



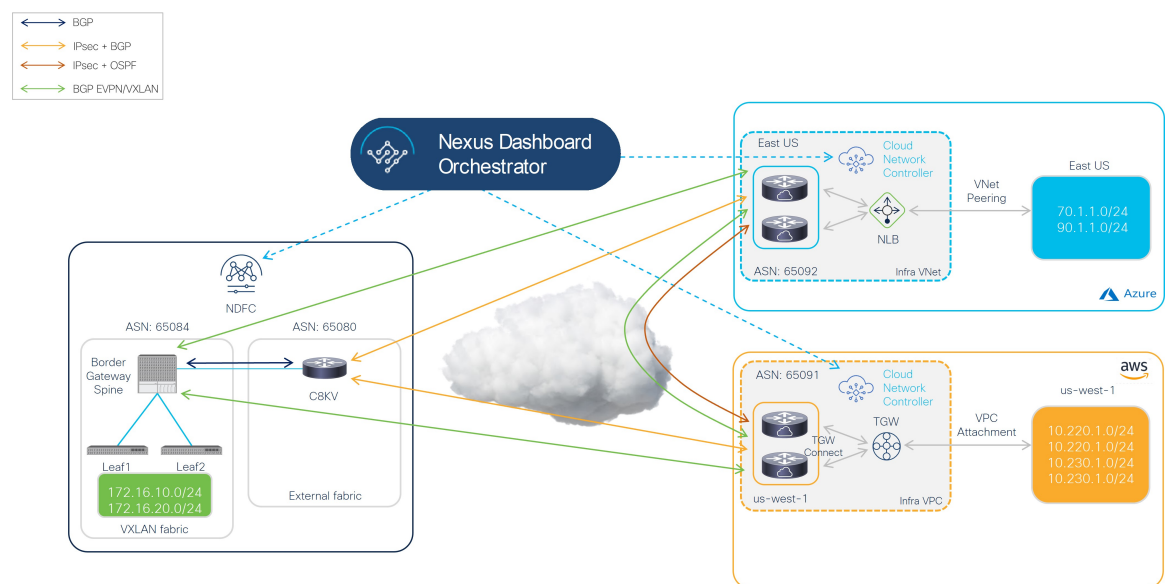
Setting Up the Infra Configuration for Hybrid Cloud and Multi-Cloud Connectivity Deployment

- [Example Topology of Infra Configuration for Hybrid Cloud and Multi-Cloud Connectivity Deployment, on page 1](#)
- [Set Up the On-Premises NDFC Fabrics, on page 2](#)
- [Deploy Cloud Network Controller on Cloud Sites, on page 19](#)
- [Onboard the NDFC and Cloud Sites into ND and NDO, on page 32](#)
- [Complete Site-to-Site Connectivity Between NDFC and Cloud Sites, on page 39](#)

Example Topology of Infra Configuration for Hybrid Cloud and Multi-Cloud Connectivity Deployment

The following figure shows one of the supported topologies that could be used for the infra configuration for hybrid cloud and multi-cloud connectivity deployment.

Figure 1:



The procedures in this document will use this topology as a specific use case, which is based on [Option 1 in Supported Topologies with IPsec \(Multi-Cloud\)](#), and will describe how to configure the hybrid cloud connectivity options specifically for this topology use case.

In this deployment procedure, you will configure multi-cloud connectivity with IPsec, where you will make certain configurations in each of these hybrid cloud connectivity areas. The overall configuration steps are as follows:

- Installing NDFC

For more detailed information, see:

- [Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide](#), Release 12.1.2 or later
- [Cisco NDFC-Fabric Controller Configuration Guide](#), Release 12.1.2 or later
- [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#), Release 12.1.2 or later

- Initial setup:

- Setting up the on-premises NDFC fabric
- Installing Cisco Cloud Network Controller
- Setting up cloud sites
- Installing NDO
- Setting up hybrid cloud connectivity using NDO

- Deploying the tenant and schema:

- Use case 1: Stretched VRF (intra-VRF)
- Use case 2: Route leaking (inter-VRF)

Set Up the On-Premises NDFC Fabrics

In this section, you will set up the two on-premises NDFC fabrics:

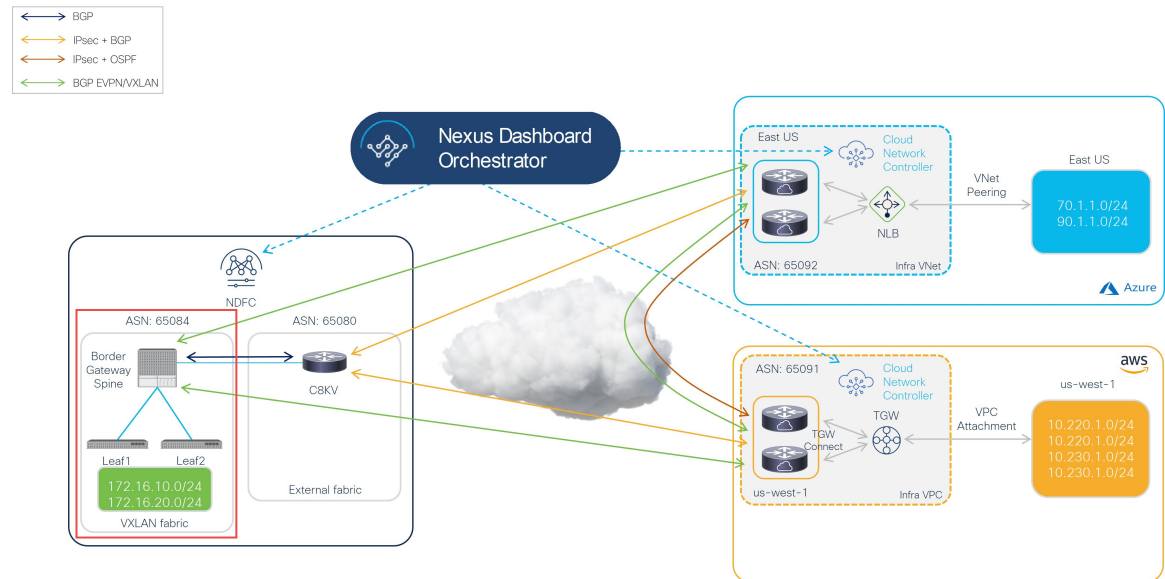
- NDFC VXLAN fabric
- NDFC external fabric

Complete the procedures in the following sections to set up the two on-premises NDFC fabrics.

Create an NDFC VXLAN Fabric

In this procedure, you will be configuring the part of the example topology highlighted below.

Figure 2:



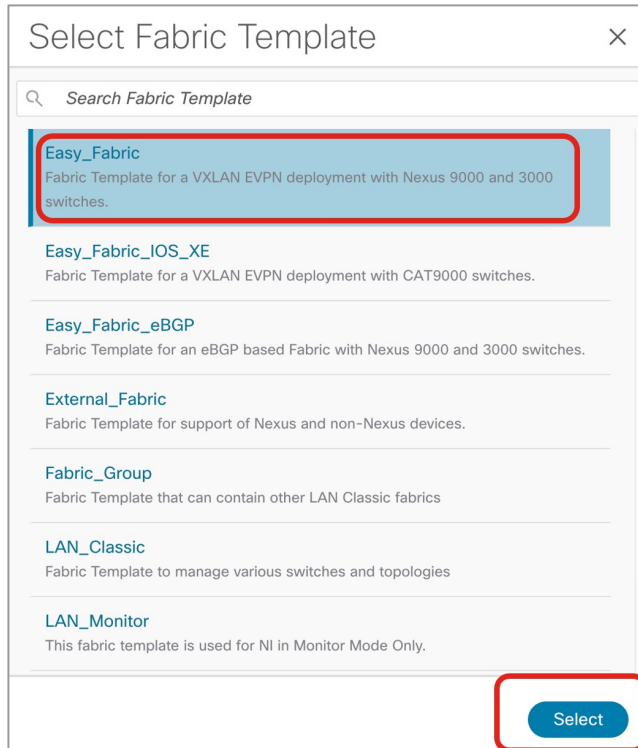
The VXLAN fabric must contain one or more Border Gateway (BGW) devices, which are used to build VXLAN Multi-Site connectivity between on-premises fabrics and the cloud sites.

Complete the procedures in the following sections to configure an NDFC VXLAN fabric.

Create an NDFC VXLAN Fabric

- Step 1** Log into the Nexus Dashboard where you have NDFC installed.
- Step 2** Log into your NDFC account.
- Step 3** Navigate to **LAN > Fabrics**.
The **LAN Fabrics** window appears.
- Step 4** Click **Actions > Create Fabric**.
The **Create Fabric** window appears.
- Step 5** Begin the process of creating an NDFC VXLAN fabric using the `Easy_Fabric` template.
 - a) In the **Fabric Name** field, enter a name for the NDFC VXLAN fabric.
 - b) In the **Pick a Template** area, click **Choose Template**.
The **Select Fabric Template** window appears.
 - c) Locate and click the `Easy_Fabric` template.
 - d) Click **Select**.

Figure 3:



Step 6 Complete the necessary general VXLAN fabric parameter configurations.

The following parameter tabs in the `Easy_Fabric` template must be completed, but they do not contain parameters that are specific to this hybrid cloud topology use case:

- **General Parameters**
- **Replication**
- **VPC**
- **Protocols**

Complete the VXLAN fabric configurations in those parameter tabs as you normally would. See [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#), Release 12.1.2 or later, for more information.

For example, using the information in the example topology, you would enter `65084` in the **BGP ASN** field in the **General Parameters** page.

Figure 4:

The screenshot displays the configuration interface for creating an NDFC VXLAN Fabric. The 'Fabric Name' is set to 'sydney'. The 'Pick Template' section shows 'Easy_Fabric' selected. The 'General Parameters' tab is active, showing the following settings:

- BGP ASN***: 65084. Note: 1-4294967295 | 1-65535(,0-65535) It is a good practice to have a unique ASN for each Fabric.
- Enable IPv6 Underlay**: . Note: If not enabled, IPv4 underlay is used.
- Enable IPv6 Link-Local Address**: . Note: If not enabled, Spine-Leaf interfaces will use global IPv6 addresses.
- Fabric Interface Numbering***: p2p. Note: Numbered(Point-to-Point) or Unnumbered.
- Underlay Subnet IP Mask***: 30. Note: Mask for Underlay Subnet IP Range.
- Underlay Subnet IPv6 Mask**: Select an Option. Note: Mask for Underlay Subnet IPv6 Range.
- Underlay Routing Protocol***: ospf. Note: Used for Spine-Leaf Connectivity.
- Route-Reflectors***: 2. Note: Number of spines acting as Route-Reflectors.

Step 7

In the **Advanced** parameter tab, make the necessary configuration specifically for this hybrid cloud topology use case.

- Locate the **Anycast Border Gateway advertise-pip** field and check the box to enable this option. This advertises the Anycast Border Gateway PIP as VTEP.

This is required when Layer 3 only connectivity (for example, no Layer 2 extension) is established across sites, which is always the case for hybrid cloud and multi-cloud deployments.

- Complete the remaining configurations in the **Advanced** parameter tab as you normally would.

Figure 5:

The screenshot shows the configuration interface for an NDFC VXLAN Fabric. The 'Advanced' tab is active, and the 'Resources' parameter tab is selected. The 'Anycast Border Gateway advertise-pip' checkbox is highlighted with a red box.

General Parameters:

- Fabric Name: sydney
- Pick Template: Easy_Fabric >
- General Parameters | Replication | VPC | Protocols | **Advanced** | Resources | Manageability | Bootstrap

Resources Parameters:

- VRF Template*: Default_VRF_Universal (Default Overlay VRF Template For Leafs)
- Network Template*: Default_Network_Universal (Default Overlay Network Template For Leafs)
- VRF Extension Template*: Default_VRF_Extension_Universal (Default Overlay VRF Template For Borders)
- Network Extension Template*: Default_Network_Extension_Universal (Default Overlay Network Template For Borders)
- Overlay Mode: config-profile (VRF/Network configuration using config-profile or CLI, default is config-profile)
- Site Id: 82 (For EVPN Multi-Site Support (Min:1, Max: 281474976710655). Defaults to Fabric ASN)
- Intra Fabric Interface MTU*: 9216 (Min:576, Max:9216). Must be an even number
- Layer 2 Host Interface MTU*: 9216 (Min:1500, Max:9216). Must be an even number

Advanced Parameters:

- VTEP HoldDown Time: 180 (NVE Source Interface HoldDown Time (Min:1, Max:1500) in seconds)
- Brownfield Overlay Network Name Format: Auto_Net_VNI\$\$VNI\$\$_VLAN\$\$VLAN_ID\$\$ (Generated network name should be < 64 characters)
- Enable CDP for Bootstrapped Switch: (Enable CDP on management interface)
- Enable VXLAN OAM: (Enable the Next Generation (NG) OAM feature for all switches in the fabric to aid in trouble-shooting VXLAN EVPN fabrics)
- Enable Tenant DHCP:
- Enable NX-API: (Enable NX-API on port 443)
- Enable NX-API on HTTP port: (Enable NX-API on port 80)
- Enable Policy-Based Routing (PBR):
- Enable Strict Config Compliance: (Enable bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config)
- Enable AAA IP Authorization: (Enable only, when IP Authorization is enabled in the AAA Server)
- Enable NDFC as Trap Host: (Configure NDFC as a receiver for SNMP traps)
- Anycast Border Gateway advertise-pip: (To advertise Anycast Border Gateway PIP as VTEP. Effective on MSD fabric 'Recalculate Config')

Step 8

Click the **Resources** parameter tab and enter the necessary values in this page.

- Enter the appropriate information in the following fields specifically for this hybrid cloud use case:
 - **Underlay Routing Loopback IP Range:** This is typically the loopback0 IP address range.
 - **Underlay VTEP Loopback IP Range:** This is typically the loopback1 IP address range.
 - **Underlay RP Loopback IP Range:** The Anycast or Phantom Rendezvous Point (RP) IP address range.
 - **Underlay Subnet IP Range:** The address range to assign numbered and peer link SVI IP addresses.
 - **VRF Lite Subnet IP Range:** The address range to assign P2P inter-fabric connections.
- Complete the remaining configurations in the **Resources** parameter tab as you normally would.

Figure 6:

The screenshot shows the configuration interface for a VXLAN fabric. The 'Fabric Name' is 'sydney' and the template is 'Easy_Fabric'. The 'Resources' tab is selected, showing various IP address ranges for underlay and overlay. The following fields are highlighted with red boxes:

- Underlay Routing Loopback IP Range*: 20.2.0.0/22
- Underlay VTEP Loopback IP Range*: 20.3.0.0/22
- Underlay RP Loopback IP Range*: 20.254.254.0/24
- Underlay Subnet IP Range*: 20.4.0.0/16
- Layer 2 VXLAN VNI Range*: 30000-49000
- Layer 3 VXLAN VNI Range*: 50000-59000
- Network VLAN Range*: 2300-2999
- VRF VLAN Range*: 2000-2299
- Subinterface Dot1q Range*: 2-511
- VRF Lite Deployment*: Manual
- Auto Deploy Both:
- VRF Lite Subnet IP Range*: 20.33.0.0/16
- VRF Lite Subnet Mask*: 30
- Service Network VLAN Range*: 3000-3199
- Route Map Sequence Number Range*: 1-65534

Step 9 Complete the necessary general VXLAN fabric parameter configurations in the **Manageability** and **Bootstrap** parameter tabs.

The configurations in the **Manageability** and **Bootstrap** parameter tabs might need to be completed, but they do not contain parameters that are specific to this hybrid cloud topology use case.

Step 10 Click the **Configuration Backup** parameter tab and check the box in the **Hourly Fabric Backup** field to enable that feature.

Complete the remaining configurations in the **Configuration Backup** parameter tab as you normally would.

Step 11 Click **Save** when you have completed the necessary configurations in the **Create Fabric** window for the VXLAN fabric.

You are returned to the **LAN Fabrics** window, with the VXLAN fabric that you just created displayed.

What to do next

Add the switches to the VXLAN fabric and set the necessary role for the switches using the procedures provided in [Add Switches to the VXLAN Fabric, on page 7](#).

Add Switches to the VXLAN Fabric

In this procedure, you will add the switches to the VXLAN fabric and set the necessary role for the switches.

Before you begin

Create an NDFC VXLAN fabric using the procedures provided in [Create an NDFC VXLAN Fabric, on page 3](#).

Step 1 In the **LAN Fabrics** window, click the VXLAN fabric that you just created.

The **Overview** window for this fabric appears.

Note The following steps describe how to manually enter the necessary information to allow NDFC to discover switches. You could also use the Power On Auto Provisioning (POAP) feature in NDFC instead, which is useful if you do not already have certain parameters, such as the management IP address, default route, and start up configurations, already configured on the switches that need to be discovered. POAP automates the process of installing configuration files on devices that are deployed on the network for the first time and allows devices to be brought up without performing any manual configuration. See [Inband POAP Management in External Fabrics and LAN Classic Fabrics](#) and [Zero-Touch Provisioning of VXLAN Fabrics using Inband POAP with NDFC](#) for more information on POAP.

Step 2 Click **Actions > Add Switches**.

The **Add Switches** window appears.

Step 3 Add the necessary information to discover the switches.

- Fill in the necessary information in this page to discover the switches, including the Seed IP, username, and password.
- Determine if you want to preserve the existing configuration on the switches:
 - If this is a brownfield deployment where you want to keep the existing configurations on the switches, check the **Preserve Config** checkbox to preserve those existing configurations.
 - If this is a greenfield deployment, uncheck the **Preserve Config** checkbox to clean up the configurations on the switches.

Step 4 Click **Discover Switches**.

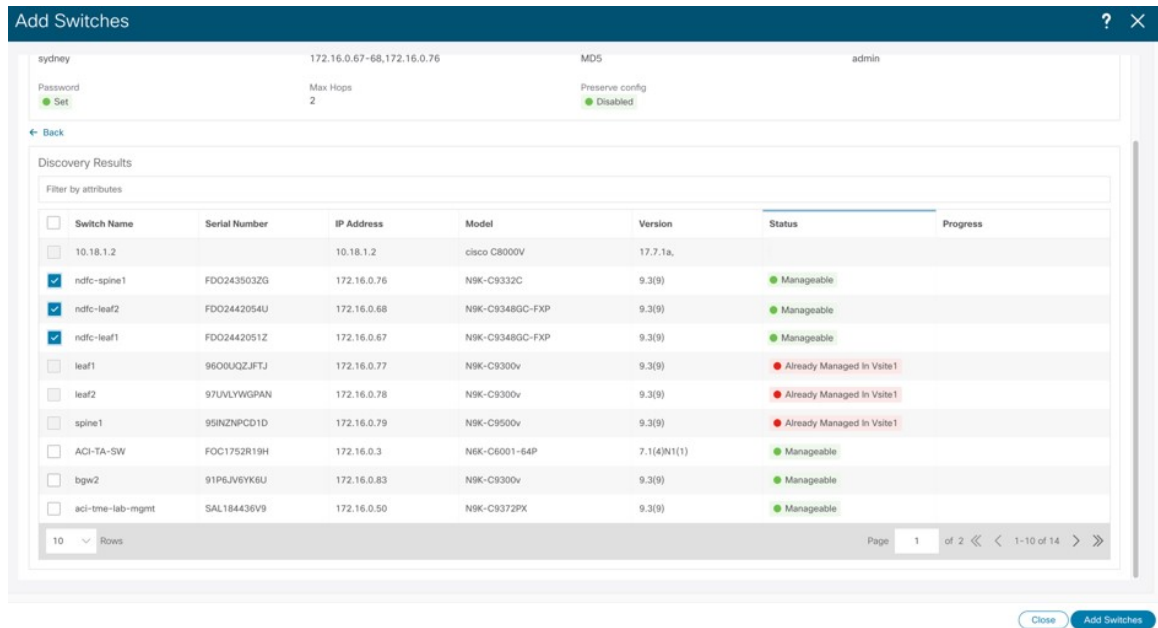
Click **Confirm** in the confirmation popup window that appears.

Step 5 Once the switches have been discovered, add the switches to the NDFC VXLAN fabric.

In the **Discovery Results** area, choose the appropriate switches (click the box next to each of the appropriate switches).

As an example, the figure below shows two leaf switches and one spine switch being added to the fabric.

Figure 7:



Step 6 Click **Add Switches**.

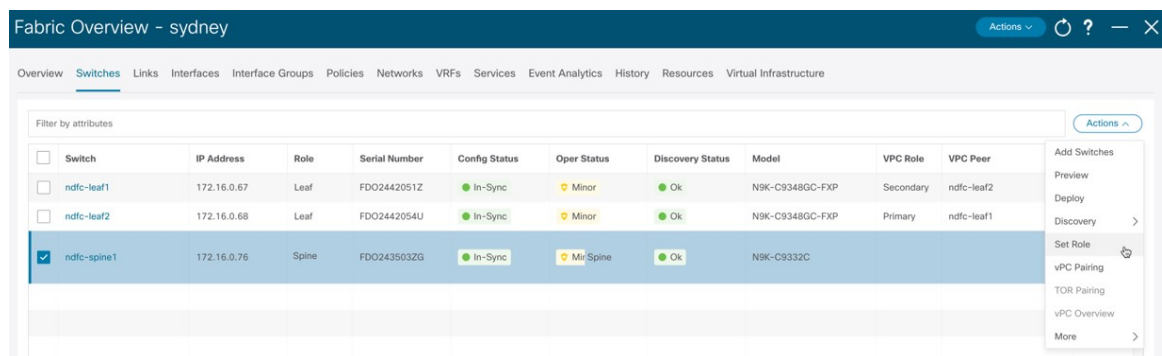
Note If the **Preserve Config** option is checked, the switches will go through a reboot after being added to the NDFC VXLAN fabric.

Step 7 Set the role for the appropriate switch to `Border Gateway Spine`.

In these example procedures, one spine switch plays the dual role of spine switch and border gateway spine switch, so we will be changing the role of the spine switch to border gateway spine switch in these example procedures. However, in your environment, you might have two separate switches, one with the role of spine switch and the other with the role of border gateway.

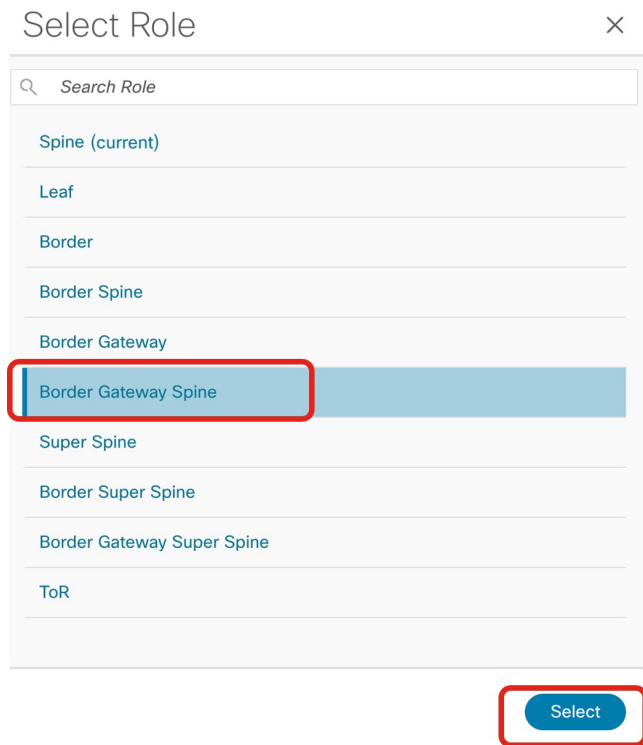
- Click the **Switches** tab in the NDFC VXLAN fabric overview window. The switches that have been added to this fabric are displayed.
- Click the box next to the spine switch to choose that switch, then click **Actions > Set Role**.

Figure 8:



- Locate and select the `Border Gateway Spine` role in the **Select Role** list, then click **Select**.

Figure 9:



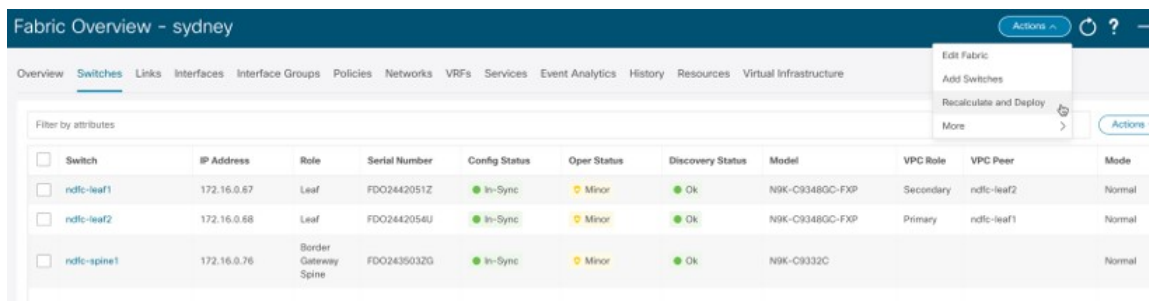
Step 8 Navigate to **LAN > Fabrics** and select the NDFC VXLAN fabric that you created.

The **Overview** page for this NDFC VXLAN fabric appears.

Step 9 Click the **Switches** tab to verify that the switches that you just added appear correctly.

Step 10 Click **Actions > Recalculate and Deploy**.

Figure 10:



As described earlier, for these procedures, one spine switch plays the dual role of spine switch and border gateway spine switch, so we changed the role of the spine switch to border gateway spine switch in these example procedures, as shown below. In these example procedures, a vPC pair has also been configured already for the two leaf switches, as shown in the figure below. For more information on configuring a vPC pair, see the [Cisco NDFC-Fabric Controller Configuration Guide](#), release 12.1.2e or later.

Figure 11:

Switch	IP Address	Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer	Mode
<input type="checkbox"/> ndfc-leaf1	172.16.0.67	Leaf	FDO2442051Z	In-Sync	Minor	Ok	N9K-C9348GC-FXP	Secondary	ndfc-leaf2	Normal
<input type="checkbox"/> ndfc-leaf2	172.16.0.68	Leaf	FDO2442054U	In-Sync	Minor	Ok	N9K-C9348GC-FXP	Primary	ndfc-leaf1	Normal
<input type="checkbox"/> ndfc-spine1	172.16.0.76	Border Gateway Spine	FDO243503ZG	In-Sync	Minor	Ok	N9K-C9332C			Normal

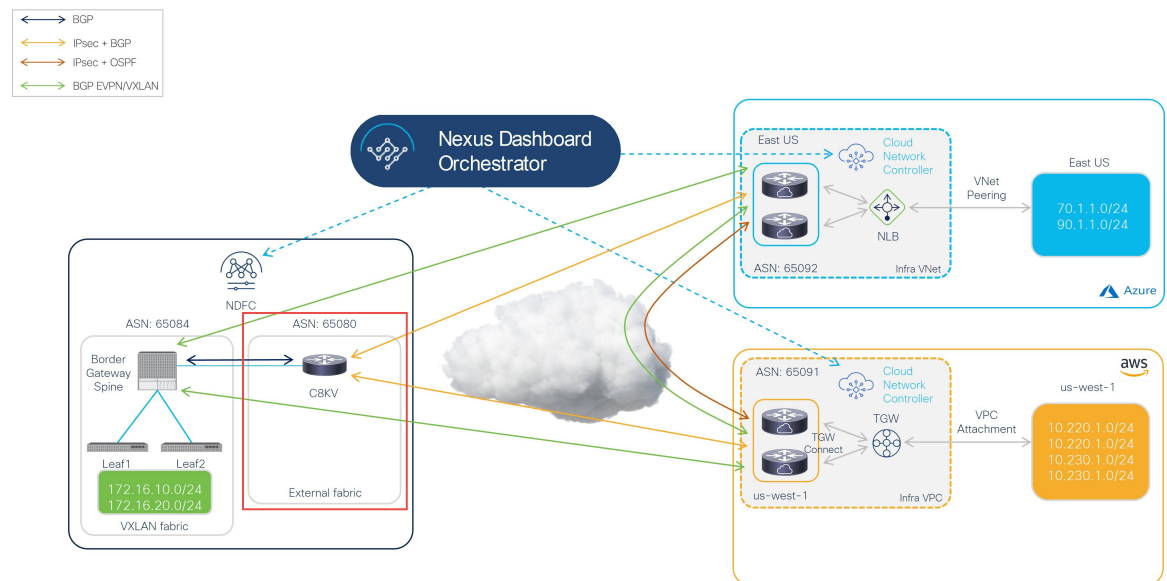
What to do next

Configure an NDFC external fabric using the procedures provided in [Configure an NDFC External Fabric, on page 11](#).

Configure an NDFC External Fabric

In this procedure, you will be configuring the part of the example topology highlighted below. In the example figure below and throughout the use case procedures, a Cisco Catalyst 8000V is used as the IPsec device in the external fabric, but there could be many different types of devices in the external fabric, as long as they support IPsec and can be managed by NDFC (for example, ASR 1000 and Catalyst 8000V).

Figure 12:



An NDFC-managed external fabric contains one or more IPsec devices. The IPsec devices have connectivity to cloud networks either via the internet (public) or by a private connection, such as Direct Connect (AWS) or ExpressRoute (Azure). If public internet is used to connect to the cloud sites, IPsec tunnels are established between on-premises IPsec devices and Catalyst 8000Vs in the cloud sites.

Complete the procedures in the following sections to configure an NDFC external fabric.

Create an NDFC External Fabric

Before you begin

Complete the procedures provided in [Create an NDFC VXLAN Fabric, on page 3](#) before proceeding with these procedures.

Step 1 Log into your NDFC account, if you are not logged in already.

Step 2 Navigate to **LAN > Fabrics**.

Step 3 Click **Actions > Create Fabric**.

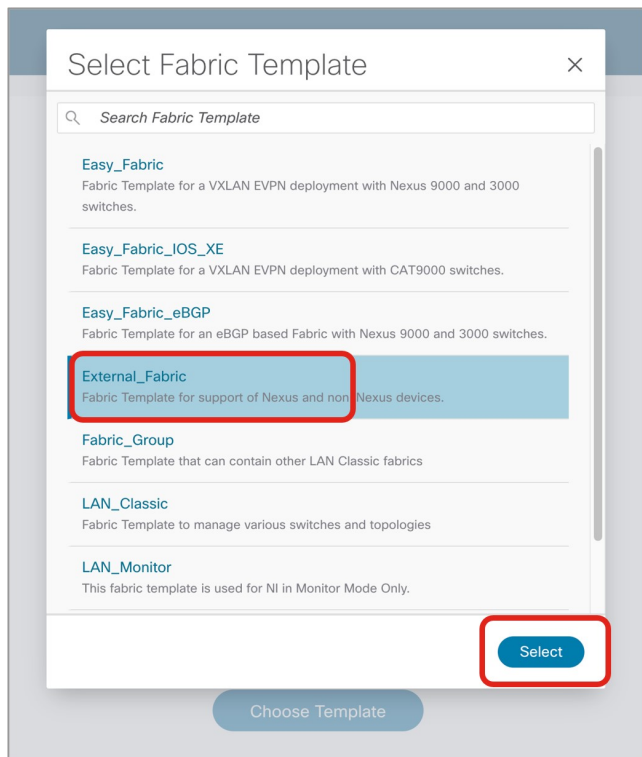
The **Create Fabric** window appears.

Step 4 Begin the process of creating an external fabric using the `External_Fabric` template.

The `External_Fabric` template is used to build traditional LAN fabrics using Nexus as well as non-Nexus devices, such as Catalyst 8000Vs.

- a) In the **Fabric Name** field, enter a name for the external fabric.
- b) In the **Pick a Template** area, click **Choose Template**.
The **Select Fabric Template** window appears.
- c) Locate and click the `External_Fabric` template.
- d) Click **Select**.

Figure 13:



Step 5 In the **General Parameters** tab, make the necessary configuration specifically for this hybrid cloud topology use case.

- In the **BGP ASN** field, define the BGP ASN.

For example, using the information in the example topology, you would enter 65080 in the **BGP ASN** field for this use case.

- Determine if you want the external fabric to be monitored or not:
 - If the on-premises IPsec device is going to be managed by NDFC, uncheck the box next to the **Fabric Monitor Mode** field to unselect this option.
 - If the on-premises IPsec device is not going to be managed by NDFC (such as a non-Cisco, third-party firewall), check the box next to the **Fabric Monitor Mode** field if the fabric is going to be monitored only.

Figure 14:

Step 6 Complete the necessary general external fabric parameter configurations.

The following parameter tabs in the `External_Fabric` template must be completed, but they do not contain parameters that are specific to this hybrid cloud topology use case:

- **Advanced**
- **Resources**
- **Configuration Backup**
- **Bootstrap**
- **Flow Monitor**

For example, in the **Configuration Backup** parameter tab, you might check the box in the **Hourly Fabric Backup** field to enable that feature.

See [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#), Release 12.1.2 or later, for more information.

- Step 7** Click **Save** when you have completed the necessary configurations in the **Create Fabric** window for the external fabric. You are returned to the **LAN Fabrics** window, with the external fabric that you just created displayed.
-

What to do next

Add the on-premises Cisco Catalyst 8000V to the external fabric and set the necessary role using the procedures provided in [Add the On-Premises Cisco Catalyst 8000V to the External Fabric, on page 14](#).

Add the On-Premises Cisco Catalyst 8000V to the External Fabric

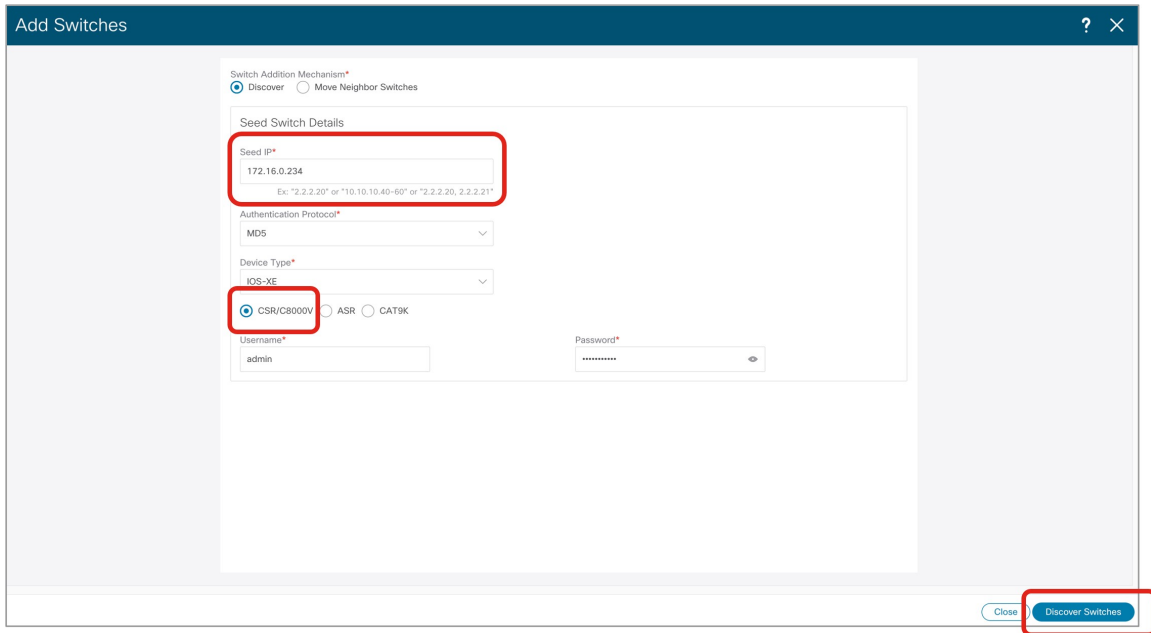
Follow these procedures to add the on-premises Cisco Catalyst 8000V to the external fabric and set the necessary role for the Cisco Catalyst 8000V.

Before you begin

Create the NDFC external fabric using the procedures provided in [Create an NDFC External Fabric, on page 12](#)

- Step 1** In the **LAN Fabrics** window, click the external fabric that you just created.
The **Overview** window for this fabric appears.
- Step 2** Click **Actions > Add Switches**.
The **Add Switches** window appears.
- Step 3** Add the necessary information to discover the Cisco Catalyst 8000V, then click **Discover Switches**.
- Enter the necessary information in the **Seed IP** field for the Cisco Catalyst 8000V.
 - In the **Device Type** field, choose `IOS-XE`.
 - Choose the `CSR/C8000V` option underneath the **Device Type** field when it appears.

Figure 15:



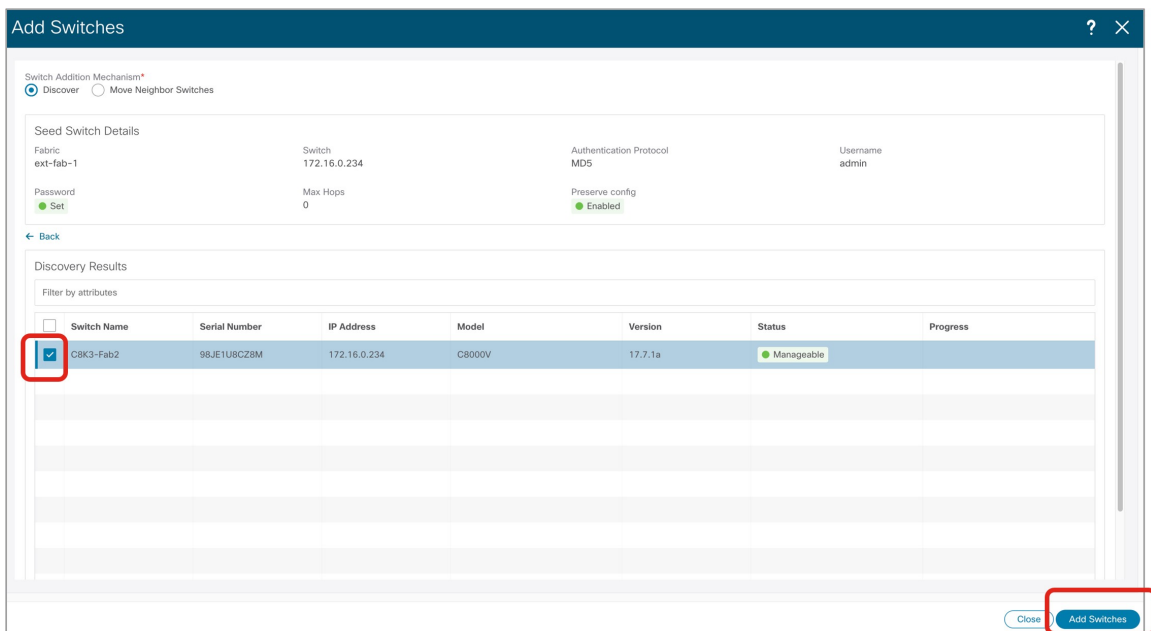
Step 4 Click **Discover Switches**.

Click **Confirm** in the confirmation pop-up window that appears.

Step 5 Once the Cisco Catalyst 8000V has been discovered, add the Cisco Catalyst 8000V to the external fabric.

In the **Discovery Results** area, choose the Cisco Catalyst 8000V (click the box next to the Cisco Catalyst 8000V) and click **Add Switches**.

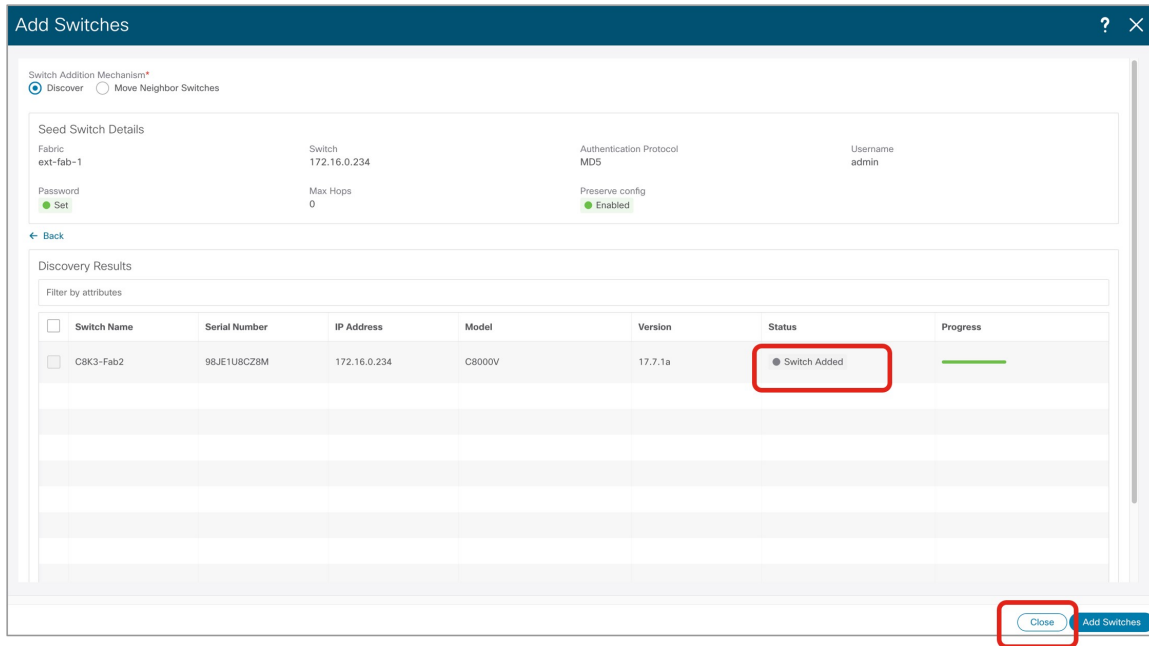
Figure 16:



Add the On-Premises Cisco Catalyst 8000V to the External Fabric

The status will change to **Switch Added**. Click **Close** to close out of this window.

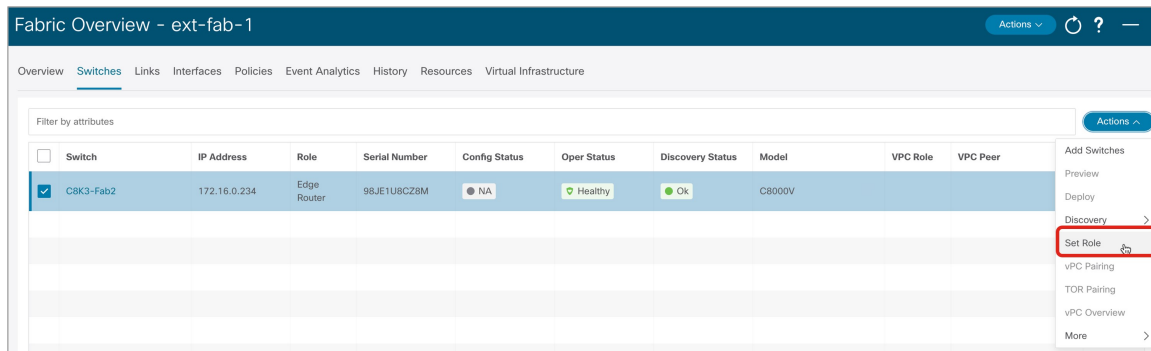
Figure 17:



Step 6 Set the role for the Cisco Catalyst 8000V to `Core Router`.

a) Click the box next to the Cisco Catalyst 8000V to choose that router, then click **Actions > Set Role**.

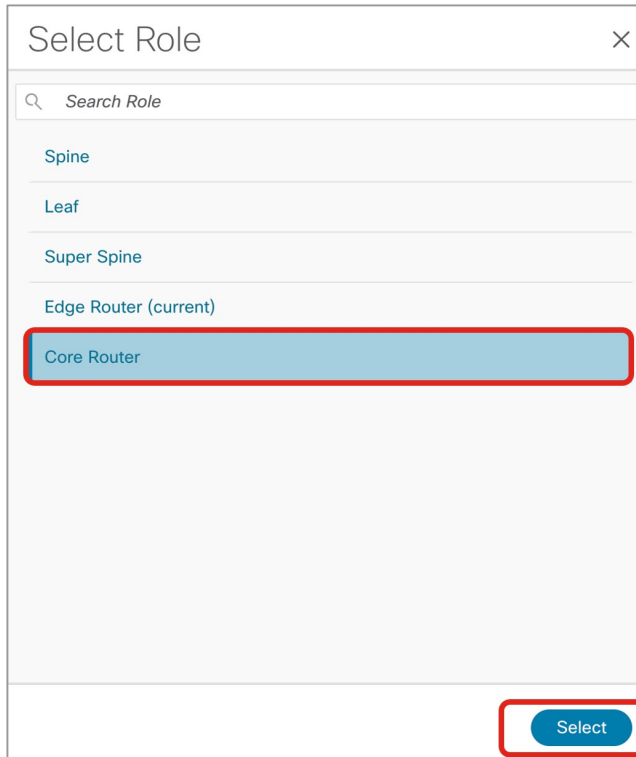
Figure 18:



b) Locate and select the `Core Router` role in the **Select Role** list, then click **Select**.

All the Catalyst 8000Vs should be set to the `Core Router` role so that NDFC automatically enables BGP protocol.

Figure 19:



Step 7 Navigate to **LAN > Fabrics** and select the external fabric that you created.

The **Overview** page for this external fabric appears.

Step 8 Click the **Switches** tab to verify that the Cisco Catalyst 8000V that you just added appears correctly.

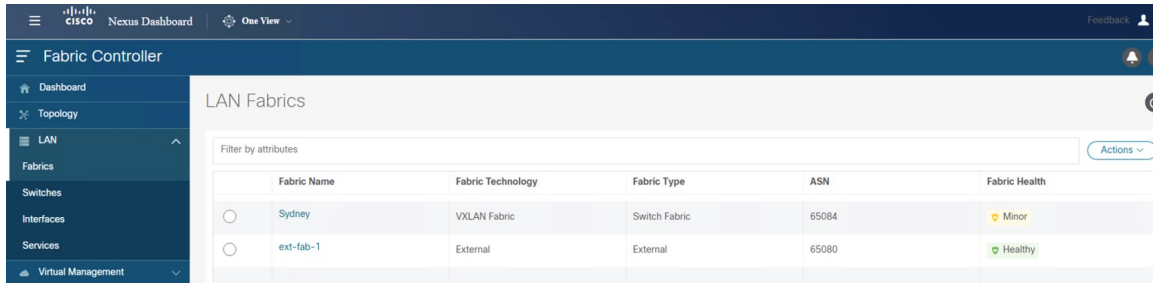
Figure 20:

Switch	IP Address	Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer	Mode
<input type="checkbox"/> CSK3-Fab2	172.16.0.234	Core Router	98JE1UBCZ8M	NA	Healthy	Ok	C8000V			Normal

Step 9 Click **Actions > Recalculate and Deploy**.

At this point in the process, the VXLAN and external fabrics are configured in NDFC, as shown when you navigate to **LAN > Fabrics**.

Figure 21:



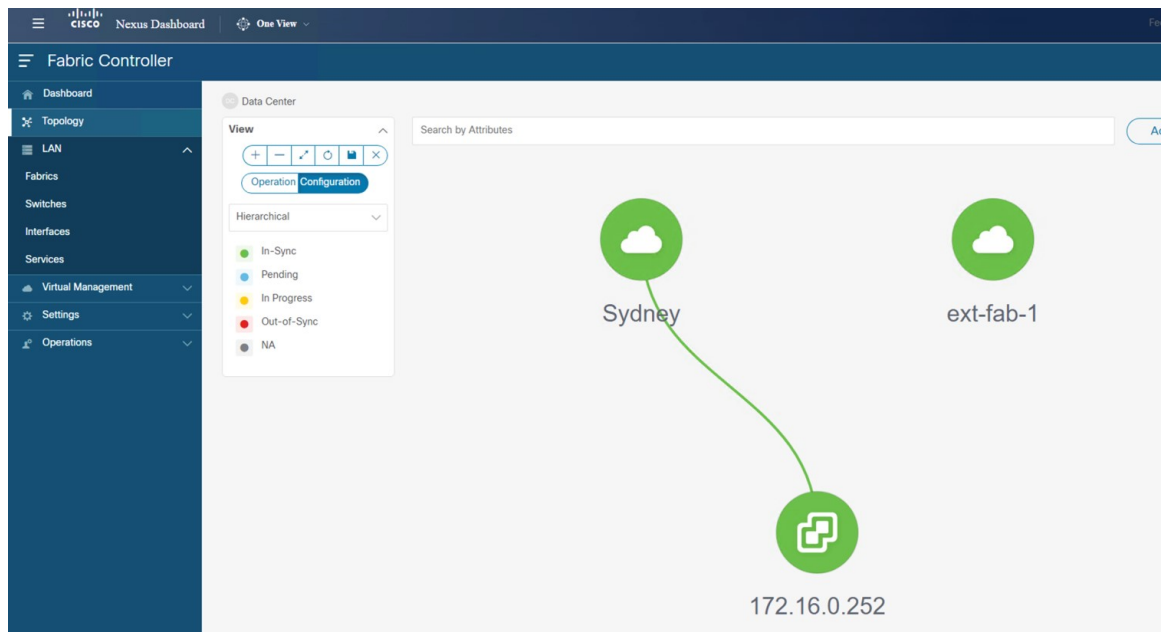
The screenshot shows the Cisco Fabric Controller interface with the 'LAN Fabrics' section selected. A table lists the fabric configurations:

Fabric Name	Fabric Technology	Fabric Type	ASN	Fabric Health
Sydney	VXLAN Fabric	Switch Fabric	65084	Minor
ext-fab-1	External	External	65080	Healthy

You can also use the **Topology** view to determine the following configurations at this point in the process:

- That there is no connectivity yet between the VXLAN and external fabrics:

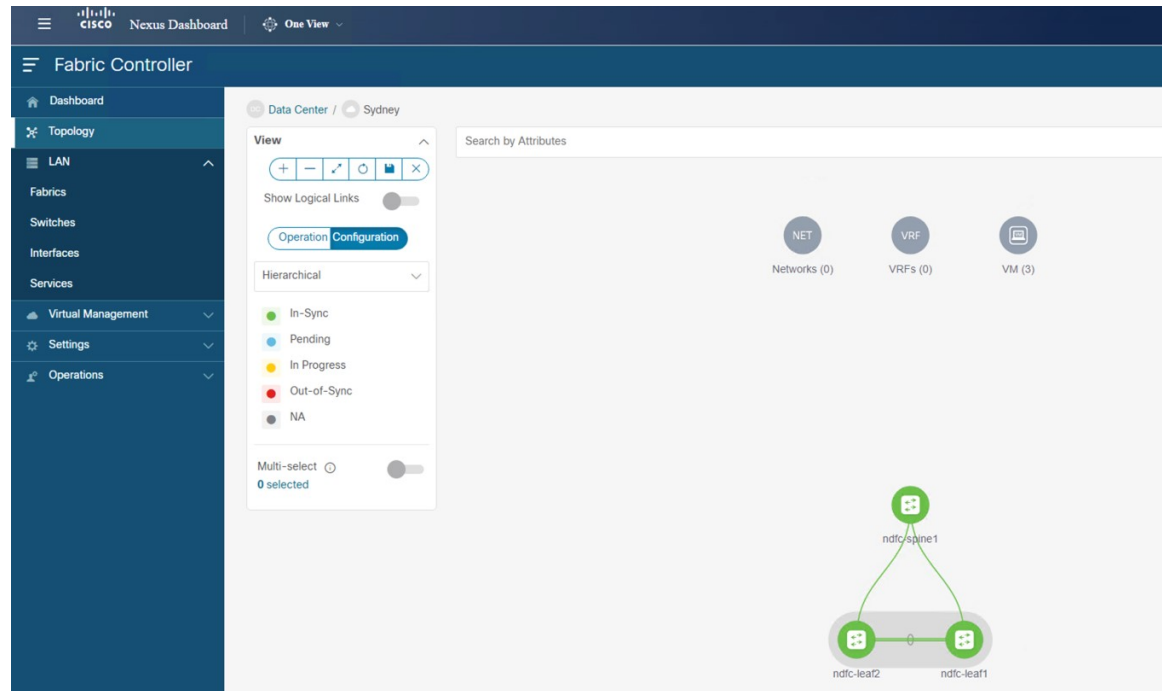
Figure 22:



This NDFC has the VMM Visualizer feature enabled, so the vCenter icon with an IP address of 172.16.0.252 is displayed in the topology view. For more information on the VMM feature, see the [Virtual Infrastructure Manager](#) chapter in the *Cisco NDFC-Fabric Controller Configuration Guide*.

- That there are no networks or VRFs created yet in the VXLAN fabric:

Figure 23:



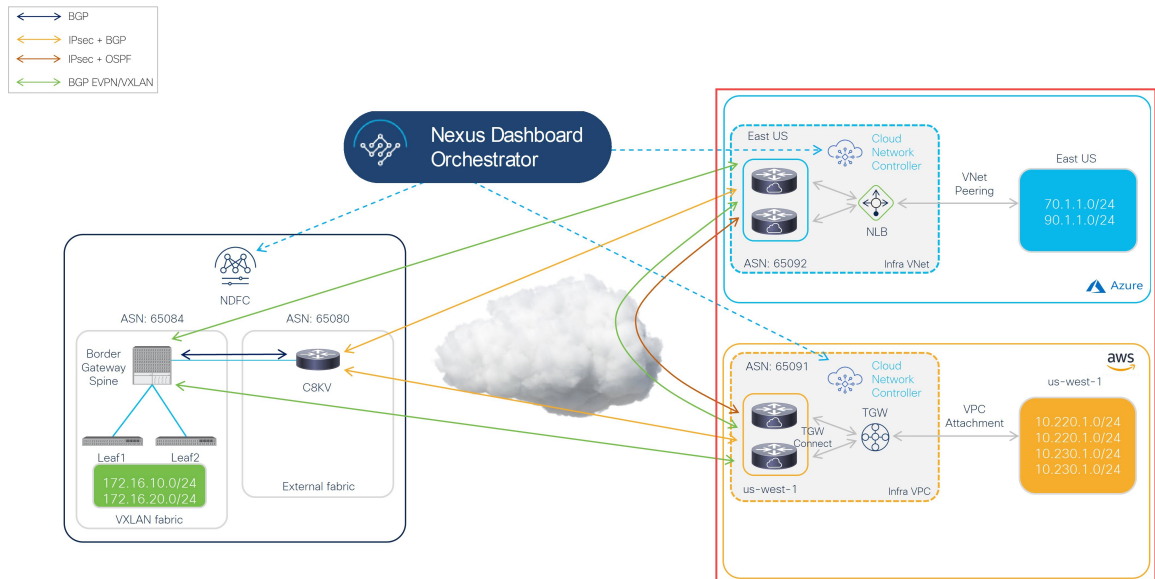
What to do next

Deploy the Cloud Network Controller on the cloud sites using the procedures provided in [Deploy Cloud Network Controller on Cloud Sites, on page 19](#).

Deploy Cloud Network Controller on Cloud Sites

In this section, you will be configuring the part of the example topology highlighted below.

Figure 24:



Based on the example hybrid cloud topology, these procedures assume that we will be setting up two cloud sites through the Cloud Network Controller (AWS and Azure cloud sites). We will therefore refer to the following documents throughout these procedures:

- [Cisco Cloud Network Controller for AWS Installation Guide](#), Release 25.1(x) or later
- [Cisco Cloud Network Controller for AWS User Guide](#), Release 25.1(x) or later
- [Cisco Cloud Network Controller for Azure Installation Guide](#), Release 25.1(x) or later
- [Cisco Cloud Network Controller for Azure User Guide](#), Release 25.1(x) or later

Complete the procedures in the following sections to deploy the Cloud Network Controller on the cloud sites.

Deploy the Cloud Network Controller on the AWS Cloud Site

Follow the procedures in these sections to deploy the Cloud Network Controller on the AWS cloud site.

Configure the Necessary Parameters in Advanced Settings for AWS

In this section, you will make the necessary configurations for the AWS cloud site in **Advanced Settings** area in the **Cloud Network Controller Setup** page specifically for this example hybrid cloud topology.

Use the procedures provided in the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the [Cisco Cloud Network Controller for AWS Installation Guide](#), but note that there are two areas in the **Cloud Network Controller Setup** page that you will have to configure specifically for this example hybrid cloud topology:

- **Contract-based routing:** Cloud Network Controller supports two types of modes:
 - Contract-based routing
 - Route map-based routing

Contract-based routing means that a contract between the EPGs will drive the routing between VRFs, but this type of contract-based routing is not available through NDFC, so for this specific example hybrid cloud topology, you will turn off contract-based routing and will use route map-based routing instead. For more information, see the "Routing Policies" and "Configuring the Global Inter-VRF Route Leak Policy" sections in the [Cisco Cloud Network Controller for AWS User Guide](#), Release 25.1(x) or later.


- **Cloud Network Controller Access Privilege:** By default, the Cloud Network Controller has Routing & Security access privilege, which means that the Cloud Network Controller can automate not only networking, it can also automate and configure security groups on the cloud. If the Cloud Network Controller automates and configures the security groups, it also has to configure the EPGs and contracts; however, EPGs and contracts are not applicable to NDFC end users who only need routing automation. To integrate well with NDO and NDFC, you should set the **Cloud Network Controller Access Privilege** option to **Routing Only**.

Step 1 Log into your Cisco Cloud Network Controller for AWS.

Step 2 Begin the process of setting up the first cloud site, the AWS cloud site, for this example hybrid cloud topology.

The first few chapters in the [Cisco Cloud Network Controller for AWS Installation Guide](#), Release 25.1(x) or later, contain generic information that is not specific to this hybrid cloud topology use case, so complete the procedures in these chapters in that document, then return here:

- Overview
- Preparing for Installing the Cisco Cloud Network Controller
- Configuring the Cloud Formation Template Information for the Cisco Cloud Network Controller

Step 3 In the Cisco Cloud Network Controller GUI, click the Intent icon () and select **Cloud Network Controller Setup**. The **Let's Configure the Basics** page appears.

Step 4 Locate the **Advanced Settings** area and click **Edit Configuration**.

Step 5 In the **Advanced Settings** page, set the following configurations:

- **Contract Based Routing:** Verify that the box is unchecked (that this feature is not enabled). This turns off contract-based routing and uses route map-based routing instead
- **Cloud Network Controller Access Privilege:** Choose the **Routing Only** option.

Step 6 Click **Save and Continue**.

You are returned to the **Let's Configure the Basics** page.

What to do next

Follow the procedures provided in [Configure the Necessary Parameters in Region Management for AWS](#), on page 22.

Configure the Necessary Parameters in Region Management for AWS

In this section, you will make the necessary configurations for the AWS cloud site in the **Region Management** area in the **Cloud Network Controller Setup** page specifically for this example hybrid cloud topology.

Before you begin

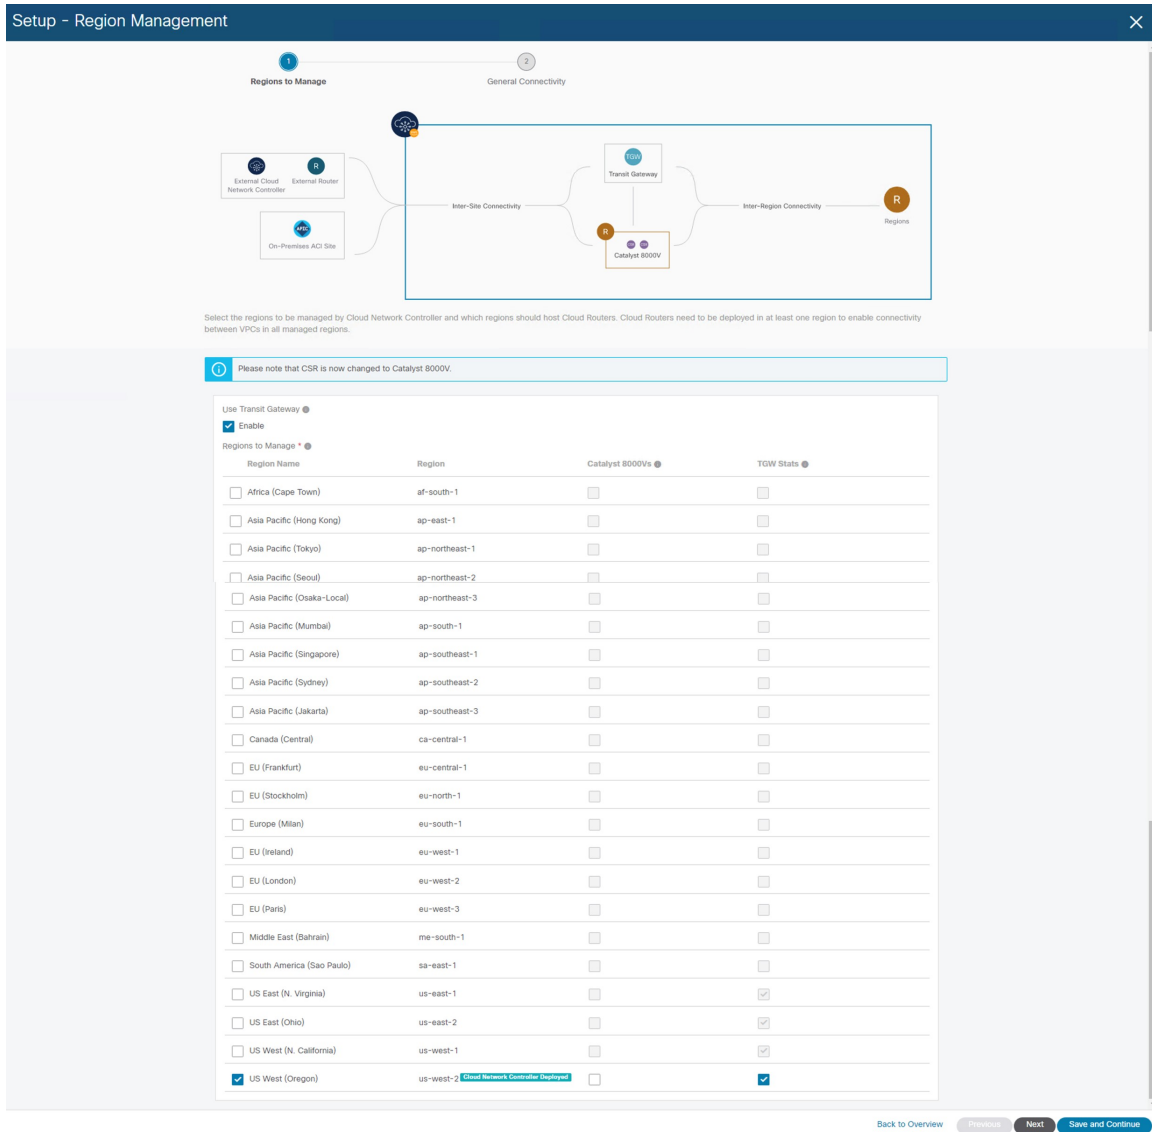
Complete the procedures provided in [Configure the Necessary Parameters in Advanced Settings for AWS, on page 20](#).

-
- Step 1** Locate the **Region Management** area and click the appropriate button.
- Click **Begin** if this is your first time setting up the Cloud Network Controller, or **Edit Configuration** if you had already configured region management in this Cloud Network Controller previously.
- Step 2** Enable AWS Transit Gateway.
- You normally use Transit Gateway to avoid using VPN tunnels for connectivity within a region and across the regions where TGW peering is supported. For more information, see the [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway or AWS Transit Gateway Connect](#) document.
- Specifically for this example hybrid cloud topology use case, in the **Use Transit Gateway** area, click the checkbox next to **Enable** to use AWS Transit Gateway. This will allow you to add a hub network later in these procedures, which is necessary to enable TGW Connect.
- Step 3** In the **Regions to Manage** area, verify that the Cisco Cloud Network Controller home region is selected.
- The region that you selected when you first deployed the Cisco Cloud Network Controller in AWS is the home region and should be selected already in this page. This is the region where the Cisco Cloud Network Controller is deployed (the region that will be managed by Cisco Cloud Network Controller), and will be indicated with the text `Cloud Network Controller deployed` in the Region column.
- Step 4** Select additional regions if you want the Cisco Cloud Network Controller to manage additional regions, and to possibly deploy Cisco Catalyst 8000Vs to have inter-VPC communication and Hybrid-Cloud, Hybrid Multi-Cloud, or Multi-Cloud connectivity on those other regions.
- The Cisco Catalyst 8000V can provide hybrid cloud and multi-cloud connectivity for up to four regions, including the home region where Cisco Cloud Network Controller is deployed.
- Step 5** To deploy cloud routers locally to a region, click to place a check mark in the **Catalyst 8000Vs** check box for that region.
- You must have at least one region with Catalyst 8000Vs deployed. However, if you choose multiple regions in this page, you do not have to have Catalyst 8000Vs in every region that you choose.
- Step 6** If you want to use AWS Transit Gateway statistics, check the box in the **TGW Stats** column for one or more regions.
- Checking the check box enables collection of AWS Transit Gateway traffic statistics for infra tenants for the specified regions.
- Note** You also need to create flow logs in order to collect AWS Transit Gateway statistics. See the section "Enabling VPC Flow Logs" in the chapter "Cisco Cloud APIC Statistics" of the *Cisco Cloud APIC for AWS User Guide*, release 25.1(x) or later.

Specifically for this example hybrid cloud topology use case:

- Place a check mark in the check boxes next to the **US East (N. Virginia)** and **US West (N. California)** regions (the **us-east-1** and **us-west-1** regions).
- Place a check mark in the check boxes in the **Catalyst 8000Vs** and **TGW Stats** columns for the Cisco Cloud Network Controller home region.

Figure 25:



Step 7 When you have selected all the appropriate regions, click **Next** at the bottom of the page. The **General Connectivity** page appears.

Step 8 Make the necessary configurations in the **General Connectivity** page. See the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the *Cisco Cloud Network Controller for AWS Installation Guide*, Release 25.1(x) or later, for more information.

Specifically for this example hybrid cloud topology use case, add a hub network using the procedures in the following steps.

In Cisco Cloud Network Controller, a collection of two or more AWS Transit Gateways is called a **hub network**. A hub network provides network isolation for VRFs. A group of VRFs can be attached to a hub network to isolate the group of VRFs from other VRFs that are attached to other hub networks. A hub network creates at least two AWS Transit Gateways for each region.

Step 9 In the **Hub Network** area, click **Add Hub Network**.

The **Add Hub Network** window appears.

Step 10 In the **Name** field, enter a name for the hub network.

Step 11 In the **BGP Autonomous System Number** field, enter a zero for AWS to choose a number, or enter a value between 64512 and 65534, inclusive, for each hub network, and then click the check mark next to the field.

For example, using the information in the example hybrid cloud topology, you would enter 65091 in this field.

Step 12 In the **TGW Connect** field, click the checkbox next to **Enable** to enable the AWS Transit Gateway Connect feature.

You will enable the AWS Transit Gateway Connect feature for this example hybrid cloud topology use case. See [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway or AWS Transit Gateway Connect](#) for more information.

Step 13 In the **CIDRs** area, click **Add CIDR**.

This will be the AWS Transit Gateway Connect CIDR block, which will be used as the connect peer IP address (the GRE outer peer IP address) on the Transit Gateway side.

- a) In the **Region** field, click **Select Region** and select the appropriate region.
- b) In the **CIDR** field, enter the CIDR block that will be used as the connect peer IP address on the Transit Gateway side.

Figure 26:

- c) Click the checkmark to accept these values for this CIDR block.
- d) For every managed region that will be using the AWS Transit Gateway Connect feature, repeat these steps to add CIDR blocks to be used for each of those managed regions.

Figure 27:

Add Hub Network ✕

Name *
hub1

BGP Autonomous System Number *
65091

TGW Connect
 Enable

⚠ Changing the use of TGW Connect will cause temporary traffic loss.

CIDR

Region *	CIDR *
US West (Oregon)	176.16.11.0/24

+ Add CIDR

TGW Route Table Association Labels ●

Name *

+ Add TGW Route Table Association Label

Add

Step 14 Complete the remaining configurations as you normally would.

- Complete the remaining configurations in the **General Connectivity** page as you normally would, then click **Save and Continue**.
- Complete the necessary configurations in the **Smart Licensing** page as you normally would.

See the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the [Cisco Cloud Network Controller for AWS Installation Guide](#), Release 25.1(x) or later, for more information.

At this point in the process, you have completed the basic configurations for the first cloud site for the Cisco Cloud Network Controller (in this example hybrid cloud topology, the AWS cloud site). Proceed with the following steps to complete the basic configurations for the second cloud site for the Cisco Cloud Network Controller (in this example hybrid cloud topology, the Azure cloud site).

Step 15 Configure Direct Connect for AWS, if necessary.

Configure Direct Connect if you want private connections for the connectivity for the Catalyst 8000V routers to the cloud networks. For information on configuring Direct Connect for AWS, see the [Cisco Cloud Network Controller for AWS User Guide](#), release 25.1(x) or later.

What to do next

Deploy the Cloud Network Controller on the second cloud site (the Azure cloud site) using the procedures provided in [Deploy the Cloud Network Controller on the Azure Cloud Site, on page 26](#).

Deploy the Cloud Network Controller on the Azure Cloud Site

Follow the procedures in these sections to deploy the Cloud Network Controller on the Azure cloud site.

Configure the Necessary Parameters in Advanced Settings for Azure

In this section, you will make the necessary configurations for the Azure cloud site in **Advanced Settings** area in the **Cloud Network Controller Setup** page specifically for this example hybrid cloud topology.

Make the same configurations for the Azure cloud site as you did for the AWS cloud site.

Use the procedures provided in the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the [Cisco Cloud Network Controller for Azure Installation Guide](#), but note that there are two areas in the **Cloud Network Controller Setup** page that you will have to configure specifically for this example hybrid cloud topology:

- **Contract-based routing:** Cloud Network Controller supports two types of modes:
 - Contract-based routing
 - Route map-based routing

Contract-based routing means that a contract between the EPGs will drive the routing between VRFs, but this type of contract-based routing is not available through NDFC, so for this specific example hybrid cloud topology, you will turn off contract-based routing and will use route map-based routing instead. For more information, see the "Routing Policies" and "Configuring the Global Inter-VRF Route Leak Policy" sections in the [Cisco Cloud Network Controller for AWS User Guide](#), Release 25.1(x) or later.

- **Cloud Network Controller Access Privilege:** By default, the Cloud Network Controller has Routing & Security access privilege, which means that the Cloud Network Controller can automate not only networking, it can also automate and configure security groups on the cloud. If the Cloud Network Controller automates and configures the security groups, it also has to configure the EPGs and contracts; however, EPGs and contracts are not applicable to NDFC end users who only need routing automation. To integrate well with NDO and NDFC, you should set the **Cloud Network Controller Access Privilege** option to **Routing Only**.

Before you begin

Deploy the Cloud Network Controller on the first cloud site (the AWS cloud site) using the procedures provided in [Deploy the Cloud Network Controller on the AWS Cloud Site, on page 20](#).


Step 1 Log into your Cisco Cloud Network Controller for Azure.

Step 2 Begin the process of setting up the second cloud site, the Azure cloud site, for this example hybrid cloud topology.

The first few chapters in the [Cisco Cloud Network Controller for Azure Installation Guide](#), Release 25.1(x) or later, contain generic information that is not specific to this hybrid cloud topology use case, so complete the procedures in these chapters in that document, then return here:

- Overview

- Preparing for Installing the Cisco Cloud Network Controller
- Deploying the Cisco Cloud Network Controller in Azure

Step 3 In the Cisco Cloud Network Controller GUI, click the Intent icon () and select **Cloud Network Controller Setup**. The **Let's Configure the Basics** page appears.

Step 4 Locate the **Advanced Settings** area and click **Edit Configuration**.

Step 5 In the **Advanced Settings** page, set the following configurations:

- **Contract Based Routing**: Verify that the box is unchecked (that this feature is not enabled). This turns off contract-based routing and uses route map-based routing instead
- **Cloud Network Controller Access Privilege**: Choose the **Routing Only** option.

Step 6 Click **Save and Continue**.

You are returned to the **Let's Configure the Basics** page.

What to do next

Follow the procedures provided in [Configure the Necessary Parameters in Region Management for Azure, on page 27](#).

Configure the Necessary Parameters in Region Management for Azure

In this section, you will make the necessary configurations for the Azure cloud site in the **Region Management** area in the **Cloud Network Controller Setup** page specifically for this example hybrid cloud topology.

Before you begin

Follow the procedures provided in [Configure the Necessary Parameters in Advanced Settings for Azure, on page 26](#).

Step 1 Locate the **Region Management** area and click the appropriate button.

Click **Begin** if this is your first time setting up the Cloud Network Controller, or **Edit Configuration** if you had already configured region management in this Cloud Network Controller previously.

Step 2 Verify that the **Virtual Network Peering** in the **Connectivity for Internal Network** area is automatically enabled.

VNet peering at the global level is set in the **Connectivity for Internal Network** area, which enables VNet peering at the Cisco Cloud Network Controller level, deploying NLBs in all the regions with a CCR. For release 5.1(2) and later, VNet peering at the global level is enabled by default and cannot be disabled. See [Configuring VNet Peering for Cloud APIC for Azure](#) for more information.

Step 3 In the **Regions to Manage** area, verify that the Cisco Cloud Network Controller home region is selected.

The region that you selected when you first deployed the Cisco Cloud Network Controller in AWS is the home region and should be selected already in this page. This is the region where the Cisco Cloud Network Controller is deployed

(the region that will be managed by Cisco Cloud Network Controller), and will be indicated with the text `Cloud Network Controller deployed` in the Region column.

Note Because Azure VNet peering is enabled automatically, you must also check the box in the **Catalyst 8000Vs** column for the Cisco Cloud Network Controller home region, if it is not checked already.

Step 4 Select additional regions if you want the Cisco Cloud Network Controller to manage additional regions, and to possibly deploy Cisco Catalyst 8000Vs to have inter-VNet communication and Hybrid-Cloud, Hybrid Multi-Cloud, or Multi-Cloud connectivity on those other regions.

The Cisco Catalyst 8000V can provide hybrid cloud and multi-cloud connectivity for up to four regions, including the home region where Cisco Cloud Network Controller is deployed.

Step 5 To deploy cloud routers locally to a region, click to place a check mark in the **Catalyst 8000Vs** check box for that region.

You must have at least one region with Catalyst 8000Vs deployed. However, if you choose multiple regions in this page, you do not have to have Catalyst 8000Vs in every region that you choose.

Specifically for this example hybrid cloud topology use case, place a check mark in the check box in the **Catalyst 8000Vs** column for the Cisco Cloud Network Controller home region.

Figure 28:

Setup - Region Management

Please note that CSR is now changed to Catalyst 8000V.

Connectivity for Internal Network

VNet Peering

Regions to Manage *

Region Name	Region	Catalyst 8000Vs
<input type="checkbox"/> Australia Central	australiacentral	<input type="checkbox"/>
<input type="checkbox"/> Australia Central 2	australiacentral2	<input type="checkbox"/>
<input type="checkbox"/> Australia East	australiaeast	<input type="checkbox"/>
<input type="checkbox"/> Australia Southeast	australiasoutheast	<input type="checkbox"/>
<input type="checkbox"/> Brazil South	brazilsouth	<input type="checkbox"/>
<input type="checkbox"/> Canada Central	canadacentral	<input type="checkbox"/>
<input type="checkbox"/> Canada East	canadaeast	<input type="checkbox"/>
<input type="checkbox"/> Central India	centralindia	<input type="checkbox"/>
<input type="checkbox"/> Central US	centralus	<input type="checkbox"/>
<input type="checkbox"/> East Asia	eastasia	<input type="checkbox"/>
<input checked="" type="checkbox"/> East US	eastus <small>Cloud Network Controller Deployed</small>	<input checked="" type="checkbox"/>

Step 6 When you have selected all the appropriate regions, click **Next** at the bottom of the page.

The **General Connectivity** page appears.

Step 7 Make the necessary configurations in the **General Connectivity** page.

See the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the [Cisco Cloud Network Controller for Azure Installation Guide](#), Release 25.1(x) or later, for more information.

Specifically for this example hybrid cloud topology use case, make the following configurations for the Cisco Catalyst 8000Vs using the procedures in the following steps.

Step 8

Under the **General** area, in the **Subnet Pools for Cloud Routers** field, click **Add Subnet Pool for Cloud Routers** to add additional subnets for the Catalyst 8000Vs.

The first subnet pool is automatically populated (shown as `System Internal`). Addresses from this subnet pool will be used for inter-region connectivity for any additional regions that are added that need to be managed by the Cisco Cloud Network Controller. Subnet pools added in this field must be a valid IPv4 subnet with mask /24.

Add additional subnets for Catalyst 8000Vs in this step in these situations:

- If you have a Catalyst 8000V deployed in the Cisco Cloud Network Controller home region, add one additional subnet pool in addition to the `System Internal` subnet pool that is automatically generated.
- If you selected additional regions to be managed by Cisco Cloud Network Controller in the previous page:
 - Add *one* additional subnet pool for every managed region with 2-4 Catalyst 8000Vs per managed region (if you enter **2**, **3**, or **4** in the **Number of Routers Per Region** field in this page)
 - Add *two* additional subnet pools for every managed region with five or more Catalyst 8000Vs per managed region (if you enter between **5** and **8** in the **Number of Routers Per Region** field in this page)

Specifically for this example hybrid cloud topology use case, add one additional subnet pool using `10.90.1.0/24` as the subnet entry.

Figure 29:

Configure the fabric infra connectivity for the Cloud Site. The Fabric Autonomous System Number is used for BGP peering inside the configuration template used for the Cloud Routers in the Cloud Site.

Please note that CSR is now changed to Catalyst 8000V.

General

Subnet Pools for Cloud Routers

Subnet *	Regions	Created By
10.90.0.0/24		System Internal
10.90.1.0/24		User

+ Add Subnet Pool for Cloud Routers

Step 9

Under the **Catalyst 8000Vs** area, in the **BGP Autonomous System Number for C8kVs** field, enter the BGP autonomous system number (ASN) that is unique to this site.

The BGP autonomous system number can be in the range of 1 - 65534. See the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the *Cisco Cloud Network Controller for Azure Installation Guide*, Release 25.1(x) or later, for additional restrictions.

Specifically for this example hybrid cloud topology use case, you would enter 65092 in the **BGP Autonomous System Number for C8kVs** field.

Figure 30:

Setup - Region Management

Catalyst 8000Vs

BGP Autonomous System Number for C8kVs *

65092

Assign Public IP to C8KV Interface

Enable

Changing C8KV connectivity from private to public (or vice versa) may cause disruption in your network.

Number of Routers Per Region

2

Username *

cisco

Password

Confirm Password

Please ensure that the license account has licenses corresponding to the Router's throughput entered below.

Pricing Type *

BYOL

Throughput of the routers

Tier1 (up to 100M throughput)

TCP MSS *

1300

License Token

Back to Overview Previous **Next**

Step 10 Click **Next**, then complete the remaining configurations as you normally would.

- Complete the remaining configurations in the **General Connectivity** page as you normally would, then click **Save and Continue**.
- Complete the necessary configurations in the **Smart Licensing** page as you normally would.

See the "Configuring Cisco Cloud Network Controller Using the Setup Wizard" chapter in the *Cisco Cloud Network Controller for Azure Installation Guide*, Release 25.1(x) or later, for more information.

Step 11 Configure ExpressRoute for Azure, if necessary.

Configure ExpressRoute if you want private connections for the connectivity for the Catalyst 8000V routers to the cloud networks. For information on configuring ExpressRoute for Azure, see the *Cisco Cloud Network Controller for Azure User Guide*, release 25.1(x) or later.

What to do next

Onboard the NDFC-managed sites (VXLAN fabric, external fabric, and cloud sites) into Nexus Dashboard (ND) and Nexus Dashboard Orchestrator (NDO) using the procedures provided in [Onboard the NDFC and Cloud Sites into ND and NDO](#), on page 32.

Onboard the NDFC and Cloud Sites into ND and NDO

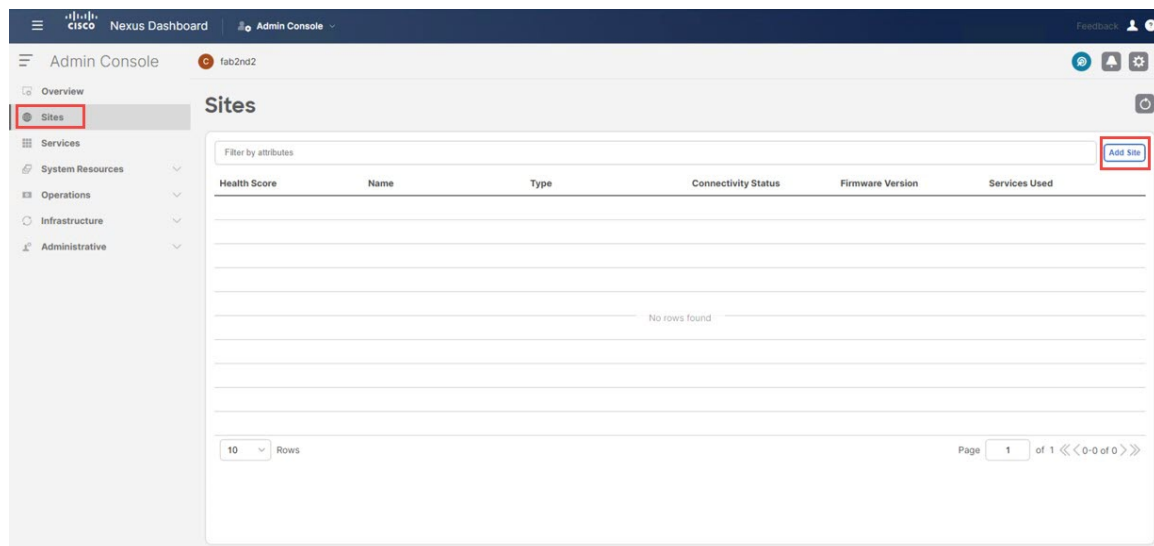
Before you begin

- Create the NDFC VXLAN fabric using the procedures provided in [Create an NDFC VXLAN Fabric, on page 3](#).
- Create the NDFC external fabric using the procedures provided in [Create an NDFC External Fabric, on page 12](#).
- Deploy the Network Cloud Controller on the first cloud site using the procedures provided in [Deploy the Cloud Network Controller on the AWS Cloud Site, on page 20](#).
- Deploy the Network Cloud Controller on the second cloud site using the procedures provided in [Deploy the Cloud Network Controller on the Azure Cloud Site, on page 26](#).

Step 1 Log into the Nexus Dashboard (ND) cluster with Nexus Dashboard Orchestrator (NDO).

Step 2 In Nexus Dashboard, click **Sites** > **Add Site**.

Figure 31:



The **Add Site** page appears.

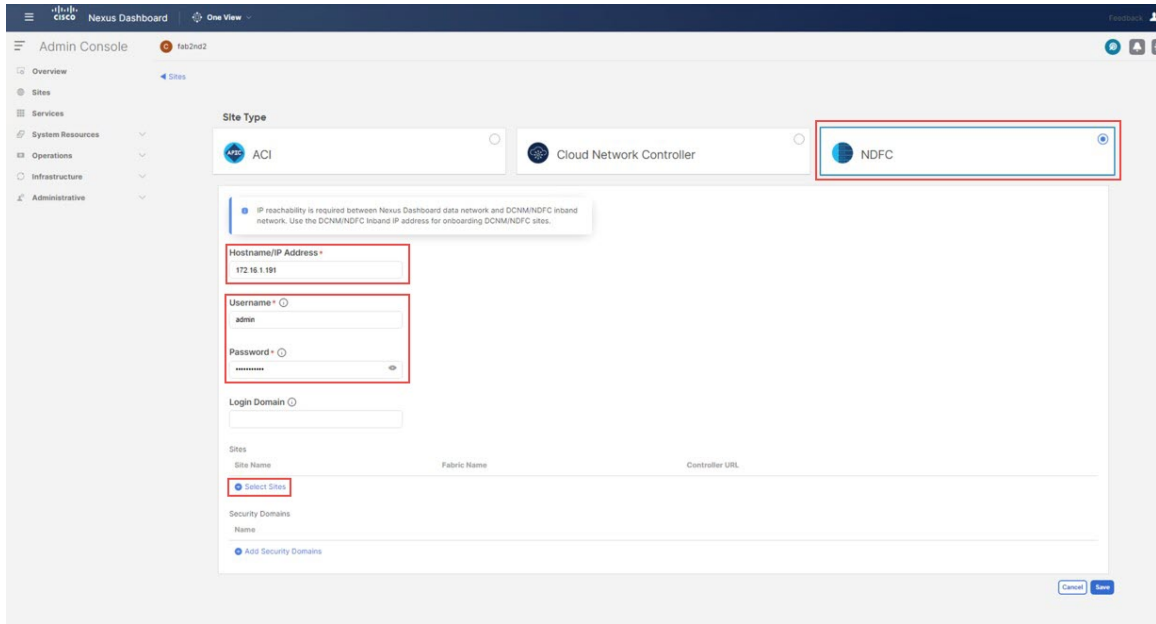
Step 3 Click the **NDFC** box in the **Add Site** page.

Step 4 Enter the necessary information to add the NDFC site.

- In the **Hostname/IP Address** field, enter the data interface IP address for your NDFC.
- In the **Username** and **Password** field, enter the username and password login information for your NDFC.

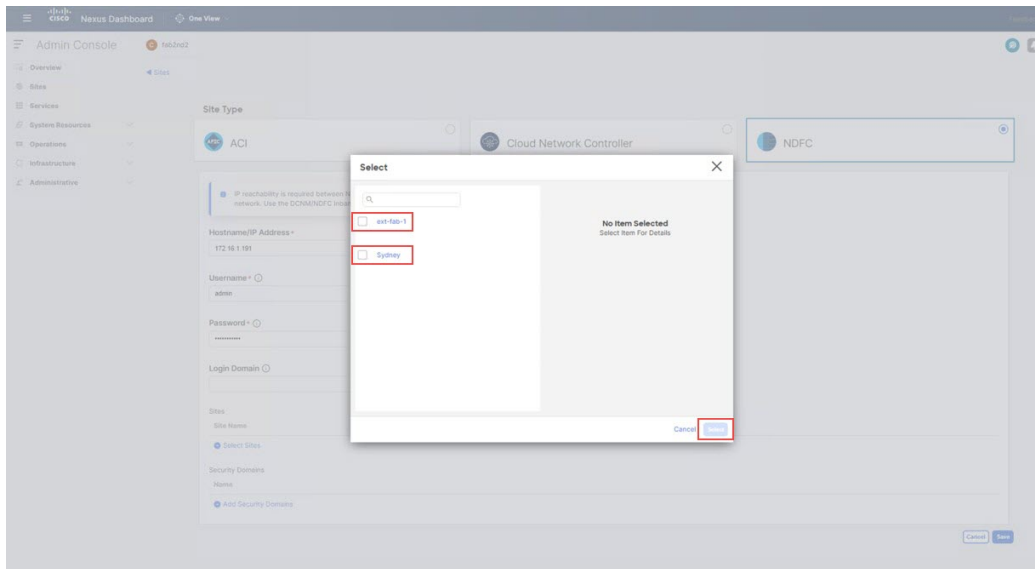
Step 5 Click **Select Sites**.

Figure 32:

**Step 6**

Click the boxes next to the two NDFC sites that you added previously (the VXLAN fabric and external fabric sites), then click **Select**.

Figure 33:



You are returned to the **Add Site** page.

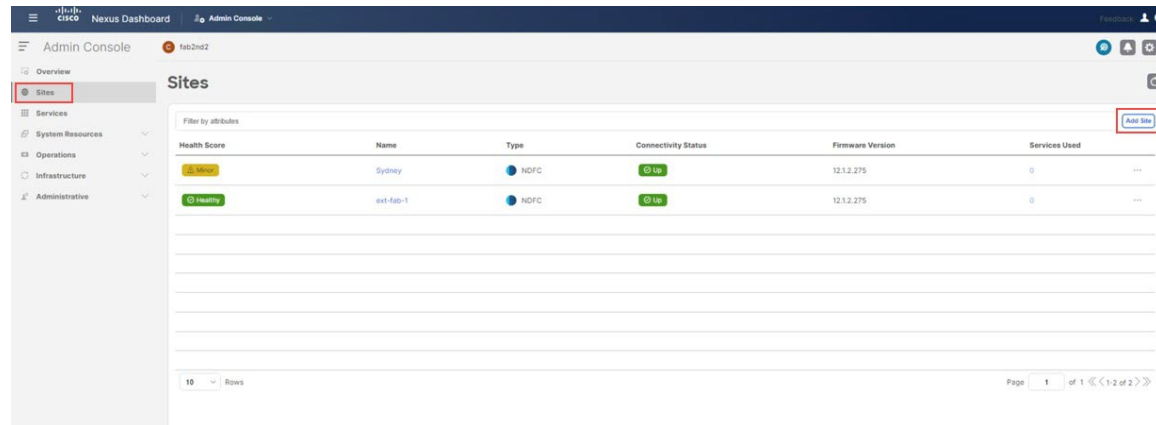
Step 7

Verify that the two NDFC sites (VXLAN fabric and external fabric sites) appear correctly in the Nexus Dashboard **Add Site** page, then click **Save**.

Step 8

In Nexus Dashboard, click **Sites** > **Add Site** again to add the first cloud site.

Figure 34:



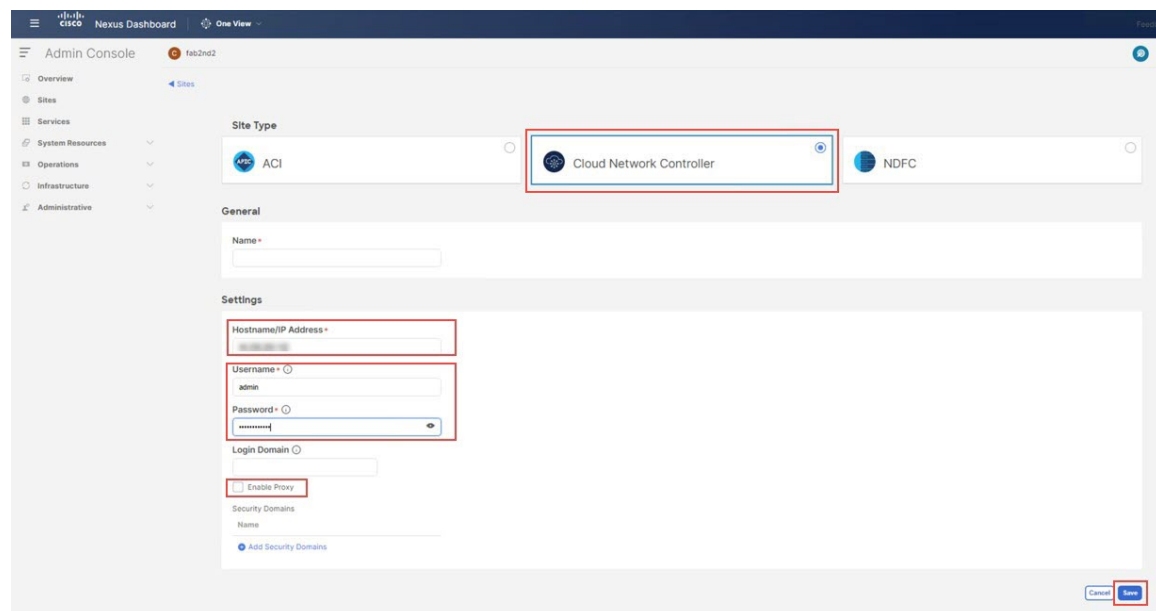
The **Add Site** page appears.

Step 9

Click the **Cloud Network Controller** box in the **Add Site** page, then enter the necessary information to add the first cloud site (the AWS site in this example topology).

- In the **Hostname/IP Address** field, enter the IP address of the Cloud Network Controller (CNC) for the first cloud site.
- In the **Username** and **Password** field, enter the username and password login information of the Cloud Network Controller (CNC) for the first cloud site.
- For Cloud Network Controller (CNC), **Enable Proxy** if the CNC is reachable via a proxy. Proxy must be already configured in your Nexus Dashboard's cluster settings. If the proxy is reachable via management network, a static management network route must also be added for the proxy IP address. For more information about proxy and route configuration, see [Nexus Dashboard User Guide](#) for your release.

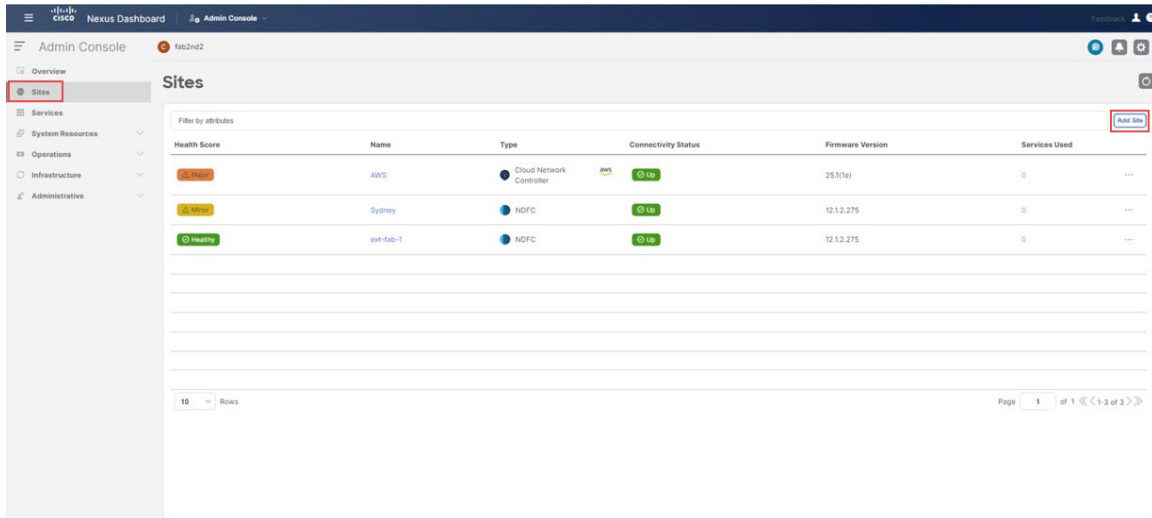
Figure 35:



Step 10 Click **Save** to add the first cloud site.

Step 11 In Nexus Dashboard, click **Sites** > **Add Site** again to add the second cloud site.

Figure 36:

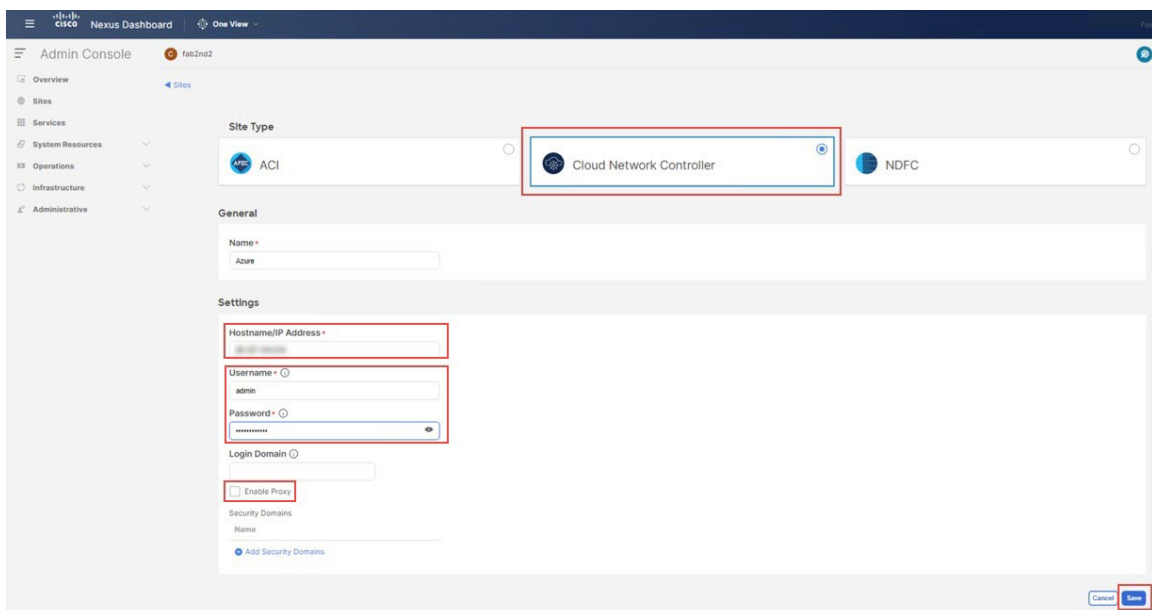


The **Add Site** page appears.

Step 12 Click the **Cloud Network Controller** box in the **Add Site** page, then enter the necessary information to add the Cloud Network Controller (CNC) for the second cloud site (the Azure site in this example topology).

Repeat the previous set of steps, this time entering the necessary information in the **Hostname/IP Address**, **Username**, and **Password** fields for the Cloud Network Controller (CNC) for the second cloud site, and clicking **Enable Proxy** if the CNC for the second cloud site is reachable via a proxy.

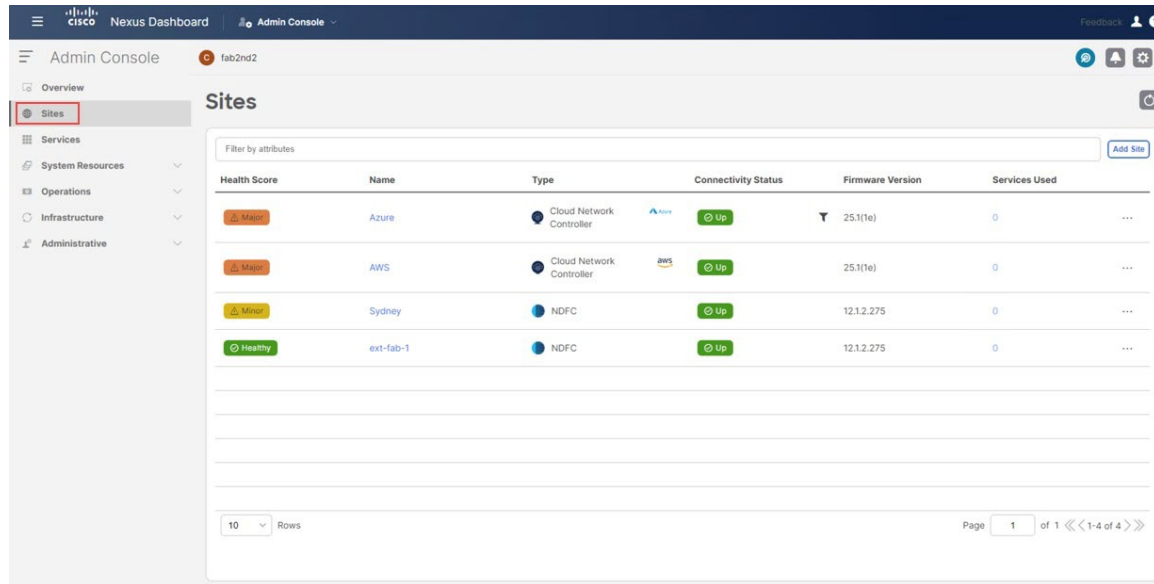
Figure 37:



Step 13 In Nexus Dashboard, click **Sites** and verify that the four sites appear correctly:

- The two sites from NDFC (the VXLAN fabric and external fabric sites)
- The cloud sites with Cloud Network Controller deployed (for this example hybrid cloud topology, the AWS and Azure cloud sites)

Figure 38:

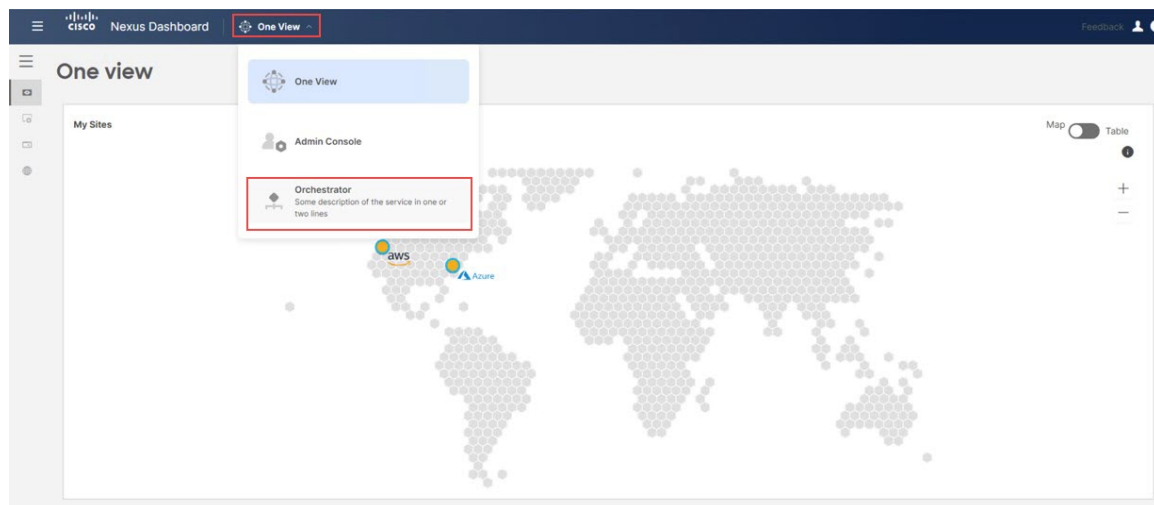


Health Score	Name	Type	Connectivity Status	Firmware Version	Services Used
Major	Azure	Cloud Network Controller	Up	25.1(1e)	0
Major	AWS	Cloud Network Controller	Up	25.1(1e)	0
Minor	Sydney	NDFC	Up	12.1.2.275	0
Healthy	ext-fab-1	NDFC	Up	12.1.2.275	0

Step 14 Access the Nexus Dashboard Orchestrator (NDO).

In Nexus Dashboard, at the top of the window, click **One View > Orchestrator**.

Figure 39:



Step 15 In NDO, click **Sites**.

The four sites that you added in ND appear but are shown in the **Unmanaged** state.

Figure 40:

The screenshot shows the 'Sites' page in the Cisco Nexus Dashboard. The table lists four sites, all with a state of 'Unmanaged'. The 'State' column is highlighted with a red box.

Controller Connectivity	Name	Type	State	Version
OK	AWS	AWS	Unmanaged	25.1(1e)
OK	Azure	Azure	Unmanaged	25.1(1e)
OK	ext-fab-1	NDFC	Unmanaged	12.1.2.275
OK	Sydney	NDFC	Unmanaged	12.1.2.275

Step 16 From NDO, manage the four sites.

Perform the following steps for each site in NDO:

- For the first site listed in NDO, under the **State** column, change the state from **Unmanaged** to **Managed**.

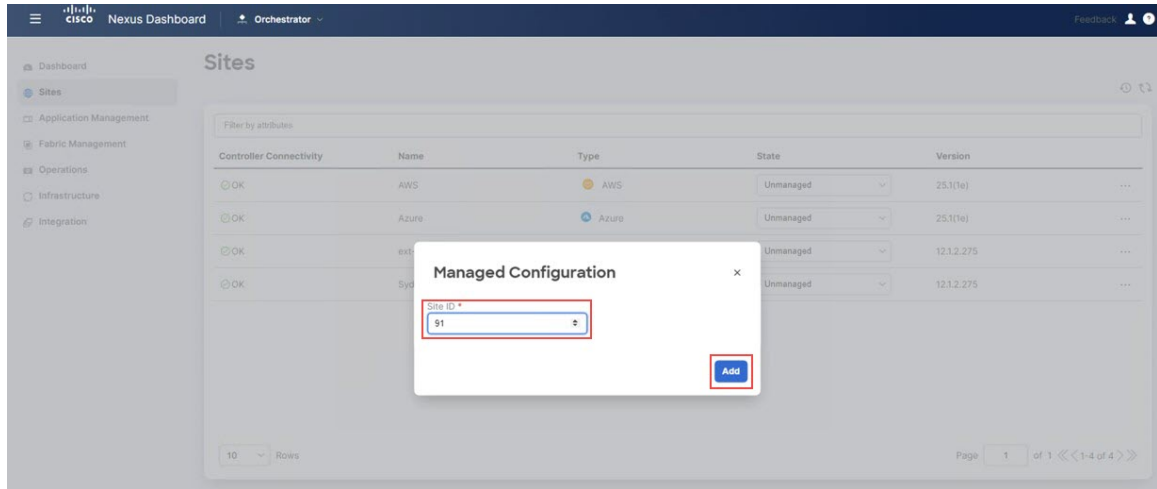
Figure 41:

The screenshot shows the 'Sites' page in the Cisco Nexus Dashboard. The first site, 'AWS', now has a state of 'Managed'. The 'State' column is highlighted with a blue box.

Controller Connectivity	Name	Type	State	Version
OK	AWS	AWS	Managed	25.1(1e)
OK	Azure	Azure	Unmanaged	25.1(1e)
OK	ext-fab-1	NDFC	Unmanaged	12.1.2.275
OK	Sydney	NDFC	Unmanaged	12.1.2.275

- Provide a site ID that is unique to this particular site (a site ID that does not conflict with site IDs for any other site being managed through this NDO), then click **Add**.

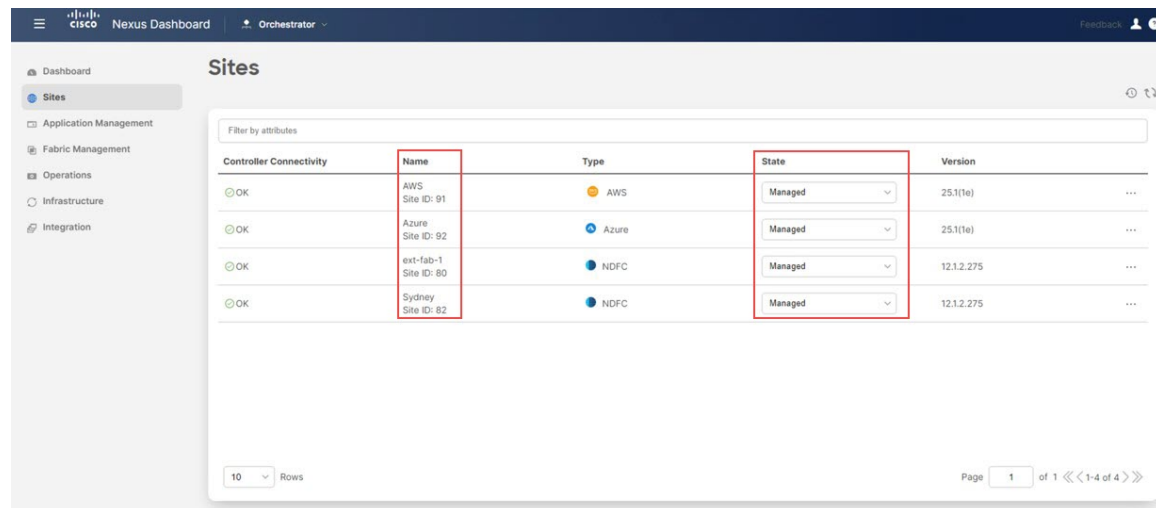
Figure 42:



- c) Repeat these steps for the remaining sites in NDO to change each site to the **Managed** state and provide a unique site ID for each site.

The following figure shows an example of all four sites (the two NDFC sites and the two cloud sites) with their states changed to **Managed** and a unique site ID provided for each site.

Figure 43:



What to do next

Complete the site-to-site connectivity between the NDFC and the cloud sites using the procedures provided in [Complete Site-to-Site Connectivity Between NDFC and Cloud Sites](#), on page 39.

Complete Site-to-Site Connectivity Between NDFC and Cloud Sites

Follow the procedures in the following sections to complete the site-to-site connectivity between the NDFC and cloud sites.

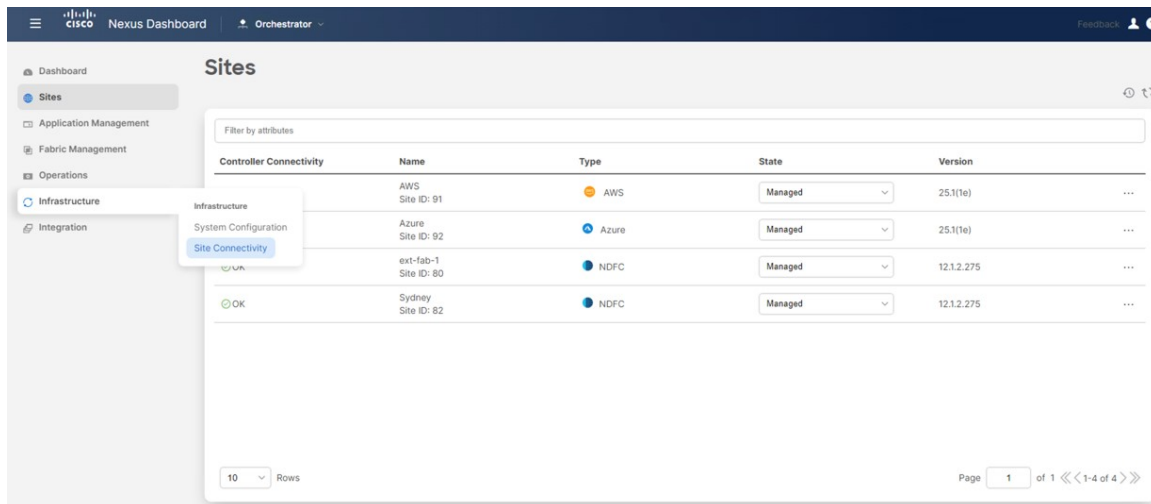
Complete the Necessary Control Plane Configurations

Before you begin

Onboard the NDFC and cloud sites in ND and NDO using the procedures provided in [Onboard the NDFC and Cloud Sites into ND and NDO, on page 32](#).

Step 1 In NDO, navigate to **Infrastructure > Site Connectivity**.

Figure 44:

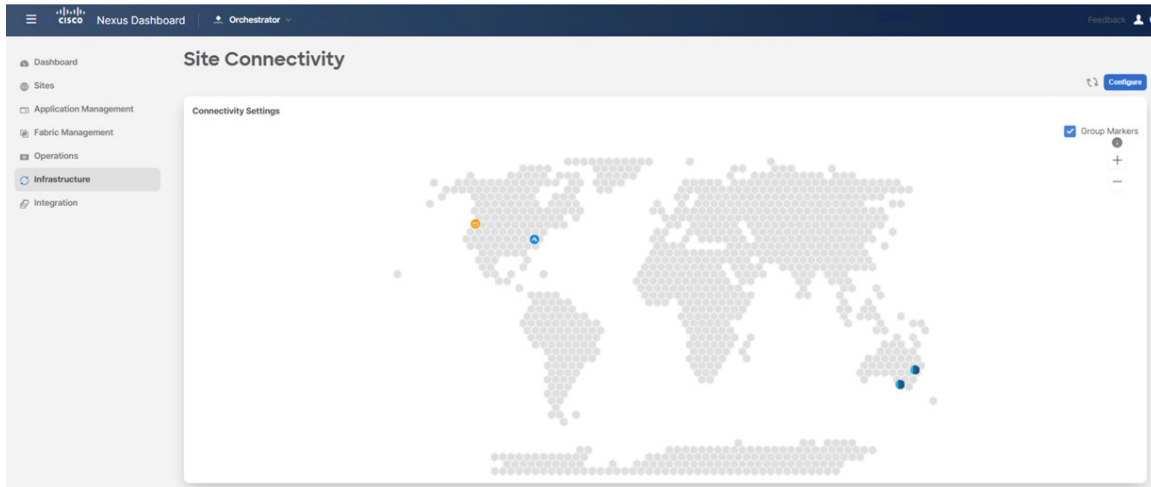


The screenshot shows the Cisco Nexus Dashboard interface. The left sidebar contains navigation options: Dashboard, Sites, Application Management, Fabric Management, Operations, Infrastructure, and Integration. The main content area is titled 'Sites' and features a 'Filter by attributes' search bar. Below the search bar is a table with the following columns: Controller Connectivity, Name, Type, State, and Version. The table contains four rows of data. At the bottom of the table, there is a '10 Rows' dropdown and a 'Page 1 of 1' indicator.

Controller Connectivity	Name	Type	State	Version
Infrastructure	AWS Site ID: 91	AWS	Managed	25.1(1e)
System Configuration	Azure Site ID: 92	Azure	Managed	25.1(1e)
	ext-fab-1 Site ID: 80	NDFC	Managed	12.1.2.275
OK	Sydney Site ID: 82	NDFC	Managed	12.1.2.275

At this point, you will see the sites on the world map but they will not have any links in between, which means that there is no connectivity between the sites at this point.

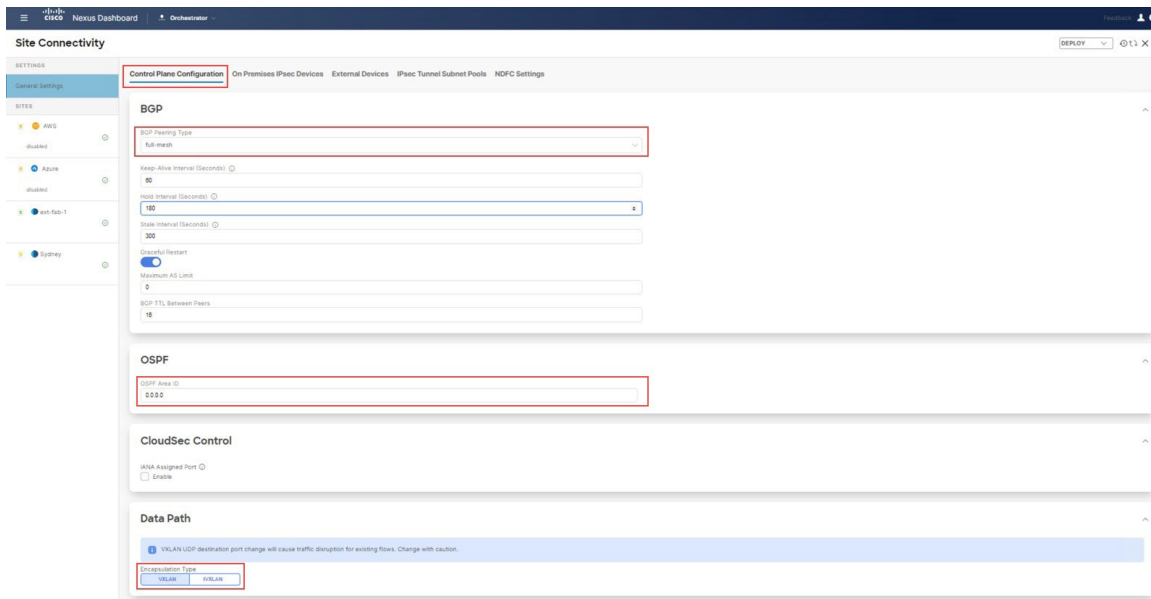
Figure 45:



Step 2 In the upper right area in the **Site Connectivity** window, click **Configure**. The **General Settings** area of the **Site Connectivity** window appears.

Step 3 In the **General Settings** area, click the **Control Plane Configuration** tab, then make the necessary configurations in this page.

Figure 46:



Note that BGP is used for underlay connectivity between on-premises and cloud sites, whereas OSPF is used for cloud-to-cloud underlay connectivity.

Note These general BGP settings apply to the use of BGP for both underlay and overlay connectivity and normally should not be changed, with the exception of the **BGP Peering Type** option in the next step that only applies to overlay peering.

- Step 4** For overlay connectivity between on-premises and cloud sites, in the **BGP Peering Type** field in the **BGP** area, choose either **full-mesh** or **route-server**.
- See [Supported Topologies](#) to see the topologies that use full mesh or route server connectivity.
- For this specific use case, we are configuring a deployment based on the [Option 1](#) topology in [Supported Topologies with IPsec \(Multi-Cloud\)](#), so we would choose **full-mesh** for this use case.
- Step 5** Define any remaining parameters in the **BGP** area, if necessary.
- Step 6** For cloud-to-cloud underlay connectivity, in the **OSPF** area, enter the appropriate value in the **OSPF Area ID** field.
- This configuration is necessary for cloud-to-cloud connectivity because the underlay routing between two cloud sites use OSPF. For this example, enter OSPF Area ID 0.0.0.0 in this field.
- Step 7** Under **Data Path**, locate the **Encapsulation Type** area and select **VXLAN**.
- By default, NDO uses standard VXLAN in data-plane for Hybrid Cloud for NDFC based on-premises fabrics. The other option is iVXLAN, which should be used when building Hybrid Cloud connectivity for ACI sites (since ACI uses iVXLAN).

What to do next

Follow the procedures provided in [Add the On-Premises IPsec Device and IPsec Tunnel Subnet Pools, on page 41](#).

Add the On-Premises IPsec Device and IPsec Tunnel Subnet Pools

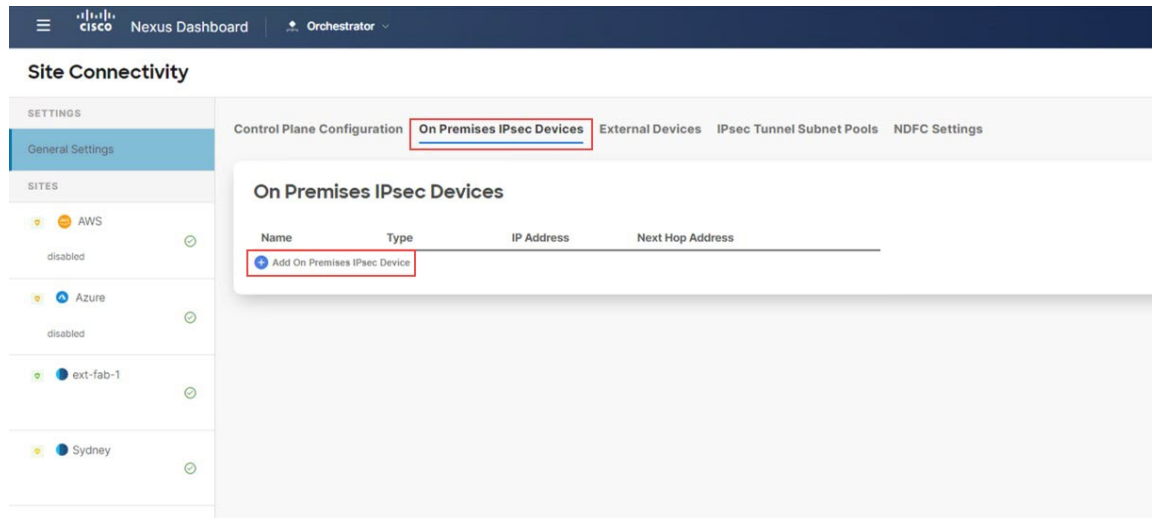
In this section, you will add the on-premises IPsec device (the Cisco Catalyst 8000V in the NDFC external fabric site) and configure the IPsec tunnel pool.

Before you begin

Follow the procedures provided in [Complete the Necessary Control Plane Configurations, on page 39](#).

-
- Step 1** In the same **General Settings** page, click the **On Premises IPsec Devices** tab.
- Step 2** Click **Add On Premises IPsec Device**.

Figure 47:



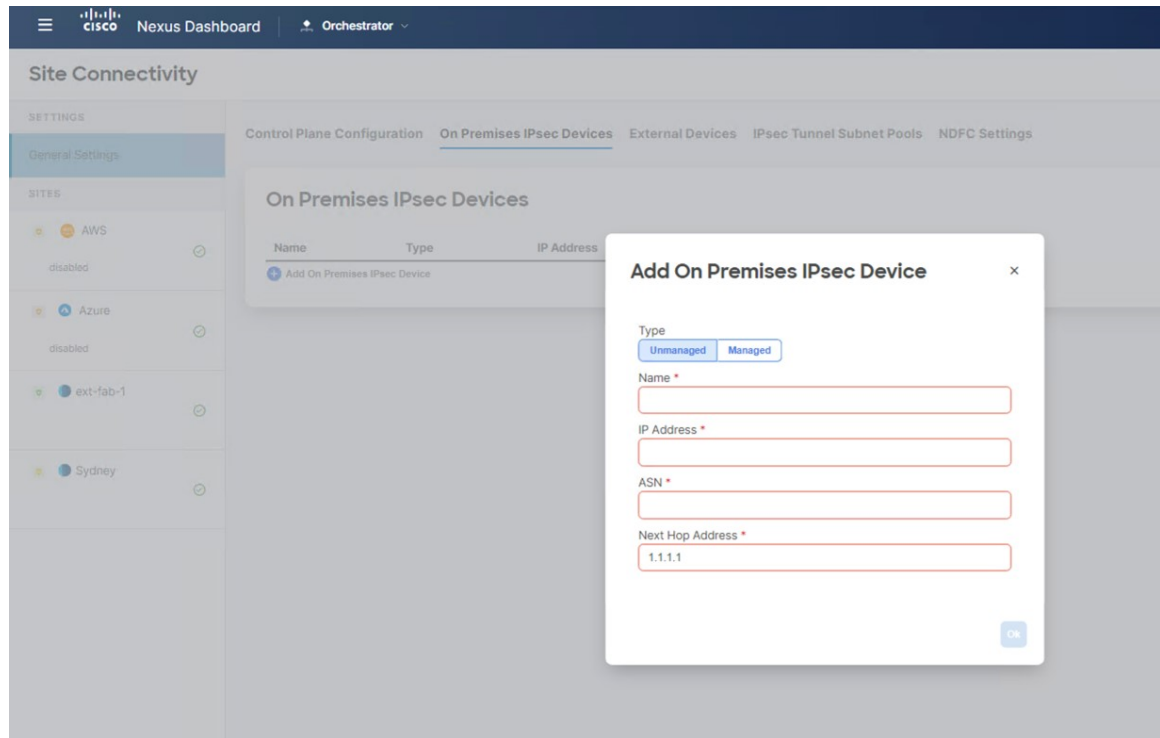
The **Add On Premises IPsec Device** page appears.

Step 3 In the **Type** field, choose either **Unmanaged** or **Managed**.

Both the **Unmanaged** and **Managed** options are supported for the on-premises IPsec device.

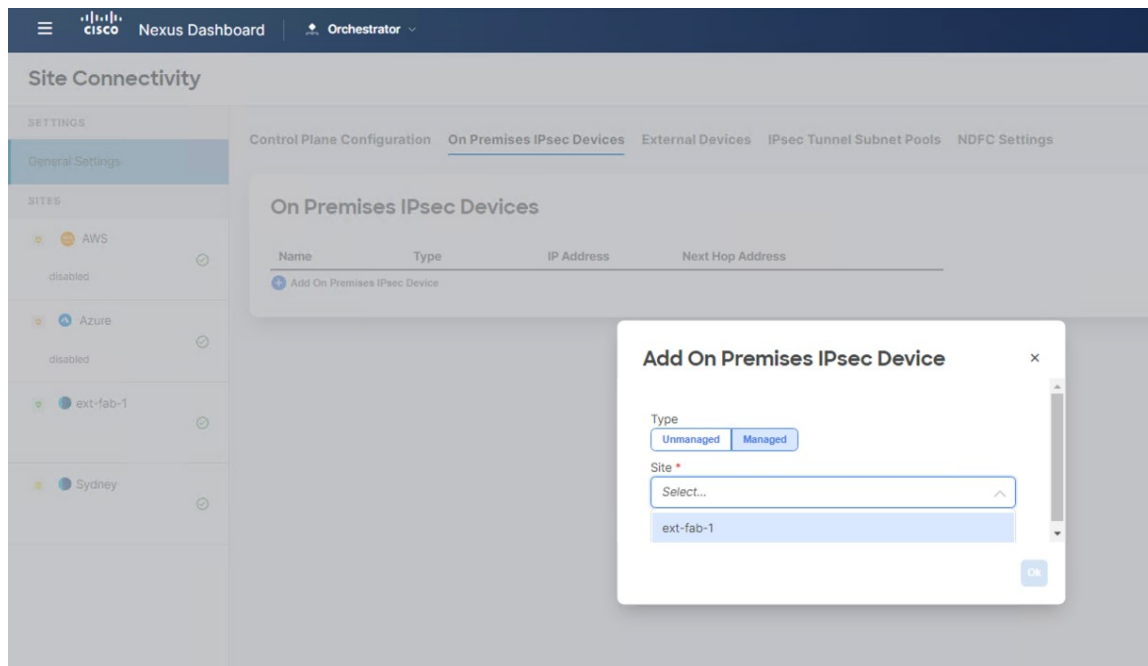
- If you choose the **Unmanaged** option for the on-premises IPsec device, you must enter the necessary information for this unmanaged on-premises IPsec device, such as the **Name**, **IP Address**, and **Next Hop Address**. Use the **Unmanaged** when the on-premises IPsec device is not being managed by NDFC (either that device is not supported by NDFC or it's a third-party device). NDO then generates the required configuration for the unmanaged IPsec device, which can be downloaded and applied on the on-premises IPsec devices manually.

Figure 48:



- If you choose the **Managed** option for the on-premises IPsec device, the **Site** field becomes available below the **Managed** option. The sites available in the **Site** field is based on information that NDO pulls from NDFC for the external fabrics configured in NDFC.

Figure 49:



Choose the NDFC external fabric with the managed on-premises IPsec device. The **ASN** field is automatically populated in this case based on the site that you chose.

For this use case example, we will choose **Managed** for the type for the on-premises IPsec device.

- a) In the **Device** field, select the on-premises IPsec device that you want to use for this deployment.

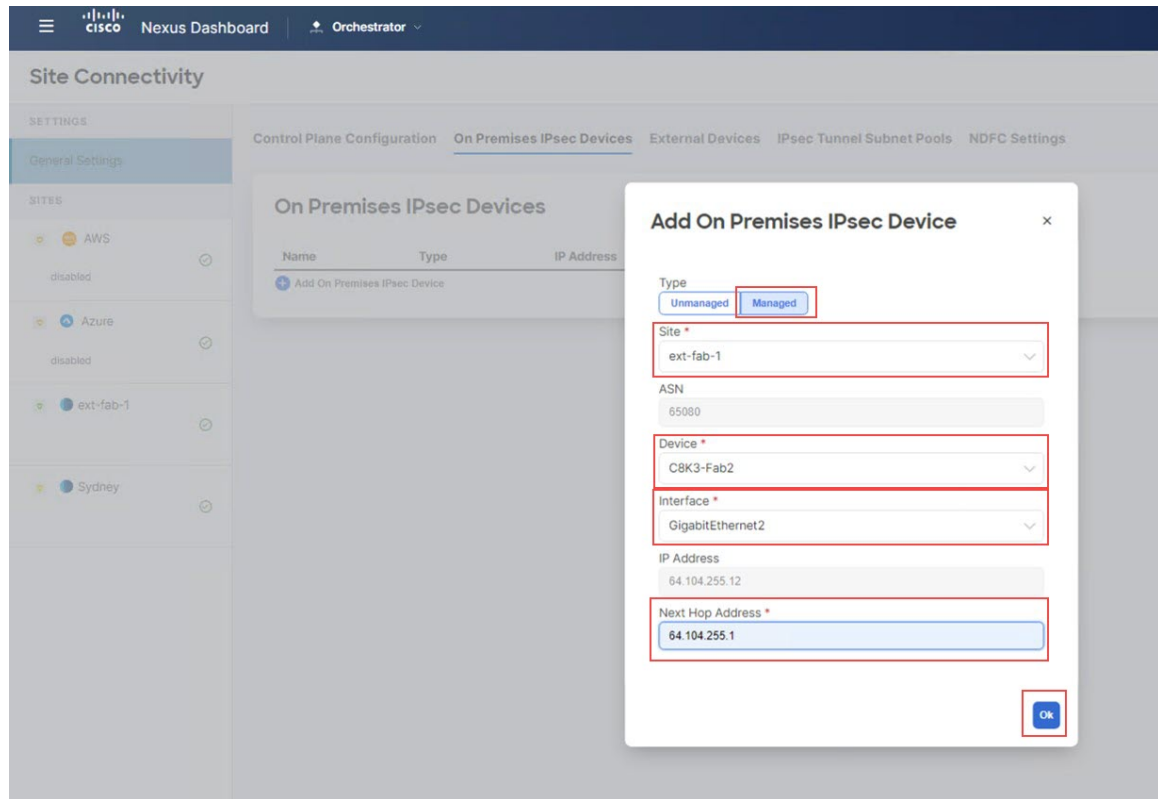
The devices available in the **Device** field is based on information that NDO pulls from NDFC for the on-premises IPsec devices configured in the NDFC site that you selected above. The **ASN** field is then automatically populated based on the on-premises IPsec device that you selected in the **Device** field.

- b) In the **Interface** field, select the appropriate interface that you want to use for the on-premises IPsec device.

The **IP Address** field for this interface is then automatically populated based on the interface that you selected in the **Interface** field.

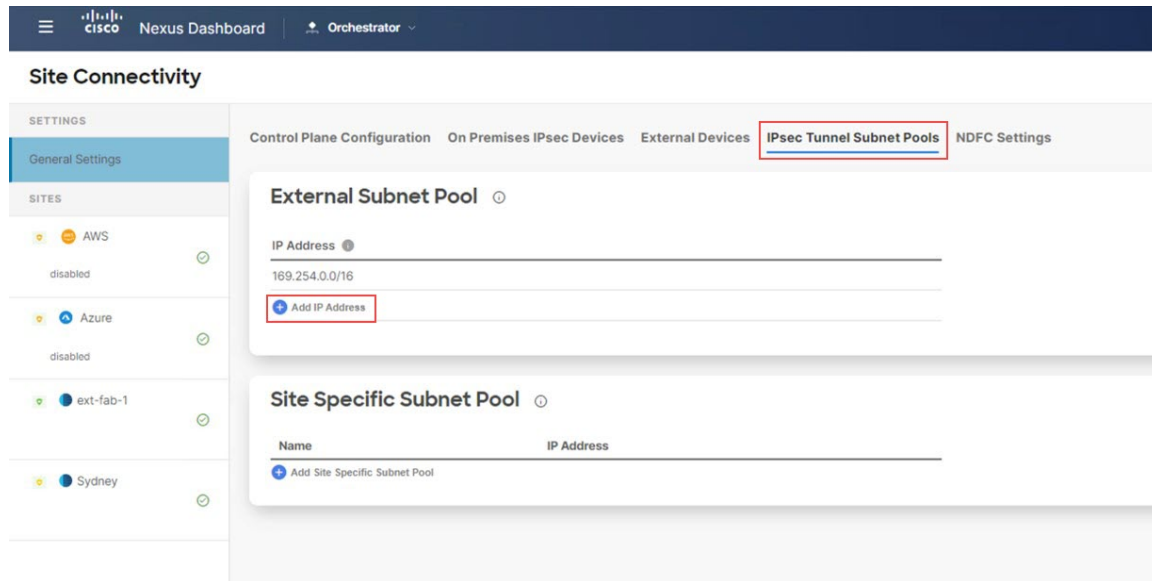
- c) In the **Next Hop Address** field, enter the address to be used for the route that you want to be configured on IPsec.

Figure 50:



- Step 4** When you have finished entering the necessary information in the **Add On Premises IPsec Device** page, click **Ok**. You are returned to the **On Premises IPsec Devices** page, which now shows the configured on-premises IPsec device.
- Step 5** Click the **IPsec Tunnel Subnet Pools** tab to configure the IPsec tunnel subnet pools. The **IPsec Tunnel Subnet Pools** information is required for the cloud tunnel IP assignment.
- Step 6** In the **External Subnet Pool** area, click **Add IP Address**.

Figure 51:

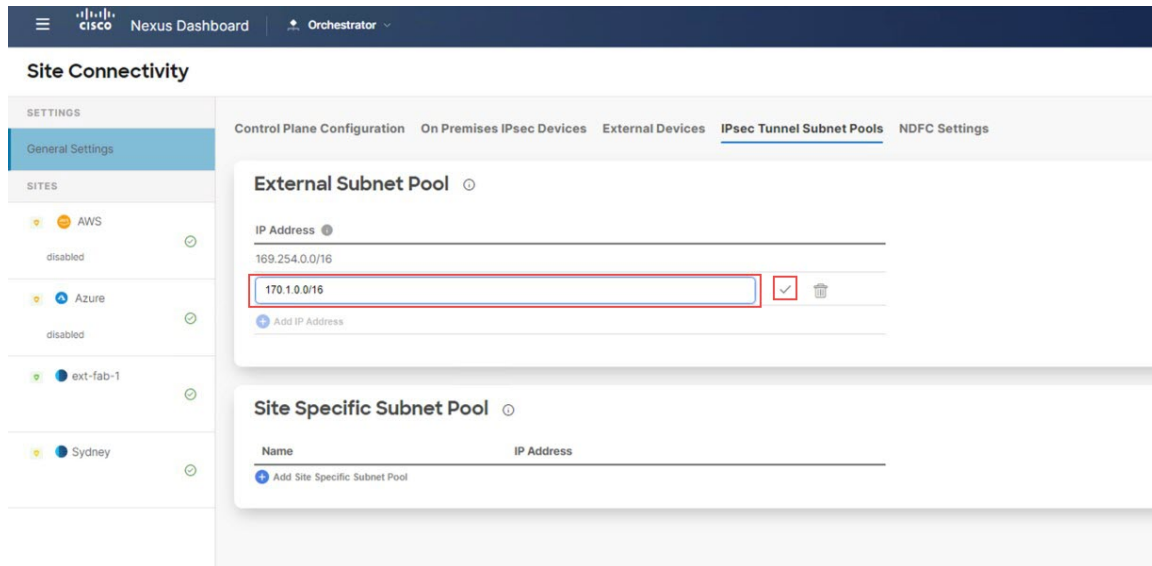


Step 7 Enter the IP subnet pool that you will use for the IPsec tunnels.

Define the IP subnet pool, using public or private IP addresses, for the IPsec tunnels. This is the pool of IP addresses for the IPsec tunnel addressing between the on-premises external device to the Cisco Catalyst 8000V, and between the Cisco Catalyst 8000Vs deployed in the cloud sites.

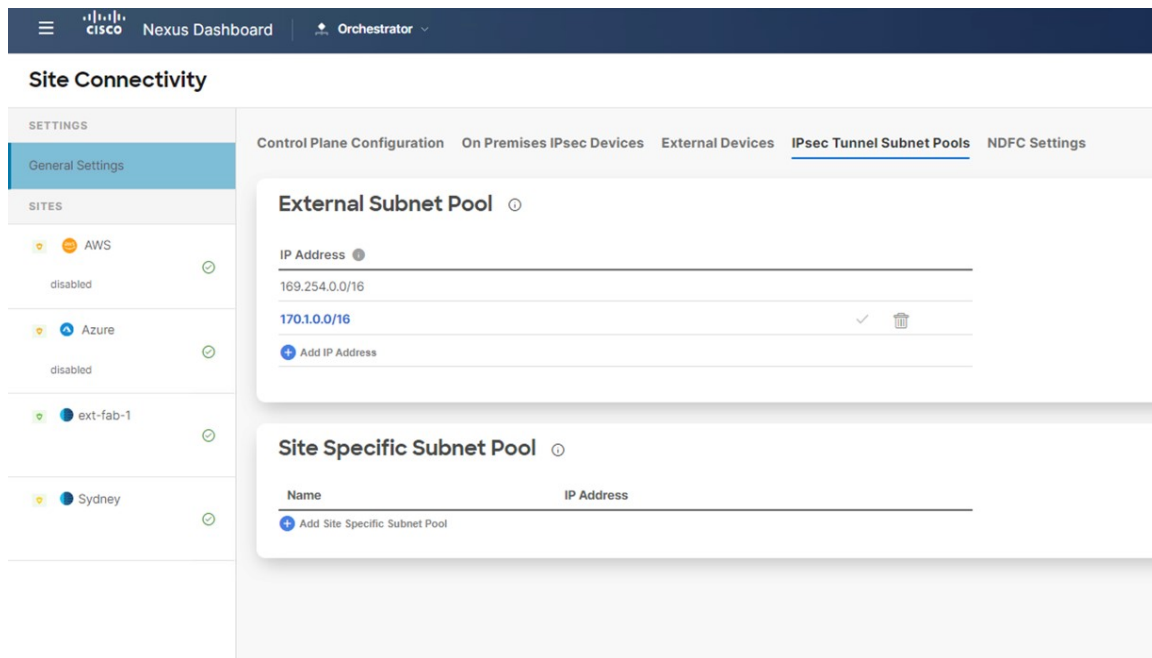
- A /30 subnet is required for each IPsec tunnel.
- The pool size should be able to accommodate all the IPsec tunnels.
- The minimum allowed pool size is of 512 addresses (/23 subnet) .
- Use a range of IP addresses (public or private) that does not overlap with other IP addresses in your environment.

Figure 52:



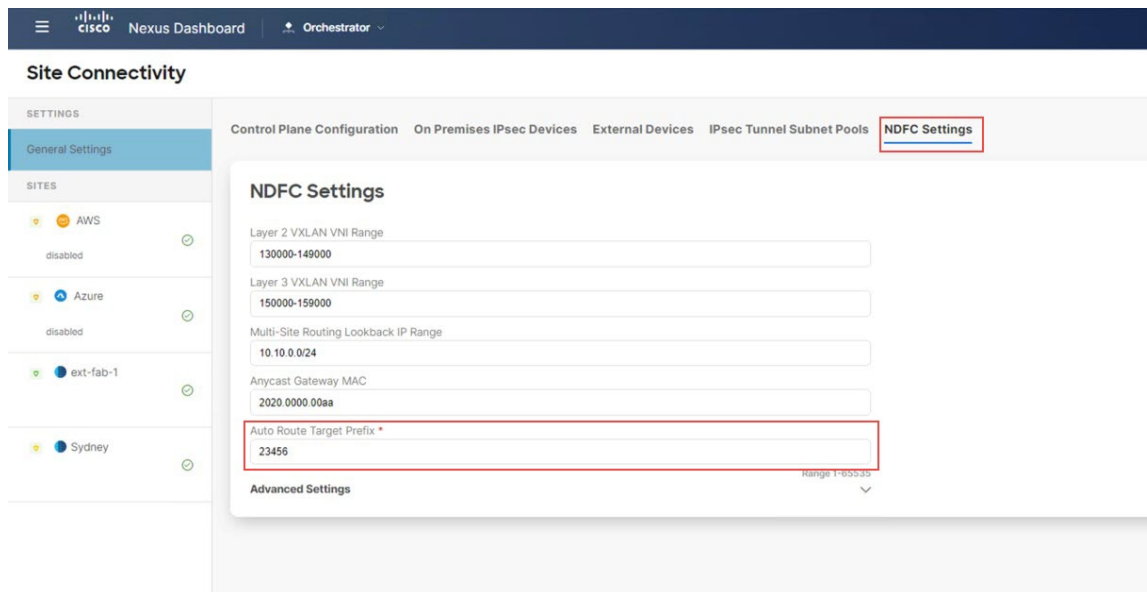
- Step 8** Click the checkbox to accept the IP subnet pool that you entered.
The IP subnet pool appears under the **External Subnet Pool** area.

Figure 53:



- Step 9** Click the **NDFC Settings** tab and enter the necessary information in the **Auto Route Target Prefix**, if necessary.

Figure 54:



Under NDFC settings in NDO, the Route Target Prefix for the Route Target generation is set with a default value of 23456 for NDFC (Cloud Network Controller has different values for this setting), so you can change this value in the **Auto Route Target Prefix** field if required to avoid any possible duplication. Setting the value in this field allows NDO to push this value out to NDFC by NDO.

What to do next

Follow the procedures provided in [Add Ports for the External Devices in the NDFC External Fabric](#), on page 48.

Add Ports for the External Devices in the NDFC External Fabric

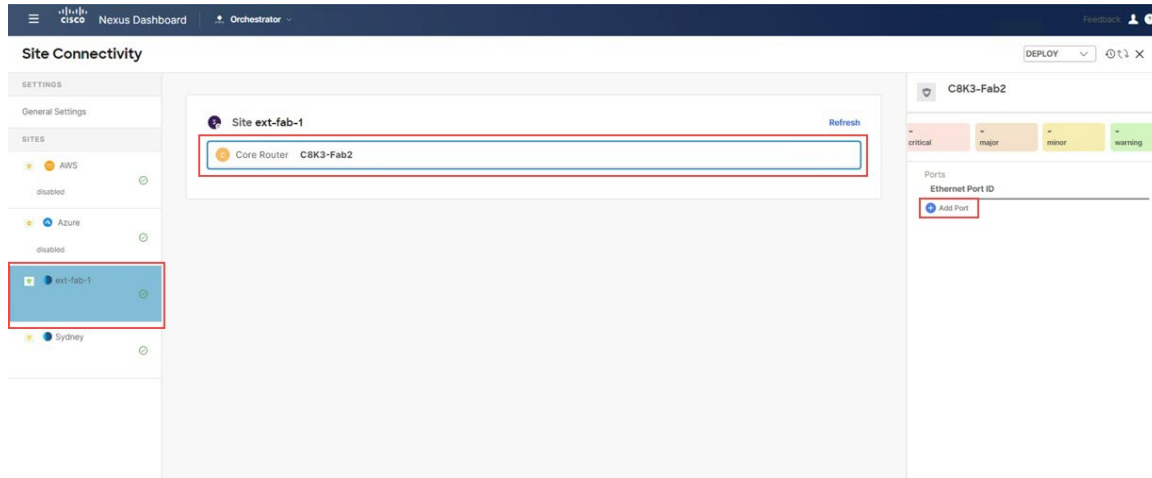
In this section, you will add and configure the necessary ports for the external devices in the NDFC external fabric. These are the interfaces connecting the core router to the BGW nodes.

Before you begin

Follow the procedures provided in [Add the On-Premises IPsec Device and IPsec Tunnel Subnet Pools](#), on page 41.

- Step 1** In the left pane under **General Settings: Sites**, click the NDFC external fabric (the ext-fab-1 site in this example).
- Step 2** In the middle pane, click on the first external device in the NDFC external fabric.
- Step 3** In the right pane, click **Add Port**.

Figure 55:



Step 4 Enter the necessary information for the port configuration, including the IP address, remote IP address, and remote ASN.

Note The **Towards Cloud Router** option is only applicable for border gateways in a hub site. You will not enable this option in this window for the following reasons:

- Because the topology that we're using for this example use case does not use a hub site, you will not enable the **Towards Cloud Router** for this example use case.
- Even if we were configuring for a topology that uses a hub site, such as [Option 3 in Supported Topologies with IPsec \(Multi-Cloud\)](#), we would not enable this option in this page for the external device in the NDFC external fabric for that hub site topology; instead, we would enable this option in the page for the BGW spine device in the NDFC VXLAN fabric, as described in [Add the Port for the BGW Spine Device in the NDFC VXLAN Fabric, on page 53](#).

Figure 56:

Add Port ×

Ethernet Port ID *
GigabitEthernet4 ✕ ▾

IP Address *
10.140.1.1/30

Description
towards on-prem Spine BGW E1/32

Remote Address *
10.140.1.2

Remote ASN *
65084

MTU *
9216

Inherit BGP Authentication and BFD ⓘ

BGP Authentication
 None Simple Cisco

Towards Cloud Router ⓘ

BFD Enabled

Ok

Step 5 Click **Ok** when you are finished.

Step 6 Repeat these steps for the remaining external devices.

What to do next

Follow the procedures provided in [Define the Multi-Site VIP for the VXLAN Fabric Site, on page 50](#).

Define the Multi-Site VIP for the VXLAN Fabric Site

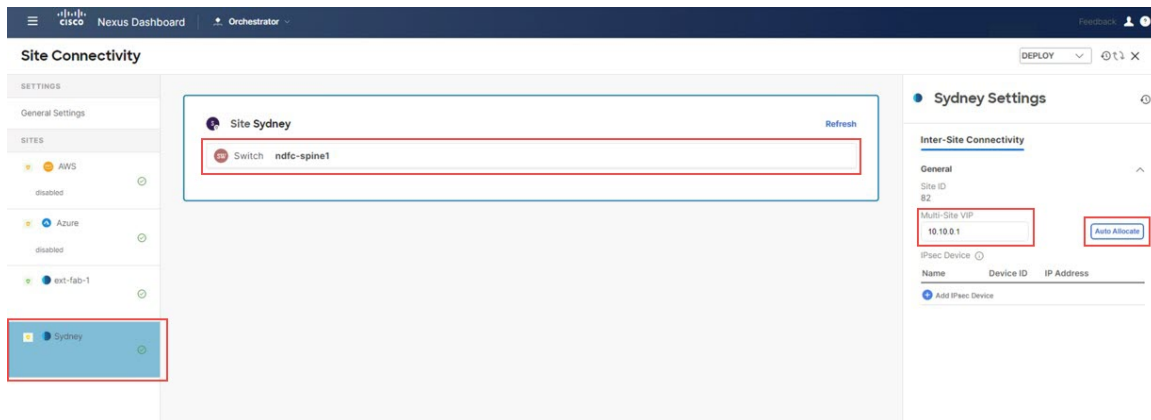
In this section, you will define the Multi-Site VIP for the VXLAN fabric site.

Before you begin

Follow the procedures provided in [Add Ports for the External Devices in the NDFC External Fabric, on page 48](#).

- Step 1** In the left pane under **General Settings: Sites**, click the NDFC VXLAN fabric site.
- Step 2** In the middle pane, click on the spine device.
- Step 3** In the right pane, under **Inter-Site Connectivity**, define the Multi-Site VIP in the **Multi-Site VIP** field. You can click **Auto Allocate** or you can explicitly define the IP address for the Multi-Site VIP.

Figure 57:



What to do next

Follow the procedures provided in [Map the IPsec Device to the VXLAN Fabric Site, on page 51](#).

Map the IPsec Device to the VXLAN Fabric Site

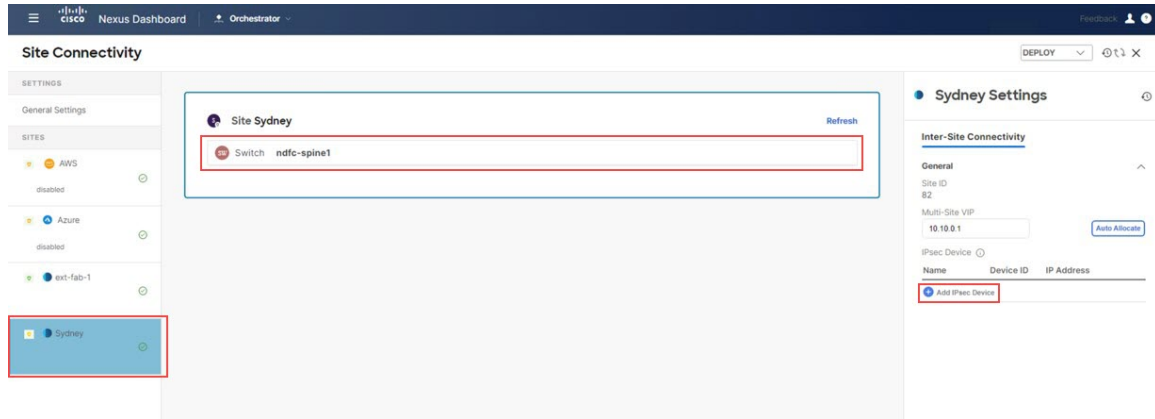
In this section, you will map the IPsec device to the VXLAN fabric site.

Before you begin

Follow the procedures provided in [Define the Multi-Site VIP for the VXLAN Fabric Site, on page 50](#).

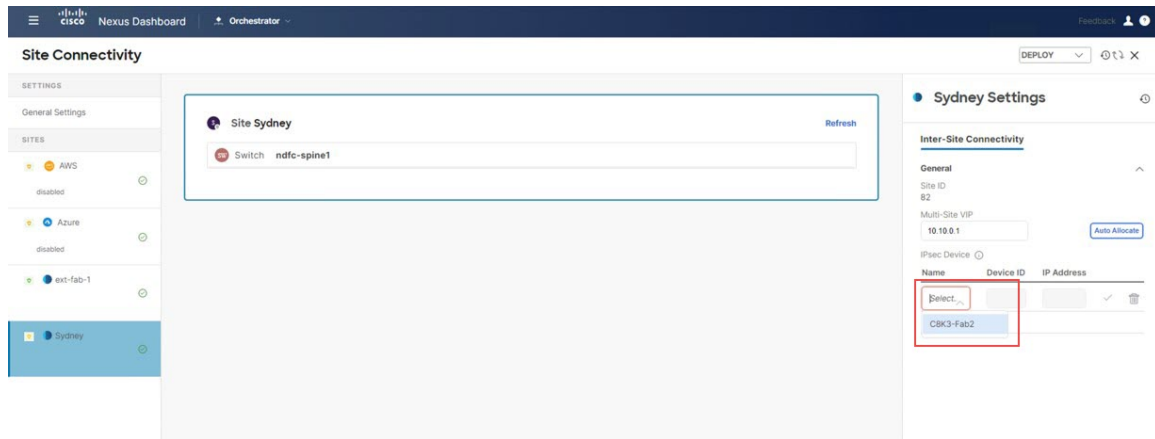
- Step 1** In the left pane under **General Settings: Sites**, click the NDFC VXLAN fabric site.
- Step 2** In the middle pane, click the spine device.
- Step 3** In the right pane, under **Inter-Site Connectivity**, click **Add IPsec Device**.

Figure 58:



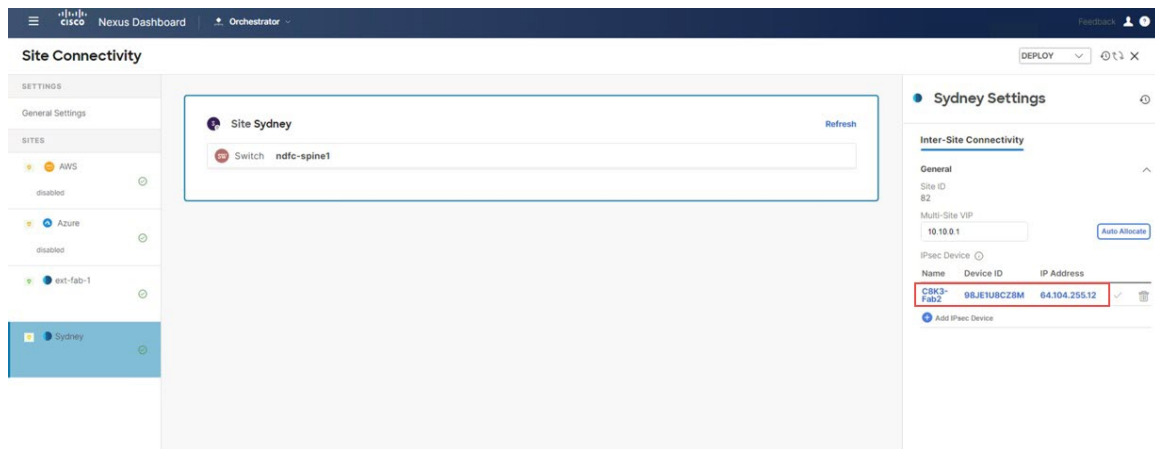
Step 4 Click **Select**, then choose the appropriate IPsec device.

Figure 59:



The on-premises IPsec device is now mapped to the VXLAN fabric site.

Figure 60:



- Step 5** Repeat this step for each on-premises IPsec device (Cisco Catalyst 8000V) that will be used to connect the NDFC VXLAN site to the cloud sites.

What to do next

Configure the ports on the BGW spine device connecting to the core router (Cisco Catalyst 8000V) using the procedures provided in [Add the Port for the BGW Spine Device in the NDFC VXLAN Fabric, on page 53](#).

Add the Port for the BGW Spine Device in the NDFC VXLAN Fabric

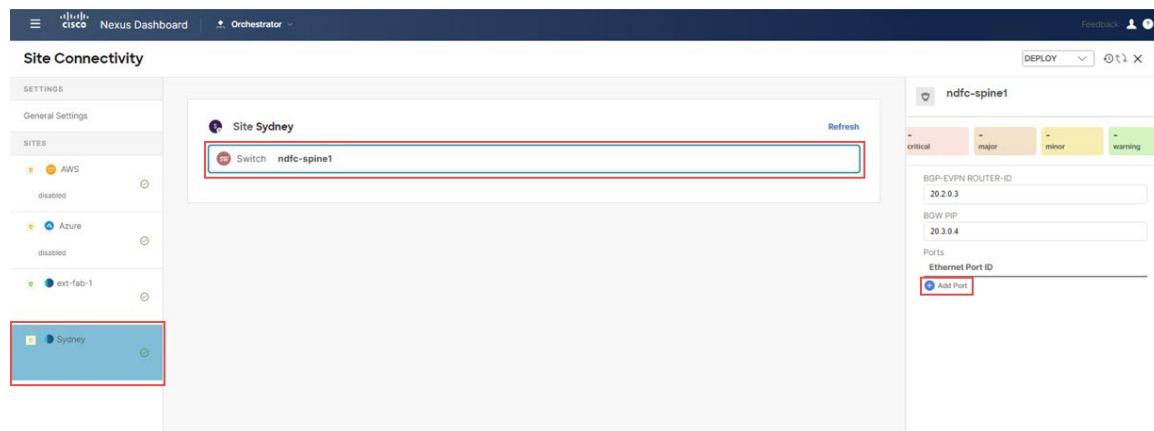
In this section, you will add and configure the necessary port for the BGW spine device in the NDFC VXLAN fabric facing towards the on-premises IPsec device.

Before you begin

Follow the procedures provided in [Map the IPsec Device to the VXLAN Fabric Site, on page 51](#).

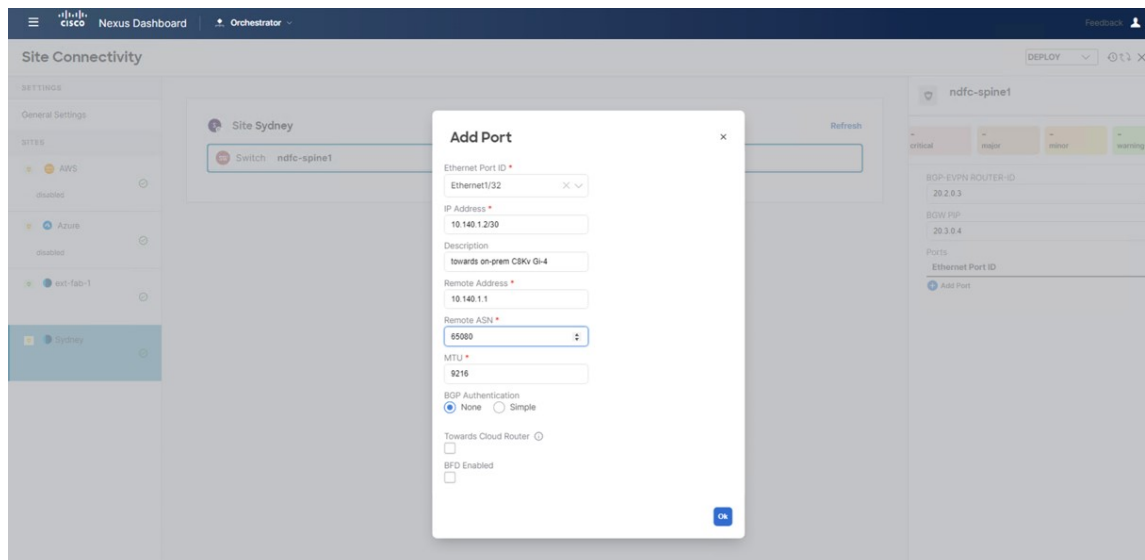
- Step 1** In the left pane under **General Settings: Sites**, click the NDFC VXLAN fabric site.
- Step 2** In the middle pane, click on the spine device.
- Step 3** In the right pane, click **Add Port**.

Figure 61:



- Step 4** Enter the necessary information in this page.
Define the port parameters in this page.

Figure 62:



- In the **Ethernet Port ID** field, select the interface that is facing toward the on-premises Cisco Catalyst 8000V.
- In the **IP Address** field, enter the IP address for this interface. Later in these procedures, Nexus Dashboard Orchestrator will configure this IP address for this interface on the BGW spine switch residing in the VXLAN fabric.
- In the **Remote Address** field, enter the IP address of the gigabit 4 interface of the on-premises IPsec device.
- In the **Remote ASN** field, enter the ASN for the on-premises IPsec device. For example, for this example use case, we would enter 65080 as the ASN for the on-premises IPsec device.

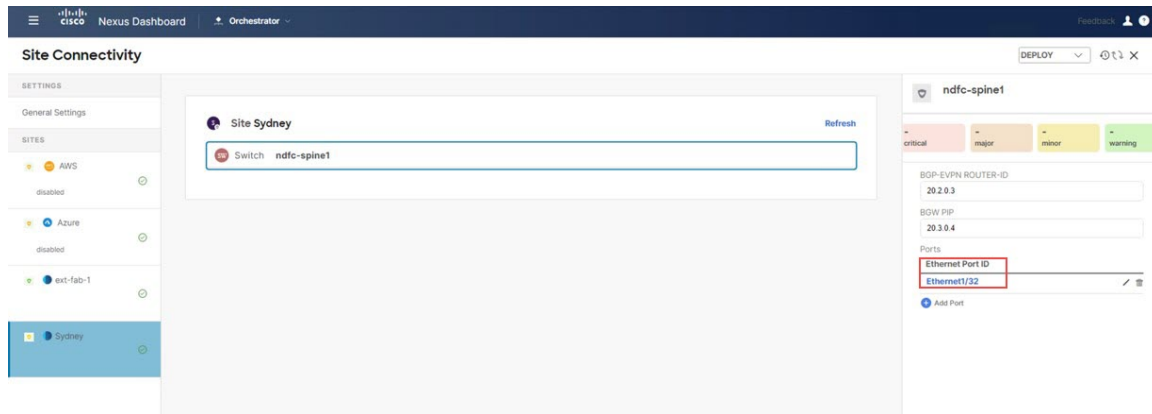
Note The **Towards Cloud Router** option is only applicable for border gateways in an on-premises hub site. This option would need to be enabled in topologies where you are using a hub site, such as [Option 3 in Supported Topologies with IPsec \(Multi-Cloud\)](#).

Because the topology that we're using for this example use case does not use a hub site, you will not enable the **Towards Cloud Router** for this example use case.

Step 5 Click **Ok**.

The port for the BGW spine device is now added in the NDFC VXLAN fabric

Figure 63:



What to do next

Follow the procedures provided in [Connect the First Cloud Site to the NDFC VXLAN Fabric Site](#), on page 55.

Connect the First Cloud Site to the NDFC VXLAN Fabric Site

In this section, you will connect the first cloud site to the NDFC VXLAN fabric site.

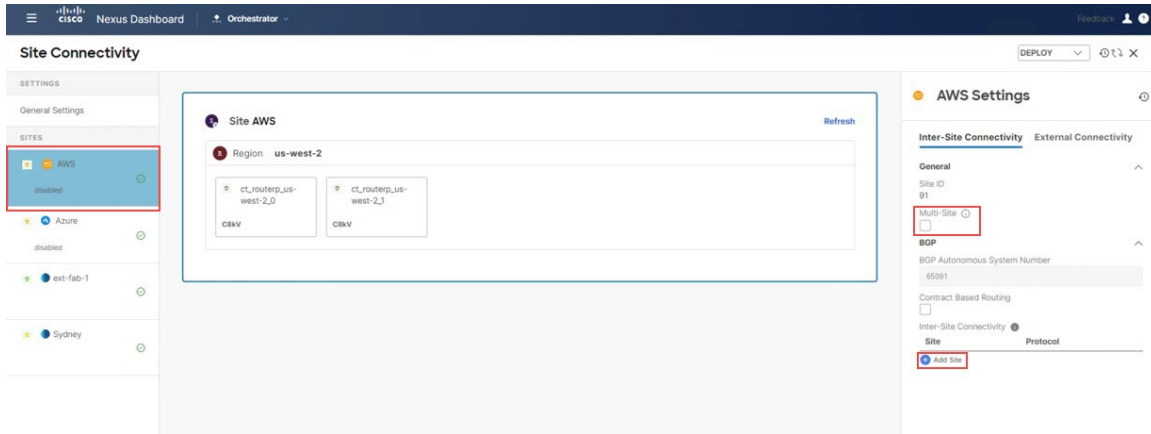
Before you begin

Follow the procedures provided in [Add the Port for the BGW Spine Device in the NDFC VXLAN Fabric](#), on page 53.

- Step 1** In the left pane under **General Settings: Sites**, click the first cloud site (for example, the AWS site).
- Step 2** In the right pane, click **Inter-Site Connectivity**, then check the box under **Multi-Site** to enable that feature. This feature is required for building VXLAN Multisite overlay tunnels between the sites.
- Step 3** In the right pane, click **Add Site**.

Connect the First Cloud Site to the NDFC VXLAN Fabric Site

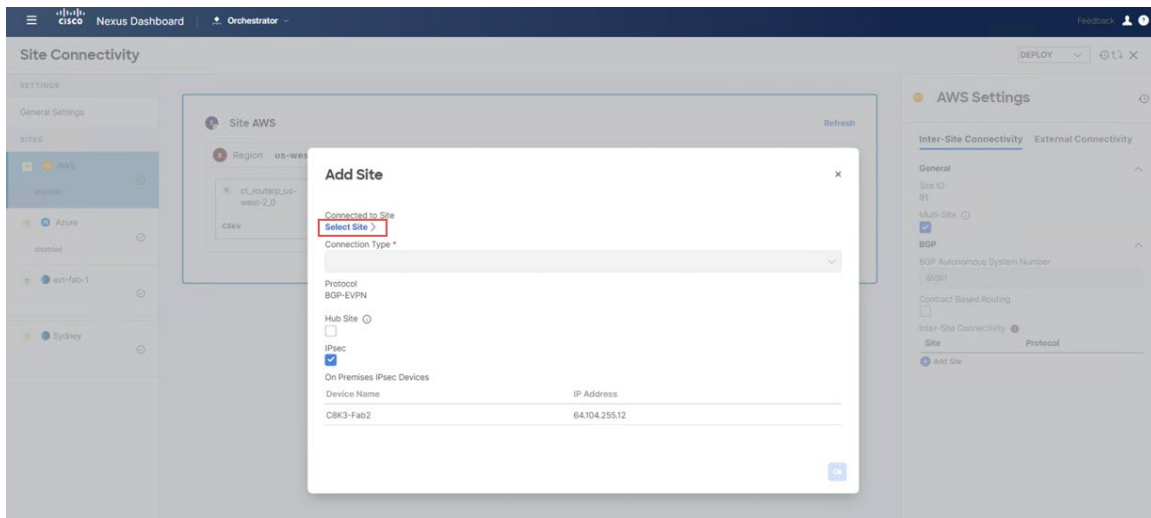
Figure 64:



The **Add Site** page appears.

Step 4 In the **Add Site** page, click **Select a Site**.

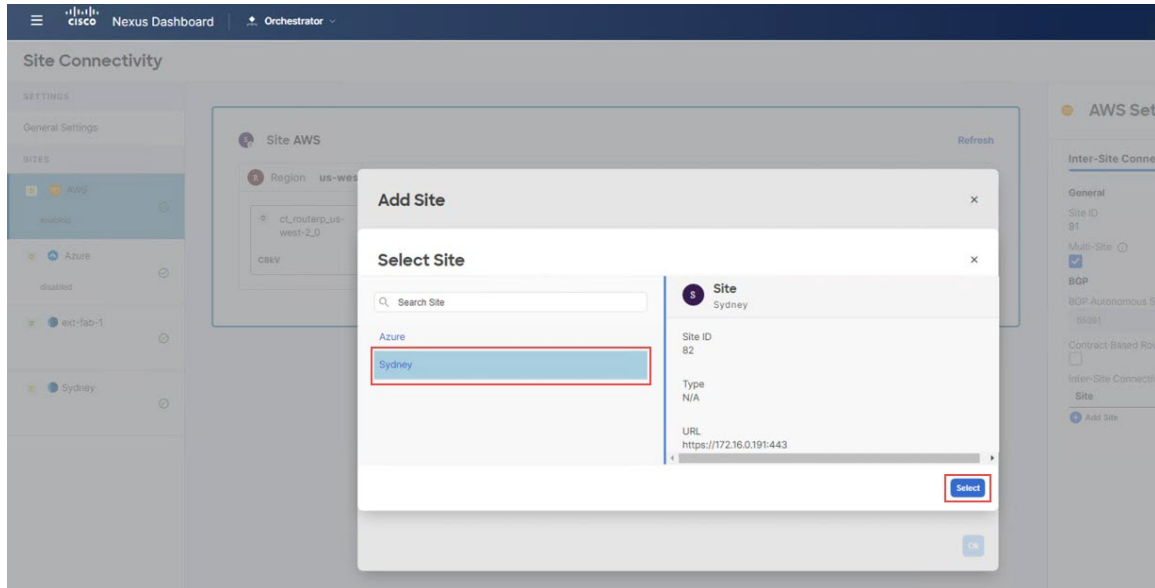
Figure 65:



The **Select a Site** page appears.

Step 5 Select the NDFC VXLAN fabric (the Sydney site in this example), then click **Select**.

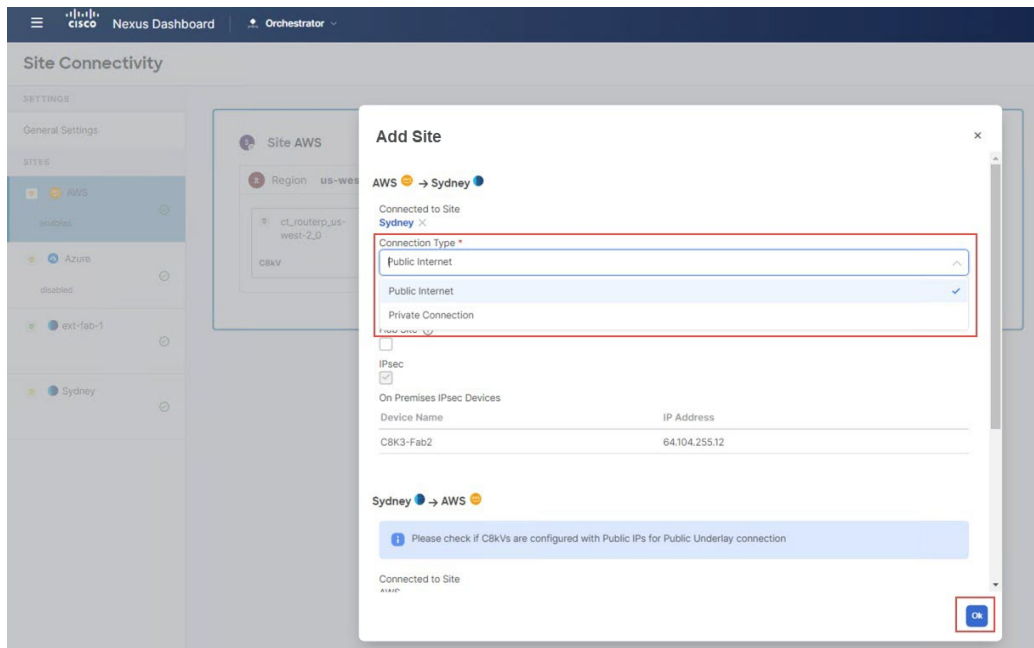
Figure 66:



You are returned to the **Add Site** page.

- Step 6** In the **Add Site** page, in the **Connection Type** field, choose the type of connection that you will use from the first cloud site to the NDFC VXLAN fabric site.

Figure 67:



You can select **Public Internet**, or you can select a **Private Connection** if you are using Direct Connect with AWS or ExpressRoute with Azure.

Connect the First Cloud Site to the Second Cloud Site

- Both **Public Internet** and **Private Connection** options are available for the on-premises site, whereas only the **Public Internet** connection option is available for the cloud sites.
- IPsec is mandatory for the **Public Internet** connection type and is automatically enabled for that connection type, whereas IPsec is optional for the **Private Connection** type.

Note The **Hub Site** option would need to be enabled in topologies where you are using a hub site, such as [Option 3 in Supported Topologies with IPsec \(Multi-Cloud\)](#).

Because the topology that we're using for this example use case does not use a hub site, you will not enable the **Hub Site** option for this example use case.

Step 7 When you have finished the configurations in this page, click **OK**.

What to do next

Follow the procedures provided in [Connect the First Cloud Site to the Second Cloud Site, on page 58](#).

Connect the First Cloud Site to the Second Cloud Site

In this section, you will connect the first cloud site to the second cloud site.

Before you begin

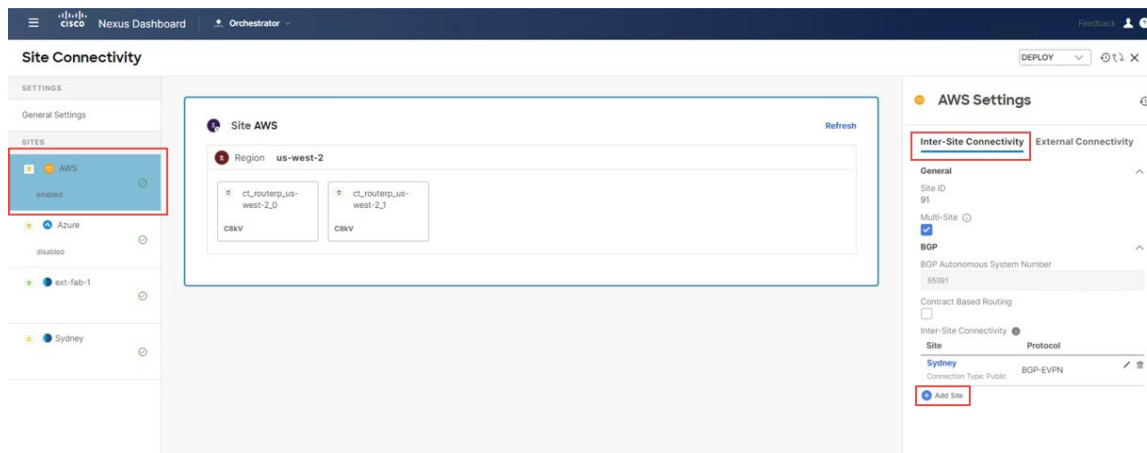
Follow the procedures provided in [Connect the First Cloud Site to the NDFC VXLAN Fabric Site, on page 55](#).

Step 1 In the left pane under **General Settings: Sites**, click the first cloud site (for example, the AWS site).

Step 2 In the right pane, click **Inter-Site Connectivity**.

Step 3 In the right pane, click **Add Site**.

Figure 68:



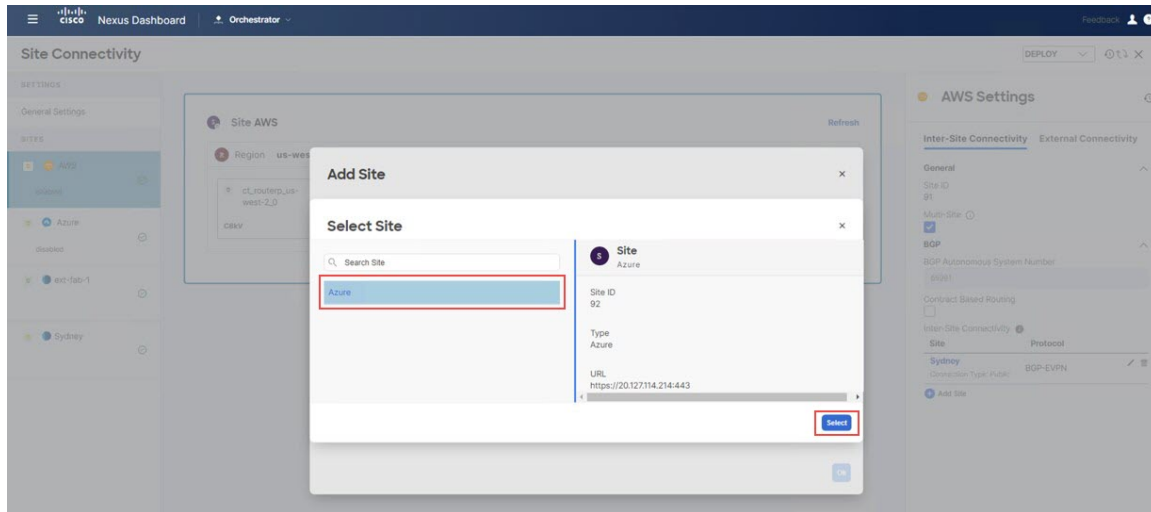
The **Add Site** page appears.

Step 4 In the **Add Site** page, click **Select a Site**.

The **Select Site** page appears.

Step 5 Select the second cloud site (for example, the Azure cloud site), then click **Select**.

Figure 69:



You are returned to the **Add Site** page.

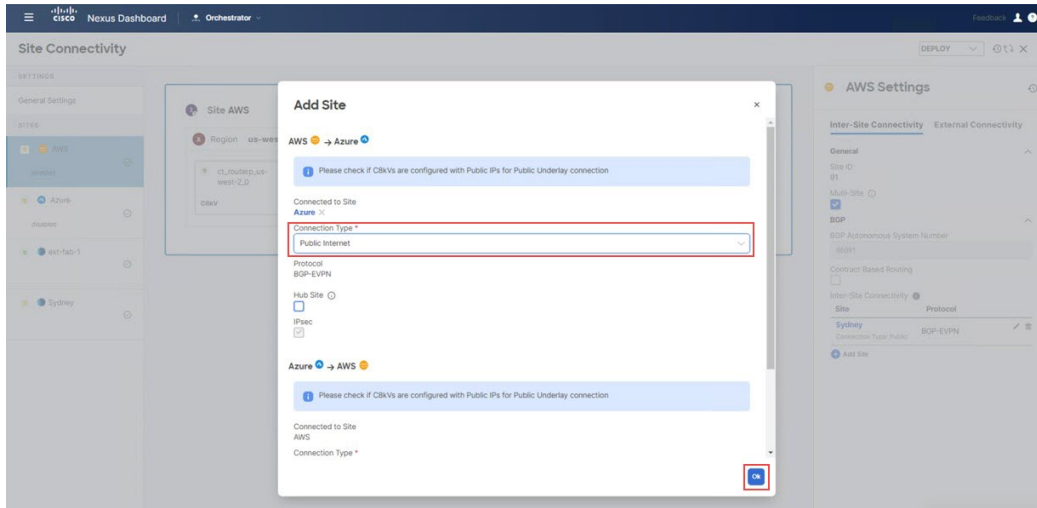
Step 6 In the **Add Site** page, in the **Connection Type** field, choose the type of connection that you will use from the first cloud site to the second cloud site.

For some types of cloud-to-cloud connectivity, you might have these options:

- **Public Internet**
- **Cloud Backbone**

Cloud Backbone can be used to establish connectivity between cloud sites of the same provider (for example, an AWS site 1 managed by one Cloud Network Controller, and an AWS site 2 managed by a second Cloud Network Controller). However, between sites of different cloud providers (for example, AWS to Azure), **Public Internet** is the only option, as shown in the following figure.

Figure 70:



When the **Public Internet** connection type is selected, the **IPsec** option is mandatory and is automatically enabled for that connection type, whereas IPsec is optional for the **Cloud Backbone** type.

Note You will not enable the **Hub Site** option for cloud-to-cloud connectivity, even if the topology uses a hub site (you would enable the **Hub Site** option when configuring connectivity between the cloud site and the NDFC VXLAN fabric site in that case).

Step 7 When you have finished the configurations in this page, click **Ok**.

What to do next

Follow the procedures provided in [Connect the Second Cloud Site to the NDFC VXLAN Fabric Site, on page 60](#).

Connect the Second Cloud Site to the NDFC VXLAN Fabric Site

In this section, you will connect the second cloud site to the NDFC VXLAN fabric site.

The procedures in this section are essentially the same steps that you performed in the previous sections, where you:

- Connected the first cloud site to the NDFC VXLAN fabric site in [Connect the First Cloud Site to the NDFC VXLAN Fabric Site, on page 55](#).
- Connected the first cloud site to the second cloud site in [Connect the First Cloud Site to the Second Cloud Site, on page 58](#).

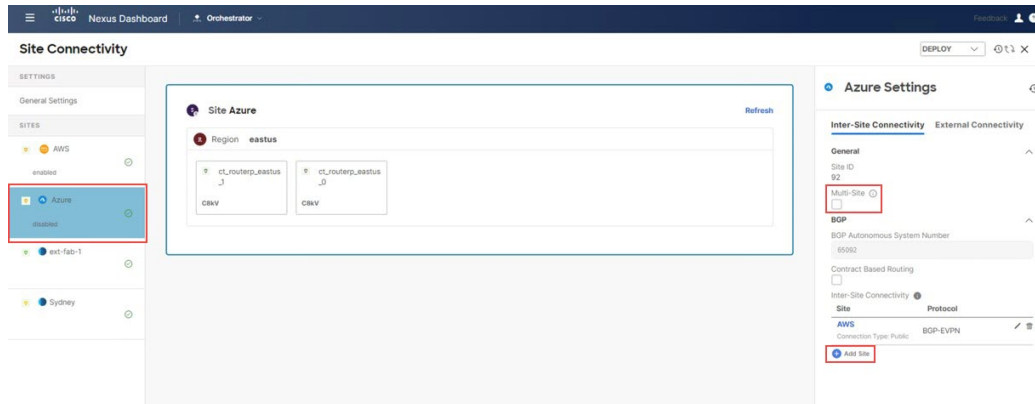
For this section, you will be connecting the second cloud site to the NDFC VXLAN fabric site. Note that because you had already configured connectivity between AWS and Azure in [Connect the First Cloud Site to the Second Cloud Site, on page 58](#), you do not have to configure connectivity from the second cloud site (Azure) back to AWS because that connectivity was already configured in that previous section.

Before you begin

Follow the procedures provided in [Connect the First Cloud Site to the Second Cloud Site, on page 58](#).

- Step 1** In the left pane under **General Settings: Sites**, click the second cloud site (for example, the Azure site).
- Step 2** In the right pane, click **Inter-Site Connectivity**, then check the box under **Multi-Site** to enable that feature.
- Step 3** In the right pane, click **Add Site**.

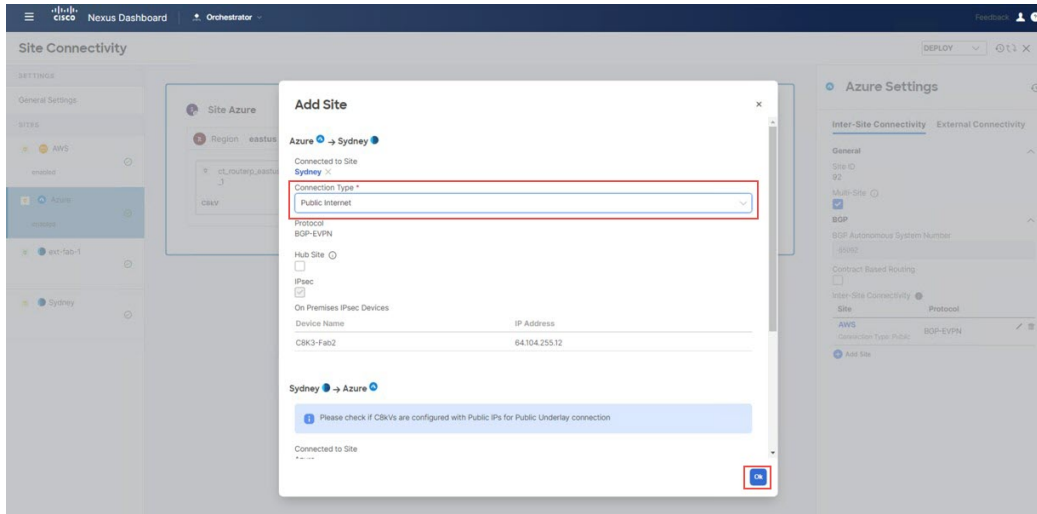
Figure 71:



The **Add Site** page appears.

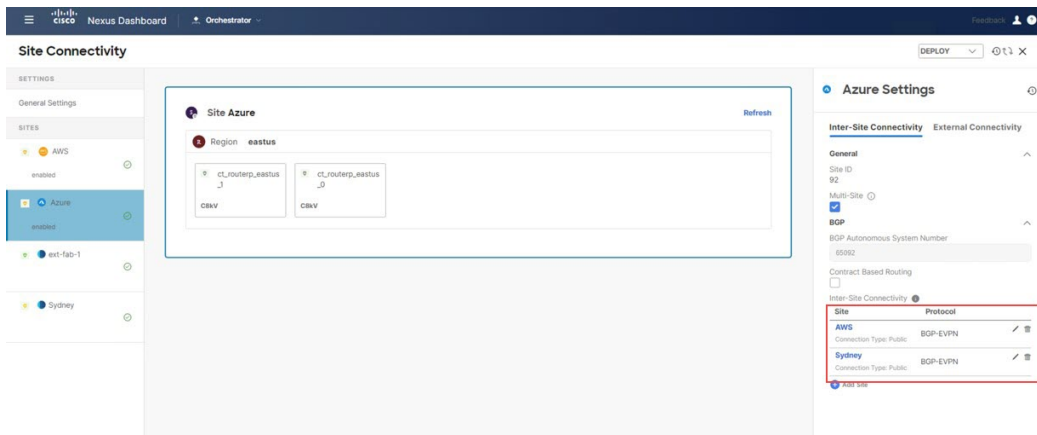
- Step 4** In the **Add Site** page, click **Select a Site**.
The **Select a Site** page appears.
- Step 5** Select the NDFC VXLAN fabric (the Sydney site in this example), then click **Select**.
You are returned to the **Add Site** page.
- Step 6** In the **Add Site** page, in the **Connection Type** field, choose the type of connection that you will use from the second cloud site to the NDFC VXLAN fabric site.

Figure 72:



Step 7 When you have finished the configurations in this page, click **OK**.
The configured sites appear.

Figure 73:



What to do next

Follow the procedures provided in [Deploy the Configuration in Nexus Dashboard Orchestrator, on page 62](#).

Deploy the Configuration in Nexus Dashboard Orchestrator

In this section, you will deploy the configuration in Nexus Dashboard Orchestrator (NDO).

Before you begin

Follow the procedures provided in [Connect the Second Cloud Site to the NDFC VXLAN Fabric Site](#), on page 60.

Step 1

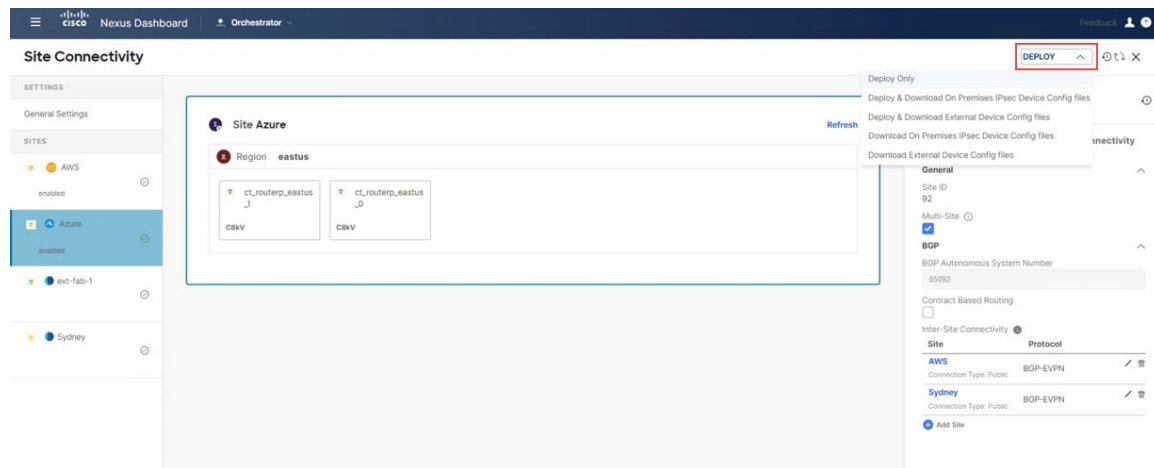
Deploy the configuration in NDO.

- If you chose the **Unmanaged** option for the on-premises IPsec device in [Add the On-Premises IPsec Device and IPsec Tunnel Subnet Pools](#), on page 41, at the top right of the page, click **Deploy** > **Deploy & Download External Device Config files**.

This option downloads a zip file that contains the necessary configuration information that you will use to configure the on-premises IPsec device. A followup screen appears that allows you to select all or some of the configuration files to download.

- If you chose the **Managed** option for the on-premises IPsec device in [Add the On-Premises IPsec Device and IPsec Tunnel Subnet Pools](#), on page 41, at the top right of the page, click **Deploy** > **Deploy Only**.

Figure 74:



Step 2

Click **Yes** in the **Confirmation** window.

NDO does the following things at this point:

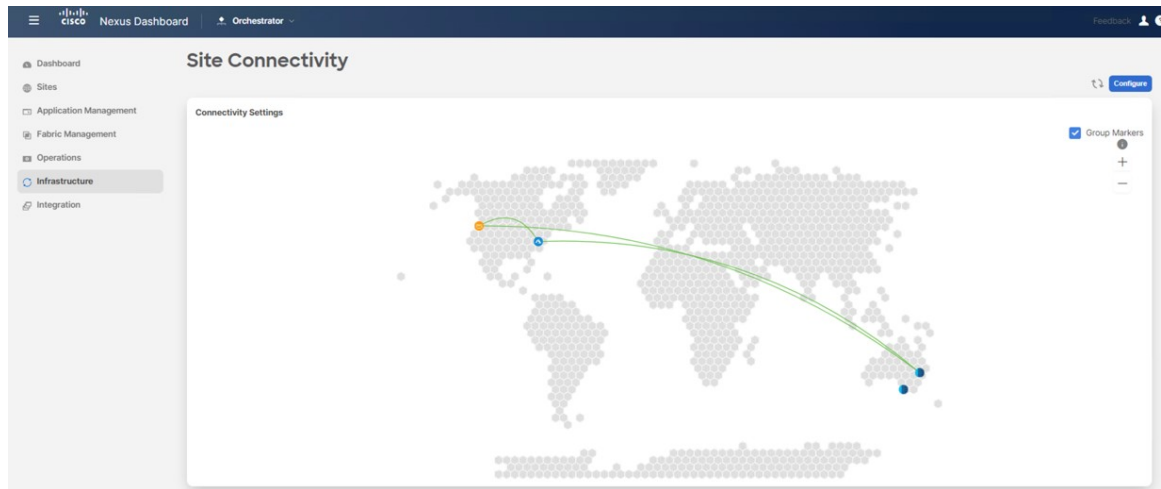
- Initiates communication with NDFC and the cloud sites (AWS and Azure) through the Cloud Network Controller to automate the IPsec tunnels.
- Configures OSPF between the Azure Catalyst 8000V and the AWS Catalyst 8000V.
- Configures eBGP between the BGW spine switch, the on-premises IPsec device, and the Azure Catalyst 8000V and the AWS Catalyst 8000V.
- Establishes BGP-EVPN peering sessions between the sites.

Step 3

Verify that the configurations were done correctly in NDO.

- In the left nav bar, click **Infrastructure** > **Site Connectivity** and verify the connectivity between sites in the **Connectivity Settings** area.

Figure 75:



- In the same page, scroll down to the area for the first cloud site (for example, the AWS site), click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

In this example, there are six IPsec tunnels because there are two Cisco Catalyst 8000Vs on the first cloud site (AWS) that have IPsec tunnels to two Cisco Catalyst 8000Vs on the second cloud site (Azure), and to one Cisco Catalyst 8000V for the on-premises external fabric.

Figure 76:

Device	Device Status	Interface Status	Peering Status	BGP Peer	Destination
ct_routerp_us-west-2_1	↑ Up	tunn-7 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_us-west-2_1	↑ Up	tunn-6 ↑ Up	BGP ↑ Up	170.1.254.6	64.104.255.12
ct_routerp_us-west-2_1	↑ Up	tunn-8 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_us-west-2_0	↑ Up	tunn-7 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_us-west-2_0	↑ Up	tunn-8 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_us-west-2_0	↑ Up	tunn-6 ↑ Up	BGP ↑ Up	170.1.254.2	64.104.255.12

- Scroll down to the area for the second cloud site (for example, the Azure site), click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

In this example, there are six IPsec tunnels because there are two Cisco Catalyst 8000Vs on the second cloud site (Azure) that have IPsec tunnels to two Cisco Catalyst 8000Vs on the first cloud site (AWS), and to one Cisco Catalyst 8000V for the on-premises external fabric.

Figure 77:

Device	Device Status	Interface Status	Peering Status	BGP Peer	Destination
ct_routerp_eastus_0	↑ Up	tunn-3 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_eastus_0	↑ Up	tunn-2 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_eastus_0	↑ Up	tunn-1 ↑ Up	BGP ↑ Up	170.1255.2	64.104.255.12
ct_routerp_eastus_3	↑ Up	tunn-2 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_eastus_3	↑ Up	tunn-3 ↑ Up	OSPF ↑ Up	-	-
ct_routerp_eastus_3	↑ Up	tunn-1 ↑ Up	BGP ↑ Up	170.1255.6	64.104.255.12

- Scroll down to the area for the NDFC external fabric site, click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

The external fabric's function is to provide underlay reachability from the on-premises IPsec devices to the VXLAN fabric and the cloud sites. The underlay protocol uses eBGP.

- Scroll down to the area for the NDFC VXLAN fabric site, click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

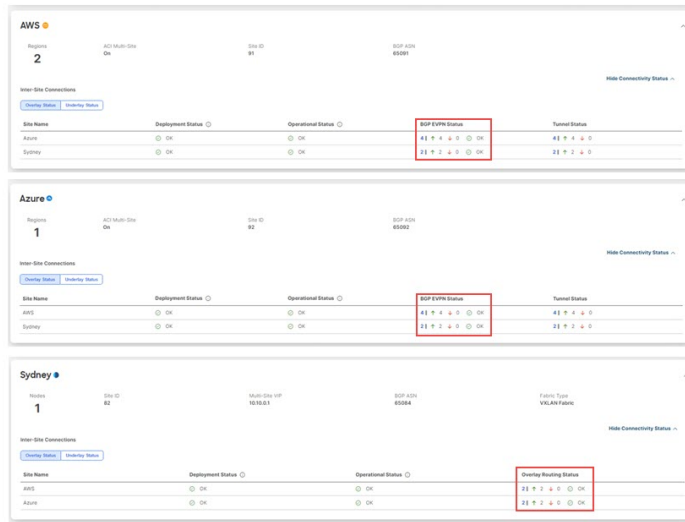
The underlay status shows the eBGP session status between the BGW spine switch and the on-premises IPsec device.

Figure 78:

Device	Device Status	Interface Status	Peering Status	BGP Peer
ndfc-spine1	↑ Up	Ethernet1/32 ↑ Up	BGP ↑ Up	10.140.1.1

- In each of those screens, click **Overlay Status** to verify the overlay status for each.

Figure 79:



- Return to the NDFC screen and verify the hybrid cloud connectivity in the **Topology** screen. In the following example, you can see the NDFC VXLAN fabric site (the Sydney site) connected to the first and second cloud sites (the AWS and Azure cloud sites).

Figure 80:

