



System Requirements

This section describes the requirements that are necessary to deploy Cisco Container Platform.

It contains the following topics:

- [Supported Version Matrix, on page 1](#)
- [Software Requirements, on page 1](#)
- [Hardware Requirements, on page 2](#)
- [Resource Management Requirements, on page 2](#)
- [Network Requirements, on page 3](#)
- [HyperFlex Integration Requirements, on page 6](#)
- [ACI Integration Requirements, on page 8](#)
- [GPU Integration Requirements, on page 11](#)

Supported Version Matrix

Cisco Container Platform uses various software and hardware components. The following table provides information on the validated versions of each component.

| Component | Validated Version |
|--------------------|-------------------|
| Kubernetes | 1.12 |
| | 1.13 |
| vSphere | vSphere 6.0 (u2)+ |
| | vSphere 6.5 |
| HyperFlex software | 4.0.1 |
| ACI | 4.1(1) |

Software Requirements

Ensure that the following software applications are installed in your deployment environment:

- VMware vCenter server 6.5

- VMware client integration plugin
- vSphere Flash client

Hardware Requirements

- If you are enabling VMware EVC Mode, you must use an Ivy Bridge or a later micro-architecture so that the CPU RDRAND instruction set is available.
- In Cisco Container Platform 1.3.0 or later, the hypervisor hosts need to run CPUs with an Ivy Bridge (UCS C220 M4) or newer micro-architecture so that the CPU RDRAND instruction set is available.
- In the Cisco Container Platform Control Plane VM, each master and worker node requires 2 vCPUs, 8 GB memory, and 40 GB HDD.
- In the Cisco Container Platform Tenant Cluster VM, each master and worker node requires 2 vCPUs, 16 GB memory, and 40 GB HDD. You can modify the vCPU and memory configurations when you deploy a new tenant cluster.

Resource Management Requirements

The following topics provide information on the necessary resource management requirements:

Enabling DRS and HA on Clusters



Note You must use the Enterprise Plus license to set up VMware clusters with HA and DRS enabled. For more information on the supported versions of VMware, see [Supported Version Matrix, on page 1](#).

It is required that you enable DRS and HA on vCenter for the following reasons:

- DRS continuously monitors resource utilization across vSphere servers and intelligently balances VMs on the servers.
- HA provides easy to use, cost-effective high availability for applications running on virtual machines.


-
- Step 1** In the vSphere Web Client, navigate to the host or cluster on which you want to deploy Cisco Container Platform.
- Step 2** Click the **Configure** tab.
- Step 3** Under **Services**, click **vSphere DRS**, and then click **Edit**.
- Step 4** In the right pane of the **Edit Cluster Settings** window, check the **Turn ON vSphere DRS** check box, and then click **OK**.
- Step 5** Under **Services**, click **vSphere Availability**, and then click **Edit**.
- Step 6** In the right pane of the **Edit Cluster Settings** window, check the **Turn ON vSphere HA** check box, and then click **OK**.
-

Enabling NTP Services

You need to enable the Time Synchronization services on each host within your vSphere environment. If you do not enable this service, errors due to timing differences between hosts may cause installation of the Cisco Container Platform to fail.

- Step 1** In the vSphere Web Client, navigate to the host or cluster on which you want to deploy Cisco Container Platform.
- Step 2** Click the **Configure** tab.
- Step 3** From the left pane, expand **System**, and then click **Time Configuration**.

Figure 1: Time Configuration pane



- Step 4** In the right pane, click **Edit**.
- Step 5** In the **Edit Time Configuration** window, check the **Turn ON vSphere DRS** check box, and then click **OK**.

Note You must ensure that each host has DNS access to enable NTP services.

Network Requirements

The following topics provide information on the necessary network requirements:

If you have chosen Contiv as the CNI, the pod-to-pod traffic across nodes is tunneled by the VXLAN protocol.

Provisioning a Port Group for Cisco Container Platform VM Deployment

Cisco Container Platform creates VMs that are attached to a Port Group on either a vSphere Standard Switch (VSS) or a Distributed Virtual Switch (DVS). The HyperFlex installer creates VSS switches in vSphere for the networks that are defined during installation. You need to create either VSS or DVS Switches for managing the VM traffic.

The following topics provide information on configuring a VSS or a DVS.

Configuring vSphere Standard Switch

- Step 1** In the vSphere Web Client, navigate to the host or cluster on which you want to deploy Cisco Container Platform.
 - Step 2** Click the **Configure** tab.
 - Step 3** Expand **Networking**, and then select **Virtual switches**.
 - Step 4** Click **Add host networking**.
 - Step 5** Choose **Virtual Machine Port Group for a Standard Switch** as the connection type for which you want to use the new standard switch and click **Next**.
 - Step 6** Select **New standard switch** and click **Next**.
 - Step 7** Add physical network adapters to the new standard switch.
 - Step 8** Under **Assigned adapters**, click **Add adapters**.
 - Step 9** Select one or more physical network adapters from the list.
 - Step 10** From the **Failover order group** drop-down list, choose from the Active or Standby failover lists.
 - Step 11** For higher throughput and to provide redundancy, configure at least two physical network adapters in the Active list.
 - Step 12** Click **OK**.
 - Step 13** Enter connection settings for the adapter or the port group as follows:
 - a) Enter a network Label or the port group, or accept the generated label.
 - b) Set the VLAN ID to configure VLAN handling in the port group.
 - Step 14** On the **Ready to Complete** screen, click **OK**.
-

Configuring Distributed Virtual Switch

- Step 1** In the **Navigation** pane, click the DVS switch.
 - Step 2** In the right pane, click the **Hosts** tab.
 - Step 3** Click the **Actions** icon and click the **Add and Manage Hosts** radio button. The **Add and Manage Hosts** wizard appears.
 - Step 4** In the **Select tasks** screen, click the **Add Hosts** radio button, and then click **Next**.
 - Step 5** In the **Select hosts** screen, click the **Add Hosts** icon.
 - Step 6** In the **Select new hosts** screen, check the check box next to the hosts that you want to add, and then click **OK**.
 - Step 7** Click **Next** in the **Select network adapter tasks** screen.
 - Step 8** In the **Manage physical network adapters** screen, click the network switch that you want to configure, and then click the **Assign** uplink.
 - Step 9** Repeat Step 8 for all the networks, and click **Next**.
 - Step 10** In the **Manage VMKernel network adapters** screen, click **Next**.
 - Step 11** In the **Analyze impact** screen, click **Next**.
 - Step 12** In the **Ready to complete** screen, click **Finish**.
-

Configuring DHCP Server

Cisco Container Platform requires a DHCP server to be present. The Cisco Container Platform installer VM and upgrade VM get their primary interface IP addresses from the DHCP server. You must ensure that you have configured a DHCP server.

If the DHCP server does not provide the location of the NTP service, enter the NTP address in the Installer UI, under **Control Plane Settings > Advanced Settings**.

Reserving IP Addresses for Static Allocation

Cisco Container Platform uses static IP addresses for all cluster nodes and the **CCP Control Plane master node VIP**, which provides worker nodes with a consistent IP address. Additionally, a load balancer VIP is used as an external IP address for NGINX Ingress in each Kubernetes cluster. These VIPs are configured using IP pools. The static IP addresses are assigned from the same subnet as the load balancer VIP addresses, and you must ensure that the static IP address pools for the subnet do not overlap with a DHCP pool.

Static and DHCP IP Address Requirements

You must ensure that the following conditions are met:

- The subnet is routable to and from the VMware vCenter server.
- The client install machine is routable to the network during the Cisco Container Platform control plane install.
- The network allows communication between Cisco Container Platform VM instances. You must not use a private LAN.

The following sections summarize the static and DHCP IP address requirements for the Cisco Container Platform components:

Calico

| Component | Static IP | DHCP IP |
|--|---|---------|
| Installer VM | 0 | 1 |
| Tenant clusters | 3 + Number of load balancer VIPs desired for applications + Number of workers | 0 |
| Control Plane and Cisco Container Platform web interface | 6 Note 1 for the Kubernetes master VIP, 1 for the Ingress LoadBalancer, and 1 each for the master and worker nodes. | 0 |

By default, the Cisco Container Platform Control Plane pod network uses the 192.168.0.0/16 subnet for Calico. If you have routed IP addresses in that space, you must assign another RFC1918 range for your VXLAN

network. It does not need to be a full /16 subnet, a /22 subnet is adequate for the Cisco Container Platform Control Plane.

Contiv

| Component | Static IP | DHCP IP |
|--|--|---------|
| Installer VM | 0 | 1 |
| Tenant clusters | 4 + Number of load balancer VIPs desired for applications + (2 x Number of workers) | 0 |
| Control Plane and Cisco Container Platform web interface | 10 Note 1 for the Kubernetes master VIP, 1 for the Ingress LoadBalancer VIP, 2 IP for the master node, and 6 for the worker nodes. | 0 |

Contiv requires a /14 or larger subnet for pods. The Cisco Container Platform web interface default of 192.168.0.0/16 is not large enough for Contiv. You must find an RFC1918 /14 subnet that is not routed within your organization, such as one from 172.16.0.0/12 or 10.0.0.0/8.

HyperFlex Integration Requirements



Note This section is applicable only if you want to use HyperFlex environment. It is not required for running VMware on UCS.

Cisco Container Platform is supported on all hardware configurations that are supported by the required HyperFlex software versions. For more information on HyperFlex hardware configurations, refer to the UCS HyperFlex product documentation.

The following topics provide information on the necessary HyperFlex integration requirements:

Configuring Shared Datastore

After HyperFlex is installed, you need to configure a shared datastore. The datastore must be accessible to hosts such as NFS or iSCSI or FC in the cluster.

The datastore is required for the following purposes:

- Provisioning persistent volume storage
- Deploying the Cisco Container Platform tenant base VM

-
- Step 1** Log in to the **HX Connect UI** using the VMware vCenter SSO administrator credentials.
For more information on installing HyperFlex and accessing the HyperFlex Connect UI, refer to the [latest HperFlex documentation](#).
- Step 2** In the left pane, click **Manage > Datastores**.
- Step 3** Perform these steps to create a datastore for provisioning the Kubernetes persistent volume storage and deploying the Cisco Container Platform tenant base VM:
- In the right pane, click **Create Datastore**.
 - In the **Name** field, enter **ds1**, and then enter a size and block size for the datastore.
Note We recommend that you use **1TB** size and **8K** block size.
 - Click **Create Datastore**.
- The newly created datastore is available on vCenter.
-

Configuring Link-local Network for HyperFlex iSCSI Communication

The FlexVolume plug-in requires a host-only link between each VM that runs Kubernetes and the Internet Small Computer System Interface (iSCSI) target on the ESX host.

For HyperFlex 3.5+

-
- Step 1** Log in to the **HX Connect UI**.
- Step 2** Choose **Settings > Integrations > Kubernetes**.
- Step 3** Click **Enable All Node** and wait until the **KUBERNETES STORAGE PROVISIONING** option is enabled.
The HyperFlex infrastructure is configured and ready to use for Cisco Container Platform with Kubernetes persistent volume support.
-

For HyperFlex 3.0.x

-
- Step 1** Open an SSH session to the HyperFlex 3.0 Platform Installer VM or one of the HyperFlex Controller VMs and log in as a root user.
- Step 2** Perform these steps to get the vCenter details that you need to enter when you run the `add_vswitch.py` script.
- Run the following command to get the vCenter datacenter name and vCenter cluster name.

```
stcli cluster info | grep -i vcenter
```
 - Run the following command to validate the reachability of vCenter IP address.

```
ping <vcenter URL>
```
- Step 3** Navigate to the following location:
`/usr/share/springpath/storfs-misc/hx-scripts/`
- Step 4** Run the `add_vswitch.py` script.

```
python add_vswitch.py --vcenter-ip <vCenter IP address>
```

When prompted, specify the vCenter credentials, datacenter name, and cluster name that you got from the output of Step 2.

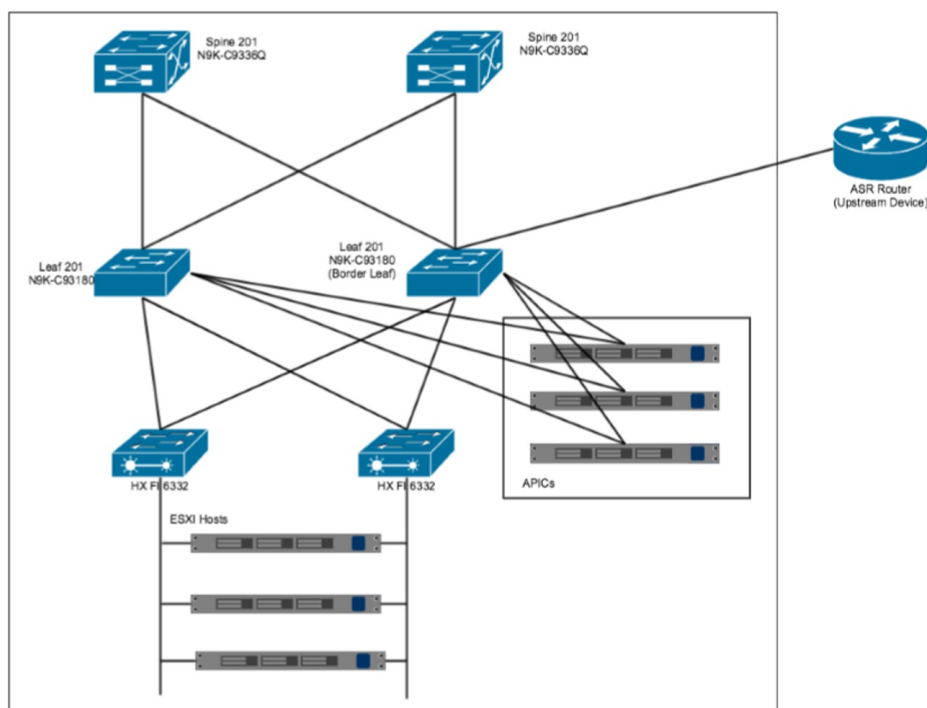
The HyperFlex infrastructure is configured and ready to use for Cisco Container Platform with Kubernetes persistent volume support.

ACI Integration Requirements

Cisco ACI enables you to group your application into End Point Groups (EPGs), define policies for the EPGs, and then deploy network policies on the ACI fabric. The policy enforcement is implemented using the spine and leaf architecture of the ACI fabric.

The following figure shows the components of a Cisco Container Platform ACI integrated network topology.

Figure 2: Cisco Container Platform ACI Integrated Network Topology



The main components of the network topology are as follows:

- **ACI Fabric** includes two spine nodes, two leaf nodes, and three APIC controllers. You can choose the number of the spine and leaf nodes and APIC controllers as per your network requirement.
- **HyperFlex Fabric Interconnect (FI)** includes two fabric interconnect switches connected between the ESXi hosts and the ACI leaf switches.
- **ESXi Hosts** includes a UCS server such as UCS C220 M4.

- **ASR router** is connected to an ACI border leaf for external internet access.

APIC Controller Requirements

If you are using ACI, ensure that you have configured the following settings on the APIC controller:

- Assign a port number other than 4094 for Infra VLAN as 4094 is reserved for provisioning HyperFlex fabric interconnect
- Create a common tenant
- Create a Virtual Route Forwarder (VRF) in the common tenant
- Create at least one L3OUT
- Create an Access Entity Profile (AEP) for the ACI tenant physical domain
- Create an AEP for L3OUT
- Create a Virtual Machine Manager (VMM) domain which connects to vSphere

For more information on configuring an APIC controller, refer to the [latest ACI documentation](#).

HyperFlex FI Requirements


Ensure that you have configured the following settings on HyperFlex FI:

- Configure QoS
 1. From the left pane, click **LAN**.
 2. From the right pane, click the **QoS** tab, and then configure QoS.



Note Using the **MTU** configuration, you must set the priority that is associated with the QoS policy of the vNIC template.

Figure 3: QoS Tab



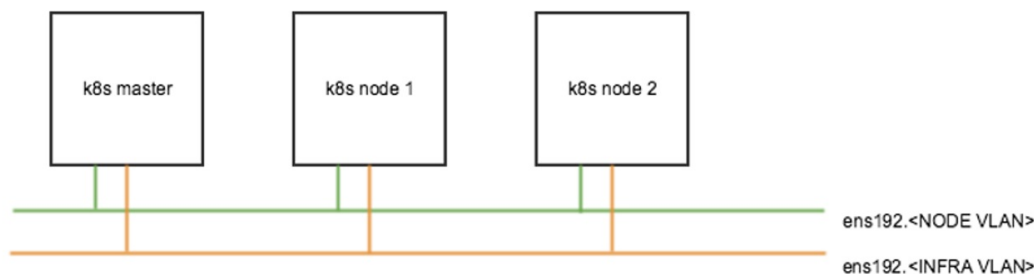
- Ensure that the tenant VLAN is allowed

Once Cisco Container Platform Control Plane and management node networking are configured, you can access the HyperFlex cluster on vSphere and install Cisco Container Platform. Each time you create a tenant cluster, the ACI constructs such as L3OUT, VRF, and AEP stored in the common tenant cluster are reused.

Tenant Cluster with ACI Deployment

With an ACI deployment, each tenant cluster is required to have its own routable subnet. The node VLAN, pod subnet, and multicast subnet range should not overlap between clusters. Cisco Container Platform ensures that the VLAN and subnet do not overlap.

Unlike other CNI, an ACI tenant cluster requires two VLAN subinterfaces, one for the Node VLAN, and another for the Infra VLAN. As shown in the following figure, Cisco Container Platform assigns unique Node VLAN IDs. You need to assign a unique Infra VLAN ID for clusters during cluster creation.



For more information on creating tenant clusters, refer to the *Creating Kubernetes Clusters* section of the *Cisco Container Platform User Guide*.

For more information on the ACI and CNI plugin, refer to the [latest documentation on Cisco ACI and Kubernetes Integration](#).

GPU Integration Requirements

Cisco Container Platform supports GPU devices in passthrough mode to enable AI/ML workloads.

This section describes the requirements on the ESXi and vCenter hosts to integrate the GPU devices with Cisco Container Platform.

- Step 1** Follow these steps to enable GPU Passthrough for the devices that you want to use:
- Access the ESXi host by typing its IP address in a web browser.
 - From the left pane, click **Manage**.
 - In the right pane, click **Hardware > PCI Devices** .
The list of available passthrough devices is displayed.
 - Select the device, and then click **Toggle Passthrough**.
- Step 2** Follow these steps to enable shared direct passthrough for the GPU device:
- Access the vCenter server by typing its IP address in a web browser.
 - From the right pane, click **Configure > Graphics > Graphics Devices**.
 - Select the device for which you want to enable shared direct passthrough.
 - In the **Edit Graphics Device Settings** dialog box, click the **Shared Direct** radio button.
 - Click **Ok**.
- Step 3** Follow these steps to allow VM access to the GPU device:
- From the right pane, click **Configure > PCI Devices**.
 - Click the **Edit** icon.
The **Edit PCI Device Availability** dialog box appears.
 - Select the device and check the checkbox next to the device.
 - Click **OK**.
-

