

Prerequisites: Orchestrator

- Requirements for Orchestrator, on page 1
- Communication Ports for Orchestrator, on page 2
- Fabric Requirements for Orchestrator, on page 3

Requirements for Orchestrator

Note This section describes *additional* requirements and guidelines if you plan to enable the Orchestrator service. Ensure that you have already satisfied the platform-level requirements described in the Prerequisites and guidelines for all enabled services section.

- Starting with Nexus Dashboard release 3.1.1, Cisco DC App Center connectivity has been removed from Nexus Dashboard because downloading the services separately is no longer required.

  To deploy Orchestrator, download the unified installation image from the Software Download page; individual services' installation images are no longer available from the Cisco DC App Center.

- If you are deploying Nexus Dashboard Orchestrator to manage Cisco ACI fabrics, you can establish connectivity from either the data interface or the management interface to either the in-band or out-of-band (OOB) interface of each fabric's APIC or both.

  Note that if the fabric connectivity is from the Nexus Dashboard's management interface, you must configure specific static routes or ensure that the management interface is part of the same IP subnet of the APIC interfaces.

- If you are deploying Nexus Dashboard Orchestrator to manage Cisco NDFC fabrics, the data network must have in-band reachability for Cisco NDFC fabrics.

- Connectivity between the nodes is required on both networks with the following additional round trip time (RTT) requirements:

Table 1: Orchestrator RTT Requirements

Connectivity	Maximum RTT
To any managed APIC fabrics	500 ms

Connectivity	Maximum RTT
To any managed NDFC fabrics	150 ms

Communication Ports for Orchestrator

In addition to the ports required by the Nexus Dashboard cluster nodes (listed in a previous section), the following ports are required by the Orchestrator service.

Table 2: Nexus Dashboard Orchestrator Ports (Management Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SCP or SFTP	22	TCP	In/Out	Remote servers for storing backups and downloading software upgrade images
HTTP	80	TCP	Out	Splunk or syslog server if external log streaming is enabled
HTTPS	443	TCP	In/Out	Splunk or syslog server if external log streaming is enabled

Table 3: Nexus Dashboard Orchestrator Ports (Data Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
HTTPS	443	TCP	Out	In-band of switches and APIC/NDFC
Kafka	3001	TCP	In/Out	In-band IP of switches and APIC/NDFC

Fabric Requirements for Orchestrator

The following additional fabric-related guidelines apply to the Orchestrator service:

- Cisco Mini ACI fabrics are supported as typical on-premises fabrics without requiring any additional configuration.

  Detailed info on deploying and configuring this type of fabrics is available in *Cisco Mini ACI Fabric and Virtual APICs*.

- If you are managing ACI fabrics that contain Remote Leaf switches, the following restrictions apply:
  - Only physical Remote Leaf switches are supported.
  - Only -EX and -FX or later switches are supported as Remote Leaf switches.
  - Remote Leaf is not supported with back-to-back connected fabrics without IPN switches.
  - Remote Leaf switches in one fabric cannot use another fabric's L3Out.
  - Stretching a bridge domain between one fabric (local leaf or remote leaf) and a Remote Leaf in another fabric is not supported.

  You must also perform the following tasks before the fabric can be added to and managed by the Nexus Dashboard Orchestrator:

  - You must enable Remote Leaf direct communication directly in the fabric's APIC.

    To enable direct communication, log in to the fabric's APIC, navigate to **System** > **System Settings** > **Fabric Wide Setting** and **Enable Remote Leaf Direct Traffic Forwarding**.

Note You cannot disable this option after you enable it.

  - You must configure external TEP pools for remote leaf switches.

    To configure one or more external TEP pools, log in to the fabric's APIC and navigate to **Fabric** > **Inventory** > **Pod Fabric Setup Policy**. Then double-click the pod where you want to configure the subnets and click **+** in the **External TEP** area. Finally, enter the **IP** and **Reserve Address Count**, set the state to `Active` or `Inactive`, then click **Update** to save the subnet.

    When configuring external TEP pools, you must provide a netmask between `/22` and `/29`. Multiple, non-contiguous, external TEP pools can be configured, including at different points in time.

  - You must add the routable IP addresses of APIC nodes (assigned from the defined external TEP pool) in the DHCP-Relay configuration applied on the interfaces of the Layer 3 routers connecting to the Remote Leaf switches.

    The routable IP address of each APIC node is listed in the **Routable IP Address** field of the **System** > **Controllers** > **<controller-name>** screen of the APIC GUI.

- You must configure Pod Profile, Policy Group, and Fabric Access policies as described in the following sections.

Pod Profile and Policy Group

In each fabric's APIC, you must have one Pod profile with a Pod policy group. If your fabric does not have a Pod policy group you must create one. Typically, these settings will already exist as you will have configured them when you first deployed the fabric.

Procedure

Step 1 Log in to the fabric's APIC GUI.

Step 2 Check that the Pod profile contains a Pod policy group.

Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**.

Step 3 If necessary, create a Pod policy group.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Policy Groups**.
b) Right-click **Policy Groups** and select **Create Pod Policy Group**.
c) Enter the appropriate information and click **Submit**.

Step 4 Assign the new Pod policy group to the default Pod profile.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**
b) Select the default profile.
c) Choose the new pod policy group and click **Update**.

Configuring Fabric Access Global Policies

This section describes the global fabric access policy configurations that must be created for each APIC fabric before it can be on-boarded on the Nexus Dashboard cluster and managed by the Nexus Dashboard Orchestrator.

Procedure

Step 1 Log in directly to the fabric's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

You must configure a number of fabric policies before the fabric can be managed by the Nexus Dashboard Orchestrator. From the APIC's perspective, this is something you do just like you would if you were connecting a bare-metal host, where you would configure domains, AEPs, policy groups, and interface selectors; you must configure the same options for connecting the spine switch interfaces to the inter-fabric network for all the fabrics that will be part of the same Multi-Site domain.

Step 3 Specify the VLAN pool.

The first thing you configure is the VLAN pool. We use Layer 3 sub-interfaces tagging traffic with VLAN-4 to connect the spine switches to the inter-fabric network.

a) In the left navigation tree, browse to **Pools** > **VLAN**.
b) Right-click the **VLAN** category and choose **Create VLAN Pool**.

In the **Create VLAN Pool** window, specify the following:

- For the **Name** field, specify the name for the VLAN pool, for example `mfabric`.
- For **Allocation Mode**, specify `Static Allocation`.
- And for the **Encap Blocks**, specify just the single VLAN 4. You can specify a single VLAN by entering the same number in both **Range** fields.

Step 4 Configure Attachable Access Entity Profiles (AEP).

a) In the left navigation tree, browse to **Global Policies** > **Attachable Access Entity Profiles**.

b) Right-click the **Attachable Access Entity Profiles** category and choose **Create Attachable Access Entity Profiles**.

In the **Create Attachable Access Entity Profiles** window, specify the name for the AEP, for example `mfabric-aep`.

c) Click **Next** and **Submit**

No additional changes, such as interfaces, are required.

Step 5 Configure the external routed domain.

The domain you configure is what you will select from the Nexus Dashboard Orchestrator when adding this fabric.

a) In the left navigation tree, browse to **Physical and External Domains** > **External Routed Domains**.

b) Right-click the **External Routed Domains** category and choose **Create Layer 3 Domain**.

In the **Create Layer 3 Domain** window, specify the following:

- For the **Name** field, specify the name the domain, for example `mfabric-l3`.
- For **Associated Attachable Entity Profile**, select the AEP you created in Step 4.
- For the **VLAN Pool**, select the VLAN pool you created in Step 3.

c) Click **Submit**.

No additional changes, such as security domains, are required.

What to do next

After you have configured the global access policies, you must still add interfaces policies as described in Configuring Fabric Access Interface Policies, on page 5.

Configuring Fabric Access Interface Policies

This section describes the fabric access interface configurations that must be done for the Nexus Dashboard Orchestrator on each APIC fabric.

Before you begin

You must have configured the global fabric access policies, such as VLAN Pool, AEP, and domain, in the fabric's APIC, as described in Configuring Fabric Access Global Policies, on page 4.

Procedure

Step 1 Log in directly to the fabric's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

In addition to the VLAN, AEP, and domain you have configured in previous section, you must also create the interface policies for the fabric's spine switch interfaces that connect to the Inter-Fabric Network.

Step 3 Configure a spine policy group.

a) In the left navigation tree, browse to **Interface Policies** > **Policy Groups** > **Spine Policy Groups**.

This is similar to how you would add a bare-metal server, except instead of a Leaf Policy Group, you are creating a Spine Policy Group.

b) Right-click the **Spine Policy Groups** category and choose **Create Spine Access Port Policy Group**.

In the **Create Spine Access Port Policy Group** window, specify the following:

- For the **Name** field, specify the name for the policy group, for example `Spine1-PolGrp`.
- For the **Link Level Policy** field, specify the link policy used between your spine switch and the ISN.
- For **CDP Policy**, choose whether you want to enable CDP.
- For the **Attached Entity Profile**, select the AEP you have configured in previous section, for example `mfabric-aep`.

c) Click **Submit**.

No additional changes, such as security domains, are required.

Step 4 Configure a spine profile.

a) In the left navigation tree, browse to **Interface Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Interface Profile**.

In the **Create Spine Interface Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1-ISN`.
- For **Interface Selectors**, click the **+** sign to add the port on the spine switch that connects to the ISN. Then in the **Create Spine Access Port Selector** window, provide the following:
  - For the **Name** field, specify the name for the port selector, for example `Spine1-ISN`.
  - For the **Interface IDs**, specify the switch port that connects to the ISN, for example `5/32`.
  - For the **Interface Policy Group**, choose the policy group you created in the previous step, for example `Spine1-PolGrp`.

  Then click **OK** to save the port selector.

c) Click **Submit** to save the spine interface profile.

Step 5 Configure a spine switch selector policy.

a) In the left navigation tree, browse to **Switch Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Profile**.

In the **Create Spine Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1`.
- For **Spine Selectors**, click the **+** to add the spine and provide the following:
  - For the **Name** field, specify the name for the selector, for example `Spine1`.
  - For the **Blocks** field, specify the spine node, for example `201`.

c) Click **Update** to save the selector.
d) Click **Next** to proceed to the next screen.
e) Select the interface profile you have created in the previous step

For example `Spine1-ISN`.

f) Click **Finish** to save the spine profile.