” next to the VI Workload Domain Host template name and select **Derive Profiles**.

Step 2. Under the Server Assignment, select **Assign Now** and pick three Cisco UCS X210c M6 compute nodes. Customers can adjust the number of servers depending on the number of profiles to be deployed.

Note: In this deployment **three** FlashStack VI Workload Domain hosts will be derived.

Server Assignment

Assign Now

Assign Server from a Resource Pool

Assign Later

🔍 Add Filter 🔗 8 items found

<input type="checkbox"/>	Name	Model	UCS Domain
<input checked="" type="checkbox"/>	AA01-6454-1-2	UCSX-210C-M6	AA01-6454
<input checked="" type="checkbox"/>	AA01-6454-1-3	UCSX-210C-M6	AA01-6454
<input checked="" type="checkbox"/>	AA01-6454-1-4	UCSX-210C-M6	AA01-6454

Step 3. Click **Next**.

Step 4. Cisco Intersight will fill in “default” information for the selected servers (only two out of three servers shown):

Derive

Profile Name Prefix	Digits Count	Start Index for Suffix
AA01-FC-Boot-Template_DERIVED-	1 ⬆️ ⬆️ >= 1	1 ⬆️ ⬆️ >= 0
1 Name *		
AA01-FC-Boot-Template_DERIVED-1	Assigned Server	AA01-6454-1-7
2 Name *		
AA01-FC-Boot-Template_DERIVED-2	Assigned Server	AA01-6454-1-8

Step 5. Adjust the Prefix name and number (if needed).

Step 6. Click **Next**.

Step 7. Verify the information and click **Derive** to create the Server Profiles.

Step 8. Cisco Intersight will start configuring the server profiles and will take some time to apply all the policies. Use the Requests tab to see the progress.



Step 9. When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK.

* All UCS Server Prof... 🔍 +

⋮ ✎ 🗑️ 🔍 Add Filter

<input type="checkbox"/>	Name	Status	UCS Server Template
<input type="checkbox"/>	AA01-FC-Boot-01	🟢 OK	AA01-FC-Boot-Template
<input type="checkbox"/>	AA01-FC-Boot-02	🟢 OK	AA01-FC-Boot-Template
<input type="checkbox"/>	AA01-FC-Boot-03	🟢 OK	AA01-FC-Boot-Template

SAN Switch Configuration

This chapter contains the following:

- [Physical Connectivity](#)
- [FlashStack Initial Switch Configuration](#)
- [Enable Features](#)
- [Add NTP Servers and Local Time Configuration](#)
- [Configure Ports](#)
- [Create VSANs](#)
- [Create Device Aliases](#)
- [Create Zones and Zoneset](#)

This chapter provides the procedure for configuring the Cisco MDS 9132T switches used for Fibre Channel (FC) switching in this solution. The switch configuration for this validated design is based on the MDS configuration explained in the FlashStack Datacenter with Cisco UCS X-Series Cisco Validated Design (CVD):

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FibreChannelSANConfiguration therefore this chapter only explains the changes from the base CVD.

Physical Connectivity

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

FlashStack Initial Switch Configuration

To set up the initial switch configuration, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#FlashStackCiscoMDSBase

Enable Features

To set up various features on Cisco MDS, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#EnableLicenses

Add NTP Servers and Local Time Configuration

To configure the NTP server and add local time configuration, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#AddSecondNTPServerandLocalTimeConfiguration

Configure Ports

To set up the port and port-channel configuration, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#ConfigureIndividualPorts

Create VSANs

To create necessary VSANs, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CreateVSANs

Create Device Aliases

To obtain the WWPN information from Cisco Intersight and Pure Storage and to configure device aliases, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CreateDeviceAliases

Create Zones and Zoneset

To configure the zones and zonesets, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CreateZonesandZoneset

At this time, two Cisco MDS switches should be fully configured and the ports and port-channels should be enabled. Zoning configuration on MDS will allow compute nodes to communicate with the Pure Storage array's FC interfaces.

Storage Configuration – Purity//FA Boot Storage Setup

This chapter contains the following:

- [Create and map Initiator Groups](#)
- [Create Boot LUNs](#)
- [Create Infrastructure LUN](#)

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

Create and map Initiator Groups

To obtain the WWPN information from Cisco Intersight, create the initiators, plus the initiator groups for the three (or more) VI workload domain hosts, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#HostRegistration

Create Boot LUNs

To create boot LUNs for all three (or more) VI workload domain ESXi servers, and map the VI workload domain hosts to their respective boot LUNs, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#PrivateBootVolumesforeachESXiHost

Create and map Infrastructure LUNs

To create the VI workload domain infrastructure LUNs, which will be used as a VMFS datastore for VM storage and to map the initiator group to the infrastructure LUN, complete the steps explained here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html#CreateInfraDatastores

Note: It is advised to create at least one large infrastructure LUN for VM storage, and one additional smaller LUN to be used for swap files.

Note: On completing this storage configuration, the VI workload domain hosts should be able to access their boot LUNs on the Pure Storage FlashArray. VMware vSphere ESXi 7.0U3 can be installed on the configured boot LUNs for all the hosts. The VI workload domain hosts will also see one or more larger infrastructure LUNs, which will be consumed during the workload domain setup.

VMware vSphere ESXi 7.0U3 Initial Setup

This chapter contains the following:

- [VMware ESXi 7.0U3](#)
- [Create the custom ESXi installation ISO](#)
- [Access Cisco Intersight and Launch KVM](#)
- [Set Up VMware ESXi Installation](#)
- [Install ESXi](#)
- [Prepare the ESXi Hosts](#)

VMware ESXi 7.0U3

This section provides detailed instructions for installing VMware ESXi 7.0 U3 on all the hosts in the environment. On successful completion of these steps, four ESXi hosts will be available to setup the management domain and three ESXi hosts will be available for the VI workload domain. ESXi software will be installed on the local drive for the management hosts and FC based boot LUNs for the VI workload domain hosts.

Several methods exist for installing ESXi in a VMware environment. This procedure focuses on using the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco Intersight to map remote installation media to individual servers.

Because VMware Cloud Foundation requires a specific version of ESXi to be used during installation, only the 7.0 U3g version can be used to install the management and workload domain hosts. In addition, Cisco specific drivers must be used, and their versions must align with the firmware in use on the servers. Therefore, a custom ESXi installation ISO file must be created using the VMware vCenter image builder to use for installations. Assuming all versions of firmware and software are in place as listed in [Table 3](#), then all components will be aligned when following the steps below to create a custom ESXi installation ISO file.

Create the custom ESXi installation ISO

Procedure 1. Download VMware ESXi ISO

Step 1. Click the following link: [VMware ESXi 7.0 U3g](#)

Step 2. Download the Offline Bundle .zip file.

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

Procedure 2. Download Cisco Drivers

Step 1. Download the Cisco nenic driver here:

<https://customerconnect.vmware.com/downloads/details?downloadGroup=DT-ESXI70-CISCO-NENIC-10450&productId=974>

Step 2. Download the Cisco nfnic driver here:

<https://customerconnect.vmware.com/downloads/details?downloadGroup=DT-ESXI70-CISCO-NFNIC-50037&productId=974>

Step 3. Both downloads are a zip file, which also contains another zip file with the driver. Extract both driver zip files to known locations.

Procedure 3. Create the custom installation ISO file

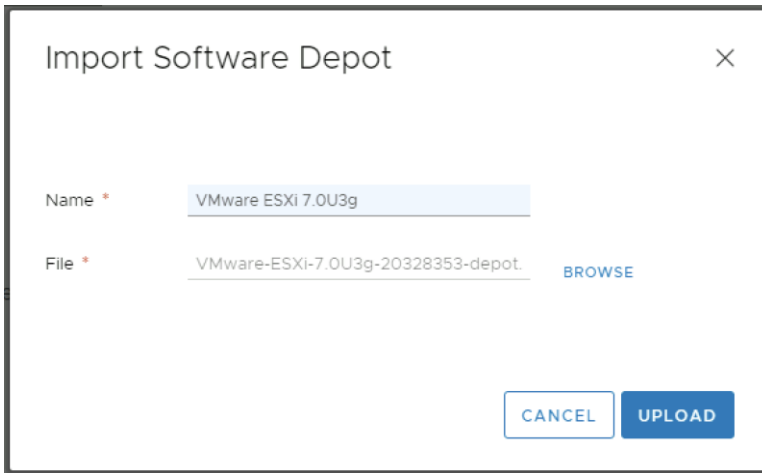
Step 1. Log in to the HTML5 management page of an existing VMware vCenter server using an account with administrative rights, for example, the vCenter server which will be later used to deploy the VMware Cloud Builder OVA for the VMware Cloud Foundation installation.

Step 2. From the upper menu, select Auto Deploy.

Step 3. If not already enabled, click the button for Enable Image Builder.

Step 4. Click Import to upload a software depot.

Step 5. Name the new depot “VMware ESXi 7.0U3g”. Click Browse and locate the base VMware ESXi 7.0U3g offline bundle .zip file downloaded earlier, then click Open.



Import Software Depot

Name * VMware ESXi 7.0U3g

File * VMware-ESXi-7.0U3g-20328353-depot. BROWSE

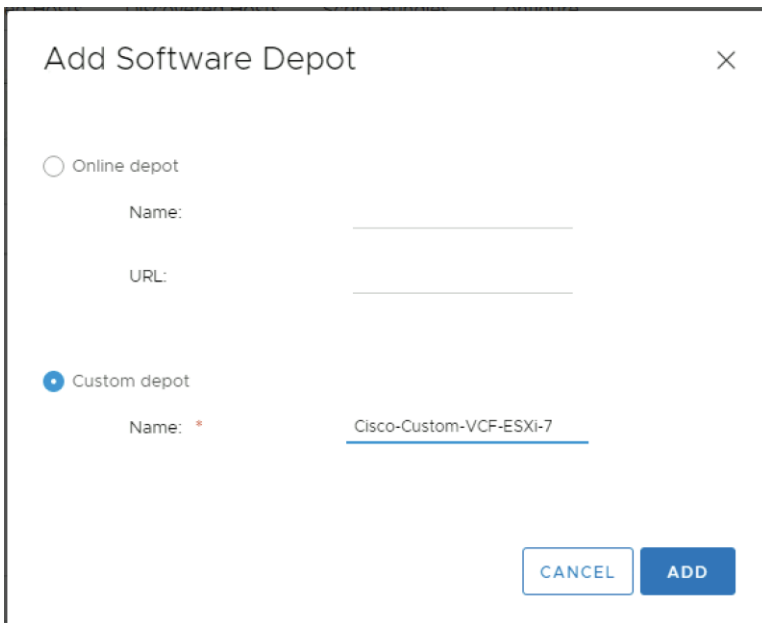
CANCEL UPLOAD

Step 6. Click Upload to upload the new software depot.

Step 7. Repeat steps 5-7 to create new software depots for the Cisco nenic and nfnis drivers, respectively, giving each a descriptive name and selecting the extracted .zip files from the downloads.

Step 8. Click New to add a custom software depot.

Step 9. Select Custom depot and name the depot “Cisco-Custom-VCF-ESXi-7.0U3g”.



Add Software Depot

Online depot

Name: _____

URL: _____

Custom depot

Name: * Cisco-Custom-VCF-ESXi-7

CANCEL ADD

Step 10. Click Add to add the custom depot.

Step 11. From the drop-down list, select the VMware ESXi 7.0U3g (ZIP) software depot. Make sure the Image Profiles tab is selected, then click the radio button to select the “ESXi-7.0U3g-20328353-standard” profile, then click Clone to clone the profile.

Step 12. Name the clone “Cisco-Custom-VCF-ESXi-7.0U3g”, for the vendor enter “VMware-Cisco”, and for the description enter “Cisco Custom ESXi 7.0U3g ISO with nenic 1.0.0.45 and nfnic 5.0.0.37”.

Step 13. Select Cisco-Custom-VCF-ESXi-7.0U3g as the software depot, then click Next.

Step 14. Under Available software packages, check the Cisco nenic 1.0.0.45 version, and the Cisco nfnic 5.0.0.37 version, then unselect any other nenic or nfnic packages.

<input type="checkbox"/>	Name	Version	Acceptance Level	Vendor	Depot
<input checked="" type="checkbox"/>	Isuv2-smartpiv2...	1.0.0-8vmw.703.0.20.19...	VMware certified	VMware	VMware ESXi 7.0U3g
<input checked="" type="checkbox"/>	mtp32xx-native	3.9.8-1vmw.703.0.20.19...	VMware certified	VMW	VMware ESXi 7.0U3g
<input checked="" type="checkbox"/>	native-misc-drive...	7.0.3-0.55.20328353	VMware certified	VMware	VMware ESXi 7.0U3g
<input checked="" type="checkbox"/>	ne1000	0.9.0-1vmw.703.0.50.2...	VMware certified	VMW	VMware ESXi 7.0U3g
<input type="checkbox"/>	nenic	1.0.33.0-1vmw.703.0.20...	VMware certified	VMW	VMware ESXi 7.0U3g
<input checked="" type="checkbox"/>	nenic	1.0.45.0-10EM.700.1.0.1...	VMware certified	Cisco	Cisco nenic 1.0.0.45
<input checked="" type="checkbox"/>	nfnic	5.0.0.37-10EM.700.1.0.1...	VMware certified	Cisco	Cisco nfnic 5.0.0.37
<input type="checkbox"/>	nfnic	4.0.0.70-1vmw.703.0.2...	VMware certified	VMW	VMware ESXi 7.0U3g
<input checked="" type="checkbox"/>	nhpsa	70.0051.0.100-4vmw.7...	VMware certified	VMW	VMware ESXi 7.0U3g

Step 15. Click Next, then click Finish to generate the new depot.

Step 16. Using the Software Depot pulldown list, select the Cisco-Custom-VCF-ESXi-7.0U3g (Custom) software depot. Under Image Profiles, select the Cisco-Custom-VCF-ESXi-7.0U3g image profile. Click Export to export the profile. ISO should be selected. Click OK to generate a bootable ISO installation image.

Step 17. After the export completes, click Download to download the ISO file. After the download finishes, you may rename the file to any descriptive name you like.

Access Cisco Intersight and Launch KVM

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


Procedure 1. Access Server KVM

Step 1. Log into the Cisco Intersight.

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

Step 3. Find the desired server and click “...” to see more options.

Step 4. Click **Launch vKVM**.



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

Note: Customers can launch the HTML5 KVM console for all the servers at the same time, but lab validation seemed to work the best when working with a couple of hosts at a time.

Note: Since the Cisco Custom ISO image will be mapped to the vKVM for software installation, it is better to use the standard vKVM (not the Tunneled vKVM) and that the Cisco Intersight is being accessed from a PC that has routable (or direct) access to the management network.

Set Up VMware ESXi Installation

Procedure 1. Prepare the Server for the OS Installation

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 customized ESXi installer ISO file created in the vCenter image builder from the previous steps.

Step 3. Click **Map Drive**.

Step 4. Select **Macros > Static Macros > Ctrl + Alt + Delete** to reboot the Server if the server is showing shell prompt. If the server is shutdown, from Intersight, select **Power > Power On System**.

Step 5. Monitor the server boot process in the KVM.

- VI workload domain servers should find the FC boot LUNs and then load the ESXi installer
- Management Domain servers should 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.

Install ESXi

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

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 M.2 local drive as installation disk for the management domain hosts or select the FC Pure Storage boot LUN as the installation disk for VI workload domain hosts 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, click on **Virtual Media** to unmap the installer ISO. Press **Enter** to reboot the server.

Step 8. Repeat these steps for installing ESXi on all the servers in the environment.

Prepare the ESXi Hosts

All the hosts in the environment need to be configured with following base configuration before a host can be onboarded in the VMware Cloud Foundation setup:

- Setup management access and enable SSH access
- Set hostname and DNS
- Set jumbo MTU on default vSwitch
- Set management VLAN for default VM Network port-group
- Configure NTP server
- Regenerate Certificates on all the ESXi hosts

Procedure 1. Set Up Management Access, enable SSH, set NTP and DNS information

Adding a management network for each VMware host is required for accessing and managing the host. Follow these steps on each ESXi host.

Step 1. After the server has finished rebooting, in the 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 SSH** and press **Enter**.

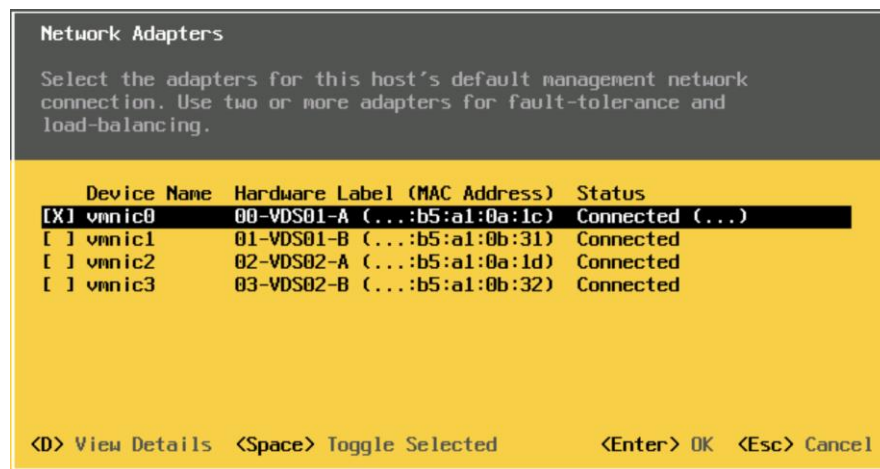
Step 5. Press **Esc** to exit the Troubleshooting Options menu.

Step 6. Select the **Configure Management Network** option and press **Enter**.

Step 7. Select **Network Adapters** and press Enter.

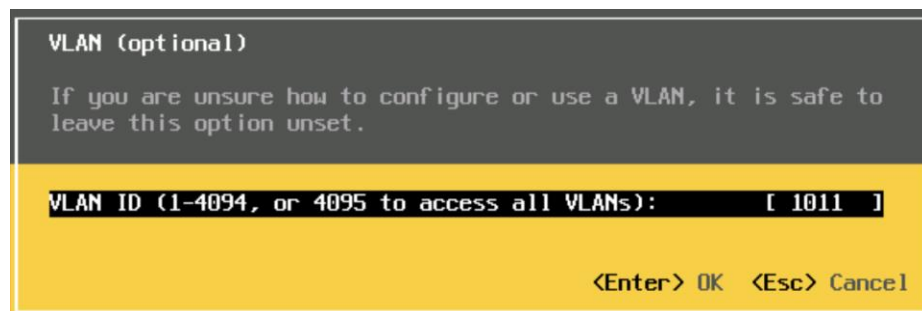
Step 8. Verify vmnic0 is selected as the only device.

Note: Do not add the second redundant NIC at this time.



Step 9. Press **Enter**.

Step 10. Under **VLAN (optional)** enter the IB-MGMT VLAN (for example, 1011) and press **Enter**.



Step 11. Select **IPv4 Configuration** and press **Enter**.

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

Step 12. Select the **Set static IPv4 address and network configuration** option by using the arrow keys and space bar.

Step 13. Under **IPv4 Address**, enter the IP address for managing the ESXi host.

Step 14. Under **Subnet Mask**, enter the subnet mask.

Step 15. Under **Default Gateway**, enter the default gateway.

Step 16. Press **Enter** to accept the changes to the IP configuration.

Step 17. Select the **IPv6 Configuration** option and press **Enter**.

Step 18. Using the spacebar, select **Disable IPv6 (restart required)** and press **Enter**.

Step 19. Select the **DNS Configuration** option and press **Enter**.

Note: If the IP address is configured manually, the DNS information must be provided. Make sure the ESXi hostnames are populated in the DNS server because VMware Cloud Foundation requires DNS resolution (both forward and reverse lookups) of all the ESXi hosts and the VCF component VMs.

Step 20. Using the spacebar, select **Use the following DNS server addresses and hostname:**

Step 21. Under **Primary DNS Server**, enter the IP address of the primary DNS server.

Step 22. Optional: Under **Alternate DNS Server**, enter the IP address of the secondary DNS server.

Step 23. Under **Hostname**, enter the fully qualified domain name (FQDN) for the ESXi host.

Step 24. Press **Enter** to accept the changes to the DNS configuration.

Step 25. Press **Esc** to exit the Configure Management Network submenu.

Step 26. Press **Y** to confirm the changes and reboot the ESXi host.

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

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, type **Ctrl-Alt-F1** to access the VMware console command line interface. In the Cisco Intersight KVM, Ctrl-Alt-F1 appears in the list of Static Macros.

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. Verify vmk1 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.

Step 10. Exit the ESXi host configuration:

Step 11. Type `exit` to log out of the command line interface.

Step 12. Type **Ctrl-Alt-F2** to return to the ESXi console menu interface.

Procedure 3. Setup Jumbo MTU on vSwitch and management VLAN on VM Network port-group

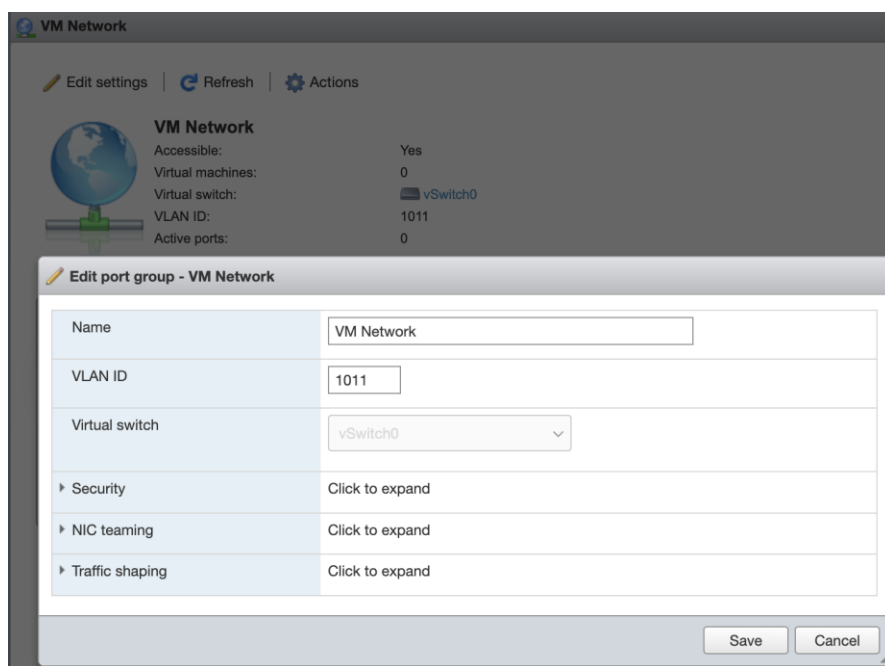
In this procedure, log into each ESXi host using a web browser and set the following options:

Step 1. Open a web browser and navigate to the first ESXi server's management IP address or DNS name.

Step 2. Enter "root" as the username.

Step 3. Enter the <root password>.

- Step 4.** Click **Log in** to connect.
- Step 5.** Decide whether to join the VMware Customer Experience Improvement Program or not and click **OK**.
- Step 6.** From the Host Client Navigator, select **Networking**.
- Step 7.** In the center pane, select the **Virtual switches** tab.
- Step 8.** Click **vSwitch0**.
- Step 9.** Click **Edit settings**.
- Step 10.** Change the MTU to **9000**.
- Step 11.** Click **Save**.
- Step 12.** Select **Networking**, then select the **Port groups** tab.
- Step 13.** In the center pane, right-click **VM Network** and select **Edit settings**.
- Step 14.** Set the **VLAN ID** to <IB-MGMT-VLAN> (for example, 1011).



- Step 15.** Click **Save** to finalize the edits for the **VM Network** port group.

Procedure 4. Configure NTP Server on the ESXi Hosts

- Step 1.** From the left pane in the ESXi web console, click **Manage** under Host.
- Step 2.** In the center pane, select **System > Time & date**.
- Step 3.** Click **Edit NTP Settings**.
- Step 4.** Select **Use Network Time Protocol (enable NTP client)**.
- Step 5.** Use the drop-down list to select **Start and stop with host**.
- Step 6.** Enter the NTP server IP address(es) in the NTP servers.

Step 7. Click **Save** to save the configuration changes.

Step 8. Select the **Services** tab.

Step 9. Right-click **ntpd** and select **Start**.

Step 10. Under **System > Time & date**, the NTP service status should now show “Running.”

Note: You might have to click **Refresh** to get the latest service status.

Current date and time	Sunday, November 20, 2022, 04:37:39 UTC
NTP service status	Running
NTP servers	1. 172.20.10.11

Procedure 5. Regenerate the ESXi self-signed certificates

After updating the ESXi host with the FQDN and setting up various parameters previously explained, regenerate the self-signed certificates:

Step 1. SSH to each VMware ESXi host and log in as **root**.

Step 2. Run the following commands on each host:

```
/sbin/generate-certificates
/etc/init.d/hostd restart && /etc/init.d/vpxa restart
```

Note: If you were logged into the web UI for any of the ESXi hosts, the session will have to be refreshed and you will need to log back in.

VMware Cloud Foundation Deployment

This chapter contains the following:

- [Prepare the Existing Infrastructure](#)
- [Deploy Cloud Builder Virtual Appliance](#)
- [Deploy the Management Domain](#)
- [Commission Workload Domain Hosts](#)
- [Deploy the VI Workload Domain](#)

The VMware Cloud Foundation deployment is comprised of the following steps:

1. Prepare the existing infrastructure
2. Deploy the Cloud Builder Virtual Appliance
3. Deploy the management domain*
4. Onboard FlashStack workload domain ESXi hosts
5. Deploy the VI workload domain

Note: *For customers who only need to onboard FlashStack VI workload domain in an existing VMware Cloud Foundation setup, the management domain setup and related infrastructure preparation steps can be skipped. This deployment guide assumes customers are setting up a new VMware Cloud Foundation from the beginning.

Prepare the Existing Infrastructure

Before starting the automated deployment of the management domain using VMware Cloud Builder, the environment must meet target prerequisites and be in a specific starting state. Make sure following elements are present in the current environment:

- The deployment environment should contain an NTP server for the ESXi hosts.
- The deployment environment should have a DNS infrastructure and all following VM hostnames should be programmed in the DNS server (both forward and reverse lookups):
 - VMware Cloud Builder VM
 - VMware SDDC Manager
 - VMware vCenter for management and VI domains
 - VMware NSX-T manager VMs and cluster VIPs for management and VI domains
 - All the ESXi hosts for management and VI domains
- The cloud builder VM is deployed in the existing customer environment. Customers should have a vSphere environment available to deploy the cloud builder OVF.
- The management network where the VCF components are being deployed should be routable.

[Table 11](#) lists the DNS information used during this deployment. Customers should validate both the forward and reverse lookups to verify DNS is working properly.

Table 11. VMware Cloud Foundation sample DNS information

FQDN	IP Address	Description
aa01-ad1.vcf.local	10.101.1.53	DNS server #1
aa01-ad2.vcf.local	10.101.1.54	DNS server #2
aa01-cloudbuilder.vcf.local	10.101.1.5	Cloud Builder VM
Management Domain		
vcf-sddc.vcf.local	10.101.1.110	SDDC manager
vcf-vc.vcf.local	10.101.1.80	Management Domain vCenter
vcf-esxi-01.vcf.local	10.101.1.81	Management Domain ESXi server #1
vcf-esxi-02.vcf.local	10.101.1.82	Management Domain ESXi server #2
vcf-esxi-03.vcf.local	10.101.1.83	Management Domain ESXi server #3
vcf-esxi-04.vcf.local	10.101.1.84	Management Domain ESXi server #4
vcf-mgmt-nsx.vcf.local	10.101.1.90	Management Domain NSX-T Cluster VIP
vcf-mgmt-nsx-1.vcf.local	10.101.1.91	Management NSX-T virtual appliance node # 1
vcf-mgmt-nsx-2.vcf.local	10.101.1.92	Management NSX-T virtual appliance node # 2
vcf-mgmt-nsx-3.vcf.local	10.101.1.93	Management NSX-T virtual appliance node # 3
Workload Domain		
aa01-vc.vcf.local	10.101.1.100	VI workload domain vCenter
aa01-esxi-01.vcf.local	10.101.1.101	VI workload domain ESXi server #1
aa01-esxi-02.vcf.local	10.101.1.102	VI workload domain ESXi server #2
aa01-esxi-03.vcf.local	10.101.1.103	VI workload domain ESXi server #3
vcf-wd-nsx.vcf.local	10.101.1.95	Workload Domain NSX-T Cluster VIP
vcf-wd-nsx-1.vcf.local	10.101.1.96	Workload Domain NSX-T virtual appliance node # 1
vcf-wd-nsx-2.vcf.local	10.101.1.97	Workload Domain NSX-T virtual appliance node # 2
vcf-wd-nsx-3.vcf.local	10.101.1.99	Workload Domain NSX-T virtual appliance node # 3

Deploy Cloud Builder Virtual Appliance

The VMware Cloud Builder virtual appliance 4.5 and the associated parameter files can be downloaded from the VMware website:

https://customerconnect.vmware.com/downloads/info/slug/datacenter_cloud_infrastructure/vmware_cloud_foundation/4_5

Download both the **VMware Cloud Builder OVA** and the **Cloud Builder Deployment Parameter Guide - Subscription** xlsx spreadsheet file.

Procedure 1. Deploy the Cloud Builder OVA

The VMware Cloud Builder OVA will be deployed on an existing VMware infrastructure.

Step 1. Log into an existing VMware vCenter and select the cluster/host where the Cloud Builder OVA will be deployed.

Step 2. Right-click the cluster or host and select **Deploy OVF Template**.

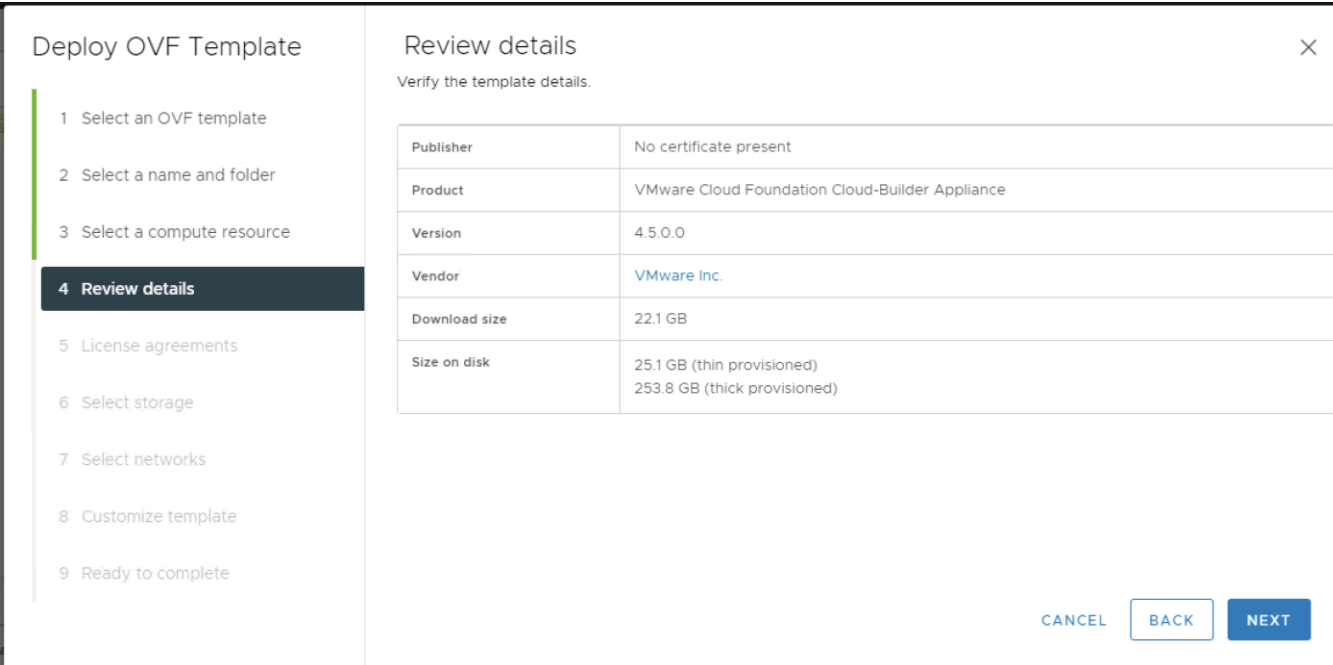
Step 3. Select **Local file** and click **UPLOAD FILES**.

Step 4. Select the VMware Cloud Builder OVA file downloaded in the last step and click **Open**.

Step 5. Click **NEXT**.

Step 6. Provide a VM name (for example, aa01-cloudbuilder) and select the location for the install. Click **NEXT**.

Step 7. Verify the template details and click **NEXT**.



The screenshot shows the 'Deploy OVF Template' wizard. The sidebar on the left lists the following steps:

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details**
- 5 License agreements
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

The main area is titled 'Review details' and contains the text 'Verify the template details.' Below this is a table with the following data:

Publisher	No certificate present
Product	VMware Cloud Foundation Cloud-Builder Appliance
Version	4.5.0.0
Vendor	VMware Inc.
Download size	22.1 GB
Size on disk	25.1 GB (thin provisioned) 253.8 GB (thick provisioned)

At the bottom right of the wizard, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

Step 8. Accept the license agreement and click **NEXT**.

Step 9. Select the datastore where the VM is deployed and click **NEXT**.

Step 10. Select the correct management network from the drop-down list and click **NEXT**.

Step 11. On the Customize template screen, enter password for the default **admin** account.

Step 12. On the same Customize template screen, enter password for the default **root** account.

Step 13. Scroll down and provide the **Hostname, IP address, Subnet Mask, Default GW, DNS Servers, DNS Domain Name** for the cloud builder appliance.

Hostname	Enter a hostname for this virtual appliance. <u>aa01-cloudbuilder</u>
Network 1 IP Address	Enter an IP Address for the interface of this virtual appliance. <u>10.101.1.5</u>
Network 1 Subnet Mask	Enter a subnet mask for the interface of this virtual appliance. Example: 255.255.255.0 <u>255.255.255.0</u>
Default Gateway	Enter a default gateway for the interface of this virtual appliance. <u>10.101.1.254</u>
DNS Servers	Enter the DNS servers for this virtual appliance (comma separated). WARNING: Do not specify more than two entries otherwise no configuration will be set. <u>10.101.1.53,10.101.1.54</u>
DNS Domain Name	Enter the domain name for this virtual appliance. Example: rainpole.local <u>vcf.local</u>

Step 14. Scroll down, provide the **DNS Domain Search Paths** and **NTP Server**. Click **NEXT**.

DNS Domain Search Paths	Enter the domain name search paths for this virtual appliance (comma separated). Example: rainpole.local, sfo01.rainpole.local <u>vcf.local</u>
NTP Servers	Enter NTP time sources for this virtual appliance (comma separated). Example: ntp0.rainpole.local,ntp1.rainpole.local <u>172.20.10.11</u>

Step 15. Verify all the information and click **FINISH** to deploy the appliance.

Step 16. When the deployment is complete, access the Cloud Builder appliance by typing the FQDN of cloud builder (for example, <https://aa01-cloudbuilder.vcf.local>) in a web browser window to verify the installation was successful.

The cloud builder appliance is now ready to deploy VMware Cloud Foundation management domain. The next step is to populate the Cloud Builder Deployment Parameter Guide with the necessary infrastructure information.

Deploy the Management Domain

The first step in deploying VMware Cloud Foundation management domain is to fill in all the deployment information in the Cloud Builder Deployment Parameter Guide. The second step is to upload this parameter file into cloud builder and start the VMware Cloud Foundation deployment process.

Procedure 1. Update Cloud Builder Deployment Parameter Guide

Step 1. Open the Cloud Builder Deployment Parameter workbook in Microsoft Excel.

The Introduction worksheet provides an overview of the workbook.

Step 2. Click to select **Credentials** worksheet.

Step 3. Add **root** password for all the ESXi hosts. This password was set at the time of ESXi installation.

Step 4. Provide the default passwords (to be set) for various roles of vCenter, NSX-T and SDDC Manager.

Credentials		
Instructions: Use the Users and Groups tab to input the default passwords used for built-in accounts for each component, these will be used to implement the Management Domain. - Grey cells are for information purposes and cannot be modified. - Red cells mean the input data is either missing and required or some type of validation of the input data has failed. Password Policy: Each password has its own password policy typically a minimum number of characters in length and atleast one uppercase, lowercase, number and special character (e.g: {}[]()/\\" *~.,;.:<>)		
Users		
Username	Default Password	Description
ESXi		
root		ESXi Host Root Account (Same for all ESXi hosts)
vCenter Server		
administrator@vsphere.local		Default Single-Sign On Domain Administrator User
root		vCenter Server Virtual Appliances Root Account
NSX-T Data Center		
root		NSX-T Virtual Appliance Root Account - NSX-T Manager and Edge Nodes
admin		NSX-T User Interface and Default CLI Admin Account - NSX-T Manager and Edge Nodes
audit		NSX-T Audit CLI Account - NSX-T Manager and Edge Nodes
SDDC Manager		
root		SDDC Manager Appliance Root Account
vcf		SDDC Manager Super User
admin@local		SDDC Manager Local Account

Step 5. Select **Hosts and Networks** worksheet.

Step 6. Provide the Management Network, vMotion Network and vSAN Network information including port-group name, IP subnet, IP gateway and MTU under **Management Domain Networks**.

Management Domain Networks					
Network Type	VLAN #	Portgroup Name	CIDR Notation	Gateway	MTU
Management Network	1011	MGMT_10_101_1_NET	10.101.1.0/24	10.101.1.254	1500
vMotion Network	3030	vds01-pg-vmotion	192.168.30.0/24	192.168.30.254	9000
vSAN Network	3001	vds01-pg-vsant	192.168.1.0/24	192.168.1.254	9000

Step 7. Provide the existing vSwitch name on the ESXi hosts under **Virtual Networking**.

Step 8. Select **Profile-3** under VDS Switch Profile

Step 9. Provide the names (to be set) for the two VDSs (for example, vds01 and vds02) and the vNICs assigned to these VDSs (for example, vmnic0, vmnic1 for vds01 and vmnic2, vmnic3 for vds02).

Virtual Networking	Value
vSphere Standard Switch Name	vSwitch0
Primary vSphere Distributed Switch	Value
Primary vSphere Distributed Switch - Name	vds01
Primary vSphere Distributed Switch - pNICs	vmnic0,vmnic1
Primary vSphere Distributed Switch - MTU Size	9000
Secondary vSphere Distributed Switch (Optional)	Value
Secondary vSphere Distributed Switch - Name	vds02
Secondary vSphere Distributed Switch - pNICs	vmnic2,vmnic3
Secondary vSphere Distributed Switch - MTU Size	9000

vSphere Distributed Switch Profile	Profile-3
vSphere Distributed Switch = Two (2) / Physical NICs = Four (4)	
Primary vDS - vds01 - Traffic for Management, vMotion, vSAN - e.g. vmnic0,vmnic1	
Secondary vDS -vds02 - Traffic for Host Overlay - e.g. vmnic2,vmnic3	

Step 10. Provide the name and IP addresses for accessing the four ESXi hosts under **Management Domain ESXi Hosts**.

Note: Cloud Builder appliance should be able to resolve the ESXi hostnames to IP addresses.

Step 11. Provide the pool range (start and end IP addresses) for both vMotion and vSAN networks. An IP address from each of these ranges will be configured on each ESXi host.

Management Domain ESXi Hosts			
vcf-esxi-01	vcf-esxi-02	vcf-esxi-03	vcf-esxi-04
10.101.1.81	10.101.1.82	10.101.1.83	10.101.1.84
vMotion Start IP	192.168.30.81	vMotion End IP	192.168.30.90
vSAN Start IP	192.168.1.81	vSAN End IP	192.168.1.90

Step 12. Select **No** for **Validate Thumbprints**.

Step 13. Under the **NSX-T Host Overlay Network**, provide the Overlay Network **VLAN ID** (for example, 3002).

Step 14. Select **Yes** for **Configure NSX-T Host Overlay Using a Static IP Pool**.

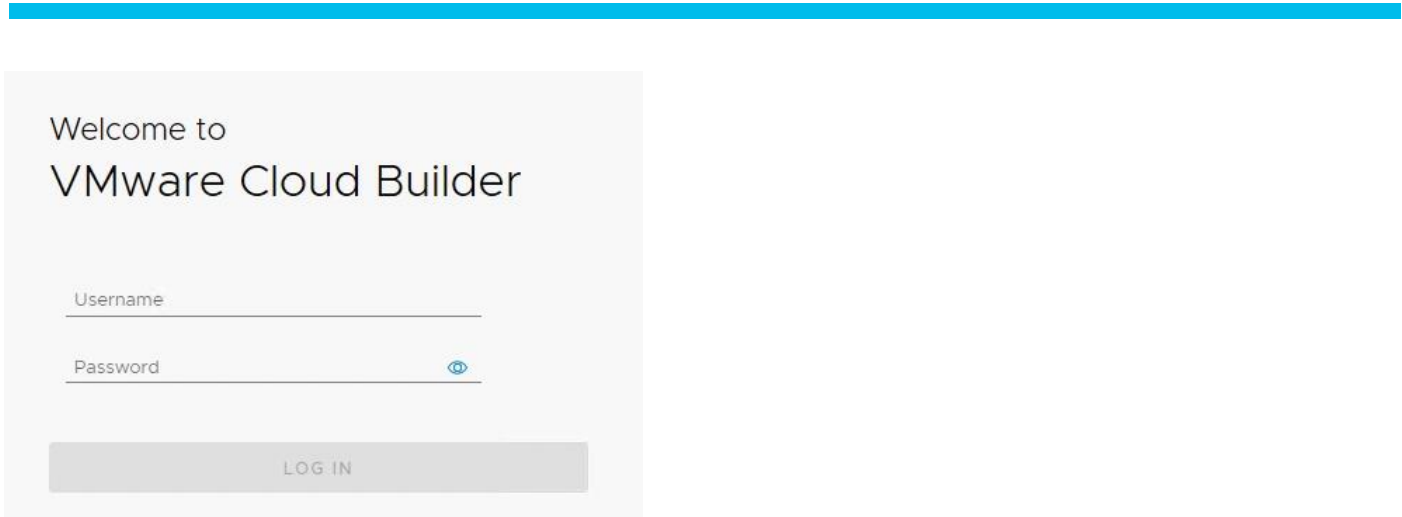
Step 15. Provide the NSX-T host overlay network values including pool name, IP subnet, IP gateway and IP range.

VLAN ID	3002		
Configure NSX-T Host Overlay Using a Static IP Pool	Yes		
Pool Description	ESXi Host Overlay TEP IP Pool		
Pool Name	tep01		
CIDR Notation	192.168.2.0/24	Gateway	192.168.2.254
NSX-T Host Overlay Start IP	192.168.2.1	NSX-T Host Overlay End IP	192.168.2.10

Step 16. Click to select **Deploy Parameters** worksheet.

Step 17. Provide the DNS and NTP server information under **Existing Infrastructure Details**

Existing Infrastructure Details	Infrastructure	Value
<input checked="" type="checkbox"/> DNS Server and DNS Zone Defined	DNS Server #1	10.101.1.53
<input checked="" type="checkbox"/> NTP Servers	DNS Server #2	10.101.1.54
	NTP Server #1	172.20.10.11
	NTP Server #2	n/a



Step 3. Read and agree to the End User License Agreement and click **NEXT**.

Step 4. Select **VMware Cloud Foundation** under the select platform and click **NEXT**.

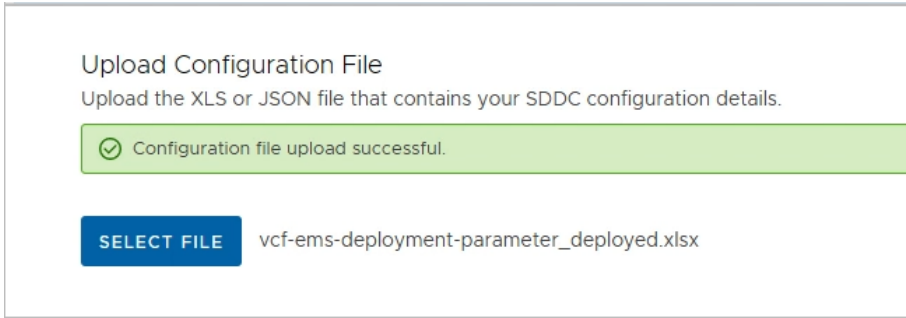


Step 5. Review the prerequisites for the SDDC deployment and agree to meeting the requirements. Click **NEXT**.

Step 6. Click **NEXT** for Download Deployment Parameter Workbook.

Step 7. Click **NEXT**. Deployment parameter workbook was already downloaded and filled.

Step 8. Click **SELECT FILE** and browse to the location of completed excel workbook and select the completed file. When the file is uploaded successfully, click **NEXT**.



Step 9. Cloud Builder appliance will take a while and verify the uploaded file. If there are any errors, fix the errors and click **RETRY**.

Step 10. On successful validation of the configuration file, click **NEXT**.

VMware Cloud Foundation

Cloud Builder will validate data provided in the configuration file and elements of the physical infrastructure.

Configuration file validated successfully.

DOWNLOAD PRINT

History	Validation Items	Status
Current	JSON Spec Validation	Success
10/31/22, 8:43 PM	Cloud Builder Configuration Validation	Success
	DNS Resolution Validation	Success
	Preparing Security Requirements for Running Validation	Success
	ESXi Host Configuration Validation	Success
	vSAN Disk Availability Validation(Hybrid)	Success
	License Key Validation	Success
	Password Validation	Success
	Network Configuration Validation	Success
	vMotion Network Connectivity Validation	Success
	vSAN Network Connectivity Validation	Success
	NSX-T Data Center Host Overlay Network Connectivity Validation	Success
	Time Synchronization Validation	Success
	Network IP Pool Validation	Success

BACK RETRY NEXT

Step 11. In the dialog box for deploying SDDC, click **DEPLOY SDDC**.

Deploy SDDC?



Select Deploy SDDC to begin deployment of VMware Cloud Foundation.
Once you begin deployment, you cannot stop the process.

If you are not yet ready, select Cancel to stay at this step until you are ready to deploy the SDDC.

CANCEL DEPLOY SDDC

The Cloud Builder appliance will take a while to deploy the vCenter, vSAN, SDDC-Manager, NSX-T appliances and adjust various parameters on the ESXi hosts.

Note: You can log into the management ESXi hosts to see vCenter, vSAN and NSX-T VMs getting deployed.

VMware Cloud Foundation

Cloud Builder will deploy your SDDC.



SDDC Bringup is in progress.

DOWNLOAD PRINT

SDDC Bringup started at 10/31/22, 5:09 PM. 0 tasks in progress

Search Tasks

Status

Tasks	Start Time	End Time	Status
▼ Validate SSH/SSL Thumbprints			⊖ Not Started
Generate Security Thumbprints Input Data			⊖ Not Started
Validate Security Thumbprints			⊖ Not Started
▼ Add Certificates in Trust-Store			⊖ Not Started
Generate input for Trust Certificates			⊖ Not Started
Trust Certificates			⊖ Not Started
▼ Import SSH Keys			⊖ Not Started
Generate input for Import SSH Keys			⊖ Not Started
Import SSH Keys			⊖ Not Started
▼ Prepare Environment for Bringup Execution			⊖ Not Started
Generate ESXi Host vSAN Configuration Input Data			⊖ Not Started
Generate ESXi Host Input Data			⊖ Not Started
Retrieve ESXi Host Lockdown Mode Configuration			⊖ Not Started
Disable Lockdown Mode on ESXi Hosts			⊖ Not Started
Generate ESXi Service Accounts Data			⊖ Not Started

BACK RETRY FINISH

Step 12. When all the configuration steps are successfully completed, Cloud Builder appliance will notify user of the deployment completed successfully. Click **FINISH**.

VMware Cloud Foundation

Cloud Builder will deploy your SDDC.



Deployment of VMware Cloud Foundation is successful.

DOWNLOAD PRINT

SDDC Bringup finished at 10/31/22, 6:43 PM. 0 tasks in progress

Search Tasks

Status

Tasks	Start Time	End Time	Status
Clear Alarms on vSAN	6:42:30 PM	6:42:31 PM	⊕ Success
Clear Alerts on Hosts	6:42:31 PM	6:42:33 PM	⊕ Success
Set SDDC Deployment Details on the Management vCenter Server	6:42:33 PM	6:42:34 PM	⊕ Success
▼ Disable Bash Shell on vCenter			⊕ Success
Generate vSphere Input Data	6:42:34 PM	6:42:35 PM	⊕ Success
Disable Bash Shell on vCenter Server	6:42:35 PM	6:42:38 PM	⊕ Success
▼ Configure NSX-T Data Center to Comply with Security Requirements			⊕ Success
Generate NSX-T Data Center Input Data	6:42:39 PM	6:42:39 PM	⊕ Success
Enable/Disable SSH on NSX-T Data Center Manager Nodes	6:42:40 PM	6:42:53 PM	⊕ Success
▼ Perform configuration changes on SDDC Manager to disable basic auth based API access			⊕ Success
Generate SDDC Manager Input Data	6:42:54 PM	6:42:54 PM	⊕ Success
Disable Basic Authentication API Access on SDDC Manager	6:42:55 PM	6:42:59 PM	⊕ Success
▼ Perform disable SSH operation on all ESXi hosts			⊕ Success
Generate SDDC Manager Input Data	6:42:59 PM	6:43:00 PM	⊕ Success
Disable SSH on ESXi host	6:43:00 PM	6:43:01 PM	⊕ Success

BACK RETRY FINISH

Step 13. In the pop up, click **LAUNCH SDDC MANAGER**.

SDDC Deployment Complete



✔ You have successfully deployed VMware Cloud Foundation.

VMware Cloud Foundation Proactive Support

Skyline proactive support helps you avoid problems before they occur and reduces the time spent on resolving active support requests. With just a few clicks you can increase team productivity and the overall reliability of your VMware environments. And, it's included in your active Production Support or Premier Services subscription. With Skyline, you've got control, and we've got your back. Please install [Skyline](#) to enable proactive support for your Cloud Foundation environment

LAUNCH SDDC MANAGER

Step 14. Use [administrator@vsphere.local](#) username and password set in the parameters workbook to log into the SDDC manager.

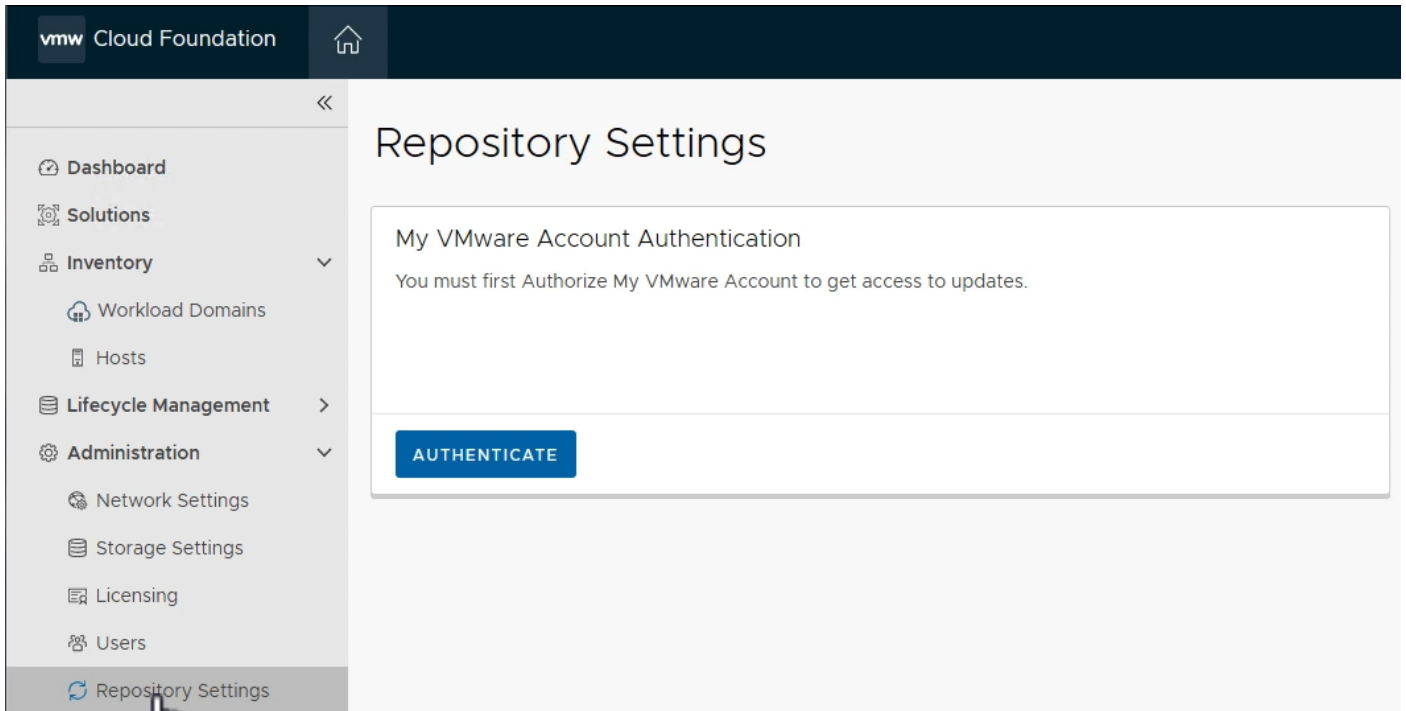
Step 15. Provide an input to Customer Experience Improvement Program question and click **APPLY**.

Now the SDDC Manager dashboard is accessible.

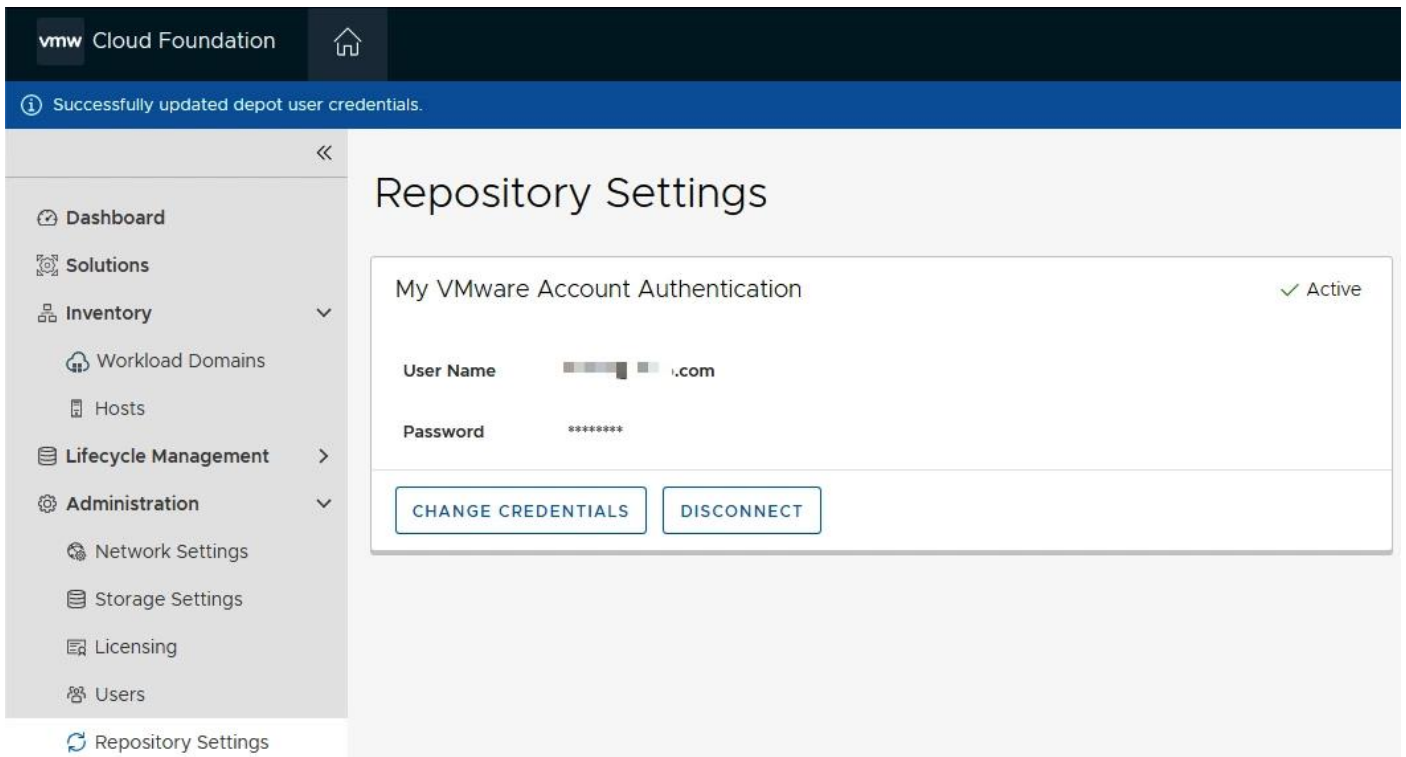
The screenshot displays the SDDC Manager Dashboard interface. The top navigation bar includes the VMware logo, 'Cloud Foundation', a home icon, a user profile icon, and the username 'administrator@vsphere.local'. The left sidebar contains a navigation menu with categories like Solutions, Inventory, Lifecycle Management, Administration, and Developer Center. The main content area is titled 'SDDC Manager Dashboard' and features several widgets: '0 Solutions' (Workload Management), '1 Workload Domains' (Management Domain, VI Domain), 'Host Type and Usage' (Hybrid Host, All Flash Host, Usage bar chart), and 'CPU, Memory, Storage Usage' (CPU usage bar chart, Top Domains in allocated CPU Usage, Memory usage bar chart, Top Domains in allocated Memory Usage, Storage usage bar chart, Top Domains in allocated Storage Usage). A 'Recent tasks' widget on the right shows 'You haven't started any tasks'. At the bottom, a 'Tasks' table is visible with columns for Task, Description, Status, and Last Occurrence, and a message 'No Tasks found'.

Step 16. Navigate to **Administration > Repository Settings**.

Step 17. Click **AUTHENTICATE** to authorize the VMware account to access updates for VMware Cloud Foundation components.

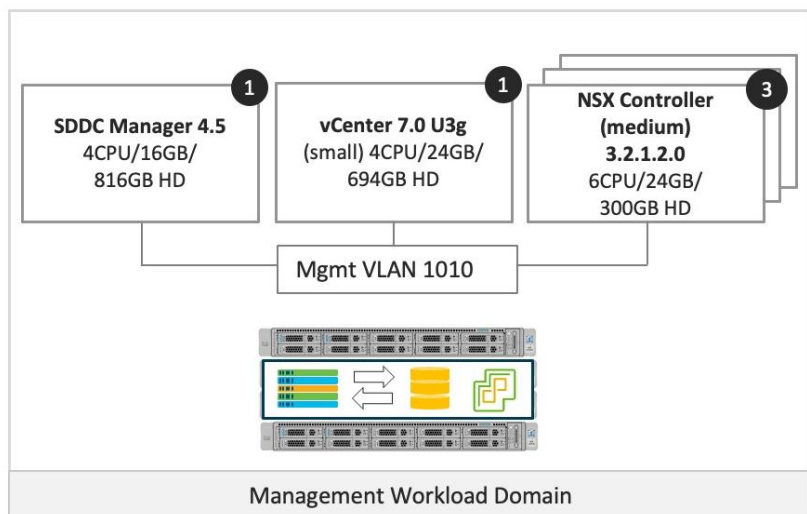


Step 18. Provide the VMware credentials and click **AUTHORIZE**.



Now the VMware Cloud Foundation management domain setup is complete. [Figure 9](#) shows various virtual machines and their deployment configuration in the management domain. Depending on the size of the deployment chosen in the deployment worksheet, the virtual machine size could be different.

Figure 9. VMWare Cloud Foundation Management Workload Domain



Customers can log into the VMware SDDC manager and find various deployment parameters, perform lifecycle management, and gather information about the vCenter and NSX manager. Customers may choose to perform the necessary lifecycle management for VMware Cloud Foundation by downloading and upgrading the software packages.

For more details on deploying the management cluster using VMware cloud builder appliance, see: <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-78EEF782-CF21-4228-97E0-37B8D2165B81.html>

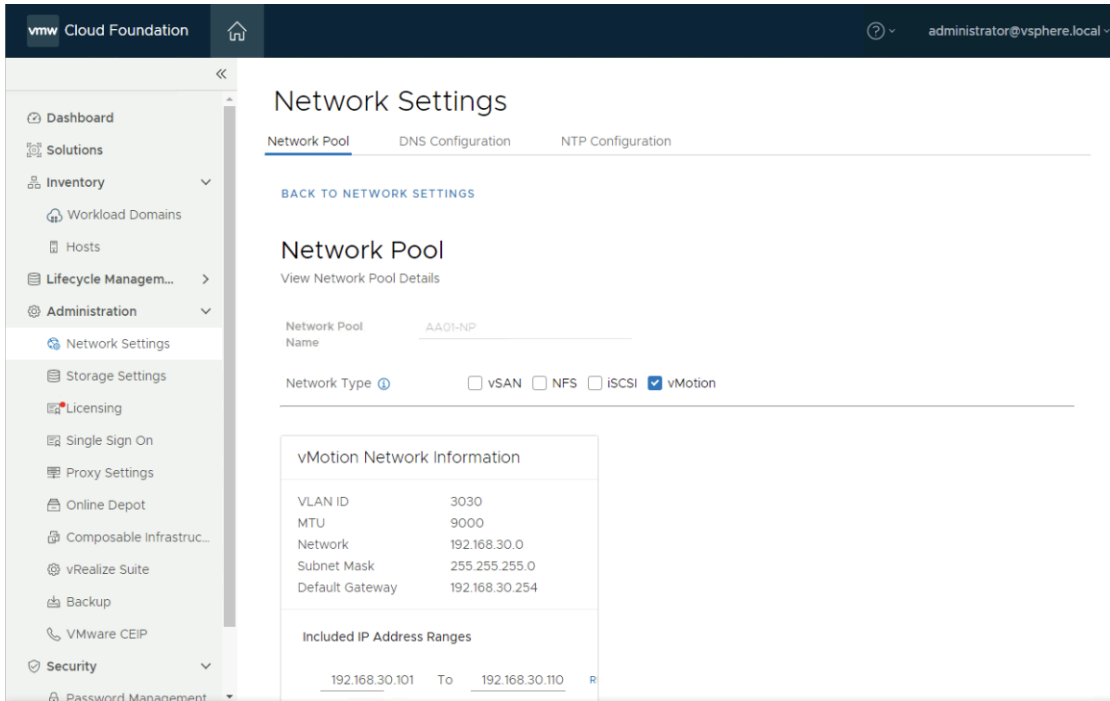
Commission Workload Domain Hosts

A VMware Cloud Foundation VI workload domain is deployed using VMware SDDC manager. When deploying a VI workload domain, the storage type, compute, and networking details are provided to the SDDC manager. The FlashStack VI workload domain ESXi hosts are commissioned in the SDDC manager before proceeding with the workload domain deployment. This section covers the ESXi host commissioning in SDDC manager.

Procedure 1. Create the Network Pool for vMotion IP addresses

In this procedure, a network pool will be created to assign IP addresses to vMotion VMkernel ports of the workload domain ESXi hosts.

- Step 1.** Log into VMware SDDC manager GUI.
- Step 2.** Navigate to **Administration > Network Settings**.
- Step 3.** Click on **CREATE NETWORK POOL**.
- Step 4.** Provide a **Network Pool Name**.
- Step 5.** Check **vMotion**.
- Step 6.** Provide various Network values including **VLAN, MTU, IP Subnet, Subnet Mask, Default Gateway**, and a range of IP addresses that will be assigned to the workload domain hosts.



Step 7. Click **ADD** to add both the IP address ranges.

Step 8. Click **SAVE** to create the network pool.

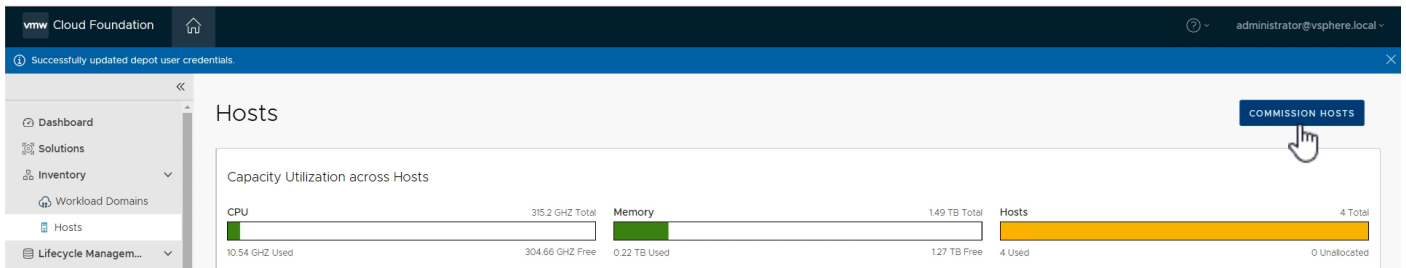
Procedure 2. Commission VI workload domain hosts

In this procedure, the VI workload domain ESXi hosts will be commissioned in SDDC manager.

Step 1. Log into VMware SDDC manager GUI.

Step 2. Navigate to **Inventory > Hosts**.

Step 3. Click on **COMMISSION HOSTS** in the main panel.



Step 4. Read and verify the checklist, check **Select All**, and click **PROCEED**.


Checklist

Commissioning a host adds it to the VMware Cloud Foundation inventory. The host you want to commission must meet the checklist criterion below.

- ✓ **Select All**
- ✓ Host for vSAN workload domain should be vSAN compliant and certified per the VMware Hardware Compatibility Guide. BIOS, HBA, SSD, HDD, etc. must match the VMware Hardware Compatibility Guide.
- ✓ Host has a standard switch with two NIC ports with a minimum 10 Gbps speed.
- ✓ Host has the drivers and firmware versions specified in the VMware Compatibility Guide.
- ✓ Host has ESXi installed on it. The host must be preinstalled with supported versions (7.0.3-19482537)
- ✓ Host is configured with DNS server for forward and reverse lookup and FQDN.
- ✓ Hostname should be same as the FQDN.
- ✓ Management IP is configured to first NIC port.
- ✓ Ensure that the host has a standard switch and the default uplinks with 10Gb speed are configured starting with traditional numbering (e.g., vmnic0) and increasing sequentially.
- ✓ Host hardware health status is healthy without any errors.
- ✓ All disk partitions on HDD / SSD are deleted.
- ✓ Ensure required network pool is created and available before host commissioning.
- ✓ Ensure hosts to be used for VSAN workload domain are associated with VSAN enabled network pool.
- ✓ Ensure hosts to be used for NFS workload domain are associated with NFS enabled network pool.
- ✓ Ensure hosts to be used for VMFS on FC workload domain are associated with NFS or VMOTION only enabled network pool.
- ✓ Ensure hosts to be used for vVol FC workload domain are associated with NFS or VMOTION only enabled network pool.
- ✓ Ensure hosts to be used for vVol NFS workload domain are associated with NFS and VMOTION only enabled network pool.
- ✓ Ensure hosts to be used for vVol iSCSI workload domain are associated with iSCSI and VMOTION only enabled network pool.

CANCEL

PROCEED



Step 5. On the Host Addition and Validation screen, select **Add new**.

Step 6. Provide the ESXi host **FQDN** (for example, aa01-esxi-01.vcf.local).

Step 7. For Storage Type, select **VMFS on FC**.

Step 8. From the **Network Pool Name** drop down, select the Network Pool create in the last procedure (for example, AA01-NP).

Step 9. For the ESXi host username, enter **root**.

Step 10. Provide the root password.

Step 11. Click **ADD**.

Commission Hosts

- 1 Host Addition and Validation
- 2 Review

Host Addition and Validation

▼ Add Hosts

You can either choose to add host one at a time or download [JSON template](#) and perform bulk commission.

Add new Import

Host FQDN:

Storage Type:
 VSAN NFS VMFS on FC vVol

Network Pool Name ⓘ:

User Name:

Password:

ADD

Step 12. Repeat steps 6 through 11 to add all three hosts.

Step 13. In the Host Added section, select all hosts, and click **Confirm FingerPrint**.

Hosts Added

Click on Confirm FingerPrint button in the below grid to enable or disable to **validate** hosts before proceeding to commission

✔ Hosts added successfully. Add more or confirm fingerprint and validate host ✕

REMOVE
VALIDATE ALL


<input checked="" type="checkbox"/>	FQDN	Network Pool	IP Address	<input checked="" type="checkbox"/> Confirm FingerPrint	Validation Status
<input checked="" type="checkbox"/>	aa01-esxi-03.vcf.local	AA01-NP ⓘ	10.101.1103	<input checked="" type="checkbox"/> SHA256:ti2Uj8d S06FLbpWKEllq ttxLyOJ6S2ZAQ YqxzPpJmTQ	⊖ Not Validated
<input checked="" type="checkbox"/>	aa01-esxi-02.vcf.local	AA01-NP ⓘ	10.101.1102	<input checked="" type="checkbox"/> SHA256:UvUJ8 pC48Pd25vkP6 ewIWATeWS6P z4oTHayqvJh/u tE	⊖ Not Validated
<input checked="" type="checkbox"/>	aa01-esxi-01.vcf.local	AA01-NP ⓘ	10.101.1101	<input checked="" type="checkbox"/> SHA256:XKEC5 KYJICUKPKJqq hRoT+8GCqUY9 BNQa5Aa+3Bplf g	⊖ Not Validated
<input checked="" type="checkbox"/> 3					3 hosts

Step 14. Click **VALIDATE ALL**.




Step 15. SDDC manager will take a while and validate host configurations. When the validation is successful, **Host Validated Successfully** message appears on the screen.


Hosts Added

Click on Confirm FingerPrint button   in the below grid to enable or disable to **validate** hosts before proceeding to commission

 Host Validated Successfully.✕

REMOVEVALIDATE ALL

<input checked="" type="checkbox"/>	FGDN	Network Pool	IP Address	<input checked="" type="checkbox"/> Confirm FingerPrint	Validation Status
<input checked="" type="checkbox"/>	aa01-esxi-03.vcf.local	AA01-NP 	10.101.1.103	<input checked="" type="checkbox"/> <div style="border: 1px solid #ccc; padding: 5px; font-family: monospace; font-size: 0.8em;">SHA256:ti2Uj8d S06FLbpWKEllq ttxLyOJ6S2ZAQ YqxzPpJmTQ</div>	<input checked="" type="checkbox"/> Valid
<input checked="" type="checkbox"/>	aa01-esxi-02.vcf.local	AA01-NP 	10.101.1.102	<input checked="" type="checkbox"/> <div style="border: 1px solid #ccc; padding: 5px; font-family: monospace; font-size: 0.8em;">SHA256:UvUJ8 pC48Pd25vkP6 ewIWATeWS6P z4oTHayqvJh/u tE</div>	<input checked="" type="checkbox"/> Valid
<input checked="" type="checkbox"/>	aa01-esxi-01.vcf.local	AA01-NP 	10.101.1.101	<input checked="" type="checkbox"/> <div style="border: 1px solid #ccc; padding: 5px; font-family: monospace; font-size: 0.8em;">SHA256:XKEC5 KYJICUKPKJqq hRoT+8GCqUY9 BNQa5Aa+3Bplf g</div>	<input checked="" type="checkbox"/> Valid

3 3 hosts

CANCEL

NEXT 

Step 16. Click NEXT.

 Commission Hosts

- 1 Host Addition and Validation
- 2 Review

Review

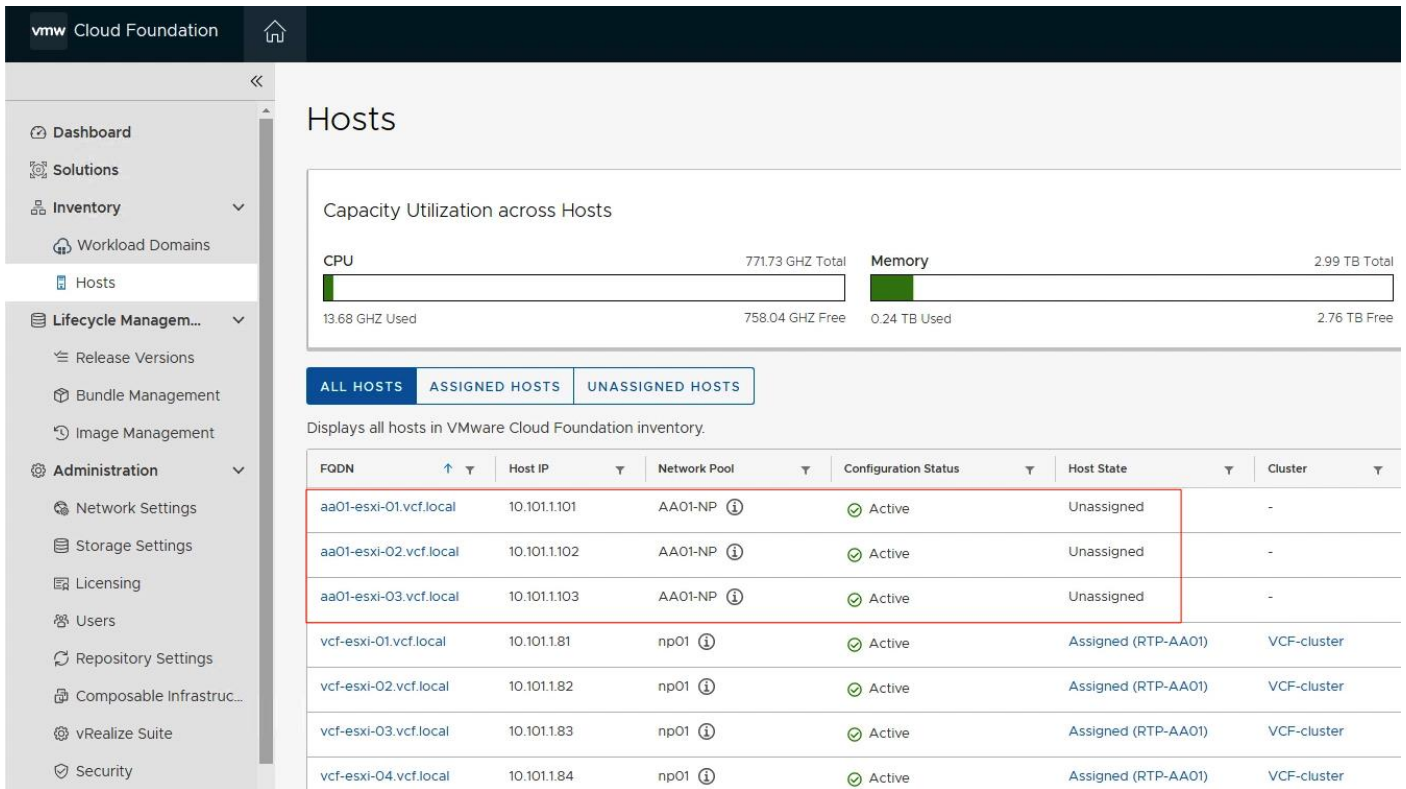
Validated Host(s)

aa01-esxi-03.vcf.local	Network Pool Name: AA01-NP IP Address: 10.101.1.103 Storage Type: VMFS on FC
aa01-esxi-02.vcf.local	Network Pool Name: AA01-NP IP Address: 10.101.1.102 Storage Type: VMFS on FC
aa01-esxi-01.vcf.local	Network Pool Name: AA01-NP IP Address: 10.101.1.101 Storage Type: VMFS on FC

Step 17. Verify the information on Review screen and click COMMISSION.

SDDC Manager will take some time to commission the hosts.

Step 18. On successful commissioning of the hosts, the hosts will appear under **Inventory > Hosts** in Active but Unassigned state.



Deploy the VI Workload Domain

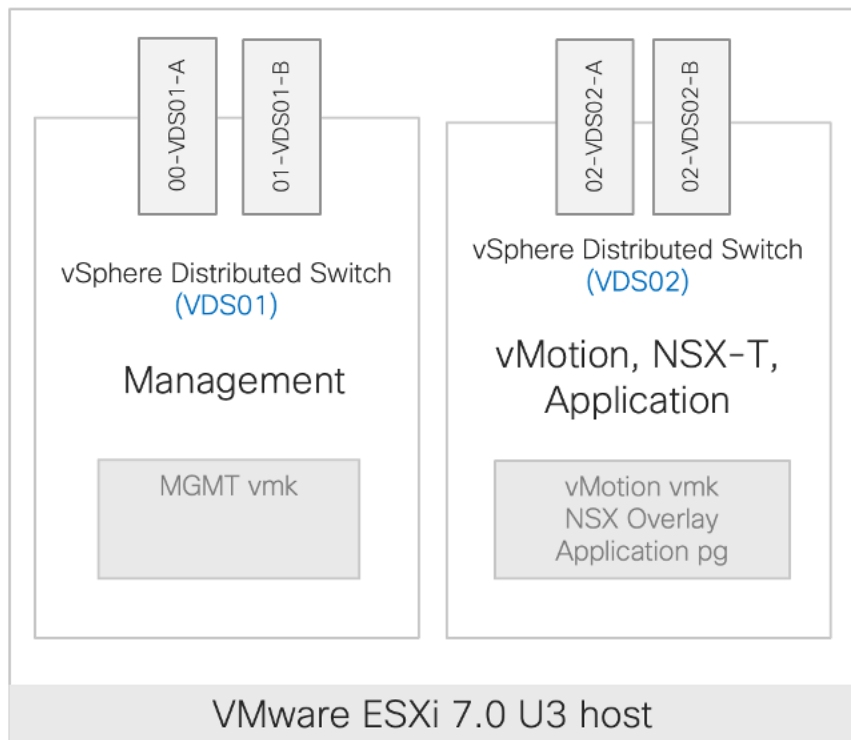
As part of VI workload onboarding, the VMware SDDC manager automatically:

- Deploys a vCenter Server Appliance for the new VI workload domain within the management domain.
- Connects the specified ESXi servers to this vCenter Server instance and groups them into a cluster. Each host is configured with the port groups applicable to the VI workload domain.
- Configures networking on each host.
- Connects to VMFS on FC storage on the ESXi hosts.

Note: By default, VI workload domains do not include any NSX Edge clusters and are isolated. To provide north-south routing and network services, utilize the traditional VLAN based application deployment or add one or more NSX Edge clusters to a VI workload domain.

VMware SDDC manager allows customers to create a new workload domain using the SDDC Manager web graphical user interface (GUI) or by creating a description file using JSON and using VMware Cloud Foundation API. The VI workload domain deployment using GUI is simpler however the GUI only supports creation of a single VDS in the ESXi host. The FlashStack ESXi hosts contain at least four vNICs and require creation of 2 VDSs so traffic can be segregated and controlled on the vNIC basis. [Figure 10](#) shows the ESXi host design including two VDS switches, vNICs assigned to each VDS, and port-groups and vmk ports created on each VDS.

Figure 10. FlashStack VI Workload Domain ESXi Host Networking Configuration

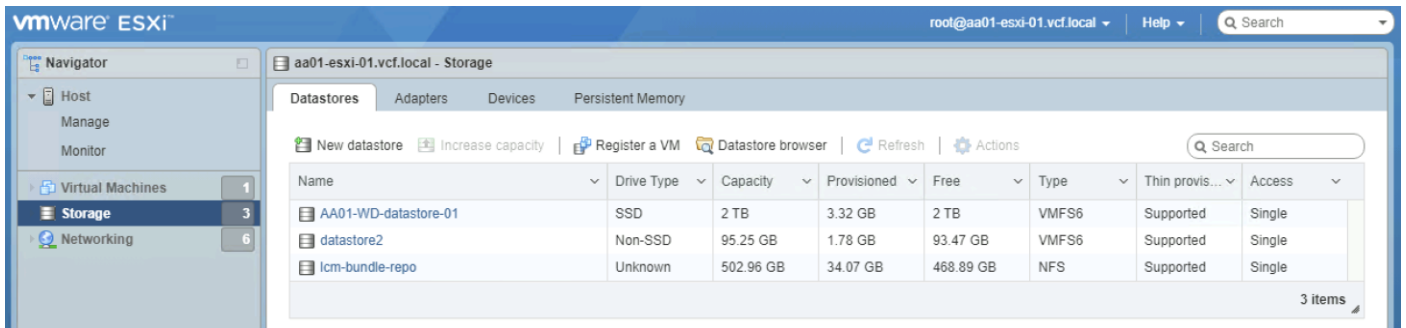


This multi-VDS configuration is completed using VMWare Cloud Foundation API. A JSON file is created with appropriate network parameters and definitions and pushed to VMWare Cloud Foundation API. Using these parameters, VMWare Cloud Foundation deploys two VDSs in the VI workload domain.

Procedure 1. Create and mount the VMFS Datastore

- Step 1.** Open a web browser and navigate to the first ESXi server’s management IP address.
- Step 2.** Enter “root” as the username.
- Step 3.** Enter the <root password>.
- Step 4.** Click **Log in** to connect.
- Step 5.** Click on **Storage** then click **New Datastore**.
- Step 6.** Click **Create new VMFS datastore**, then click **Next**.
- Step 7.** Give the new datastore a name, for example “AA01-WD-datastore-01”, then select the large unused FC LUN for VM storage created earlier on the Pure Storage FlashArray.
- Step 8.** Click **Next**.
- Step 9.** Accept the default partitioning to use the entire LUN, then click **Finish**.

The new VMFS datastore should be visible on the first VI Workload Domain host.



Step 10. Open a web browser and navigate to the second ESXi server’s management IP address.

Step 11. Enter “root” as the username.

Step 12. Enter the <root password>.

Step 13. Click **Log in** to connect.

Step 14. Click on **Storage** then click the **Adapters** tab, then click **Rescan**.

Step 15. Click on the **Datastores** tab, the new VMFS datastore should be shown. If not, return to the **Adapters** tab and click **Refresh**. It may be necessary to click on the **Devices** tab and perform a **Rescan** or **Refresh** from there before the VMFS datastore is discovered. Alternatively, the host can be rebooted.

Step 16. Repeat steps 11–15 for the remaining VI Workload Domain hosts until all of them are connected to the VMFS datastore.

Procedure 2. VI Workload Domain JSON file

In this procedure, a VI workload domain JSON description file is created. This file contains all the information necessary to deploy FlashStack VI workload domain.

Step 1. Copy the JSON file from **Appendix A** and edit the file in a text editor.

Step 2. Under the vcenterSpec section, provide following information:

- Name of workload domain (to be configured on SDDC manager)
- vCenter IP address
- vCenter FQDN
- vCenter IP gateway
- vCenter Subnet Mask
- vCenter root password
- vCenter datacenter name
- Size of VCenter deployment.

```
{
  "domainName": "AA01-WD",
  "vcenterSpec": {
    "name": "aa01-vc",
    "networkDetailsSpec": {
      "ipAddress": "10.101.1.100",
      "dnsName": "aa01-vc.vcf.local",
      "gateway": "10.101.1.254",
      "subnetMask": "255.255.255.0"
    },
    "rootPassword": "<####>",
    "datacenterName": "AA01-WD-DC",
    "vmSize": "small"
  }
}
```

```
},
```

Step 3. Obtain the ESXi host ID ("id") for all the workload domain servers from SDDC manager.

Step 4. Log into the SDDC manager and navigate to **Developer Center > API Explorer**.

Step 5. Click **APIs for managing Hosts**.

Step 6. Click **GET /v1/hosts**.

Step 7. In status field, add **UNASSIGNED_USEABLE**.

The screenshot shows the API Explorer interface for the endpoint `GET /v1/hosts`. The 'status' parameter is highlighted, and a text input field is open, containing the value `UNASSIGNED_USEABLE`. The table below shows the parameter details:

Parameter	Value	Type	Description/Data Type
status	UNASSIGNED_USEABLE	Query	Status Of The Host.One Among: ASSIGNED, UNASSIGN Data Type: String

Step 8. Scroll down and click **EXECUTE**.

Step 9. Click the response **PageOfHost**.

The screenshot shows the response for the `PageOfHost` endpoint. The response is a JSON object with the following structure:

```
{
  "elements": [
    "The list of elements included in this page"
  ],
  "Host (fa67...)",
  "Host (5a5...)",
  "Host (941b...)"
]
```

Step 10. Note the Host IDs and use these host IDs in the JSON file below.

Note: Expand each host id by clicking it to see which host the ID belongs to.

Step 11. Back in the JSON file editing, under the `ComputeSpec` section, provide following information for **all** workload domain ESXi hosts:

- Name of vCenter cluster
- Host ID (obtained from the SDDC manager API earlier)
- ESXi license key (must be present in SDDC Manager)
- VDS Switch name and the vNIC assignment

```
"computeSpec": {
  "clusterSpecs": [
    {
```

```

"name": "AA01-WD-Cluster",
"hostSpecs": [
  {
    "id": "<###>",
    "licenseKey": "<###>",
    "hostNetworkSpec": {
      "vmNics": [
        {
          "id": "vmnic0",
          "vdsName": "vds01"
        },
        {
          "id": "vmnic1",
          "vdsName": "vds01"
        },
        {
          "id": "vmnic2",
          "vdsName": "vds02"
        },
        {
          "id": "vmnic3",
          "vdsName": "vds02"
        }
      ]
    }
  }
],

```

Step 12. Under `datastoreSpec` section, enter the values in the JSON file, including the name of the VMFS datastore that was manually mounted on the hosts:

```

"datastoreSpec" : {
  "vmfsDatastoreSpec" : {
    "fcSpec": [ {
      "datastoreName" : "AA01-WD-datastore-01"
    } ]
  }
},

```

Step 13. Under `networkSpec`, provide the name of two VDS switches (`vds01` and `vds02`) and port-groups associated with each VDS. VDS switch `vds01` is used for management traffic while VDS switch `vds02` is used for NSX-T host overlay and vMotion traffic.

```

"networkSpec": {
  "vdsSpecs": [
    {
      "name": "vds01",
      "portGroupSpecs": [
        {
          "name": "vds01-pg-management",
          "transportType": "MANAGEMENT"
        }
      ]
    },
    {
      "name": "vds02",
      "isUsedByNsxt": true,
      "portGroupSpecs": [
        {
          "name": "vds02-pg-vmotion",
          "transportType": "VMOTION"
        }
      ]
    }
  ]
},

```

Step 14. Under `nsxClusterSpec`, provide the following information:

- Host Overlay VLAN (3003)
- IP address pool name (tep-pool)
- IP address pool range including IP subnet and IP gateway

```

    "nsxClusterSpec": {
      "nsxTClusterSpec": {
        "geneveVlanId": 3003,
        "ipAddressPoolSpec": {
          "name": "tep-pool",
          "subnets": [
            {
              "ipAddressPoolRanges": [
                {
                  "start": "192.168.3.101",
                  "end": "192.168.3.110"
                }
              ],
              "cidr": "192.168.3.0/24",
              "gateway": "192.168.3.254"
            }
          ]
        }
      }
    }
  ],
},

```

Step 15. Under nsxTSpec, provide the following information for all three NSX-T appliances:

- Name of the appliance (VM name)
- IP address
- FQDN
- IP gateway
- Subnet Mask

```

"nsxTSpec": {
  "nsxManagerSpecs": [
    {
      "name": "vcf-wd-nsx-1",
      "networkDetailsSpec": {
        "ipAddress": "10.101.1.96",
        "dnsName": "vcf-wd-nsx-1.vcf.local",
        "gateway": "10.101.1.254",
        "subnetMask": "255.255.255.0"
      }
    }
  ],
},

```

Step 16. Also, under the nsxTSpec, provide the following additional information:

- NSX-T VIP
- FQDN for NSX-T VIP
- NSX-T License Key
- Admin password for NSX manager
- NSX-T deployment size

```

"vip": "10.101.1.95",
"vipFqdn": "vcf-wd-nsx.vcf.local",
"licenseKey": "<###>",
"nsxManagerAdminPassword": "<###>"
"formFactor": "medium"
}
}

```

When the JSON file is updated and saved, move to the next procedure to start workload domain deployment.

Procedure 3. VI Workload Domain Creation using VMware Cloud Foundation API

In this procedure, using the VMware Cloud Foundation API and the JSON file created in the last step, the VI workload domain will be deployed.

Step 1. Log into the SDDC manager and navigate to **Developer Center > API Explorer**.

Step 2. Click on **APIs for managing Domains**.

Step 3. Click **POST /v1/domains/validations**.

APIs for managing Domains		
> GET	/v1/domains	Get the Domains
> POST	/v1/domains	Create a Domain
> GET	/v1/domains/{id}	Get a Domain
> DELETE	/v1/domains/{id}	Delete a Domain if it has been previously initialized for deletion
> PATCH	/v1/domains/{id}	Update a Domain
> GET	/v1/domains/{id}/tags	Get Tags assigned to Domain
> POST	/v1/domains/validations	Validate the input spec for domains operations
> GET	/v1/domains/{id}/endpoints	Get Endpoints of a Domain

Step 4. Copy and paste the JSON file in the domainCreationSpec box.

> POST /v1/domains/validations Validate the input spec for domains operations

> Description
No description

> Response Types

> Try it out

Parameter	Value	Type	Description/Data Type
domainCreationSpec (required)	<pre>1 { 2 "domainName": "AA01-WD", 3 "vcenterSpec": { 4 "name": "aa01-vc", 5 "networkDetailsSpec": { 6 "ipAddress": "10.101.1.100", 7 "dnsName": "aa01-vc.vcf.local", 8 "gateway": "10.101.1.254", 9 "subnetMask": "255.255.255.0" 10 }, 11 "rootPassword": "■■■■■■■■", 12 "datacenterName": "AA01-WD-DC", 13 "vmSize": "small" 14 }, 15 "computeSpec": { 16 "clusterSpecs": [17 { 18 "name": "AA01-WD-Cluster", 19 "hostSpecs": [20 { 21 {</pre>	Body	Domain Creation Spec DomainCreationSpec{ ... }

EXECUTE COPY JSON DOWNLOAD

Step 5. Click **EXECUTE** to validate the specification file.

Step 6. Click **CONTINUE** to proceed with validation.

Step 7. SDDC manager will take some time to validate the specification file. When the validation is complete, click on the **Validation** link under Response.

Step 8. Verify that the validation was successful.

Response

Validation (2c2819da-d775-4bf2-8c0e-b4d2d40817cf)   {

"description":

Description of the validation

"Validating Domain Creation Spec",

"executionStatus":

Execution status of the validation

"COMPLETED",

"id":

ID of the validation

"2c2819da-d775-4bf2-8c0e-b4d2d40817cf",

"resultStatus":

Result status of the validation after it has completed its execution

"SUCCEEDED",

"validationChecks":

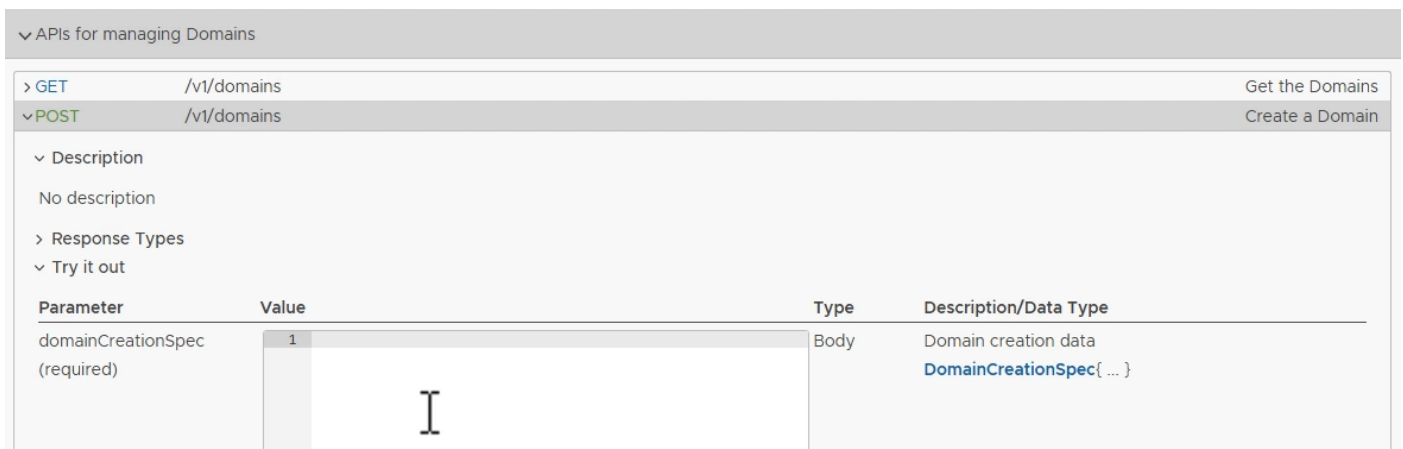
List of one or more validation checks that are performed as part of the validation

[

ValidationCheck   { ... },

]

Step 9. Click to expand **POST /v1/domains**.



API Explorer interface showing the POST /v1/domains endpoint expanded. The 'domainCreationSpec' parameter is highlighted in a table.

Parameter	Value	Type	Description/Data Type
domainCreationSpec (required)	1	Body	Domain creation data DomainCreationSpec { ... }

Step 10. Copy and paste the specification JSON file in the domainCreationSpec box.

Step 11. Click **EXECUTE**.

Step 12. Click **CONTINUE** to proceed with domain creation.

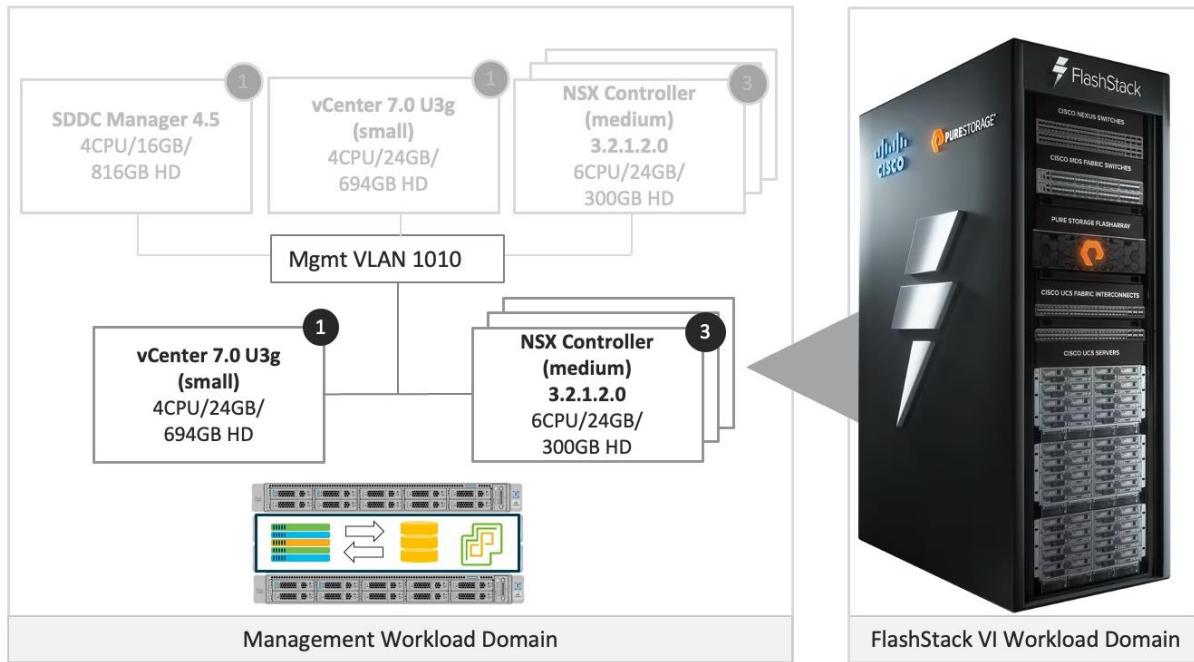
It will take a while for SDDC manager to deploy the VI workload domain. Customers can monitor the Task panel in SDDC manager to check the current task being executed.

Customers can also log into the management vCenter to monitor the tasks related to vCenter and NSX-T VM deployment.

On the successful deployment of the VI workload domain, a new vCenter and 3 NSX controller VMs will be deployed on the VMware Cloud Foundation management domain for management of the new workload domain as shown in [Figure 11](#).

Note: In a customer environment, the size of the VMs shown in [Figure 11](#) Figure 11. can be different depending on the size of the deployment selected during workload domain deployment.

Figure 11. VMware Cloud Foundation VI workload Domain



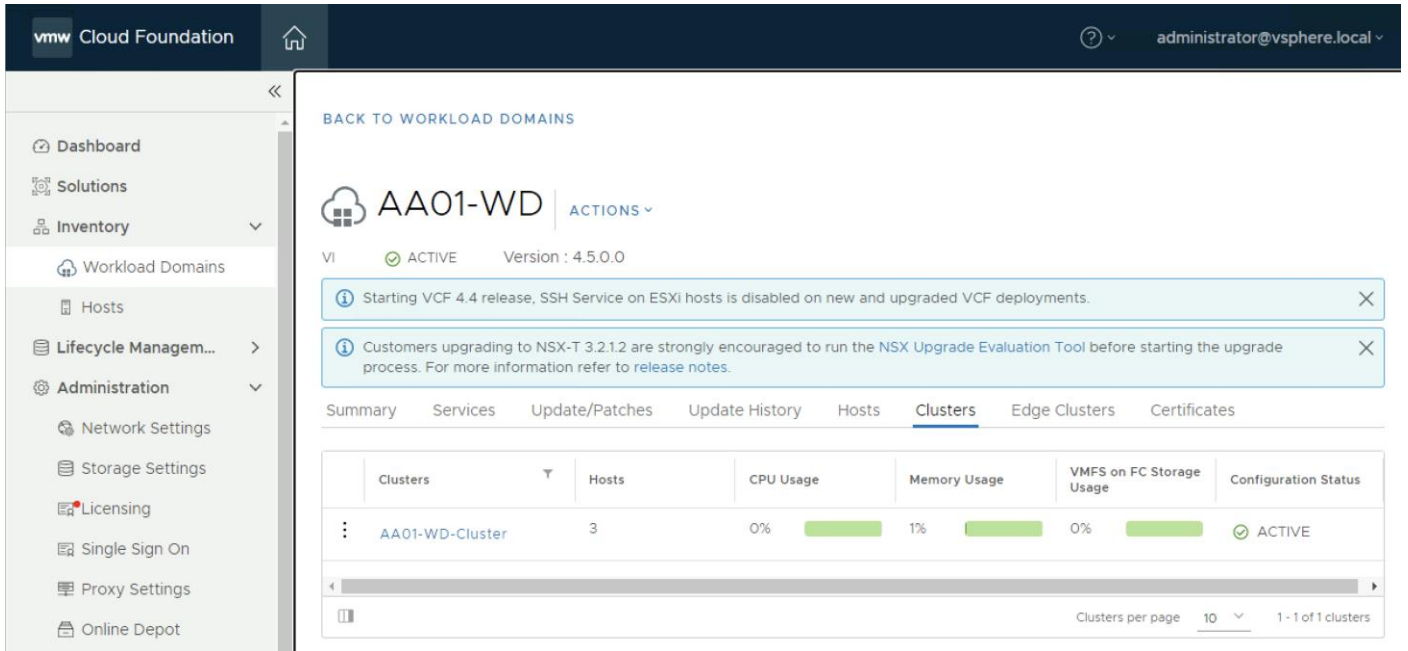
Step 13. Log into the VMware SDDC manager and navigate to **Inventory > Workload Domains** to find various deployment parameters for the newly created VI workload domain.

The screenshot shows the VMware Cloud Foundation SDDC Manager interface. The main window displays the 'Workload Domain' page, which includes a navigation sidebar on the left and a main content area. The main content area shows 'Capacity Utilization across Domains' with four progress bars for CPU, Memory, vSAN Storage, and NFS Storage. Below this, a table lists the workload domains.


Domain	Type	CPU Usage	Memory Usage	vSAN Storage Usage	NFS Storage Usage	VMFS on FC Storage Usage	vVol Storage Usage	Configuration Status	Owner	Cluster	Hosts	Updates Available
RTP-AA01	MANAGEMENT	6%	21%	2%	-	-	-	ACTIVE	administrator@vsphere.local	1 Cluster	4	Up-to-da
AA01-WD	VI	0%	2%	-	0%	-	-	ACTIVE		1 Cluster	3	Up-to-da

Step 14. Click the workload domain name to gather more information about NSX Manager IP address and host information.

Step 15. From the VI workload domain page, click on **Clusters** in the main window and select the WD cluster.



Step 16. Click **ACTIONS** next to the cluster name and select **Open in vSphere Client**.



Step 17. Log into the vCenter deployed for the VI workload domain.

Now the VMware Cloud Foundation deployment is complete, and the VI workload domain is onboarded.

FlashStack VI Workload Domain Configuration

This chapter contains the following:

- [Finalize the VI Workload Domain ESXi Host Configuration](#)

After successfully adding the FlashStack hosts as VMware Cloud Foundation VI workload domain, following additional configuration steps need to be completed on ESXi hosts and the Pure Storage controllers.

Finalize the VI Workload Domain ESXi Host Configuration


The following configuration steps need to be completed on all the workload domain ESXi hosts:

- Mount additional VMFS datastore for swap files
- Change the swap file location on the ESXi hosts
- Configure the ESXi hosts power policy
- Add application port-groups to VDS
- Backup ESXi host keys for migration or host restoration after failure

Procedure 1. Mount additional VMFS datastore(s) on the VI Workload Domain Hosts

Step 1. Log in to the vSphere HTML client of the VI workload domain vCenter server.

Step 2. From the Web Navigator left navigation pane, select the correct data center, and select the **Datstores** tab.



Step 3. Right click on data center and select **Storage > New Datastore...** to add new datastore.

Step 4. In the New datastore popup, select **VMFS** and click **NEXT**.

Step 5. Enter `infra_swap` for the datastore name, select one of the workload domain hosts from the dropdown menu, then select the unused FC LUN from the list. Click **NEXT**.

New Datastore

1 Type

2 Name and device selection

3 VMFS version

4 Partition configuration

5 Ready to complete

Name and device selection

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

Name

The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

Select a host

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Clustered VMDK Supported
<input checked="" type="radio"/>	PURE Fibre Channel Disk (...)	253	500.00 GB	Supported	Flash	512n	Yes
<input type="radio"/>	PURE Fibre Channel Disk (...)	1	20.00 GB	Supported	Flash	512n	Yes

2 Items

CANCEL BACK NEXT

Step 6. Select VMFS6 then click **NEXT**.

Step 7. Accept the default partition configuration to use the entire disk, then click **NEXT**.

Step 8. Verify the information and click **FINISH**.

The datastore should now appear in the datastore list.

Step 9. Repeat this procedure for any additional datastores.

Procedure 2. Configure System Swap Location on the ESXi Host(s)

Step 1. In the VI workload domain vCenter Interface, under **Hosts and Clusters** select the ESXi host.

Step 2. In the center pane, select the **Configure** tab.

Step 3. In the list under **System**, select **System Swap**.

Step 4. In the right pane, click **EDIT**.

Step 5. Select **Can use datastore** and select `infra_swap` from the drop-down list.

Edit System Swap Settings | aa01-esxi-02.vcf.lo X
cal

- Can use datastore:
- Can use host cache
- Can use datastore specified by host for swap files

CANCEL OK

Step 6. Click **OK** to save the configuration changes.

Step 7. Repeat this procedure for all the ESXi hosts.

Procedure 3. Configure VM Swap File Location

- Step 1.** In the VI workload vCenter Interface, under **Hosts and Clusters** select the Workload Domain Cluster.
- Step 2.** In the center pane, select the **Configure** tab.
- Step 3.** Next to the section named “Swap file location”, click **EDIT**.
- Step 4.** Change the selection to “Datastore specified by host”, then click **OK**.
- Step 5.** In the VI workload vCenter Interface, under **Hosts and Clusters** select the ESXi host.
- Step 6.** In the center pane, select the **Configure** tab.
- Step 7.** In the list under Virtual Machines, select **Swap File Location**.
- Step 8.** In the window on the right, click **EDIT**.
- Step 9.** Select **Use a specific datastore** and select infra_swap.

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
<input type="radio"/>	infra_datastore...	1000 GB	179.56 GB	989.44 GB	NFS	Supported
<input type="radio"/>	lcm-bundle-repo	502.96 GB	45.06 GB	457.9 GB	NFS	Supported
<input type="radio"/>	nfs_ds_01	10 GB	38.25 MB	9.96 GB	NFS	Supported
<input type="radio"/>	datastore3	223.25 GB	1.79 GB	221.46 GB	VMFS	Supported
<input checked="" type="radio"/>	infra_swap	300 GB	166.37 MB	299.84 GB	NFS	Supported

5 items

- Step 10.** Click **OK** to save the configuration changes.
- Step 11.** Repeat this procedure for all the ESXi hosts.

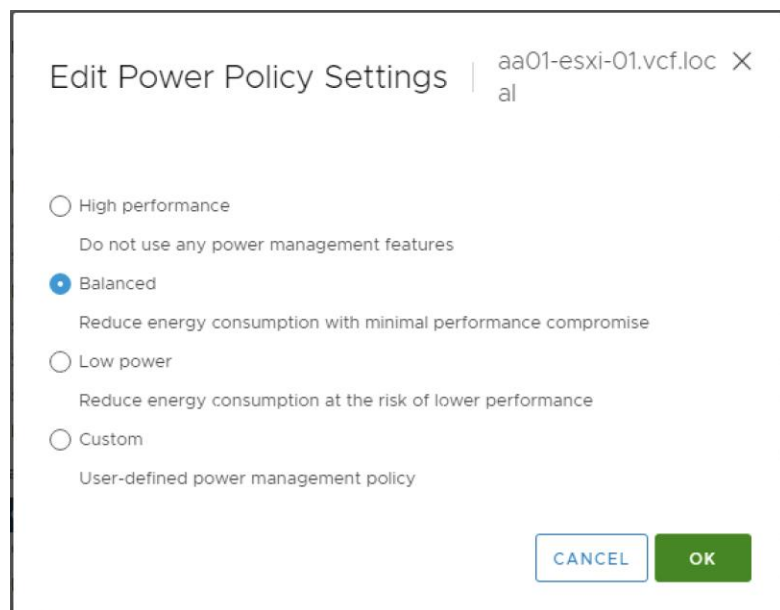
Procedure 4. Configure Host Power Policy on the ESXi Host

Note: Implementation of this policy is recommended in the Performance Tuning Guide for Cisco UCS M6 Servers: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html> for maximum VMware ESXi performance. In the past, the High-Performance setting was recommended. The steps below show using the Balanced performance policy selection which generally results in high performance levels with moderate levels of power consumption. Customers can adjust this policy based on their requirements. Selecting the High-Performance setting will increase performance slightly, but also result in significant increases in power consumption and heat generation.

- Step 1.** In the VI workload vCenter Interface, under **Hosts and Clusters** select the ESXi host.
- Step 2.** In the center pane, select the **Configure** tab.
- Step 3.** In the center pane, select **Hardware > Overview**.

Step 4. Scroll down and click **EDIT POWER POLICY** under Power Management.

Step 5. Select **Balanced**.



Step 6. Click **OK** to save the configuration changes.

Step 7. Repeat these steps for all the ESXi hosts.

Procedure 5. Configure the Application port-group on VDS (optional)

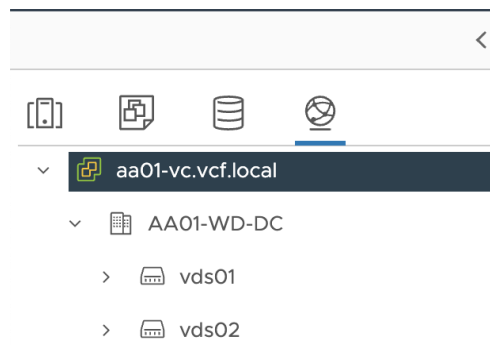
VMware Cloud Foundation deploys NSX-T controllers and integrates NSX with the VDS (vds02) as part of VCF deployment. Customers can start using NSX-T to route and control their application traffic.

Note: NSX-T configuration is not explained as part of this deployment guide.

This procedure explains setting up traditional VLAN based application networking. Customers can utilize application port-groups, defined on the VDS (vds02), to attach their application VMs to various VLANs. In the procedure below, a single application VLAN (1012) will be added to a port-group defined on VDS vds02. This port-group will allow application VMs to communicate with the enterprise network using the gateway defined on the Cisco Nexus switches.

Step 1. Log into the VI workload domain vCenter.

Step 2. Select **Networking** under Menu and expand the appropriate datacenter.



Step 3. Click the VDS **vds02**.

Step 4. Right-click and select **Distributed Port Group > New Distributed Port Group**.

- Step 5.** Provide a Name (for example, App-PG) and click **NEXT**.
- Step 6.** For the VLAN type, select **VLAN** from the drop-down list.
- Step 7.** Enter the Application VLAN ID (for example, 1012) and click **NEXT**.

New Distributed Port Group

- 1 Name and location
- 2 Configure settings
- 3 Ready to complete

Configure settings

Set general properties of the new port group.

Port binding Static binding v

Port allocation Elastic v ⓘ

Number of ports 8

Network resource pool (default) v

VLAN

VLAN type VLAN v

VLAN ID 1012

Advanced

Customize default policies configuration

- Step 8.** Review the configuration and click **FINISH**.
- Step 9.** Repeat this procedure for all the application VLANs.

Procedure 6. Backup the ESXi Recovery Keys

FlashStack ESXi hosts are configured for boot from SAN using Fibre Channel which allows stateless compute setup. The stateless compute allows a server profile to move from one compute node to another seamlessly in case of failure or hardware upgrade. Starting with ESXi 7.0 Update 2, compute nodes containing a Trusted Platform Module (TPM) and configured for UEFI boot save the sensitive information in the TPM and require a recovery key to successfully migrate or recover the ESXi host on a new/different compute node. This procedure explains backing up of the recovery keys from all the VI workload domain ESXi hosts.

- Step 1.** Log into Cisco Intersight and select **Infrastructure Service**.
- Step 2.** Click on **Servers**.
- Step 3.** Select the VI workload domain ESXi server and click **...** and select **Launch vKVM**.
- Step 4.** Click through the certificate prompts and lunch the KVM.
- Step 5.** Press **F2** and log into the ESXi host using root.
- Step 6.** Scroll down to **Troubleshooting Mode Options** and select **Enable SSH**.
- Step 7.** Connect to the management IP address of the ESXi host using an SSH client.
- Step 8.** Use root as username and password set for the host.
- Step 9.** Run the following command on the ESXi host CLI:

```
[root@aa01-esxi-01:~] esxcli system settings encryption recovery list
Recovery ID                               Key
-----
{54B47EDC-EEE3-4949-86B6-758633DA312B} 240691-xxxxxx-112774-307101-xxxxxx-339487-xxxxxx-362831-xxxxxx-
354968-xxxxxx-091124-xxxxxx-312259-xxxxxx-390449
```

Step 10. Save the recovery key in a safe location.

Step 11. Exit the SSH session.

Step 12. Log back into the KVM console for the host and disable SSH.

Step 13. Close the KVM console.

Procedure 7. Using the Recovery Keys for Server Profile Migration or Recovery on a New Compute Node

To recover the ESXi configuration when migrating ESXi host from one compute node to another, refer to this VMware article: <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-23FFB8BB-BD8B-46F1-BB59-D716418E889A.html#GUID-23FFB8BB-BD8B-46F1-BB59-D716418E889A>.

Step 1. After associating the server profile with new compute node, log into Intersight and launch the KVM console for the server.


Step 2. Boot the server and stop the boot process by pressing **Shift + O** at the ESXi boot screen.

Step 3. Type **encryptionRecoveryKey=** in the KVM console.

Step 4. Retrieve the recovery key from the previously stored safe location and copy it to clipboard.

Step 5. On the KVM console, select **File > Paste Clipboard Text**.

Step 6. Paste the contents of the recovery key to add them immediately after **encryptionRecoveryKey=**.



Step 7. Once the recovery key is posted correctly, press **Enter** to continue.

Step 8. If the new compute node needs to be permanently associated with the server profile, SSH into the ESXi host (might need to enable SSH as explained above) and issue the following command:

```
[root@aa01-esxi-01:~] /sbin/auto-backup.sh
```

Note: If a recovery key is not provided during server profile migration, following output is observed on the KVM console:

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.vmu.com/security-violation>

No port for remote debugger.

FlashStack Management Tools Setup

This chapter contains the following:

- [Cisco Data Center Network Manager \(DCNM\)-SAN](#)
- [Cisco Intersight Assist](#)
- [Cisco Intersight Cloud Orchestration](#)
- [Pure Storage vSphere Client Appliance](#)

This chapter explains the various management tools that will be installed for configuring and managing FlashStack VI workload domain hosts and infrastructure.

Cisco Data Center Network Manager (DCNM)-SAN

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.

Note: Cisco DCNM-SAN is available as SAN Controller persona in Nexus Dashboard Fabric Controller (NDFC) and available exclusively on the Cisco Nexus Dashboard (ND) as an App. You can now enable the features you want at runtime (Fabric Controller (LAN), SAN Controller, and Fabric Discovery) which allows your clusters to scale better.

With the introduction of NDFC Release 12, users get a consistent experience across NDFC, and other services host-ed on Nexus Dashboard including Insights and Orchestrator. As of publishing date of the document, Cisco DCNM 11.5(4) was used in the document as it was the suggested release. The future FlashStack design documents will use NDFC version 12 or higher.

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 the Cisco Nexus C93360YC-FX2 for SAN switching, it does not support SAN Analytics.

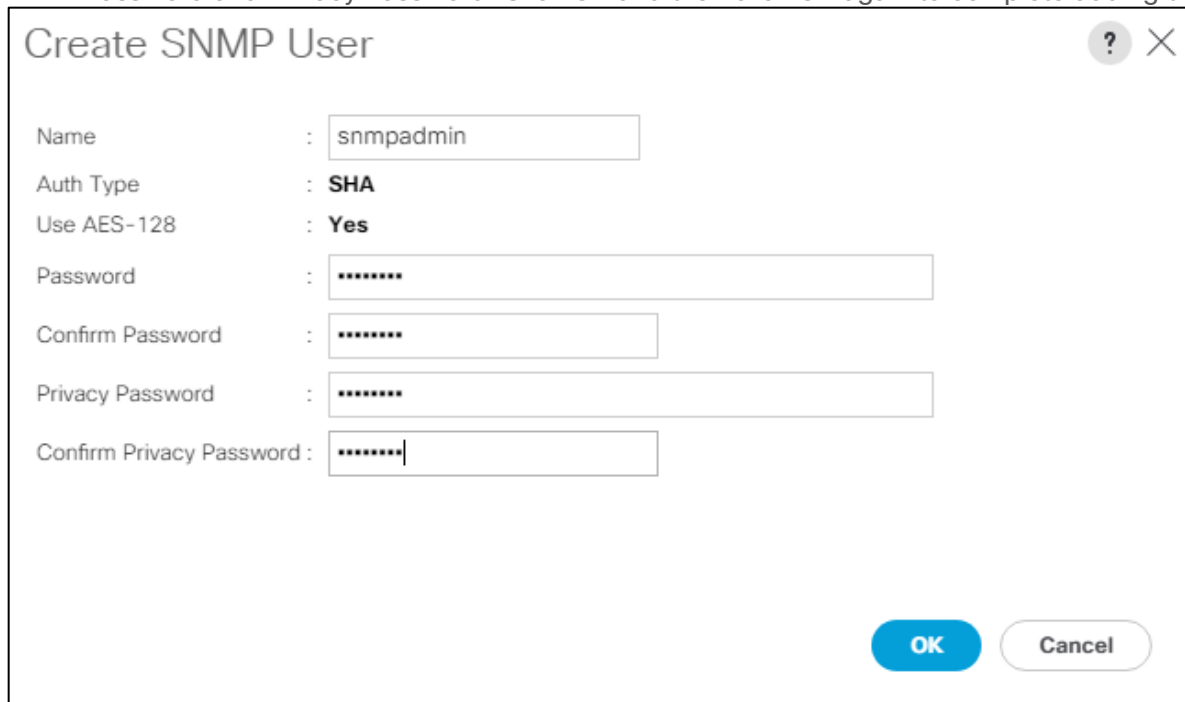
- **Passwords:** Cisco DCNM-SAN passwords should adhere to the following password requirements:
 - It must be at least eight characters long and contain at least one alphabet and one numeral.
 - It can contain a combination of alphabets, numerals, and special characters.
 - Do not use any of these special 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 snmpadmin 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.

- **DCNM SNMPv3 user in UCSM:** An SNMPv3 user needs to be added to UCSM to allow DCMN to query the LAN side of the fabric interconnects. In Cisco UCS Manager, click Admin. Navigate to All > Communication Management > Communication Services. Under SNMP, click Enabled, click Save Changes, and then click OK. Under SNMP Users, click Add. Enter the username and enter and confirm the Password and Privacy Password. Click OK and then click OK again to complete adding the user.



The screenshot shows a 'Create SNMP User' dialog box with the following fields and values:

- Name: snmpadmin
- Auth Type: SHA
- Use AES-128: Yes
- Password: [Redacted]
- Confirm Password: [Redacted]
- Privacy Password: [Redacted]
- Confirm Privacy Password: [Redacted]

Buttons: OK (blue), Cancel (grey)

Procedure 1. Deploy the Cisco DCMN-SAN OVA

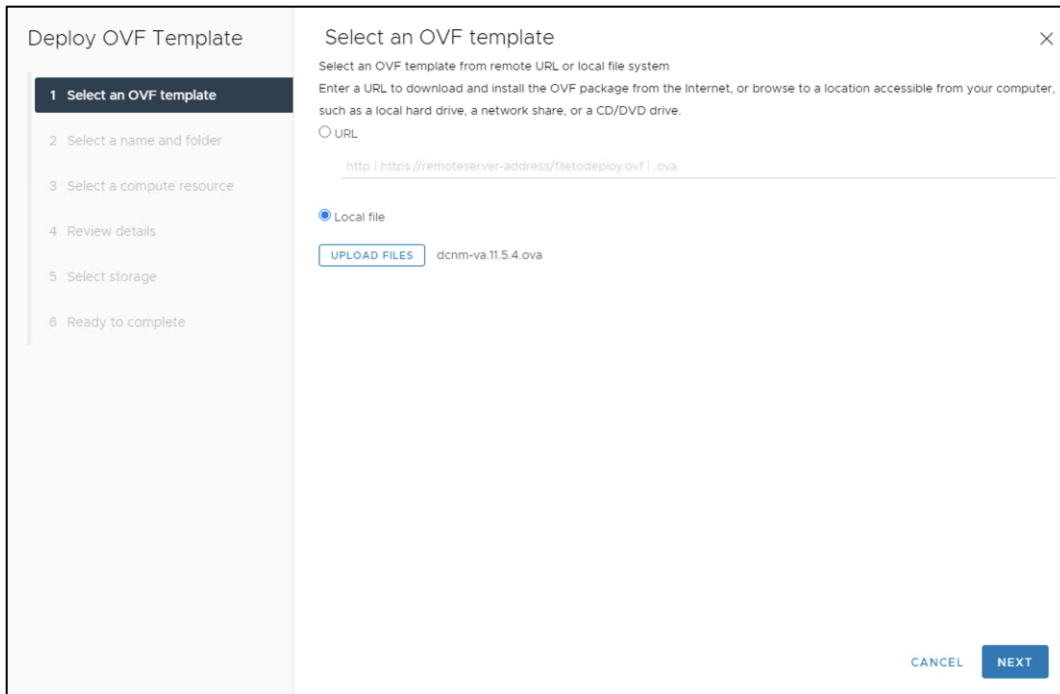
Step 1. Download the Cisco DCMN 11.5.1 Open Virtual Appliance for VMware from [https://software.cisco.com/download/home/281722751/type/282088134/release/11.5\(4\)](https://software.cisco.com/download/home/281722751/type/282088134/release/11.5(4))

Step 2. Extract dcmn-va.11.5.4.ova from the ZIP file.

Step 3. In the VMware vCenter HTML5 interface, click **Menu > Hosts and Clusters**.

Step 4. Right-click the **FlashStack-Management** cluster and click **Deploy OVF Template**.

Step 5. Select **Local file** then click **UPLOAD FILES**. Navigate to select **dcmn-va.11.5.4.ova** and click **Open**. Click **NEXT**.



Step 6. Name the virtual machine and select the **FlashStack-DC** datacenter. Click **NEXT**.

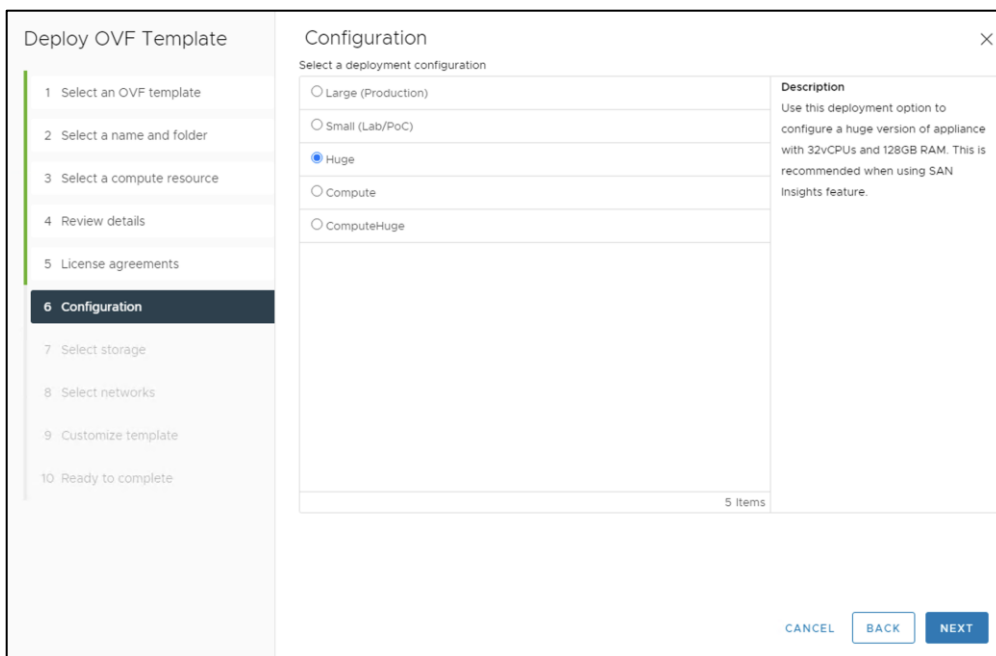
Step 7. Select the **FlashStack-Management cluster** and click **NEXT**.

Step 8. Review the details and click **NEXT**.

Step 9. Scroll through and accept the license agreements. Click **NEXT**.

Step 10. 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 11. Select **Infra-DataStore1** and the **Thin Provision virtual disk** format. Click **NEXT**.

Step 12. Select **IB-MGMT Network** for all three Source Networks. 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 License agreements
- 6 Configuration
- 7 Select storage
- 8 Select networks
- 9 Customize template
- 10 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
dcnm-mgmt	IB-Mgmt ▼
enhanced-fabric-mgmt	IB-Mgmt ▼
enhanced-fabric-inband	IB-Mgmt ▼

3 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL
BACK
NEXT

Step 13. 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 FlashStack deployment, set this field to 32. Click **NEXT**.

Step 14. 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 ✕

Click Finish to start creation.

Name	DCNM
Template name	dcnm
Download size	5.3 GB
Size on disk	Unknown
Folder	FlashStack-DC
Resource	FlashStack-Management
Storage mapping	1
All disks	Datastore: Infra-DataStore1; Format: Thick provision lazy zeroed
Network mapping	3
dcnm-mgmt	IB-Mgmt
enhanced-fabric-mgmt	IB-Mgmt
enhanced-fabric-inband	IB-Mgmt
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual
Properties	1.IP Address = 10.1.164.41 2.Subnet Mask = 255.255.255.0

CANCEL
BACK
FINISH

Step 15. 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.

Edit Settings ✕

Virtual Hardware VM Options

[ADD NEW DEVICE](#)

▼ CPU *	32 ▼	i
Cores per Socket	16 ▼ Sockets: 2	
CPU Hot Plug	<input type="checkbox"/> Enable CPU Hot Add	
Reservation	0 <input type="text"/> ▼ MHz ▼	
Limit	Unlimited <input type="text"/> ▼ MHz ▼	
Shares	Normal ▼ 32000	
CPUID Mask	Expose the NX/XD flag to guest ▼ Advanced...	
Hardware virtualization	<input type="checkbox"/> Expose hardware assisted virtualization to the guest OS	
Performance Counters	<input type="checkbox"/> Enable virtualized CPU performance counters	

CANCEL
OK

Step 16. Click **OK** to complete the change.

Step 17. Right-click the newly deployed DCNM VM and click **Open Remote Console**. Once the console is up, click the **play** icon to power on the VM. Once the VM has powered up, point a web browser to the URL displayed on the console.

```

*****
* PREPARING THE APPLIANCE... *
*****

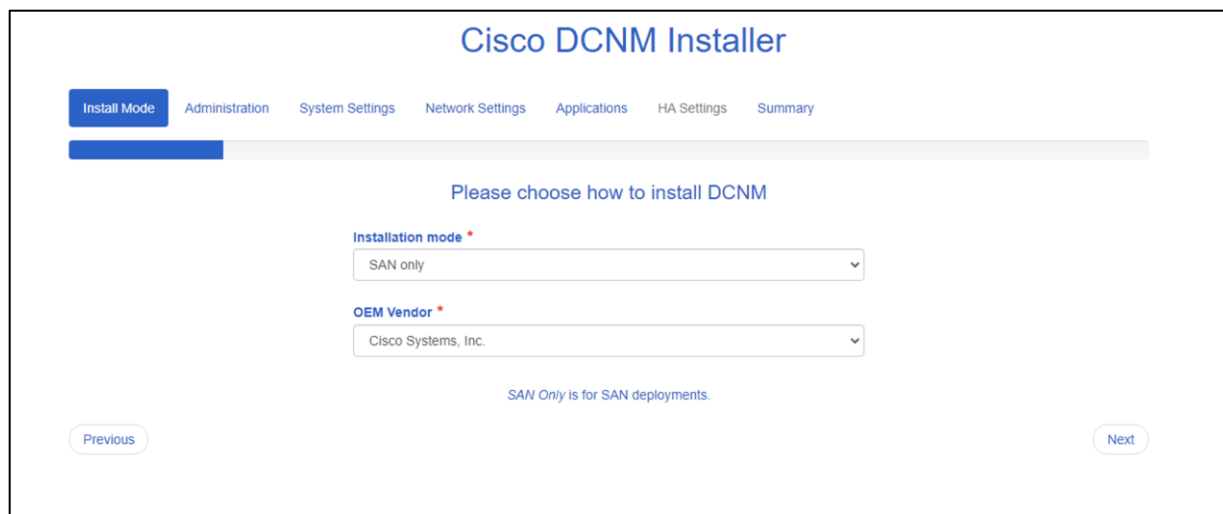
*****
Please point your web browser to
https://10.103.1.154:2443
to complete the installation
*****

```

Step 18. Navigate the security prompts and click **Get started**.

Step 19. Make sure Fresh installation – Standalone is selected and click **Continue**.

Step 20. Select **SAN only** for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click **Next**.



Step 21. Enter and repeat the administrator and database passwords and click **Next**.

Step 22. 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**.

Step 23. The Management Network settings should be filled in. For Out-of-Band Network, a different IP address in the same subnet as the management address should be used. Only input the IPV4 address with prefix. Do not put in the Gateway IPv4 Address. Scroll down and click **Next**.

Step 24. Leave Internal Application Services Network set at the default setting and click **Next**.

Step 25. Review the Summary details and click **Start installation**.

Step 26. When the Installation status is complete, click **Continue**.

Step 27. 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

Note: When the DCNM installation is complete, the browser should redirect to the DCNM management URL.

Step 1. Log in as admin with the password entered above.

Step 2. On the message that appears, select **Do not show this message again** and click **No**.

Step 3. 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 4. In the menu on the left, click **Inventory > Discovery > LAN Switches**.

Step 5. Click the **plus icon** to add LAN switches. In the Add LAN Devices window, enter the mgmt0 IP address of Cisco Nexus switch A in the Seed Switch box. Enter the snmpadmin user name and password set up in the Prerequisites section. Set Auth-Privacy to SHA_AES. Click **Next**.

Add LAN Devices

Discovery Type: Hops from seed switch Switch list

Seed Switch:

Max Hops from Seed:

User Name:

Password:

Auth-Privacy:

Add Switches To Group:

Scan Time:

Step 6. The LAN switch discovery will take a few minutes. In the LAN Discovery list that appears, the two Cisco Nexus switches and two fabric interconnects that are part of this FlashStack should appear with a status of “manageable”. Using the checkboxes, select the two Cisco Nexus switches and the two fabric interconnects that are part of this FlashStack. Click **Add**.

Step 7. After a few minutes, click **Refresh**, the two Cisco Nexus switches and two fabric interconnects that are part of this FlashStack will appear with detailed information. The SSH warning under SNMP Status can be ignored since only SNMP can be used to monitor fabric interconnects.

The screenshot shows the Cisco Data Center Network Manager interface. The left sidebar contains navigation options: Dashboard, Topology, Inventory (selected), Monitor, Configure, Administration, and Applications. The main content area displays the 'Inventory / Discovery / LAN Switches' page. A table lists four discovered switches with their details.

	Switch	IP Address	Serial No	Managed	Group	User	Auth/Priv...	Role	Last Updated Time
1	<input type="checkbox"/> aa02-6536-A	10.102.0.18	FDO25370AM2	true	Default_LAN	snmpadmin	SHA_AES		2022-10-06 23:32:48
2	<input type="checkbox"/> aa02-6536-B	10.102.0.19	FDO25370AN0	true	Default_LAN	snmpadmin	SHA_AES		2022-10-06 23:32:53
3	<input type="checkbox"/> aa03-93360-a	10.103.0.3	FDO262304Y8	true	Default_LAN	snmpadmin	SHA_AES		2022-10-06 23:32:57
4	<input type="checkbox"/> aa03-93360-b	10.103.0.4	FDO26230JUX	true	Default_LAN	snmpadmin	SHA_AES		2022-10-06 23:32:57

Step 8. In the menu, click **Inventory > Discovery > SAN Switches**.

Step 9. Click the **plus icon** to add a switching fabric.

Step 10. 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 the Cisco Nexus C93360YC-FX2 switches are being used for SAN switching, substitute them here for MDS 9132Ts. They will need to be added again under SAN switches since LAN and SAN switching are handled separately in DCNM.

Add Fabric

Fabric Seed Switch:

SNMP: Use SNMPv3/SSH

Auth-Privacy:

User Name:

Password:

Limit Discovery by VSAN

Enable NPV Discovery in All Fabrics

Step 11. Repeat steps 1–9 to add the second Cisco MDS 9132T and Fabric Interconnect. The two SAN fabrics appear in the Inventory.

Name	SeedSwitch	Status	SNMPv3/SSH	User/Cmnty	Auth/P...	Included VSAN List	Excluded VSAN List
Fabric_AA03-9132T-1	10.103.0.7	managedContinuously	true	snmpadmin	SHA_AES		
Fabric_AA03-9132T-2	10.103.0.8	managedContinuously	true	snmpadmin	SHA_AES		

Step 12. Select **Inventory > Discovery > Virtual Machine Manager**.

Step 13. Click the **plus icon** to add the vCenter.

Step 14. In the Add VCenter window, enter the IP address of the vCenter VCSA. Enter the administrator@vsphere.local user name and password. Click **Add**.

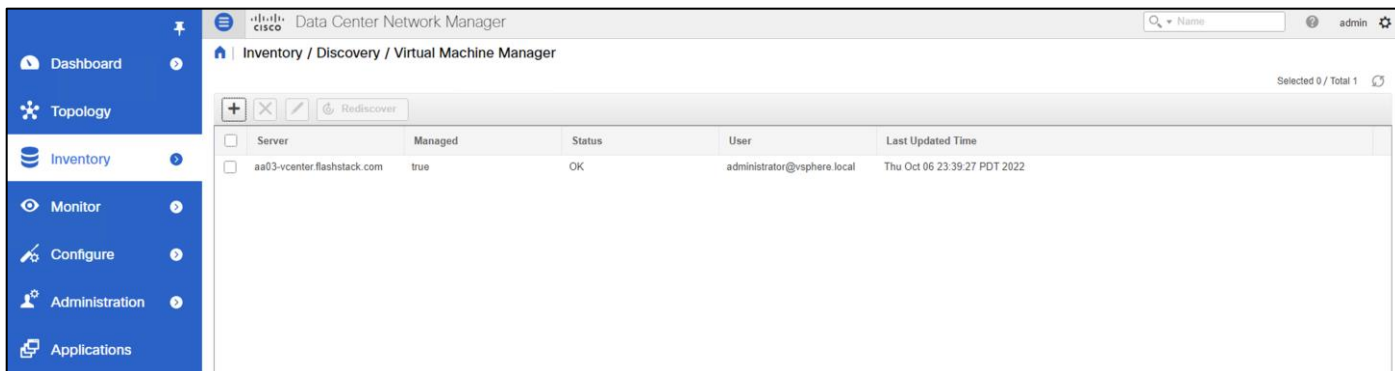
Add VCenter

Virtual Center Server:

User Name:

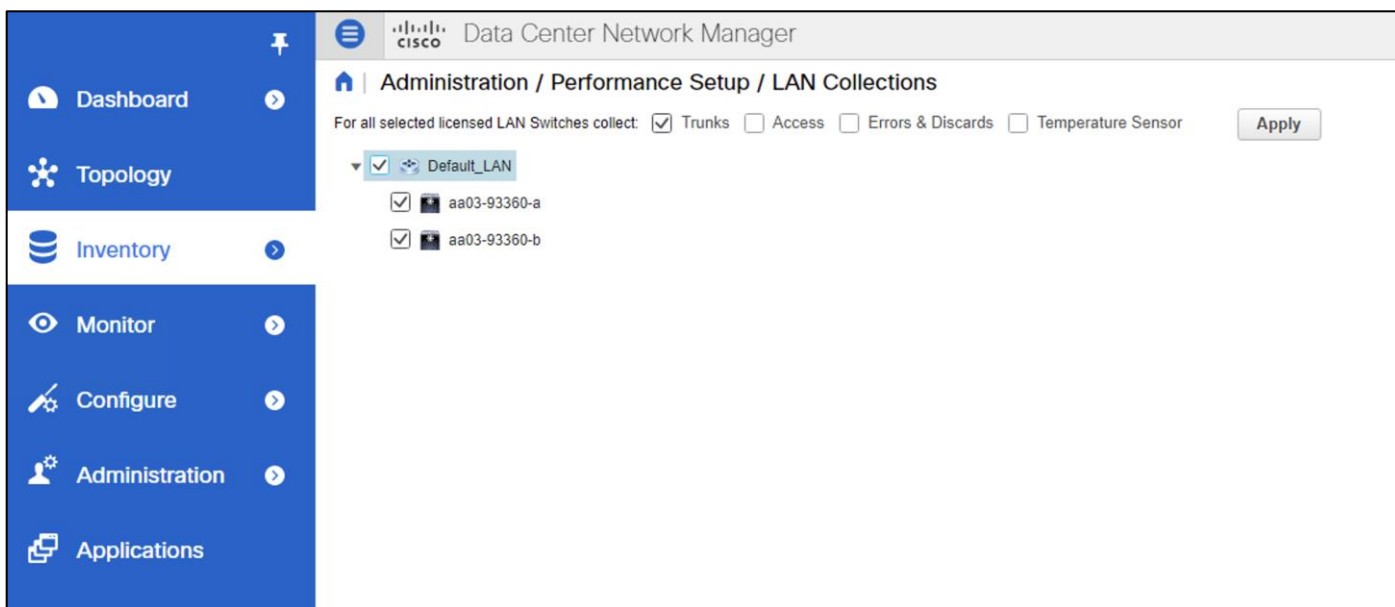
Password:

The vCenter server should now appear in the inventory.



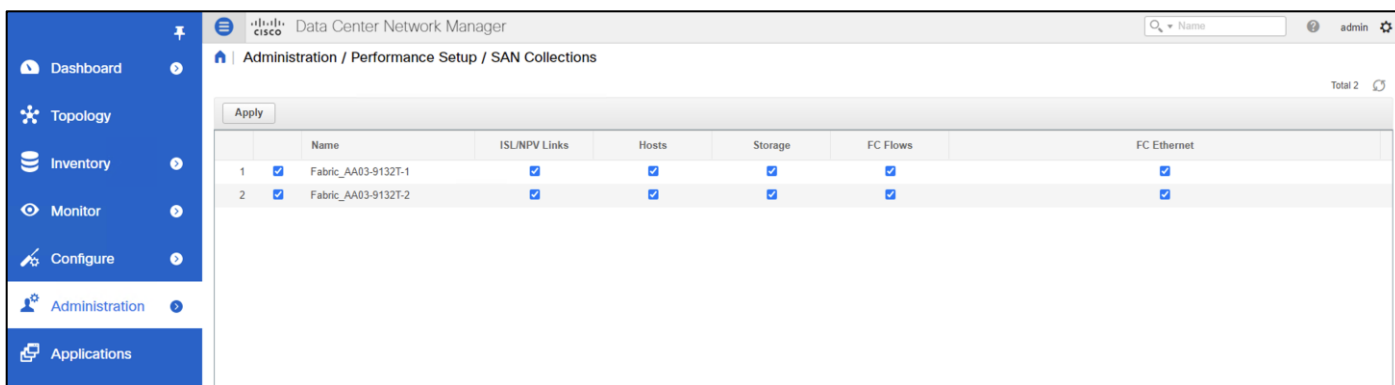
Step 15. Select **Administration > Performance Setup > LAN Collections**.

Step 16. Select the **Default_LAN** group and all information you would like to collect. Click **Apply**. Click **Yes** to restart the Performance Collector.

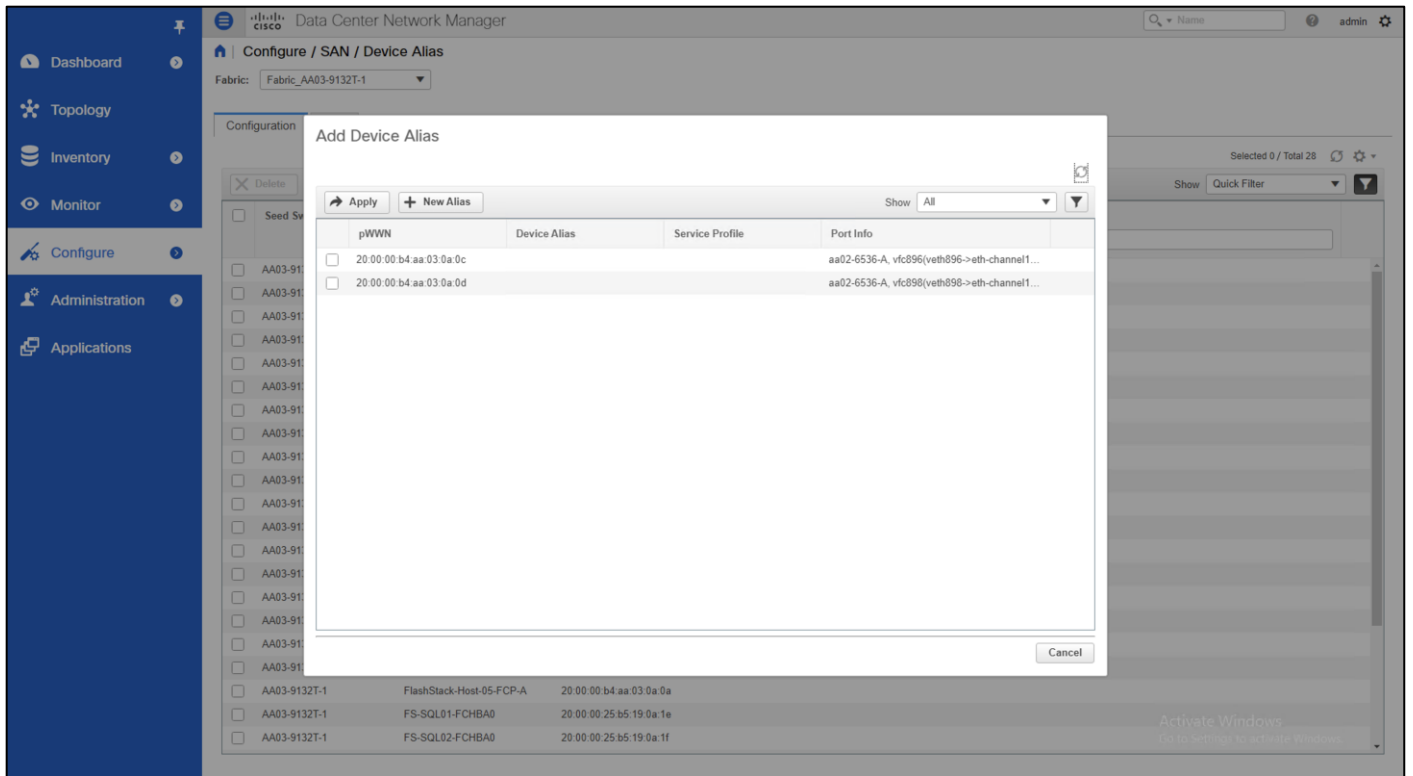


Step 17. Select **Administration > Performance Setup > SAN Collections**.

Step 18. Select both fabrics. Select the information you would like to collect and click **Apply**. Click **Yes** to restart the Performance Collector.



Step 19. Select **Configure > SAN > Device Alias**. Since the 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 20. Select **Configure > SAN > Zoning**. Just as the Device Aliases can be created and deleted from DCNM, zones can be created, deleted, and modified in DCNM and pushed to the MDS switches. Remember to enable Smart Zoning and to Zone by Device Alias.

You can now explore all 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 provides more comprehensive end-to-end flow-based data from the 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 the installation of a switch-based 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.
- For more information on SAN Insights, see the SAN Insights sections: [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/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-co-mds9000-san-analytics-telemetry-streaming-config-guide-8x.html.

Procedure 3. 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**. Optionally, under Receiver, select the second IP address in the In-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_aa13-9132t-a

DCNM server time: 10:06:10.494 EDT Tuesday August 11 2020

Selected 1 / Total 1

Disable Analytics...									
Show Quick Filter									
<input type="checkbox"/>	Switch	Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Receiver	
<input checked="" type="checkbox"/>	aa13-9132t-a	DS-C9132T-K9	8.4(1a)	Yes	10:06:12.790 EDT Tue Aug 11 2020	SCSI	Storage	10.1.156.210	

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 Pure Storage FlashArray//X50 R3. Click **Continue**.

4. Select Interfaces

Choose the switch interfaces that will generate analytics data within Fabric_BB08-MDS-9132T-A

Total Top Level Rows 1

Switch	Module	Interface	Connected To	Type	Analytics Status	Enable / Disable SCSI Telemetry	Enable / Disable NVMe Telemetry
▼ BB08-MDS-9132...	1 module(s)	4 interface(s)		Storage			
	▼ DS-C9132T-K9-S...	4 interface(s)					
		fc1/1	FlashArray-CT0FC0	both	disabled	<input checked="" type="checkbox"/> <input type="checkbox"/> pending enable	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/2	52:4a:93:77:de:d7:21:01	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/3	FlashArray-CT1FC0	both	disabled	<input checked="" type="checkbox"/> <input type="checkbox"/> pending enable	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/4	52:4a:93:77:de:d7:21:11	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	<input type="checkbox"/> <input checked="" type="checkbox"/>

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.

After approximately two hours, you can view SAN Analytics data under the Dashboard and Monitor.

Cisco Intersight Assist

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. FlashStack environments include multiple devices that may not connect directly with Cisco Intersight. Any device that is supported by Cisco

Intersight but does not connect directly with it will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism and helps you add devices into Cisco Intersight.

Cisco Intersight Assist is available within the Cisco Intersight Virtual Appliance, which is distributed as a deployable virtual machine contained within an Open Virtual Appliance (OVA) file format. You can install the appliance on an ESXi server. For more information, see the [Cisco Intersight Virtual Appliance Getting Started Guide](#).

After claiming Intersight Assist into Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option.

The following sections describe some of the sample FlashStack Orchestration and lifecycle management tasks that can be performed using Cisco Intersight.

Procedure 1. Configure Cisco Intersight Assist

Step 1. To install Intersight Assist from an Open Virtual Appliance (OVA) in your VMware FlashStack-Management Cluster, first download the latest release of the OVA from:

<https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-499>

Step 2. Refer to https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html and set up the DNS entries for the Intersight Assist hostname as specified under Before you begin.

Step 3. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the **FlashStack-Management cluster** and click **Deploy OVF Template**.

Step 4. Specify a URL or browse to the **intersight-virtual-appliance-1.0.9-499.ova** file. Click **NEXT**.

The screenshot shows the 'Deploy OVF Template' wizard in VMware vCenter. The left sidebar lists the steps: 1. Select an OVF template (active), 2. Select a name and folder, 3. Select a compute resource, 4. Review details, 5. Select storage, and 6. Ready to complete. The main area is titled 'Select an OVF template' and contains the following text: '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.' There are two radio buttons: 'URL' (unselected) and 'Local file' (selected). Below the radio buttons is a text input field with the URL 'http | https://remoteserver-address/filetoinstall.ovf | .ova'. Under the 'Local file' section, there is an 'UPLOAD FILES' button and a file entry 'intersight-appliance-installer-vsphere-1.0.9-342.ova'. At the bottom right, there are 'CANCEL' and 'NEXT' buttons.

Step 5. Name the Intersight Assist VM and select the location. Click **NEXT**.

Step 6. Select the **FlashStack-Management cluster** and click **NEXT**.

Step 7. Review details and click **NEXT**.

Step 8. Select a deployment configuration and click **NEXT**.

Note: The Tiny (8 vCPU, 16 GiB RAM) deployment option is applicable only for Intersight Assist deployment without Workload Optimizer. Small deployment is the minimum requirement for Workload Optimizer.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Configuration**
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Configuration [Close]

Select a deployment configuration

<input checked="" type="radio"/> Tiny	Description 8 vCPU, 16GiB Memory, 500GB Storage. Note: Only viable for Intersight Assist used with Intersight Orchestrator.
<input type="radio"/> Small	
<input type="radio"/> Medium	
<input type="radio"/> Large	

4 Items

CANCEL BACK NEXT

Step 9. Select **Infra-DataStore1** for storage and select the Thin Provision virtual disk format. Click **NEXT**.

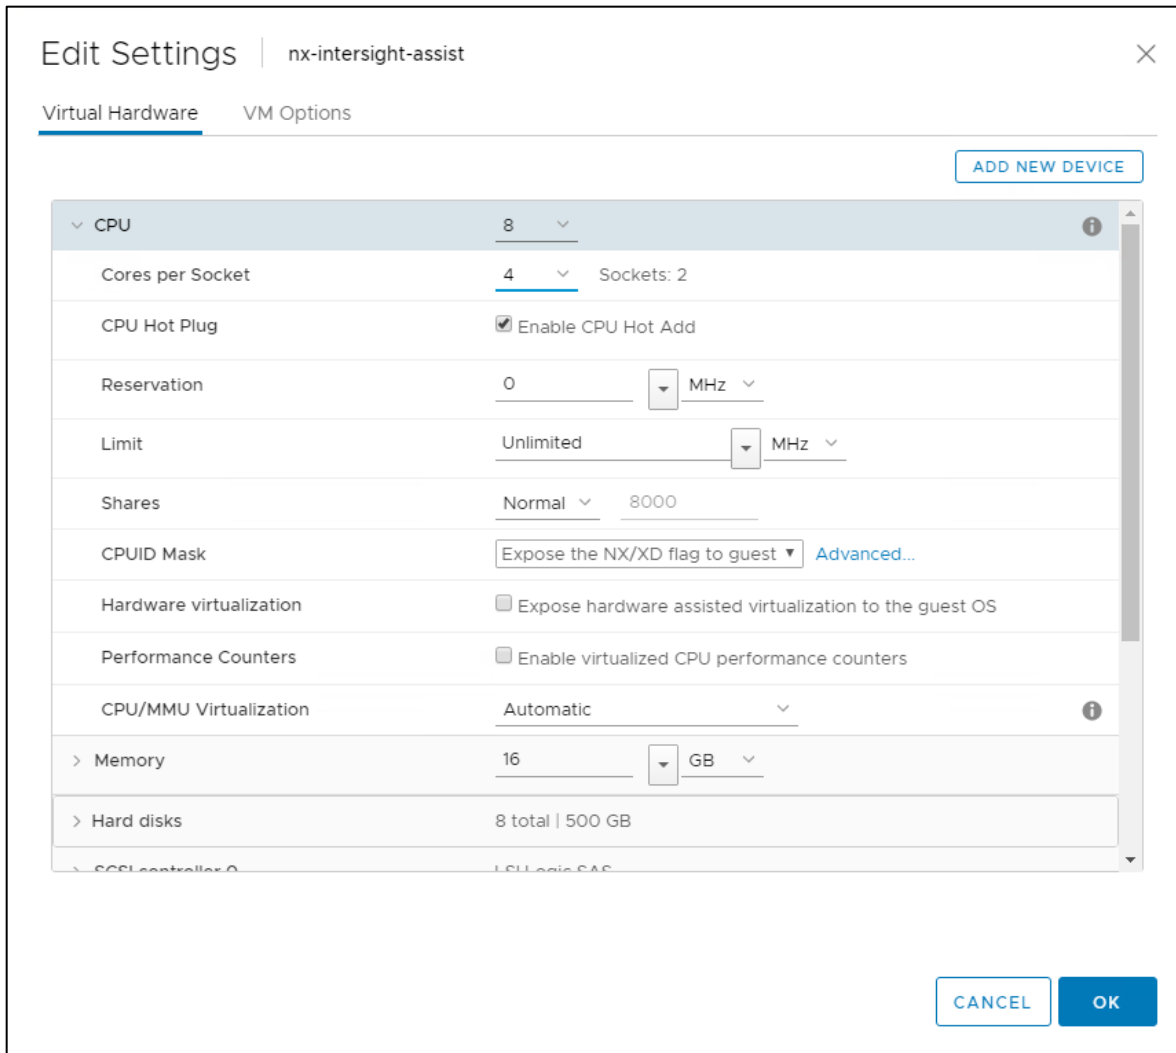
Step 10. Select **IB-MGMT Network** for the VM Network. Click **NEXT**.

Step 11. Fill in all values to customize the template. Click **NEXT**.

Step 12. Review the deployment information and click **FINISH** to deploy the appliance.

Step 13. Once the OVA deployment is complete, right-click the **Intersight Assist VM** and click **Edit Settings**.

Step 14. Expand CPU and adjust the Cores per Socket so that 2 Sockets are shown. Click **OK**.



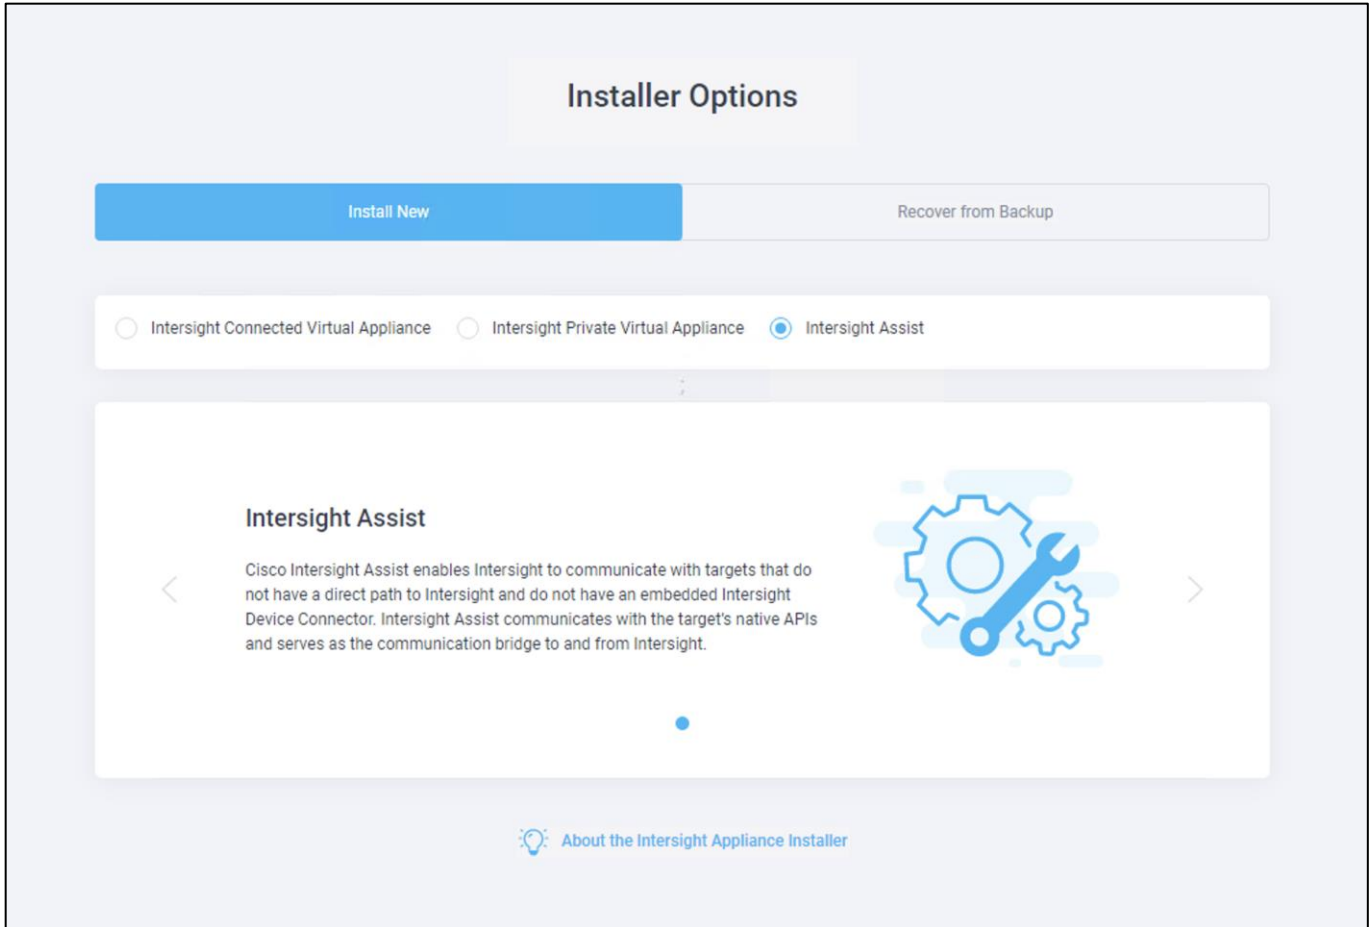
Step 15. Right-click the **Intersight Assist VM** and select **Open Remote Console**.

Step 16. Click the **play icon** to power on the VM.

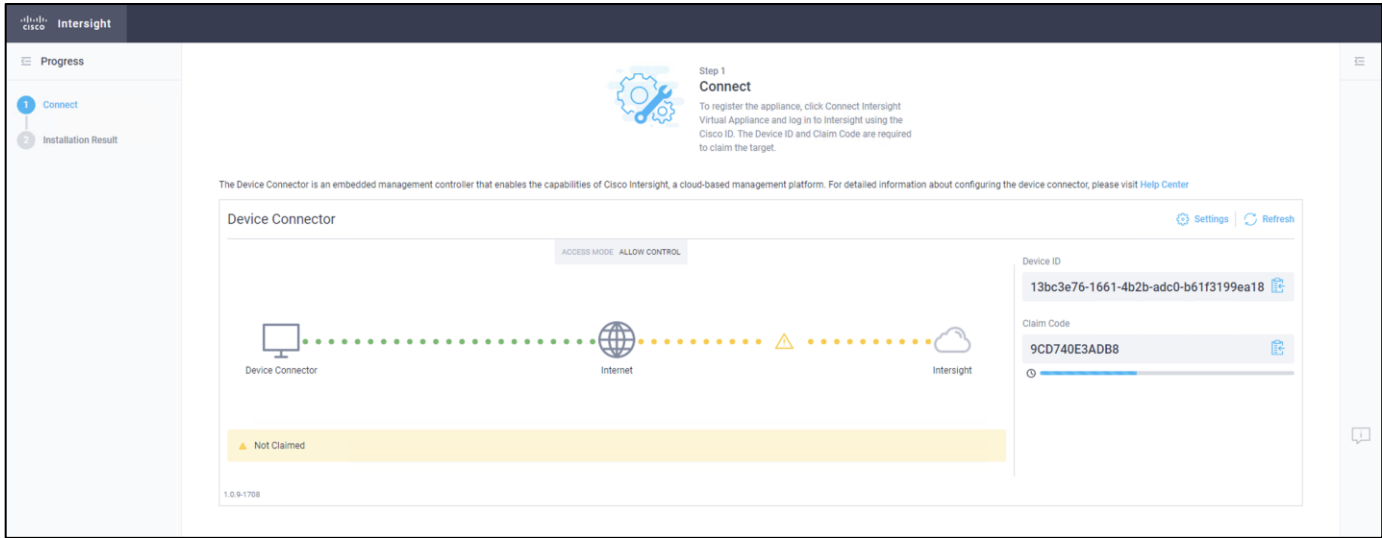
Step 17. When you see the login prompt, close the Remote Console, and connect to <https://intersight-assist-fqdn>.

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

Step 18. Navigate the security prompts and select **Intersight Assist**. Click **Proceed**.



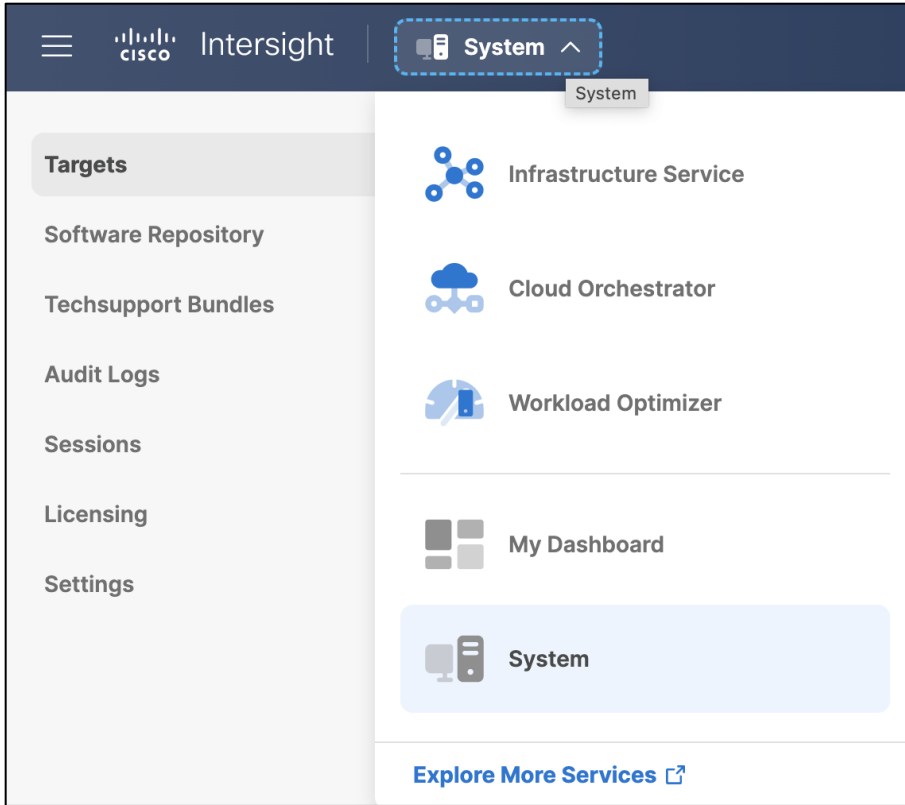
Step 19. Copy the **Device ID** and **Claim Code** shown in the Intersight Assist web interface to claim it in Cisco Intersight.



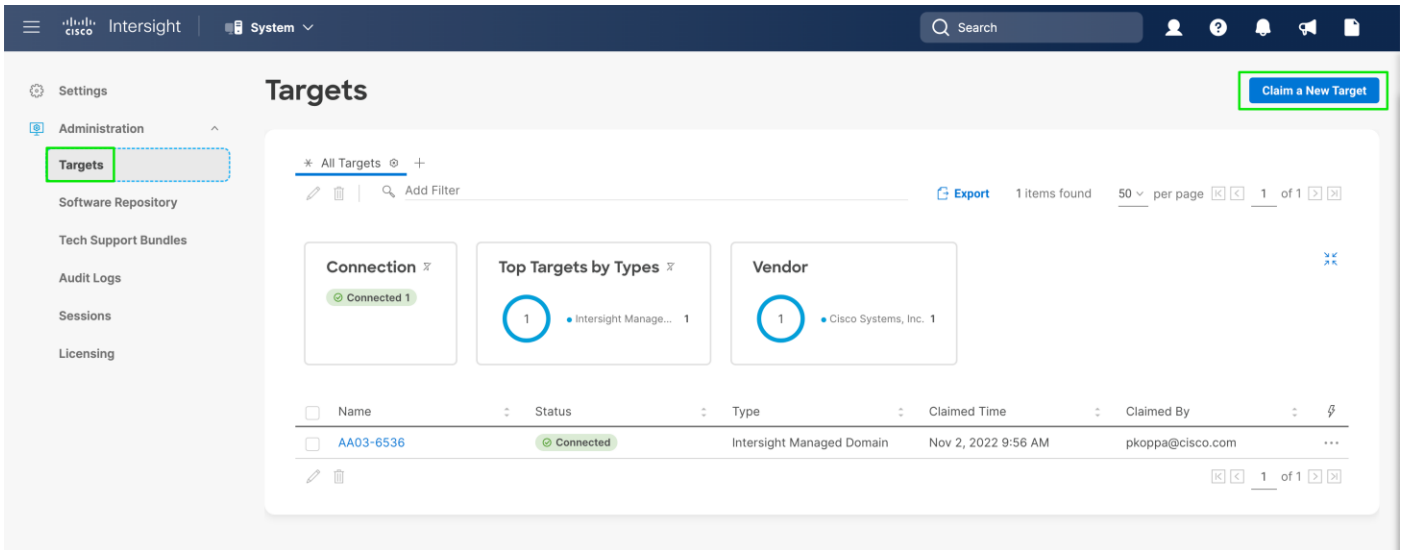
Step 20. Open a browser to Cisco Intersight, <https://intersight.com>.

Step 21. Log into your **Cisco Intersight** account.

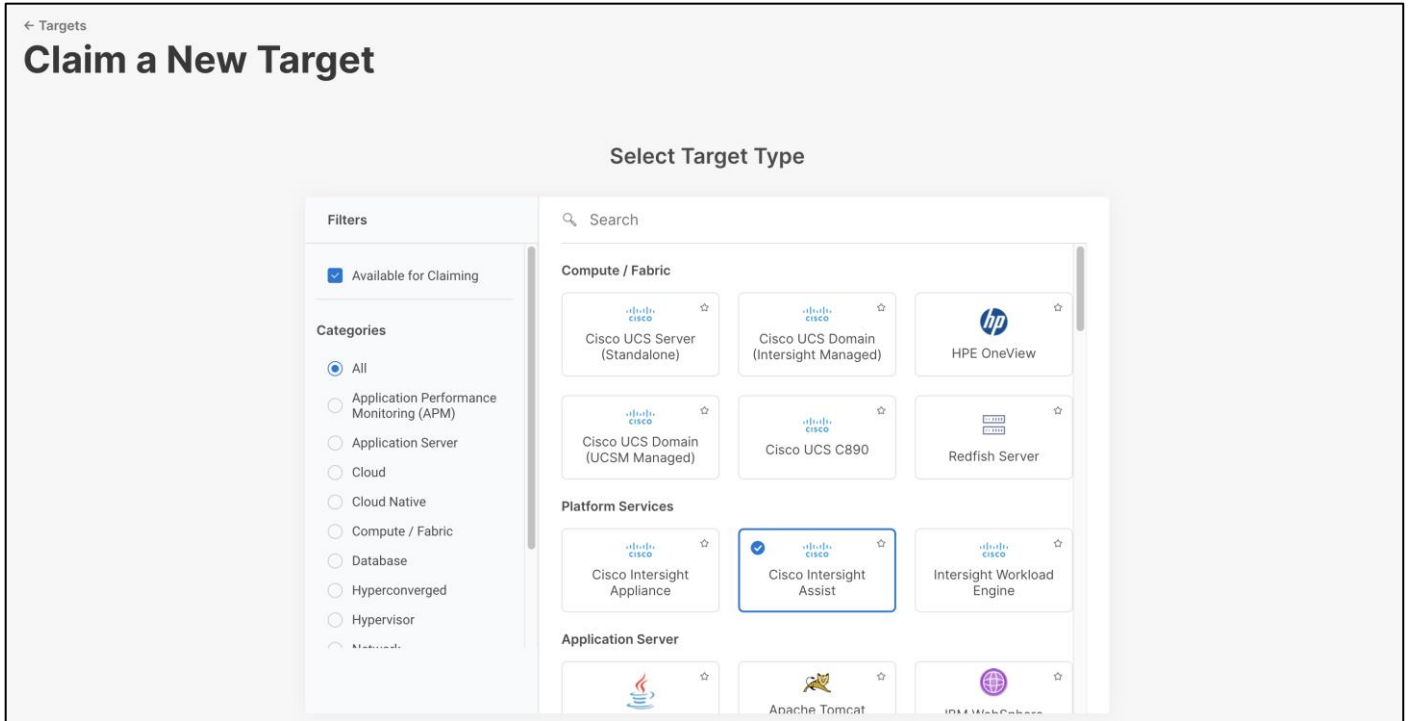
Step 22. From Service Selector, select **System**



Step 23. From the left navigation pane, select **Targets** and click **Claim a New Target**.



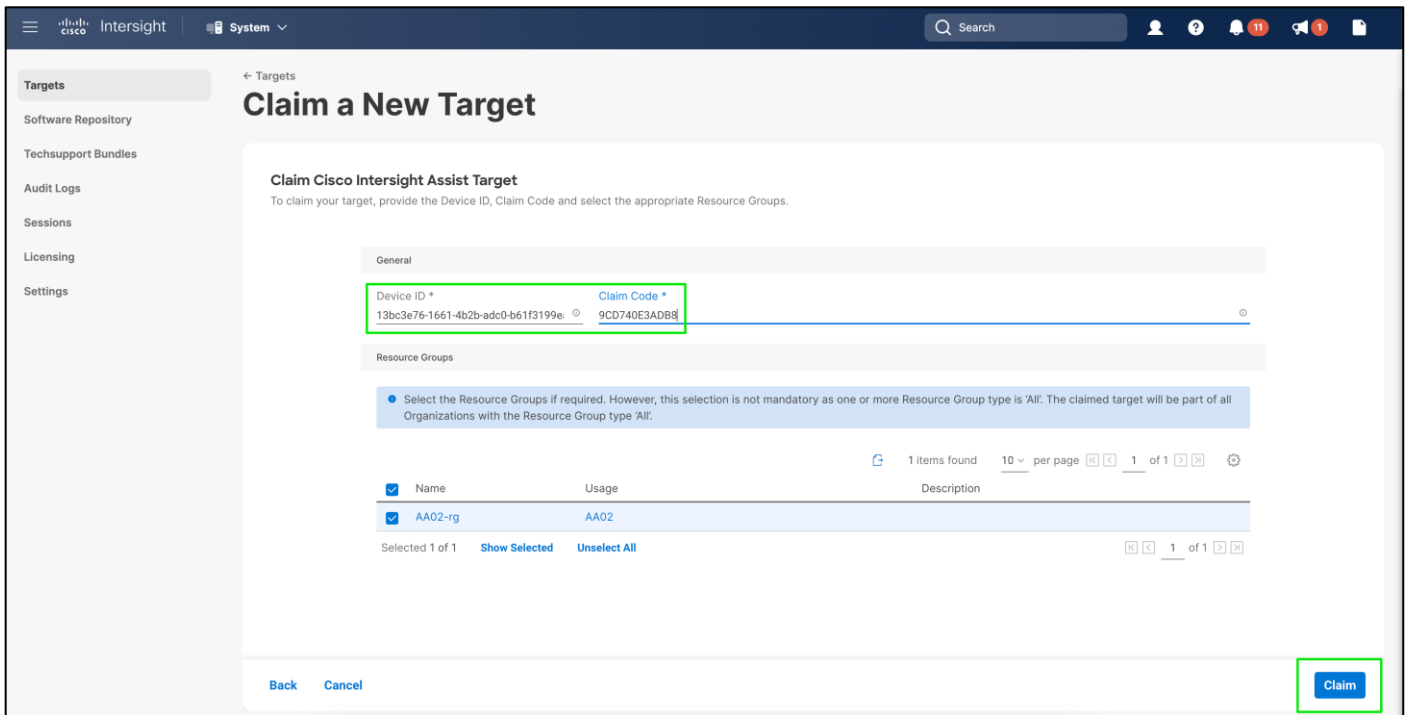
Step 24. Select **Cisco Intersight Assist** and click **Start**.




Step 25. Paste the copied Device ID and Claim Code from Intersight Assist web interface.

Step 26. Select the appropriate resource group (for example: FlashStack-rg).

Step 27. Click **Claim**.



Step 28. In the Intersight Assist web interface, click **Continue**.



The 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 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 29. When the software download is complete, navigate the security prompts and an Intersight Assist login screen will appear. 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.

VMware vCenter

Claiming a Pure VMWare vCenter requires the use of an Intersight Assist virtual machine. Deploy an Intersight assist appliance using the above-described procedure if one doesn't exist.


Procedure 2. Claim VMware vCenter Server

Step 1. Open a browser to Cisco Intersight, <https://intersight.com>, and log into your Intersight account.

Step 2. From Service Selector, select **System**.

Step 3. From the left navigation pane, select **Targets** and click **Claim a New Target**.

Step 4. Filter to Hypervisor and select VMware vCenter.



Step 5. Click **Start**.

Step 6. Fill in the vCenter information. If multiple Intersight Assist instances are deployed, make sure the Intersight Assist which has connectivity to the target is correctly selected.

Step 7. Make sure Secure option is enabled to indicate connection to the target should be established using HTTPS.

Step 8. Datastore Browsing controls whether Workload Optimizer scans vCenter datastores to identify files which are not used and can be deleted to reclaim space and improve actual disk utilization. For example, orphaned VMDK files.

Step 9. Enable retrieval of advanced memory metrics by Workload Optimizer Service. Only supported on vCenter Server version 6.5U3 or later. Guest VMs must run VMWare Tools 10.3.2 Build 10338 or later.

Step 10. Enabling Hardware Support Manager (HSM) allows vCenter to perform firmware operations on UCS servers claimed in the vCenter cluster. HSM is supported only from vCenter version 7.0 and above.

← 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 *
 aa03-assist-prod.flashstack.com

Hostname/IP Address *
 aa03-vcenter.flashstack.com

Port
 443

Username *
 administrator@vsphere.local

Password *

Secure

Enable Datastore Browsing

Enable Guest Metrics

Enable HSM

▲ Enabling HSM will give escalated privileges to the vCenter target to perform firmware operations on UCS servers claimed in Cisco Intersight.

Back Cancel Claim

Step 11. Click **Claim**.

Step 12. After a few minutes, the VMware vCenter will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

Intersight System

Targets

Claim a New Target

* All Targets +

Export 3 items found 50 per page 1 of 1

Connection 3 Connected

Top Targets by Types

- Intersight Assist 1
- VMware vCenter 1
- Intersight Manage... 1

Vendor

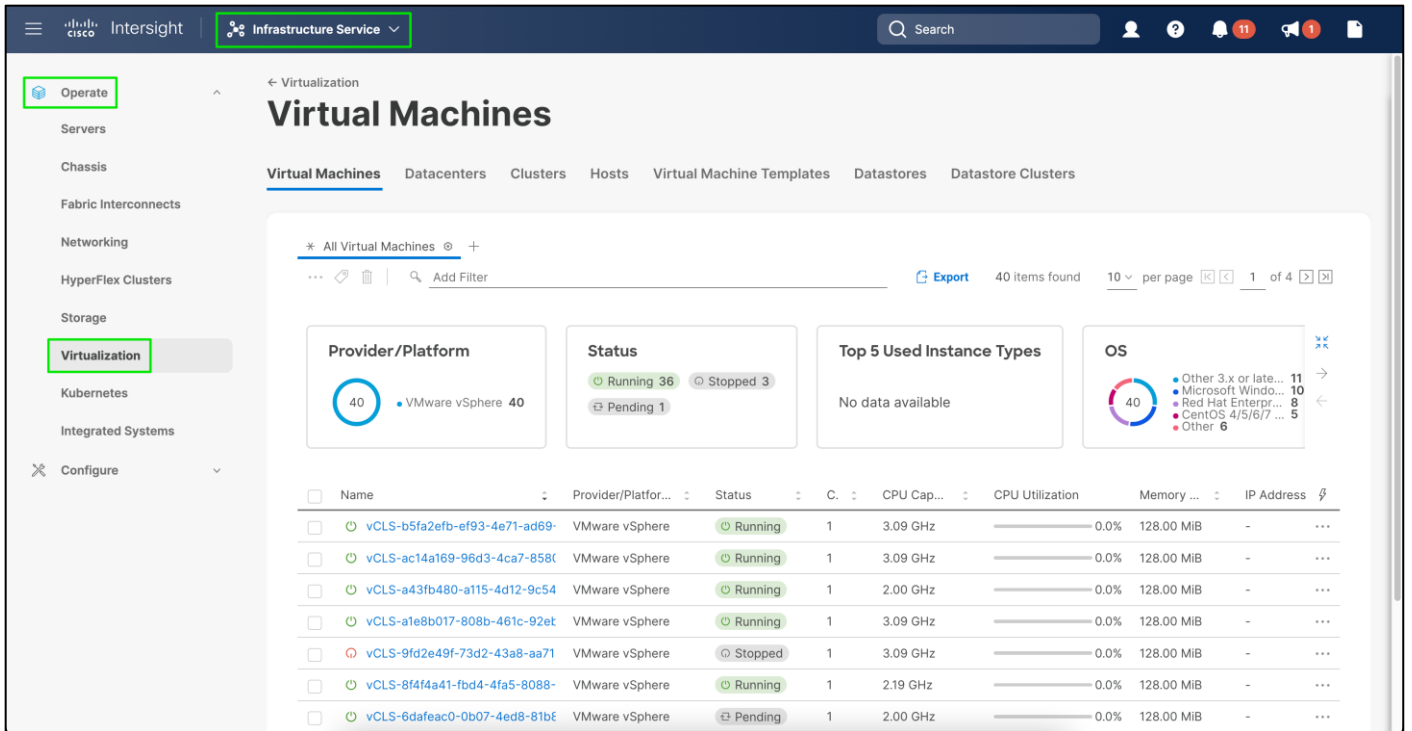
- Cisco Systems, Inc. 2
- VMware 1

Name	Status	Type	Claimed By
aa03-vcenter.flashstack.com	Connected	VMware vCenter	pkoppa@cisco.com
aa03-assist-prod.flashstack.com	Connected	Intersight Assist	pkoppa@cisco.com
AA03-6536	Connected	Intersight Managed Domain	pkoppa@cisco.com

Step 13. Repeat steps 1 - 12 for all VMware vCenter targets.

Step 14. Detailed information obtained from the vCenter can now be viewed by clicking **Virtualization** from the menu. From Service Selector, select **Infrastructure Service**.

Step 15. From the left navigation pane, select **Operate** and click **Virtualization**.




Claim Pure Storage FlashArray//X and Pure Storage FlashArray//X in Cisco Intersight

Claiming a Pure Storage FlashArray also requires the use of an Intersight Assist virtual machine.

Procedure 3. Claim a Pure Storage FlashArray using Cisco Intersight

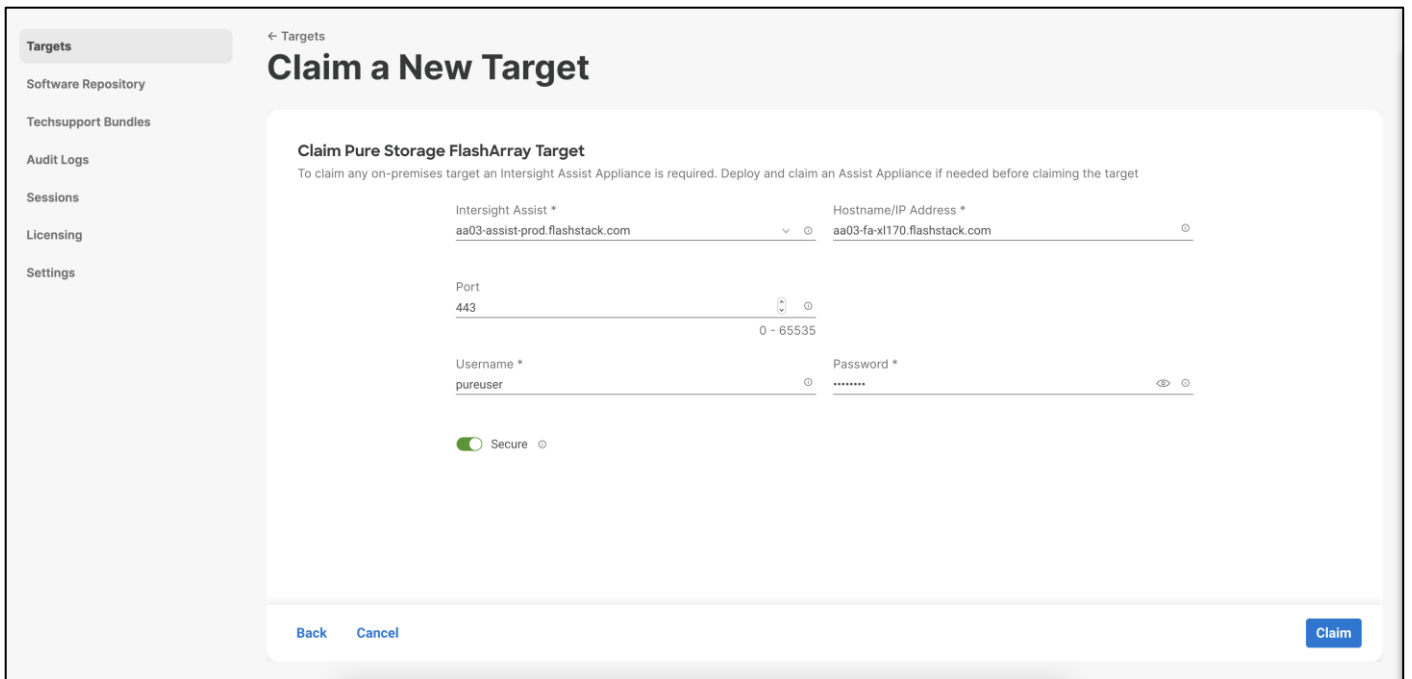
- Step 1.** Open a browser to Cisco Intersight, <https://intersight.com>, and log into your Intersight account.
- Step 2.** From Service Selector, select **System**.
- Step 3.** From the left navigation pane, select **Targets** and click **Claim a New Target**.
- Step 4.** In search bar, type Pure and select **Pure Storage FlashArray**.



Step 5. Click **Start**.

Step 6. Fill in the FlashArray information. If multiple Intersight Assist instances are deployed, make sure the Intersight Assist which has connectivity to the FlashArray management is correctly selected.

Step 7. Make sure Secure option is enabled to indicate connection to the target should be established using HTTPS.



Step 8. After a few minutes, the FlashArray target will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

Step 9. Repeat step 1 – 8 for all FlashArray storages present.

The screenshot shows the Cisco Intersight interface. The top navigation bar includes the Cisco logo, 'Intersight', and 'System'. A search bar and user profile icons are on the right. The left sidebar contains 'Settings', 'Administration', 'Targets' (selected), 'Software Repository', 'Tech Support Bundles', 'Audit Logs', 'Sessions', and 'Licensing'. The main 'Targets' section has a 'Claim a New Target' button. It displays a summary of 5 connected targets, a 'Top Targets by Types' chart showing 2 Pure Storage FlashArrays, 1 Intersight Assist, and 1 Intersight Managed Domain, and a 'Vendor' chart showing 2 Pure Storage, 2 Cisco Systems, Inc., and 1 VMware. Below is a table of targets:

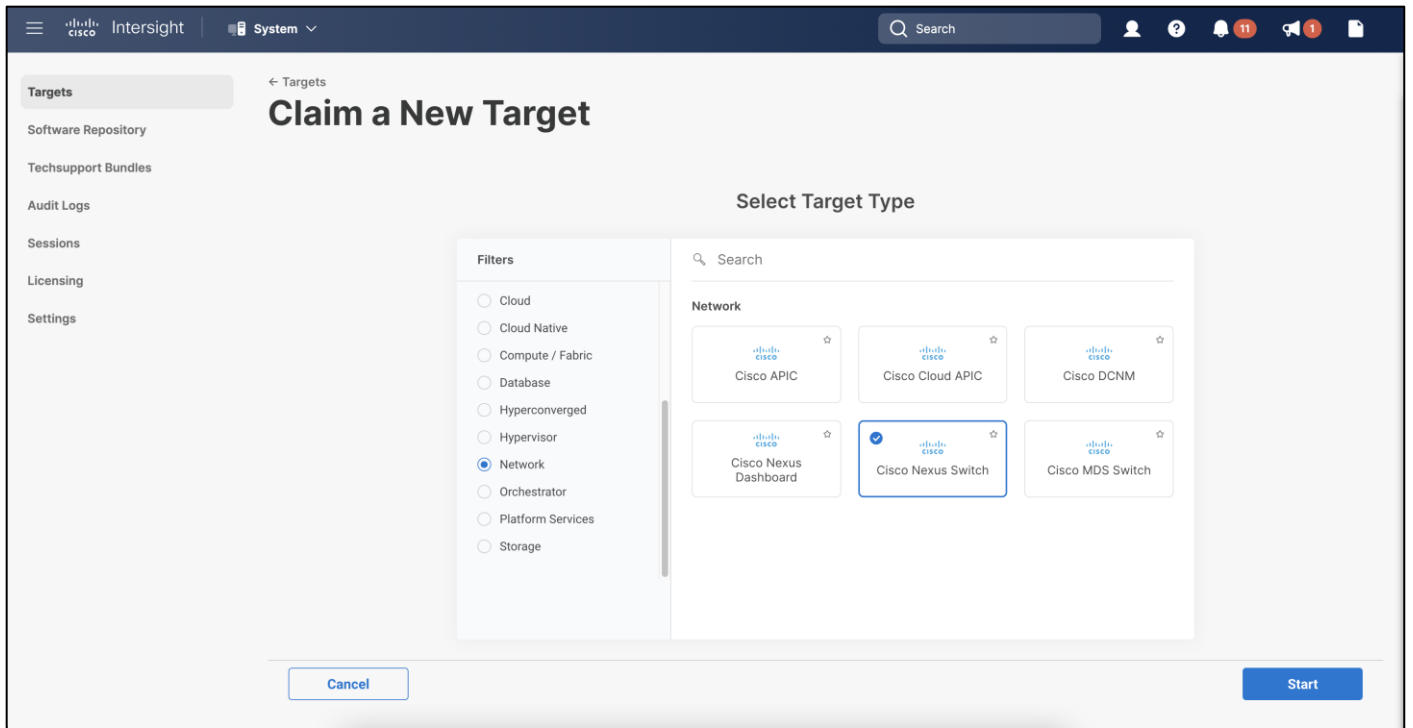
Name	Status	Type	Claimed By
aa03-fa-x170.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-fa-x50.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-vcenter.flashstack.com	Connected	VMware vCenter	pkoppa@cisco.com
aa03-assist-prod.flashstack.com	Connected	Intersight Assist	pkoppa@cisco.com
AA03-6536	Connected	Intersight Managed Domain	pkoppa@cisco.com

Claim Nexus and MDS switches in Cisco Intersight

Claiming a Cisco Nexus 9000 and Cisco MDS switches also requires the use of an Intersight Assist virtual machine.

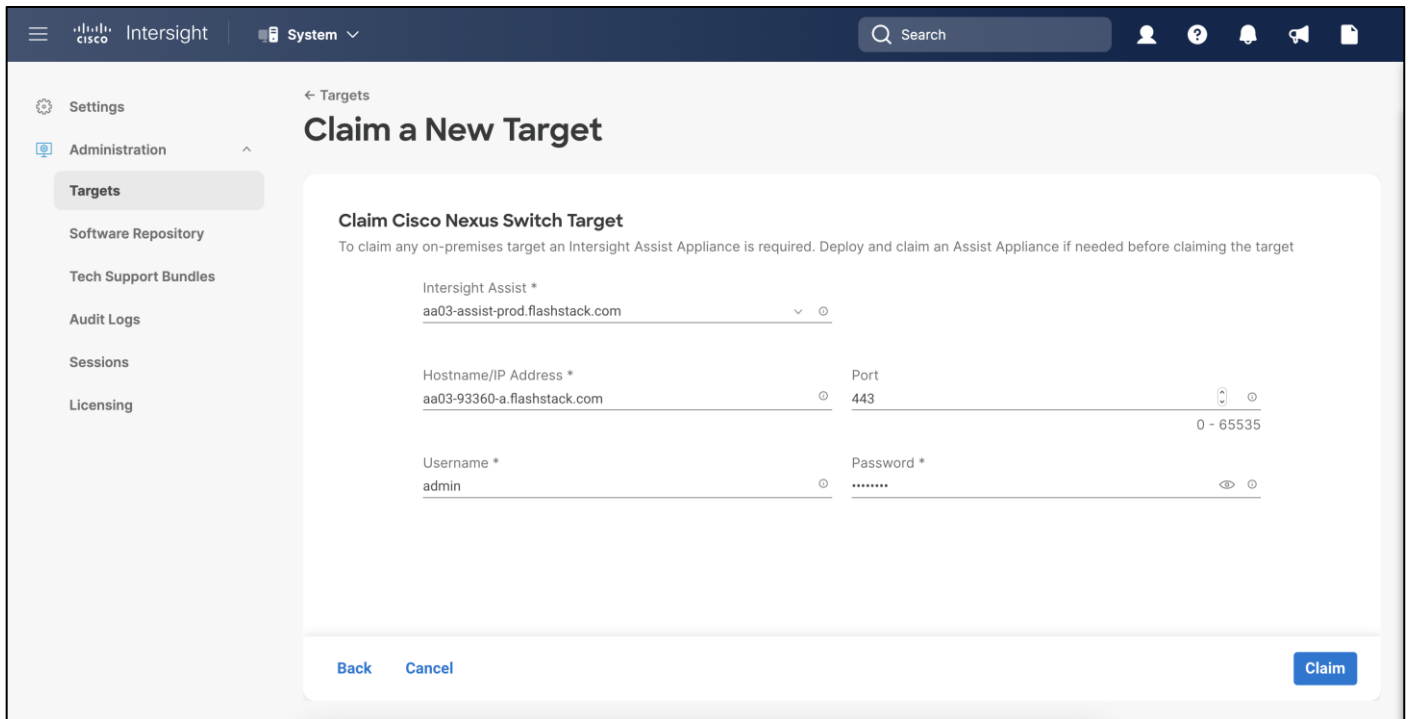
Procedure 4. Claim Cisco Nexus and MDS switches

- Step 1.** Open a browser to Cisco Intersight, <https://intersight.com>.
- Step 2.** Log into your Intersight account.
- Step 3.** From Service Selector, select **System**.
- Step 4.** From the left navigation pane, select **Network**.
- Step 5.** Click **Cisco Nexus Switch**.



Step 6. Click **Start**.

Step 7. Fill in the Cisco Nexus switch information. If multiple Intersight Assist instances are deployed, make sure the Intersight Assist which has connectivity to the Nexus switch management is correctly selected.



Step 8. After a few minutes, the Cisco Nexus 93360 switch will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

Step 9. Repeat step 1 – 8 for all Cisco Nexus 93360 switches present.

The screenshot shows the 'Targets' page in the Cisco Intersight interface. The left sidebar contains navigation options: Settings, Administration (with 'Targets' selected), Software Repository, Tech Support Bundles, Audit Logs, Sessions, and Licensing. The main content area is titled 'Targets' and includes a 'Claim a New Target' button. Below the title, there are summary cards for 'Connection' (7 Connected), 'Top Targets by Types' (Cisco Nexus Switch: 2, Pure Storage FlashArray: 2, Intersight Managed Domain: 1, VMware vCenter: 1, Other: 1), and 'Vendor' (Cisco Systems, Inc.: 4, Pure Storage: 2, VMware: 1). A table below lists the targets with columns for Name, Status, Type, and Claimed By. Two rows are highlighted with a green border.

Name	Status	Type	Claimed By
aa03-93360-b.flashstack.com	Connected	Cisco Nexus Switch	pkoppa@cisco.com
aa03-93360-a.flashstack.com	Connected	Cisco Nexus Switch	pkoppa@cisco.com
aa03-fa-xl170.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-fa-x50.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-vcenter.flashstack.com	Connected	VMware vCenter	pkoppa@cisco.com
aa03-assist-prod.flashstack.com	Connected	Intersight Assist	pkoppa@cisco.com
AA03-6536	Connected	Intersight Managed Domain	pkoppa@cisco.com

Step 10. Click **Claim a New Target**.

Step 11. In the Claim a New Target Wizard, click **Cisco MDS Switch**.

The screenshot shows the 'Claim a New Target' wizard in the Cisco Intersight interface. The left sidebar is the same as in the previous screenshot. The main content area is titled 'Claim a New Target' and shows the 'Select Target Type' step. On the left, there are filters: Database, Hyperconverged, Hypervisor, Network (selected), Orchestrator, Platform Services, and Storage. On the right, there is a search bar and a grid of target types under the 'Network' category: Cisco APIC, Cisco Cloud APIC, Cisco DCNM, Cisco Nexus Dashboard, Cisco Nexus Switch, and Cisco MDS Switch (which is selected with a blue checkmark). At the bottom, there are 'Cancel' and 'Start' buttons.

Step 12. Click **Start**.

Step 13. Fill in the MDS switch information. If multiple Intersight Assist instances are deployed, make sure the Intersight Assist which has connectivity to the Cisco Nexus switch management is correctly selected.

The screenshot shows the Intersight web interface. The top navigation bar includes the Cisco logo, 'Intersight', and a 'System' dropdown menu. A search bar and several utility icons are on the right. The left sidebar contains a menu with 'Settings', 'Administration', 'Targets', 'Software Repository', 'Tech Support Bundles', 'Audit Logs', 'Sessions', and 'Licensing'. The main content area is titled 'Targets' and 'Claim a New Target'. Below this is a sub-header 'Claim Cisco MDS Switch Target' and a note: 'To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target'. The form contains the following fields:

- Intersight Assist ***: A dropdown menu with the selected value 'aa03-assist-prod.flashstack.com'.
- Hostname/IP Address ***: A text input field containing 'aa03-9132t-a.flashstack.com'.
- Port**: A text input field containing '8443', with a range indicator '0 - 65535' below it.
- Username ***: A text input field containing 'snmpadmin'.
- Password ***: A password input field with masked characters '.....' and a visibility toggle icon.

At the bottom of the form, there are 'Back' and 'Cancel' links on the left, and a blue 'Claim' button on the right.

Step 14. After a few minutes, the MDS switch will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

Step 15. Repeat steps 1 - 14 for the other MDS switches present.

The screenshot displays the Cisco Intersight 'Targets' management interface. On the left, a navigation pane includes 'Settings', 'Administration', 'Targets' (selected), 'Software Repository', 'Tech Support Bundles', 'Audit Logs', 'Sessions', and 'Licensing'. The main area shows a list of 9 targets, all with a 'Connected' status. Two targets are highlighted with a green border:

Name	Status	Type	Claimed By
aa03-9132t-b.flashstack.com	Connected	Cisco MDS Switch	pkoppa@cisco.com
aa03-9132t-a.flashstack.com	Connected	Cisco MDS Switch	pkoppa@cisco.com
aa03-93360-b.flashstack.com	Connected	Cisco Nexus Switch	pkoppa@cisco.com
aa03-93360-a.flashstack.com	Connected	Cisco Nexus Switch	pkoppa@cisco.com
aa03-fa-x170.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-fa-x50.flashstack.com	Connected	Pure Storage FlashArray	pkoppa@cisco.com
aa03-vcenter.flashstack.com	Connected	VMware vCenter	pkoppa@cisco.com
aa03-assist-prod.flashstack.com	Connected	Intersight Assist	pkoppa@cisco.com
AA03-6536	Connected	Intersight Managed Domain	pkoppa@cisco.com

Cisco Intersight Cloud Orchestration

Cisco Intersight contains powerful orchestration tools and workflows which enable administrators to automate many common and complex sequences of tasks. This additional feature can be used to speed the deployment and configuration of many environments. As an example, a new FC based host in a FlashStack environment can be added using these orchestration tools, as shown below.

Procedure 1. FC Host Registration using Cisco Intersight

Step 1. Open a browser to Cisco Intersight, <https://intersight.com>, and log into your Intersight account.

Step 2. From Service Selector, select **Cloud Orchestrator**.

Step 3. From the left navigation pane, select **Workflows > All Workflows**.

Validation Status
Valid 30

Last Execution Status
Success 3

Top 5 Workflows by Execution Count
31
• Create iSCSI Data... 26
• Create Terraform... 4
• New Storage H... 1

Top 5 Workflow Categories
24
• Storage 9
• Virtualization 7
• Terraform Cloud 6
• Compute 1
• IWE 1

System
Yes 28
No 2

Display Name	Description
Update VMFS Datastore	Expand a datastore on hypervisor manager by extending the backing storage volume to specified capacity, and then grow the datastore to utilize the a...
Update Storage Host	Update the storage host details. If the inputs for a task are provided then the task is run, else it is skipped.
Update NAS Datastore	Update NAS datastore by expanding capacity of the underlying NFS volume. The expanded capacity is visible to all hosts connected to the datastore.
Remove VMFS Datastore	Remove VMFS datastore and remove the backing volume from the storage device. When a datastore is removed from a host, it is destroyed and will di...
Remove Storage Host Group	Remove storage host group. If hosts are provided as input, the workflow will remove the hosts from the host group.
Remove Storage Host	Remove storage host. If host group name is provided as input, the workflow will also remove the host from the host group.
Remove Storage Export Policy	Remove the NFS volume and the export policy attached to the volume.
Remove NAS Datastore	Remove the NAS datastore and the underlying NFS storage volume. When a datastore is removed from a host, it is destroyed and will disappear from a...
New VMFS Datastore	Create a storage volume and build VMFS datastore on the volume.
New Virtual Machine	Create a new virtual machine on the hypervisor from an OVA or OVF file. Datastore, Host/Cluster, and Image URL fields are mandatory. All other inputs...
New Storage Virtual Machine	Create a storage virtual machine.
New Storage Interface	Create a storage IP or FC interface.
New Storage Host Group	Create a new storage host group. If hosts are provided as inputs, the workflow will add the hosts to the host group.
New Storage Host	Create a new storage host. If host group is provided as input, then the host will be added to the host group.
New NAS Datastore	Create a NFS storage volume and build NAS datastore on the volume.

Step 4. Click **New Storage Host** and click **Execute**.

New Storage Host Valid

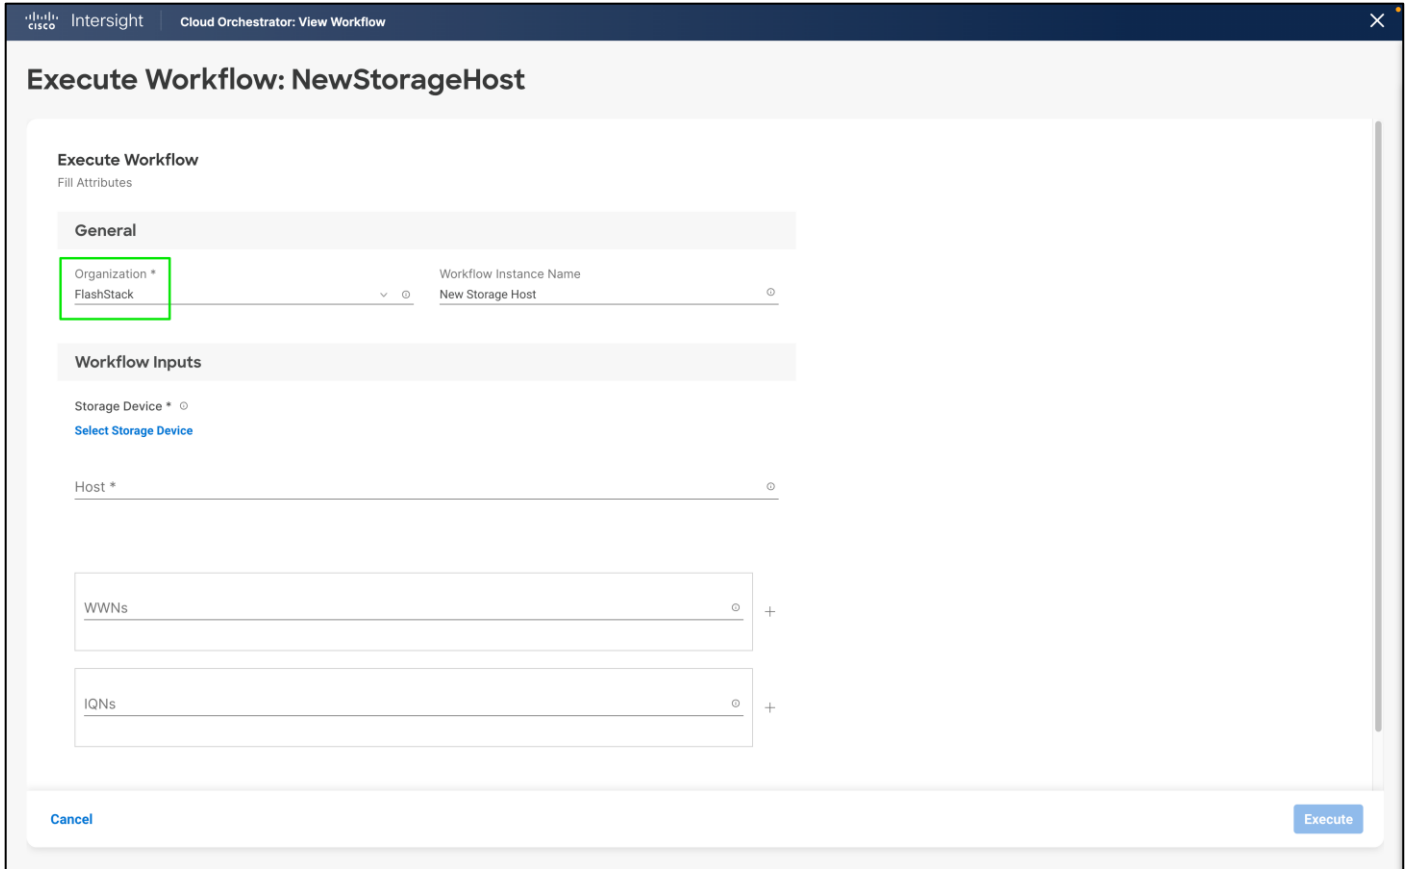
Last saved 18 hours ago **Actions** **Execute**

General **Designer** Mapping Code History

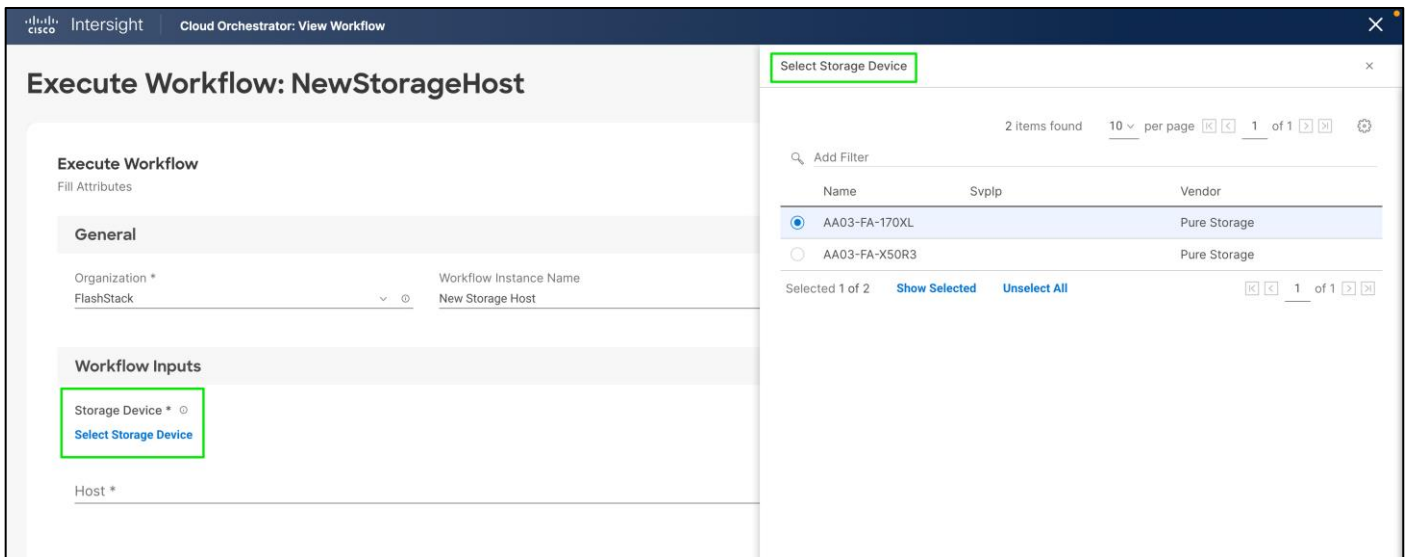
```

graph TD
    Start([Start]) --> NewStorageHost[New Storage Host  
Storage]
    NewStorageHost --> AddHost[Add Host to Storage Host Group  
Storage]
    AddHost --> Success([Success])
    AddHost --> Failed([Failed])
  
```

Step 5. Select the appropriate Organization (**default** by default).



Step 6. Click **Select Storage Device** and select the appropriate Pure Storage device.



Step 7. Enter the name of the Host name and WWNs for host VM-Host-Infra-FCP-01.

Intersight Cloud Orchestrator: View Workflow

Execute Workflow: NewStorageHost

Execute Workflow
Fill Attributes

General

Organization *
FlashStack

Workflow Instance Name
New Storage Host

Workflow Inputs

Storage Device *
Selected Storage Device AA03-FA-170XL

Host Group
[Select Host Group](#)

Host *
VM-Host-Infra-FCP-01

WWNs
20:00:00:b4:aa:03:0a:00

WWNs
20:00:00:b4:aa:03:0b:00

Cancel Execute

Step 8. Click **Select Host Group** and select the Host group (Optional).

Intersight Cloud Orchestrator: View Workflow

Execute Workflow: NewStorageHost

FlashStack New Storage Host

Workflow Inputs

Storage Device *
Selected Storage Device AA03-FA-170XL

Host Group
[Select Host Group](#)

Host *
VM-Host-Infra-FCP-01

WWNs
20:00:00:b4:aa:03:0a:00

WWNs
20:00:00:b4:aa:03:0b:00

IQNs

Cancel

Select Host Group

12 items found 10 per page 1 of 2

Add Filter

Name

- AMD-FCP
- AMD-NVMe
- AMD-iSCSI
- FC-NVMe
- FCP
- FS-SQL-ESXiCius-FCNVMe
- FS-SQL-ESXiCluster
- VM-Host-Infra-FCP-Host-Group
- VM-Infra-FCP-Host-Group
- VM-Infra-NVMe-Host-Group

Selected 1 of 12 [Show Selected](#) [Unselect All](#) 1 of 2

Cancel Select

Step 9. Click **Execute**.

Step 10. Confirm that the execution is successful.

New Storage Host Valid

General Designer Mapping Code History

[Rollback](#) [Clone Execution](#)

```

graph TD
    Start([Start]) --> NewStorageHost[New Storage Host]
    NewStorageHost --> AddHost[Add Host to Storage Host Group]
    AddHost --> Success([Success])
    AddHost --> Failed([Failed])
    
```

Execution
New Storage Host - Oct 11, 2022 3:35 PM

Organization
AA03

Status
Success

- Logs
- Inputs
- Outputs
- 2 Add Host to Storage Host Group Oct 11, 2022 03:36:07 PM
 - Logs
 - Inputs
 - Outputs
- Success Oct 11, 2022 03:36:11 PM

Step 11. Repeat steps 1 - 10 for all hosts.

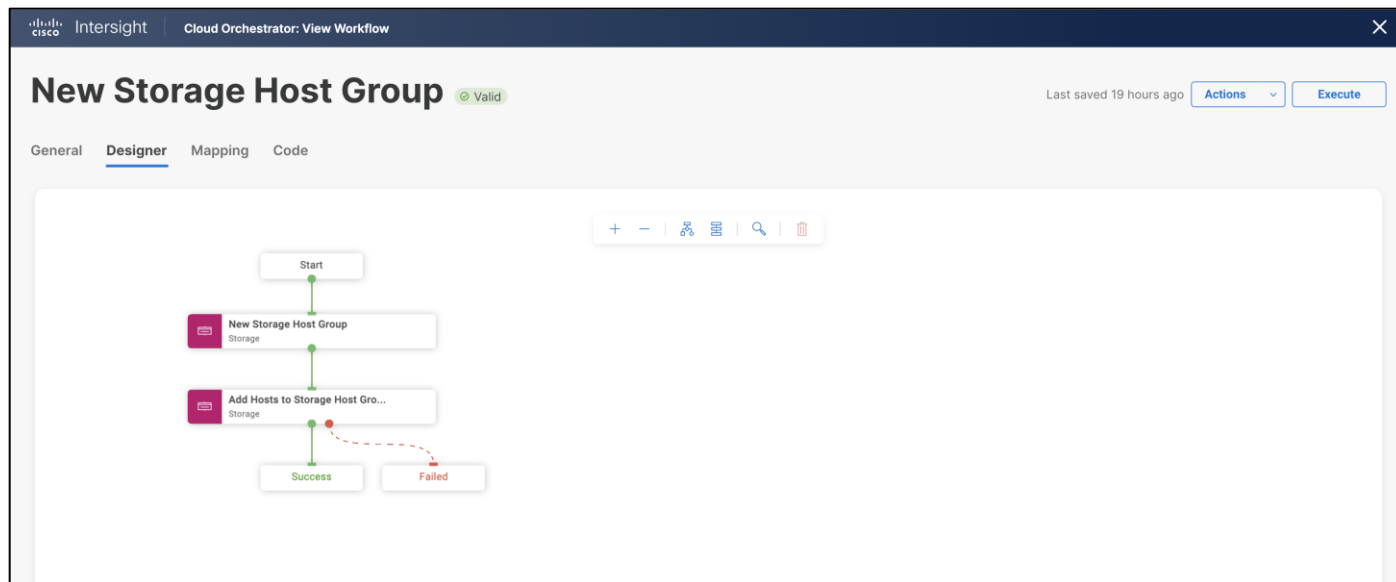
Procedure 2. Create FC Host Group using Cisco Intersight

Step 1. Open a browser to Cisco Intersight, <https://intersight.com>, and log into your Intersight account.

Step 2. From Service Selector, select **Cloud Orchestrator**.

Step 3. From the left navigation pane, Select **Workflows > All Workflows**.

Step 4. Click **New Storage Host Group** and click **Execute**.



Step 5. Select the appropriate Organization (default by default).

Step 6. Click **Select Storage Device** and select the appropriate Pure Storage device.

Step 7. Enter the name of the Host Group and enter all the host names for the new host group (Optional).

InterSight Cloud Orchestrator: View Workflow

Execute Workflow: NewStorageHostGroup

Execute Workflow
Fill Attributes

General

Organization *
FlashStack

Workflow Instance Name
New Storage Host Group

Workflow Inputs

Storage Device *
Selected Storage Device AA03-FA-170XL

Host Group *
VM-Host-Infra-FCP-Host-Group

Hosts

Cancel Execute

Step 8. Click **Execute**.

Step 9. Confirm that the execution is successful.

InterSight Cloud Orchestrator: View Workflow

New Storage Host Group Valid

Last saved 19 hours ago Execute

General Designer Mapping Code History

Rollback Clone Execution

```

graph TD
    Start([Start]) --> Step1[New Storage Host Group]
    Step1 --> Step2[Add Hosts to Storage Host Group]
    Step2 --> Success([Success])
    Step2 --> Failed([Failed])
  
```

Execution
New Storage Host Gro... - Oct 11, 2022 3:49 PM

Organization: AA03 Status: Success

Step	Time
Start	Oct 11, 2022 03:49:19 PM
1 New Storage Host Group	Oct 11, 2022 03:49:21 PM
2 Add Hosts to Storage Host Group	Oct 11, 2022 03:49:22 PM
Success	Oct 11, 2022 03:49:22 PM

Pure Storage vSphere Client Plugin

The Pure Storage Plugin for the vSphere Client provides the ability to VMware users to have insight into and control of their Pure Storage FlashArray environment while directly logged into the vSphere Client. The Pure Storage plugin extends the vSphere Client interface to include environmental statistics and objects that underpin the VMware objects in use and to provision new resources as needed.

The Pure Storage vSphere Client Plugin will be accessible through the vSphere Client after registration through the Pure Storage Web Portal.

Procedure 1. Install the Pure Storage vSphere Client Plugin

Step 1. The Pure Storage VMware Appliance is an OVA that can be deployed from Pure Storage using the link provided in the user guide here: https://support.purestorage.com/Solutions/VMware_Platform_Guide/User_Guides_for_VMware_Solutions/Using_the_Pure_Storage_Plugin_for_the_vSphere_Client/vSphere_Plugin_User_Guide%3A_Installing_the_Remote_vSphere_Plugin_with_the_Pure_Storage_VMware_Appliance. It is required to download the OVA for both the online and offline deployment types.

Step 2. Follow the steps in the user guide opened in Step 1.

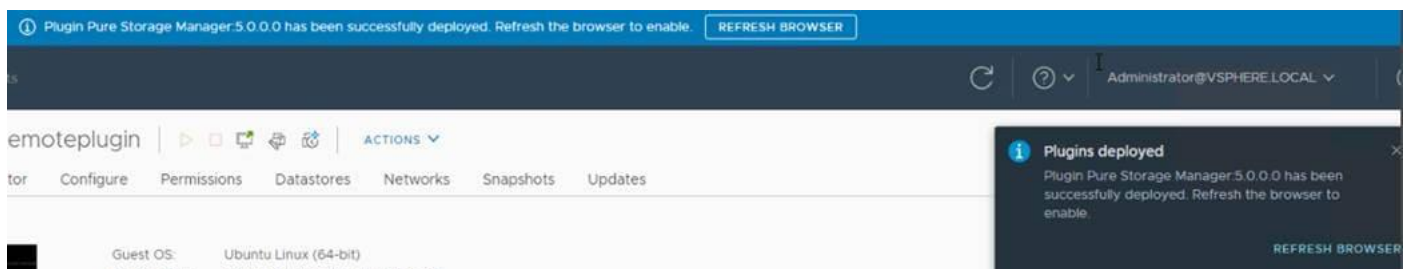
Step 3. Once the Pure Storage VMware appliance is deployed, connect it to your vCenter:

- Open an SSH connection to the appliance using the OVA VM's DNS name or IP address displayed in vCenter.
- In the SSH shell login with the pureuser account. On the first login the password needs to be changed.
- Login again with the new password for the pureuser account.

Step 4. Once the appliance has been deployed and configured, then the vSphere Remote Plugin can be configured.

Step 5. Use the pureplugin register command and positional arguments to register the remote plugin's extension with the vCenter server: `pureplugin register --host <IP or FQDN of vCenter> --user <vSphere account>`


Step 6. A successful registration message will appear within vCenter soon after.



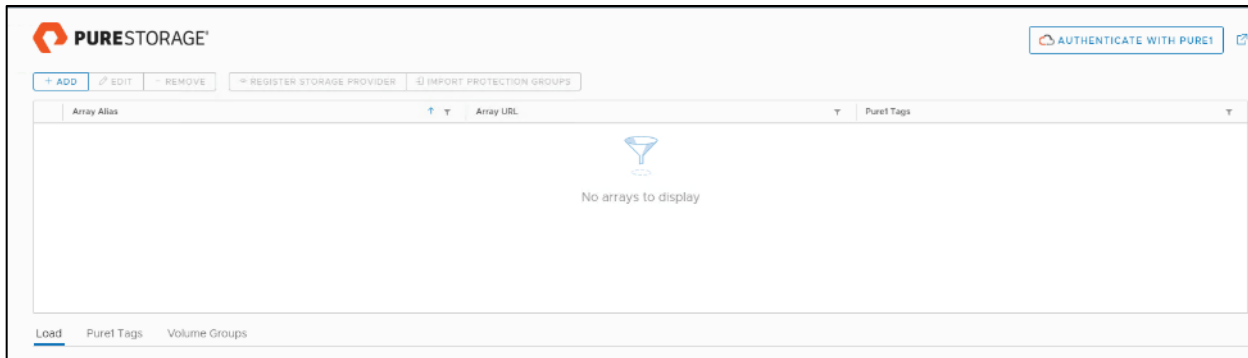
Step 7. After the discovery completes, click **Install**.



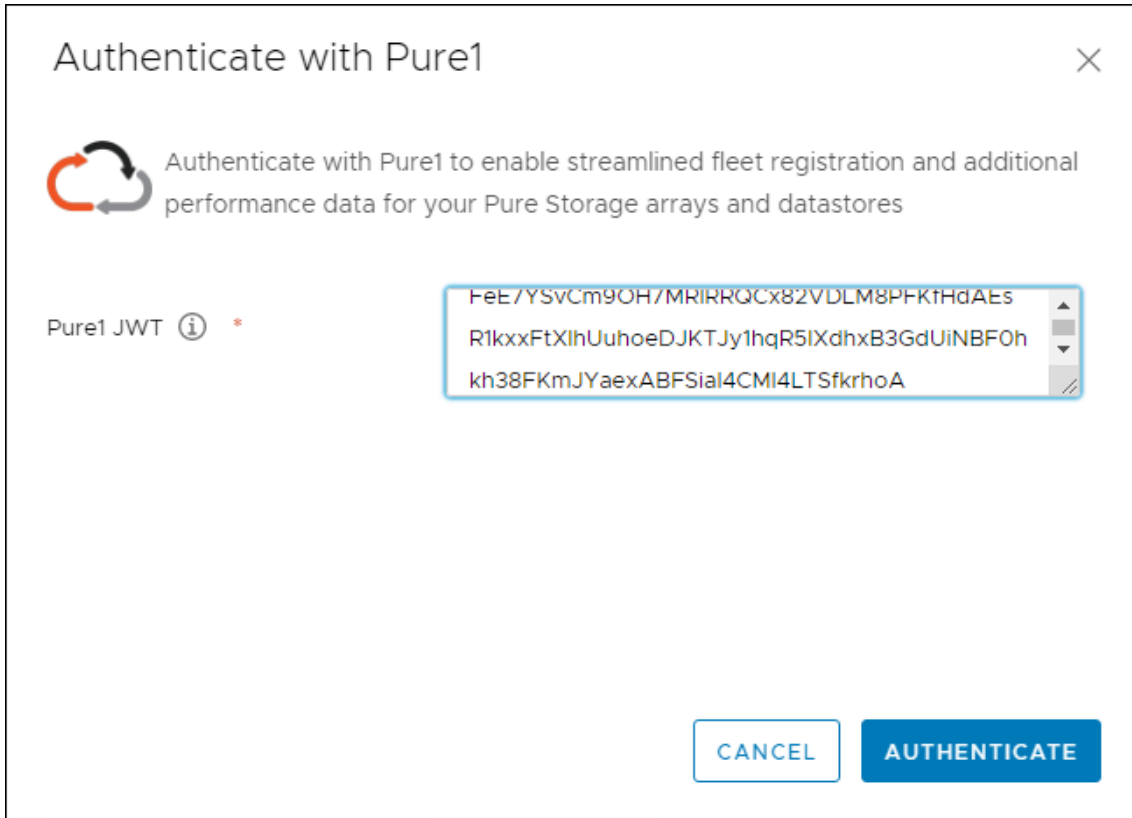
Step 8. In the vCenter HTML5 client, select **Pure Storage** from the Menu.



Step 9. Click **Authenticate with Pure1.**



Step 10. Input your Pure1 JWT (link).

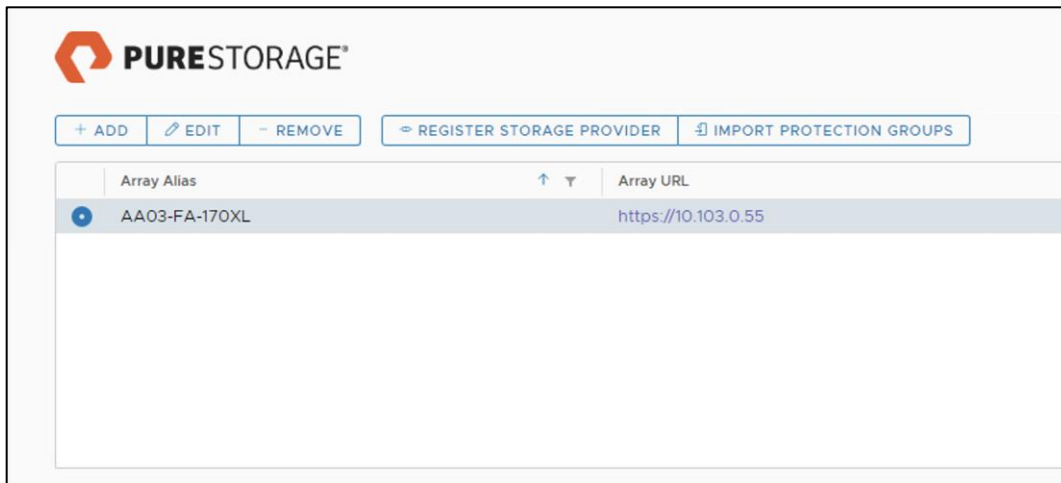


Step 11. Click **Authenticate**.

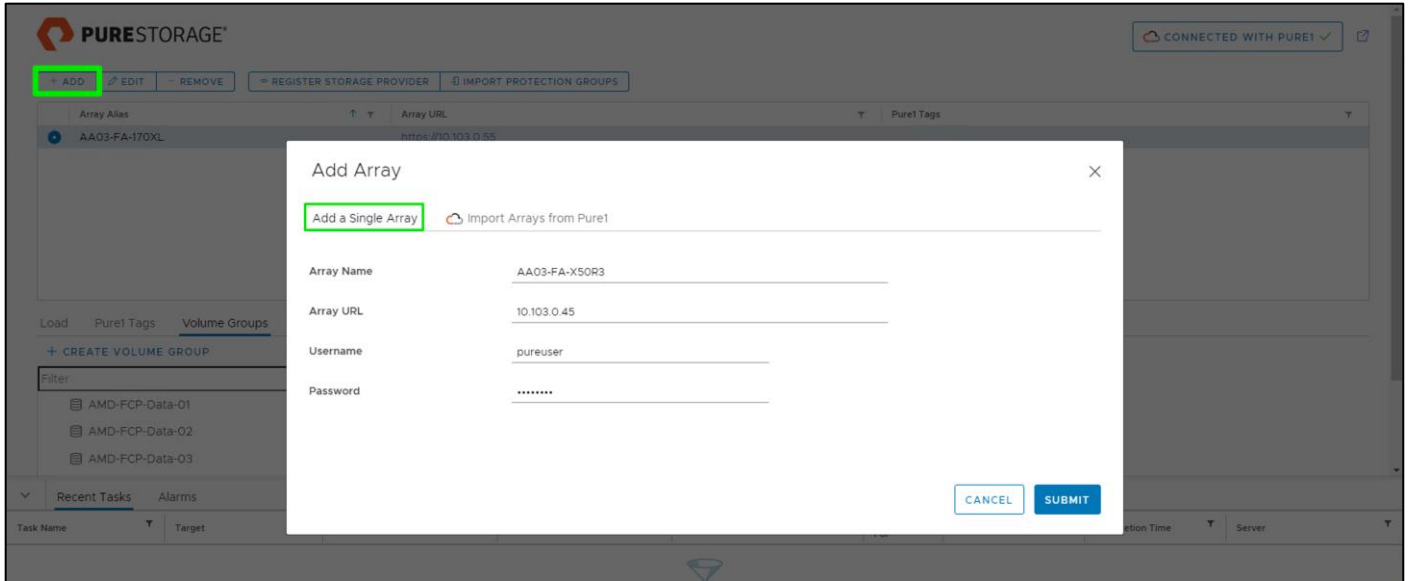
Step 12. Click **Add**.

Step 13. Click **Import Arrays from Pure1** and input the Username and Password.

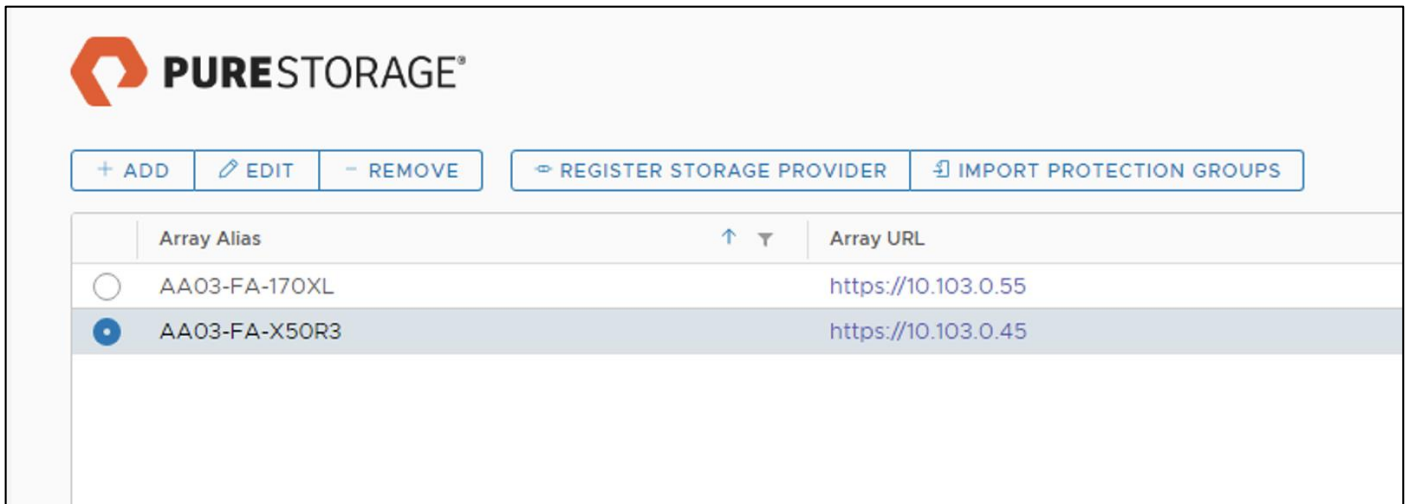
Step 14. Click **Done**.



Step 15. Alternatively, provide array details in the Add a Single Array tab to add the Array manually.



Step 16. Select the newly added array and click **Register Storage Provider**.



Step 17. Enter Username and Password.

Register Storage Provider ×

Registering the storage provider requires a valid username and password.

Username *

Password *

Step 18. Click **Register**. There is also an option to import from Pure1.

Step 19. Select **Import Arrays from Pure1** in Add option.

Add Array ×

Add a Single Array [Import Arrays from Pure1](#)

Use the same credentials for all arrays

<input checked="" type="checkbox"/>	Array Alias	Online	Array URL	Username	Password
<input checked="" type="checkbox"/>	AA03-FA-X50R3		10.103.0.45	pureuser	*****

1 1 - 1 of 1 arrays

Step 20. Select the array and click **Add**.

Add Array ✕

Add a Single Array 🔄 Import Arrays from Pure1

Use the same credentials for all arrays

Array Alias	Online	Array URL	Username	Password
<input type="checkbox"/> AA03-FA-X50R3	📶	10.103.0.45	pureuser	*****

1 - 1 of 1 arrays

Arrays successfully registered: 1

- AA03-FA-X50R3 ✔

Arrays with errors: 0

DONE
ADD


Add VMFS Datastores to the VCF Workload Domain

Additional VMFS datastores for use as principal storage of virtual machines can be added to the VCF workload domain after the initial domain has been configured. However, please note the principal storage type of an existing workload domain cannot be changed. New datastores can be added directly from vCenter using the Pure Storage vSphere Client plugin to simplify operations.

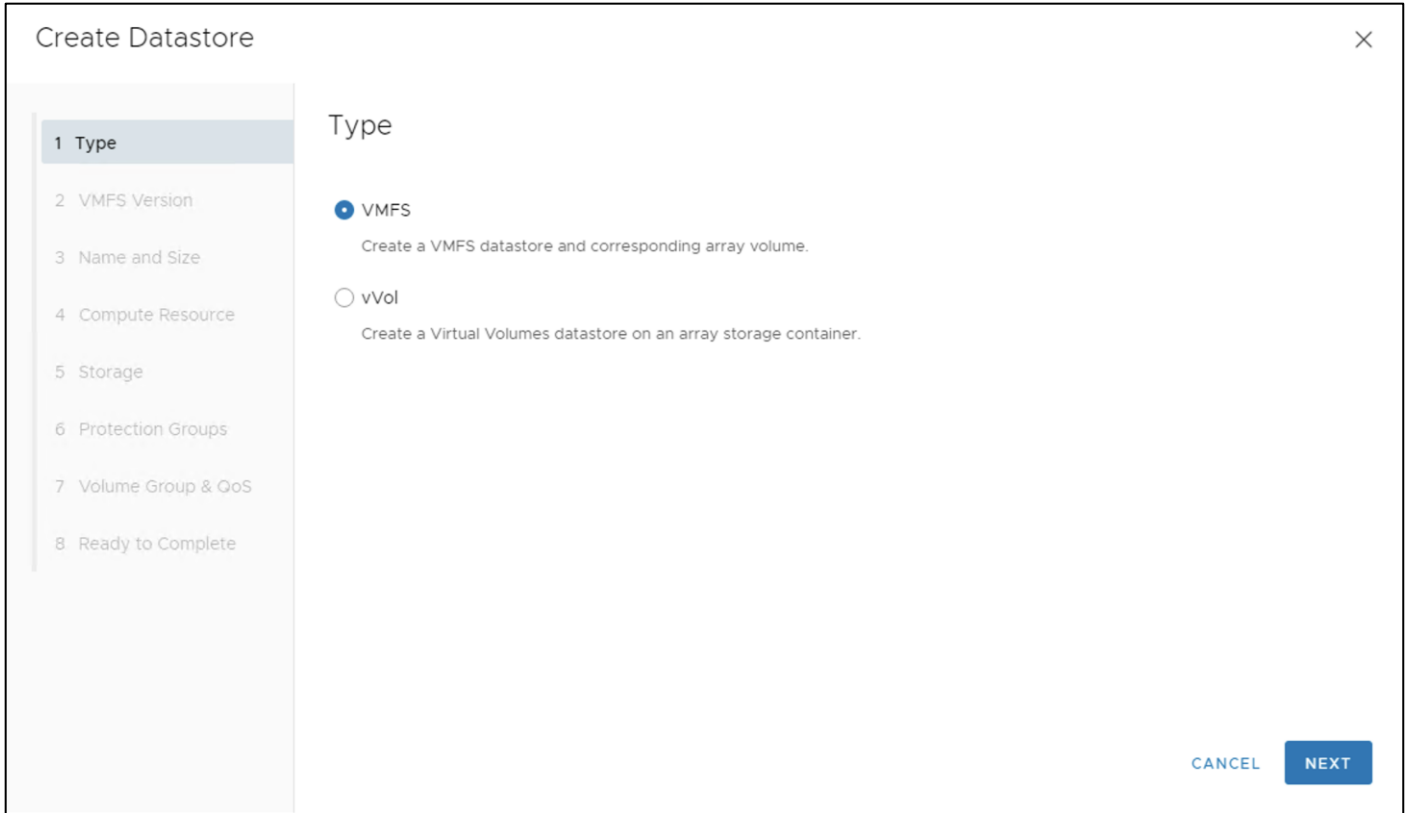
Procedure 2. Create VMFS Datastore using Pure Storage vSphere Client Plugin

Step 1. In the vCenter HTML5 client, click **Host and Clusters**.

Step 2. Right-click the **FlashStack Cluster** and select **Pure Storage > Create Datastore**.



Step 3. Click **VMFS**.

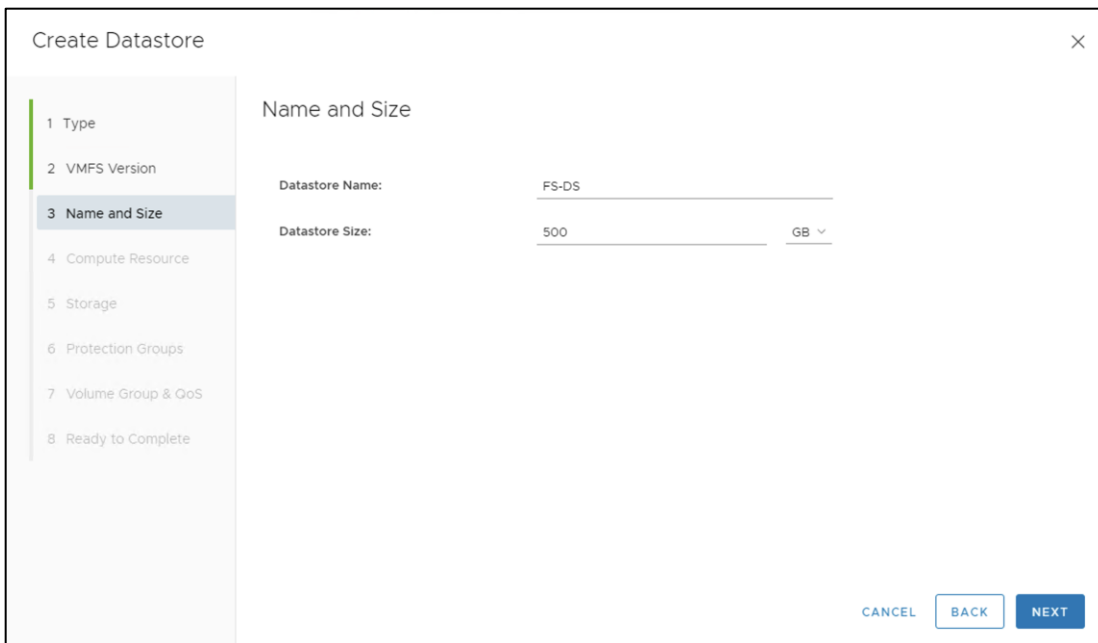


Step 4. Click **Next**.

Step 5. Keep VMFS 6 selected.

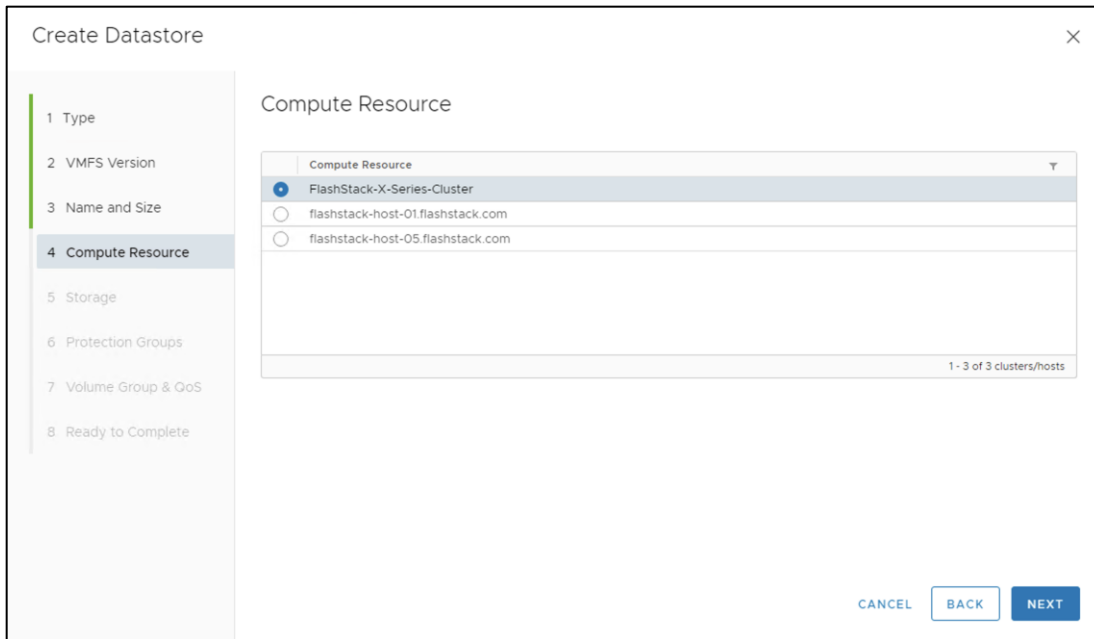
Step 6. Click **Next**.

Step 7. Enter a Datastore Name and Datastore Size.



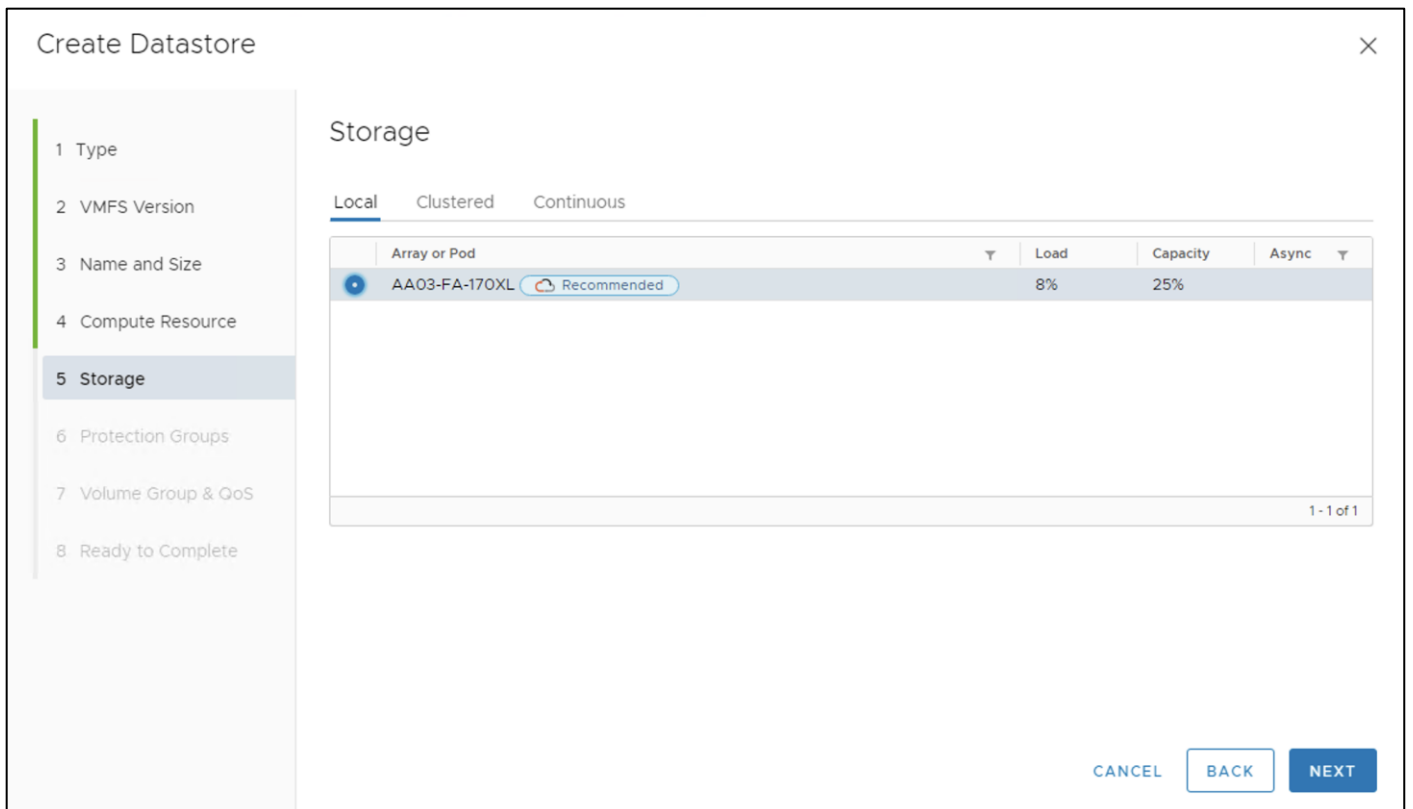
Step 8. Click **Next**.

Step 9. Select the cluster under Compute Resources.

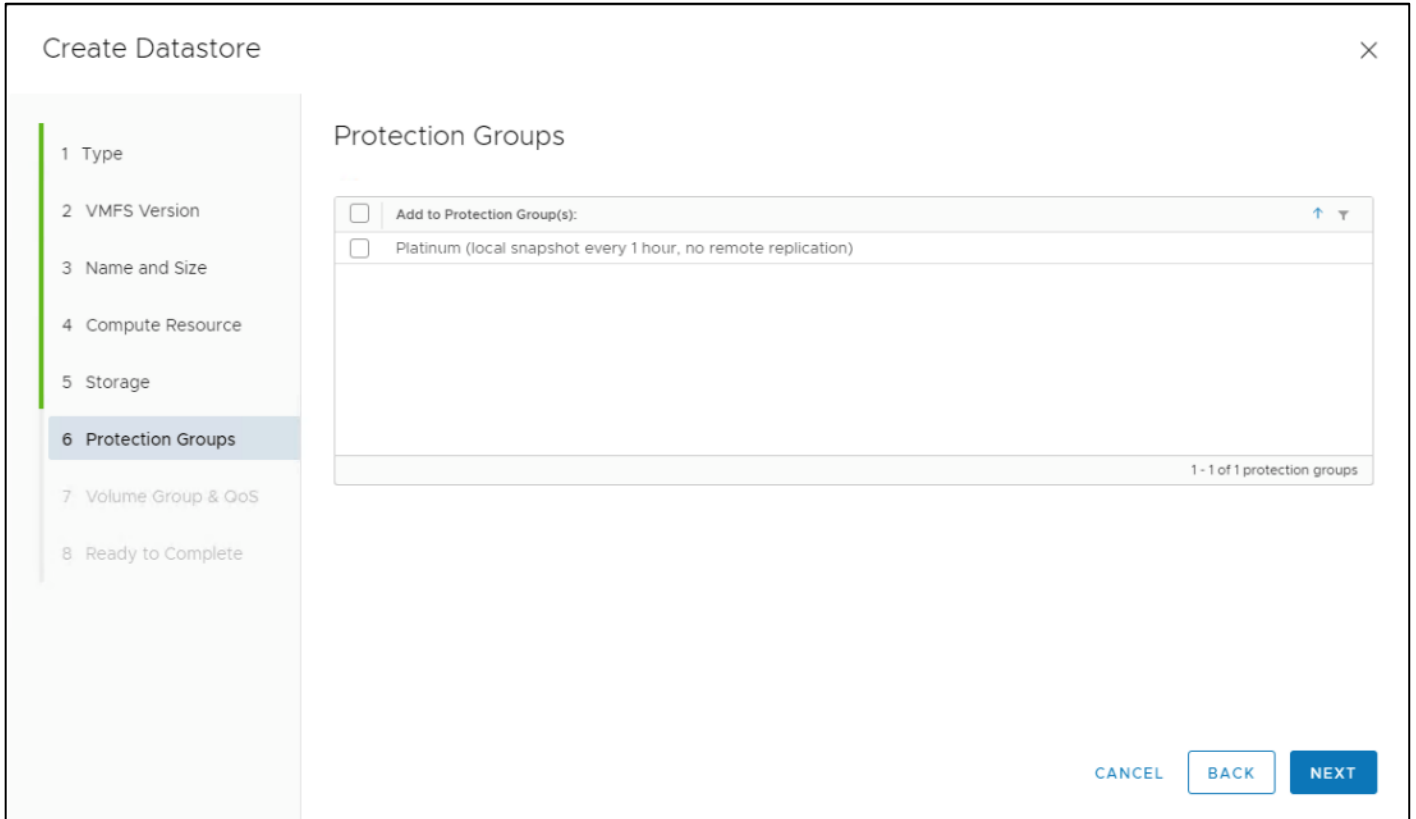


Step 10. Click Next.

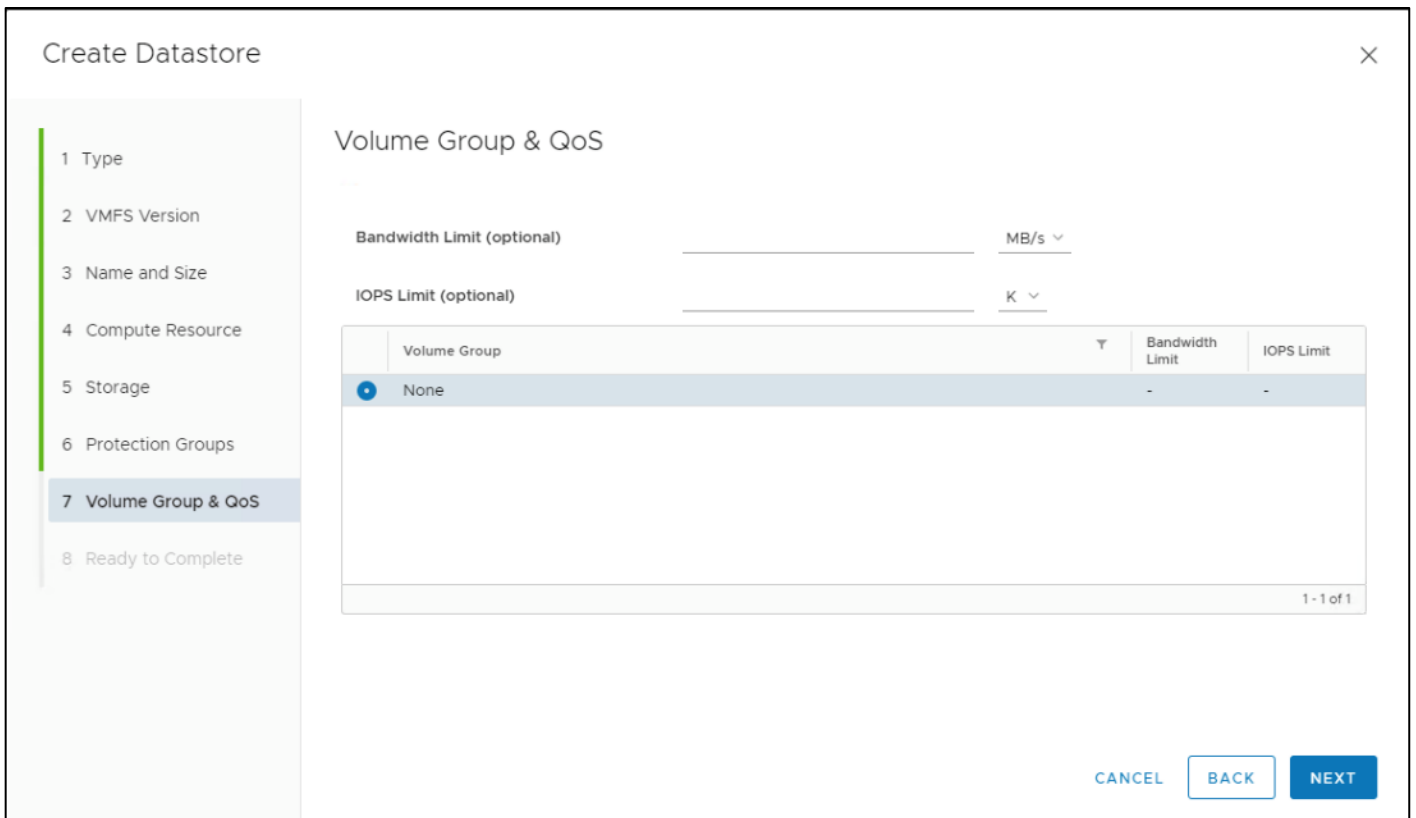
Step 11. Select the Registered FlashArray.



Step 12. Optionally, add to the protection group created earlier and click **Next**.



Step 13. Click **Next** on the Volume Group & QoS page.



Step 14. Review the information and click **Finish**.

Create Datastore ✕

- 1 Type
- 2 VMFS Version
- 3 Name and Size
- 4 Compute Resource
- 5 Storage
- 6 Protection Groups
- 7 Volume Group & QoS
- 8 Ready to Complete

Ready to Complete

Datastore Name:	FS-DS
Type:	VMFS
VMFS Version:	VMFS 6
Datastore Size:	500 GB
Compute Resource:	FlashStack-X-Series-Cluster
Array:	AA03-FA-170XL
Pod:	None
Volume Bandwidth Limit:	-
Volume IOPS Limit:	-
Volume Group:	None
Protection Groups:	None

CANCEL
BACK
FINISH

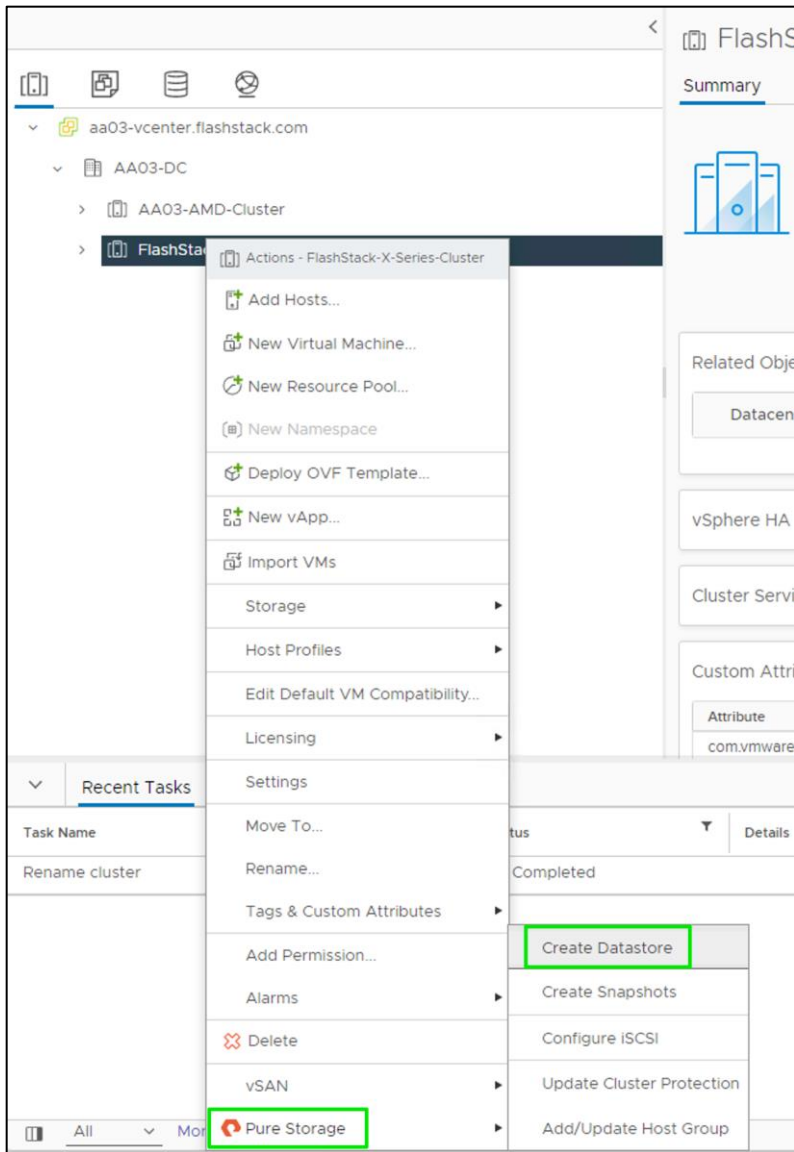
Add vVol Datastores to the VCF Workload Domain

Additional vVol datastores for use as secondary storage of virtual machines can be added to the VCF workload domain after the initial domain has been configured. New datastores can be added directly from vCenter using the Pure Storage vSphere Client plugin to simplify operations.

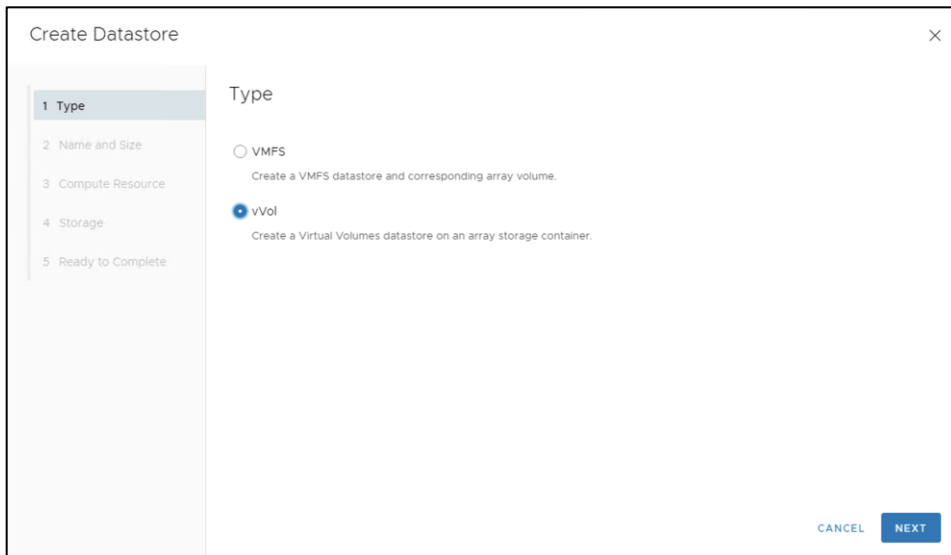
Procedure 3. Create vVol Datastore

Step 1. In the vCenter HTML5 client, select **Host and Clusters**.

Step 2. Right-click the **FlashStack Cluster** and select **Pure Storage > Create Datastore**.



Step 3. Click vVol.



Step 4. Click **Next**.

Step 5. Enter a Datastore Name.

The screenshot shows the 'Create Datastore' dialog box with the 'Name and Size' step selected in the left sidebar. The main area is titled 'Name and Size' and contains a 'Datastore Name' field with the text 'FlashStack-VSI-vVol' entered. Below the field, a note states: 'FlashArray Virtual Volume Datastores are automatically created using the maximum size.' At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

Step 6. Click **Next**.

Step 7. Click the Cluster under Compute Resources.

The screenshot shows the 'Create Datastore' dialog box with the 'Compute Resource' step selected in the left sidebar. The main area is titled 'Compute Resource' and displays a list of compute resources. The first item, 'FlashStack-X-Series-Cluster', is selected with a blue radio button. Below it are two other items: 'flashstack-host-01.flashstack.com' and 'flashstack-host-05.flashstack.com', both with unselected radio buttons. At the bottom right of the list, it says '1 - 3 of 3 clusters/hosts'. At the bottom right of the dialog, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

Step 8. Click **Next**.

Step 9. Click the Registered FlashArray.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource
- 4 Storage
- 5 Ready to Complete

Storage

Array
AA03-FA-170XL

1 - 1 of 1

CANCEL BACK NEXT

Step 10. Click **Next**.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource
- 4 Storage
- 5 Ready to Complete

Ready to Complete

Datastore Name:	pani-test
Type:	vVol
Compute Resource:	AA03-AMD-Cluster
Array:	AA03-FA-170XL
Pod:	None
Storage Provider:	✓ 2 / 2
Storage Container:	✓ default_storage_container
Protocol Endpoint Verified:	✓ Yes

CANCEL BACK FINISH

Step 11. Review the information and click **Finish**.

Conclusion

The FlashStack Datacenter solution is a validated approach for deploying Cisco and Pure Storage technologies and products for building shared private and public cloud infrastructure. VMware Cloud Foundation enables data center administrators to provision an application environment in a quick, repeatable, and automated manner. FlashStack as a workload domain for VMware Cloud Foundation provides the following benefits in any data center environment:

- Integrated solution that supports entire VMware software defined stack
- Standardized architecture for quick, repeatable, error free deployments of FlashStack based workload domains
- Automated life cycle management to keep all the system components up to date
- Simplified cloud-based management of various FlashStack components
- Hybrid-cloud-ready, policy-driven modular design
- Highly available, flexible, and scalable FlashStack architecture
- Cooperative support model and Cisco Solution Support
- Easy to deploy, consume, and manage design which aligns with Cisco, Pure Storage and VMware best practices and compatibility requirements
- Support for component monitoring, solution automation and orchestration, and workload optimization

The success of the FlashStack solution is driven through its ability to evolve and incorporate both technology and product innovations in the areas of management, compute, storage, and networking and this document highlights the deployment details of incorporating FlashStack as a workload domain for VMware Cloud Foundation.

About the Authors

Brian Everitt, Technical Marketing Engineer, Cisco Systems, Inc.

Brian Everitt is an IT industry veteran with over 25 years of experience deploying server, network, and storage infrastructures for companies around the world. During his tenure at Cisco, he has been a lead Advanced Services Solutions Architect for Microsoft solutions, virtualization, and SAP Hana on Cisco UCS. Currently his role covers industry solutions development for Cisco's Converged Infrastructure and Hyperconverged Infrastructure products, focusing on performance evaluation and product quality. Brian has earned multiple certifications from Microsoft, Cisco, and VMware.

Acknowledgements

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

- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Paniraja Koppa, Technical Marketing Engineer, Cisco Systems, Inc.
- Joe Houghes, Senior Solutions Architect, Pure Storage Inc.
- Simranjit Singh, Solutions Architect, Pure Storage Inc.
- Craig Waters, Technical Director, Cisco Solutions, Pure Storage Inc.

Appendices

This appendix contains the following:

- [Appendix A – Workload Domain Description JSON File](#)
- [Appendix B – References Used in Guide](#)
- [Appendix C – Terms Glossary](#)
- [Appendix D – Acronym Glossary](#)
- [Appendix E – Recommended for You](#)

Appendix A – Workload Domain Description JSON File

This appendix contains a complete JSON file for workload domain deployment.

Note: The elements marked by “<#####>” have been retracted.

```
{
  "domainName" : "AA01-WD",
  "vcenterSpec" : {
    "name" : "aa01-vc",
    "networkDetailsSpec" : {
      "ipAddress" : "10.101.1.100",
      "dnsName" : "aa01-vc.vcf.local",
      "gateway" : "10.101.1.254",
      "subnetMask" : "255.255.255.0"
    },
    "rootPassword" : "<#####>",
    "datacenterName" : "AA01-WD-DC",
    "vmSize" : "medium",
    "storageSize" : "1storage"
  },
  "computeSpec" : {
    "clusterSpecs" : [ {
      "name" : "AA01-WD-Cluster",
      "hostSpecs" : [ {
        "id" : "<#####>",
        "licenseKey" : "<#####>",
        "username" : "root",
        "hostNetworkSpec" : {
          "vmNics" : [ {
            "id" : "vmnic0",
            "vdsName" : "vds01"
          }, {
            "id" : "vmnic1",
            "vdsName" : "vds01"
          }, {
            "id" : "vmnic2",
            "vdsName" : "vds02"
          }, {
            "id" : "vmnic3",
            "vdsName" : "vds02"
          } ]
        }
      } ]
    }, {
      "id" : "<#####>",
      "licenseKey" : "<#####>",
      "username" : "root",
      "hostNetworkSpec" : {
        "vmNics" : [ {
          "id" : "vmnic0",
          "vdsName" : "vds01"
        }, {
          "id" : "vmnic1",
          "vdsName" : "vds01"
        } ]
      }
    } ]
  }
}
```

```

    }, {
      "id" : "vmnic2",
      "vdsName" : "vds02"
    }, {
      "id" : "vmnic3",
      "vdsName" : "vds02"
    } ]
  } ]
}, {
  "id" : "<####>",
  "licenseKey" : "<####>",
  "username" : "root",
  "hostNetworkSpec" : {
    "vmNics" : [ {
      "id" : "vmnic0",
      "vdsName" : "vds01"
    }, {
      "id" : "vmnic1",
      "vdsName" : "vds01"
    }, {
      "id" : "vmnic2",
      "vdsName" : "vds02"
    }, {
      "id" : "vmnic3",
      "vdsName" : "vds02"
    } ]
  } ]
}, {
  "datastoreSpec" : {
    "vmfsDatastoreSpec" : {
      "fcSpec" : [ {
        "datastoreName" : "AA01-WD-datastore-01"
      } ]
    }
  }
}, {
  "networkSpec" : {
    "vdsSpecs" : [ {
      "name" : "vds01",
      "portGroupSpecs" : [ {
        "name" : "vds01-pg-mgmt",
        "transportType" : "MANAGEMENT"
      } ]
    }, {
      "name" : "vds02",
      "portGroupSpecs" : [ {
        "name" : "vds02-pg-vmotion",
        "transportType" : "VMOTION"
      } ]
    } ],
    "isUsedByNsxt" : true
  } ],
  "nsxClusterSpec" : {
    "nsxTClusterSpec" : {
      "geneveVlanId" : 3003,
      "ipAddressPoolSpec" : {
        "name" : "AA01-tep-pool",
        "subnets" : [ {
          "ipAddressPoolRanges" : [ {
            "start" : "192.168.3.101",
            "end" : "192.168.3.110"
          } ],
          "cidr" : "192.168.3.0/24",
          "gateway" : "192.168.3.254"
        } ]
      }
    }
  }
} ]
}, {
  "nsxISpec" : {
    "nsxManagerSpecs" : [ {
      "name" : "vcf-wd-nsx-1",
      "networkDetailsSpec" : {

```

```

    "ipAddress" : "10.101.1.96",
    "dnsName" : "vcf-wd-nsx-1.vcf.local",
    "gateway" : "10.101.1.254",
    "subnetMask" : "255.255.255.0"
  }
}, {
  "name" : "vcf-wd-nsx-2",
  "networkDetailsSpec" : {
    "ipAddress" : "10.101.1.97",
    "dnsName" : "vcf-wd-nsx-2.vcf.local",
    "gateway" : "10.101.1.254",
    "subnetMask" : "255.255.255.0"
  }
}, {
  "name" : "vcf-wd-nsx-3",
  "networkDetailsSpec" : {
    "ipAddress" : "10.101.1.98",
    "dnsName" : "vcf-wd-nsx-3.vcf.local",
    "gateway" : "10.101.1.254",
    "subnetMask" : "255.255.255.0"
  }
} ],
"vip" : "10.101.1.95",
"vipFqdn" : "vcf-wd-nsx.vcf.local",
"licenseKey" : "<####>",
"nsxManagerAdminPassword" : "<####>",
"formFactor" : "medium"
}
}

```

Appendix B - References Used in Guide

Compute

Cisco Intersight: <https://www.intersight.com>

Cisco Intersight Managed Mode:

https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide.html

Cisco Unified Computing System: <http://www.cisco.com/en/US/products/ps10265/index.html>

Cisco UCS 6400 Series Fabric Interconnects: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html>

Network

Cisco Nexus 9000 Series Switches: <http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>

Cisco MDS 9132T Switches: <https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html>

Storage

Pure Storage FlashArray//X: <https://www.purestorage.com/products/nvme/flasharray-x.html>

Pure Storage Purity//FA: <https://www.purestorage.com/products/storage-software/purity.html>

Pure Storage Pure1: <https://www.purestorage.com/products/aiops/pure1.html>

Virtualization

VMware Cloud Foundation 4.5 release notes: <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/rn/vmware-cloud-foundation-45-release-notes/index.html>

VMware Cloud Foundation 4.5 Deployment Guide: <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-F2DCF1B2-4EF6-444E-80BA-8F529A6D0725.html>

VMware vCenter Server: <http://www.vmware.com/products/vcenter-server/overview.html>

VMware vSphere: <https://www.vmware.com/products/vsphere>

Interoperability Matrix

Cisco UCS Hardware Compatibility Matrix: <https://ucshcltool.cloudapps.cisco.com/public/>

VMware and Cisco Unified Computing System: <http://www.vmware.com/resources/compatibility>

Pure Storage Interoperability Matrix Tool: <http://support.Pure Storage.com/matrix/>

Appendix C - Terms Glossary

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)	<p>Some IT capability, X, provided as a service (XaaS). Some benefits are:</p> <ul style="list-style-type: none">• 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. <p>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.</p> <p>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.</p> <p>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.</p>
Ansible	<p>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).</p> <p>https://www.ansible.com</p>
AWS	Provider of IaaS and PaaS.

(Amazon Web Services)	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)	<p>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).</p> <p>https://www.docker.com</p> <p>https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html</p>
DevOps	<p>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.</p> <p>https://en.wikipedia.org/wiki/DevOps</p> <p>https://en.wikipedia.org/wiki/CI/CD</p>
Edge compute	<p>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.</p> <p>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.</p> <p>https://en.wikipedia.org/wiki/Mobile_edge_computing</p>
IaaS (Infrastructure as-a-Service)	<p>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).</p>
IaC (Infrastructure as-Code)	<p>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.</p> <p>https://en.wikipedia.org/wiki/Infrastructure_as_code</p>
IAM (Identity and Access Management)	<p>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.</p> <p>https://en.wikipedia.org/wiki/Identity_management</p>
IBM (Cloud)	<p>IBM IaaS and PaaS.</p> <p>https://www.ibm.com/cloud</p>
Intersight	<p>Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.</p> <p>https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html</p>

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-Service)	End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider.
SAML (Security Assertion Markup Language)	Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by the aaS for access control decisions. 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

Appendix D - Acronym Glossary

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

VCF—VMware Cloud Foundation

vHBA—virtual Host Bus Adapter

VLAN—Virtual Local Area Network

VM—Virtual Machine

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

VNI—Virtual Network Identifier (VXLAN)

vNIC—virtual Network Interface Card

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)

Appendix E - Recommended for You

FlashStack for Virtual Server Infrastructure with End-to-End 100G Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_vsi_ucs_xseries_5gen_design.html

FlashStack for Virtual Server Infrastructure with End-to-End 100G Deployment Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_xseries_e2e_5gen.html

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

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DE-SIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WAR-RANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_P2)

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)