Hybrid Cloud Connectivity Deployment for Cisco NX-OS

First Published: 2023-01-31

Last Modified: 2023-03-21

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: http://www.cisco.com/go/softwareterms.Cisco product warranty information is available at http://www.cisco.com/go/warranty. US Federal Communications Commission Notices are found here http://www.cisco.com/c/en/us/products/us-fcc-notice.html.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com go trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

Trademarks

CONTENTS

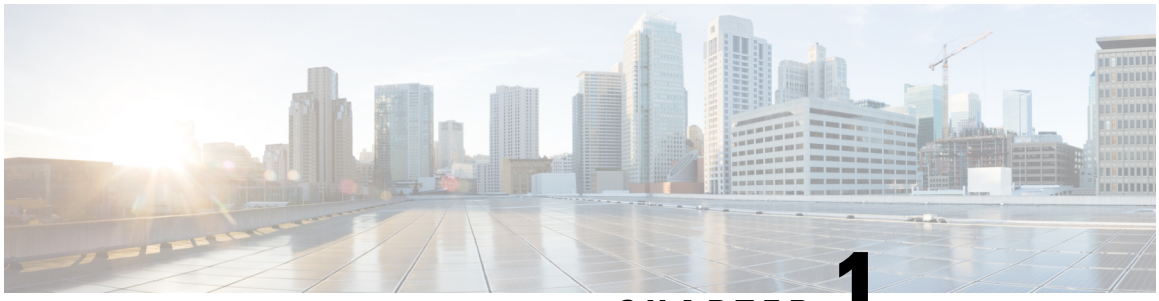

CHAPTER 1

New and Changed Information

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
NDFC release 12.1.2e	Initial release of this use case document.	Initial release of this use case document.

New and Changed Information

PART I

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

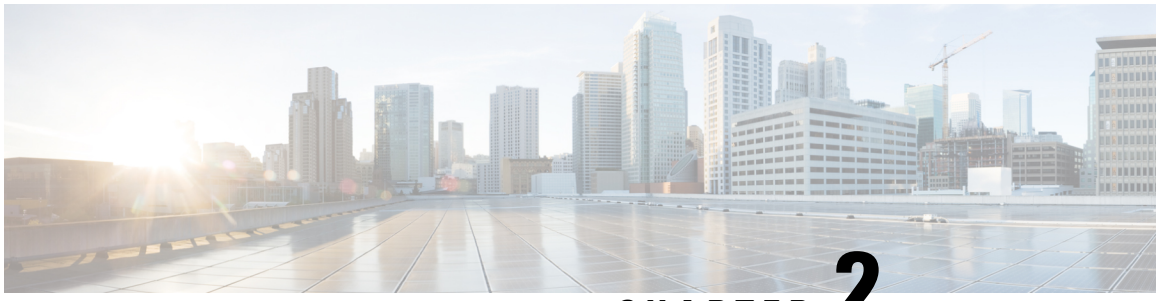

CHAPTER 2

Overview

Understanding Components of Hybrid Cloud Connectivity

This document describes deployment steps for the Cisco Hybrid Cloud Networking Solution powered by Cisco Nexus Dashboard Orchestrator (NDO) with a Cisco Nexus 9000 NX-OS based fabric managed by Nexus Dashboard Fabric Controller (NDFC) and public cloud sites managed by Cisco Cloud Network Controller (CNC).

The Cisco Nexus Dashboard Orchestrator (NDO) based Hybrid Cloud solution offers seamless connectivity between on-premises and cloud networks. This solution uses NDFC to manage on-premises VXLAN-based fabric and on-premises Cisco Catalyst 8000Vs, while cloud sites (AWS or Microsoft Azure) are managed by the Cisco Cloud Network Controller (CNC). NDO is used to orchestrate connectivity between on-premises and cloud sites, and between two or more cloud sites. VXLAN is used to build overlay tunnels between the sites.

The following figure shows an example topology for hybrid cloud connectivity using these components. See Supported Topologies, on page 13 for more information.

Figure 1:

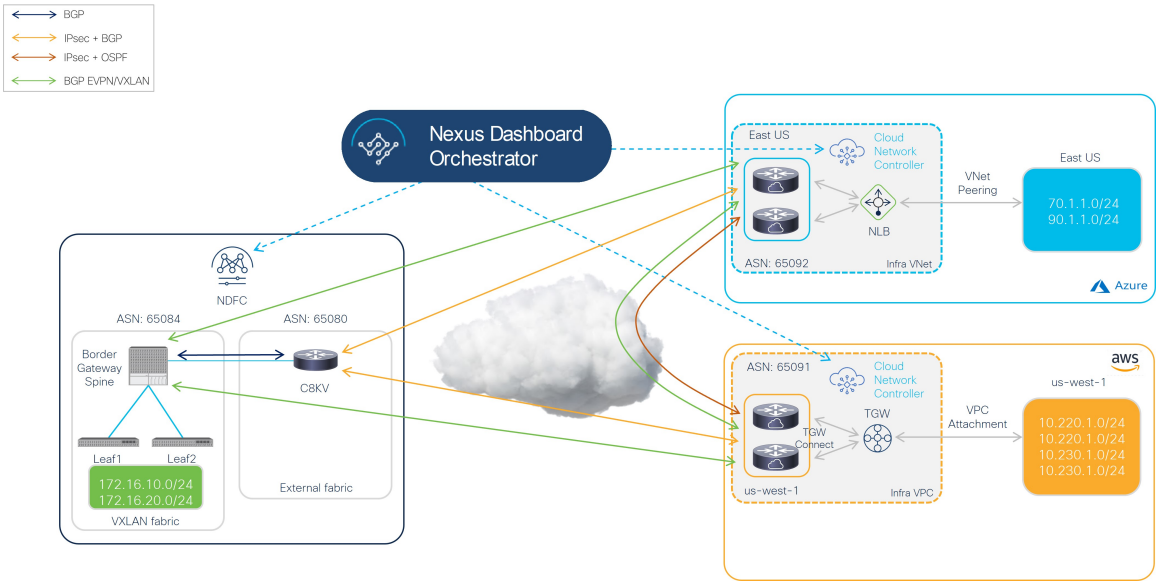

In this example topology, the on-premises site managed by NDFC has a secure connection setup to AWS and Azure cloud sites, where Cisco Catalyst 8000Vs sitting on the infra VPC/VNet serve as the cloud gateway for all traffic to and from the on-premises data centers.

On the on-premises site, Border Gateways (BGWs), which support seamless Layer-2/Layer-3 DCI extensions between different on-premises VXLAN EVPN sites, also support Layer-3 extension to the public cloud.

BGP-EVPN is used for the control plane between the BGWs and the Cisco Catalyst 8000Vs in the cloud, and VXLAN is used for the data plane.

As shown in the previous figure, the Cisco Hybrid Cloud Networking Solution consists of the following components:

- **Cisco Nexus Dashboard Orchestrator (NDO)**: NDO acts as a central policy controller, managing policies across multiple on-premises fabrics managed by different NDFC instances, with each cloud site being abstracted by its own Cisco Cloud Network Controller. NDO runs as a service on top of Nexus Dashboard, where Nexus Dashboard can be deployed as a cluster of physical appliances or virtual machines running on VMware ESXi, Linux KVM, Amazon Web Services or Microsoft Azure. Inter-version support was introduced previously, so NDO can manage Cisco Cloud Network Controller running different software versions.
- **Cisco Nexus Dashboard Fabric Controller (NDFC)**: NDFC is a network automation and orchestration tool for building LAN, VXLAN, SAN and Cisco IP Fabric for Media (IPFM) fabrics. NDFC runs as a service on top of Nexus Dashboard cluster that can be either a physical or a virtual cluster. For the Hybrid Cloud Networking Solution, NDFC manages the on-premises VXLAN fabric and on-premises Cisco Cloud Routers (Catalyst 8000V).
- **On-premises VXLAN fabric**: The on-premises VXLAN fabric is built with Nexus 90000/3000 switches managed by NDFC. The fabric should have one or more Border Gateway (BGW) devices that are responsible for originating and terminating VXLAN Multisite Overlay tunnels between on-premises and cloud sites. NDFC has pre-built templates for creating a VXLAN fabric; this document uses the `External_Fabric` template for the VXLAN fabric.

- **On-premises Cisco Cloud Router (CCR)**: The CCR is used to provide reachability between the on-premises VXLAN fabric and the cloud sites. The CCR provides connectivity to the cloud sites using either public internet or private connections (such as AWS Direct Connect or Azure ExpressRoute). The on-premises CCRs are managed by NDFC using a pre-built `External_Fabric` template and need to be assigned the `Core Router` role.

  The Cisco Catalyst 8000V is used as the on-premises CCR for the Cisco Hybrid Cloud Networking Solution.

- **Cisco Cloud Network Controller (CNC)**: Cisco Cloud Network Controller runs as a virtual instance on a supported public cloud to provide automated connectivity, policy translation, and enhanced visibility of workloads in the public cloud. The Cisco Cloud Network Controller translates all the policies received from NDO and programs them into cloud-native constructs, such as VPCs and security groups on AWS, and VNets on Microsoft Azure. Cisco Cloud Network Controller is deployed through the public cloud Marketplace, such as AWS Marketplace and Azure Marketplace.

- **Cisco Catalyst 8000V**: The Cisco Catalyst 8000V is an important component in the public cloud platforms. Cisco Catalyst 8000Vs are used for inter-site communication to on-premises sites and the public cloud platforms. In addition, Cisco Catalyst 8000Vs are used for on-premises cloud connectivity and for connectivity between different cloud providers (for example, Azure to AWS).

Building Hybrid Cloud Connectivity

This section describes the process used to build hybrid cloud connectivity.

Starting Point

The following figure shows the starting point for the hybrid cloud connectivity, where we have the various pieces described in Understanding Components of Hybrid Cloud Connectivity, on page 5:

- Nexus Dashboard Fabric Controller (NDFC) fabrics:
  - On-premises VXLAN fabric
  - External fabric
- Cloud sites (AWS and Azure) managed by Cloud Network Controller

Figure 2:

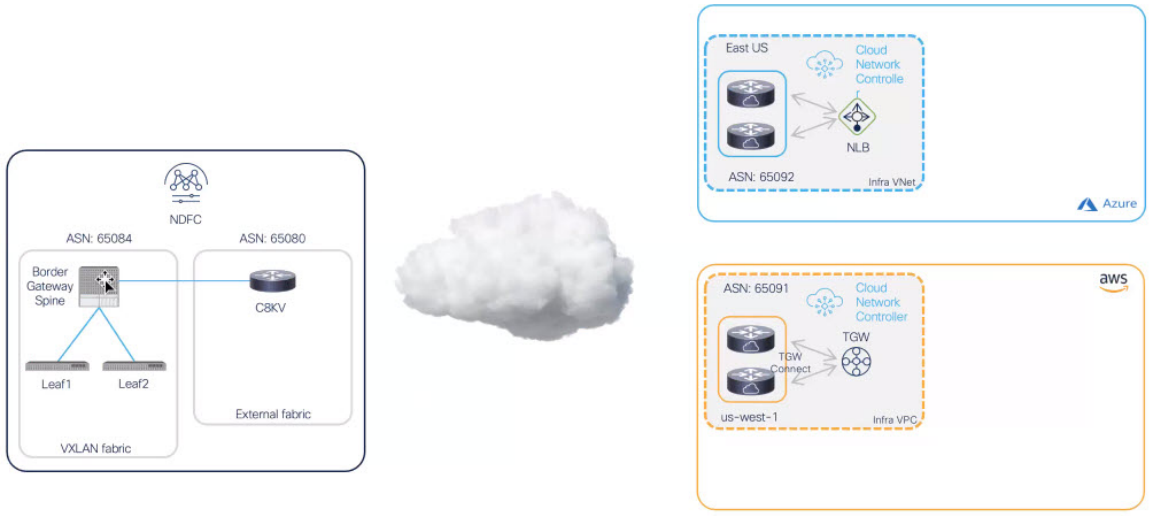

Building the Underlay Layer

Next, we will show how the underlay later is built:

1. First, a BGP connection is established between the border gateway spine switch in the VXLAN fabric and the Cisco Catalyst 8000V in the external fabric.
2. Then, BGP peering is used to establish the underlay connectivity between the on-premises Cisco Catalyst 8000V in the external fabric to each of the cloud routers in the cloud sites.
3. Finally, OSPF is used between the cloud sites for cloud-to-cloud underlay connectivity.

Figure 3:

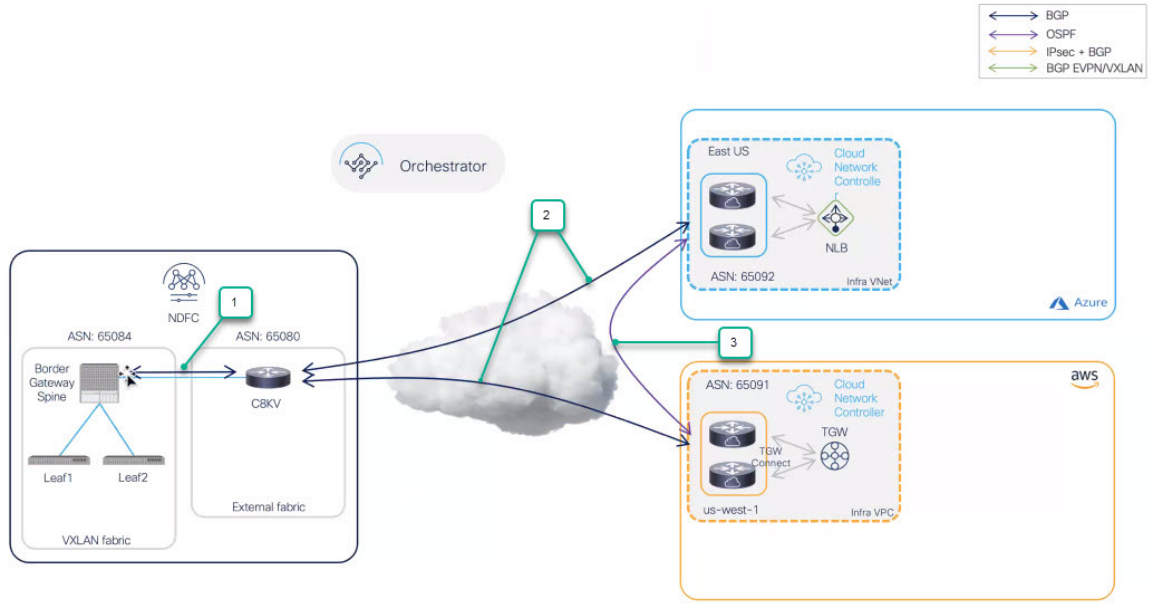

Building Overlay

Finally, we show how to establish the VXLAN Multisite Overlay on top of underlay connectivity established in previous step:

1. A VXLAN multi-site is established, which originates from the border gateway spine switch in the VXLAN fabric and terminates at the Cisco Catalyst 8000Vs in the cloud sites.
2. If you select Public Internet as the connection type, then IPsec and BGP are used to connect between the NDFC VXLAN fabric site and the cloud sites.

Figure 4:

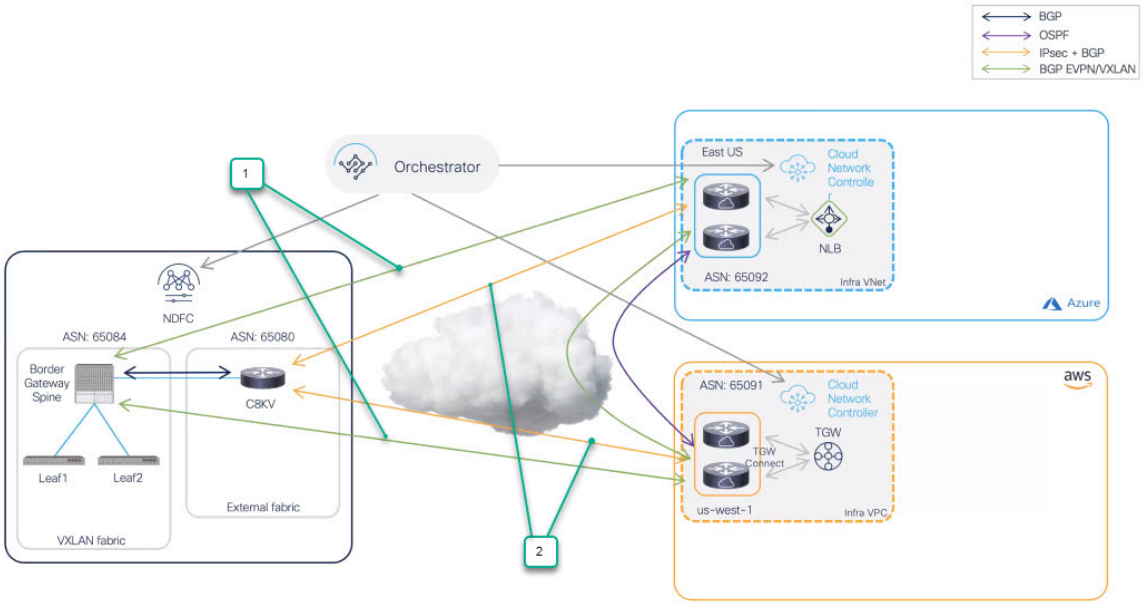

Terminology

The following terms are used throughout this document.

Term	Acronym	Definition
Border Gateway	BGW	One of the supported switch roles in an NDFC Easy Fabric (for example, a VXLAN EVPN fabric). The BGW is used to extend Layer 2/Layer 3 DCI connectivity between on-premises fabrics and Layer 3 connectivity toward public cloud sites (for example, hybrid cloud connectivity).

Term	Acronym	Definition
Core Router		One of the supported roles in an NDFC external fabric. The core router is used to establish Layer 3 connectivity (Underlay) on one side with the VXLAN EVPN fabric, and on the other with the Catalyst 8000Vs in cloud sites.
Direct Connect		Used in the AWS cloud. AWS Direct Connect is a cloud service that links your network directly to AWS to deliver consistent, low-latency performance.
ExpressRoute		Used in the Azure cloud. You can use Azure ExpressRoute to create private connections between Azure datacenters and infrastructure on premises or in a co-location environment.
Inter-Site Network	ISN	The Layer 3 infrastructure used to interconnect on-premises VXLAN fabrics, between the on-premises VXLAN fabrics and with the public cloud (also referred to as the "underlay"). As such, the ISN could also include the Internet or the Direct Connect and ExpressRoute dedicated circuits.
IP Security Router	IPsec router	A router capable of Internet Protocol Security (IPsec) is required to establish IPsec connections between the on-premises site and the cloud sites Cisco Cloud Network Controller.

Term	Acronym	Definition
Route Server	RS	The control plane node used to facilitate the establishment of EVPN adjacencies between on-premises BGW devices, alleviating the need of creating full-mesh peering between all of them. The Route Server runs BGP protocol and is used to pass routes between two or more BGP peers. The Route Server function is the eBGP equivalent of the "Route Reflector" function traditionally used for iBGP sessions; it helps in reducing the number of BGP peering required.
Virtual Network	VNet	Used in the Azure cloud. Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many types of Azure resources, such as Azure Virtual Machines (VMs), to securely communicate with each other, the internet, and on-premises networks. As related to the Cloud Network Controller, the VRF in the Cloud Network Controller maps to a VNet in Azure.
Virtual Private Cloud	VPC	Used in the AWS cloud. Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS. As related to the Cloud Network Controller, the VRF in the Cloud Network Controller maps to a VPC in AWS.

Prerequisites

The following software versions are required:

- Cisco Nexus Dashboard (ND) version 2.3.1c or later (physical or virtual cluster)
- Cisco Nexus Dashboard Fabric Controller (NDFC) version 12.1.2e or later
- Cisco Nexus Dashboard Orchestrator (NDO) version 4.1(1) or later
- Cisco Cloud Network Controller (CNC) version 25.1(1e) or later for AWS site and Microsoft Azure site

Guidelines and Limitations

Following are certain guidelines and limitations that you should understand when deploying the hybrid cloud connectivity solution:

- Currently, each Cisco Cloud Network Controller can manage up to sixteen regions in AWS and Azure clouds. If you want to manage more than sixteen regions, you will have to deploy additional Cisco Cloud Network Controllers. For more information, see the "Understanding Limitations for Number of Sites, Regions and CCRs" section in the *Cisco Cloud Network Controller for AWS Installation Guide* or *Cisco Cloud Network Controller for Azure Installation Guide*, Release 25.1(x) or later.

Related Documentation

You can find documentation for the components that make up the Cisco Hybrid Cloud Networking Solution in the following locations:

- Cisco Nexus Dashboard Orchestrator (NDO) documentation
- Cisco Nexus Dashboard Fabric Controller (NDFC) documentation
- Cisco Cloud Network Controller (CNC) documentation
- Cisco Catalyst 8000V documentation
- Amazon Web Services (AWS) documentation
- Microsoft Azure documentation

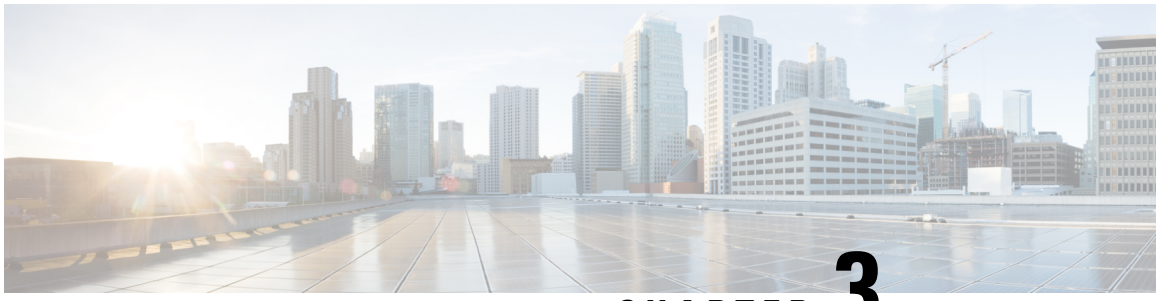

CHAPTER 3

Supported Topologies

Connection Options

You can use these connection options for the Cisco Hybrid Cloud Networking Solution:

- **With IPsec**: If the connectivity from the on-premises data center to the cloud is over the public Internet, then an IPsec tunnel is required for establishing a secure channel. In this situation, the border gateway (BGW) will be connected to an on-premises IPsec-capable device, such as an ASR 1000 or a Cisco Catalyst 8000V. This device establishes IPsec tunnels with the Catalyst 8000Vs in the cloud. The on-premises BGWs can then leverage this "IPsec secured underlay" to build VXLAN tunnels with the Catalyst 8000Vs in the cloud.
- **Without IPsec**: If the BGWs are connected to the public cloud using Direct Connect (AWS) or ExpressRoute (Azure), then enabling IPsec is optional. In this case, a VXLAN connection is employed between the on-premises VXLAN EVPN data centers and the Cisco Catalyst 8000Vs on top of those dedicated circuits.

The following sections provide more detailed information on the supported topologies available using either of these connection options:

Supported Topologies with IPsec (Single-Cloud)

The following table shows how BGP EVPN control plane adjacencies can be established between on-premises sites and on-premises to a cloud site, and how IPsec is leveraged to establish underlay connectivity between on-premises sites and a single cloud site.

Note Each of the following figures show a simple example. In a real life scenario, there might be redundant devices deployed for each role.

BGP EVPN Between On-Premises Sites	**BGP EVPN and IPsec to the Cloud Site**		
	Full-Mesh	**Through Hub Site Only**	• **BGP EVPN to the Cloud Site: Full-Mesh** • **IPsec to the Cloud Site: Through Shared IPsec Router Only**
Full-Mesh	Option 1, on page 14	Option 3, on page 16	Option 5, on page 17
With Route Server	Option 2, on page 15	Option 4, on page 17	N/A

Option 1

The following figure shows an example of a single-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud site establishes IPsec tunnels with core routers deployed in each on-premises site and full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.

Figure 5:

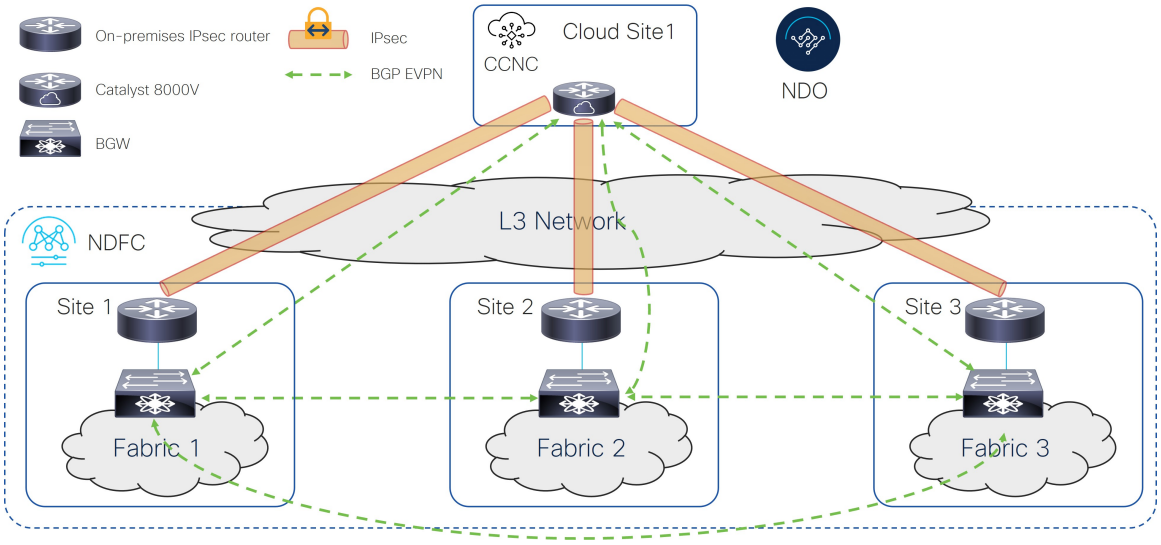

Option 2

The following figure shows an example of a single-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server (RS) control plane node.
- The Cisco Catalyst 8000V in the cloud site establishes full-mesh IPsec tunnels with core routers deployed in each on-premises site and BGP EVPN adjacencies with all the BGW devices on the on-premises sites.

Note It is currently not supported to peer the Cisco Catalyst 8000Vs with the Route Server control node.

Figure 6:

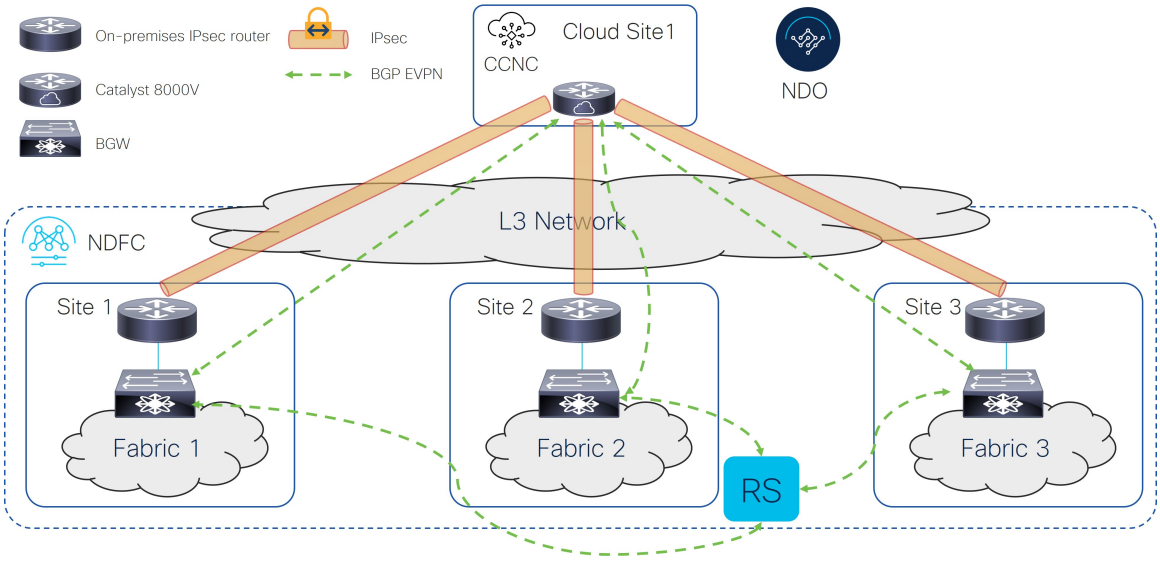

Option 3

The following figure shows an example of a single-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud site establishes an IPsec tunnel only with the core router deployed in a specific on-premises Hub Site and BGP EVPN adjacency only with the BGW device on the Hub Site.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and the cloud site.

Figure 7:

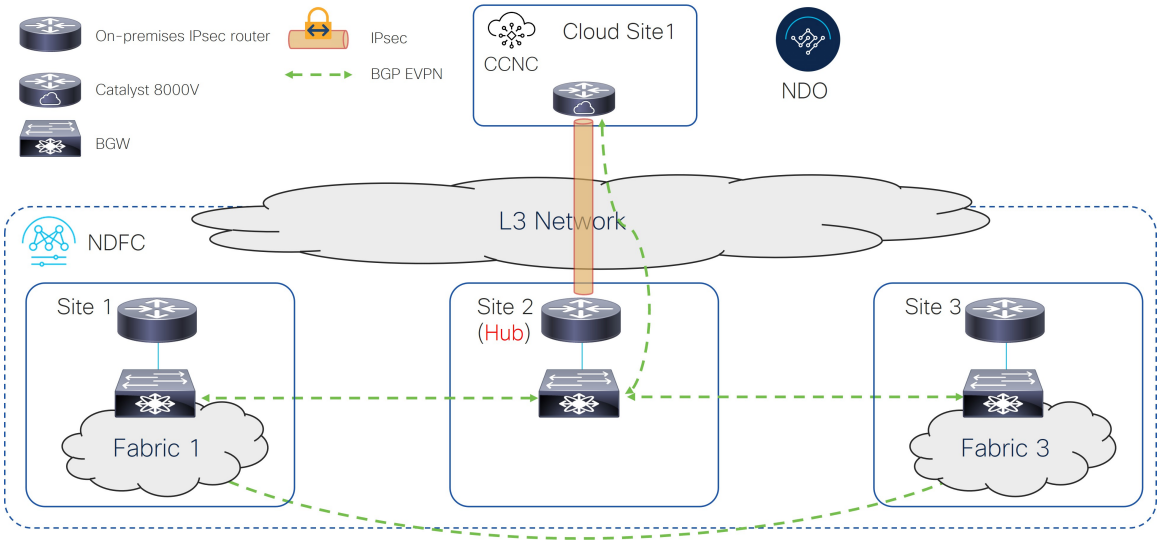

Option 4

The following figure shows an example of a single-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000V in the cloud site establishes an IPsec tunnel only with the core router deployed in a specific on-premises Hub Site and EVPN adjacency only with the BGW device on the Hub Site.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and the cloud site.

Figure 8:

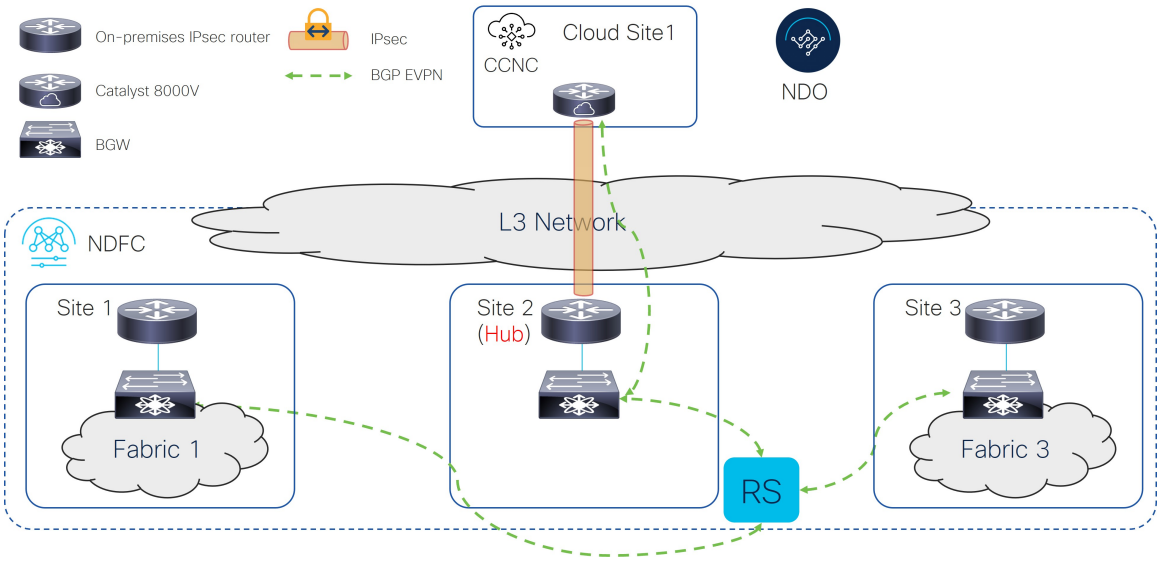

Option 5

The following figure shows an example of a single-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud site establishes full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.
- The IPsec connection to the cloud site is through a shared IPsec router only.

Figure 9:

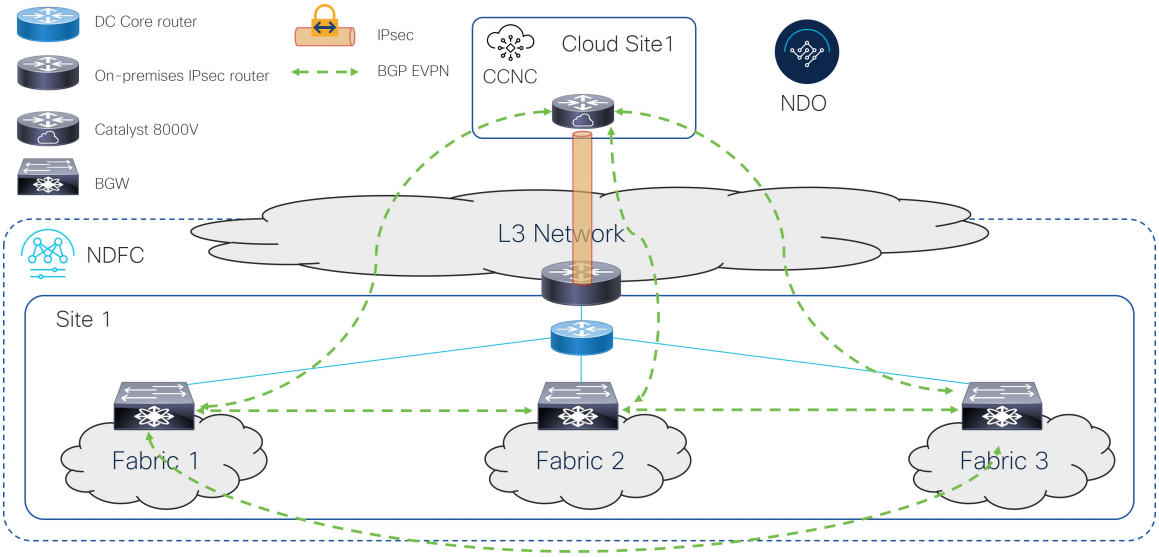

Supported Topologies with IPsec (Multi-Cloud)

The following table shows how BGP EVPN control plane adjacencies can be established between on-premises sites and on-premises to cloud sites, and how IPsec is leveraged to establish underlay connectivity between on-premises sites and multiple cloud sites.

Note Each of the following figures show a simple example. In a real life scenario, there might be redundant devices deployed for each role.

BGP EVPN Between On-Premises Sites	BGP EVPN and IPsec to the Cloud Sites			BGP EVPN and IPsec between Cloud Sites
	Full-Mesh	**Through Hub Site Only**	• **BGP EVPN to the Cloud Site: Full-Mesh** • **IPsec to the Cloud Site: Through Hub Site Only**	
Full-Mesh	Option 1, on page 19	Option 3, on page 20	Option 5, on page 22	**Full-Mesh**
With Route Server	Option 2, on page 19	Option 4, on page 21	N/A	

Option 1

The following figure shows an example of a multi-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000Vs in the cloud sites establish IPsec tunnels with core routers deployed in each on-premises site and full-mesh EVPN adjacencies with all the BGW devices on the on-premises sites.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh IPsec tunnels and EVPN adjacencies between them.

Figure 10:

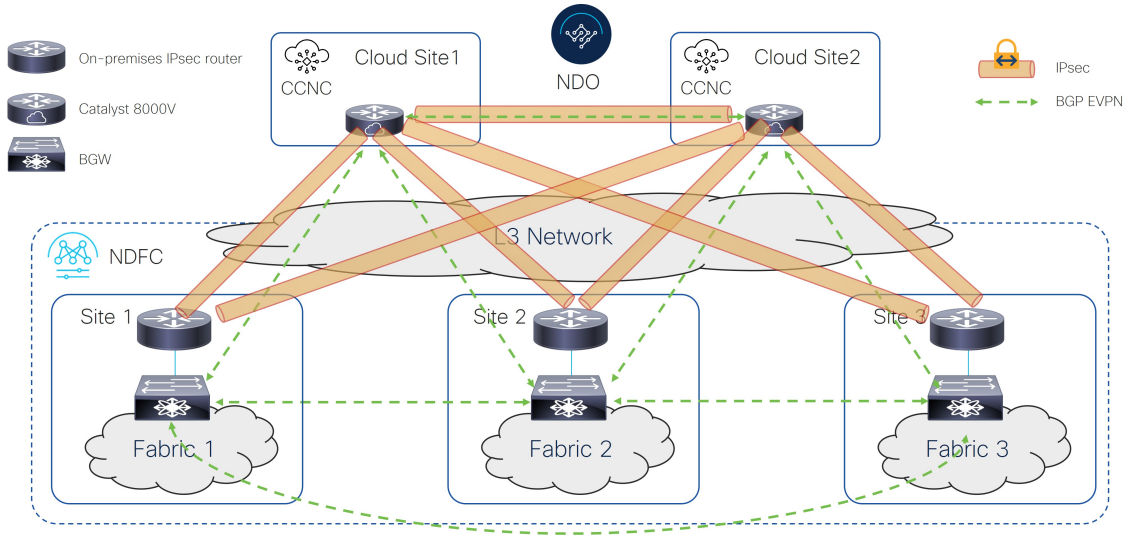

Option 2

The following figure shows an example of a multi-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000Vs in the cloud sites establish IPsec tunnels with core routers deployed in each on-premises site and full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.
- The cloud routers peer BGP EVPN with the BGW on the Hub Site.

Figure 11:

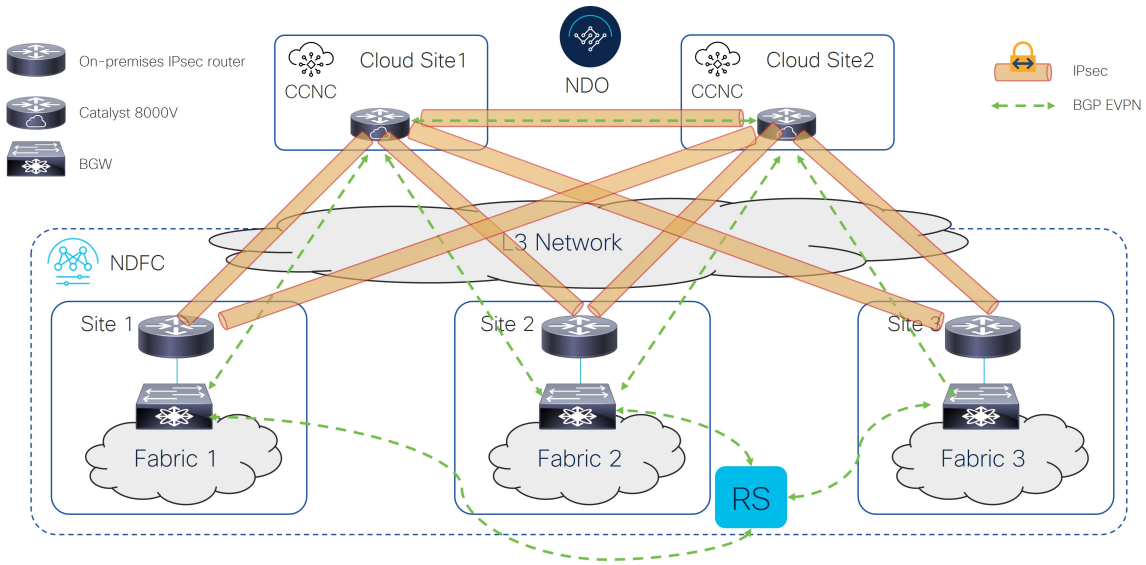

Option 3

The following figure shows an example of a multi-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh EVPN adjacencies between them.
- The Cisco Catalyst 8000Vs in the cloud sites establish IPsec tunnels only with the core router deployed in a specific on-premises Hub Site and EVPN adjacency only with the BGW device on the Hub Site.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh IPsec tunnels and EVPN adjacencies between them.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and cloud sites.

Figure 12:

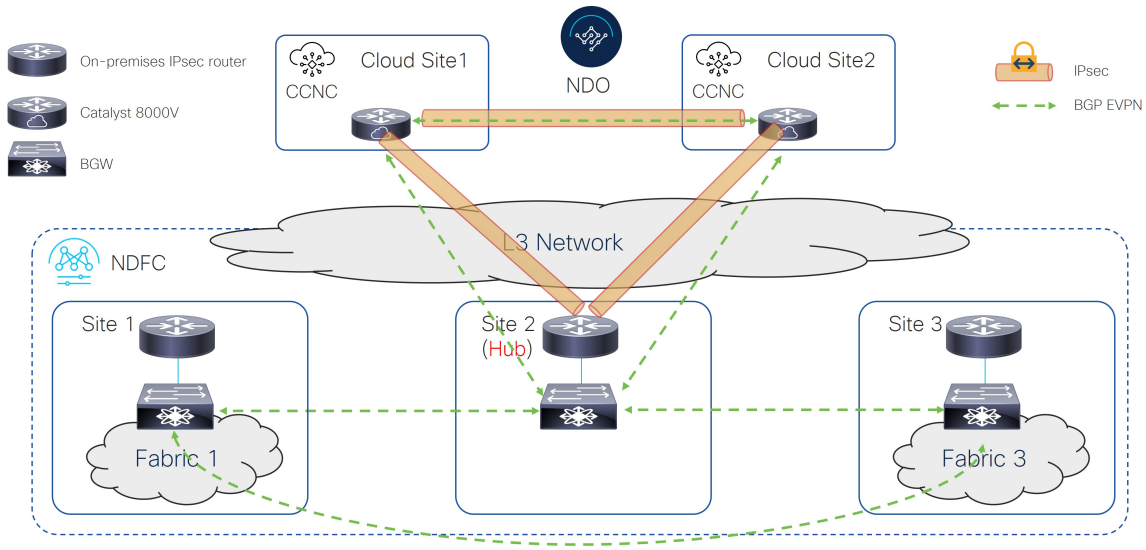

Option 4

The following figure shows an example of a multi-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000Vs in the cloud sites establish IPsec tunnels only with the core router deployed in a specific on-premises Hub Site and BGP EVPN adjacency only with the BGW device on the Hub Site.
- The cloud routers peer BGP EVPN with the BGW on the Hub Site.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and cloud sites.

Figure 13:

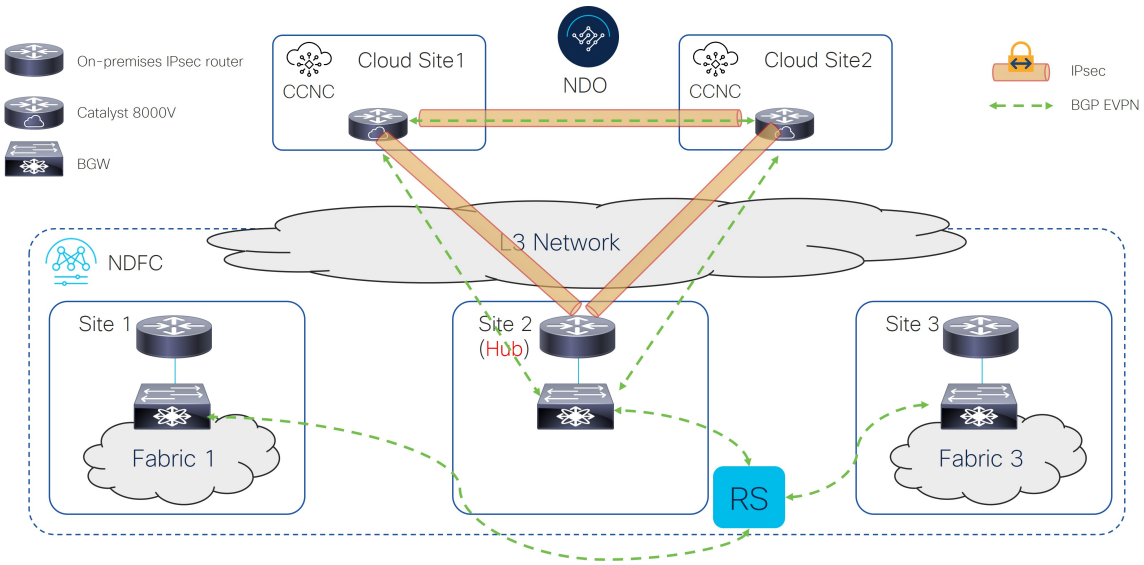

Option 5

The following figure shows an example of a multi-cloud connection using IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud sites establishes full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.
- The Cisco Catalyst 8000Vs in the cloud sites establish IPsec tunnels only with the core router deployed in a specific on-premises Hub Site.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh IPsec tunnels and EVPN adjacencies between them.

Figure 14:

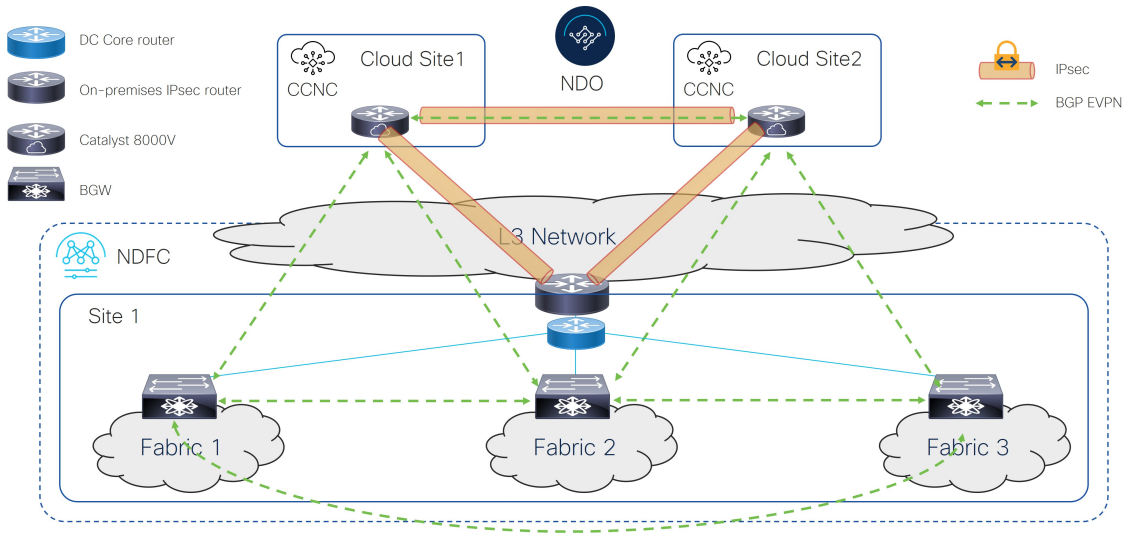

Supported Topologies without IPsec (Single Cloud)

The following table shows how BGP EVPN control plane adjacencies can be established between on-premises sites or on-premises to a cloud site.

BGP EVPN Between On-Premises Sites	BGP EVPN to the Cloud Site	
	Full-Mesh	**Through Hub Site**
Full-Mesh	Option 1, on page 23	Option 3, on page 24
With Route Server	Option 2, on page 24	Option 4, on page 25

Note Each of the following figures show a simple example. In a real life scenario, there might be redundant devices deployed for each role.

Option 1

The following figure shows an example of a single-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud site establishes full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.

Figure 15:

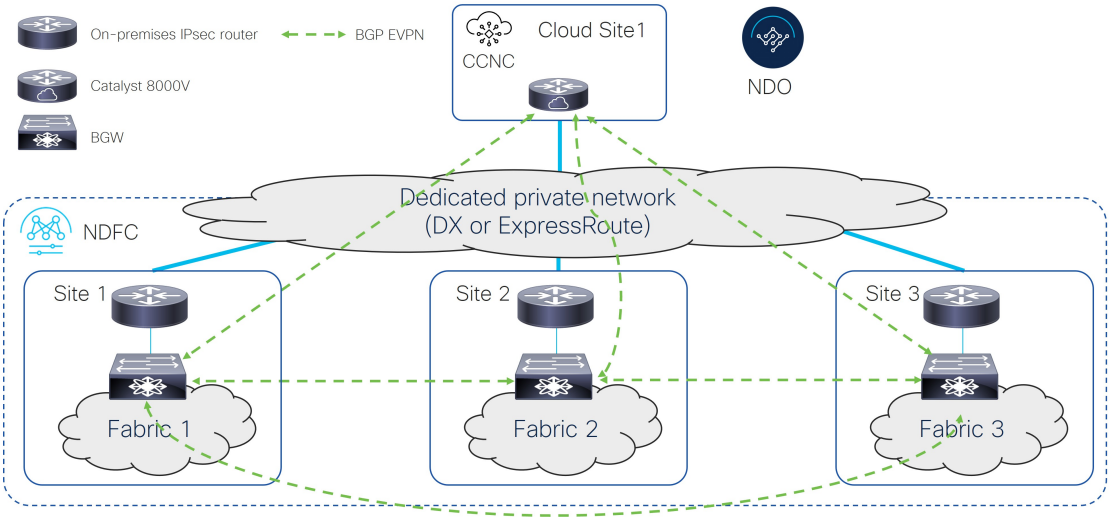

Option 2

The following figure shows an example of a single-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server (RS) control plane node.
- The Cisco Catalyst 8000V in the cloud site establishes full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.

Figure 16:

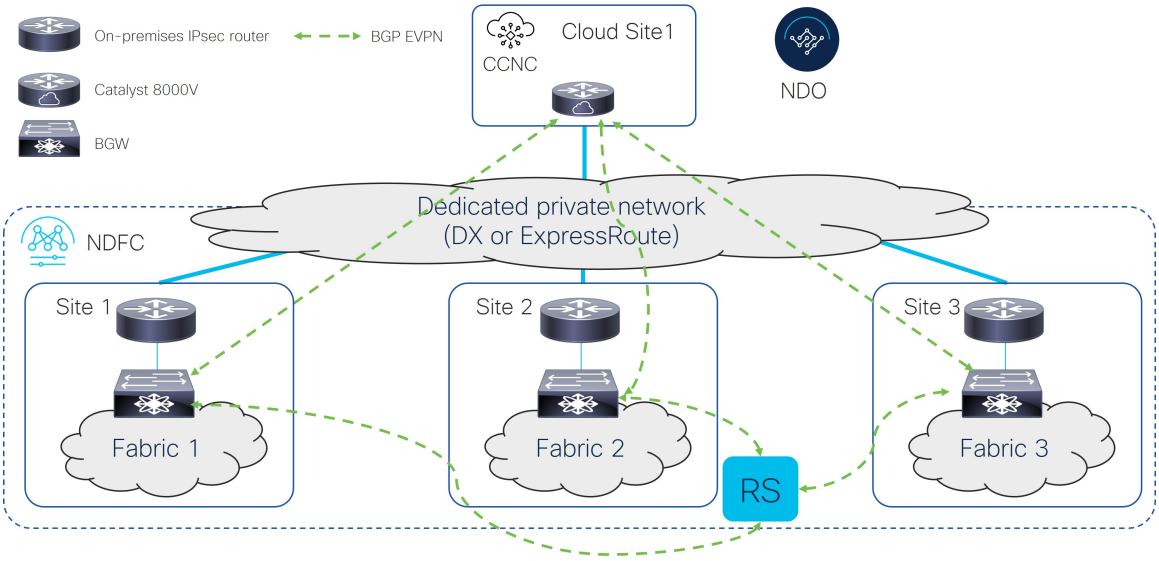

Option 3

The following figure shows an example of a single-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000V in the cloud site establishes a BGP EVPN adjacency only with the BGW device on the Hub Site.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and the cloud site.

Figure 17:

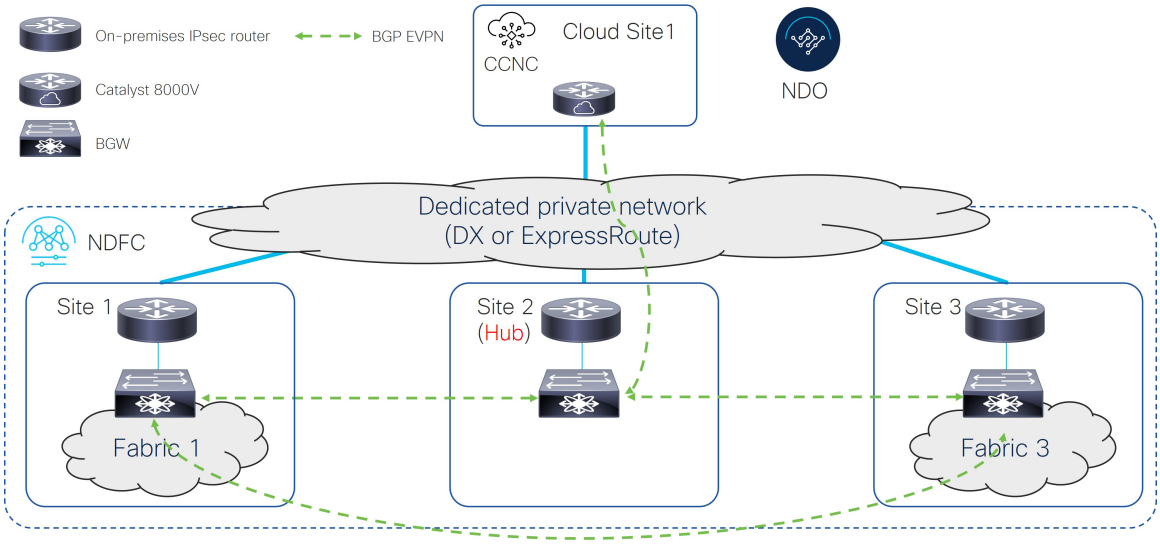

Option 4

The following figure shows an example of a single-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000V in the cloud site establishes a BGP EVPN adjacency only with the BGW device on the Hub Site.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and the cloud site.

Figure 18:

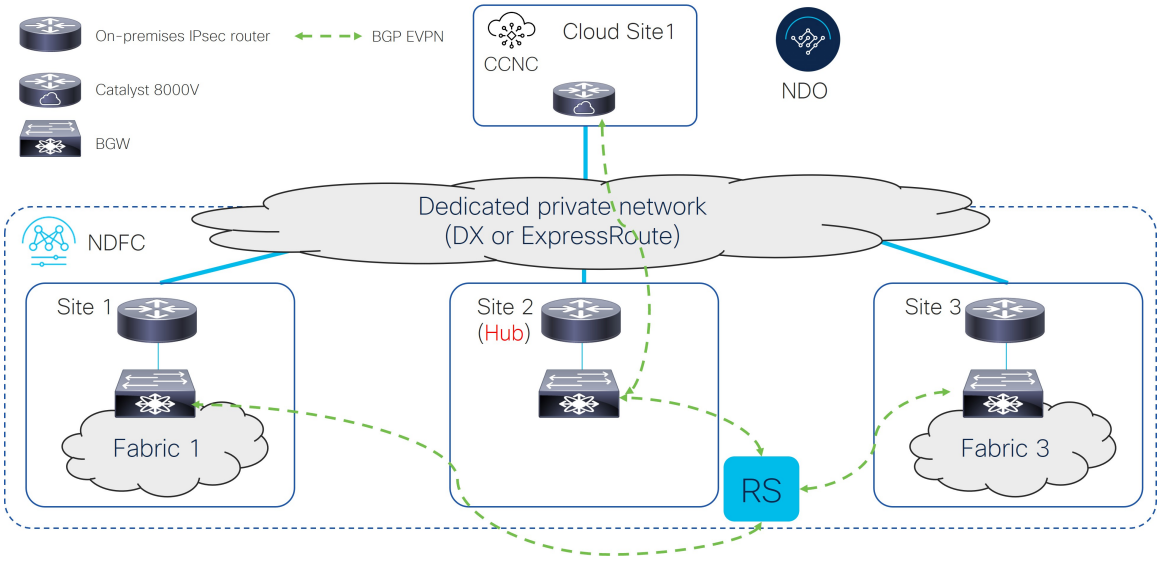

Supported Topologies without IPsec (Multi-Cloud)

The following table shows how BGP EVPN control plane adjacencies can be established between on-premises sites or on-premises to cloud sites.

BGP EVPN Between On-Premises Sites	BGP EVPN to the Cloud Sites		BGP EVPN between Cloud Sites
	Full-Mesh	**Through Hub Site**	
Full-Mesh	Option 1, on page 26	Option 3, on page 28	**Full-Mesh**
Route Server	Option 2, on page 27	Option 4, on page 28	

Note Each of the following figures show a simple example. In a real life scenario, there might be redundant devices deployed for each role.

Option 1

The following figure shows an example of a multi-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000Vs in the cloud sites establish full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh BGP EVPN adjacencies between them.

Figure 19:

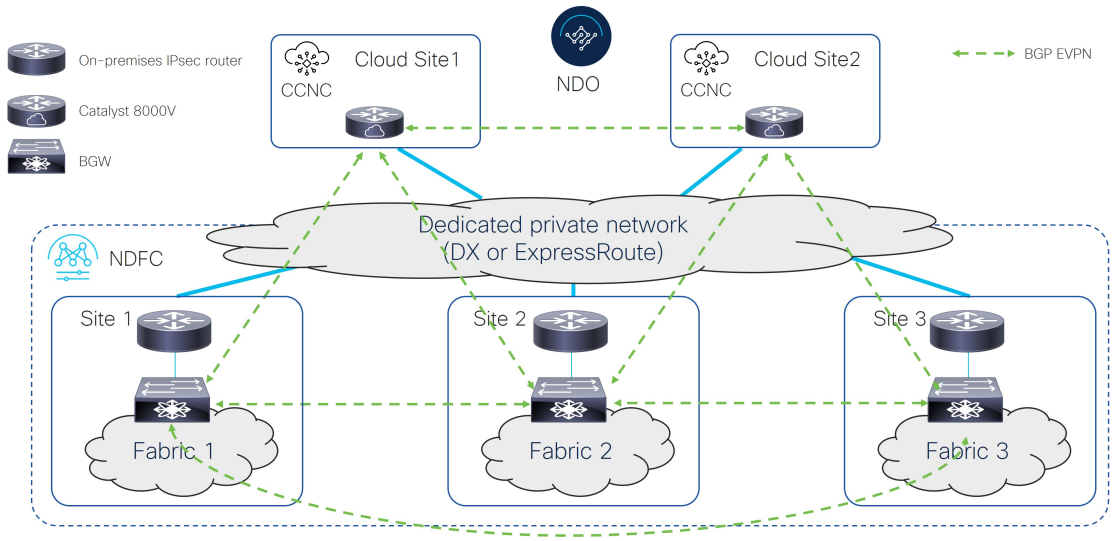

Option 2

The following figure shows an example of a multi-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000Vs in the cloud sites establish full-mesh BGP EVPN adjacencies with all the BGW devices on the on-premises sites.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh BGP EVPN adjacencies between them.

Figure 20:

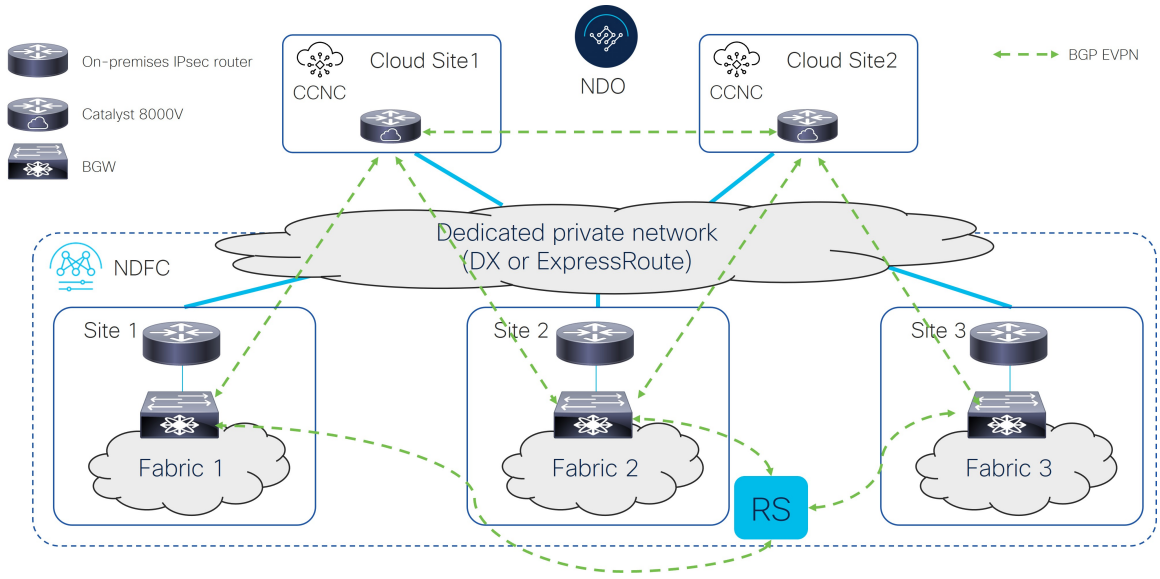

Option 3

The following figure shows an example of a multi-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish full-mesh BGP EVPN adjacencies between them.
- The Cisco Catalyst 8000Vs in the cloud sites establish BGP EVPN adjacencies only with the BGW device on the Hub Site.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh BGP EVPN adjacencies between them.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and cloud sites.

Figure 21:

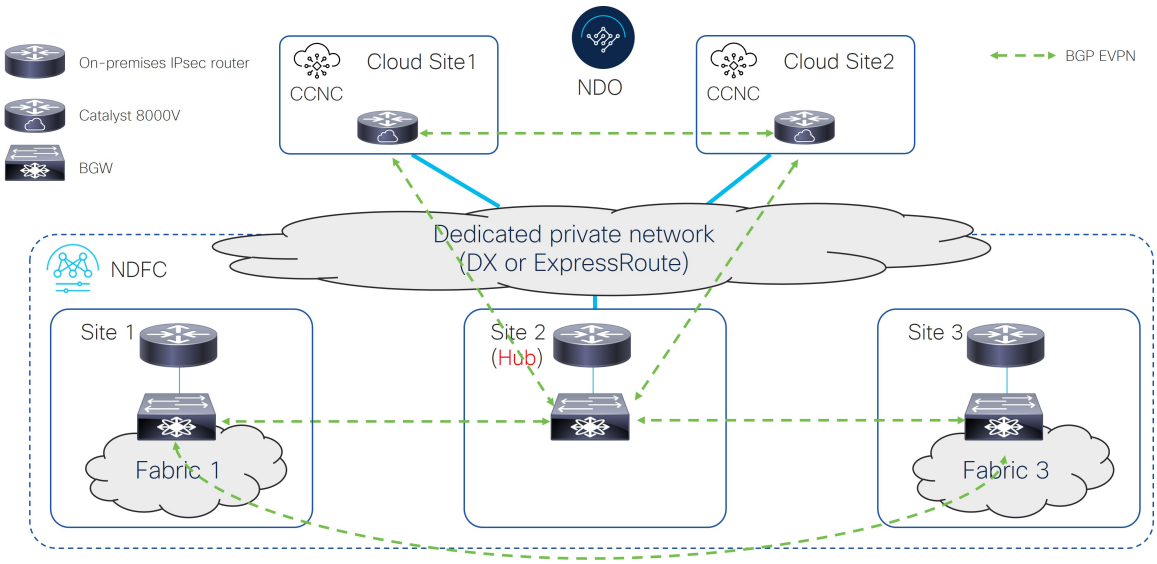

Option 4

The following figure shows an example of a multi-cloud connection without IPsec, where:

- The BGW nodes on all the on-premises sites establish BGP EVPN adjacencies with a Route Server control plane node.
- The Cisco Catalyst 8000Vs in the cloud sites establish BGP EVPN adjacencies only with the BGW device on the Hub Site.
- The Cisco Catalyst 8000Vs in different cloud sites establish full-mesh BGP EVPN adjacencies between them.
- The BGW deployed in Site 2 (to which the Cisco Catalyst 8000V peers EVPN) cannot have a fabric behind it. It is only used to exchange prefixes between the on-premises and cloud sites.

Figure 22:

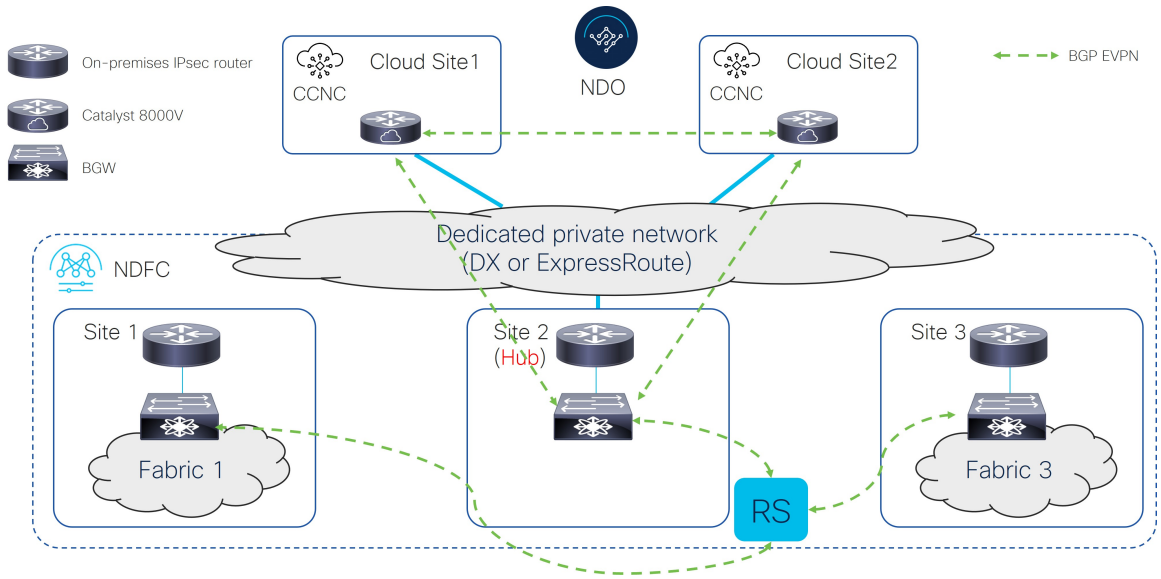
On-premises IPsec router
Catalyst 8000V
BGW
Cloud Site1
CCNC
NDO
Cloud Site2
CCNC
BGP EVPN
Dedicated private network
(DX or ExpressRoute)
NDFC
Site 1
Fabric 1
Site 2
(Hub)
RS
Site 3
Fabric 3

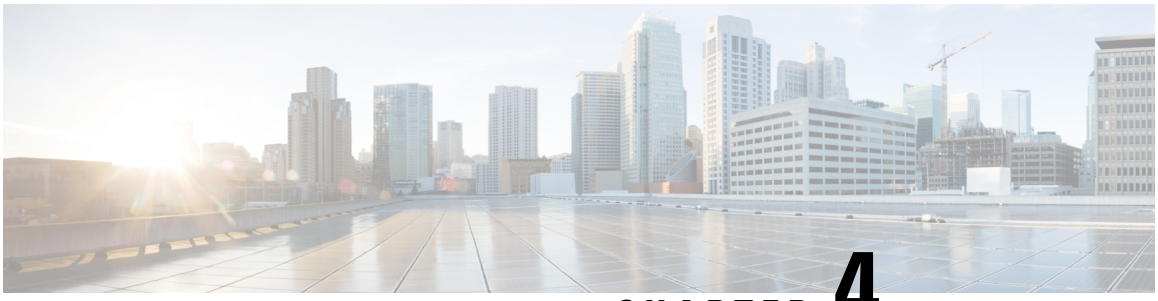

CHAPTER 4

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

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 23:

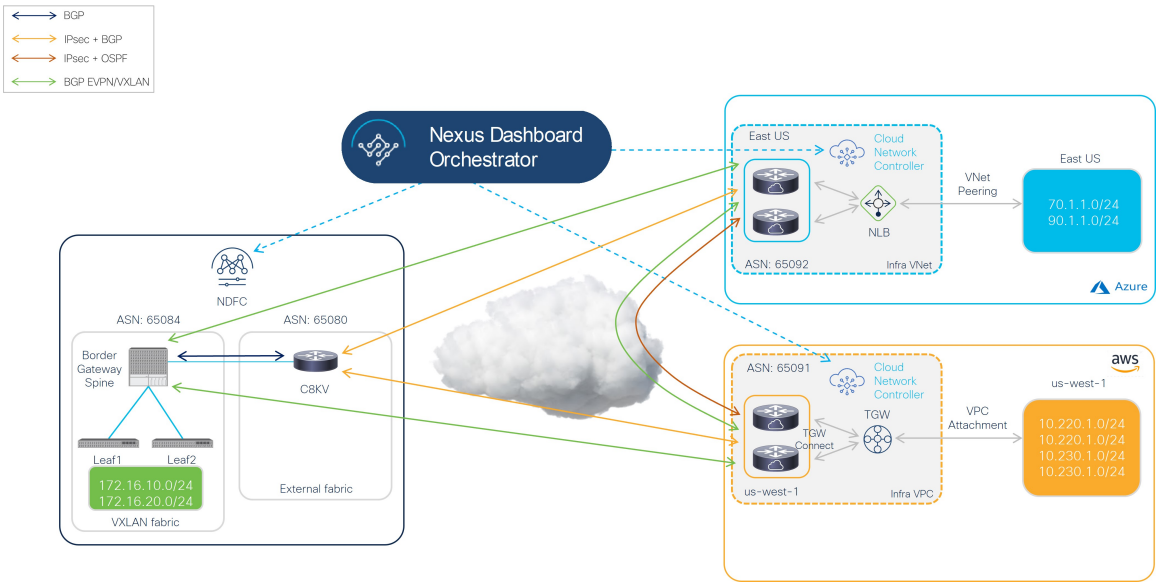

The procedures in this document will use this topology as a specific use case, which is based on Option 1, on page 19 in Supported Topologies with IPsec (Multi-Cloud), on page 18, and will describe how to configure the hybrid cloud connectivity options specifically for this topology use case.

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 24:

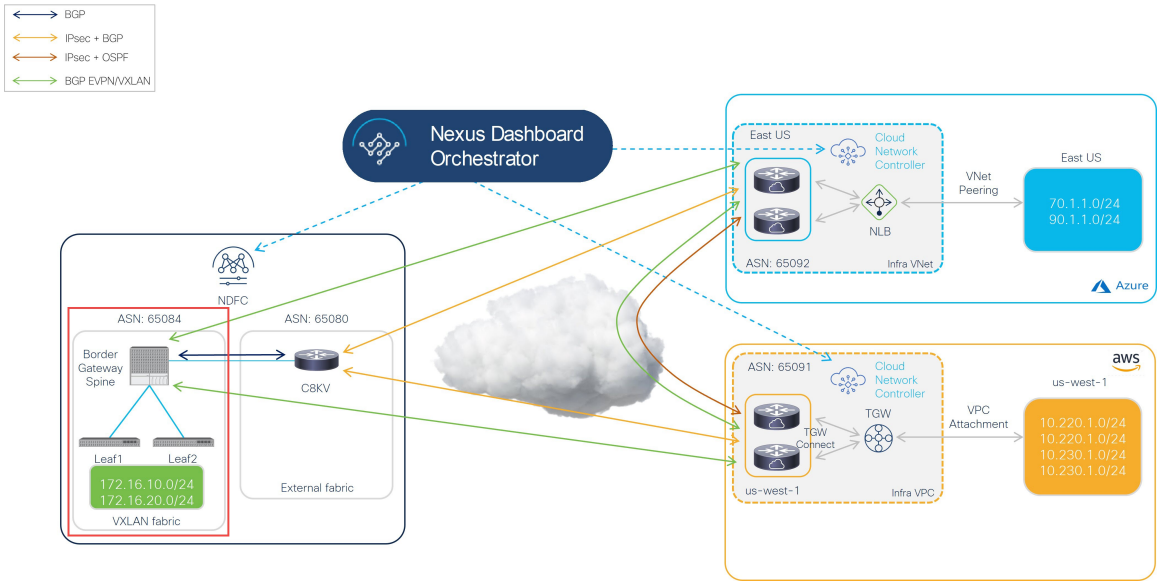

The VXLAN fabric must contain one or more Border Gateway (BGW) devices, which are used to build VXLAN Multi-Site connectivity between on-premises fabrics and the cloud sites.

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 25:

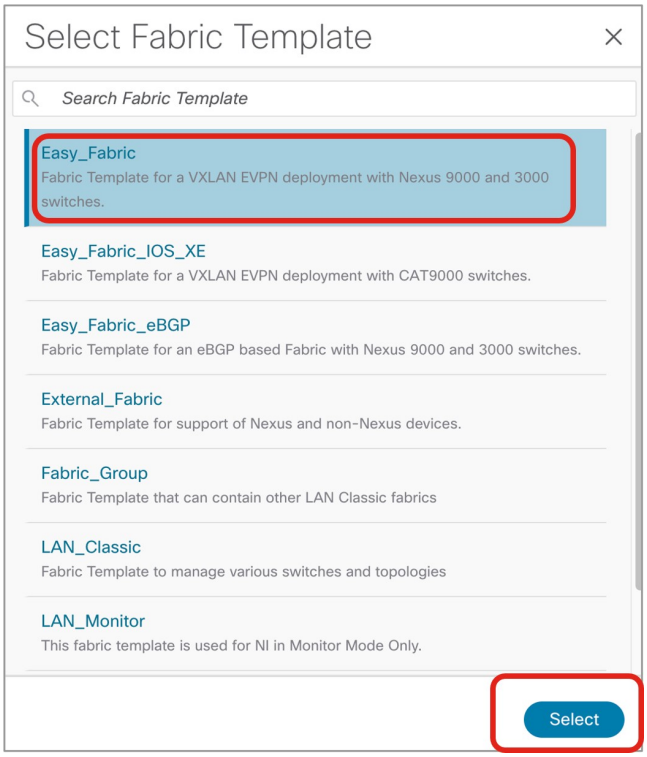

Step 6 Complete the necessary general VXLAN fabric parameter configurations.

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 26:

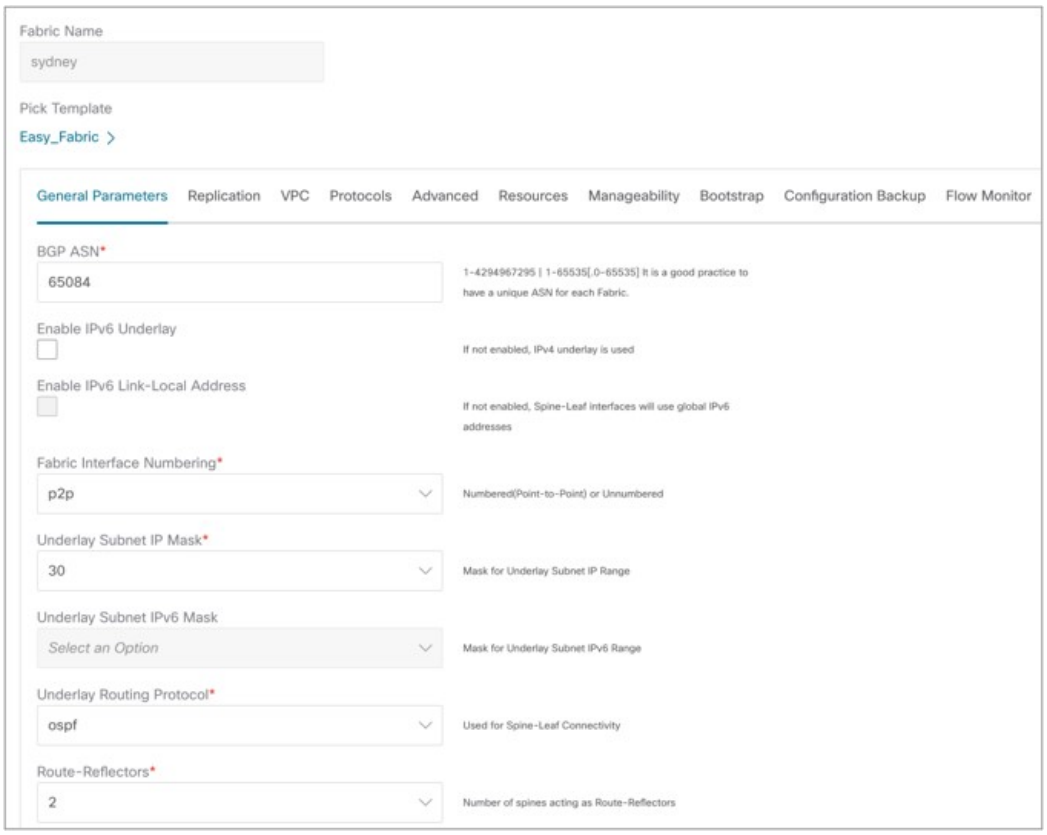

Step 7 In the **Advanced** parameter tab, make the necessary configuration specifically for this hybrid cloud topology use case.

- 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 27:

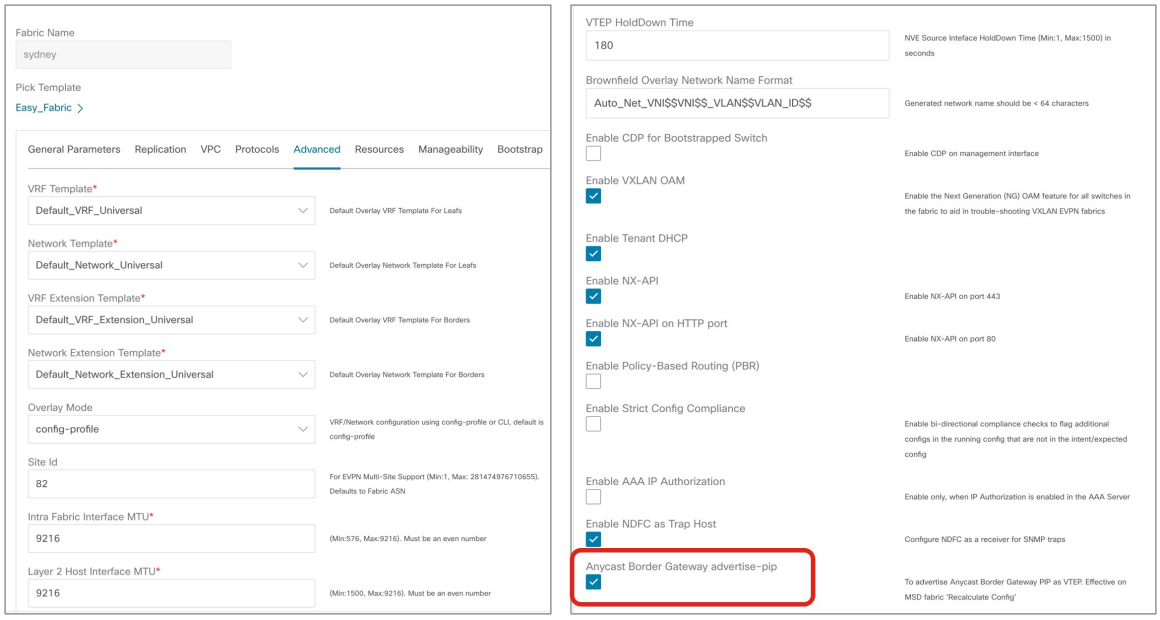

Step 8 Click the **Resources** parameter tab and enter the necessary values in this page.

- 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 28:

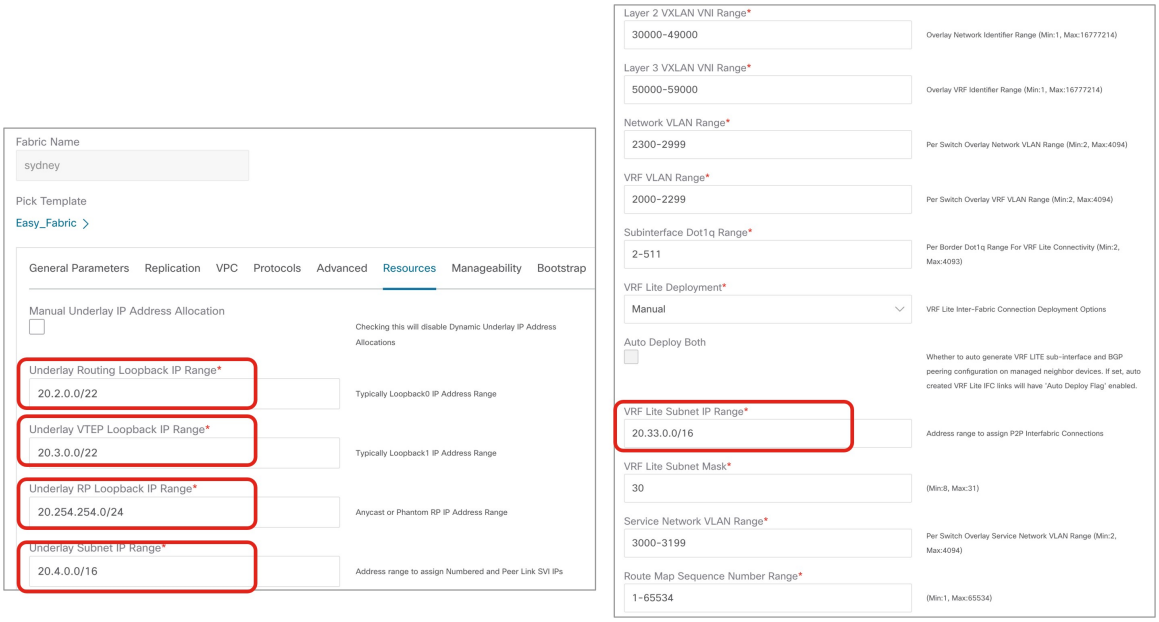

Step 9 Complete the necessary general VXLAN fabric parameter configurations in the **Manageability** and **Bootstrap** parameter tabs.

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 37.

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 33.

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 29:

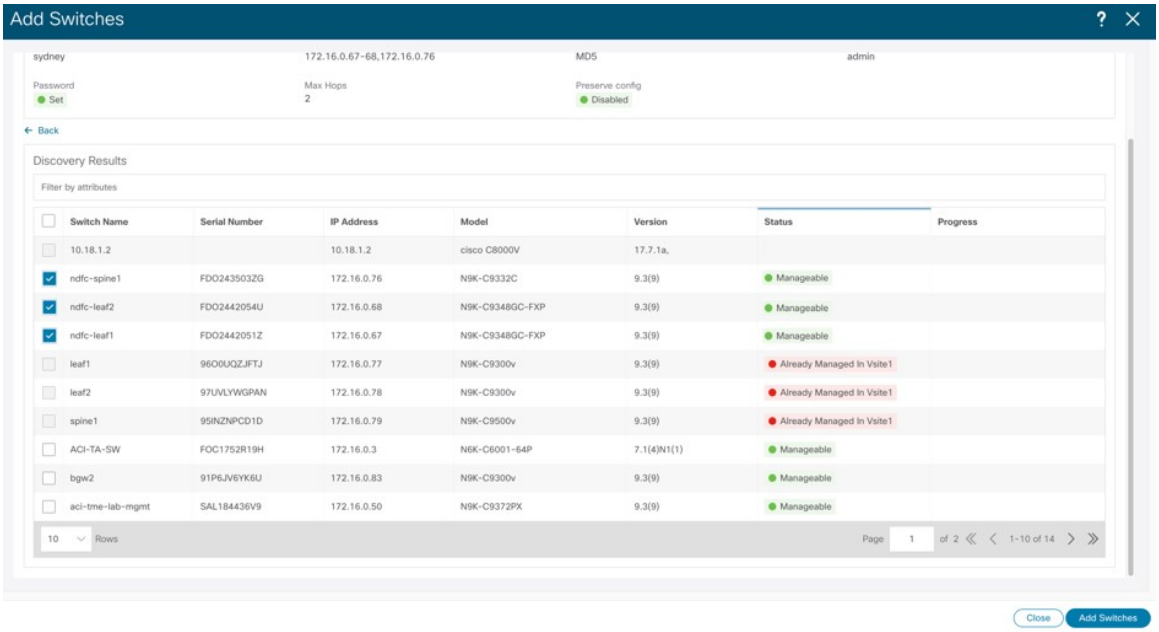

Step 6 Click **Add Switches**.

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.

a) Click the **Switches** tab in the NDFC VXLAN fabric overview window.
The switches that have been added to this fabric are displayed.

b) Click the box next to the spine switch to choose that switch, then click **Actions** > **Set Role**.

Figure 30:

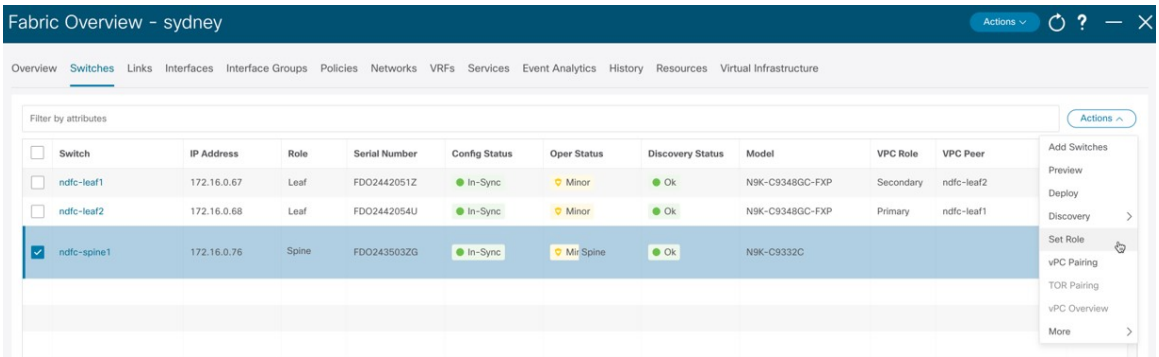

c) Locate and select the `Border Gateway Spine` role in the **Select Role** list, then click **Select**.

Figure 31:

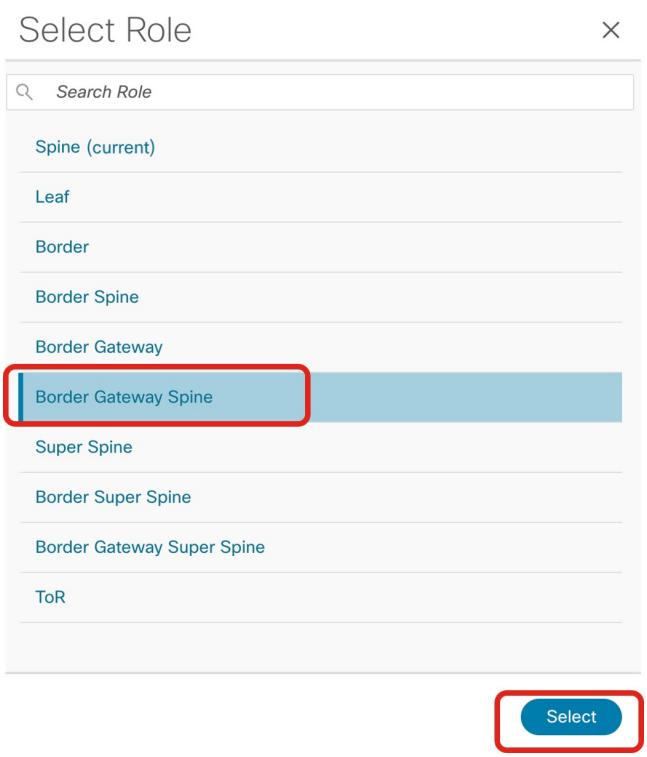

Step 8 Navigate to **LAN** > **Fabrics** and select the NDFC VXLAN fabric that you created.

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 32:

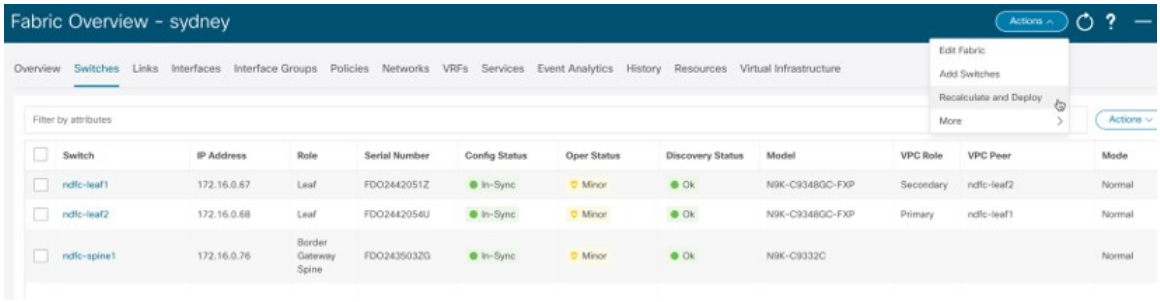

As described earlier, for these procedures, one spine switch plays the dual role of spine switch and border gateway spine switch, so we changed the role of the spine switch to border gateway spine switch in these example procedures, as shown below. In these example procedures, a vPC pair has also been 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 33:

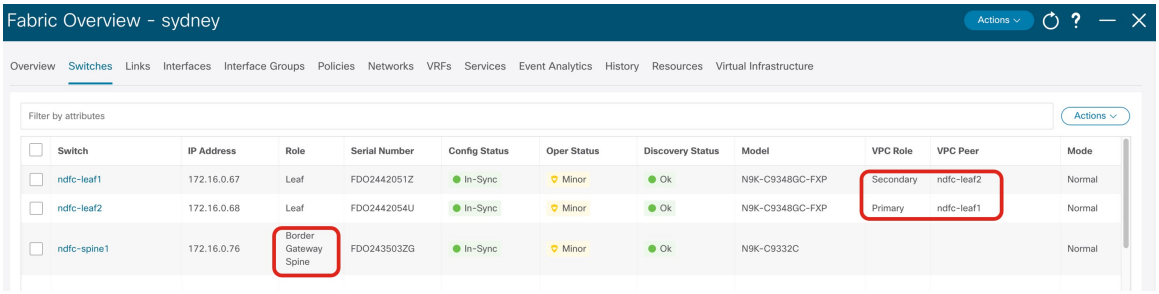

What to do next

Configure an NDFC external fabric using the procedures provided in Configure an NDFC External Fabric, on page 41.

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 34:

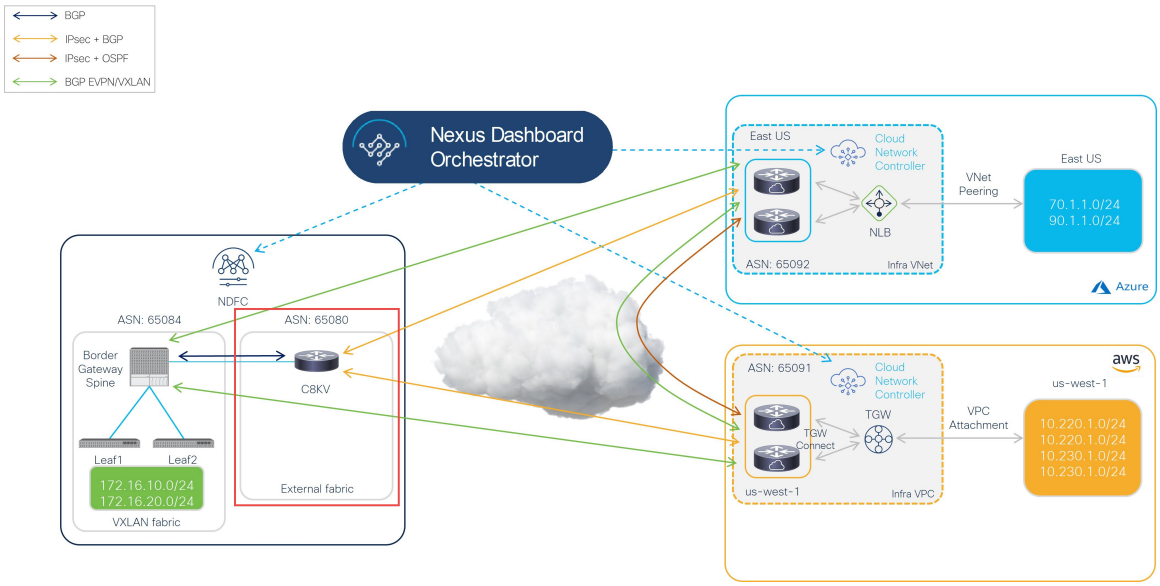

An NDFC-managed external fabric contains one or more IPsec devices. The IPsec devices have connectivity to cloud networks either via the internet (public) or by a private connection, such as Direct Connect (AWS) or ExpressRoute (Azure). If public internet is used to connect to the cloud sites, IPsec 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 33 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 35:

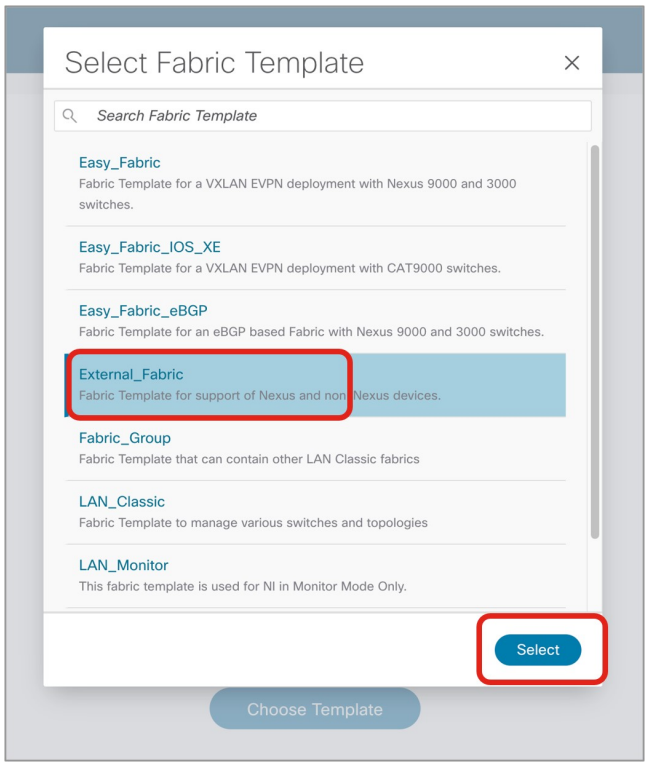

Step 5 In the **General Parameters** tab, make the necessary configuration specifically for this hybrid cloud topology use case.

- 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 36:

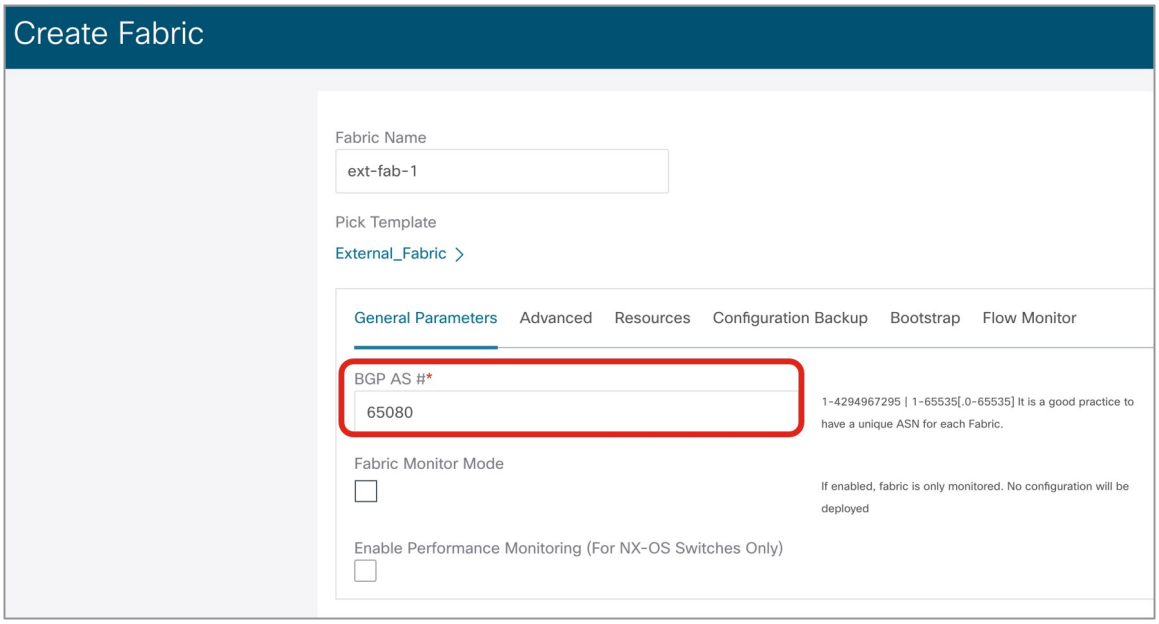

Step 6 Complete the necessary general external fabric parameter configurations.

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 44.

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 42

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 37:

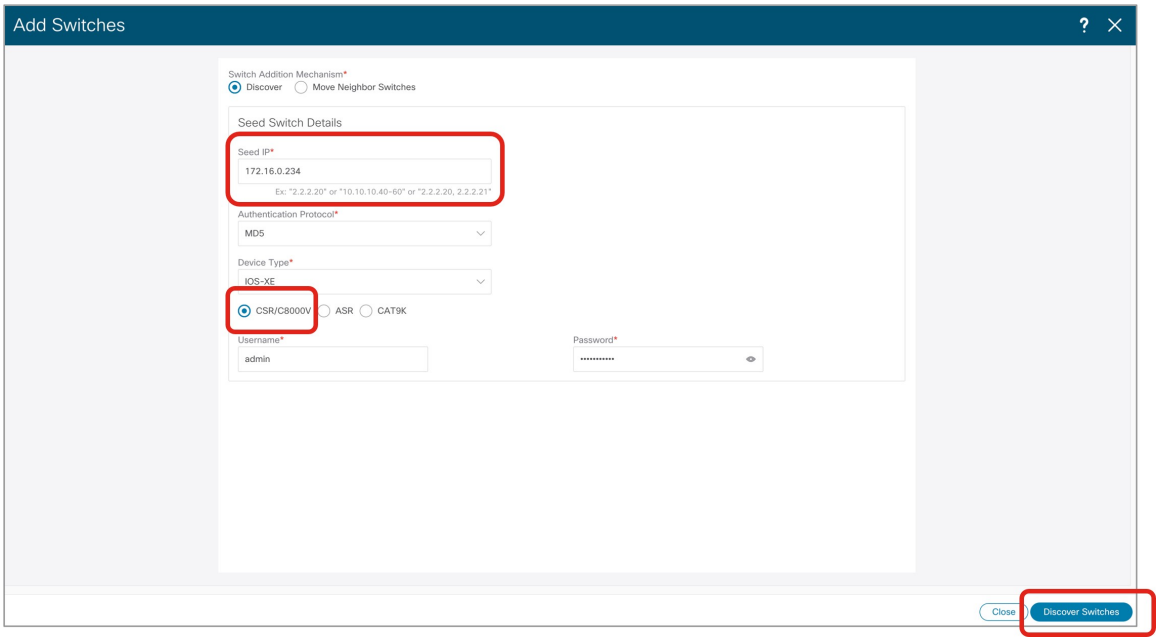

Step 4 Click **Discover Switches**.

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 38:

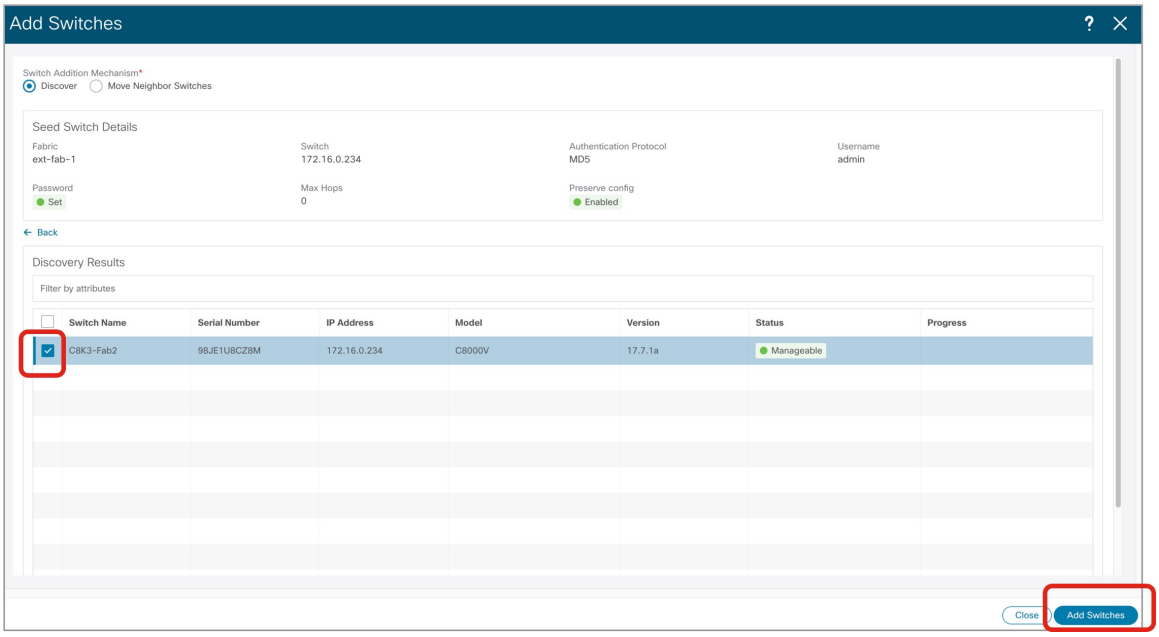

The status will change to **Switch Added**. Click **Close** to close out of this window.

Figure 39:

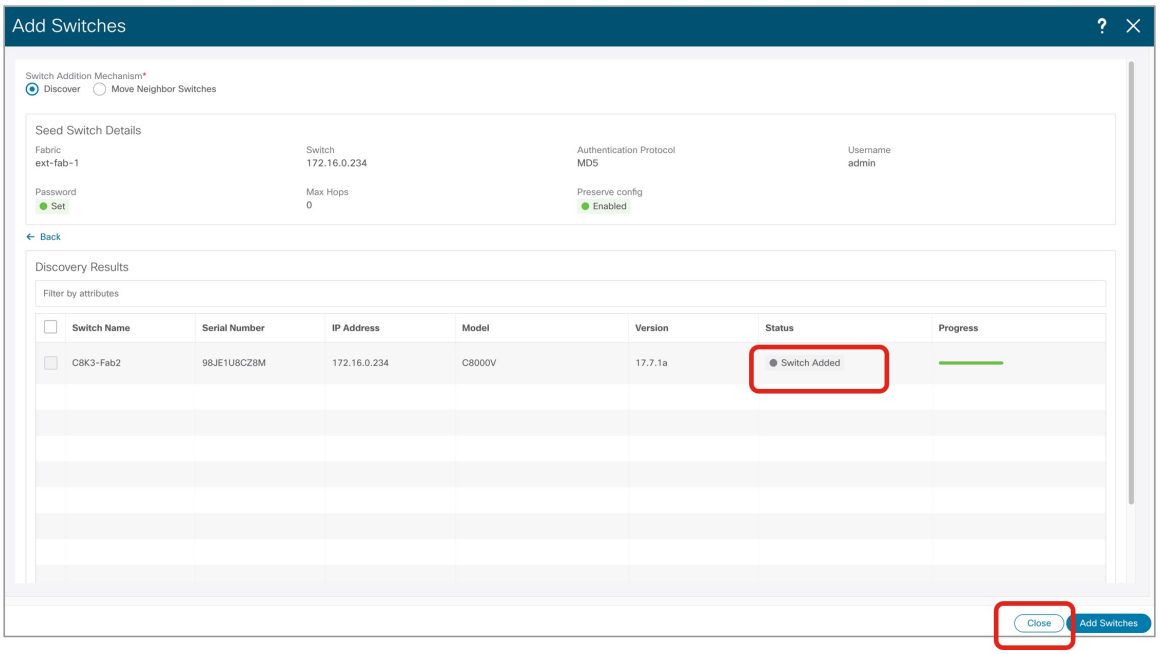

Step 6 Set the role for the Cisco Catalyst 8000V to `Core Router`.

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

Figure 40:

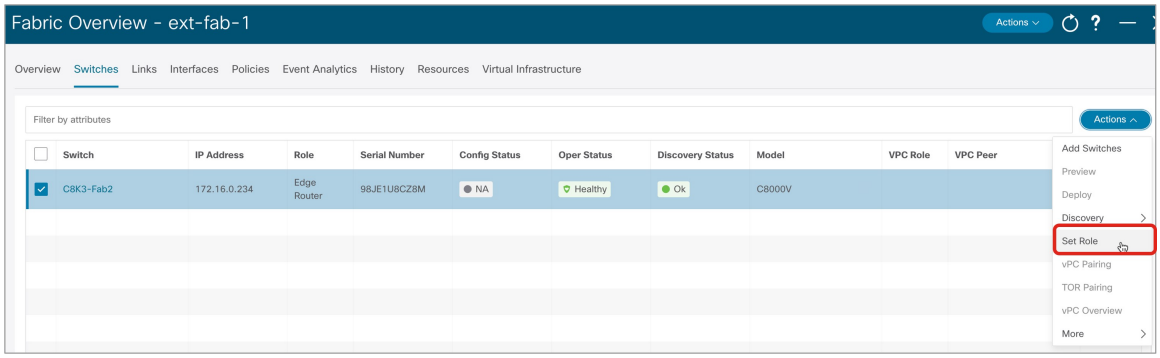

b) Locate and select the `Core Router` role in the **Select Role** list, then click **Select**.

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

Figure 41:

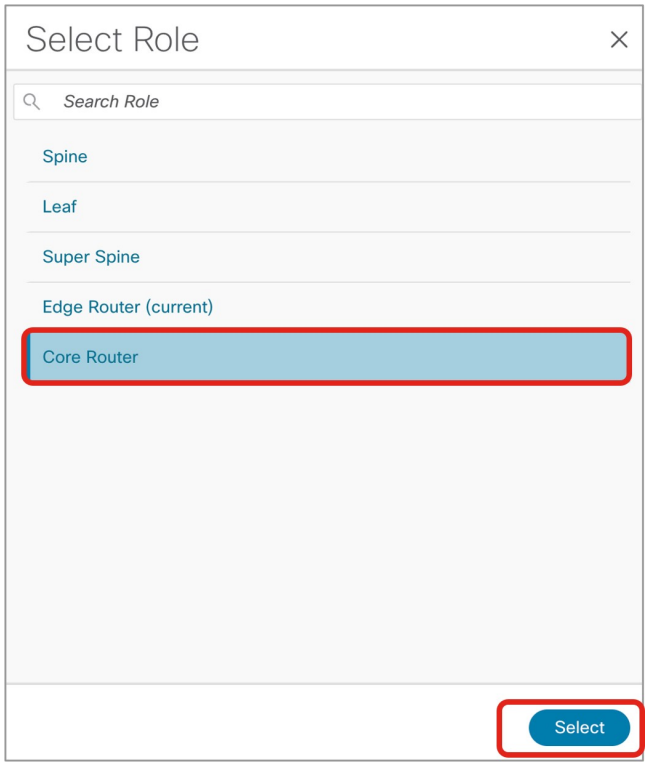

Step 7 Navigate to **LAN** > **Fabrics** and select the external fabric that you created.

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 42:

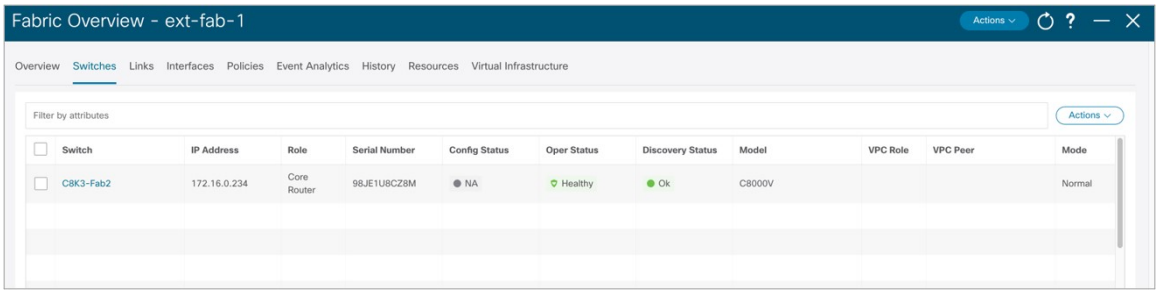

Step 9 Click **Actions** > **Recalculate and Deploy**.

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

Figure 43:

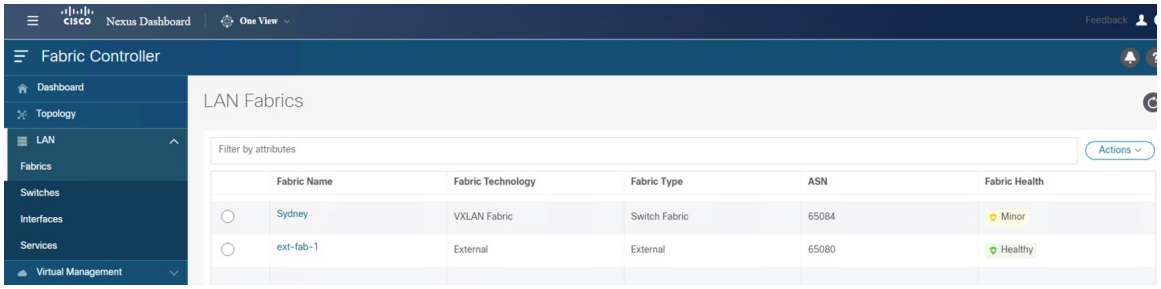

You can also use the **Topology** view to determine the following configurations at this point in the process:

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

  Figure 44:

  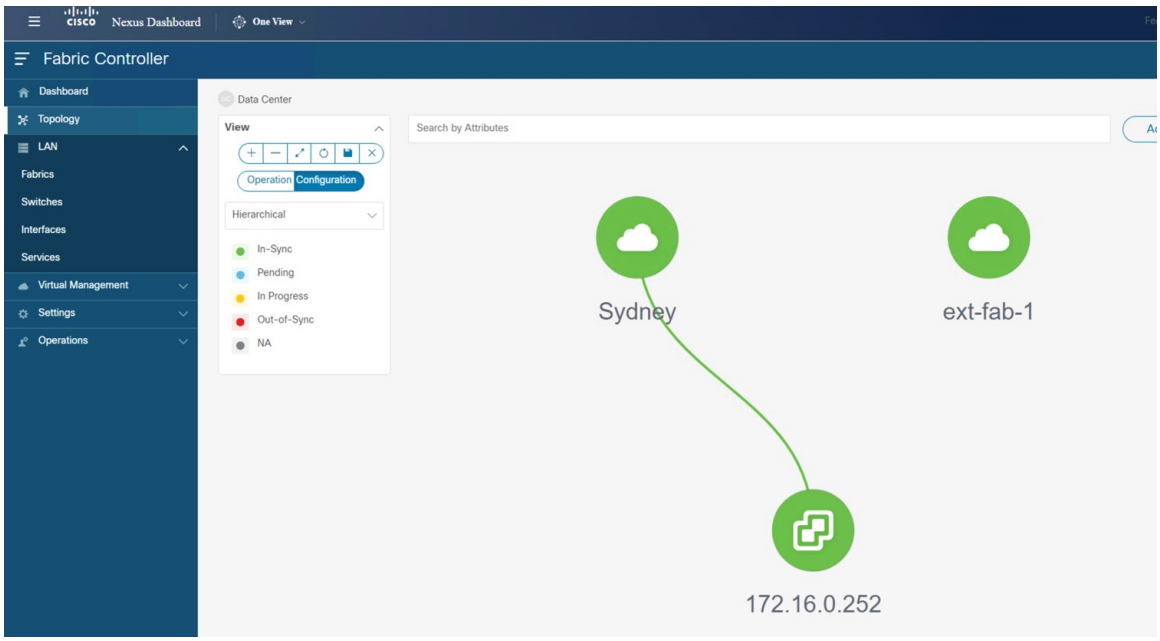

  This NDFC has the VMM Visualizer feature enabled, so the vCenter icon with an IP address of 172.16.0.252 is displayed in the topology view. For more information on the VMM feature, see the Virtual Infrastructure Manager chapter in the *Cisco NDFC-Fabric Controller Configuration Guide*.

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

Figure 45:

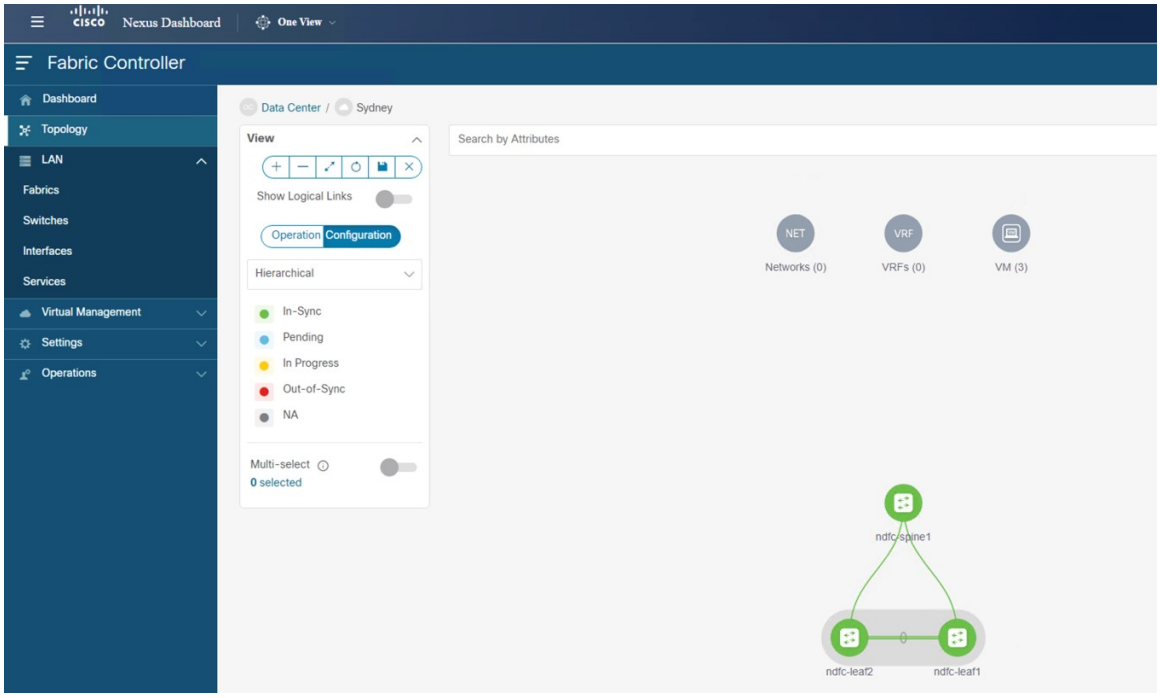

What to do next

Deploy the Cloud Network Controller on the cloud sites using the procedures provided in Deploy Cloud Network Controller on Cloud Sites, on page 49.

Deploy Cloud Network Controller on Cloud Sites

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

Figure 46:

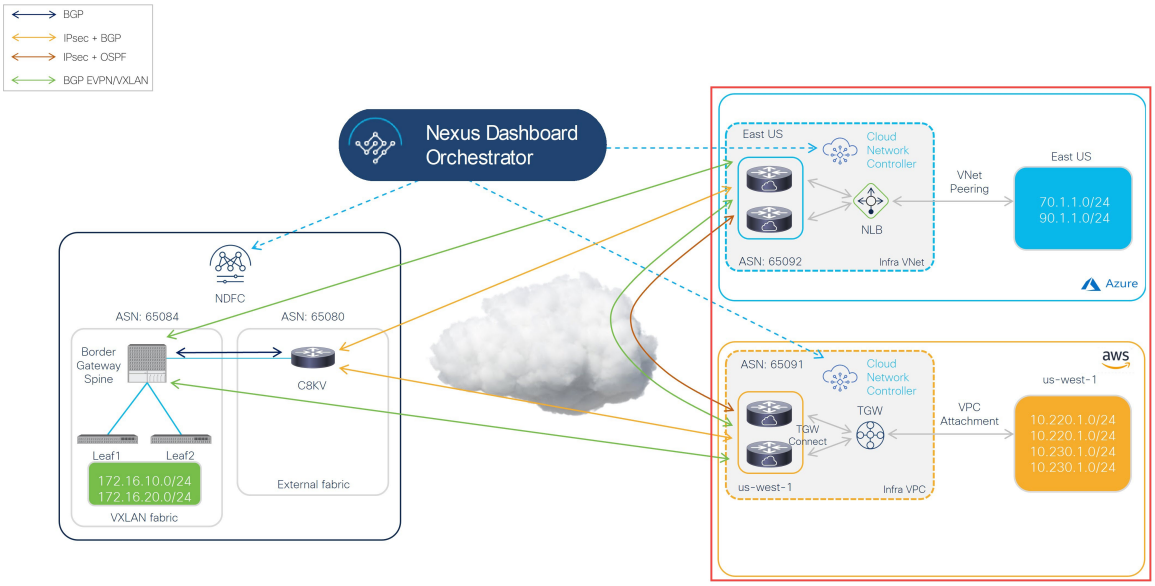

Based on the example hybrid cloud topology, these procedures assume that we will be setting up two cloud sites through the Cloud Network Controller (AWS and Azure cloud sites). We will therefore refer to the following documents throughout these procedures:

- *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 52.

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 50.

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 47:

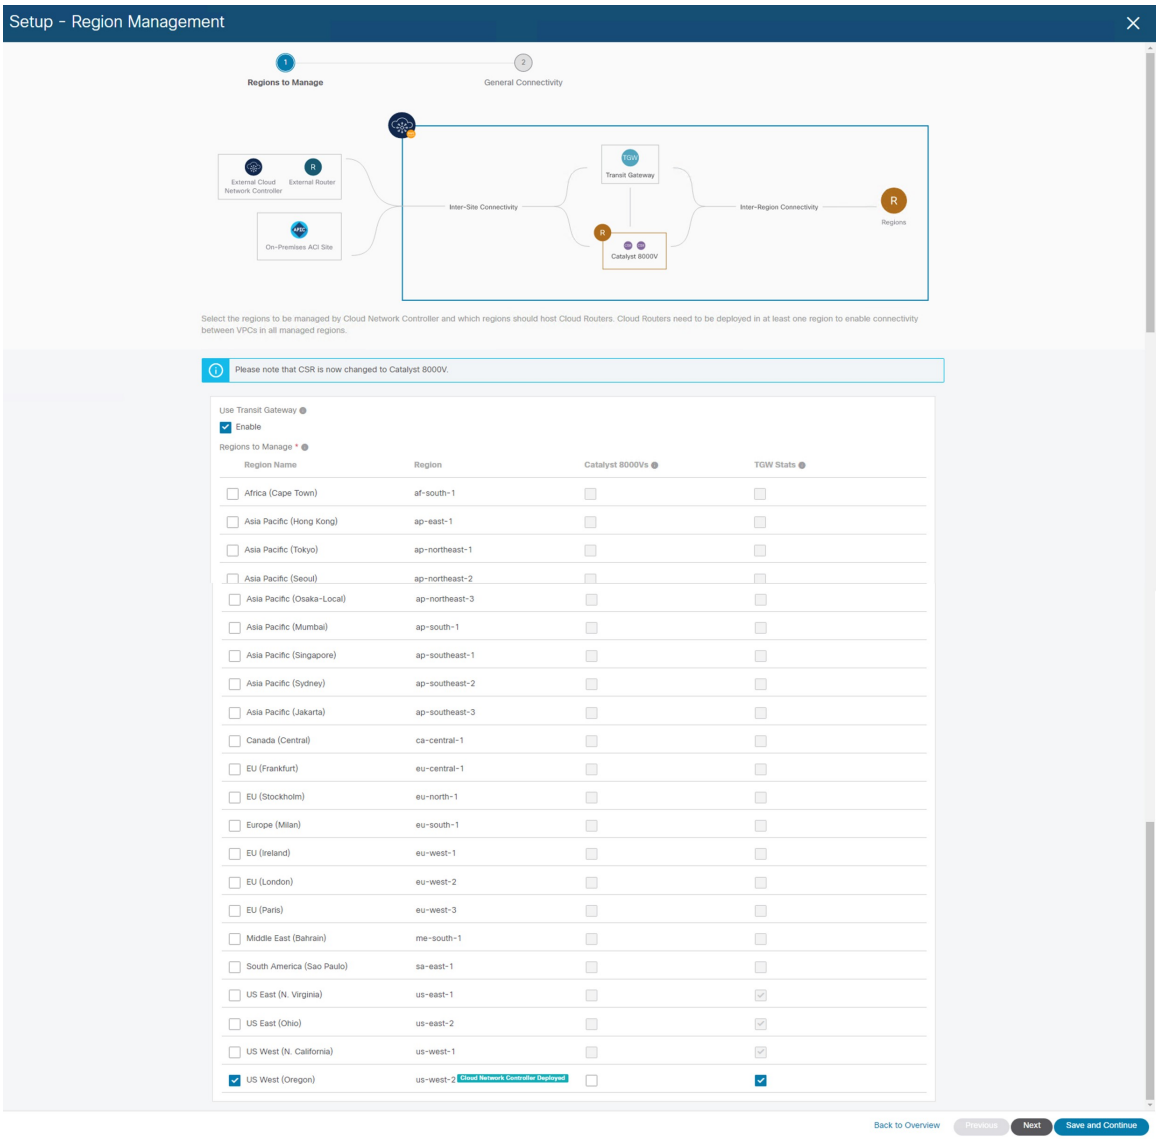

Step 7 When you have selected all the appropriate regions, click **Next** at the bottom of the page.

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 48:

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 49:

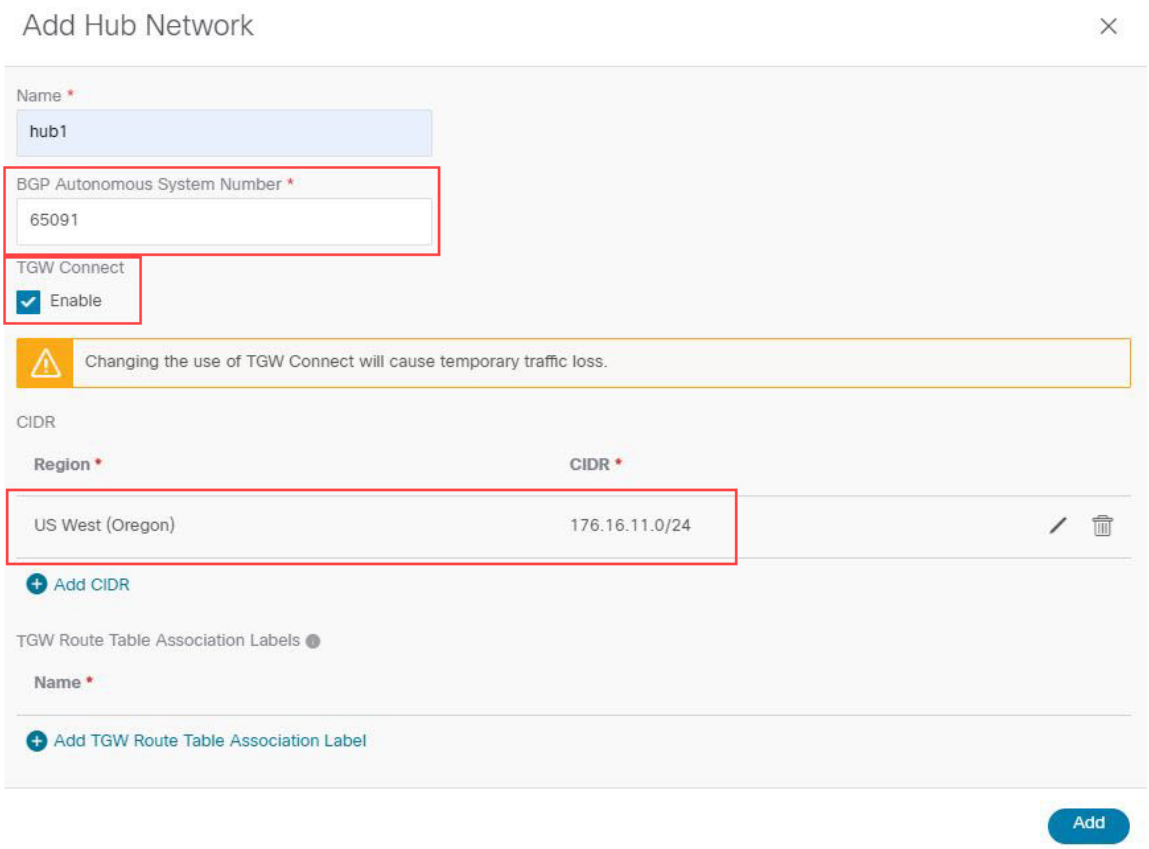

Step 14 Complete the remaining configurations as you normally would.

- 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 56.

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 50.

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 57.

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 56.

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 50:

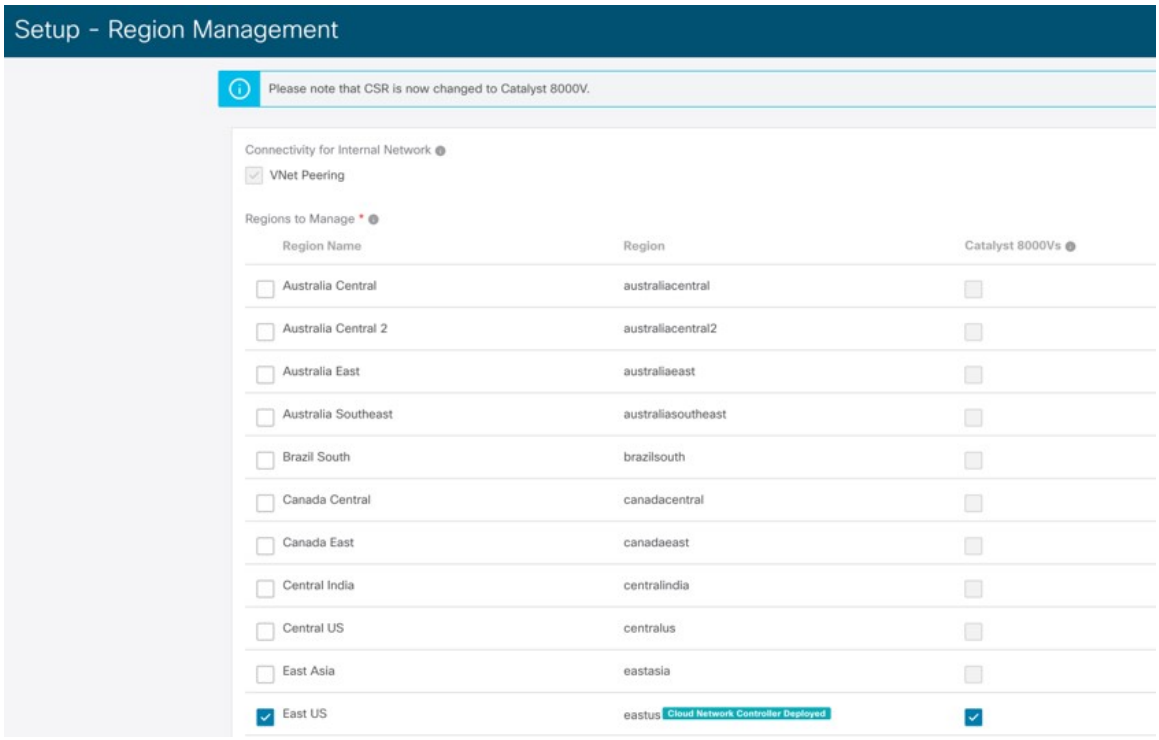

Step 6 When you have selected all the appropriate regions, click **Next** at the bottom of the page.

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 51:

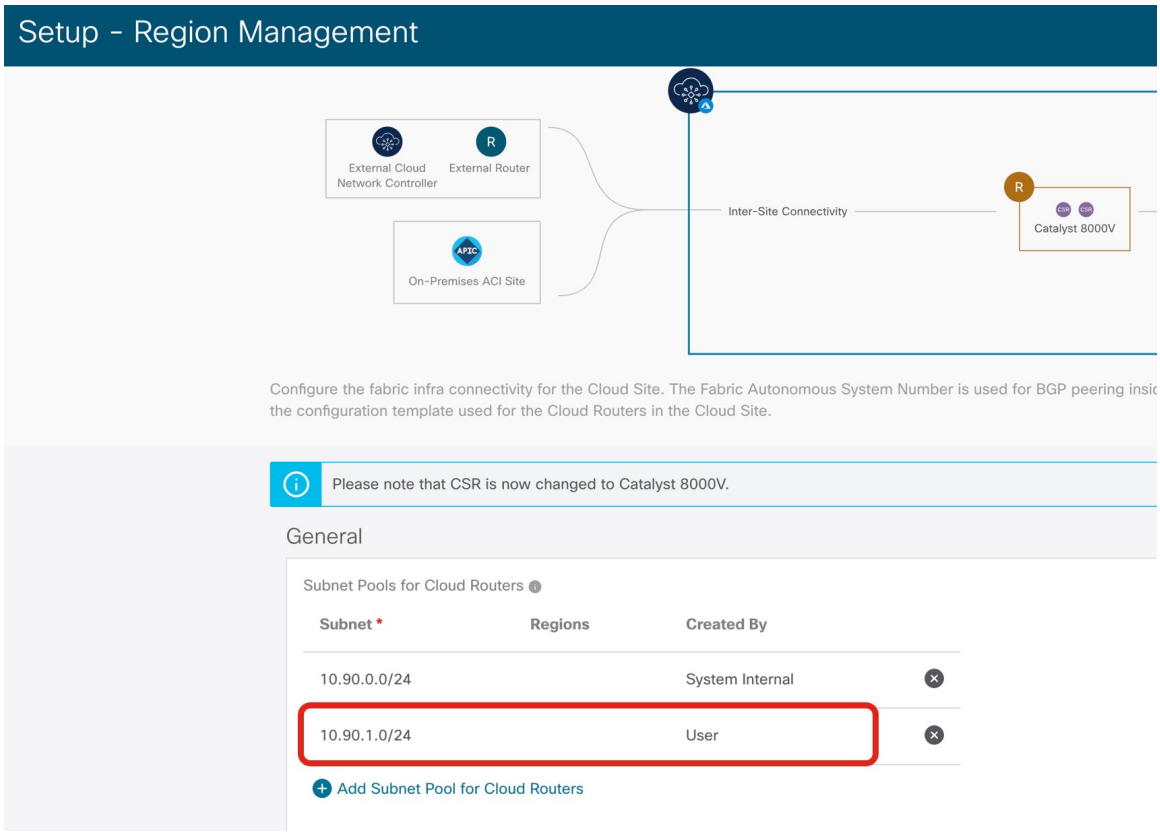

Step 9 Under the **Catalyst 8000Vs** area, in the **BGP Autonomous System Number for C8kVs** field, enter the BGP autonomous system number (ASN) that is unique to this site.

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 52:

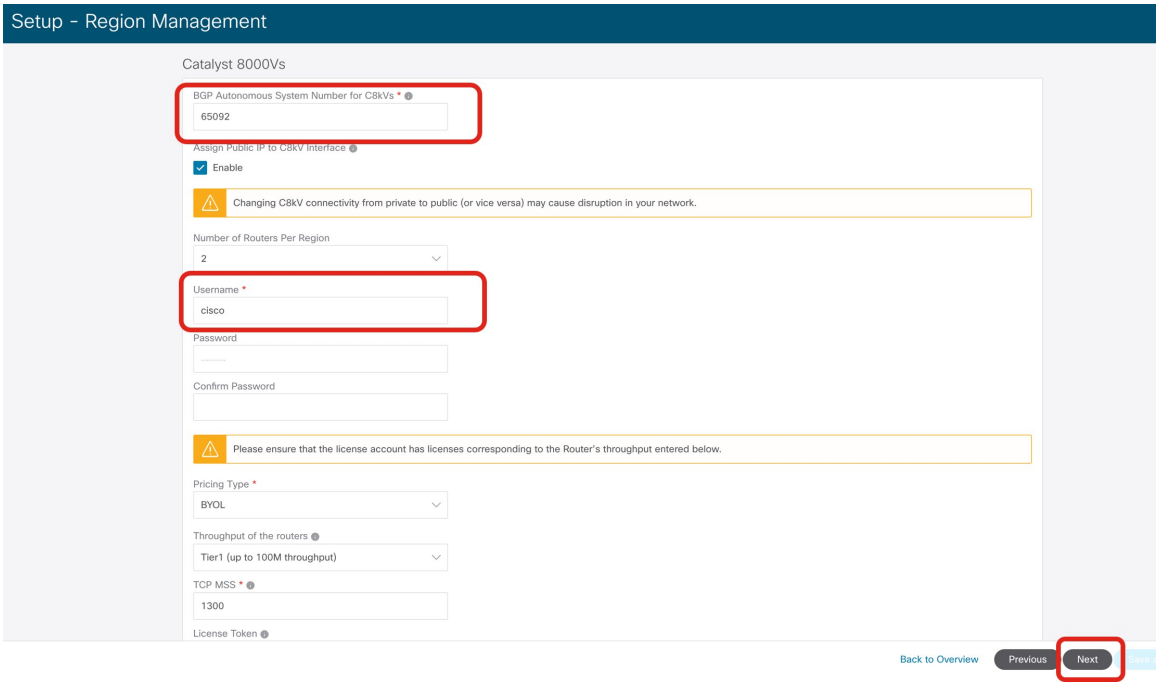

Step 10 Click **Next**, then complete the remaining configurations as you normally would.

- 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 62.

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 33.
- Create the NDFC external fabric using the procedures provided in Create an NDFC External Fabric, on page 42.
- 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 50.
- 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 56.

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

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

Figure 53:

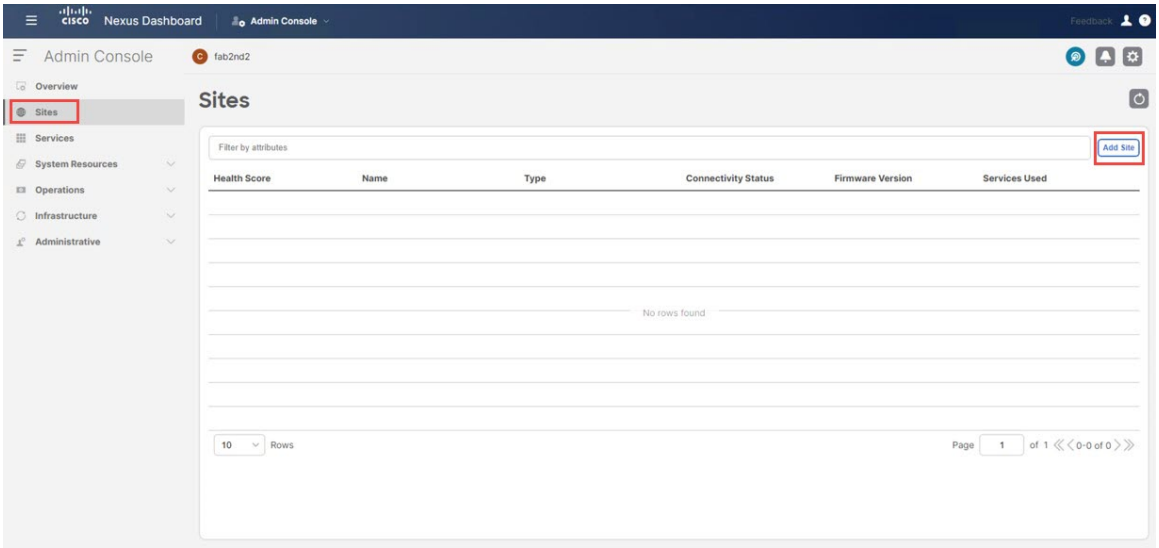

The **Add Site** page appears.

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 54:

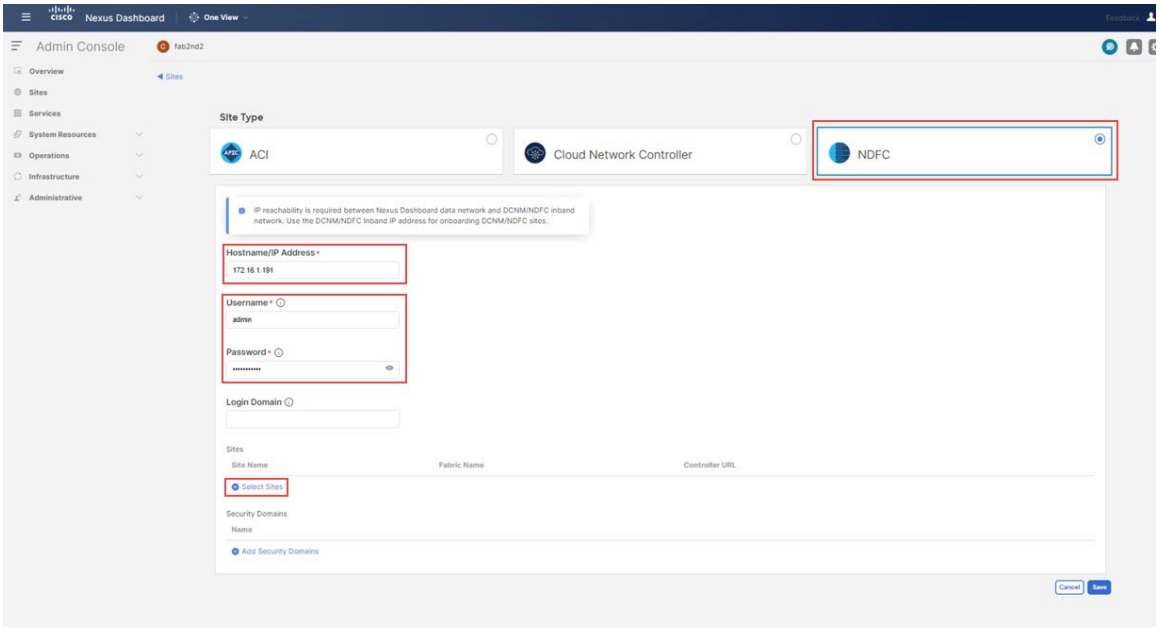

Step 6 Click the boxes next to the two NDFC sites that you added previously (the VXLAN fabric and external fabric sites), then click **Select**.

Figure 55:

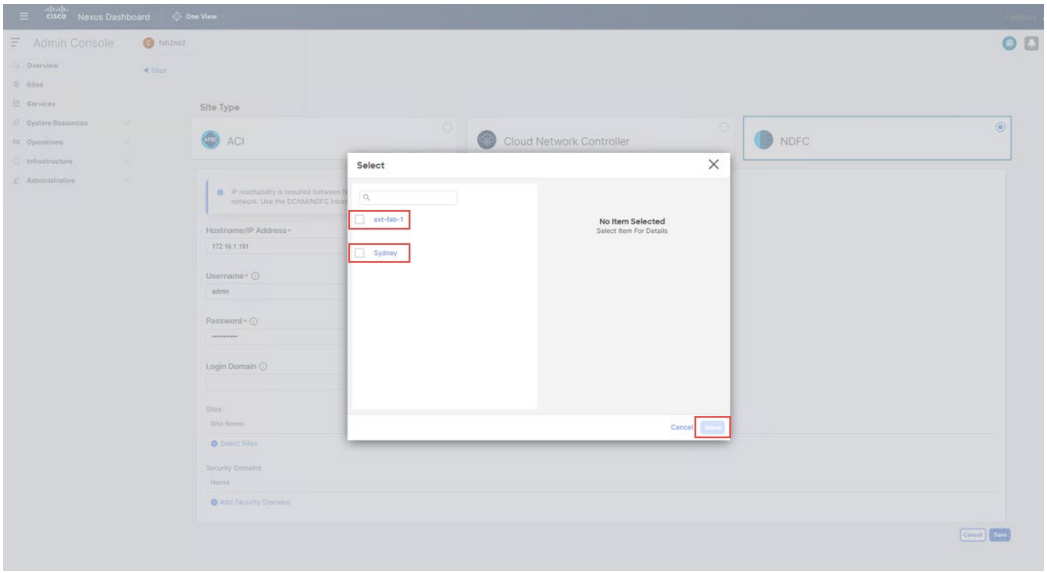

You are returned to the **Add Site** page.

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 56:

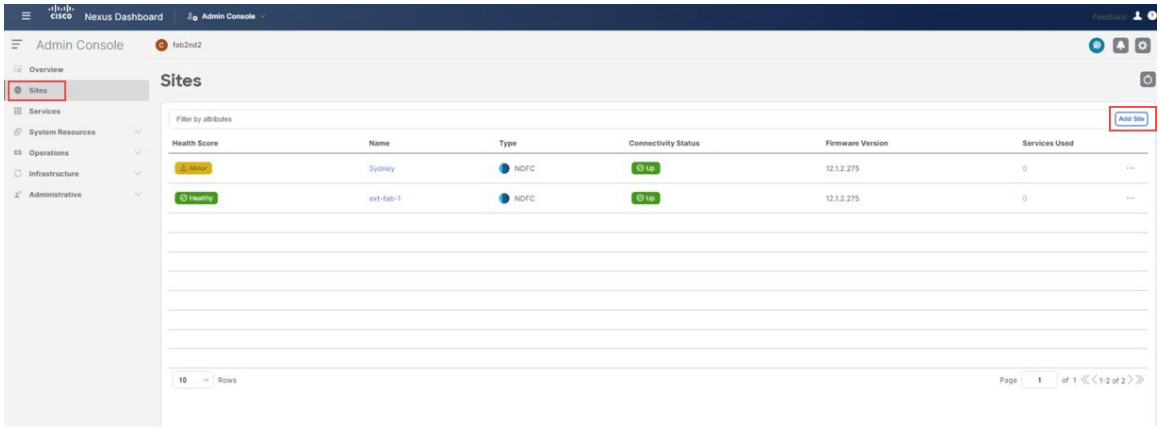

The **Add Site** page appears.

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 57:

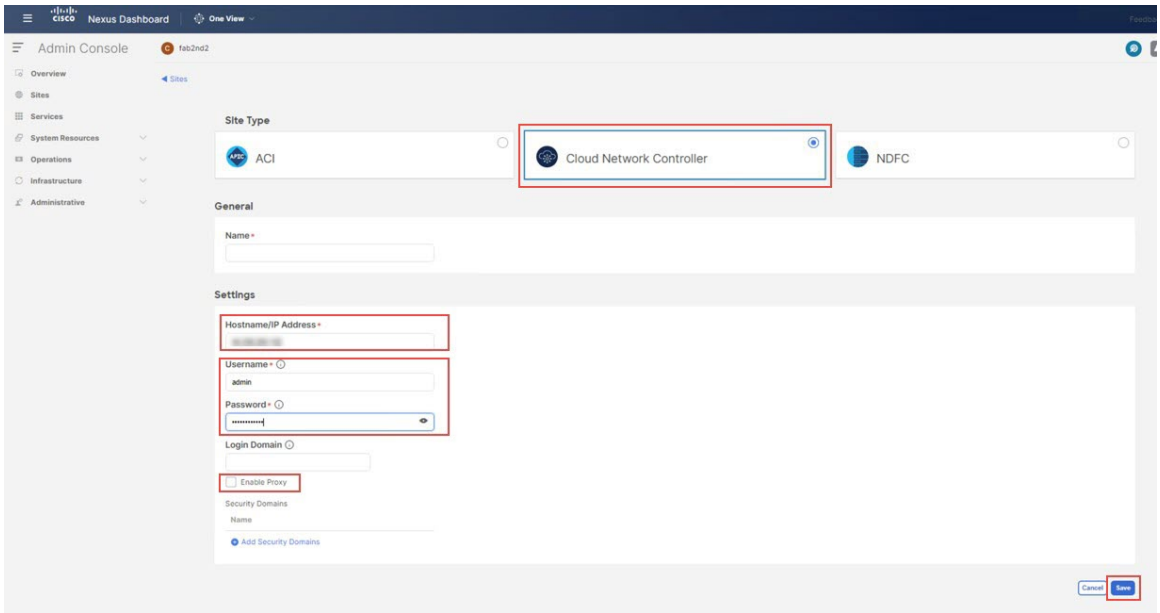

Step 10 Click **Save** to add the first cloud site.

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

Figure 58:

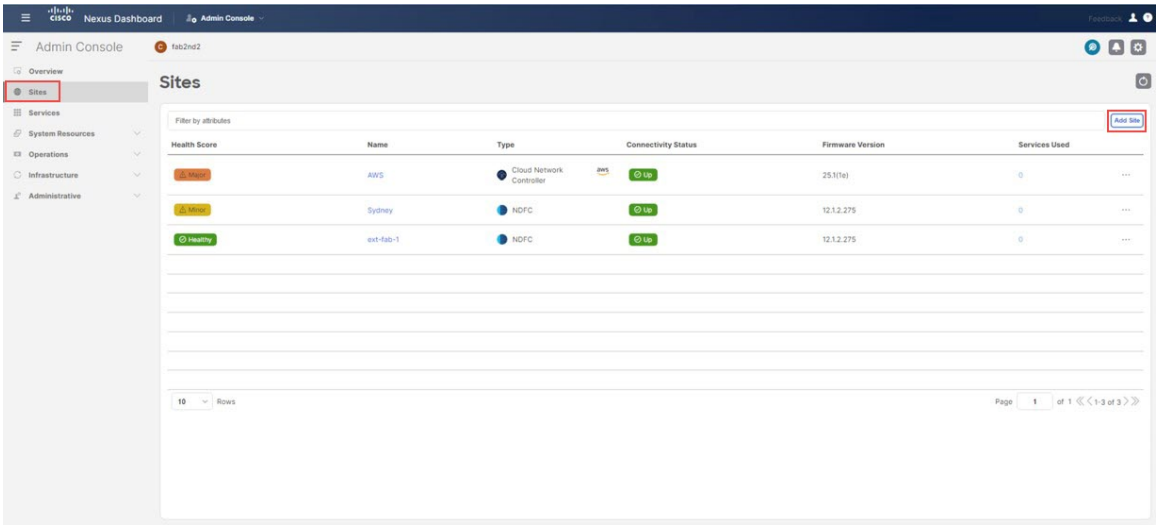

The **Add Site** page appears.

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 59:

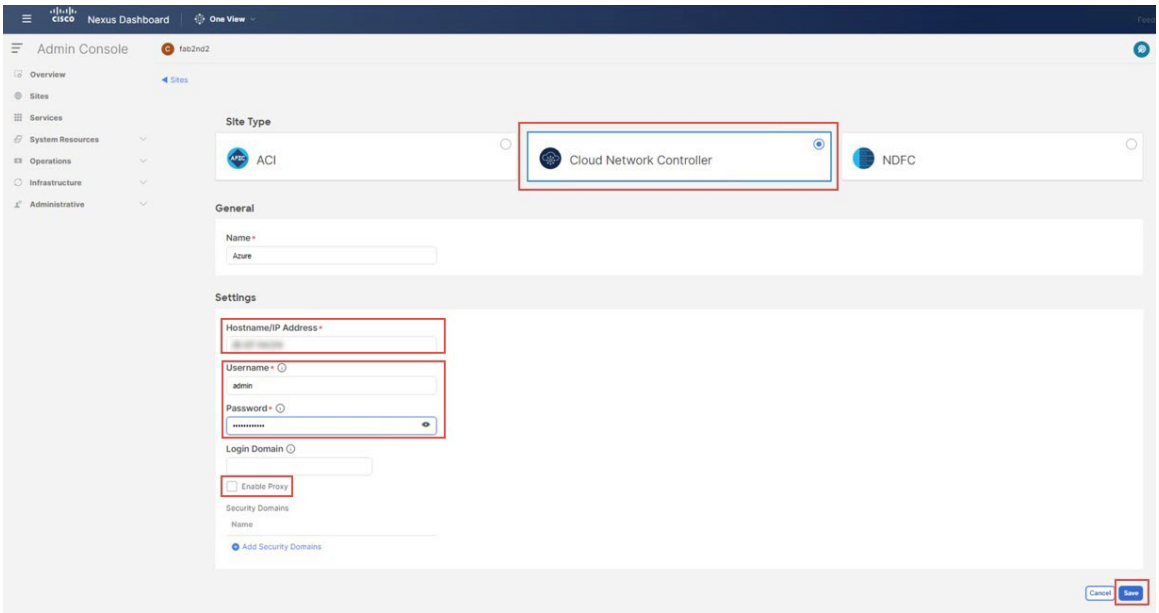

Step 13 In Nexus Dashboard, click **Sites** and verify that the four sites appear correctly:

- 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 60:

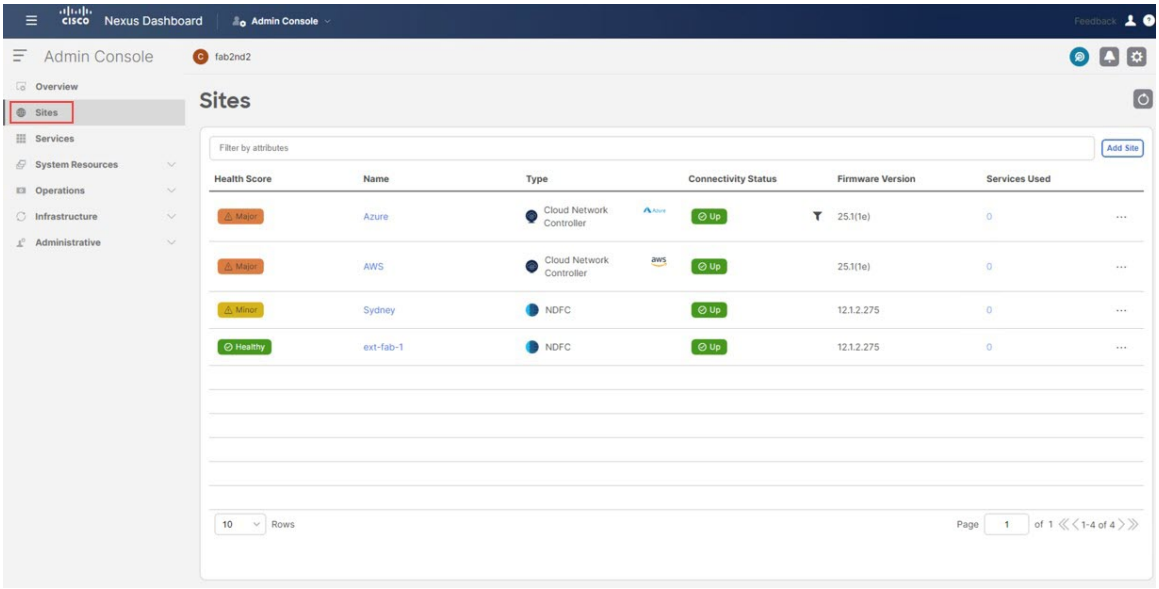

Step 14 Access the Nexus Dashboard Orchestrator (NDO).

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

Figure 61:

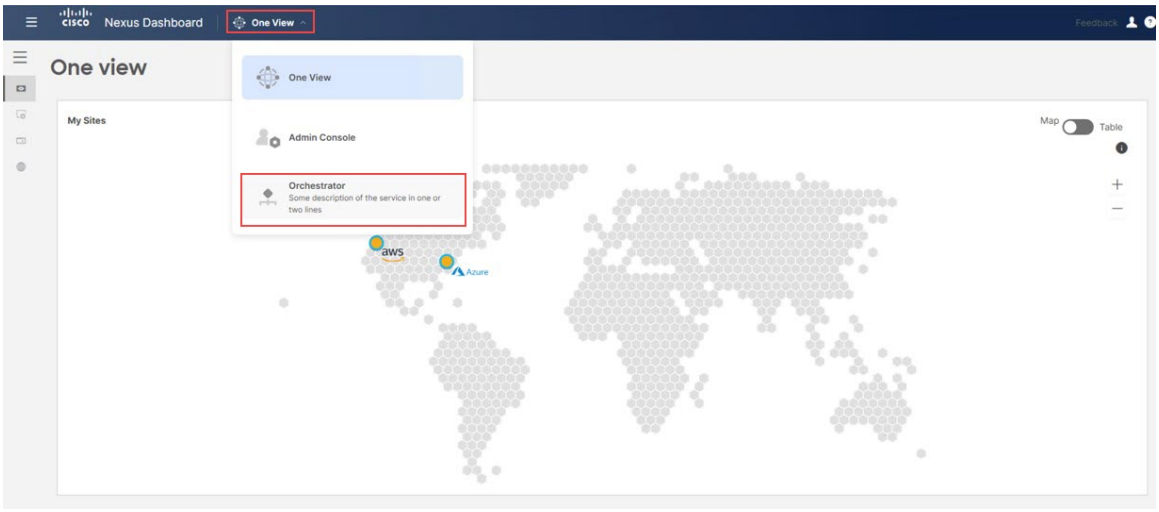

Step 15 In NDO, click **Sites**.

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

Figure 62:

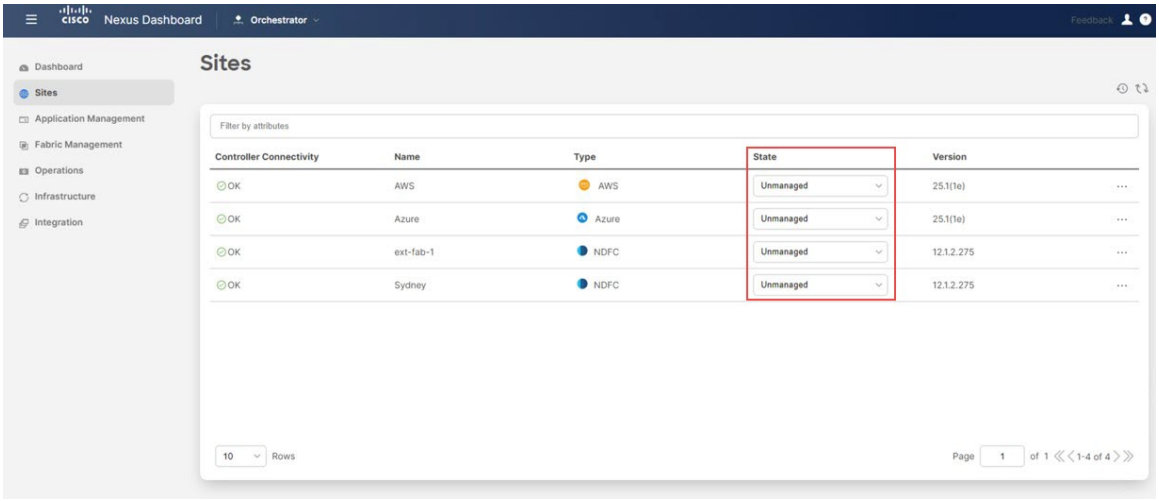

Step 16 From NDO, manage the four sites.

Perform the following steps for each site in NDO:

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

Figure 63:

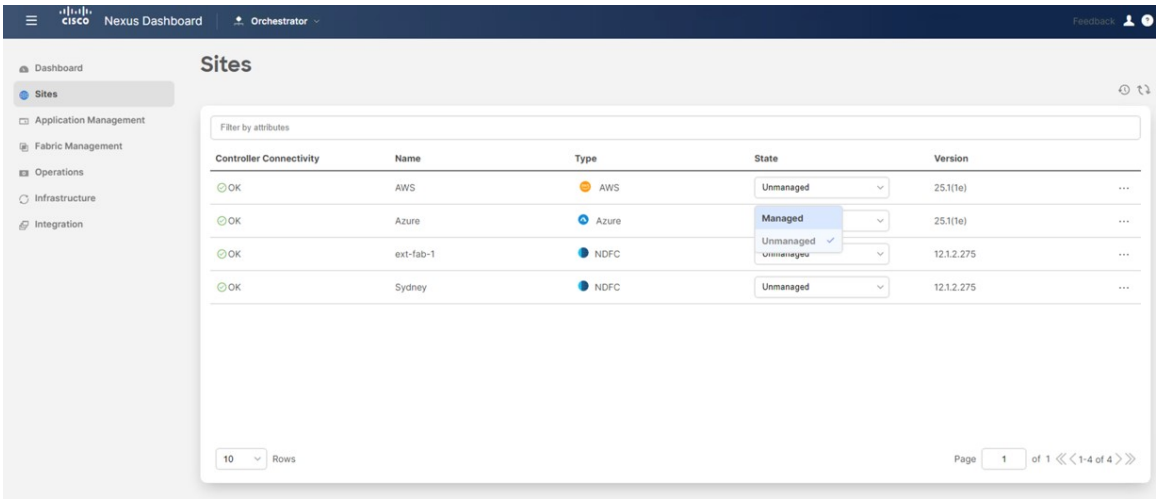

b) Provide a site ID that is unique to this particular site (a site ID that does not conflict with site IDs for any other site being managed through this NDO), then click **Add**.

Figure 64:

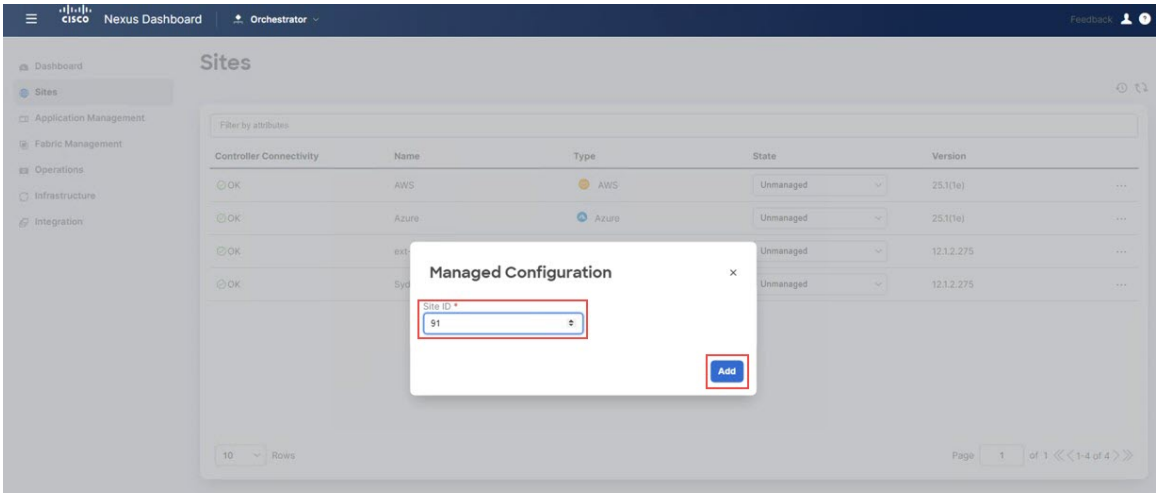

c) Repeat these steps for the remaining sites in NDO to change each site to the **Managed** state and provide a unique site ID for each site.

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 65:

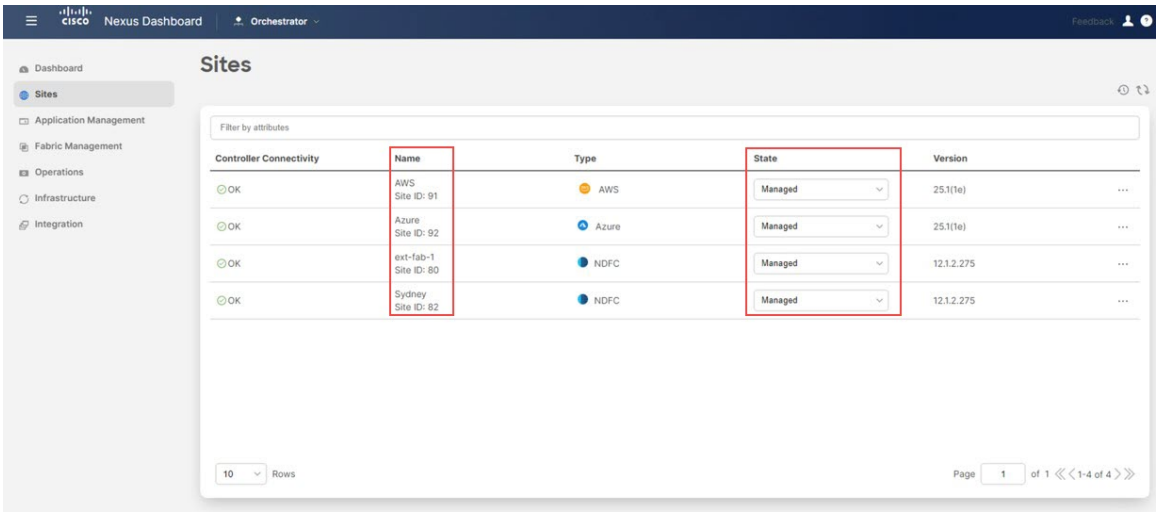

What to do next

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 69.

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 62.

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

Figure 66:

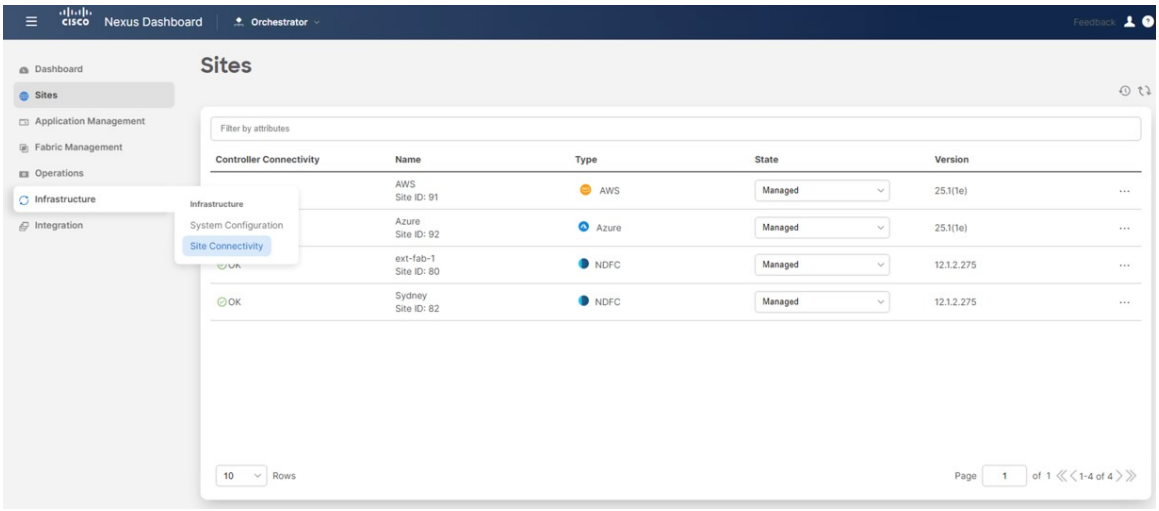

At this point, you will see the sites on the world map but they will not have any links in between, which means that there is no connectivity between the sites at this point.

Figure 67:

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 68:

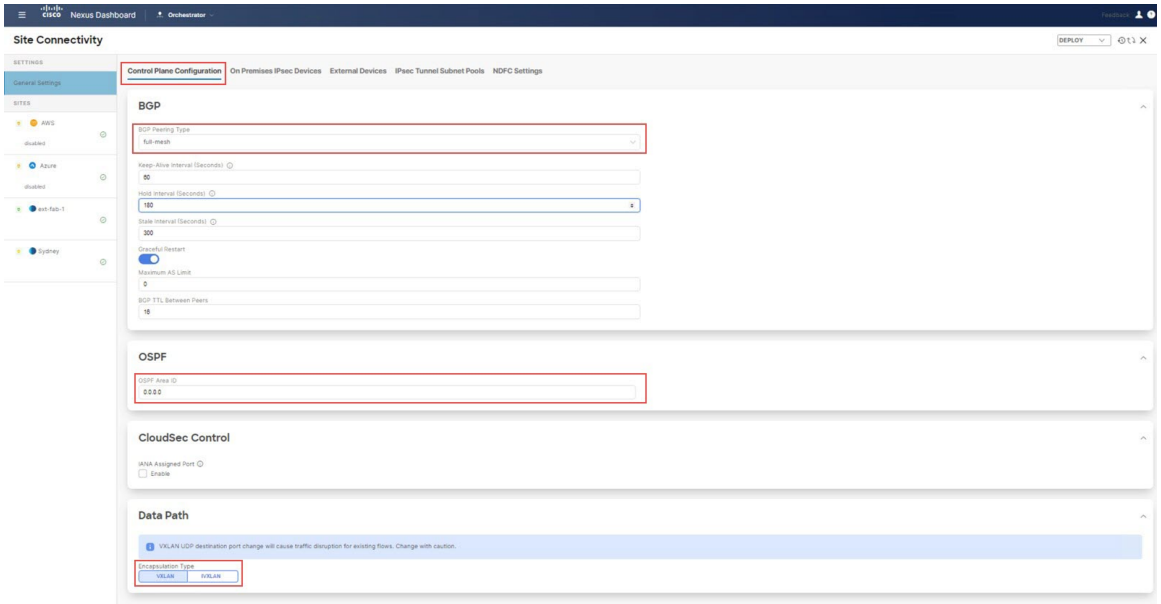

Note that BGP is used for underlay connectivity between on-premises and cloud sites, whereas OSPF is used for cloud-to-cloud underlay connectivity.

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, on page 13 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, on page 19 topology in Supported Topologies with IPsec (Multi-Cloud), on page 18, 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 OSFP 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 71.

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 69.

Step 1 In the same **General Settings** page, click the **On Premises IPsec Devices** tab.

Step 2 Click **Add On Premises IPsec Device**.

Figure 69:

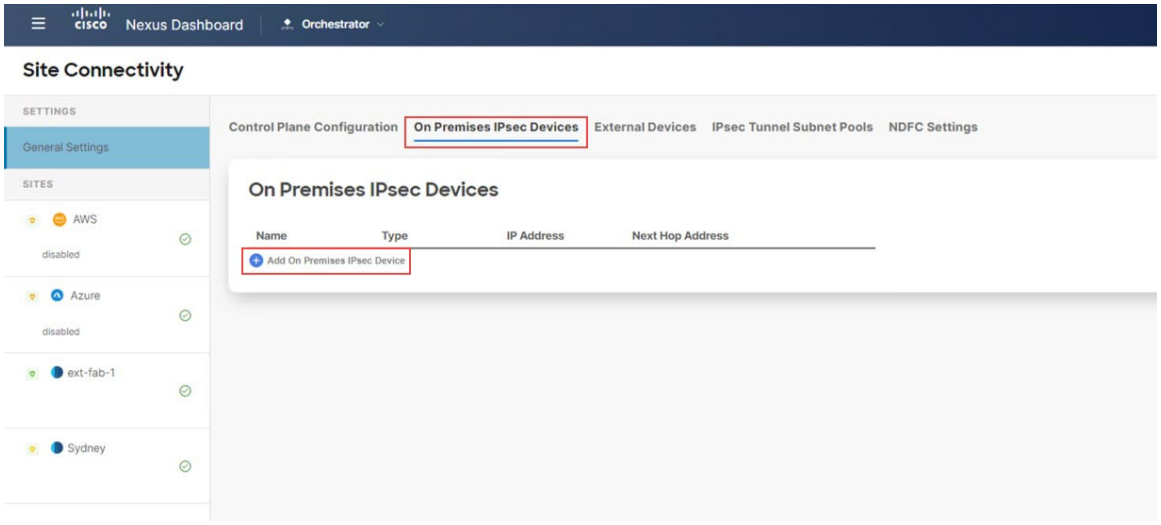

The **Add On Premises IPsec Device** page appears.

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-premies 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 70:

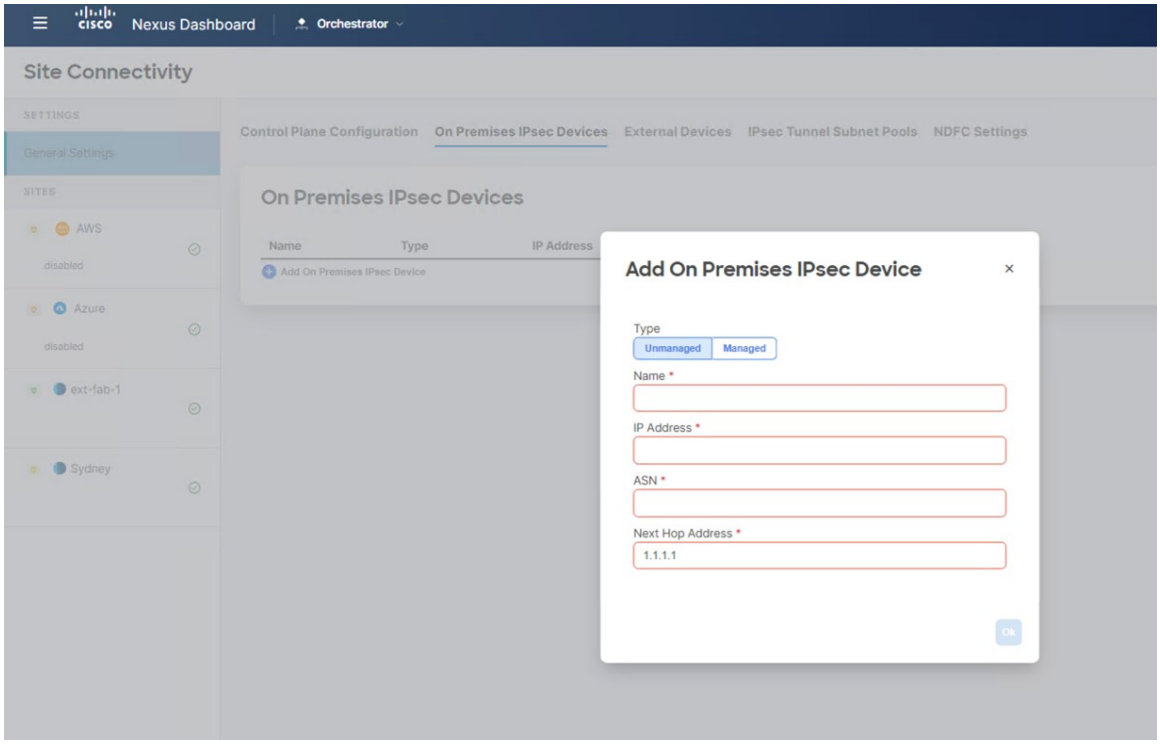

- If you choose the **Managed** option for the on-premises IPsec device, the **Site** field becomes available below the **Managed** option. The sites available in the **Site** field is based on information that NDO pulls from NDFC for the external fabrics configured in NDFC.

Figure 71:

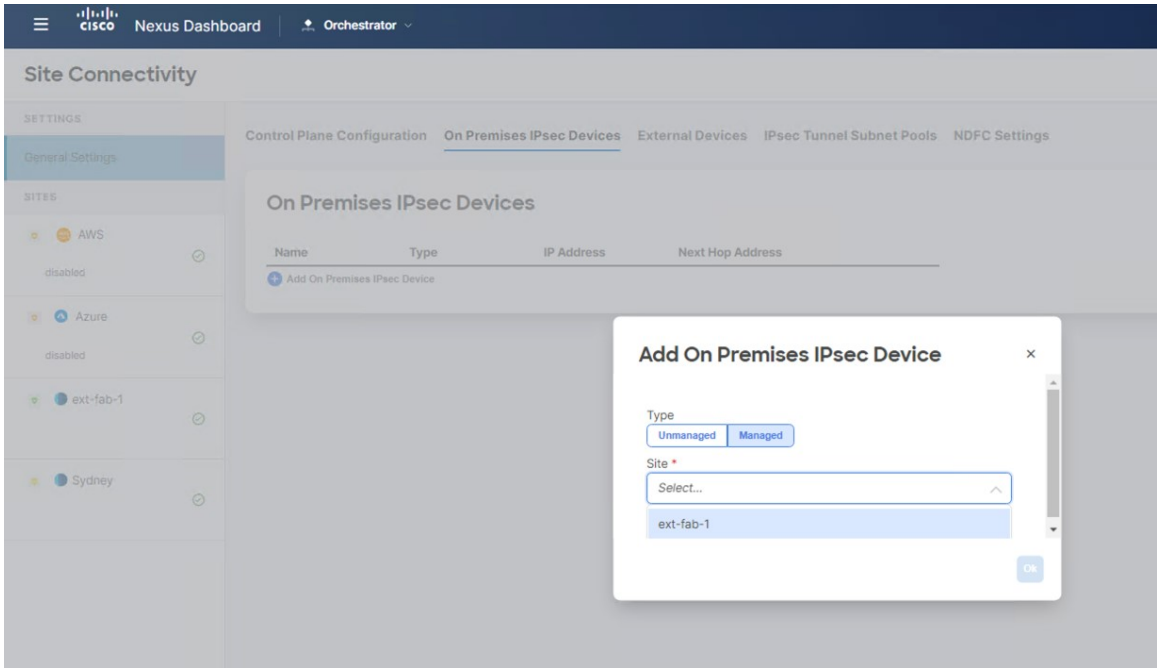

Choose the NDFC external fabric with the managed on-premises IPsec device. The **ASN** field is automatically populated in this case based on the site that you chose.

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 72:

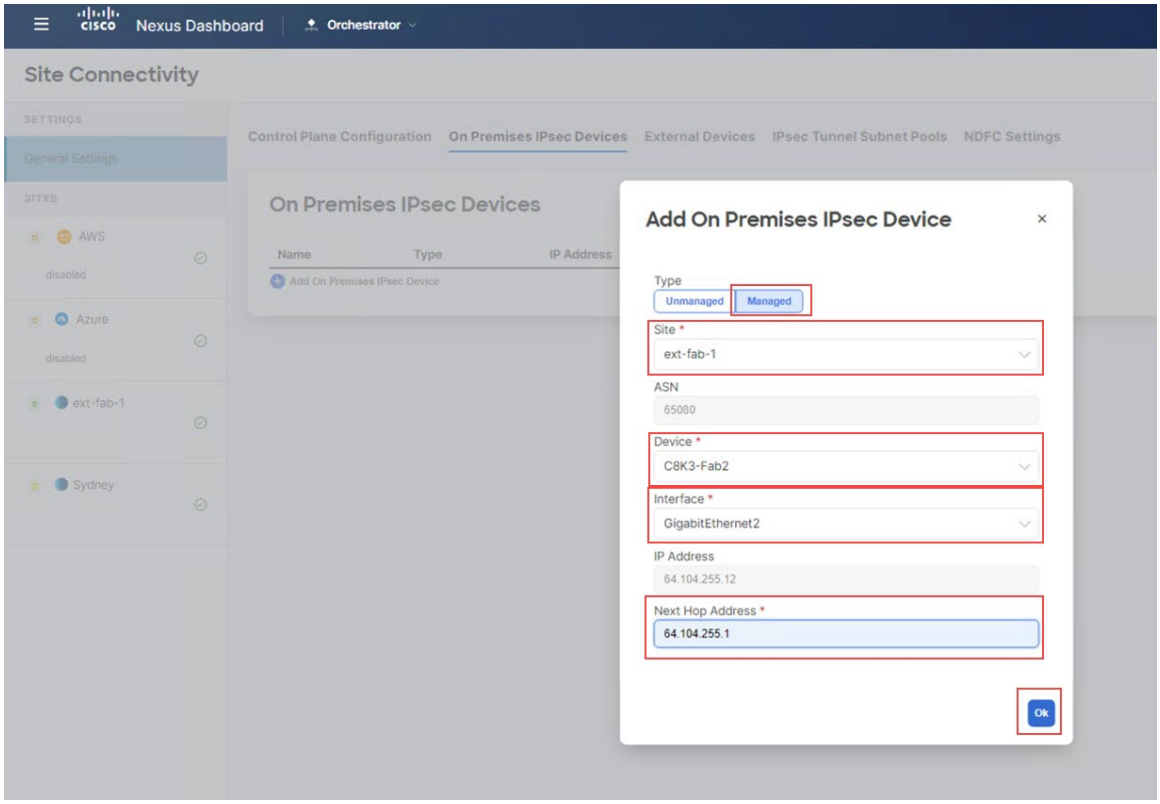

Step 4 When you have finished entering the necessary information in the **Add On Premises IPsec Device** page, click **Ok**.

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 73:

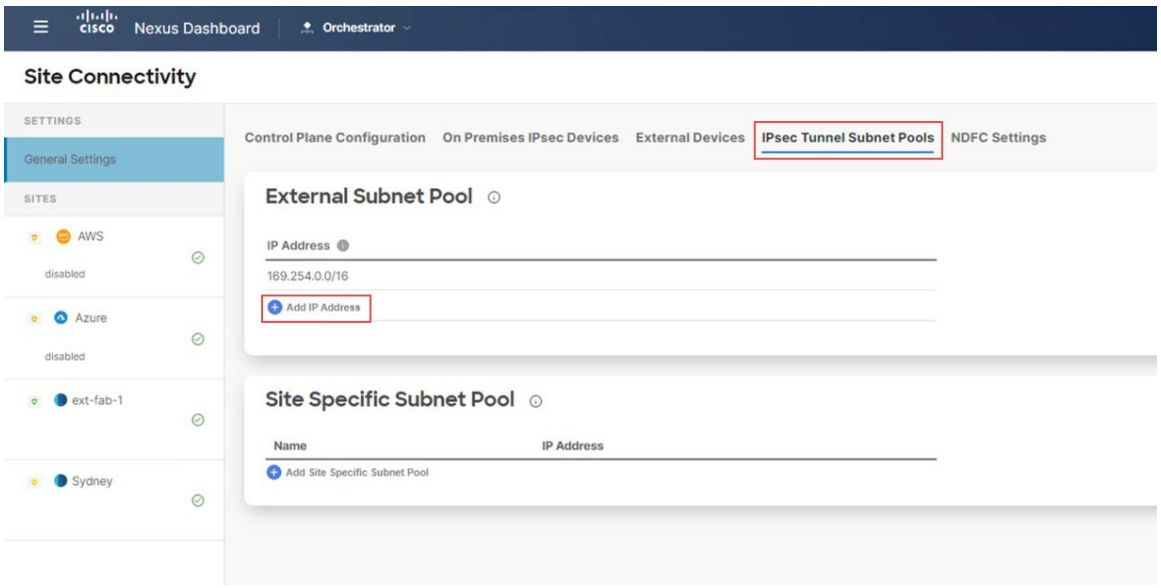

Step 7 Enter the IP subnet pool that you will use for the IPsec tunnels.

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 74:

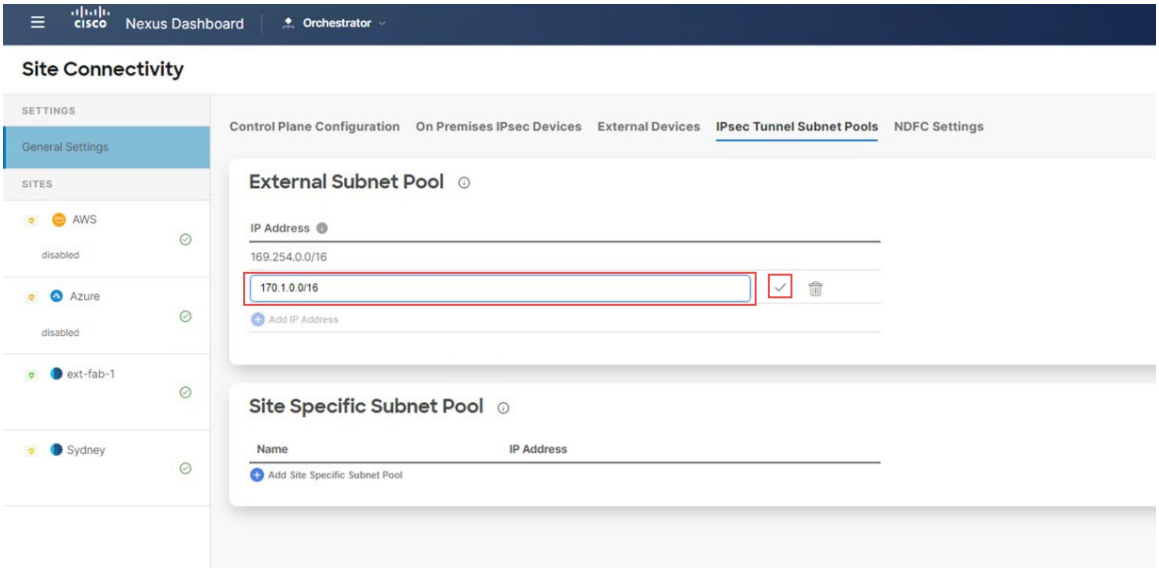

Step 8 Click the checkbox to accept the IP subnet pool that you entered.

The IP subnet pool appears under the **External Subnet Pool** area.

Figure 75:

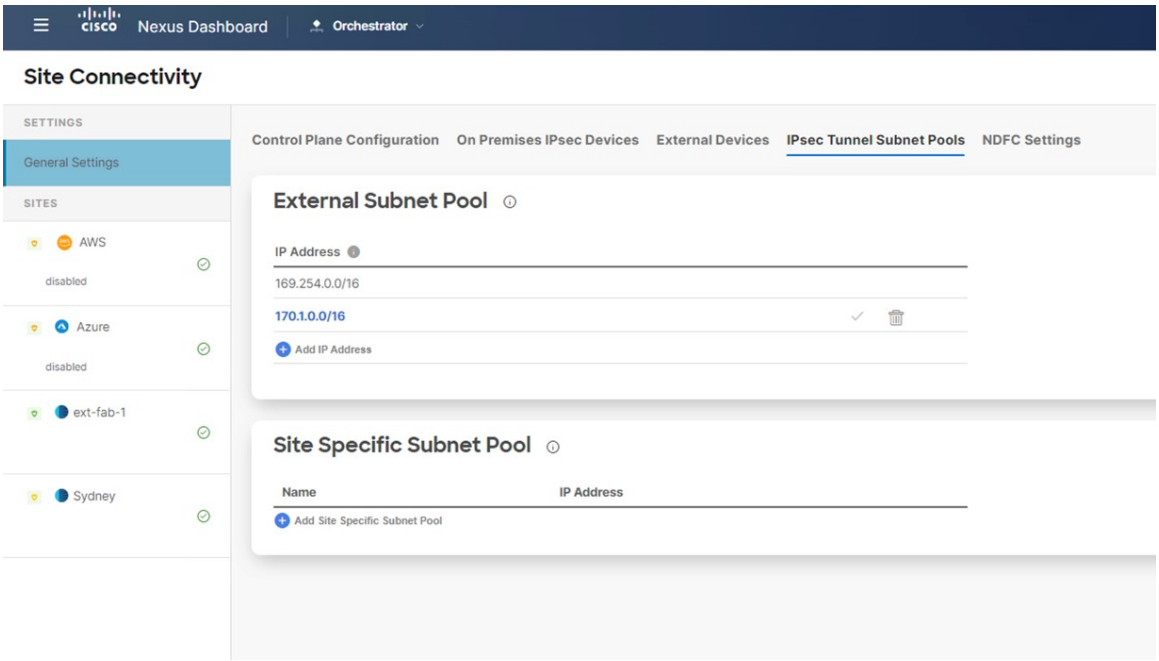

Step 9 Click the **NDFC Settings** tab and enter the necessary information in the **Auto Route Target Prefix**, if necessary.

Figure 76:

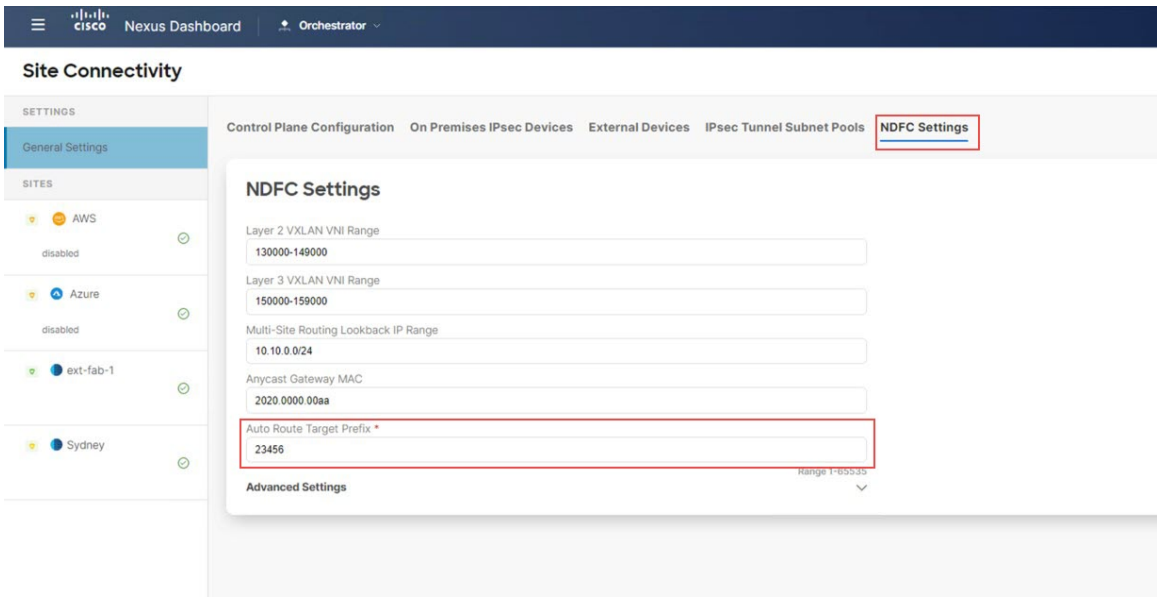

Under NDFC settings in NDO, the Route Target Prefix for the Route Target generation is set with a default value of `23456` for NDFC (Cloud Network Controller has different values for this setting), so you can change this value in the **Auto Route Target Prefix** field if required to avoid any possible 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 78.

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 71.

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 77:

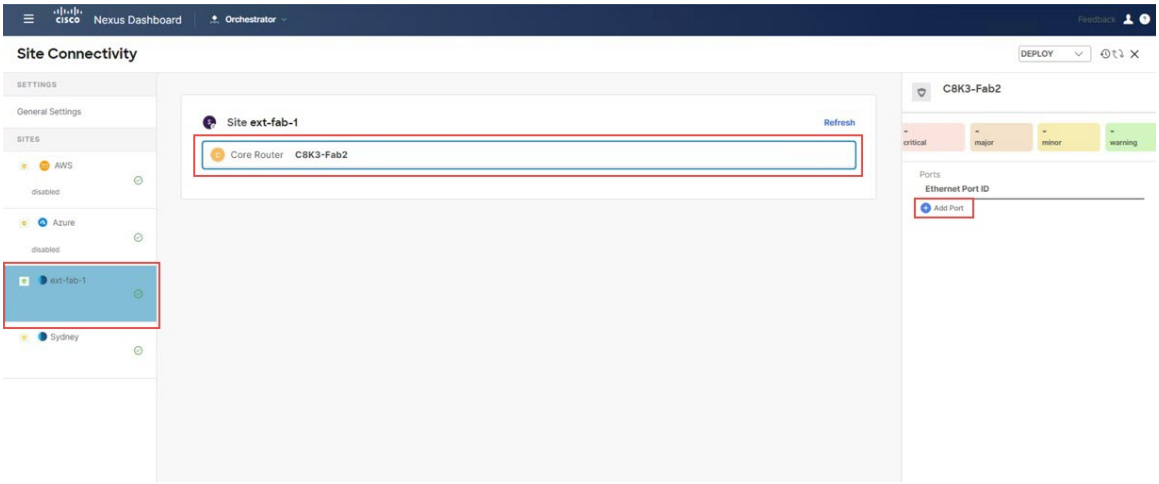

Step 4 Enter the necessary information for the port configuration, including the IP address, remote IP address, and remote ASN.

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, on page 20 in Supported Topologies with IPsec (Multi-Cloud), on page 18, 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 83.

Figure 78:

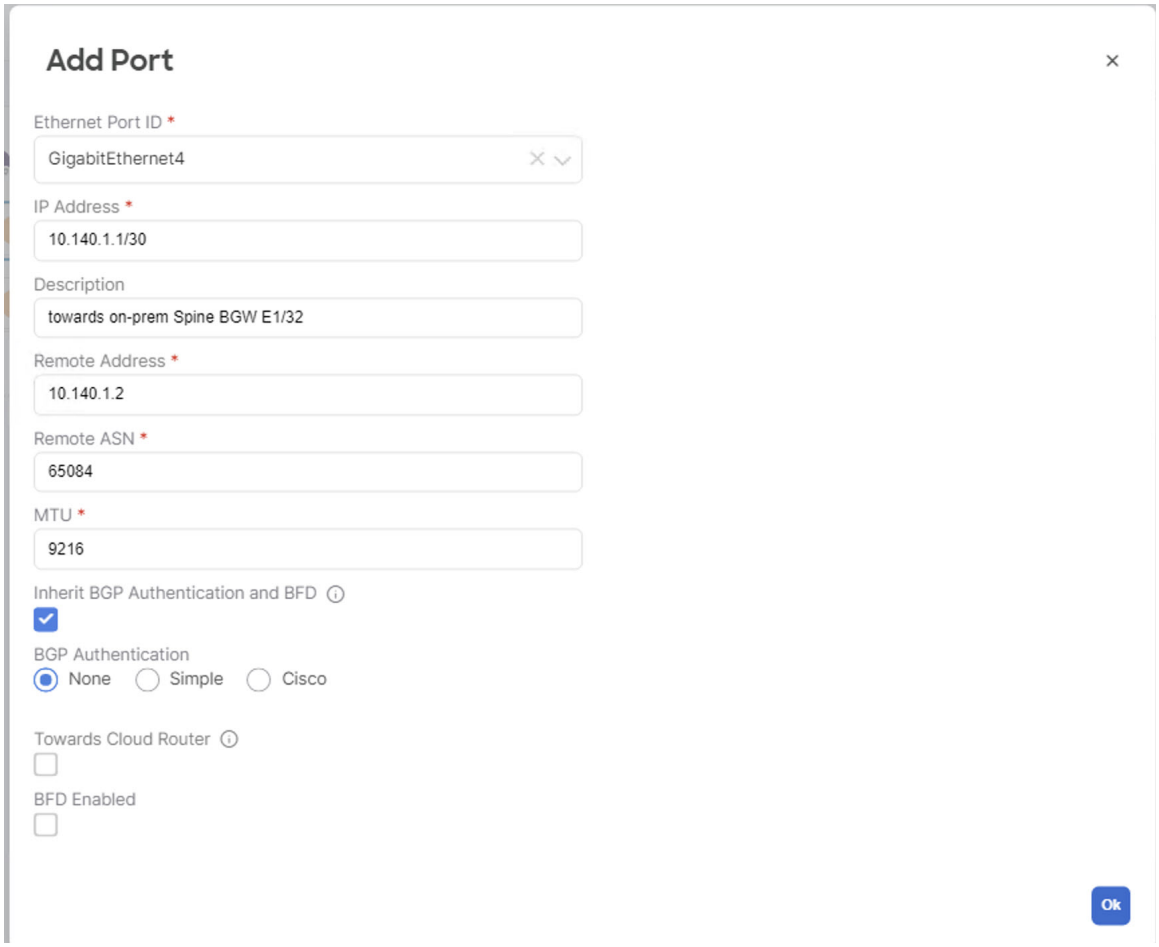

Step 5 Click **Ok** when you are finished.

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 80.

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 78.

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 79:

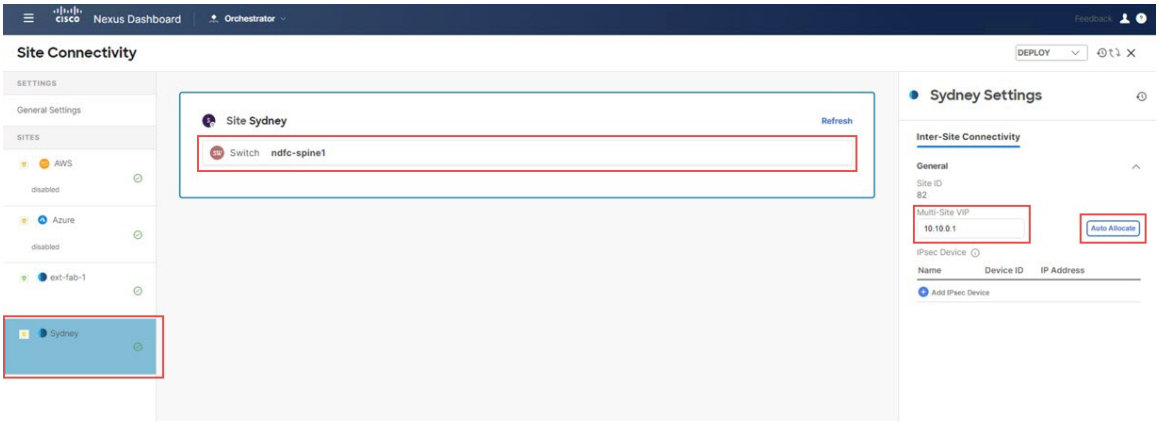

What to do next

Follow the procedures provided in Map the IPsec Device to the VXLAN Fabric Site, on page 81.

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 80.

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 80:

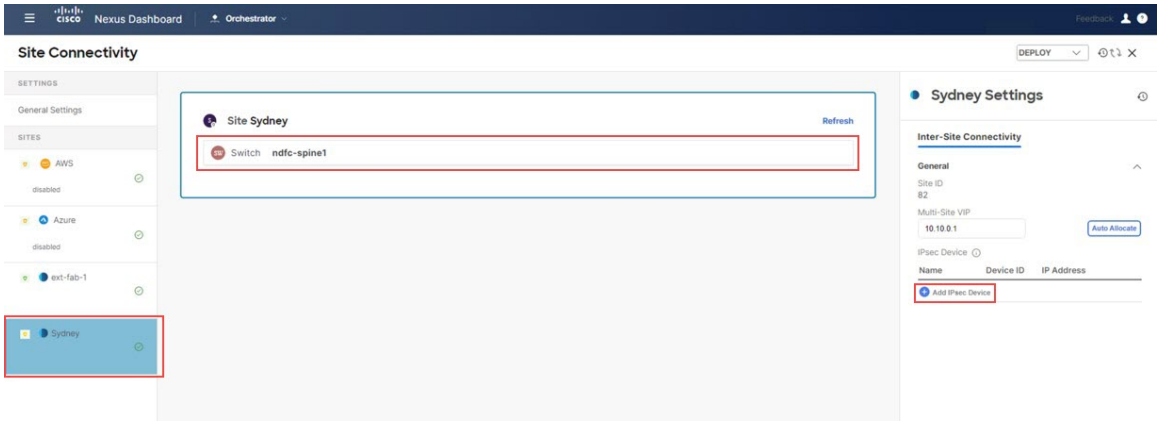

Step 4 Click **Select**, then choose the appropriate IPsec device.

Figure 81:

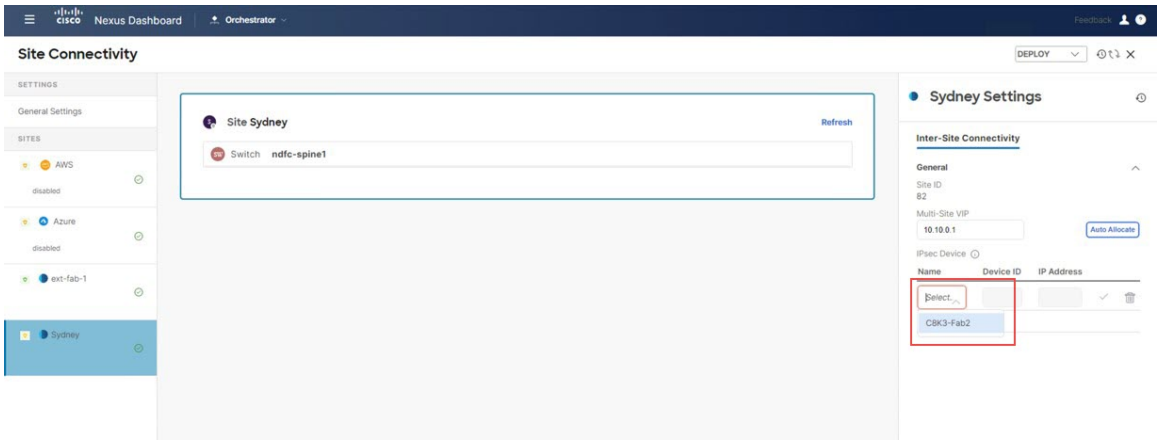

The on-premises IPsec device is now mapped to the VXLAN fabric site.

Figure 82:

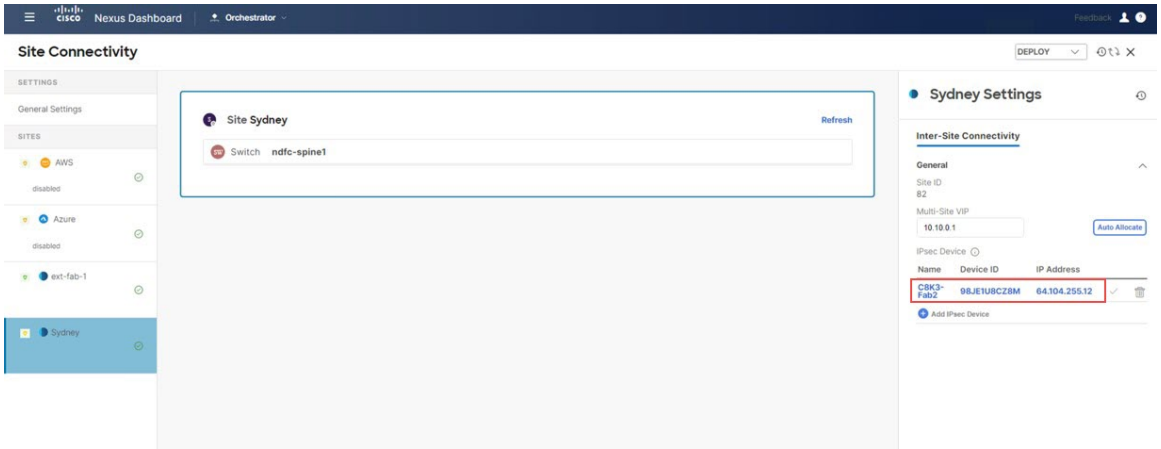

Step 5 Repeat this step for each on-premises IPsec device (Cisco Catalyst 8000V) that will be used to connect the NDFC VXLAN site to the cloud sites.

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 83.

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 81.

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 83:

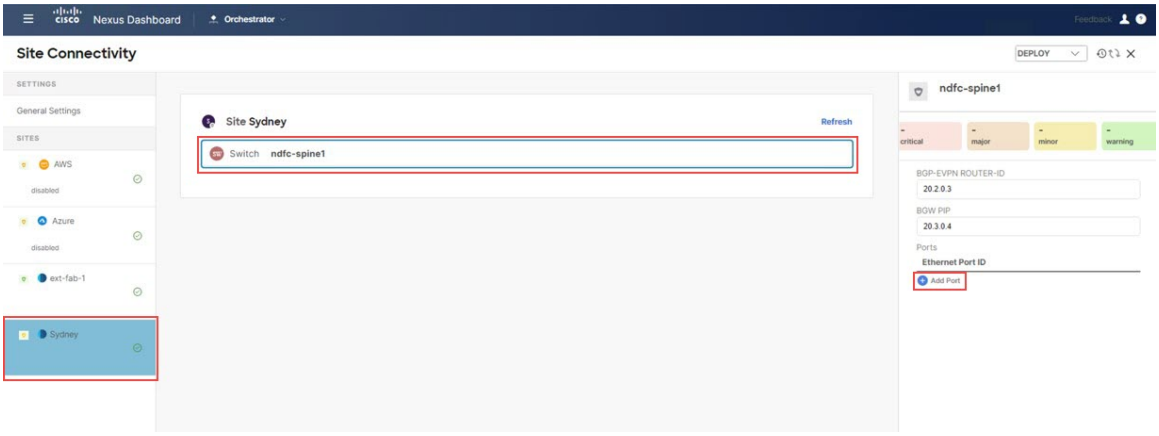

Step 4 Enter the necessary information in this page.

Define the port parameters in this page.

Figure 84:

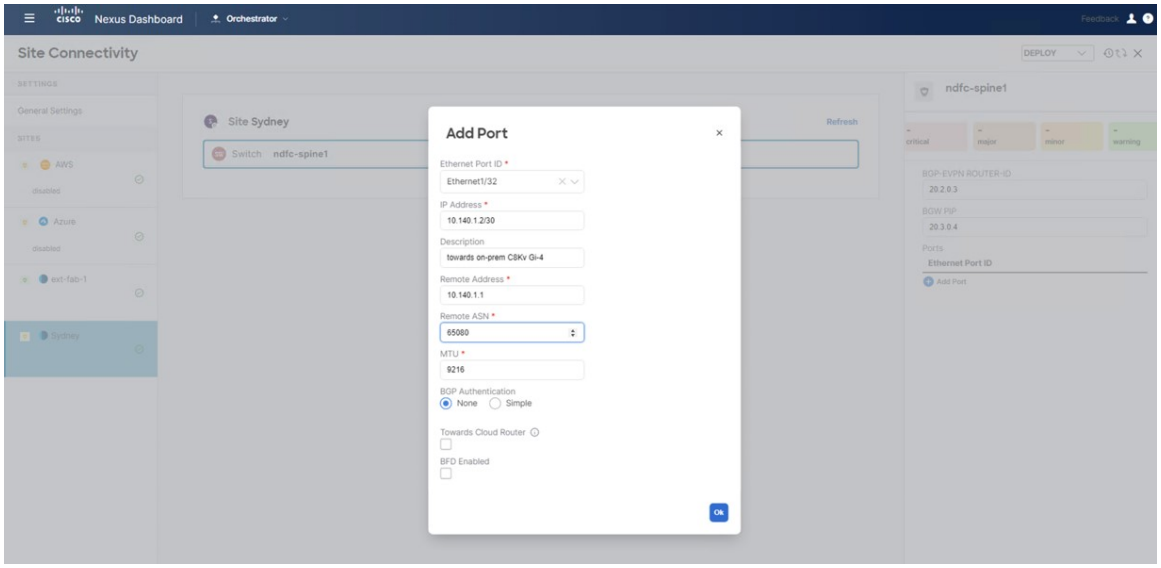

- In the **Ethernet Port ID** field, select the interface that is facing toward the on-premises Cisco Catalyst 8000V.
- In the **IP Address** field, enter the IP address for this interface. Later in these procedures, Nexus Dashboard Orchestrator will configure this IP address for this interface on the 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, on page 20 in Supported Topologies with IPsec (Multi-Cloud), on page 18.

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 85:

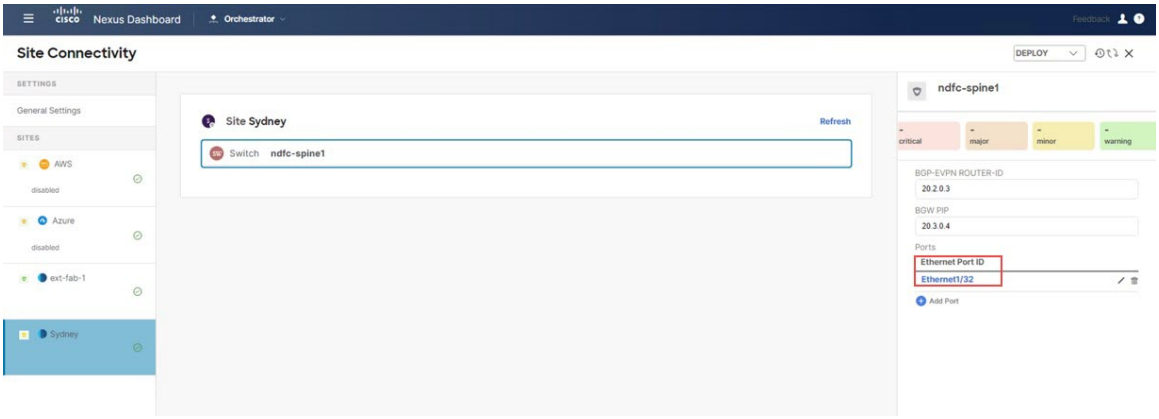

What to do next

Follow the procedures provided in Connect the First Cloud Site to the NDFC VXLAN Fabric Site, on page 85.

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 83.

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**.

Figure 86:

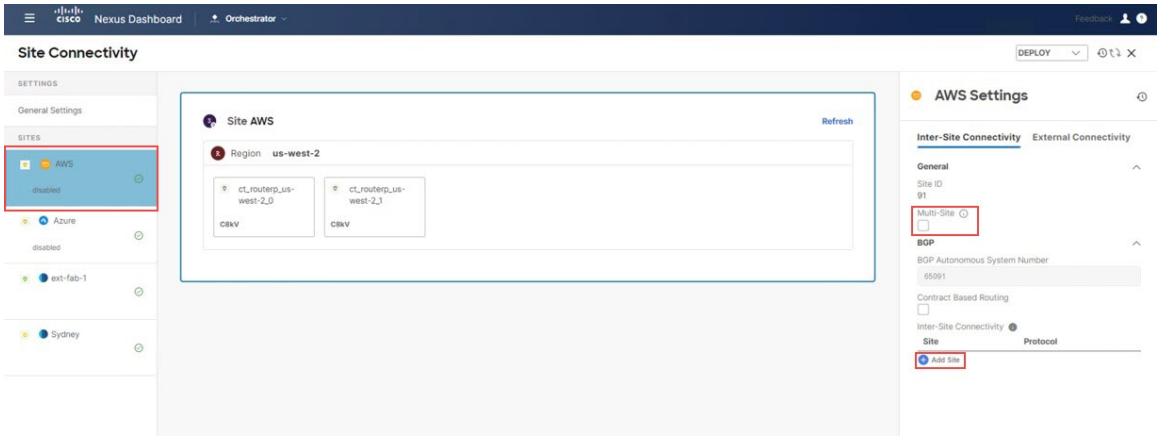

The **Add Site** page appears.

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

Figure 87:

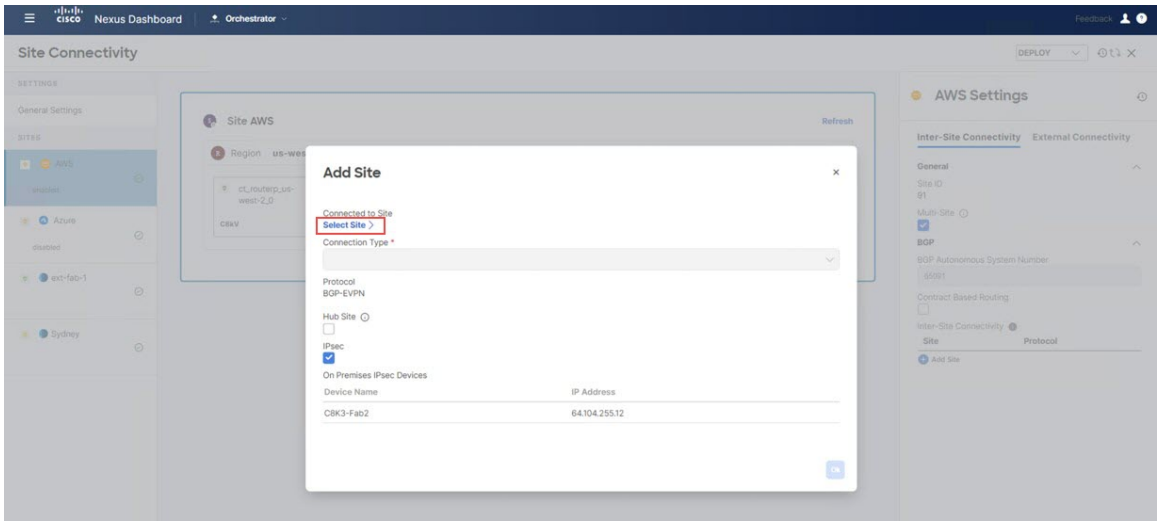

The **Select a Site** page appears.

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

Figure 88:

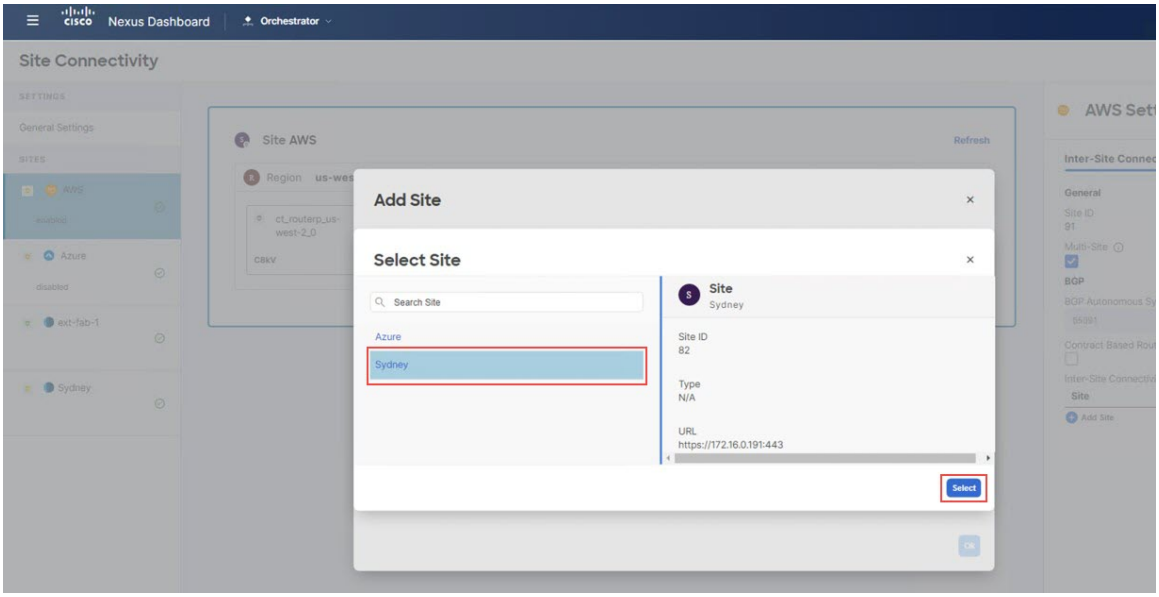

You are returned to the **Add Site** page.

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 89:

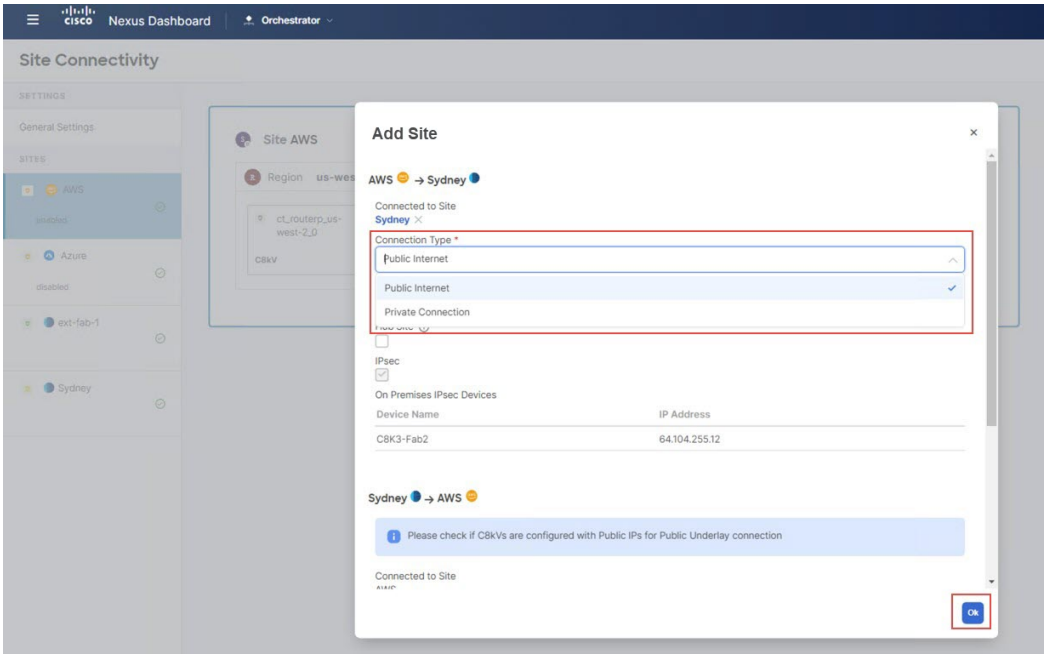

You can select **Public Internet**, or you can select a **Private Connection** if you are using Direct Connect with AWS or ExpressRoute with Azure.

- 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, on page 20 in Supported Topologies with IPsec (Multi-Cloud), on page 18.

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 88.

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 85.

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 90:

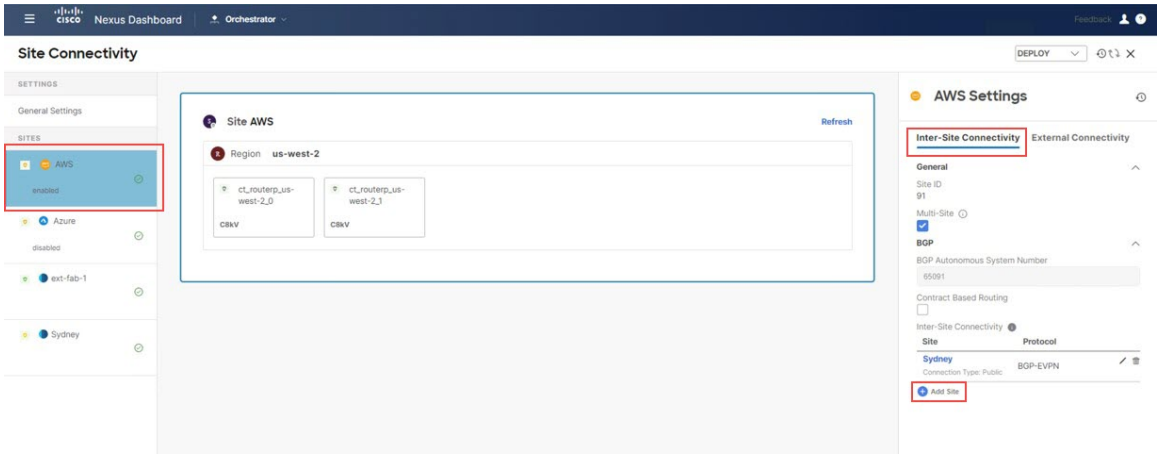

The **Add Site** page appears.

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 91:

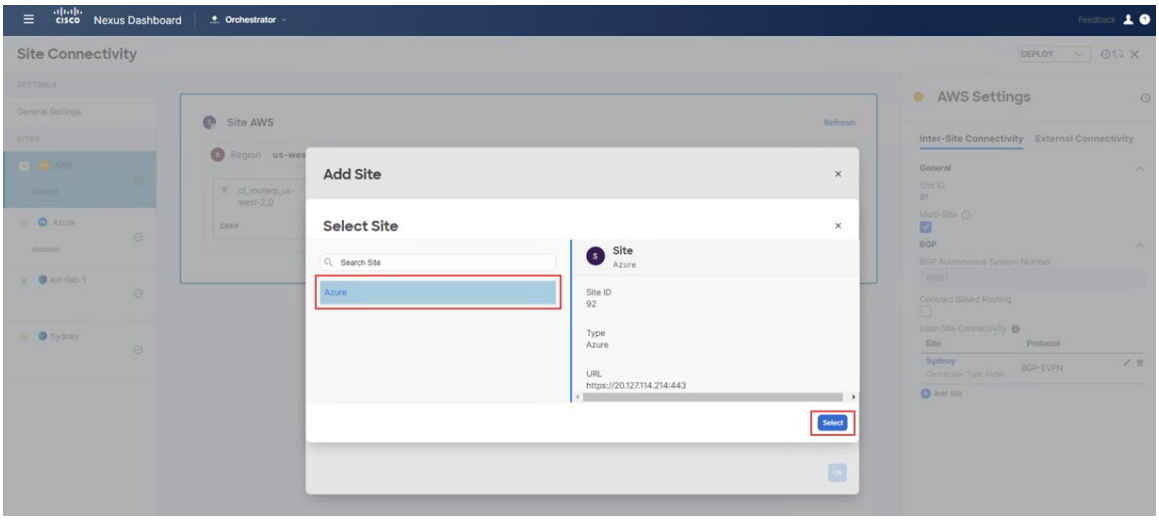

You are returned to the **Add Site** page.

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 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 92:

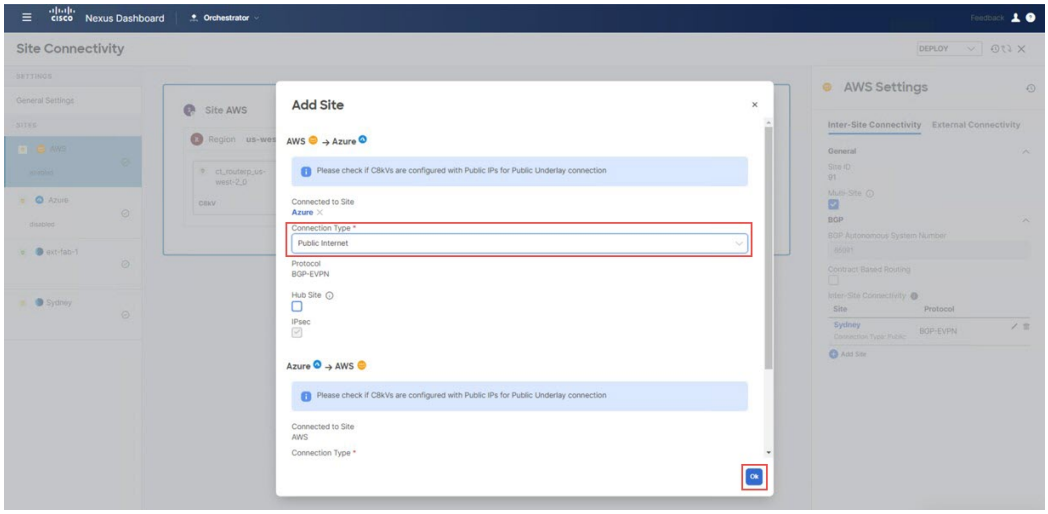

When the **Public Internet** connection type is selected, the **IPsec** option is mandatory and is automatically enabled for that connection type, whereas IPsec is optional for the **Cloud Backbone** type.

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 90.

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 85.
- Connected the first cloud site to the second cloud site in Connect the First Cloud Site to the Second Cloud Site, on page 88.

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 88, 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 88.

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 93:

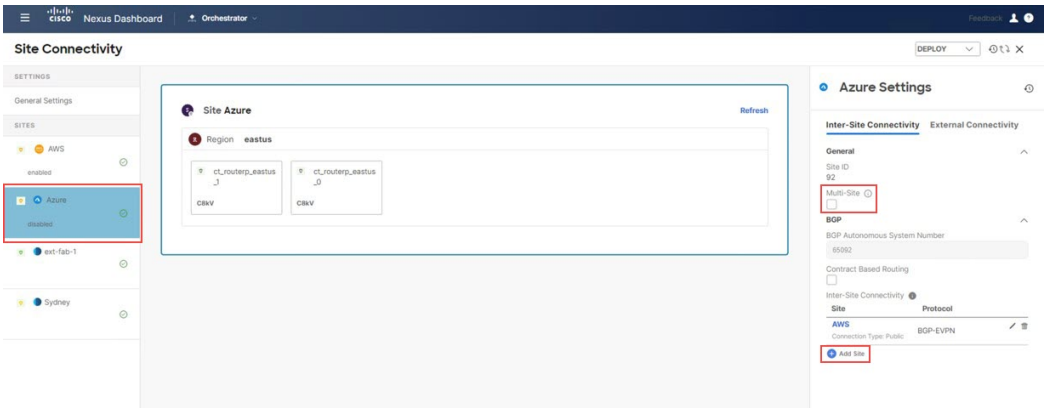

The **Add Site** page appears.

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 94:

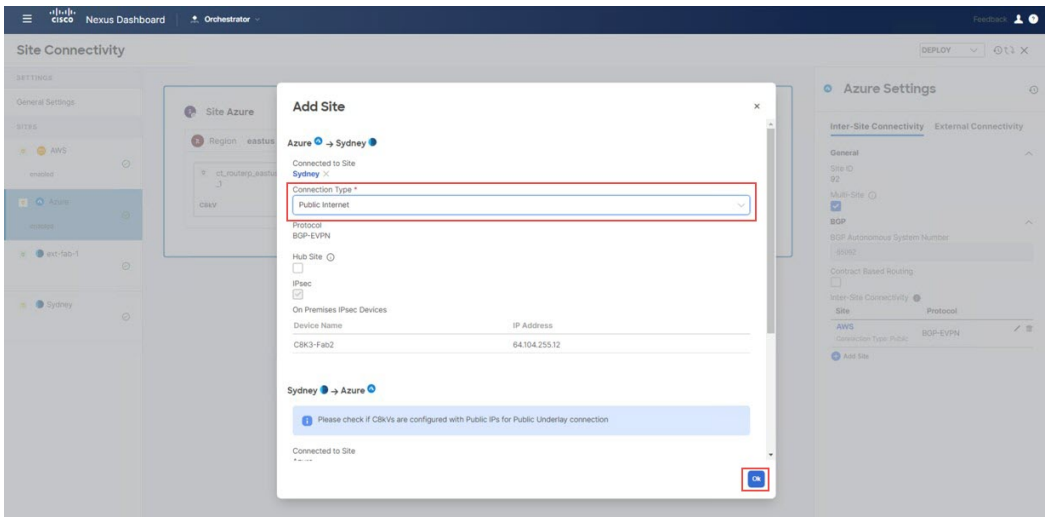

Step 7 When you have finished the configurations in this page, click **OK**.

The configured sites appear.

Figure 95:

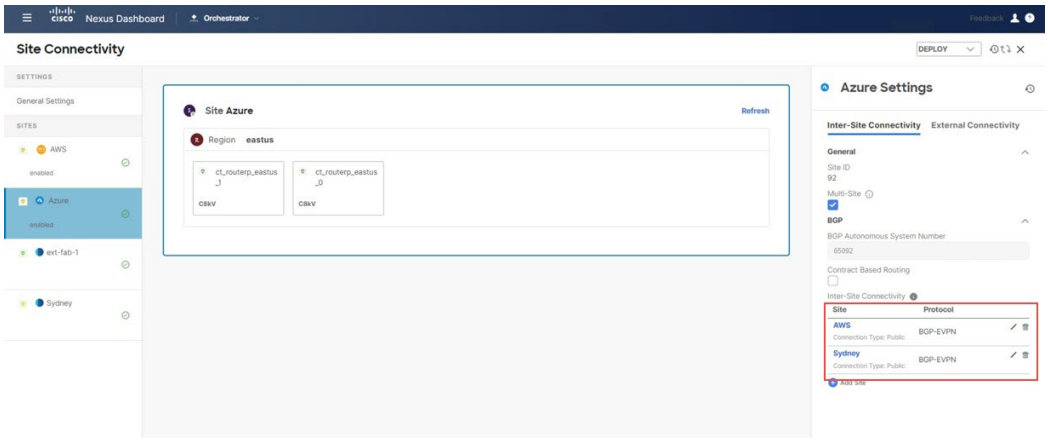

What to do next

Follow the procedures provided in Deploy the Configuration in Nexus Dashboard Orchestrator, on page 92.

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 90.

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 71, 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 71, at the top right of the page, click **Deploy** > **Deploy Only**.

Figure 96:

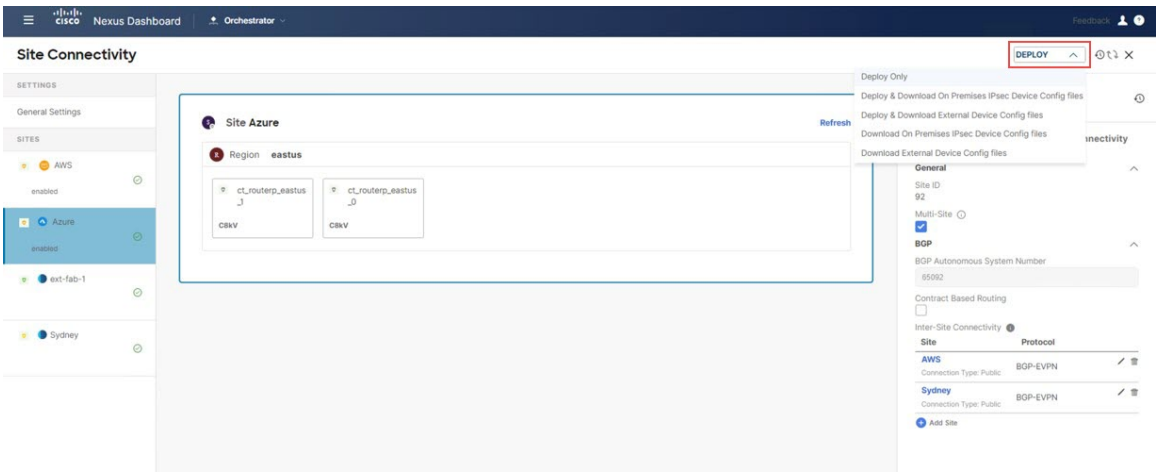

Step 2 Click **Yes** in the **Confirmation** window.

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 97:

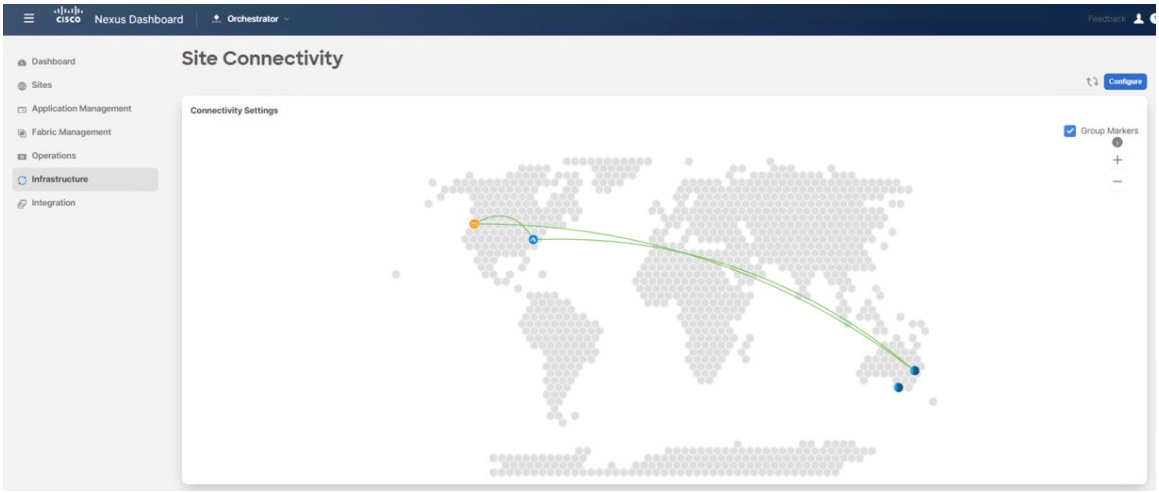

- In the same page, scroll down to the area for the first cloud site (for example, the AWS site), click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

  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 98:

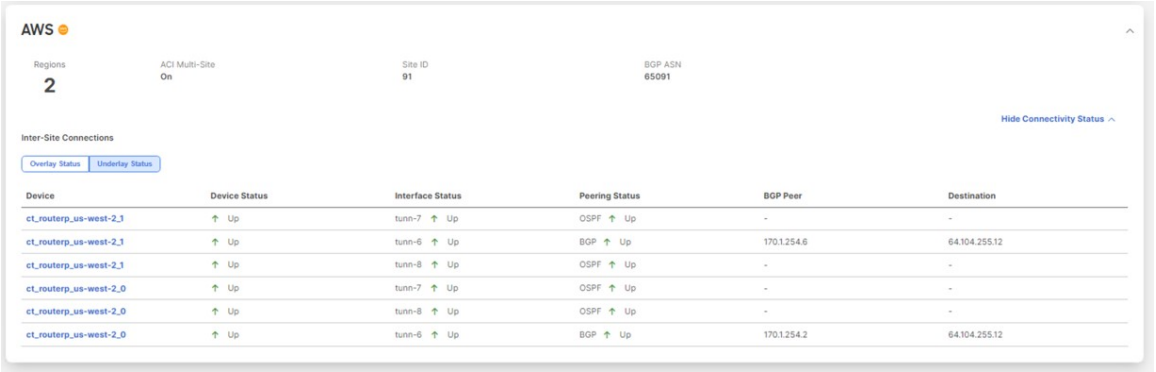

- Scroll down to the area for the second cloud site (for example, the Azure site), click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

  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 99:

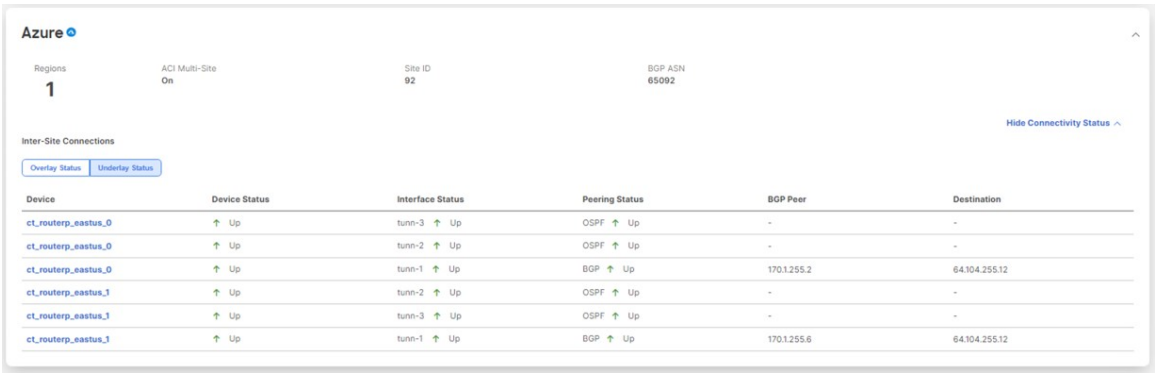

- Scroll down to the area for the NDFC external fabric site, click **Show Connectivity Status**, then click **Underlay Status** in the **Inter-Site Connections** area to verify the underlay status.

  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 100:

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

Figure 101:

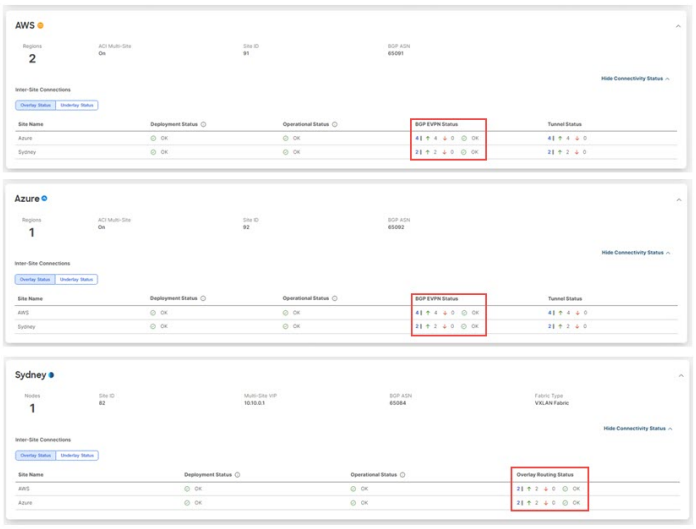

- Return to the NDFC screen and verify the hybrid cloud connectivity in the **Topology** screen. In the following example, you can see the NDFC VXLAN fabric site (the Sydney site) connected to the first and second cloud sites (the AWS and Azure cloud sites).

Figure 102:

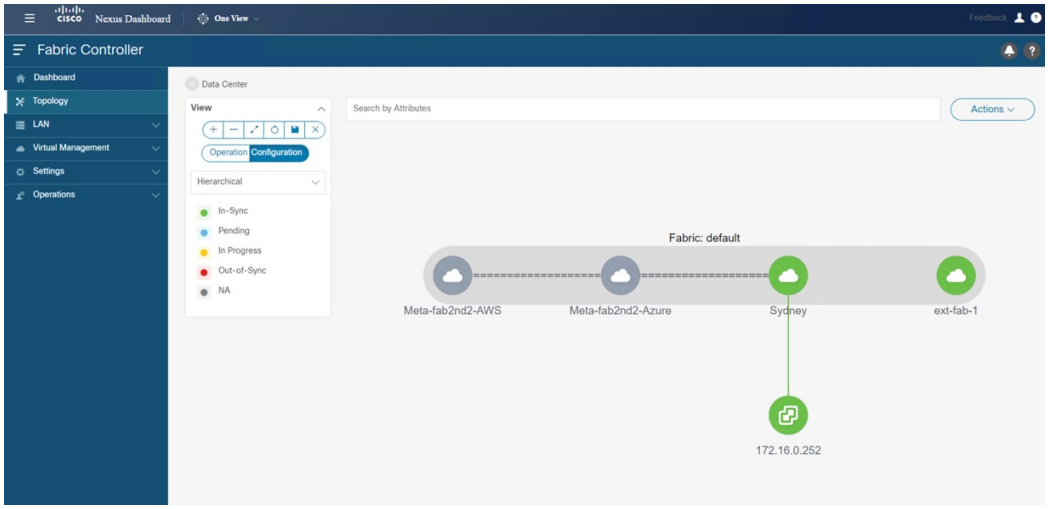

PART II

Use Cases

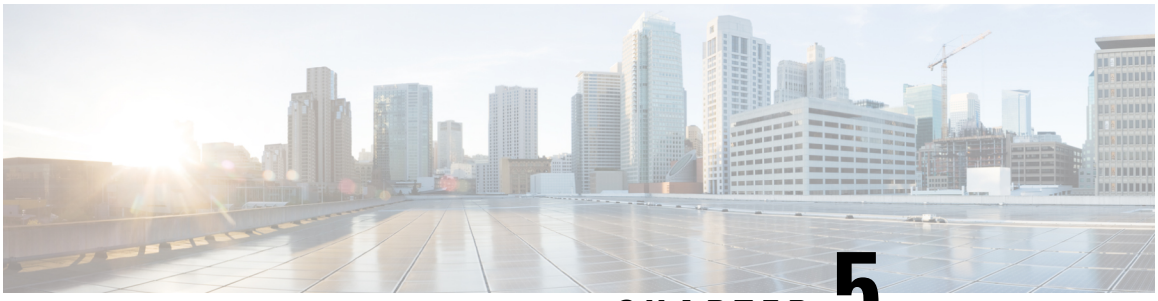

CHAPTER 5

Deploying the Tenant

Deploying the Tenant

Once the underlay and overlay connectivity is established between the sites, you must then deploy the endpoint network/VPC/VNet to establish communication between tenant endpoints deployed in the on-premises and in the cloud sites.

NDO uses the notions of schemas and templates for defining VRFs and networks. In the context of NDFC, VRFs are used to isolate one tenant from another. All the endpoint networks (subnets) of one tenant are mapped to the respective VRF. The same notion of VRFs can also be extended to the cloud, where a VRF corresponds to a VPC in AWS and a VNet in Azure.

The following procedures for deploying the tenant applies to all the topologies previously described and leverage the specific infra config deployed, and also applies for any of the following use cases.

Note NDO has a pre-built `dcnm-default-tn` tenant, which can be associated with on-premises sites as well as cloud sites. We recommend that you associate this pre-built `dcnm-default-tn` tenant with the NDFC and cloud sites when deploying hybrid cloud connectivity, but you can also create your own tenant from scratch, if necessary.

Step 1 In NDO, navigate to **Application Management** > **Tenants**.

Figure 103:

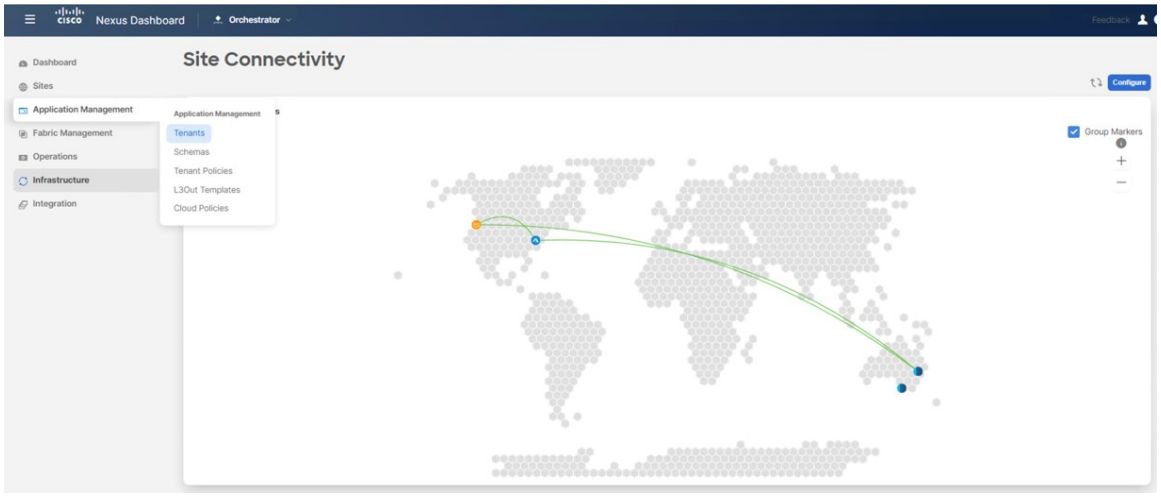

The Tenants window appears.

Figure 104:

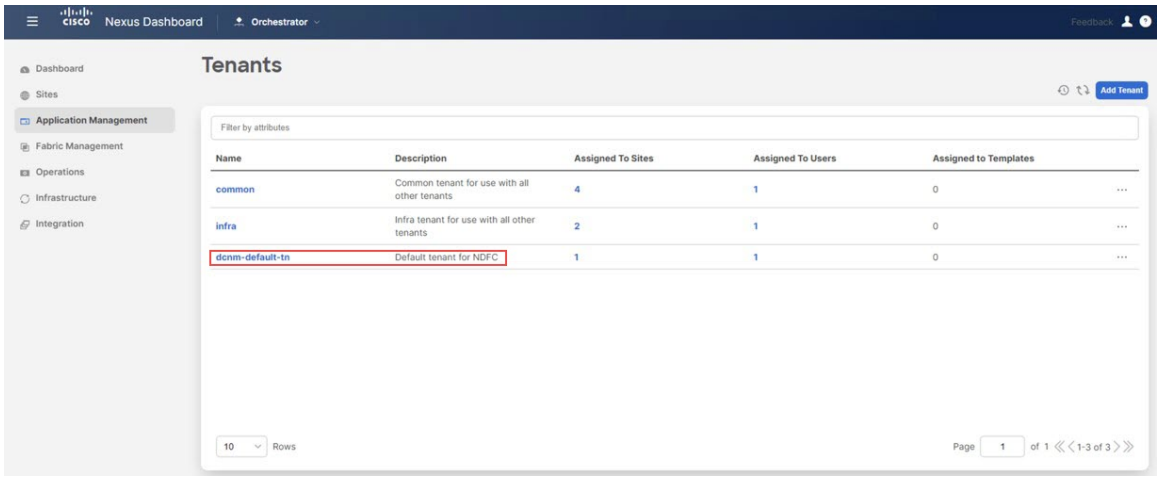

Step 2 Click the `dcnm-default-tn` tenant.

The **Update Tenant** page for the `dcnm-default-tn` tenant appears.

Figure 105:

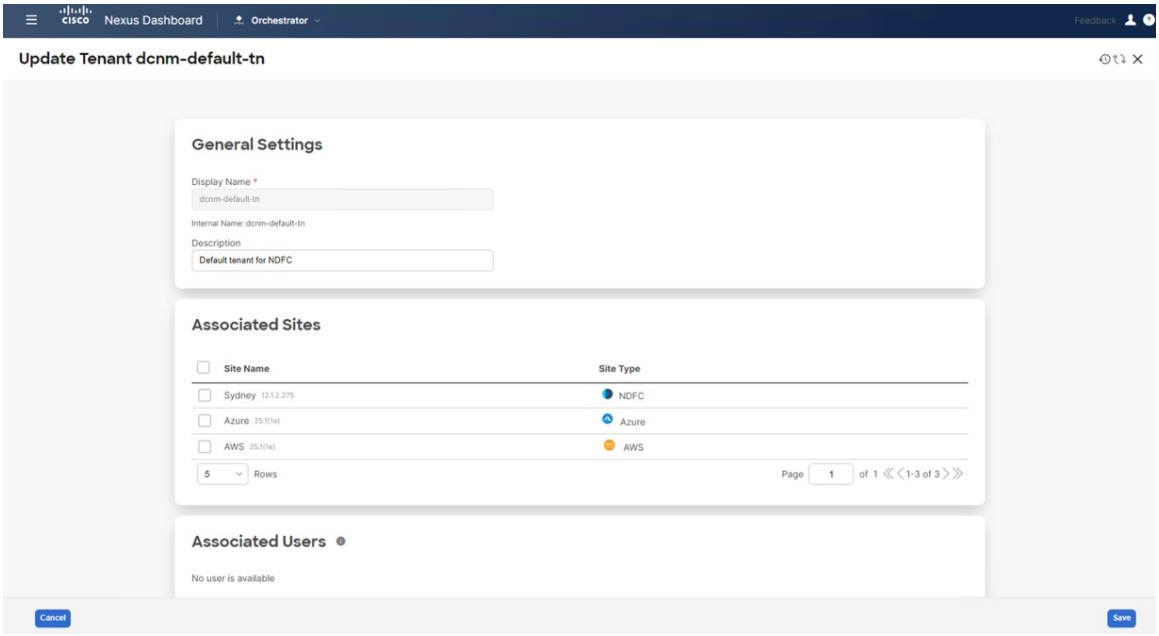

Step 3 Select the sites shown in the screen.

Note that the external fabric site does not appear in the list. The external site is only used to provide connectivity between the on-premises site to the cloud sites and there are no end hosts in the external fabric, so no tenant deployment required for the external fabric.

Figure 106:

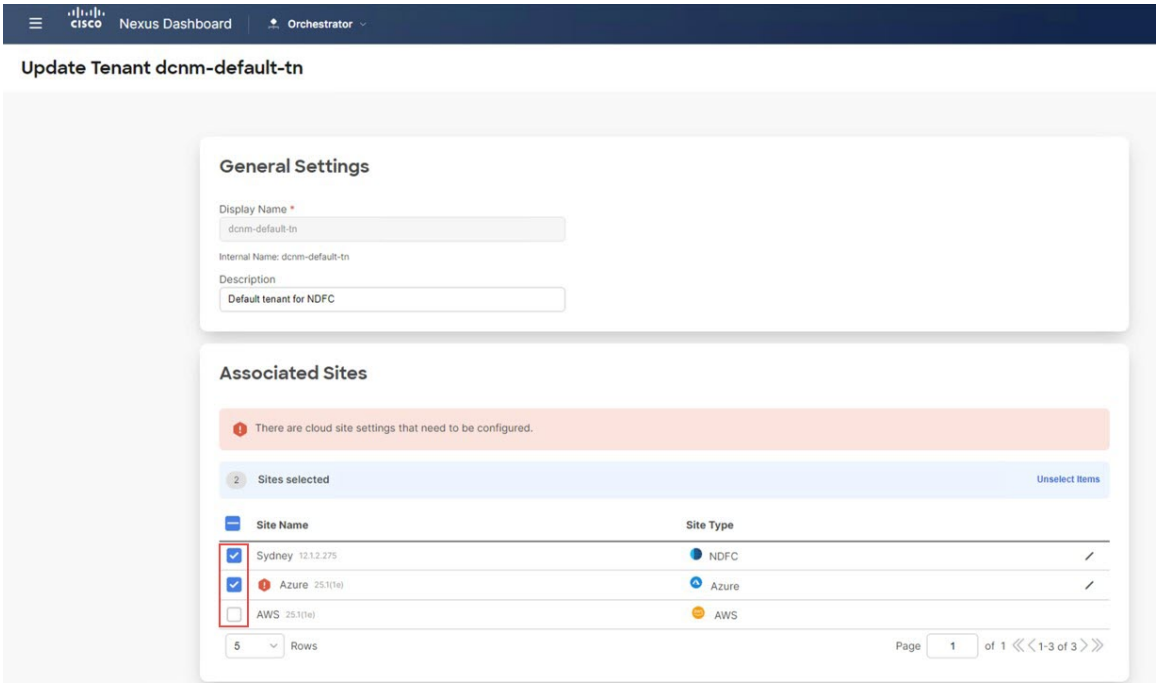

Step 4 For the cloud sites, click the Edit button (the pencil icon) and provide the necessary information for each cloud account.

You need an additional account for AWS for the user tenant, but for Azure, you can use the same subscription as the Azure infra tenant.

- For example, after clicking the Edit button for the AWS cloud site, in the **AWS Account Setting** area, you might click **Trusted** for the **Access Type** and enter the associated AWS account ID in that field.

Figure 107:

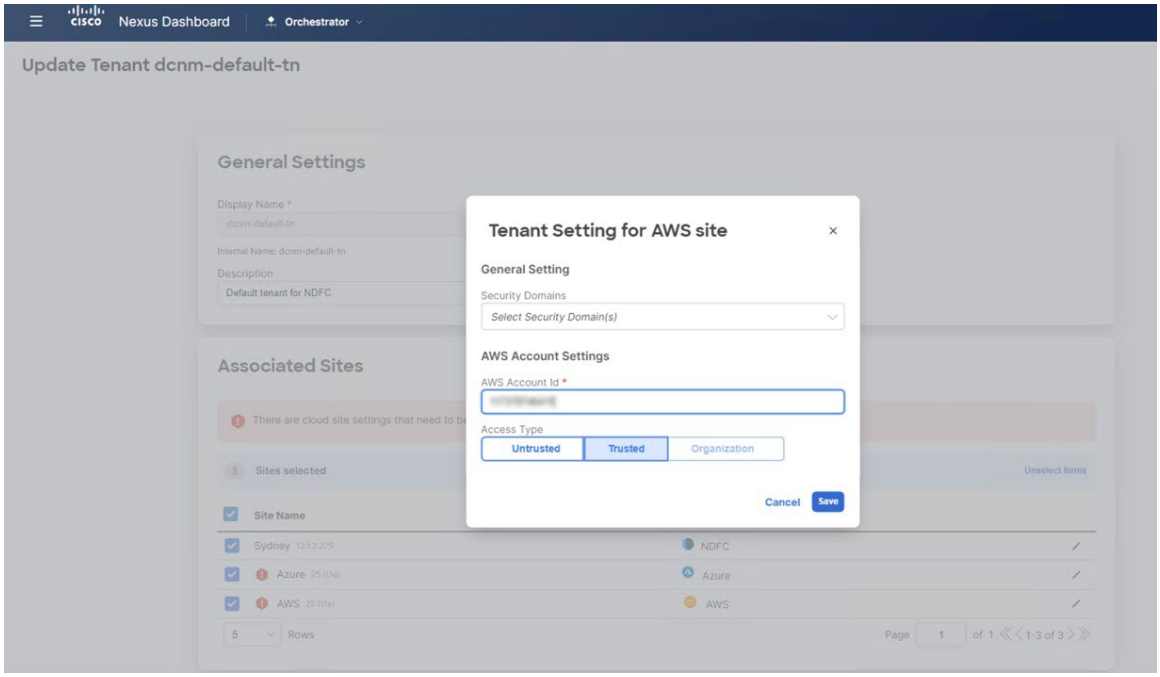

See the section "Setting Up the AWS Account for the User Tenant" in the *Cisco Cloud Network Controller for AWS Installation Guide*, Release 25.1(1) or later, for more information on the different access types for the tenants in AWS.

- Similarly, after clicking the Edit button for the Azure cloud site, you would enter the necessary information, depending on whether the tenant is managed or unmanaged.

Figure 108:

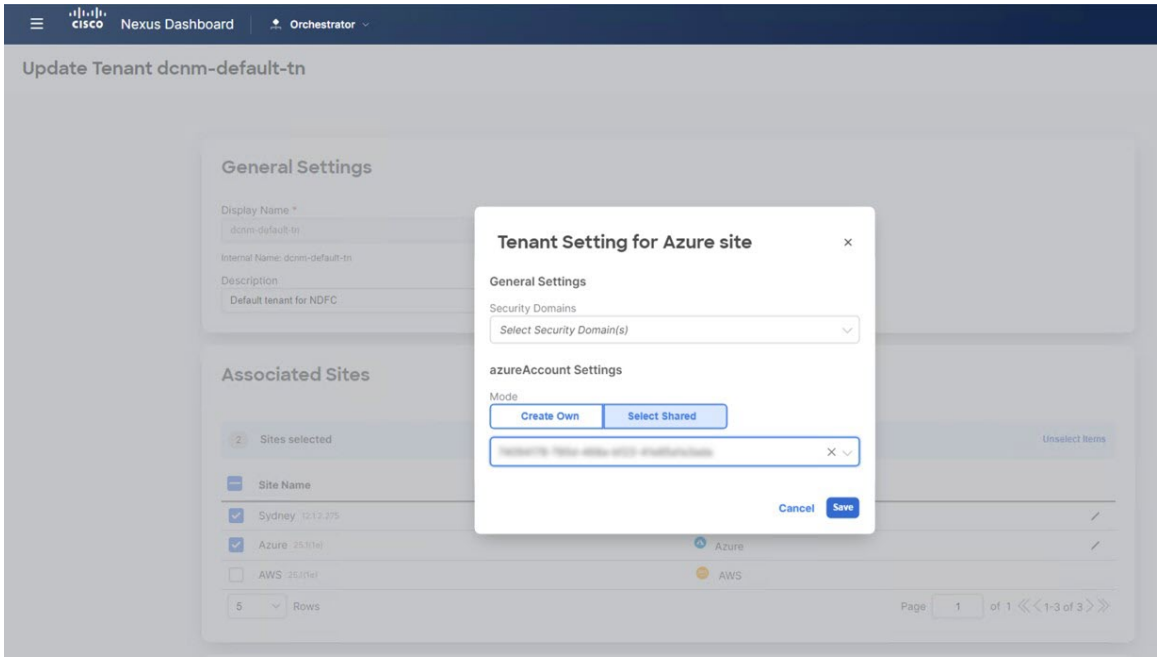

See the section "Adding a Role Assignment" in the *Cisco Cloud Network Controller for Azure Installation Guide*, Release 25.1(1) or later, for more information on the different access types for the tenants in Azure.

Step 5 Verify the tenants were deployed correctly.

For example, in the figure below, the `dcnm-default-tn` tenant has three sites mapped (one on-premises NDFC site and the two cloud sites).

Figure 109:

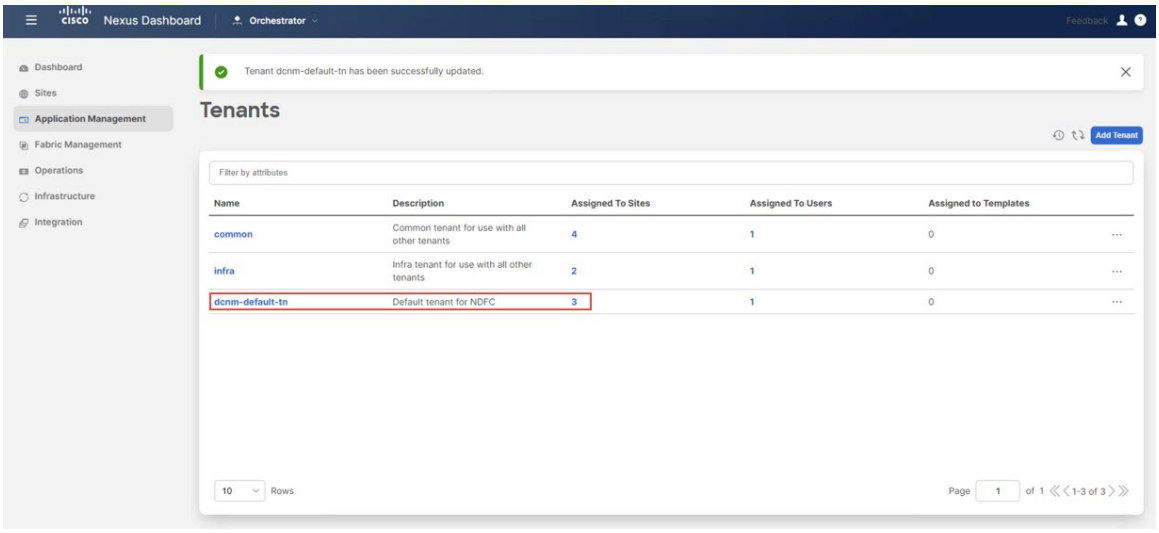

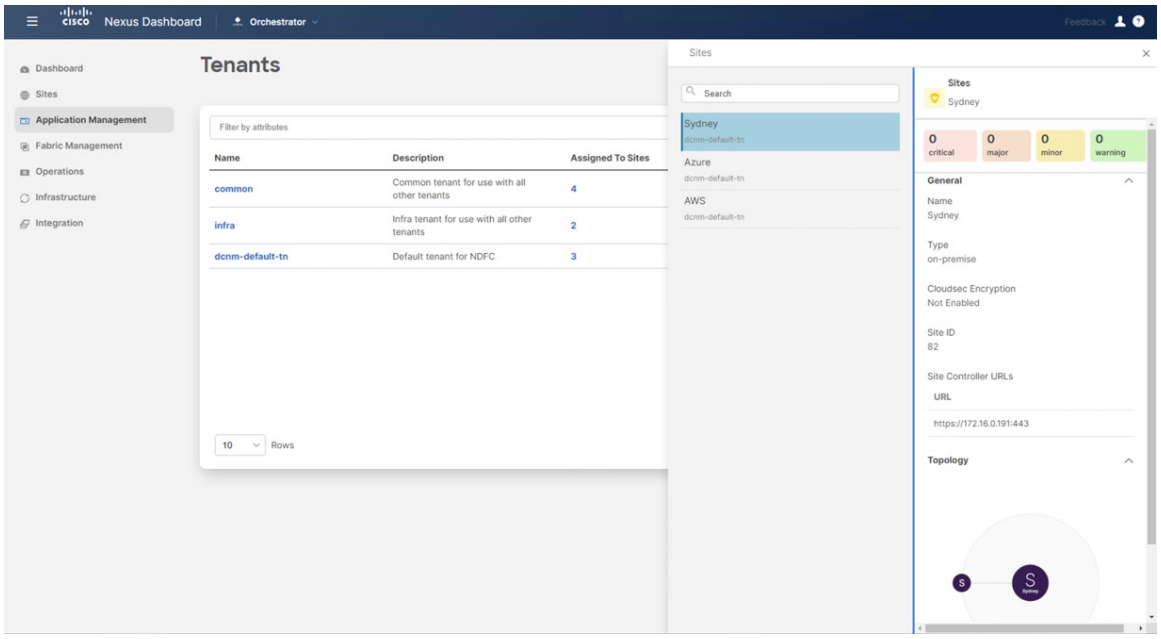

You can also check the `dcnm-default-tn` tenant deployed in the Cisco Cloud Network Controllers for the cloud sites.

Figure 110:

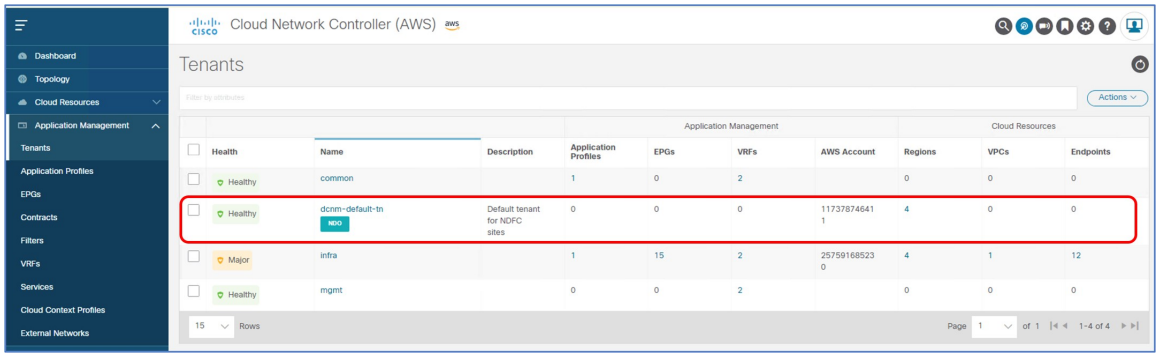

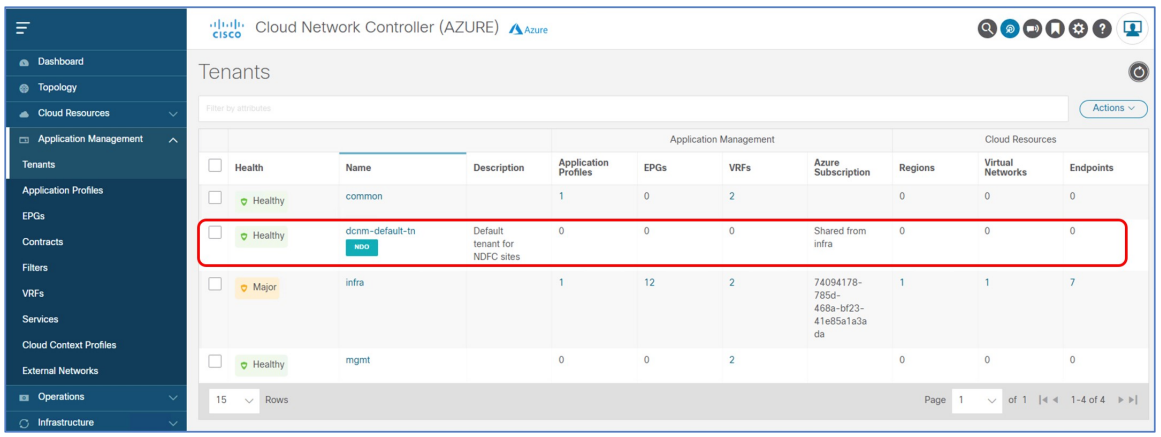

What to do next

Configure one or both of the following use cases:

- Stretched VRF Use Case, on page 107
- Route Leaking Use Case, on page 143

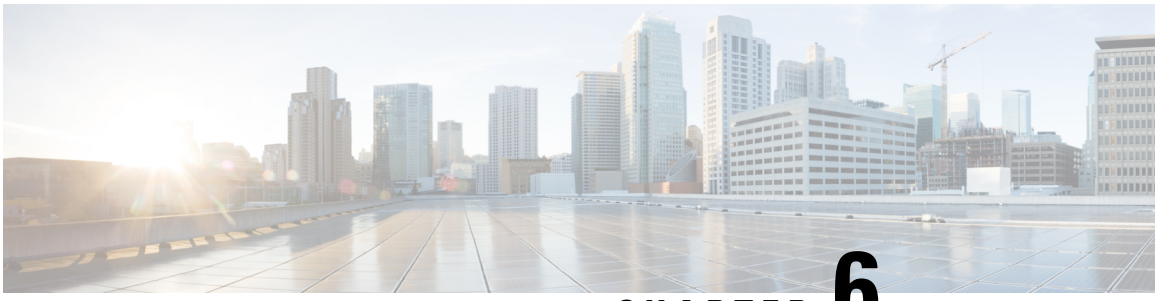

CHAPTER 6

Stretched VRF Use Case

About the Stretched VRF Use Case

Stretched VRF (intra-VRF) is a common use case where a single (common) VRF is defined in a template that is associated to all the sites (on-premises and cloud sites). A separate template is used to deploy networks for the on-premises site since it is not possible to stretch networks between on-premises and cloud sites.

Stretching the same VRF to all the sites enables the exchanging of prefixes between the sites without having the requirement of any additional routing configuration. CIDR blocks (used to provision subnets in cloud VPCs/VNets) are mapped to this stretched VRF.

Note Stretching a Layer 2 subnet across on-premises and cloud sites or between cloud sites is not supported.

The following figure shows two templates being created under the Demo schema:

- The `Stretched Template`, which defines the VRF to be deployed to all three sites. For cloud sites, we define the regions and CIDR blocks under the VRF.
- The `On_Prem Template`, which contains the networks to be deployed to the on-premises VXLAN fabric.

Figure 111:

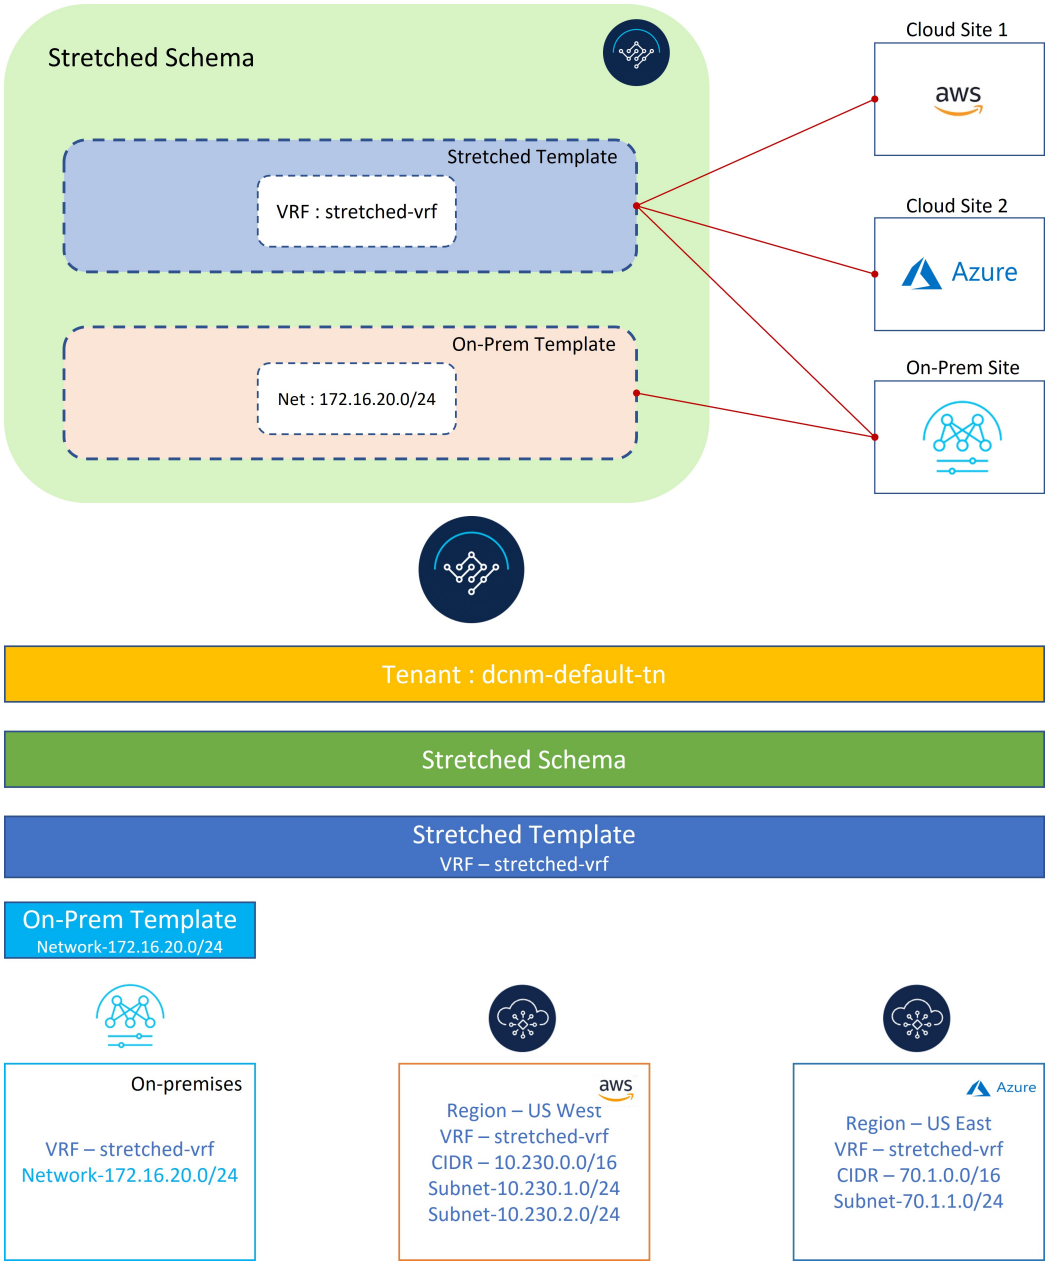

Configure the Stretched VRF Use Case

Step 1 In NDO, navigate to **Application Management** > **Schemas** and click **Add Schema**.

Figure 112:

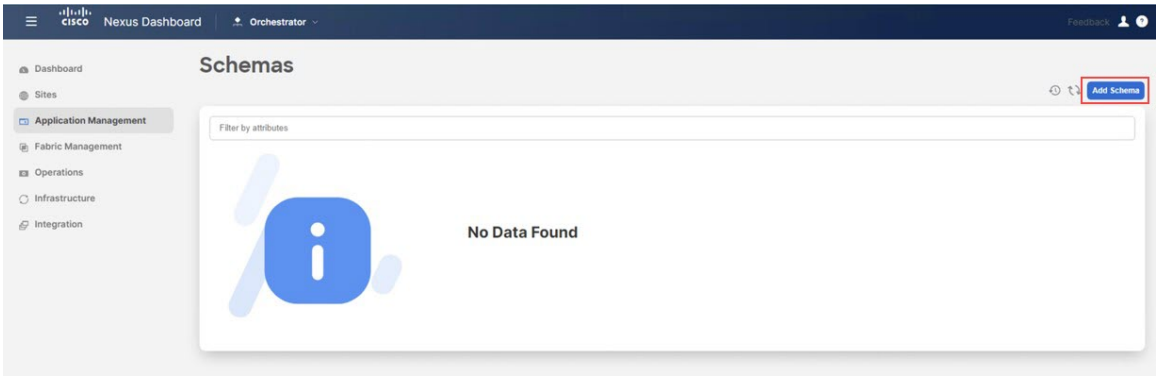

Step 2 Provide the schema name and click **Add**.

For this use case, we will name the new schema `Stretched Schema`.

Figure 113:

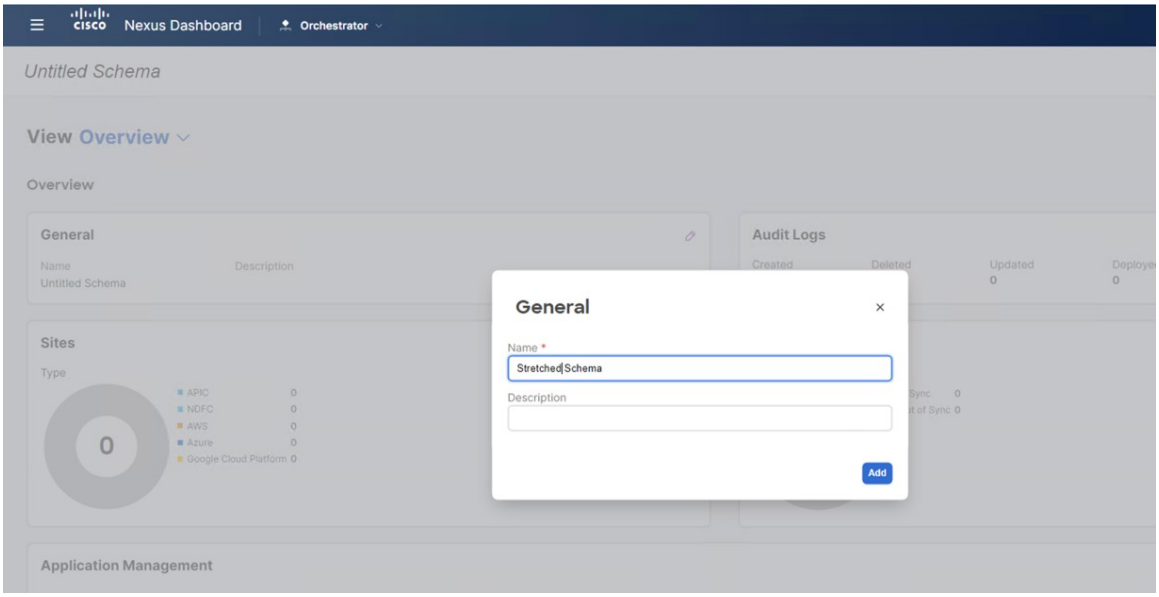

You are returned to the **Overview** page for the new `Stretched Schema` schema.

Step 3 Click **Add New Template**.

Figure 114:

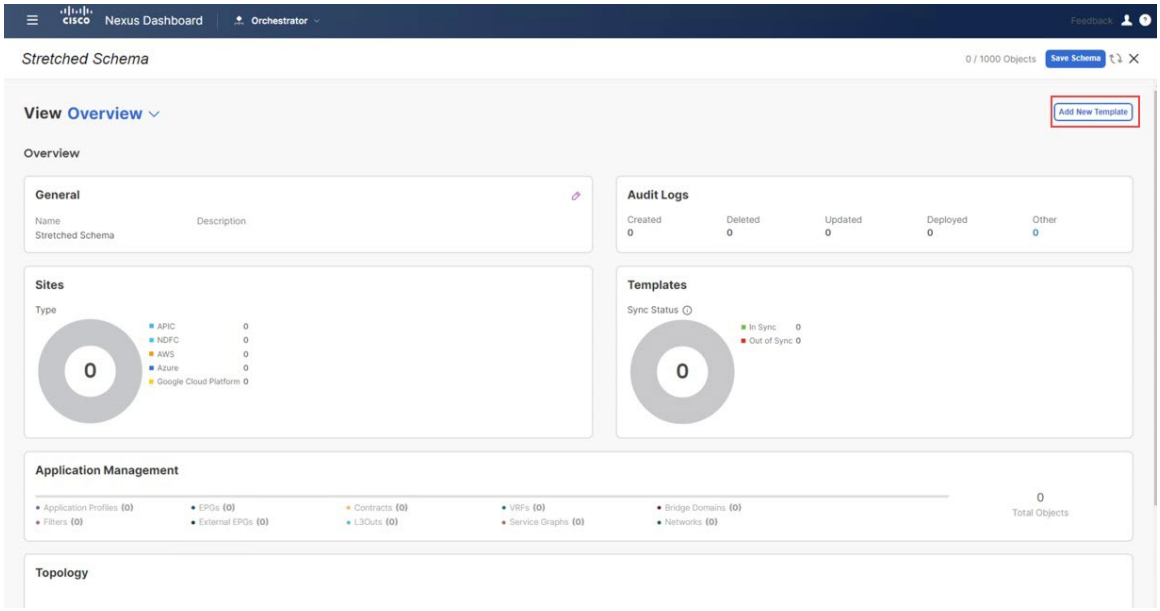

Step 4 Choose the NDFC template, then click **Add**.

You should use the NDFC template type for on-premises as well as cloud sites.

Figure 115:

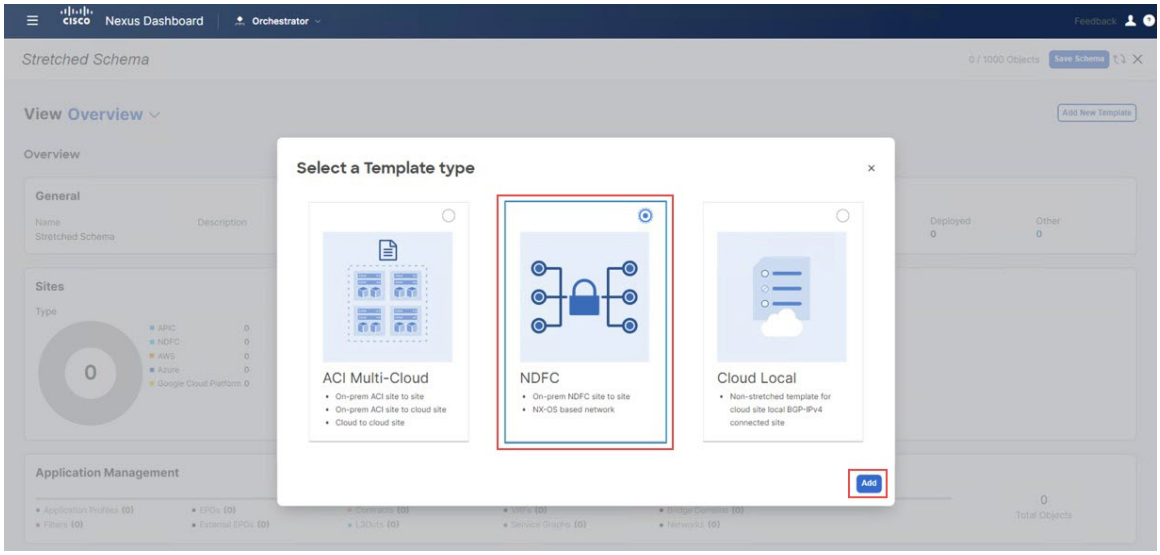

Step 5 Enter a name in the **Display Name** field to create an NDFC-type template (for example, `Stretched Template`) and select the `dcnm-default-tn` tenant in the **Select a Tenant** field to map the template to that tenant.

Figure 116:

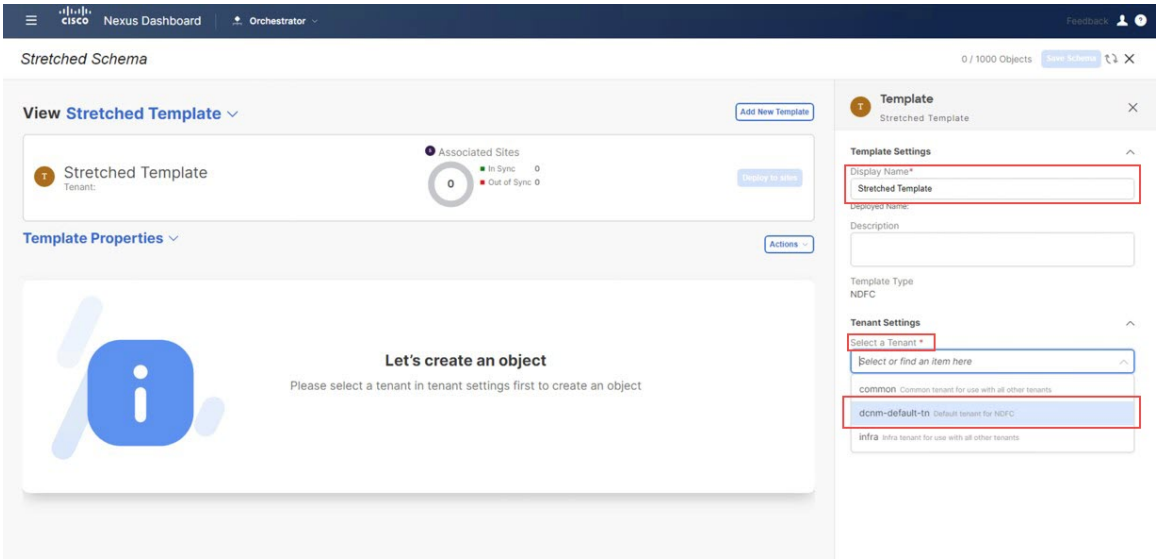

Step 6 Under **Template Properties**, click **Create Object** and choose **VRF** to create a VRF that will be stretched to all the sites.

Figure 117:

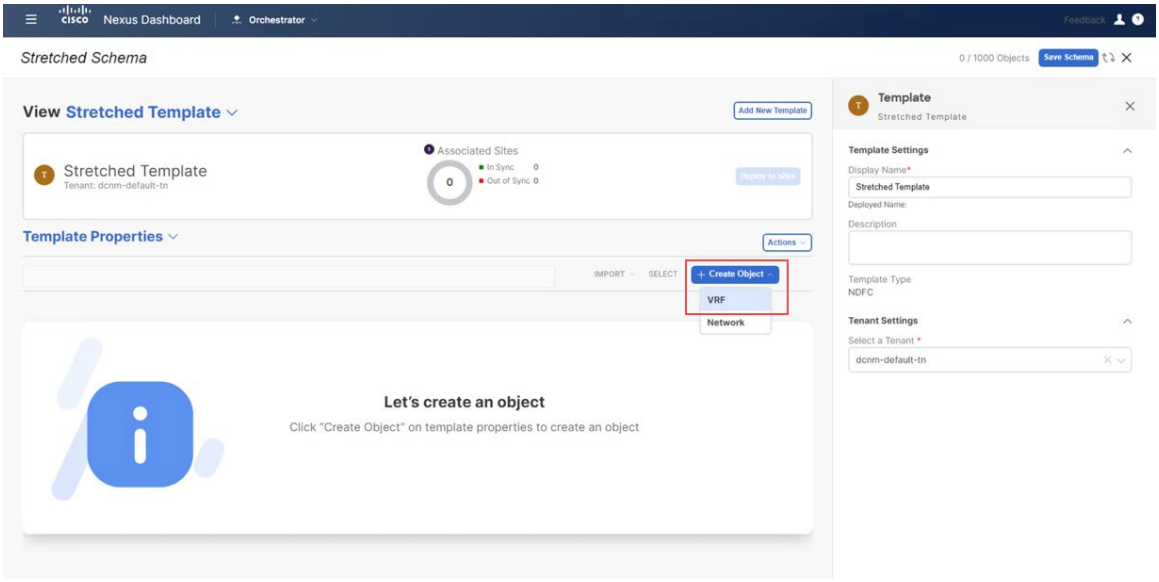

Note If you have an on-premises VRF already created that you want to use instead of creating a new VRF, under **Template Properties**, click **Import**, then import the already-created VRF.

Currently, we only support importing VRFs and networks from on-premises sites.

Step 7 Enter a name in the **Display Name** field for the stretched VRF (for example, `stretched-vrf`).

Figure 118:

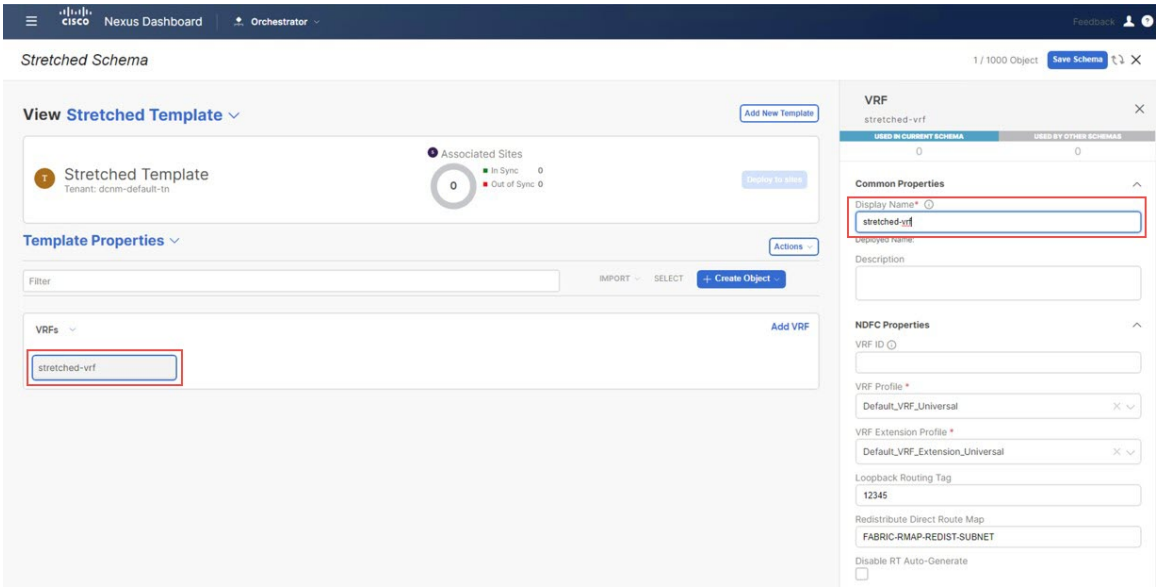

Step 8 Associate all the sites (on-premises and cloud sites) to `Stretched Template` for the stretched VRF use case.

a) In the **Template Properties** area, click **Actions** > **Sites Association**.

Figure 119:

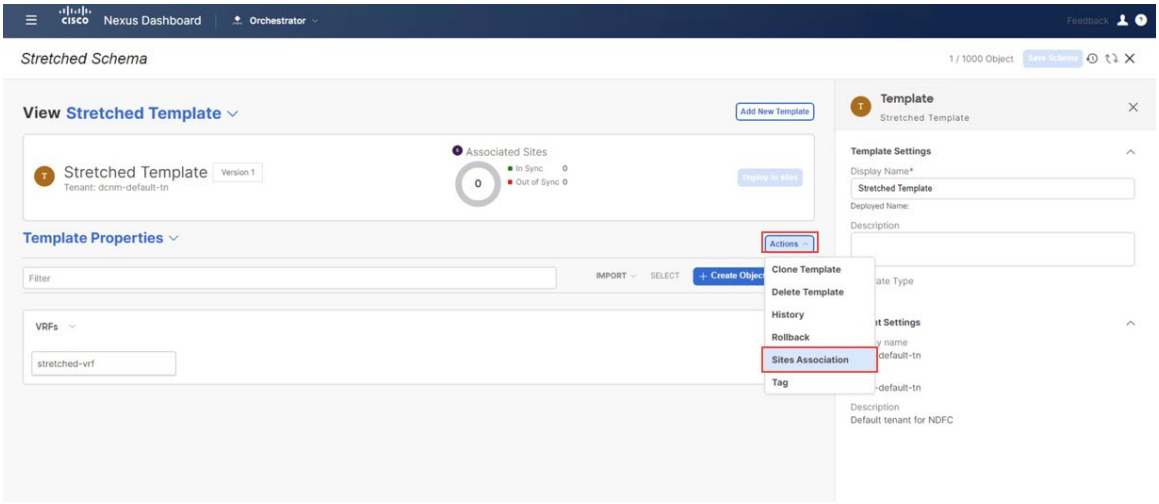

b) Select all the sites, then click **Ok**.

This also allows you to select each site individually to provision site-level configurations for the objects defined in this template (in this specific case, just the stretched VRF).

Figure 120:

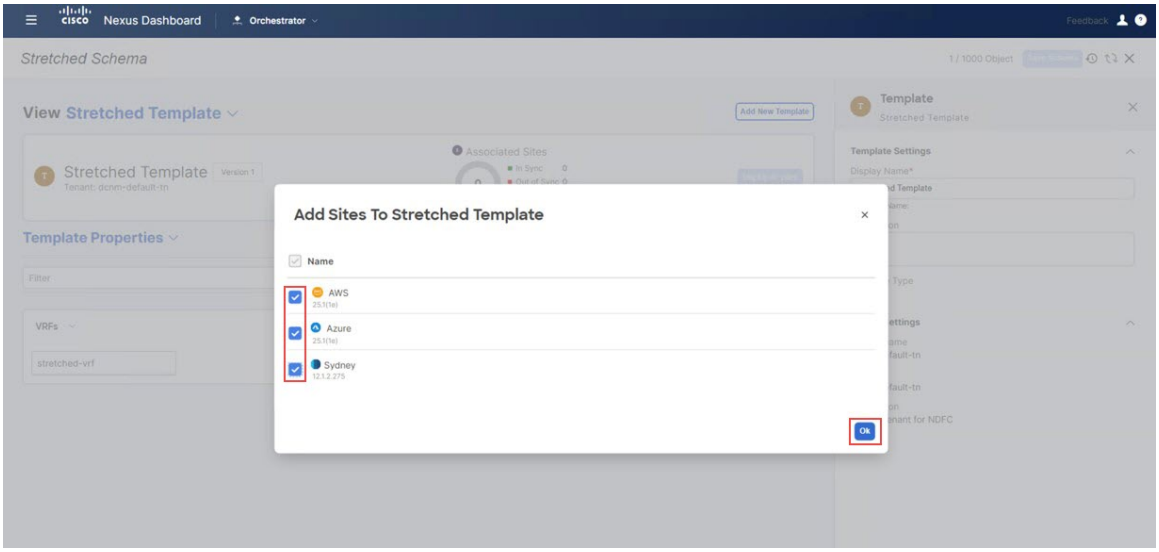

Once the sites are associated with the template, they will appear under **Template Properties**.

Figure 121:

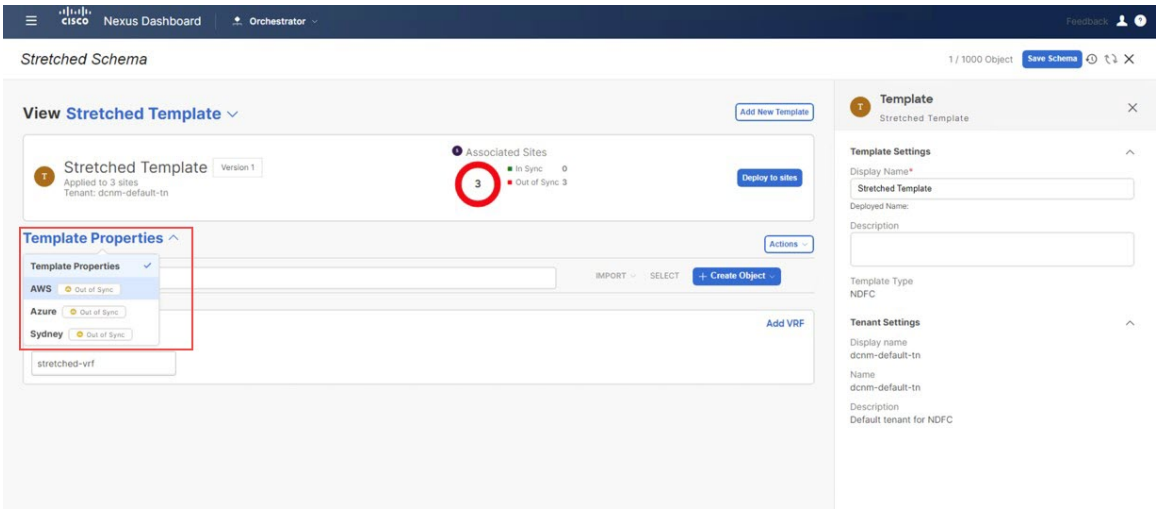

Step 9 Click **Template Properties** and select the first cloud site (the AWS site in this example use case), then associate the VRF to the appropriate regions to create the VPC.

a) Click the VRF, then click **Add Region** to create the VPC in the selected region.

Figure 122:

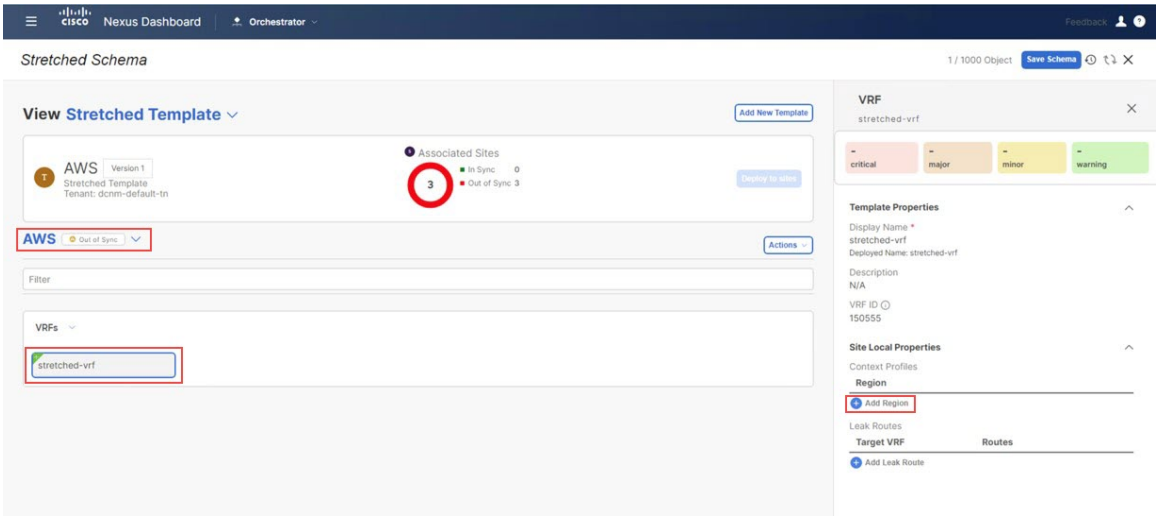

The **Add Cloud Region CIDRs** window appears.

b) In the **Region** field, choose the region where you want to create the VPC.

Figure 123:

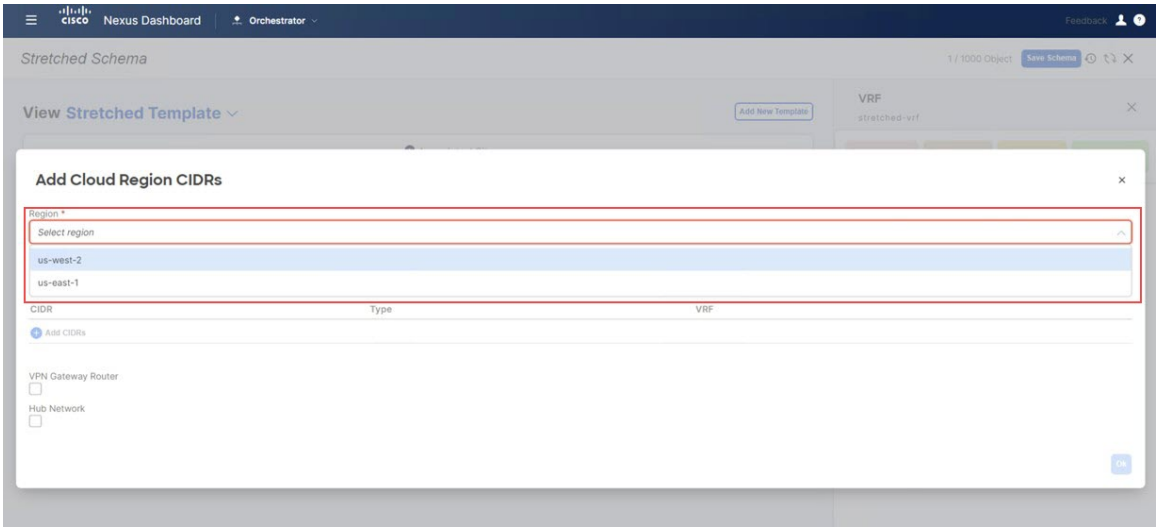

c) In the **CIDR** field, click **Add CIDRs** and define a CIDR block for the VPC.

d) Click **Add Subnet** to create the subnets and map them to the availability zones, then click **Save**.

Figure 124:

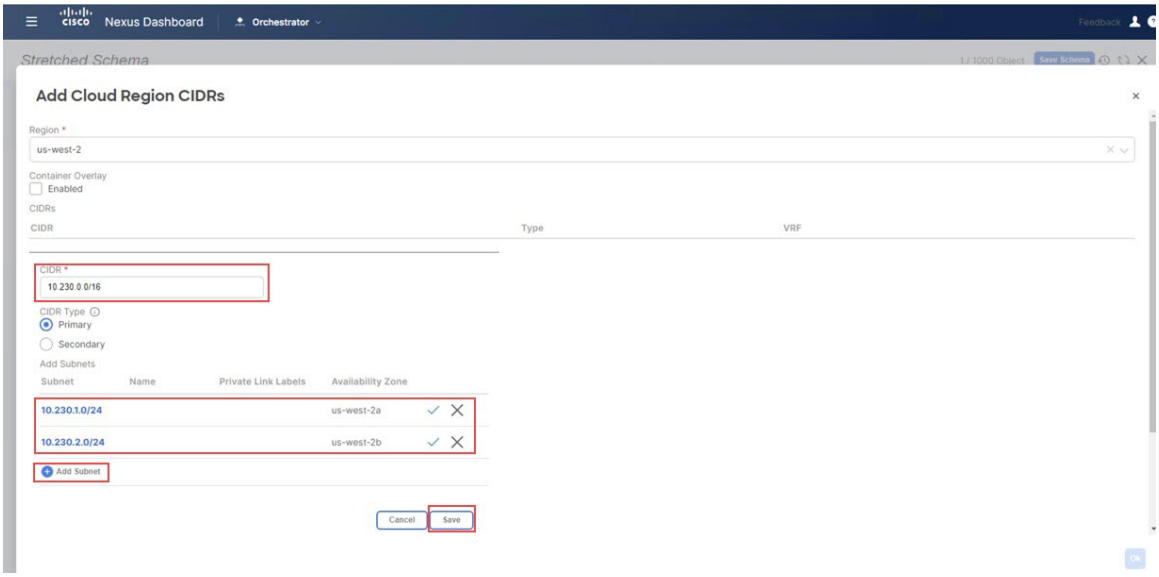

e) Check the box under the **Hub Network** field, then select the hub network that was created on the Cisco Cloud Network Controller for AWS.

This allows the Cisco Cloud Network Controller to attach the subnets onto the transit gateway, which builds the connectivity from those subnets to the transit gateway, where the transit gateway already has the connectivity to the Cisco Catalyst 8000Vs in the cloud.

f) In the **Subnets** field, map the subnets that will be used for the transit gateway.

It is best practice to have a dedicated subnet that will be used for the transit gateway.

Figure 125:

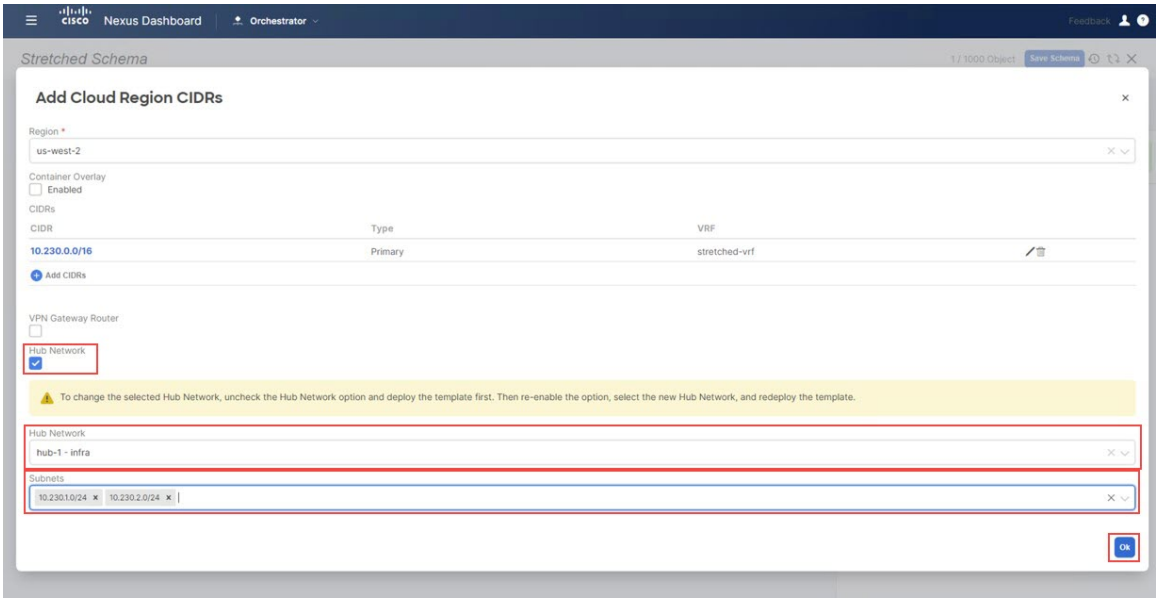

Note Alternatively, a dedicated /25 subnet per availability zone can be used for connectivity to a hub network (TGW). This will allow the entire end-point subnets to be used for end hosts.

g) Click **Ok**.

You are returned to the AWS template window.

When this configuration is deployed, a VPC with CIDR 10.230.0.0/16 will be created in the AWS cloud, stretching between the `us-west-2a` and `us-west-2b` availability zones, with the 10.230.1.0/24 and 10.230.2.0/24 subnets created respectively.

Figure 126:

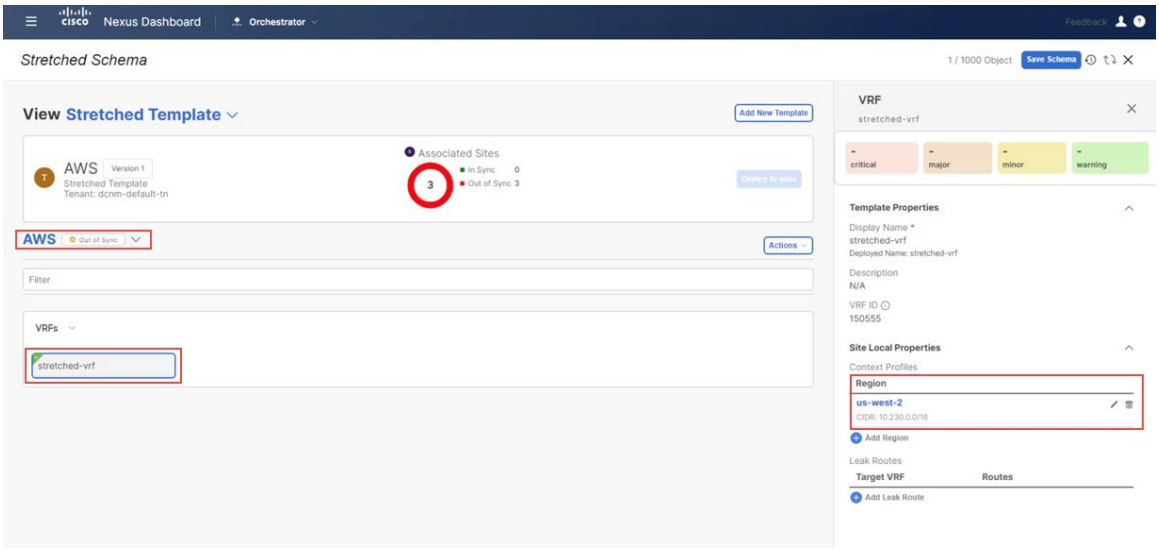

Step 10 Click **Template Properties** and select the second cloud site (the Azure site in this example use case), then associate the VRF to the appropriate region to create the VNet.

a) Click the VRF, then click **Add Region** to create the VNet in the selected region.

Figure 127:

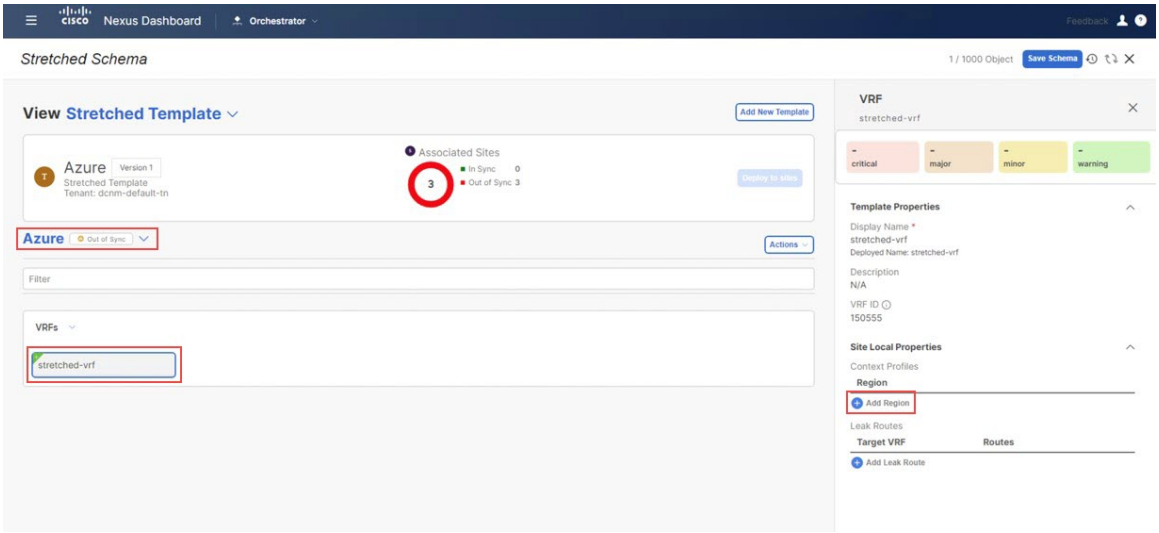

The **Add Cloud Region CIDRs** window appears.

b) In the **Region** field, choose the region where you want to create the VNet.

c) In the **CIDR** field, click **Add CIDRs** and define a CIDR block for the VNet.

d) Click **Add Subnet** to create the subnets, then click **Save**.

Figure 128:

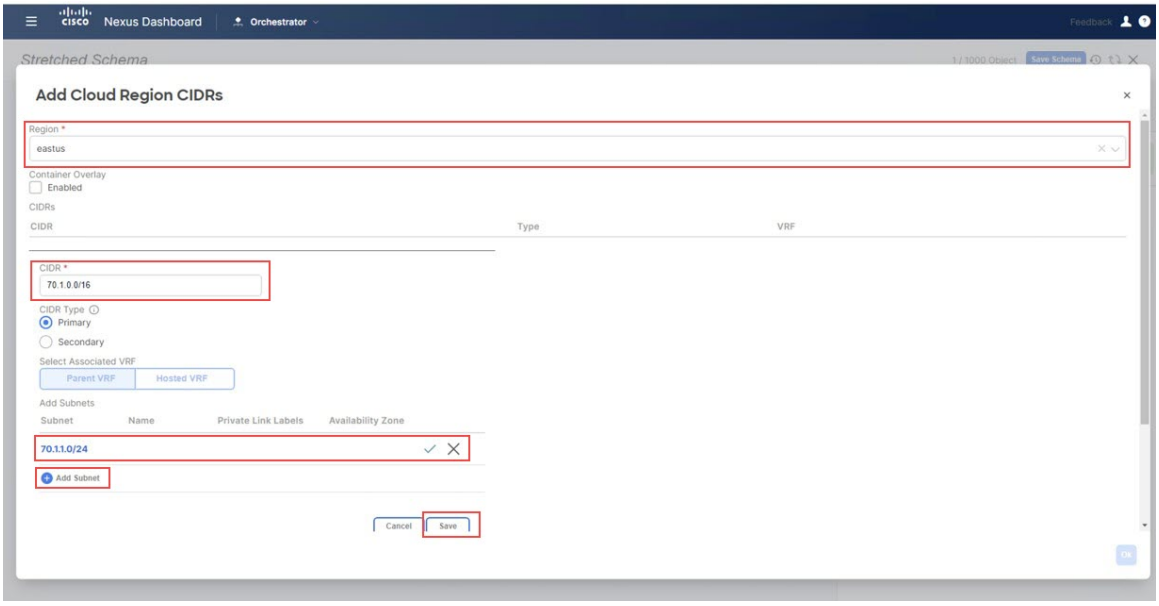

e) Check the box under the **VNet Peering** field, then select the `Default` hub network that was created on the Cisco Cloud Network Controller for Azure.

Figure 129:

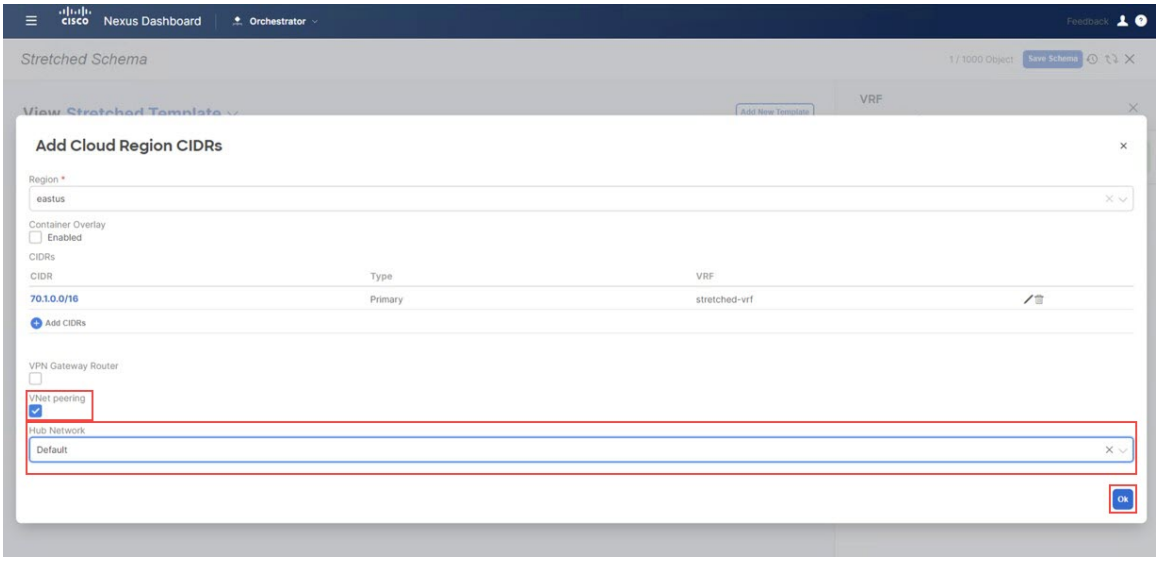

f) Click **Ok**.

When this configuration is deployed, the VNet that you configured (in this example, 70.1.0.0/16) will be created on the appropriate region in Azure (in this example, the eastus Azure region) and VNet peering is configured to the infra VNet in the infra tenant in Azure.

Figure 130:

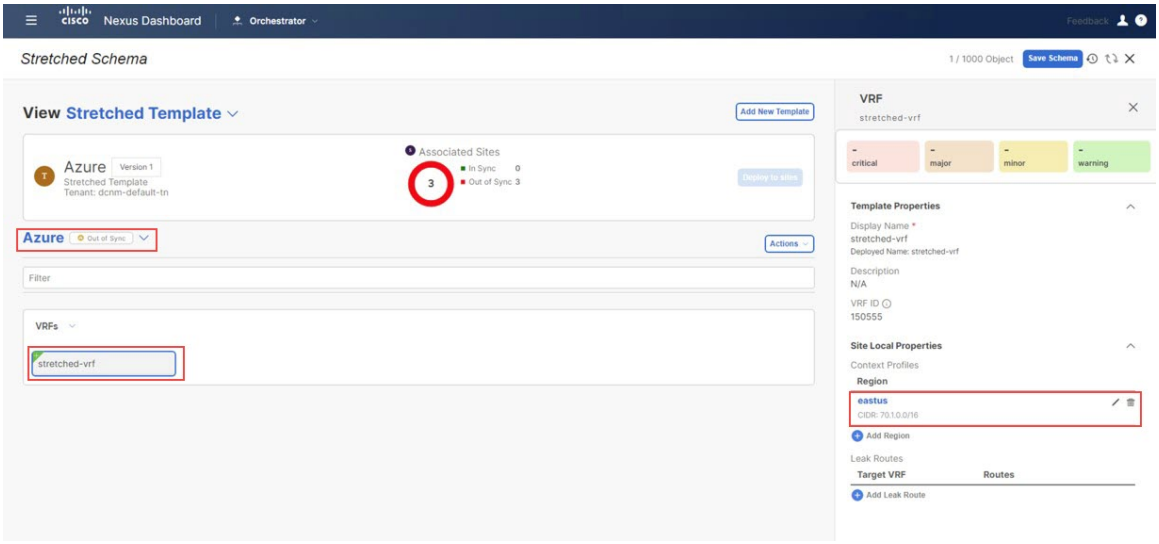

Step 11 Click **Template Properties** and select the on-premises site (the Sydney site in this example use case), then select the `stretched-vrf` VRF.

Step 12 In the right pane, click **Add Static Leaf**.

Figure 131:

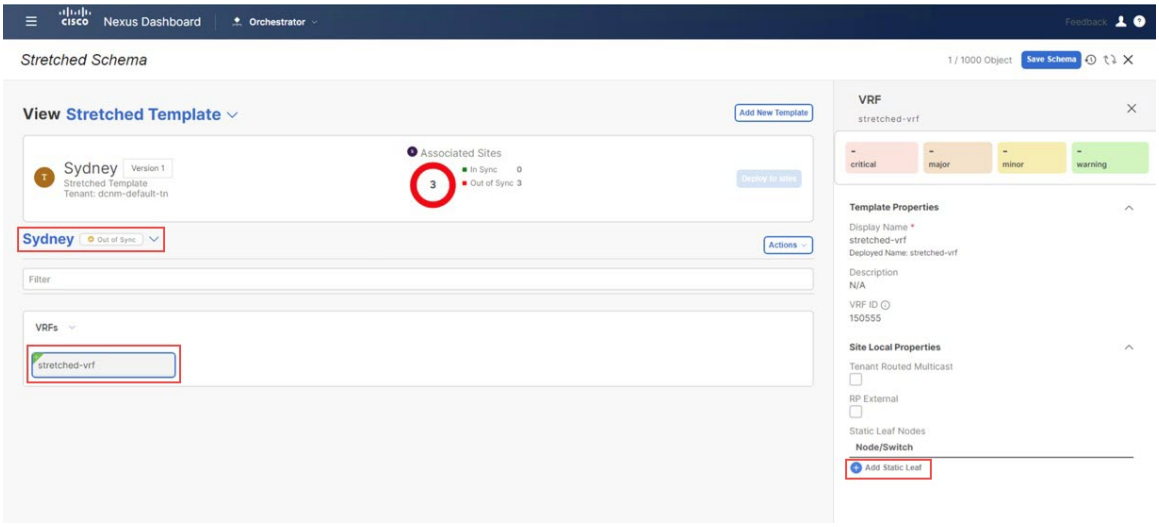

The **Add Static Leaf** window appears.

Step 13 In the **Leaf** field, select the leaf/border/border gateway device where this VRF is to be deployed and click **Ok**.

Figure 132:

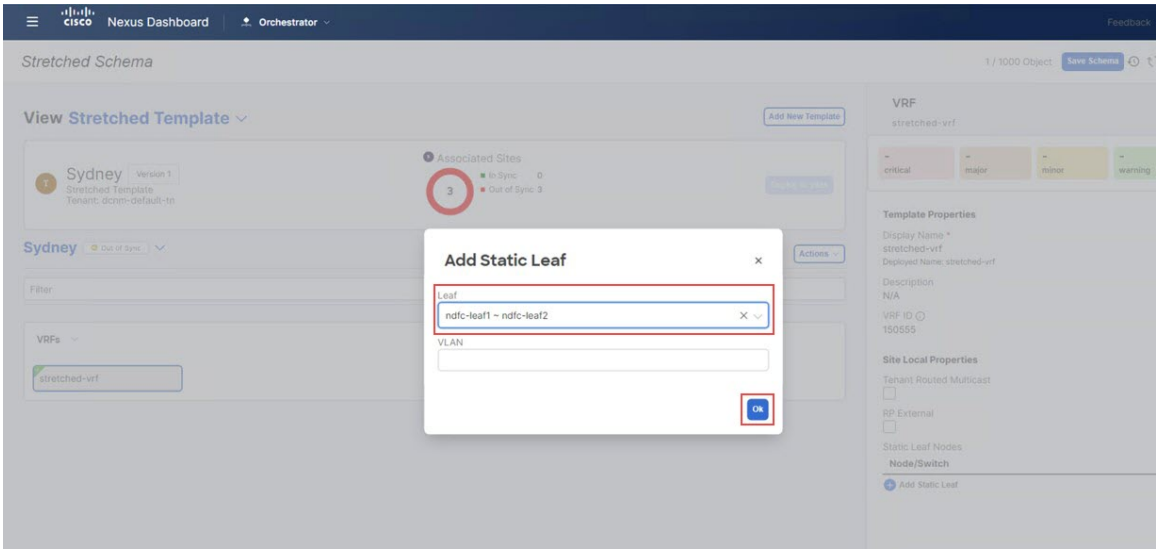

You are returned to the `Stretched Template` page.

Step 14 Click **Add Static Leaf** again to add additional leaf/border/border gateway devices where this VRF is to be deployed.

In this example, you need to deploy the VRF on the leaf nodes (where the endpoints part of the network mapped to the VRF will be connected) and on the BGW spine node to be able to extend the Layer 3 connectivity for the VRF towards the cloud sites.

Figure 133:

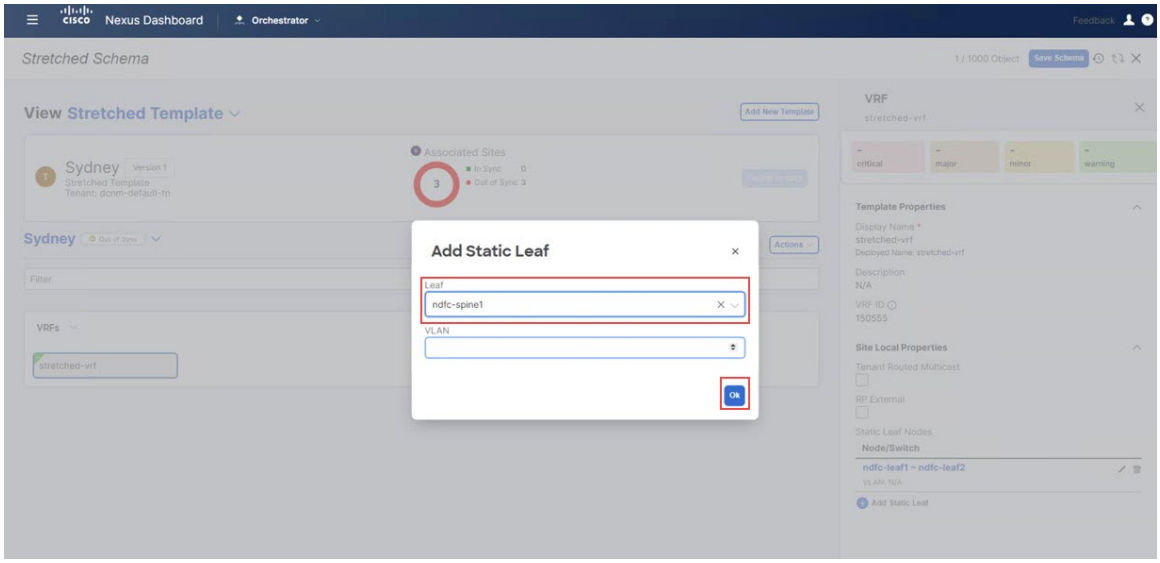

When you have added all of the leaf/border/border gateway devices where this VRF is to be deployed, they will appear in the **Stretched Template** page.

Figure 134:

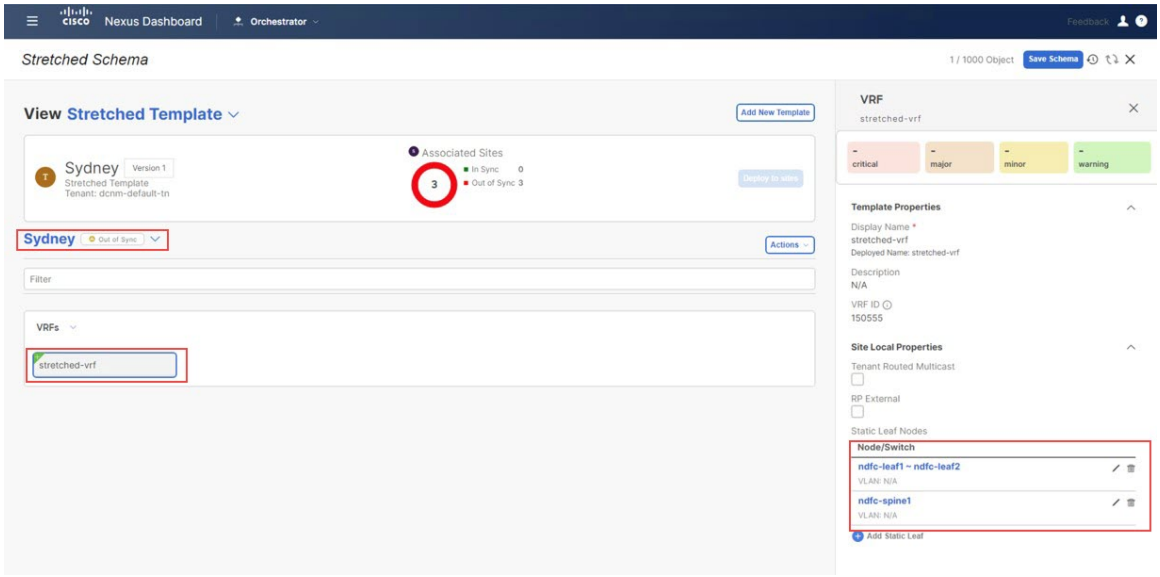

Step 15 Click the arrow next to the Sydney site, and from the drop-down menu, select **Template Properties**.

Step 16 Click **Deploy to sites**.

Figure 135:

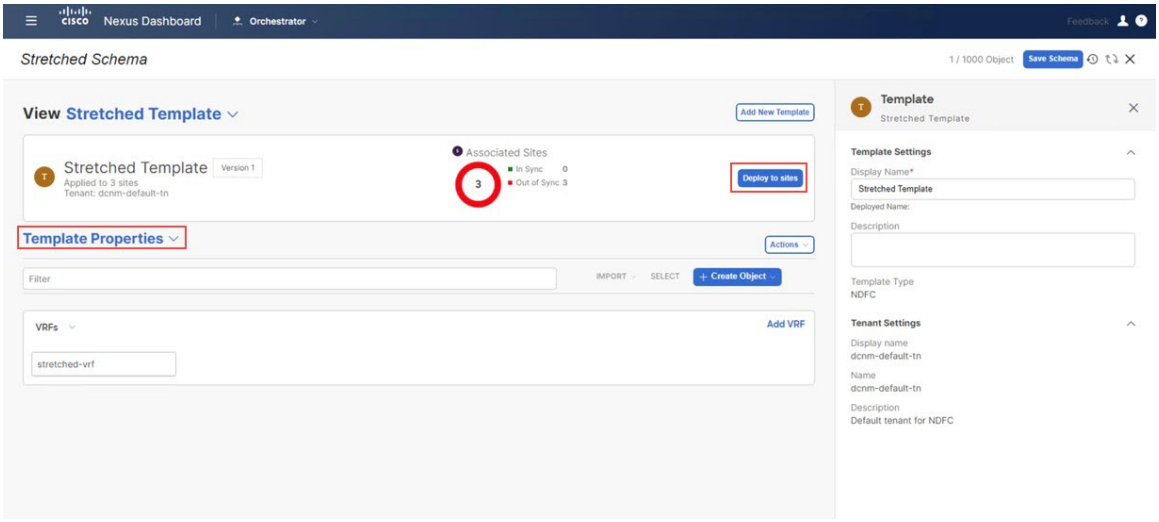

The **Deploy to Sites** window appears, showing the three sites where the stretched template will be deployed.

Figure 136:

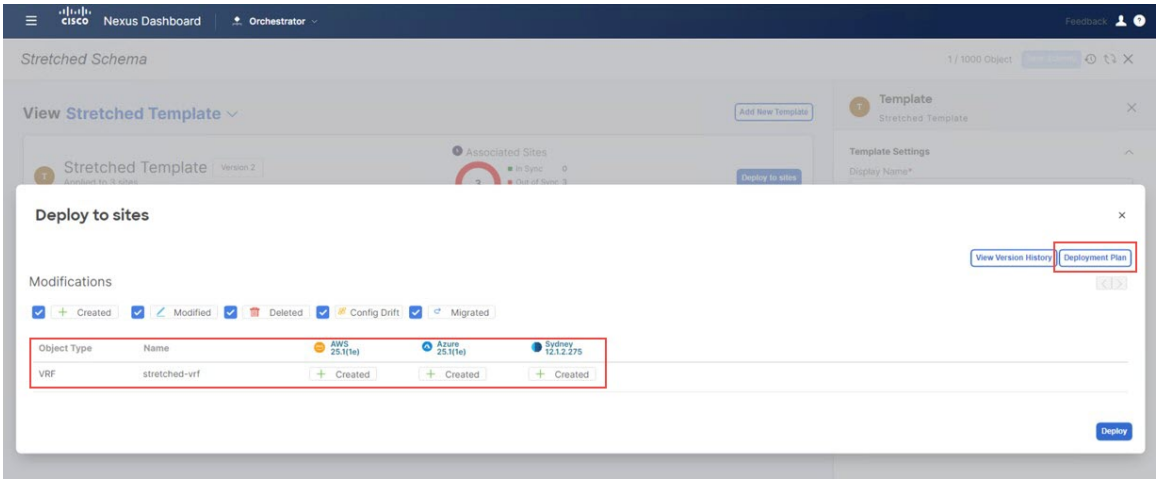

Step 17 Click **Deployment Plan** for additional verification, then click on each site to see the deployment plan for that specific site.

Figure 137:

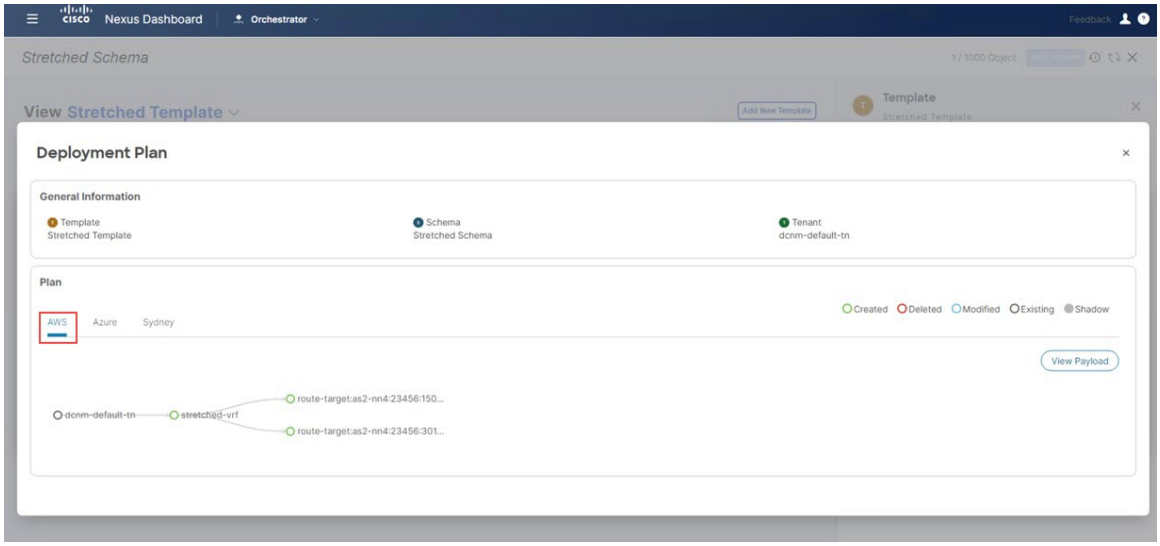

Figure 138:

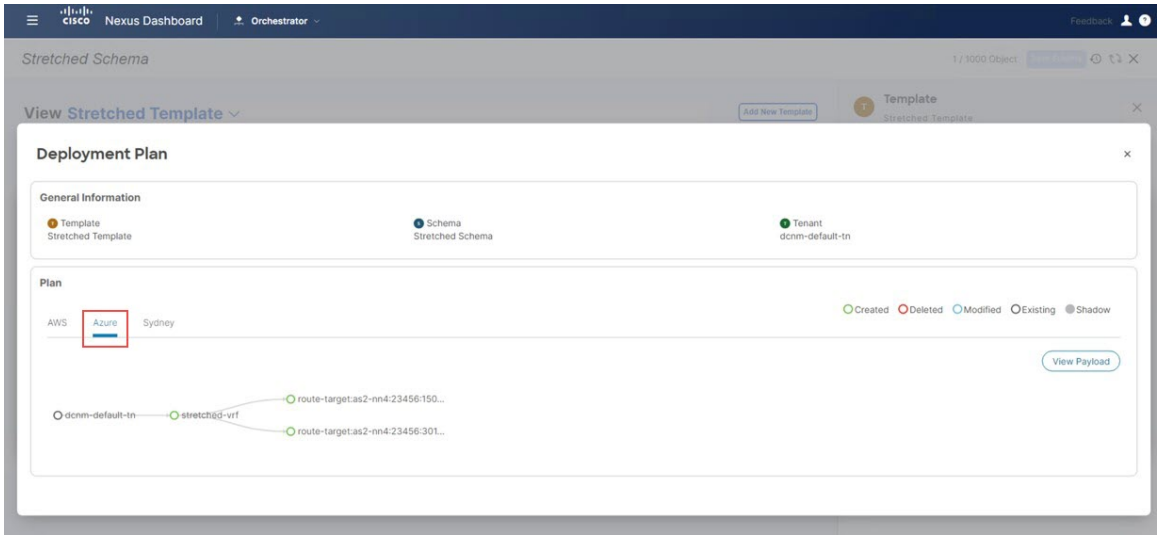

Figure 139:

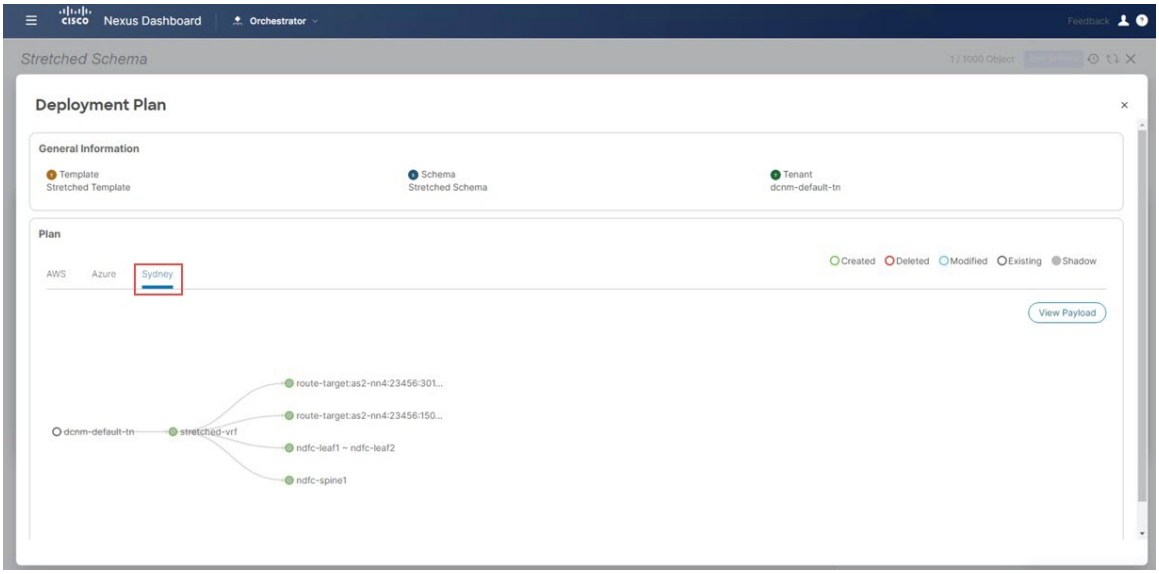

Step 18 Click **Deploy** to have NDO push the configurations to the site specific controllers (NDFC and Cloud Network Controller).

Figure 140:

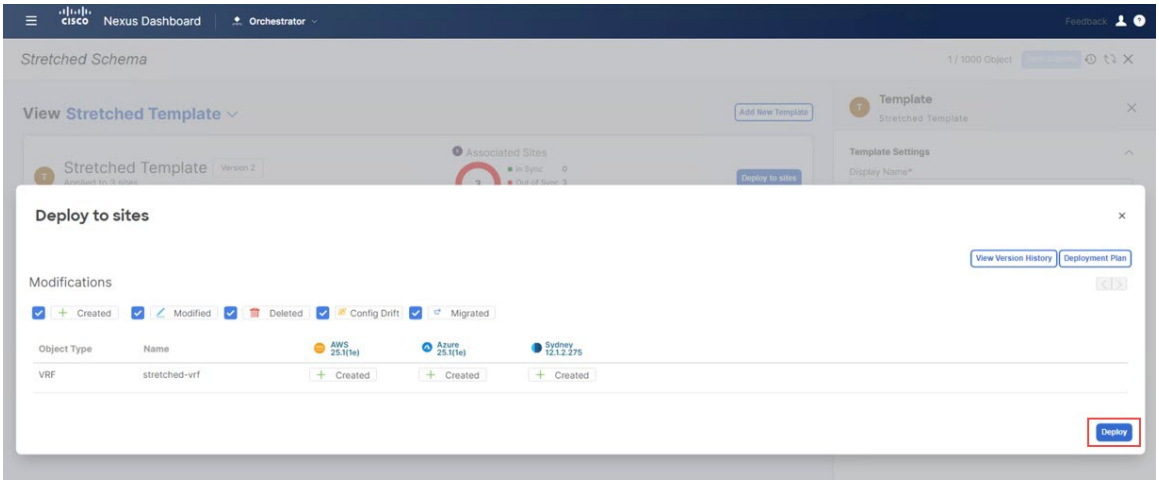

Step 19 Verify that the configurations were deployed successfully.

- To view the VRF deployment on NDFC, go to the **Topology** view, select the on-premises fabric **Sydney** > **VRFs**, then select `stretched-vrf`.

Figure 141:

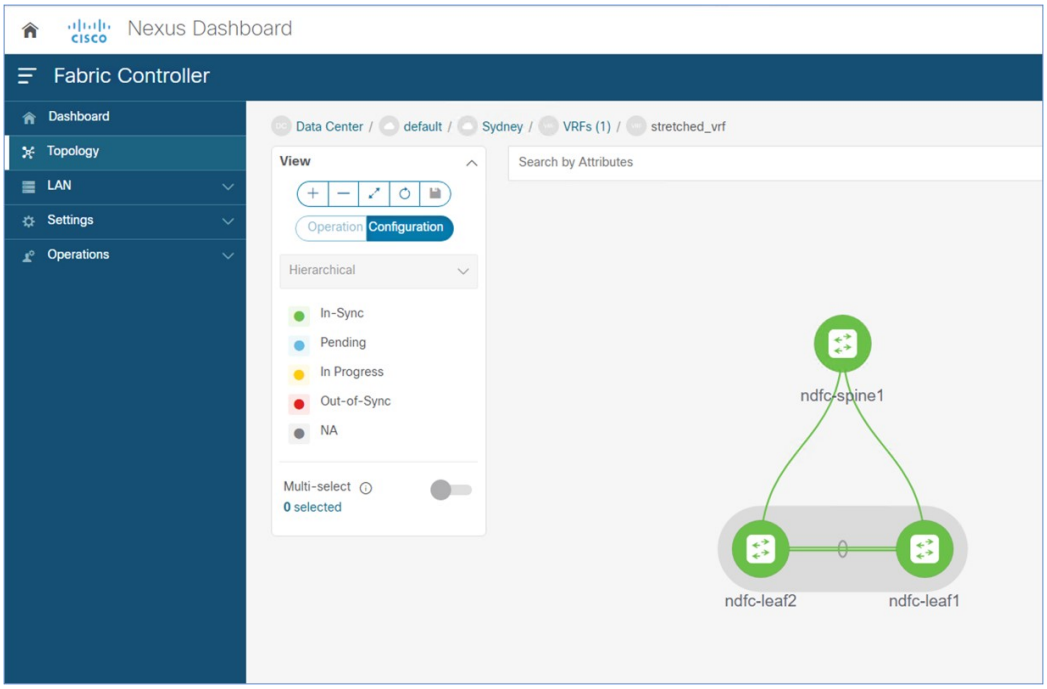

- Connect to the Cloud Network Controller deployed on AWS to verify that the configurations for the first cloud site (AWS) were deployed successfully.

  Go to **Application Management** > **VRFs**, locate `stretched-vrf` and click under the column **VPCs**, then go to the **Overview** page and click under **Subnets**.

- Connect to the Cloud Network Controller deployed on Azure to verify that the configurations for the second cloud site (Azure) were deployed successfully.

  Go to **Application Management** > **VRFs**, locate `stretched-vrf` and click under the column **Virtual Networks**, then go to the **Overview** page and click under **Subnets**.

Step 20 Create another template under `Demo Schema` for deploying networks on the on-premises site.

a) Under the `Demo Schema` template, click **Add New Template**.
b) Choose the NDFC template.
c) Enter a name in the **Display Name** field to create an NDFC-type template (for example, `On-Prem Template`) and select the `dcnm-default-tn` tenant in the **Select a Tenant** field to map the template to that tenant.

Figure 142:

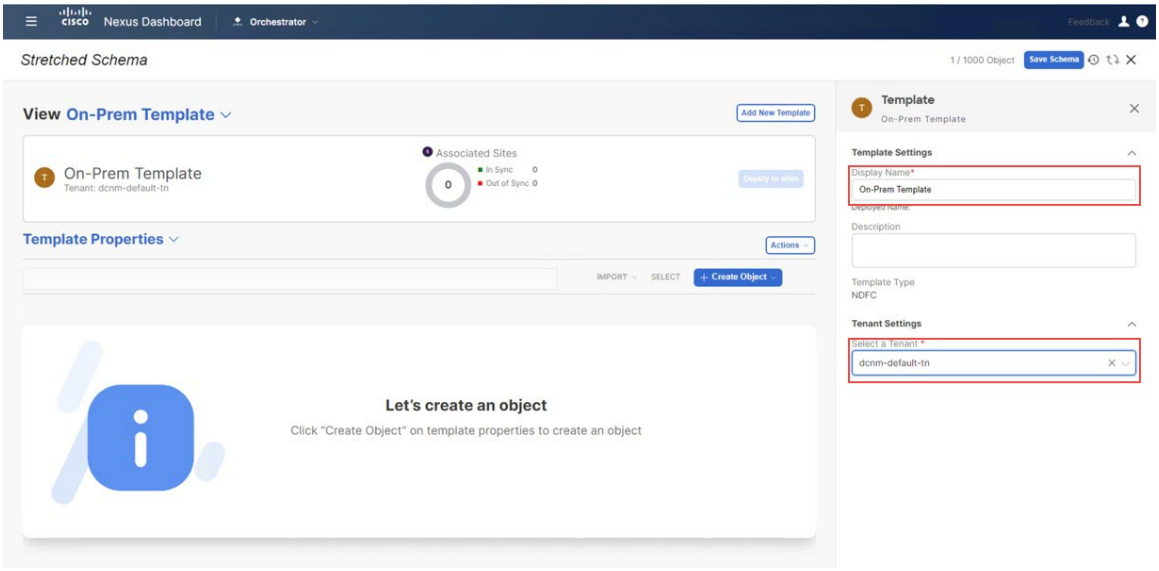

Step 21 Create the `net20` network under the VRF on `On-Prem Template`.

Note If you have a network already created that you want to use instead of creating a new network, under **Template Properties**, click **Import**, then import the already-created network.

a) Under **Template Properties**, click **Create Object** and choose **Network** to create a network.

Figure 143:

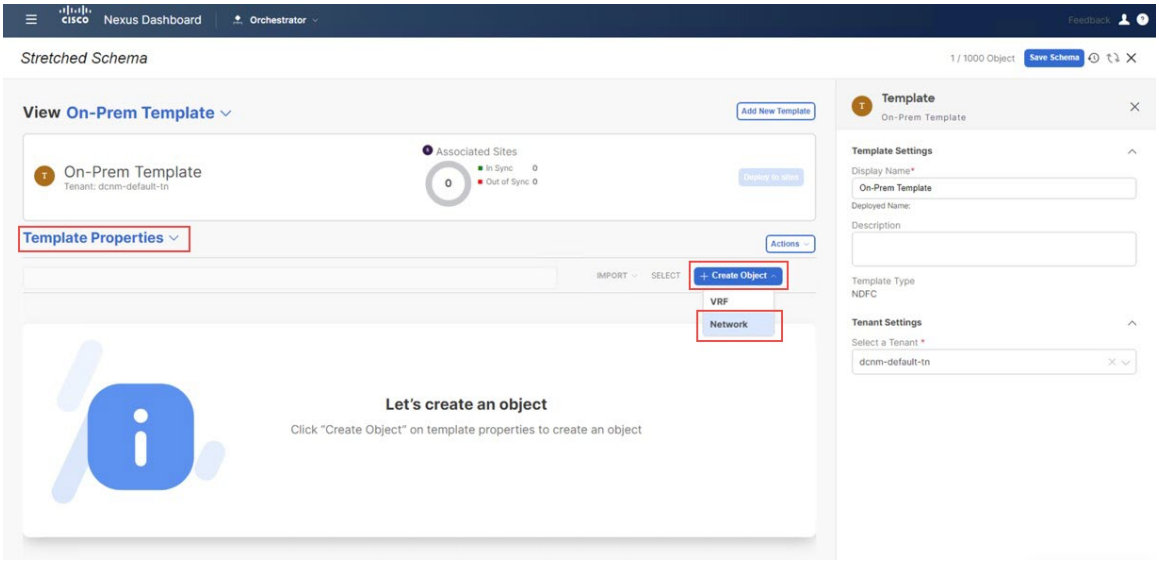

b) Enter a name in the **Display Name** field for the network (for example, `net20`).

c) In the **Virtual Routing & Forwarding** field, choose the `stretched-vrf` VRF to map `net20` to that VRF.

Figure 144:

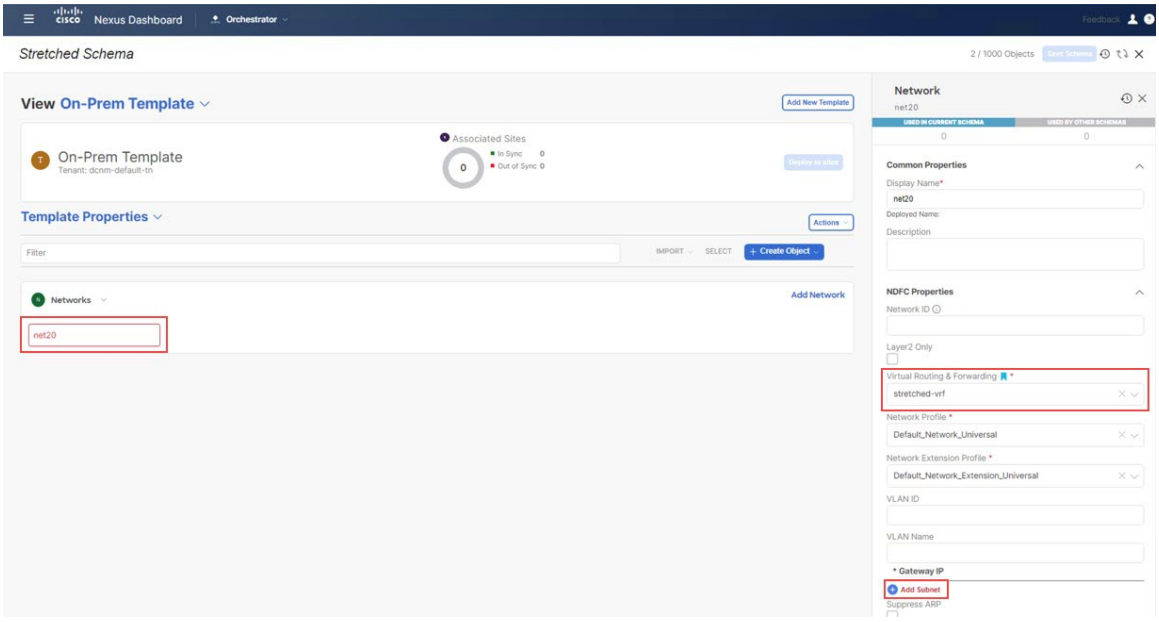

d) In the **Gateway IP** field, click **Add Subnet**.

The Add Subnet window appears.

e) Click **Add Gateway IP** and provide the gateway IP address, then click the checkmark to accept the value and click **Add**.

Figure 145:

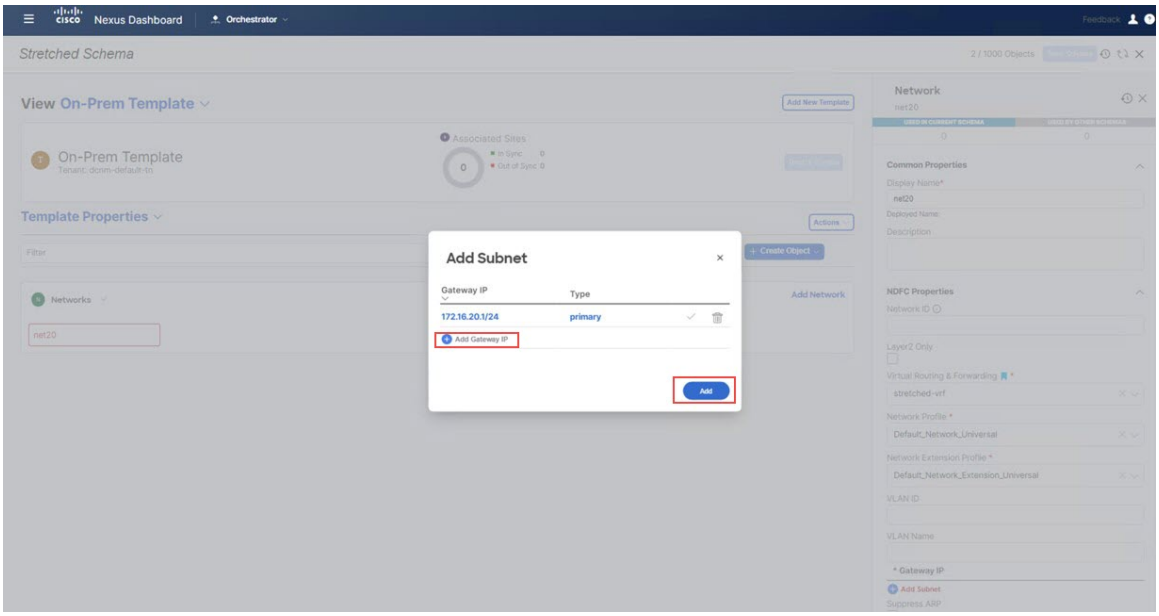

The gateway IP address is now displayed in the **Gateway IP** field.

Figure 146:

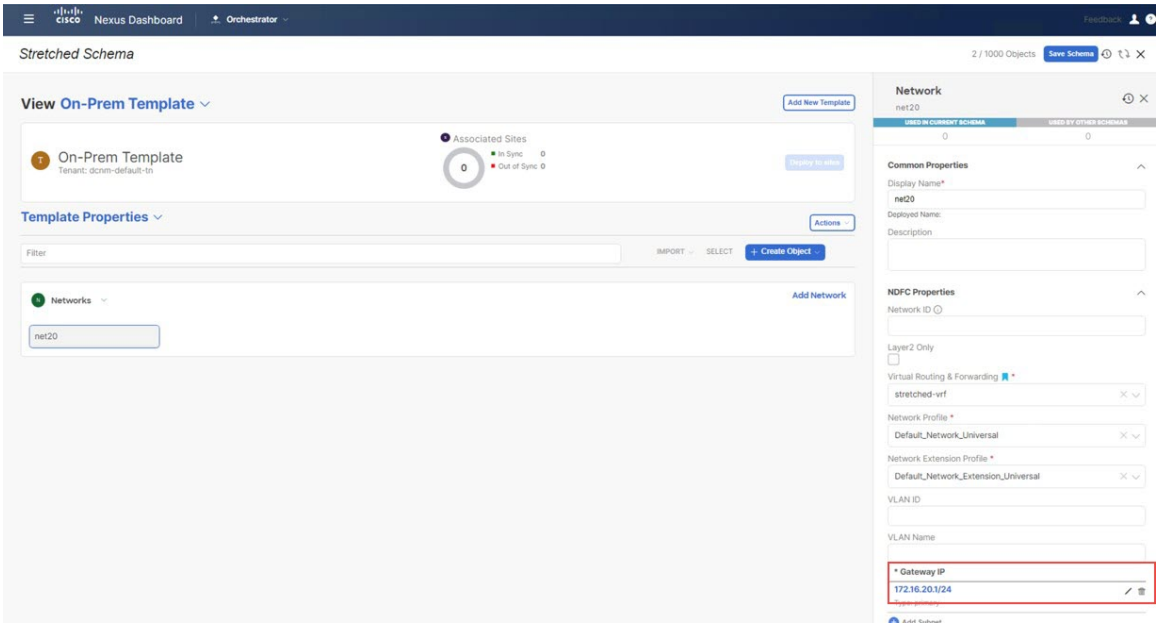

f) Define other optional parameters for this network, if necessary.

Step 22 In the **Template Properties** area, click **Actions** > **Sites Association**.

Figure 147:

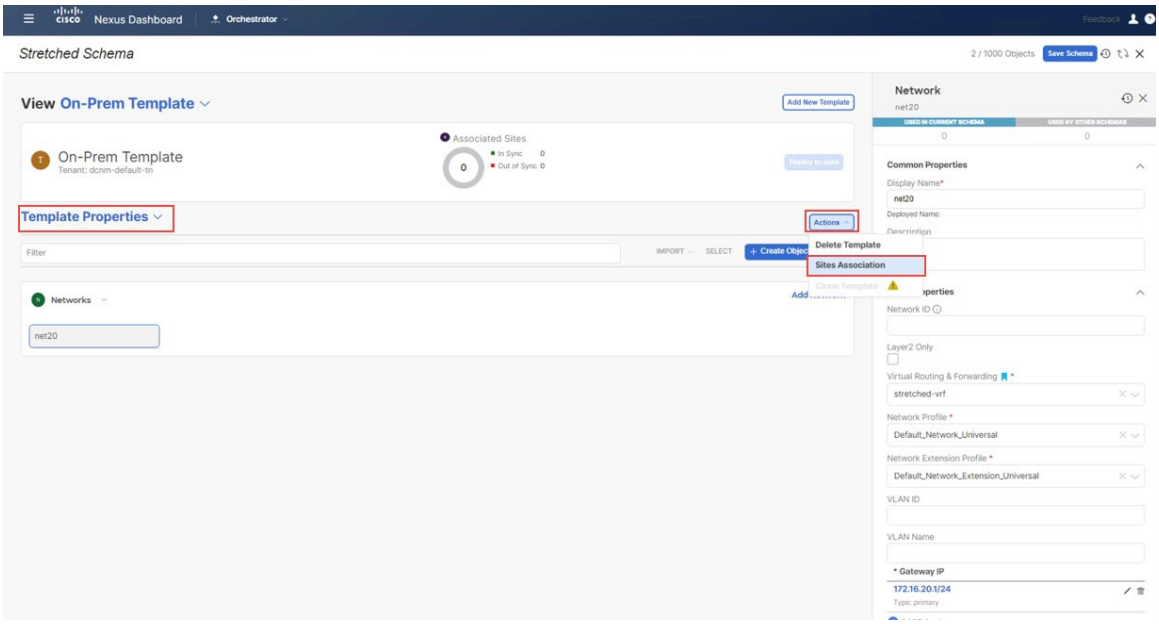

Step 23 Associate this template only to the on-premises site (the Sydney site in this example use case), then click **Ok**.

Figure 148:

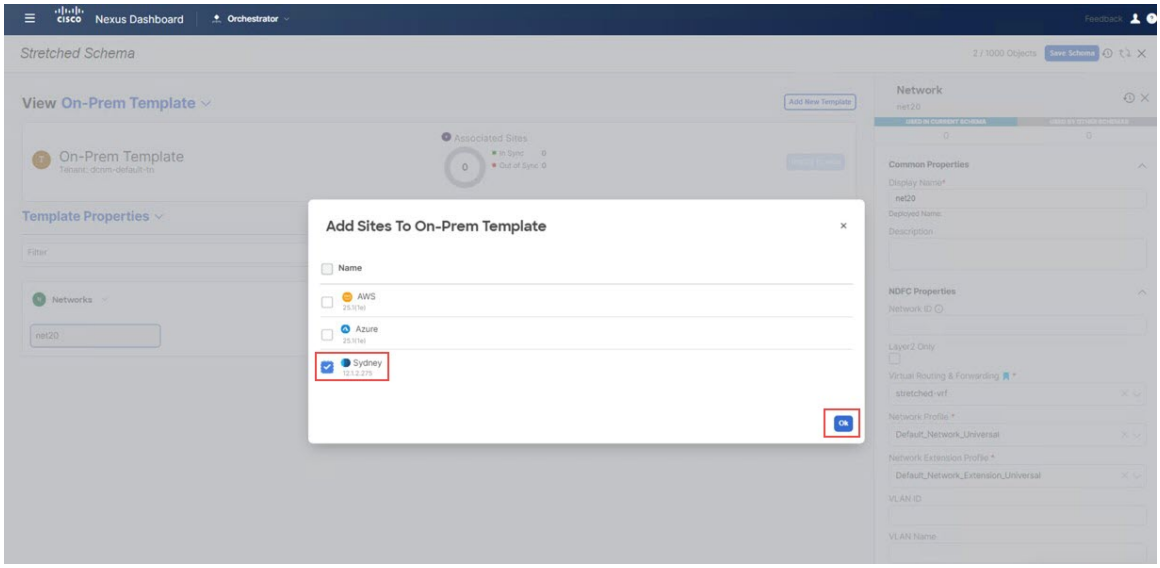

You are returned to the `On-Prem Template` window.

Step 24 From the **Template Properties** drop-down, select the on-premises site (the `Sydney` site in this example use case), click the `net20` network, then click **Add Static Port** to add the ports where you want to deploy this network.
The **Add Static Port** window appears.

Figure 149:

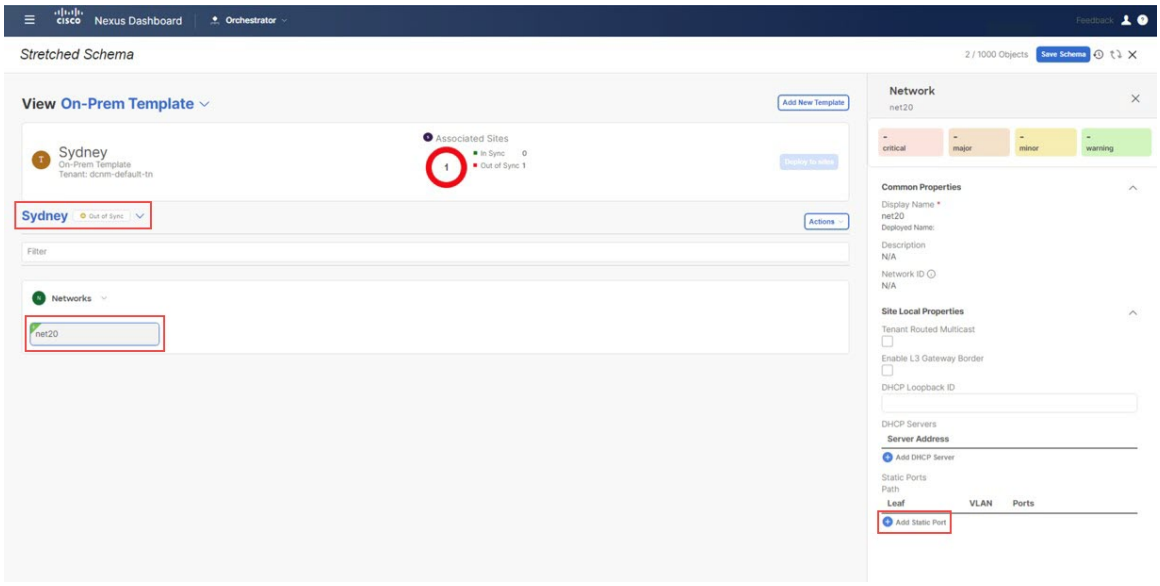

Step 25 In the **Add Static Port** window, click **Add Path**.
The **Add Static Port** window appears.

Step 26 In the **Leaf** field, select the device where you want to deploy this network.

Step 27 (Optional) Enter the necessary information in the **VLAN** field.

Step 28 In the **Ports** field, select the ports where you want to deploy this network.

Step 29 Click **Save**.

Figure 150:

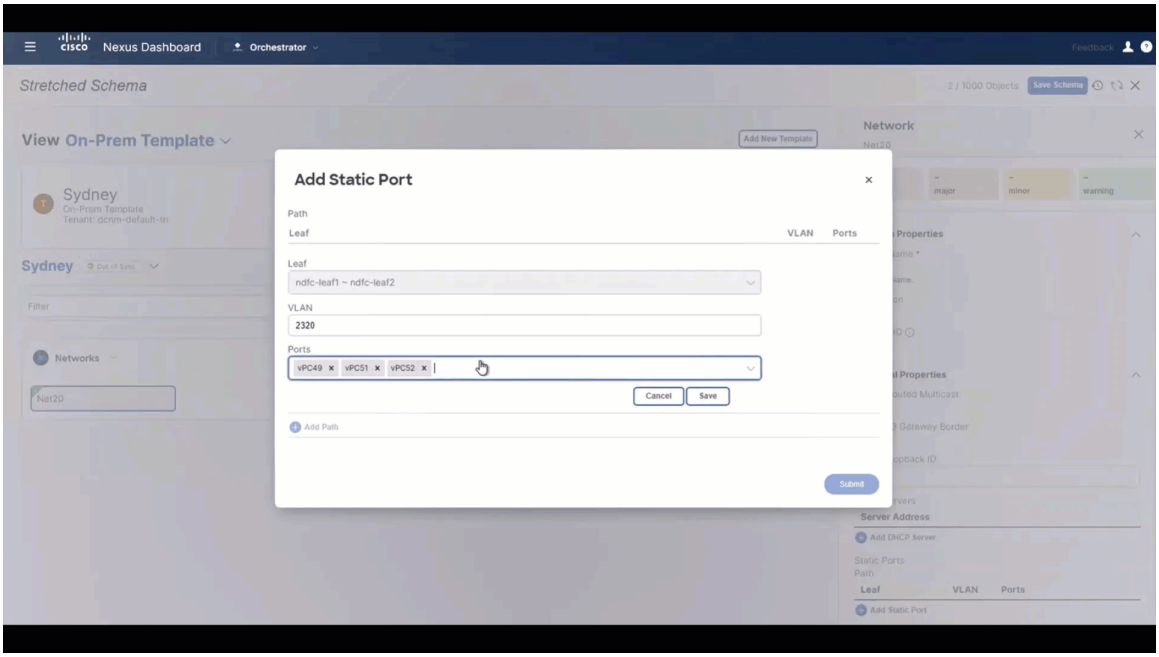

You are returned to the **Add Static Port** window.

Step 30 In the **Add Static Port** window, click **Submit**.

Figure 151:

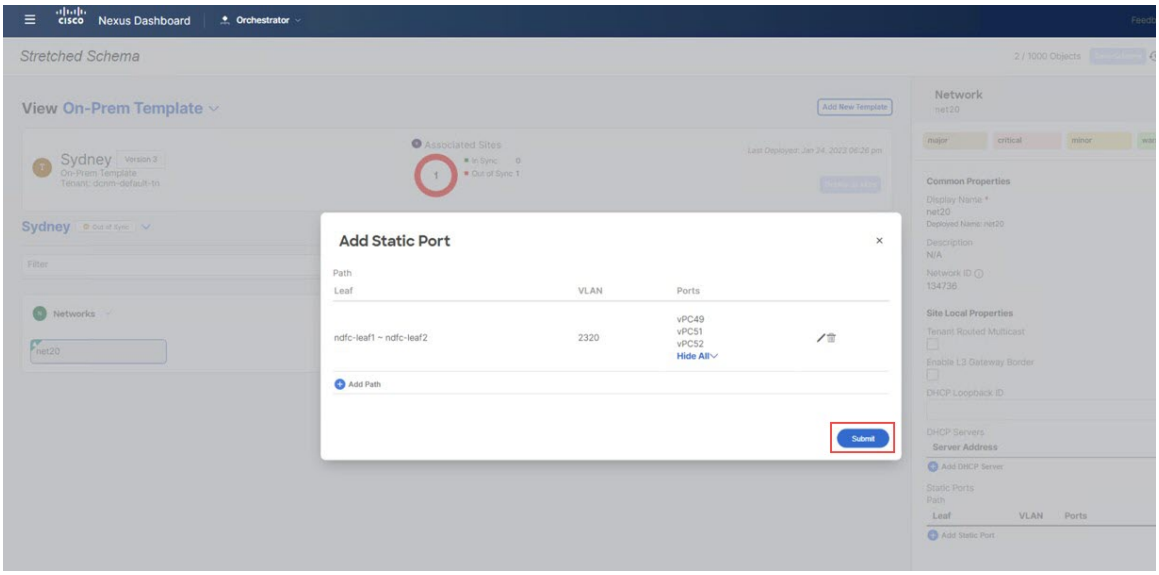

You are returned to the **On-Prem Template** window.

Step 31 Click the arrow next to the on-premises site (the `Sydney` site in this example use case), and from the drop-down menu, select **Template Properties**.

Step 32 Click **Deploy to Sites**.

Figure 152:

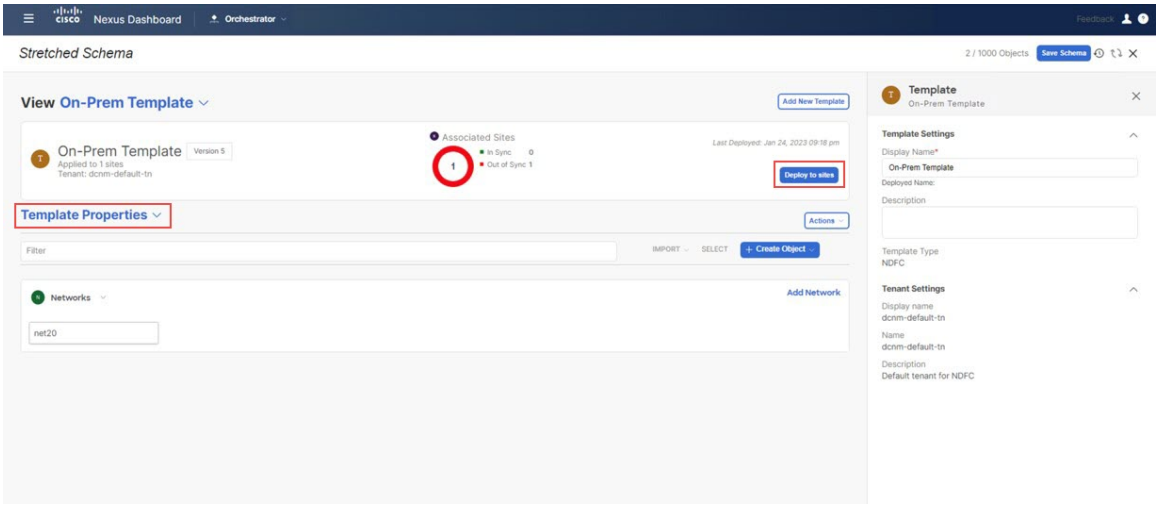

The **Deploy to Sites** window appears, showing the site where the template will be deployed.

Figure 153:

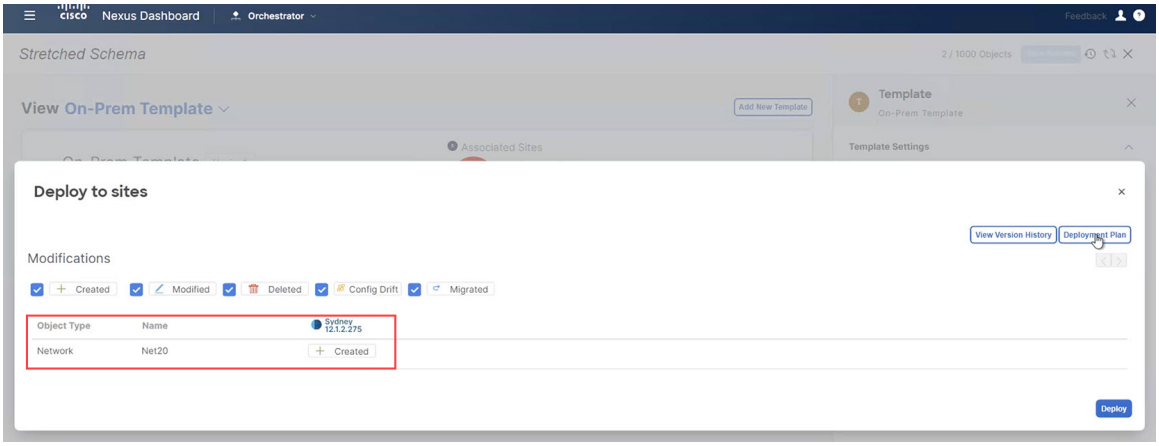

Step 33 Click **Deployment Plan** for additional verification, then click on the on-premises site to see the deployment plan for that specific site.

Figure 154:

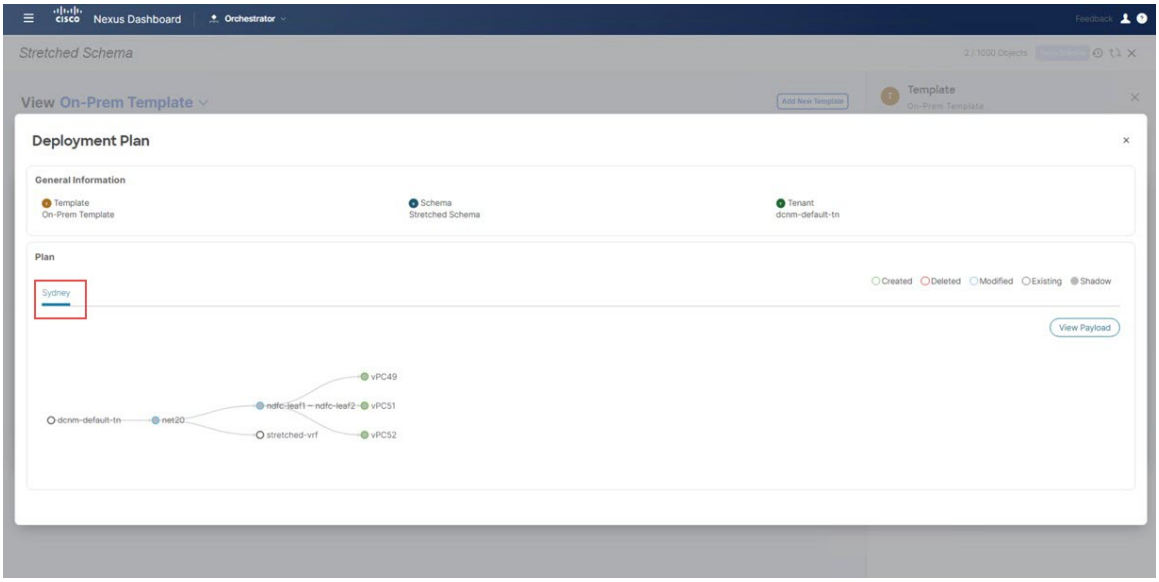

Step 34 Click **Deploy** to have NDO push the configurations to NDFC.

Figure 155:

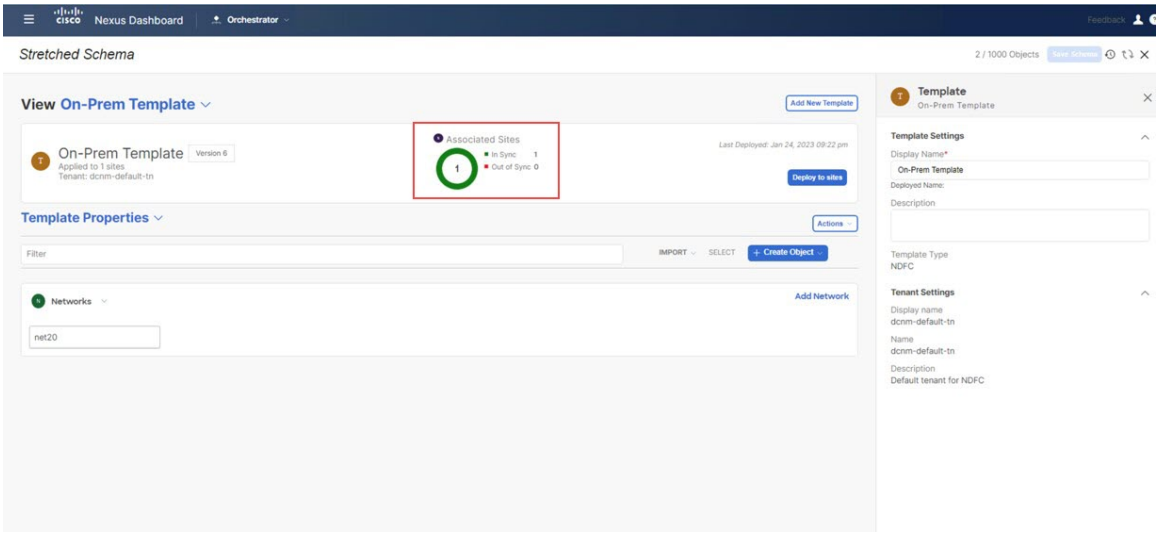

Step 35 Verify that the configurations were deployed successfully.

Note that for each of these verification steps, the exact command that would be used specifically for the configurations in this use case are shown. Replace the appropriate variables in each command based on your configuration.

a) In NDO, verify that the configurations were deployed successfully.

- Verify that the `Stretched Template` was deployed successfully.

Figure 156:

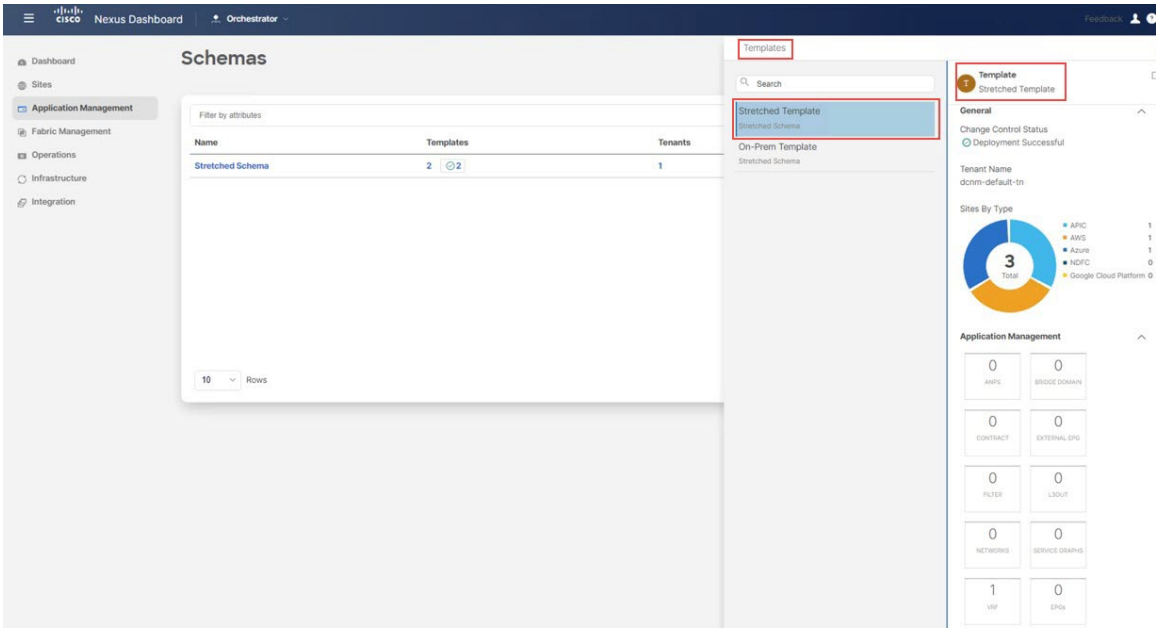

- Verify that the `On-Prem Template` was deployed successfully.

Figure 157:

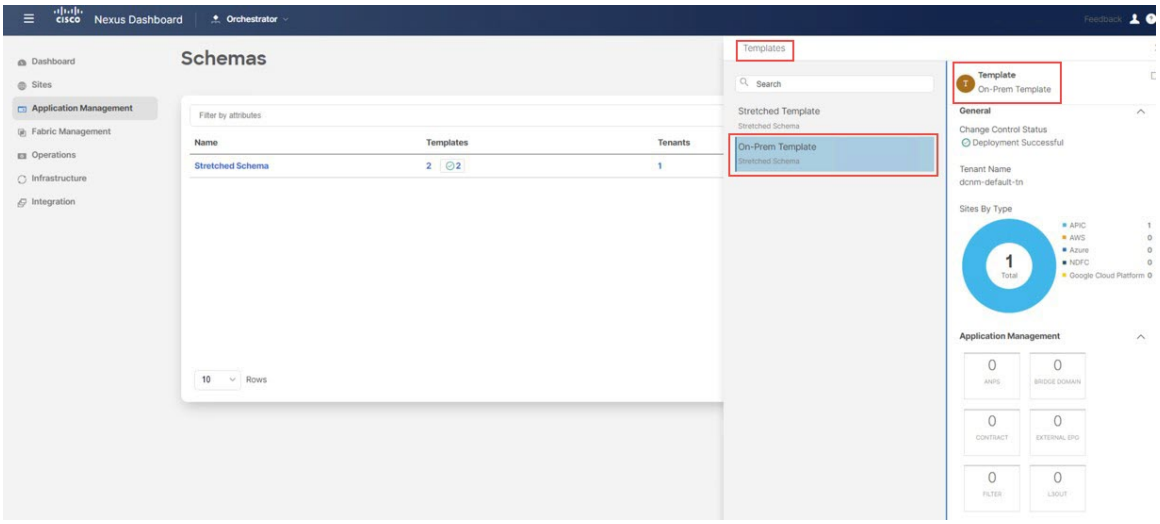

- Verify that the `dcn-default-tn` tenant was deployed successfully.

Figure 158:

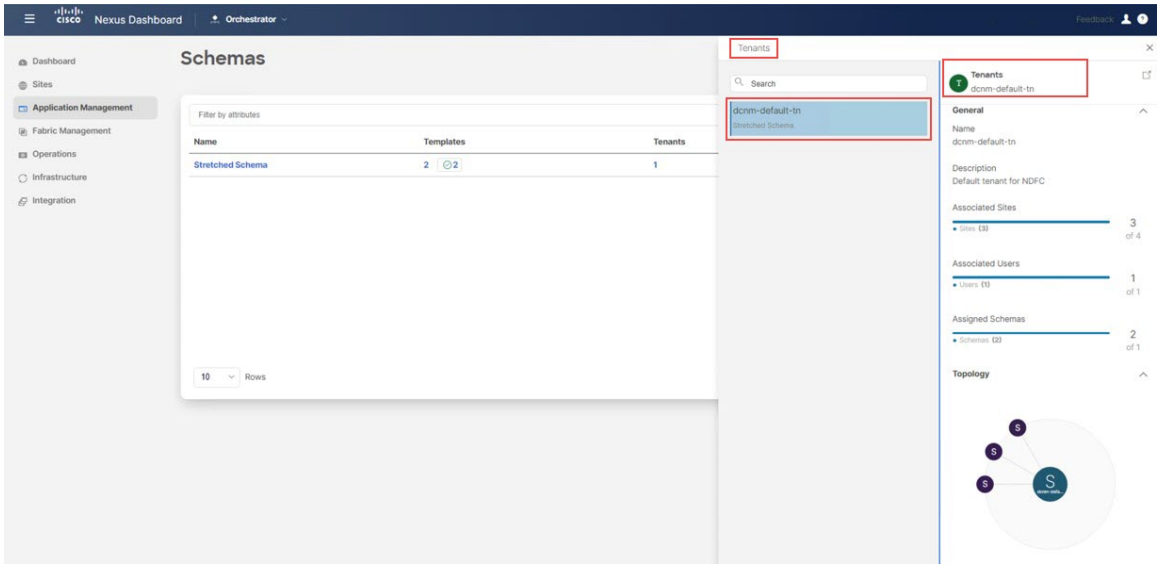

b) In NDFC, verify that the following were done successfully:

- Verify that one vrf and one network has been created.

Figure 159:

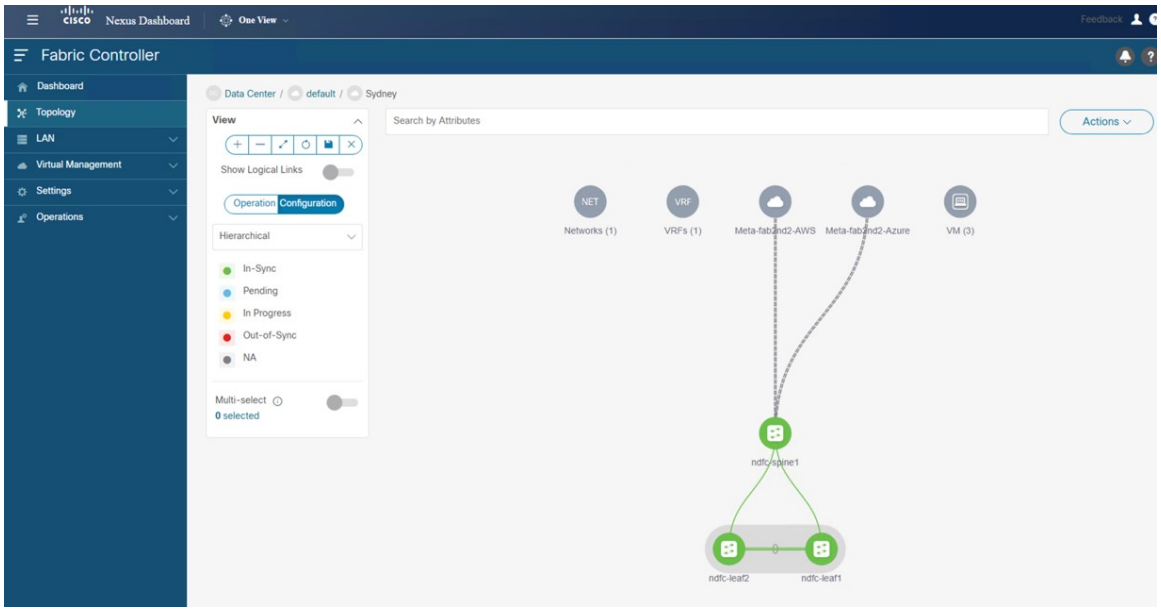

- Verify that the VRF was deployed successfully.

Figure 160:

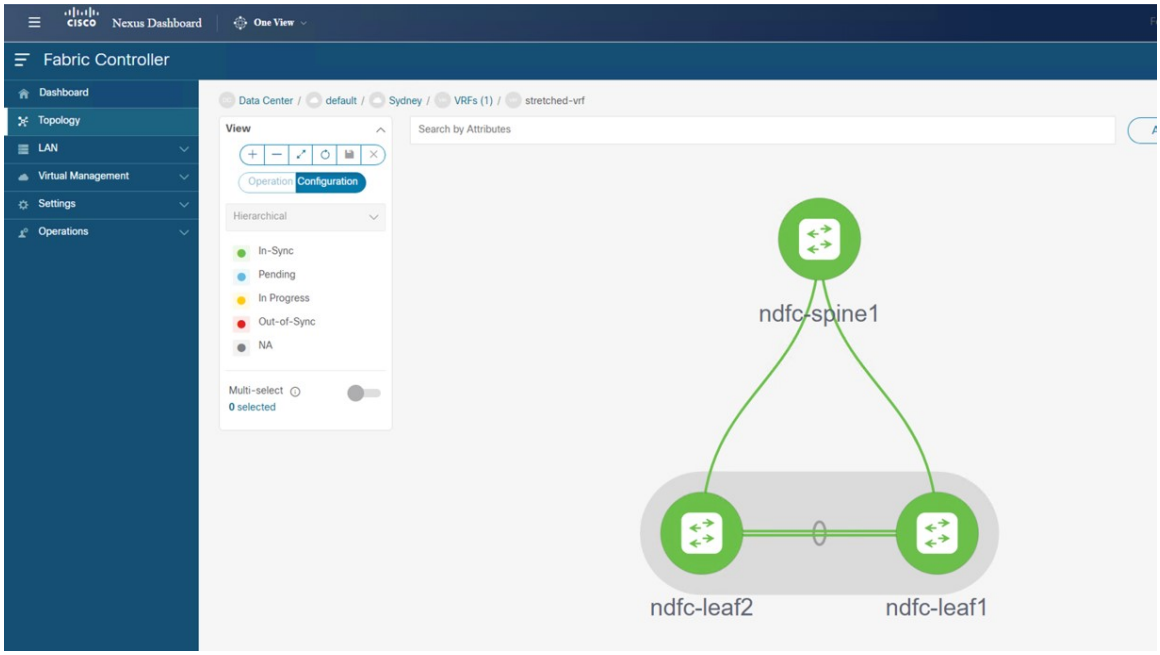

- Verify that the network was deployed successfully.

Figure 161:

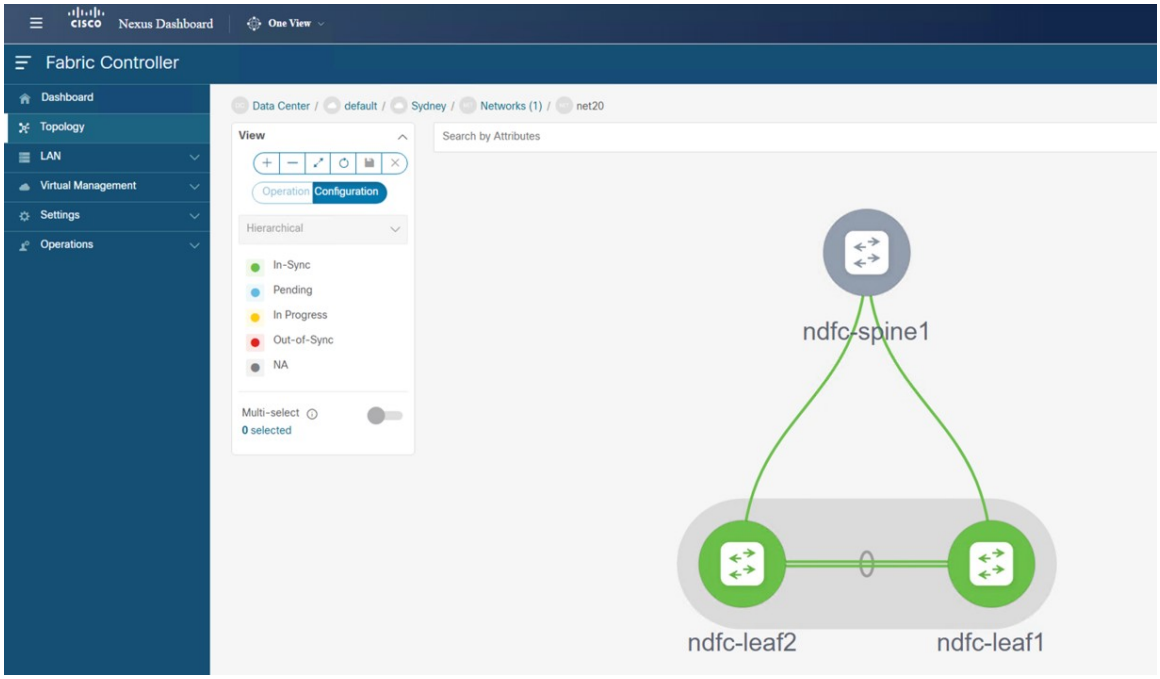

c) Enter **sh ip route vrf stretched-vrf** on the on-premises Border Gateway Spine device:

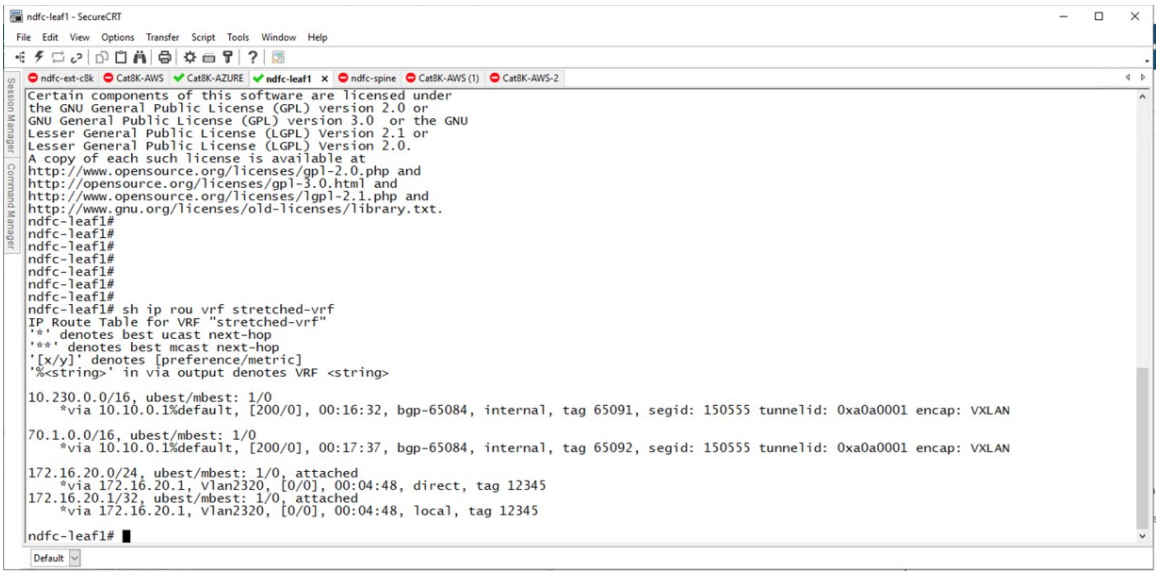

For this use case, using the routing table, you can verify that the NDFC leaf switch can reach out to the following subnets:

- **AWS**: 10.230.0.0/16
- **Azure**: 70.1.0.0/16

d) Connect to the Cloud Network Controller deployed on AWS and make the following verifications:

- Verify that the `dcnm-default-tn` tenant is created and one VPC is deployed:

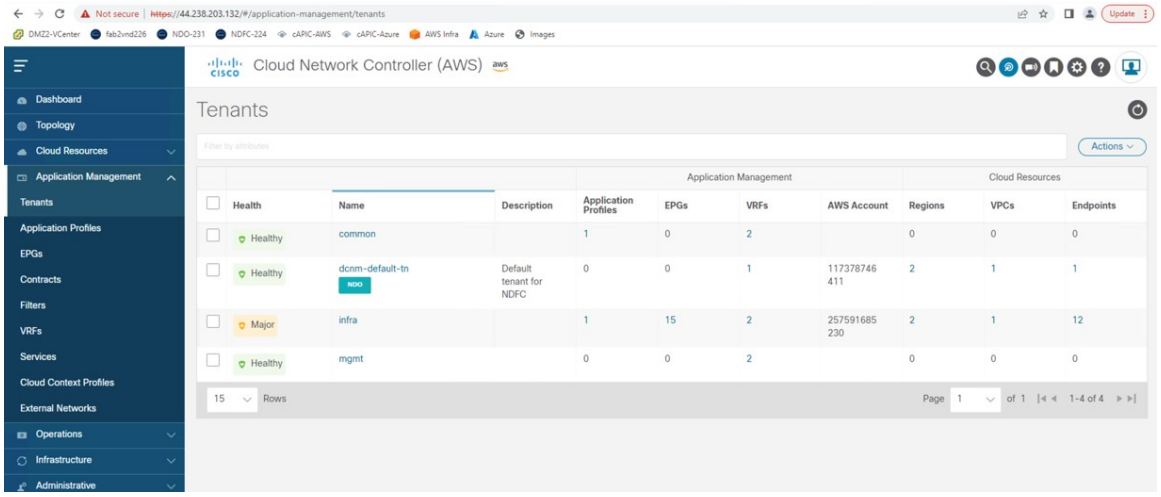

- Verify that the VPC is deployed:

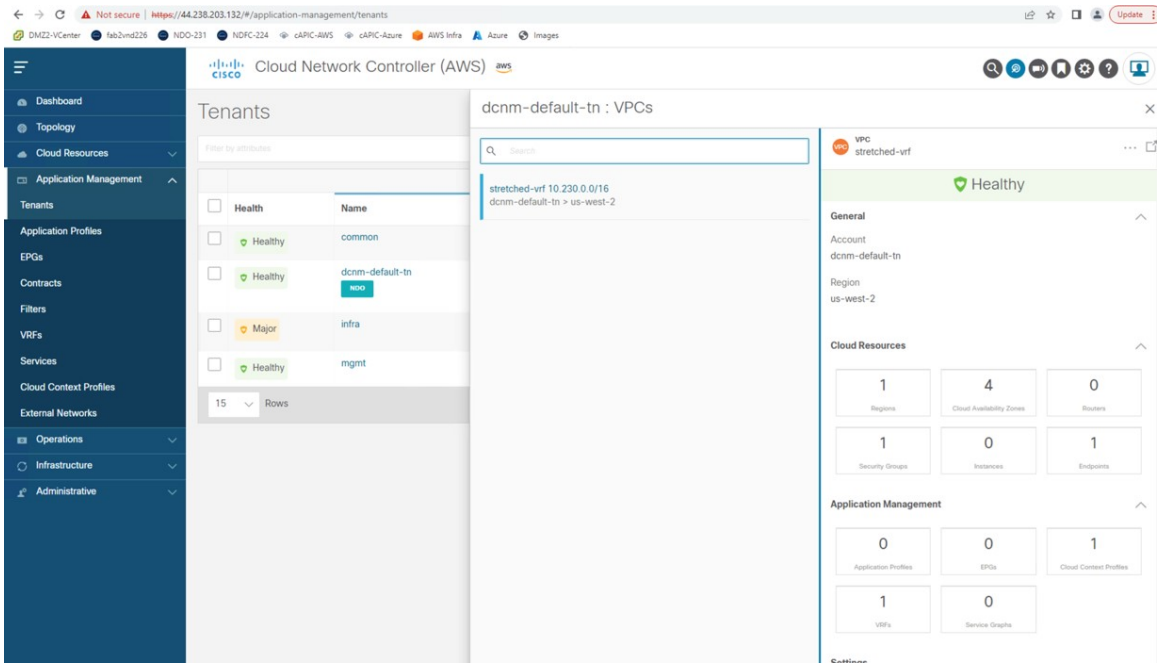

- Using the routing table view from the Cloud Network Controller deployed on AWS, verify that the reachable subnets are:
  - **NDFC**: 172.16.20.0/24
  - **Azure**: 70.1.0.0/16

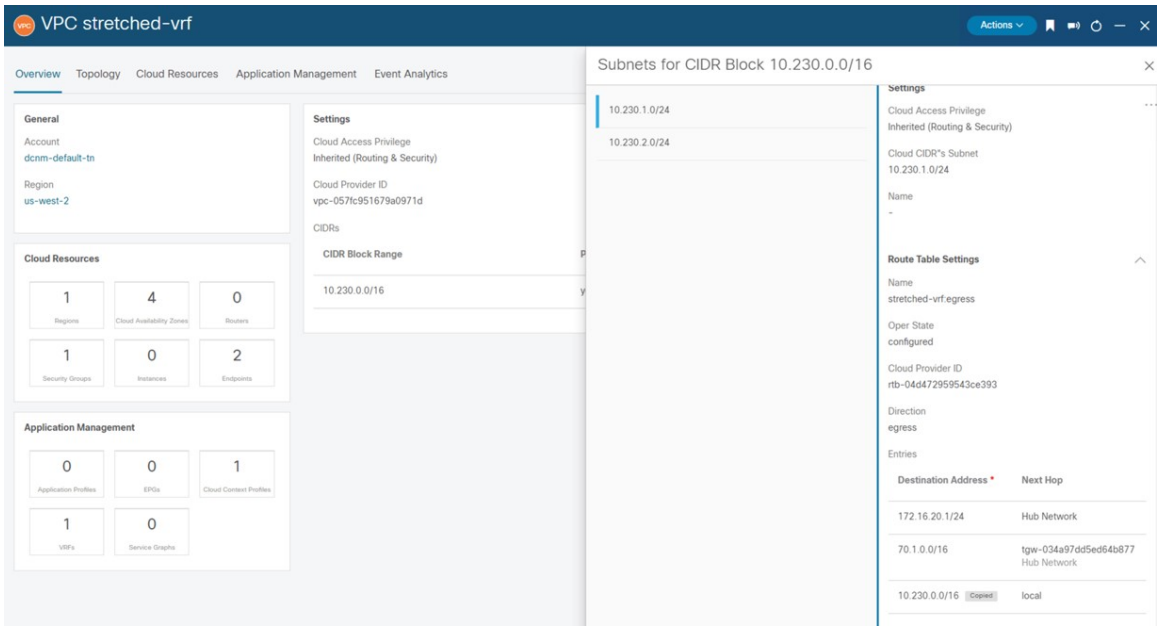

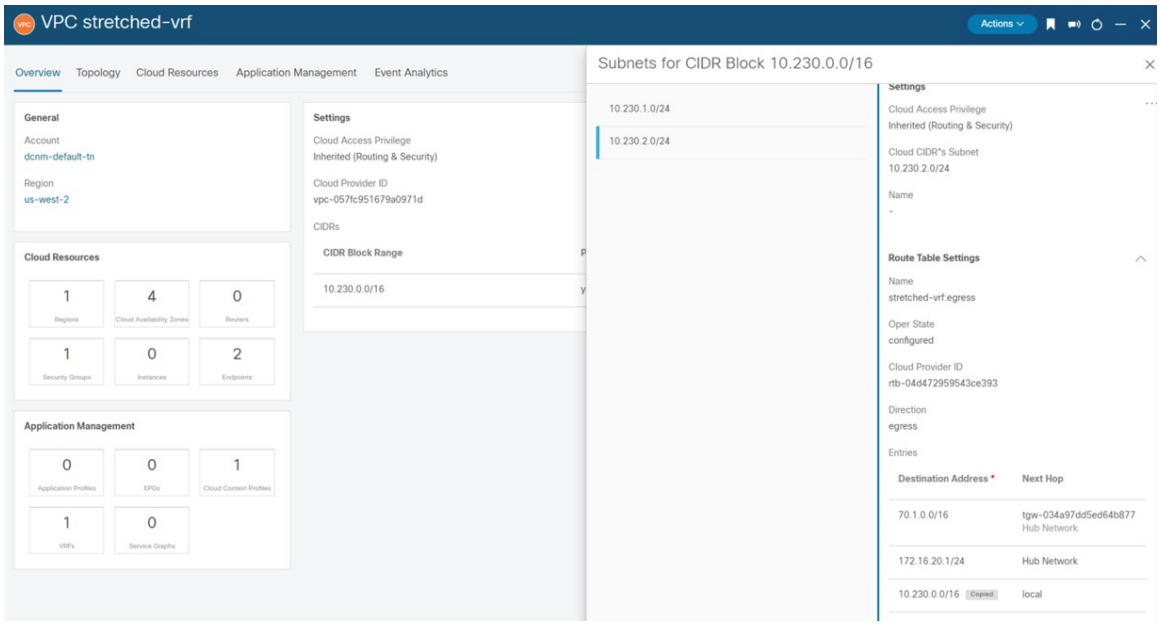

e) In the AWS console, verify the following:

- Verify that you see one VPC and two subnets.

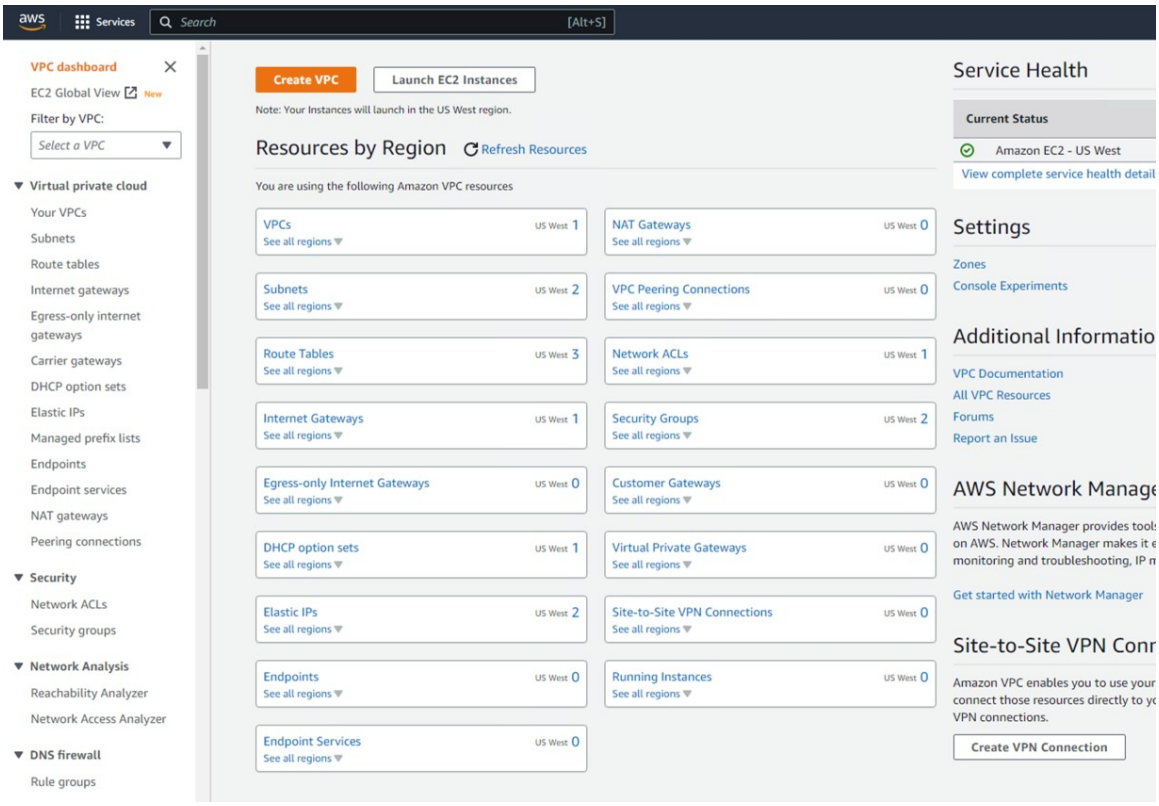
aws
Services
Search
[Alt+S]
VPC dashboard
EC2 Global View
New
Filter by VPC:
Select a VPC
Virtual private cloud
Your VPCs
Subnets
Route tables
Internet gateways
Egress-only internet gateways
Carrier gateways
DHCP option sets
Elastic IPs
Managed prefix lists
Endpoints
Endpoint services
NAT gateways
Peering connections
Security
Network ACLs
Security groups
Network Analysis
Reachability Analyzer
Network Access Analyzer
DNS firewall
Rule groups
Create VPC
Launch EC2 Instances
Note: Your Instances will launch in the US West region.
Resources by Region
Refresh Resources
You are using the following Amazon VPC resources
VPCs US West 1 See all regions
NAT Gateways US West 0 See all regions
Subnets US West 2 See all regions
VPC Peering Connections US West 0 See all regions
Route Tables US West 3 See all regions
Network ACLs US West 1 See all regions
Internet Gateways US West 1 See all regions
Security Groups US West 2 See all regions
Egress-only Internet Gateways US West 0 See all regions
Customer Gateways US West 0 See all regions
DHCP option sets US West 1 See all regions
Virtual Private Gateways US West 0 See all regions
Elastic IPs US West 2 See all regions
Site-to-Site VPN Connections US West 0 See all regions
Endpoints US West 0 See all regions
Running Instances US West 0 See all regions
Endpoint Services US West 0 See all regions
Service Health
Current Status
Amazon EC2 - US West
View complete service health detail
Settings
Zones
Console Experiments
Additional Informatio
VPC Documentation
All VPC Resources
Forums
Report an Issue
AWS Network Manage
AWS Network Manager provides tool
on AWS. Network Manager makes it e
monitoring and troubleshooting, IP n
Get started with Network Manager
Site-to-Site VPN Conn
Amazon VPC enables you to use your
connect those resources directly to yo
VPN connections.
Create VPN Connection

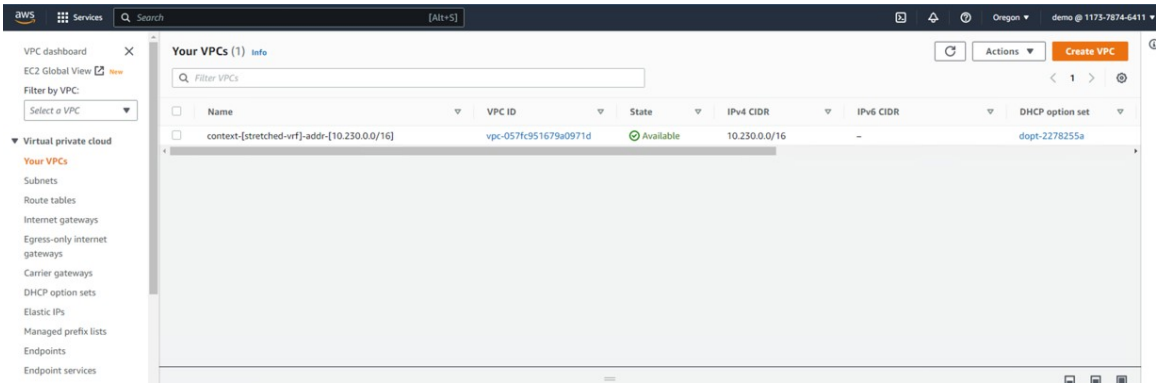
aws
Services
Search
[Alt+S]
Oregon
demo @ 1173-7874-6411
VPC dashboard
EC2 Global View
New
Filter by VPC:
Select a VPC
Virtual private cloud
Your VPCs
Subnets
Route tables
Internet gateways
Egress-only internet gateways
Carrier gateways
DHCP option sets
Elastic IPs
Managed prefix lists
Endpoints
Endpoint services
Your VPCs (1) Info
Filter VPCs
Actions
Create VPC
Name
VPC ID
State
IPv4 CIDR
IPv6 CIDR
DHCP option set
context-[stretched-vrf]-addr-[10.230.0.0/16]
vpc-057fc951679a0971d
Available
10.230.0.0/16
–
dopt-2278255a

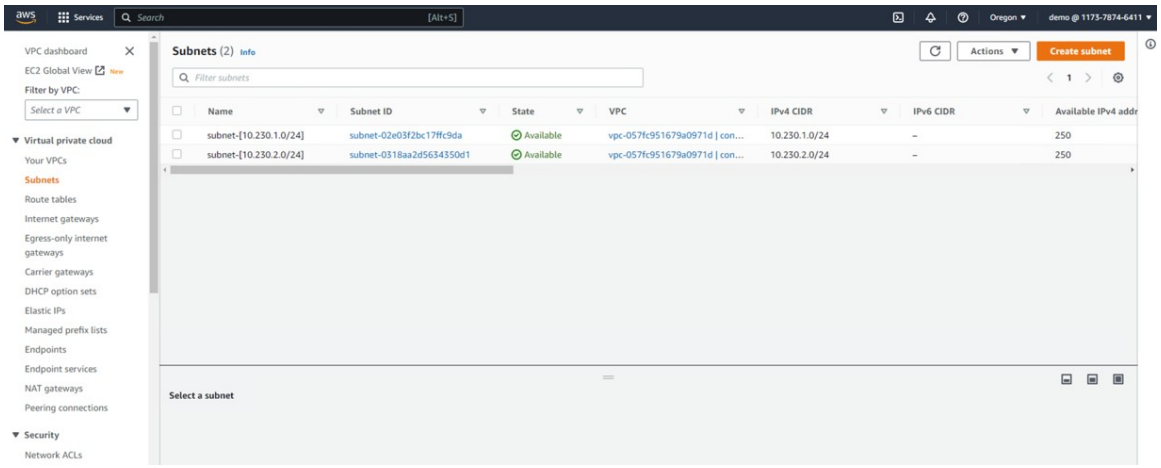

- Verify that you see the routing table.

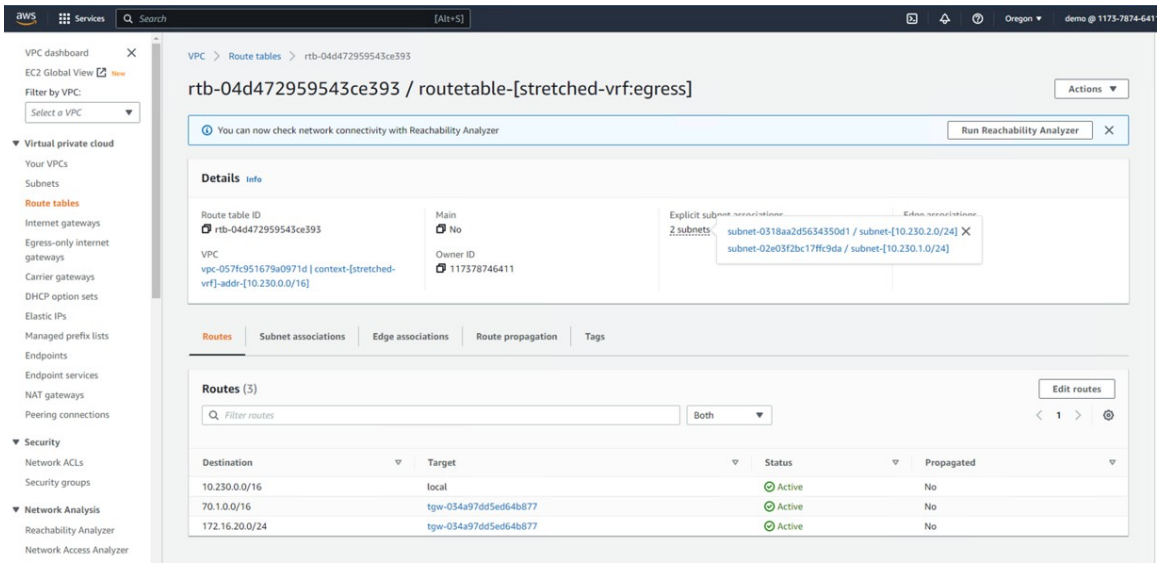

f) Connect to the Cloud Network Controller deployed on Azure and make the following verifications

- Verify that the `dcnm-default-tn` tenant is created:

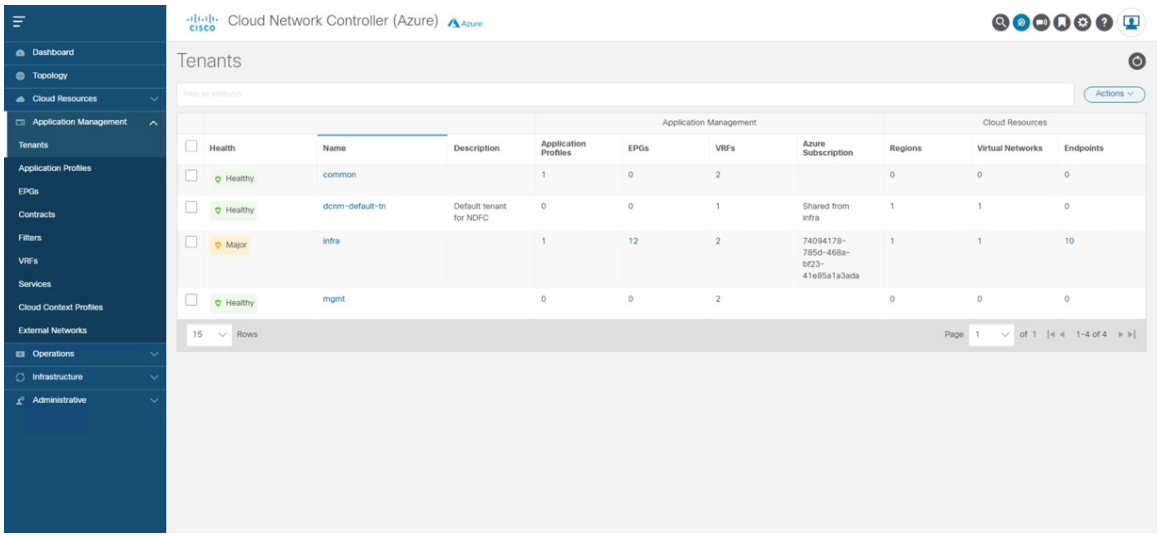

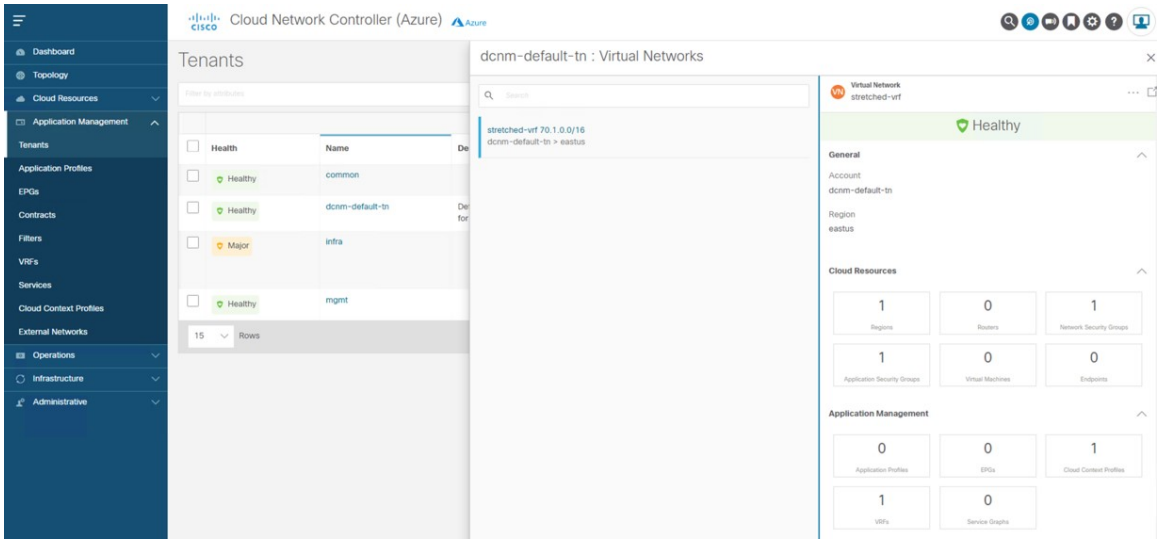

- Verify that the VRF is deployed:

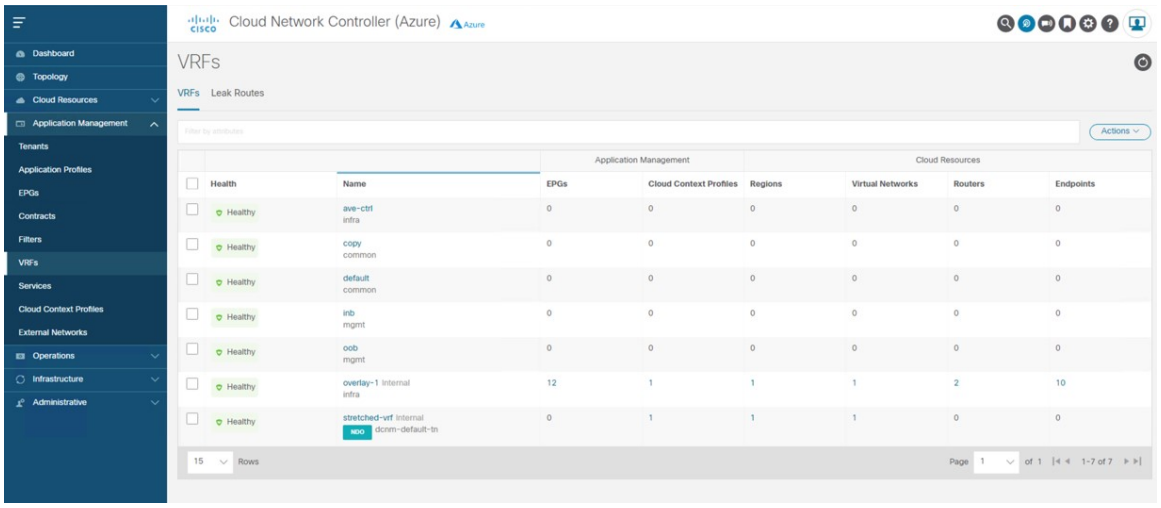

- Using the routing table view from the Cloud Network Controller deployed on AWS, verify that the reachable subnets are:
  - **NDFC**: 172.16.20.0/24
  - **AWS**: 10.230.0.0/16

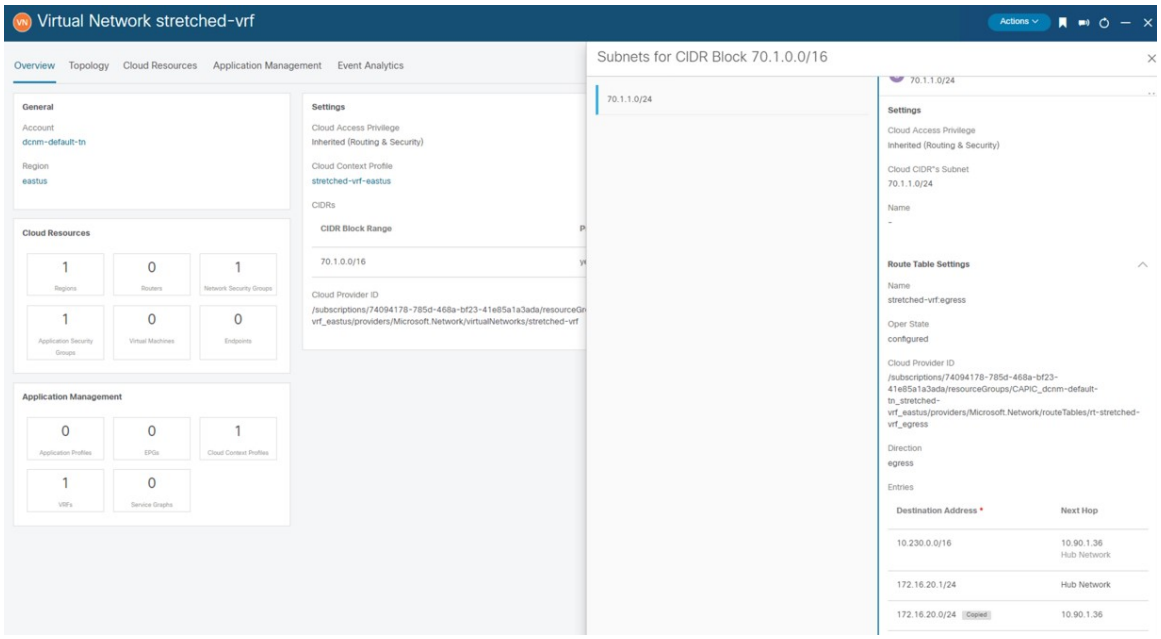

g) In the Azure console, verify that you can see the subnets:

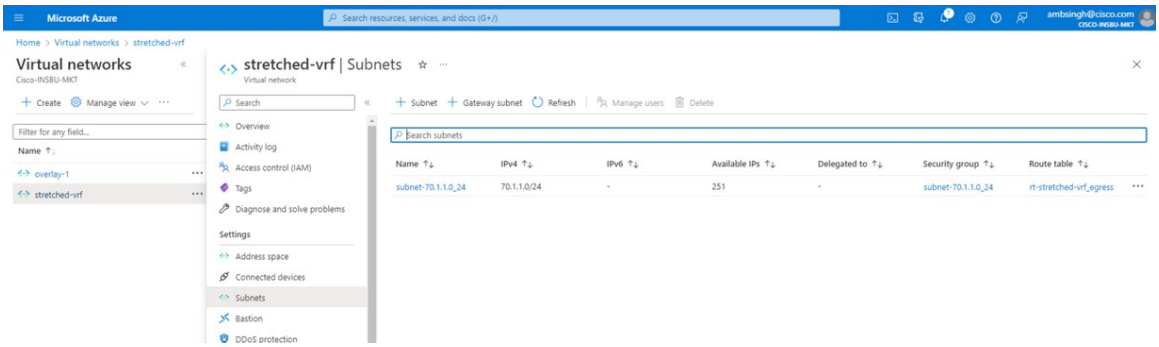
Microsoft Azure
Search resources, services, and docs (G+/)
ambsingh@cisco.com
CISCO-INSBU-MKT
Home > Virtual networks > stretched-vrf
Virtual networks
Cisco-INSBU-MKT
Create
Manage view
Filter for any field...
Name
overlay-1
stretched-vrf
stretched-vrf | Subnets
Virtual network
Search
Overview
Activity log
Access control (IAM)
Tags
Diagnose and solve problems
Settings
Address space
Connected devices
Subnets
Bastion
DDoS protection
Subnet
Gateway subnet
Refresh
Manage users
Delete
Search subnets
Name
IPv4
IPv6
Available IPs
Delegated to
Security group
Route table
subnet-70.1.1.0_24
70.1.1.0/24
-
251
-
subnet-70.1.1.0_24
rt-stretched-vrf_egress

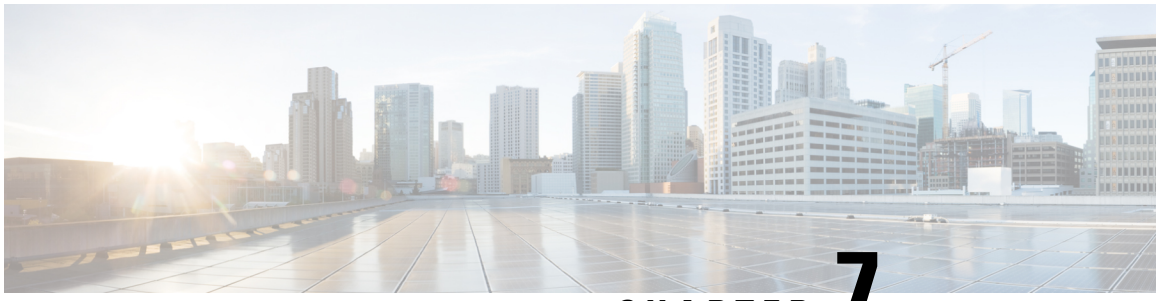

CHAPTER 7

Route Leaking Use Case

About the Route Leaking Use Case

This route leaking use case uses separate templates for each site, which contains VRF and network definitions for the on-premises site, whereas for cloud sites these templates only contain the VRF definition. Unlike the stretched VRF (intra-VRF) use case described in Stretched VRF Use Case, on page 107, which does not require any configurations for exchanging prefixes between the sites because the same VRF is stretched to all sites, you must configure VRF leaking for this use case because each site uses a different VRF.

To propagate the prefixes between the sites (on-premises as well as cloud sites), you must explicitly configure route leaking on the respective templates associated with the sites.

Figure 162:

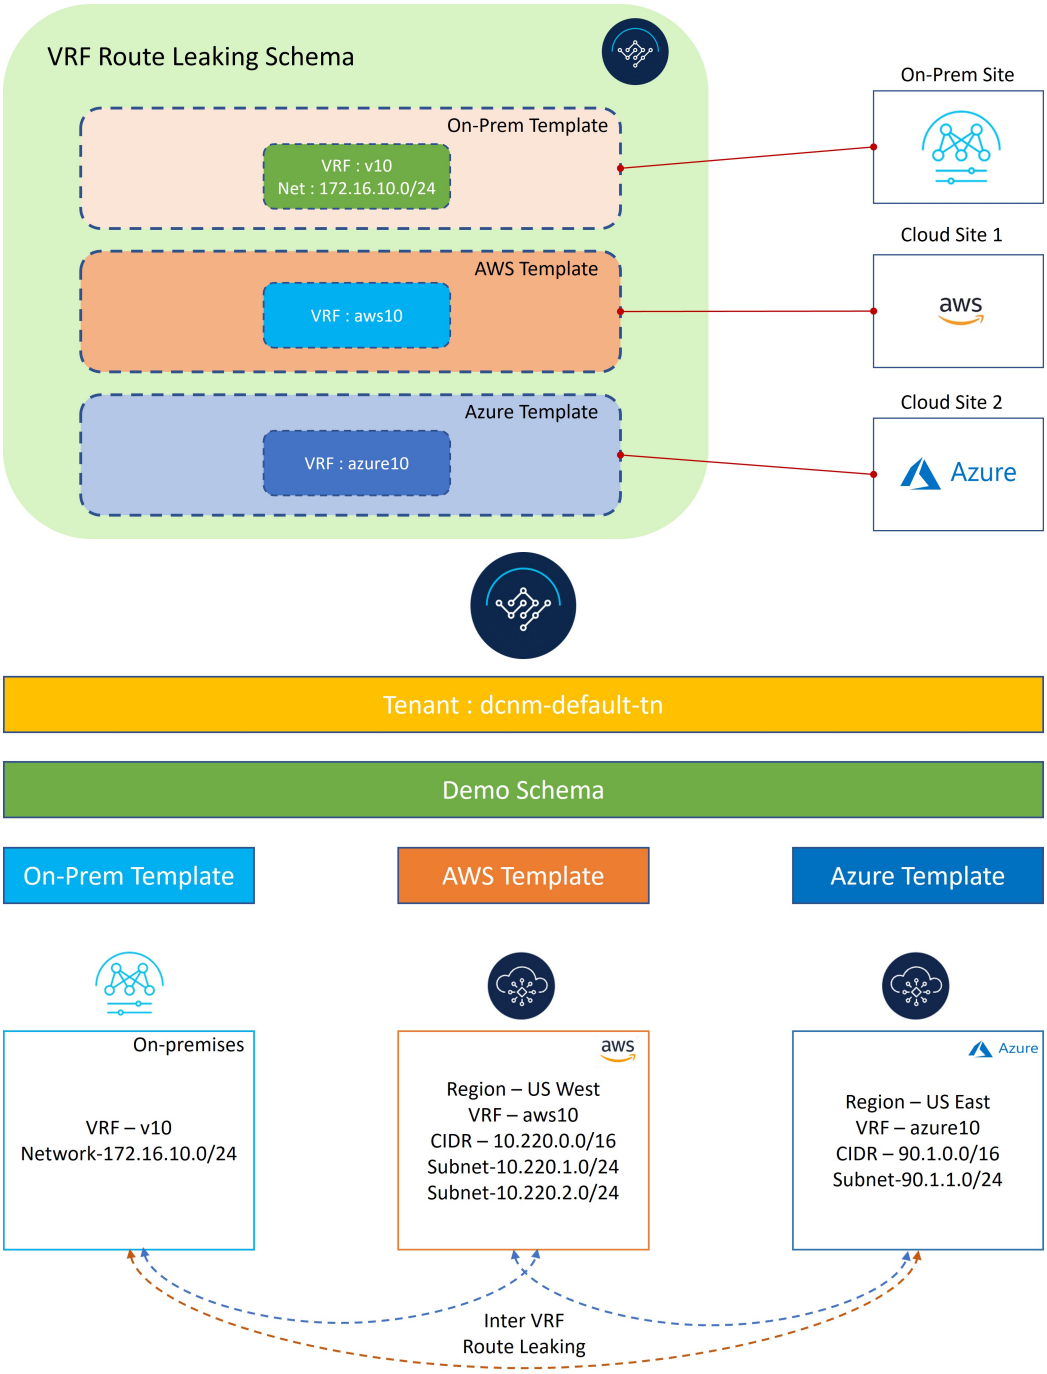

As shown in the figure above, each site has a separate associated template, which contains VRF/network definitions specific to that site only. `On-Prem Template` is associated to the NDFC managed on-premises site, whereas `AWS Template` and `Azure Template` are associated to the AWS and Azure cloud sites, respectively. Inter-VRF route leaking is configured explicitly between different VRFs to allow communication between the sites.

Configure the Necessary Templates

Use the procedures in the following sections to configure the templates that you will need for the route leaking use case.

Configure the On-Premises Site Template

In this section, you will configure the `On-Prem Template` that will be associated to the NDFC managed on-premises site.

Step 1 In NDO, navigate to **Application Management** > **Schemas** and click **Add Schema**.

Step 2 Provide the schema name and click **Add**.

For this use case, we will name the new schema `VRF Route Leaking Schema`.

Figure 163:

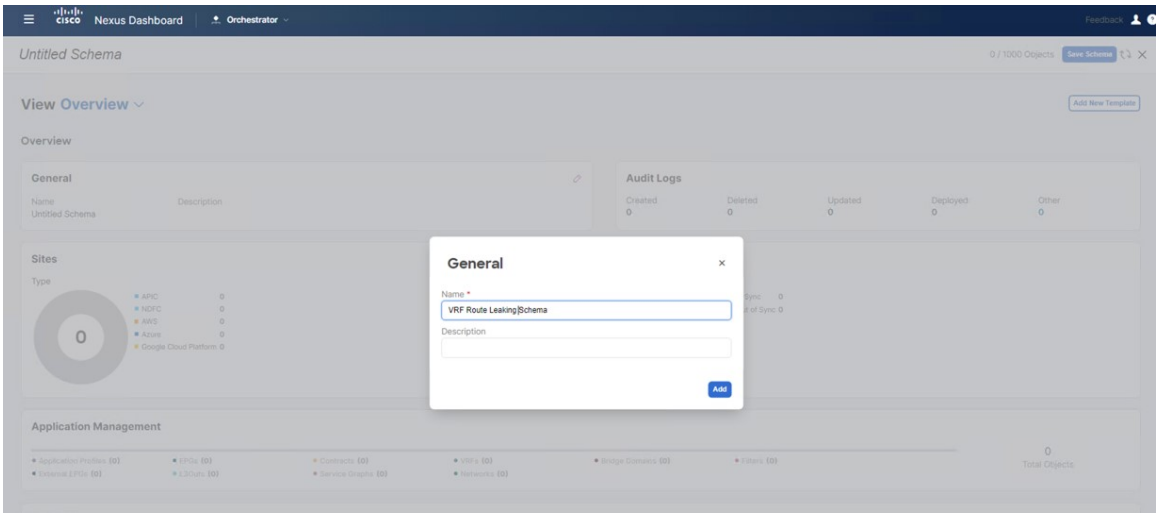

You are returned to the **Overview** page for the new `VRF Route Leaking Schema` schema.

Step 3 Under the `VRF Route Leaking Schema` schema, click **Add New Template**.

Step 4 Choose the NDFC template.

Step 5 Enter a name in the **Display Name** field to create an NDFC-type template (for example, `On-Prem Template`).

Step 6 Select the `dcnm-default-tn` tenant in the **Select a Tenant** field to map the template to that tenant.

Figure 164:

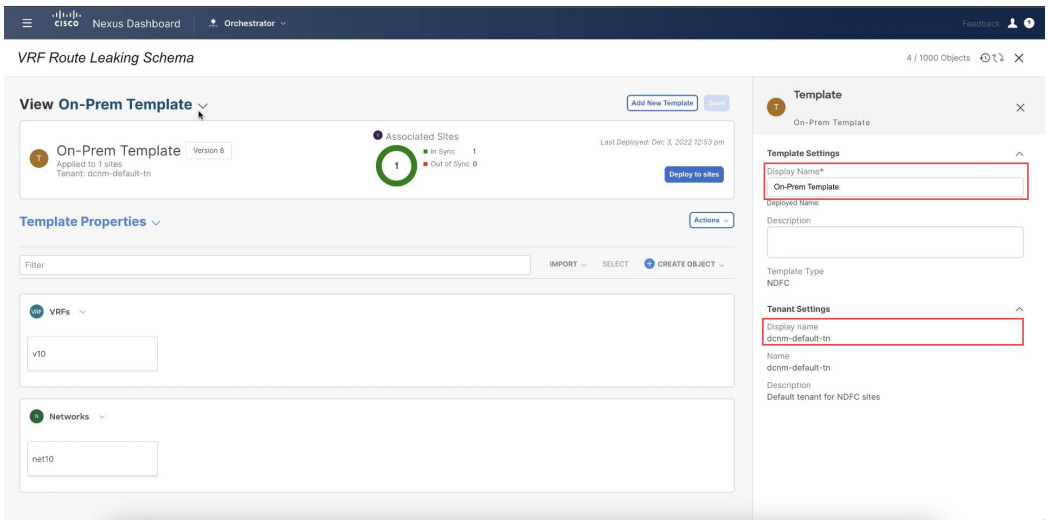

Step 7 Under **Template Properties**, click **Create Object** and choose **VRF** to create a VRF that will be used with the NDFC managed on-premises site.

Note If you have an on-premises VRF already created that you want to use instead of creating a new VRF, under **Template Properties**, click **Import**, then import the already-created VRF.

Currently, support is only available for importing VRFs and networks from on-premises sites.

Step 8 Enter a name in the **Display Name** field for this VRF (for example, `v10`).

Figure 165:

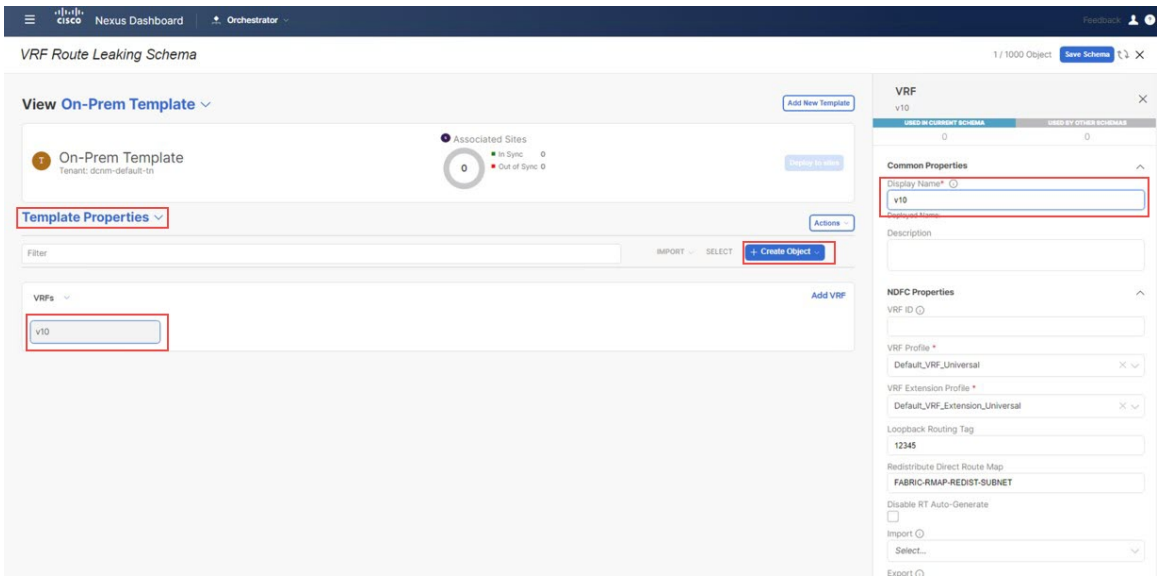

Step 9 Under **Template Properties**, click **Create Object** and choose **Network** to create a network.

Note If you have a network already created that you want to use instead of creating a new network, under **Template Properties**, click **Import**, then import the already-created network.

Step 10 Enter a name in the **Display Name** field for the network (for example, `net10`).

Step 11 In the **Virtual Routing & Forwarding** field, choose the `v10` VRF to map the `net10` network to that VRF.

Figure 166:

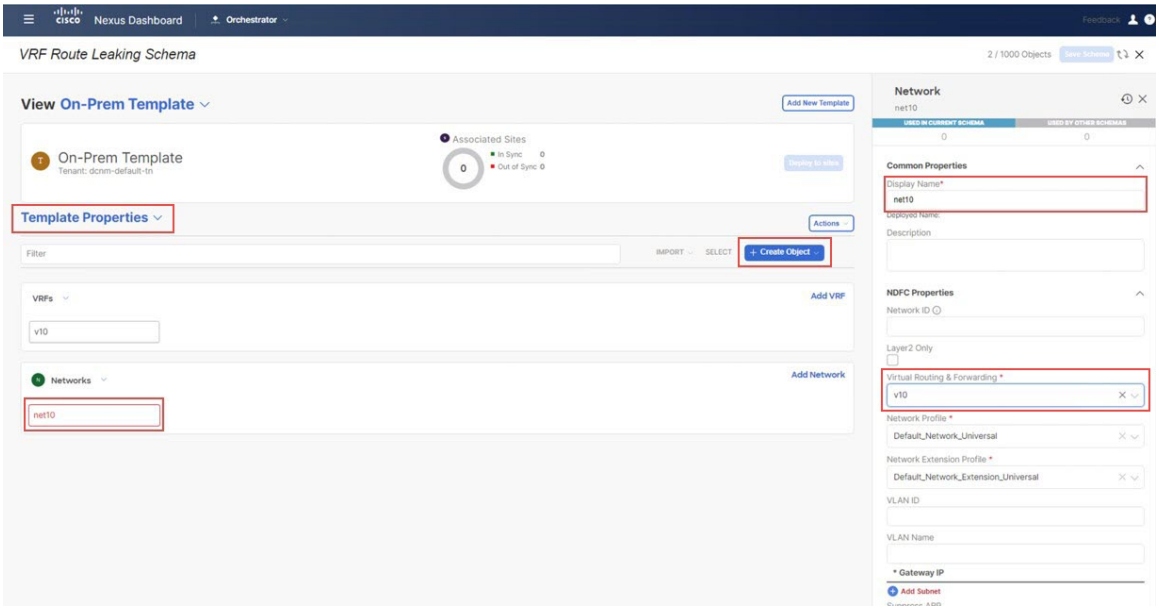

Step 12 In the **Gateway IP** field, click **Add Subnet** and provide the gateway IP address, then click **Add**.

Figure 167:

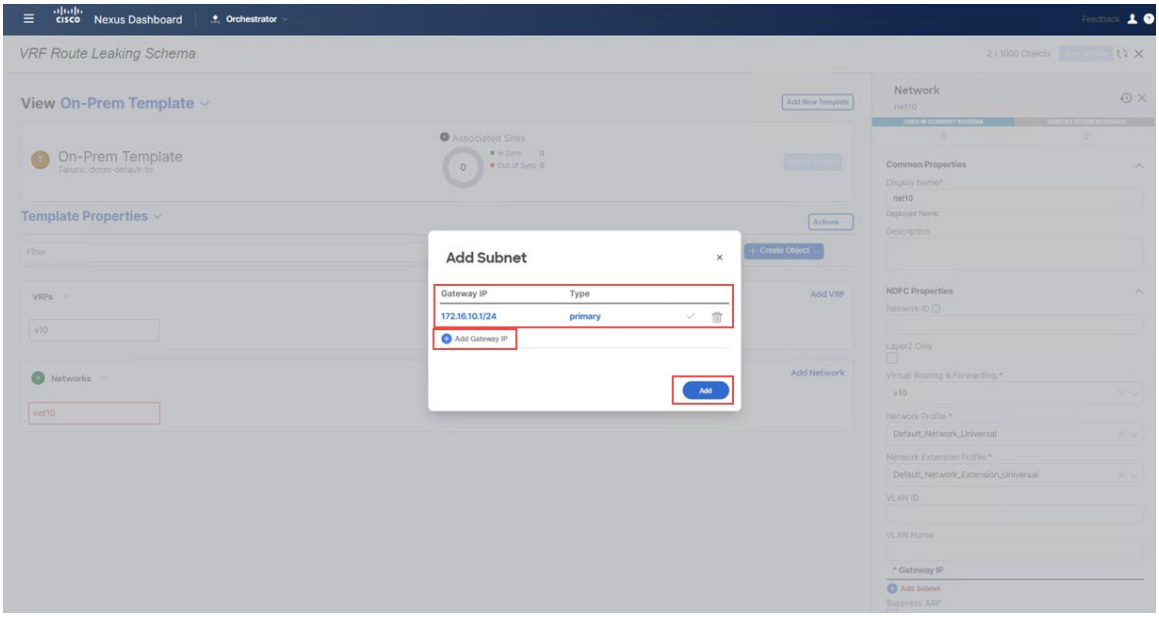

The gateway IP address is now displayed in the **Gateway IP** field.

Figure 168:

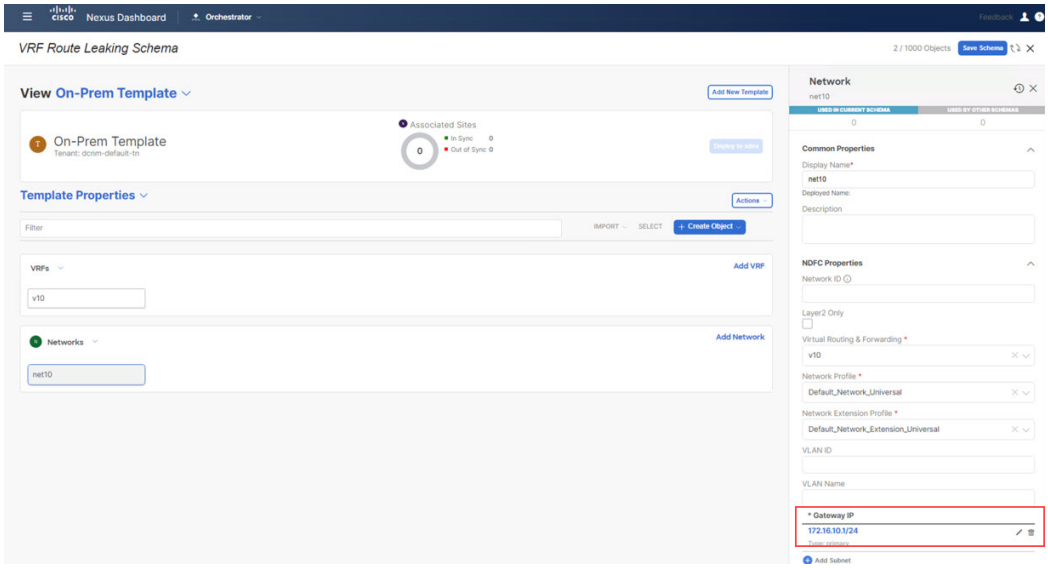

Step 13 Define other optional parameters for this network, if necessary.

Step 14 In the **Template Properties** area, click **Actions** > **Sites Association**.

Figure 169:

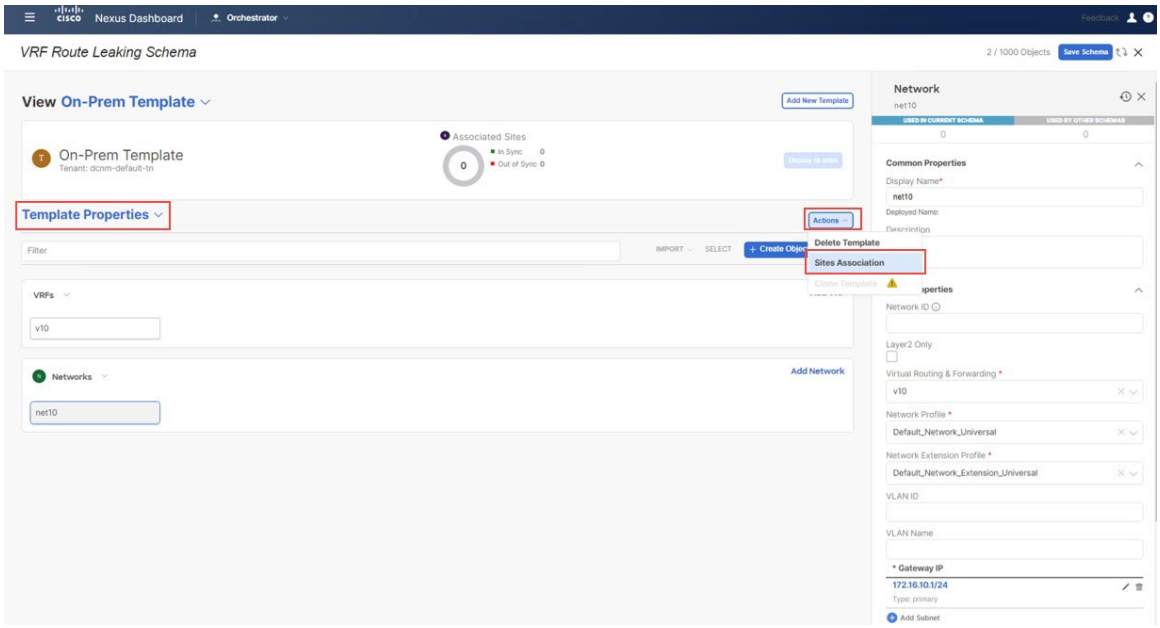

Step 15 Associate this template only to the on-premises site (the `Sydney` site in this example use case), then click **Ok**.

Figure 170:

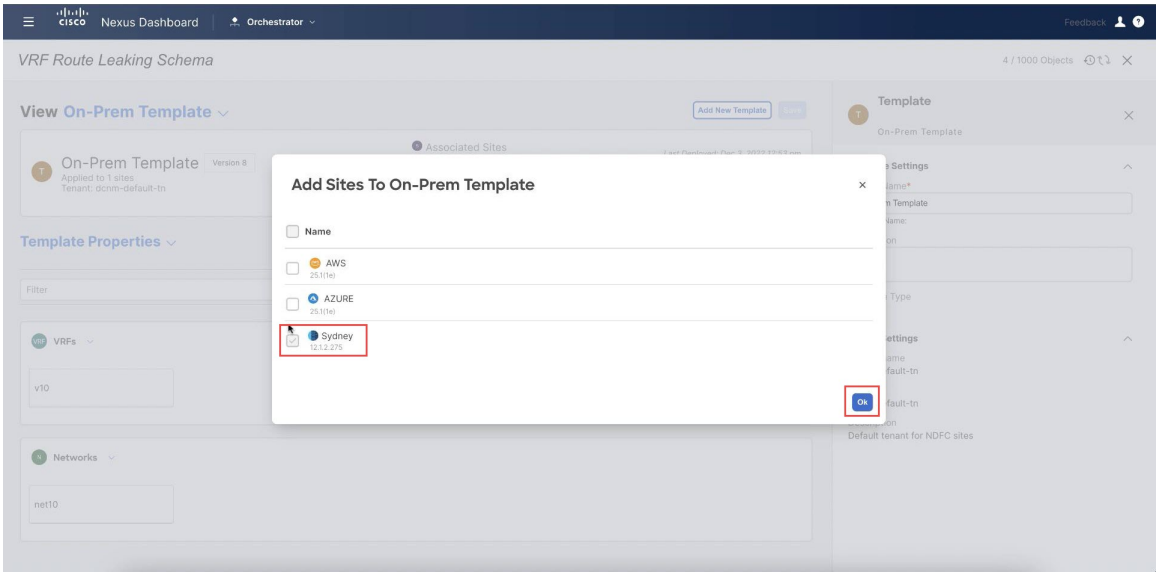

Step 16 Click **Template Properties** and select the on-premises site (the `Sydney` site in this example use case), then select the `v10` VRF.

Step 17 In the right pane, click **Add Static Leaf**.

Figure 171:

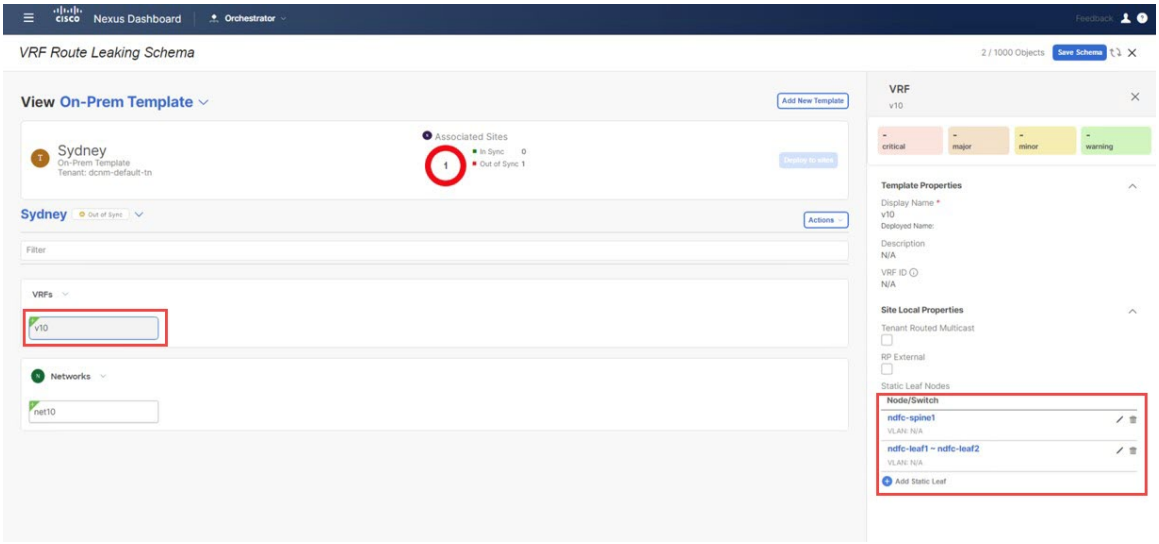

The **Add Static Leaf** window appears.

Step 18 In the **Leaf** field, select the leaf/border/border gateway device where this VRF is to be deployed and click **Ok**.

In this example, you need to deploy the VRF on the leaf nodes (where the endpoints part of the network mapped to the VRF will be connected) and on the BGW spine node to be able to extend the Layer 3 connectivity for the VRF towards the cloud sites.

Step 19 To attach the network to the leaf switches, click the `net10` network, then click **Add Static Port** to add the ports where you want to deploy this network.
The **Add Static Port** window appears.

Step 20 In the **Add Static Port** window, click **Add Path**.
The **Add Static Port** window appears.

Step 21 In the **Leaf** field, select the device where you want to deploy this network.

Step 22 (Optional) Enter the necessary information in the **VLAN** field.

Step 23 In the **Ports** field, select the ports where you want to deploy this network.

Step 24 Click **Save**.

Figure 172:

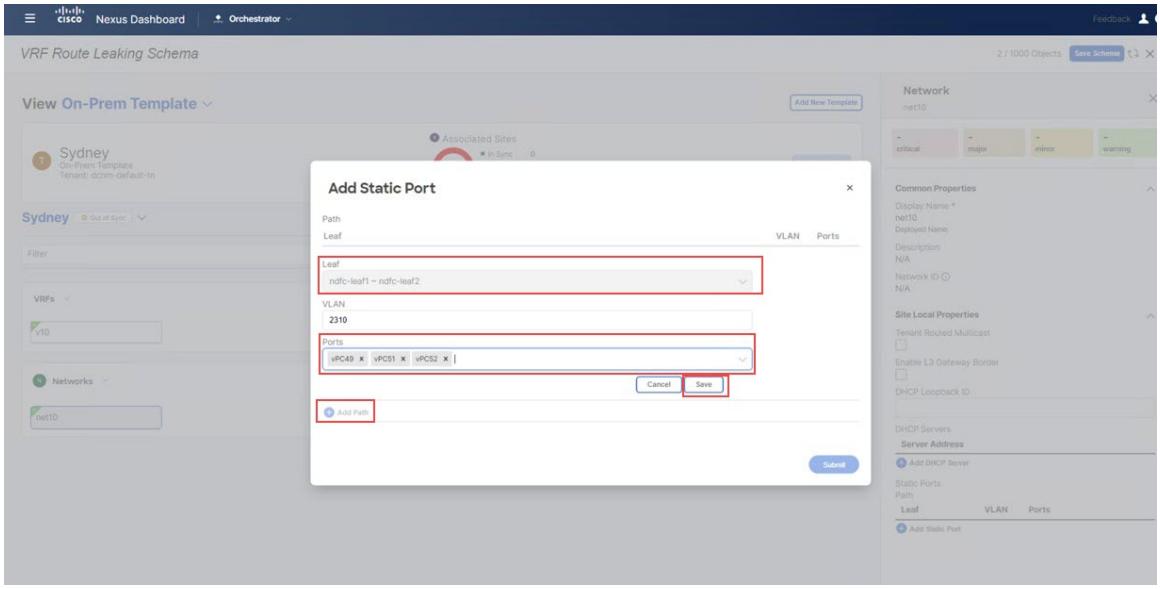

You are returned to the **Add Static Port** window.

Step 25 In the **Add Static Port** window, click **Submit**.

Figure 173:

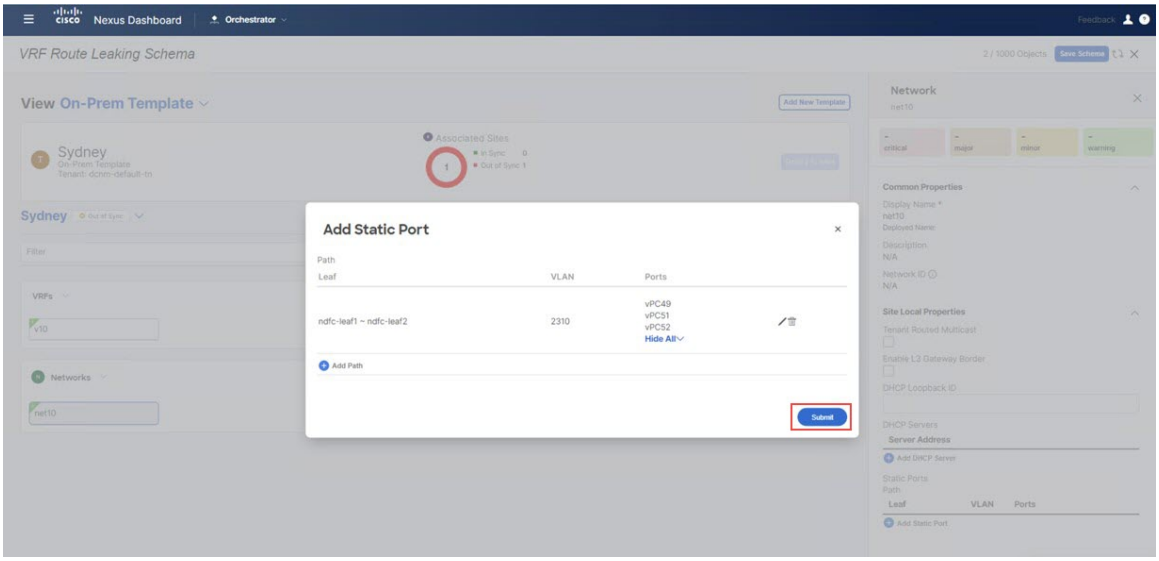

You are returned to the on-premises template window.

Figure 174:

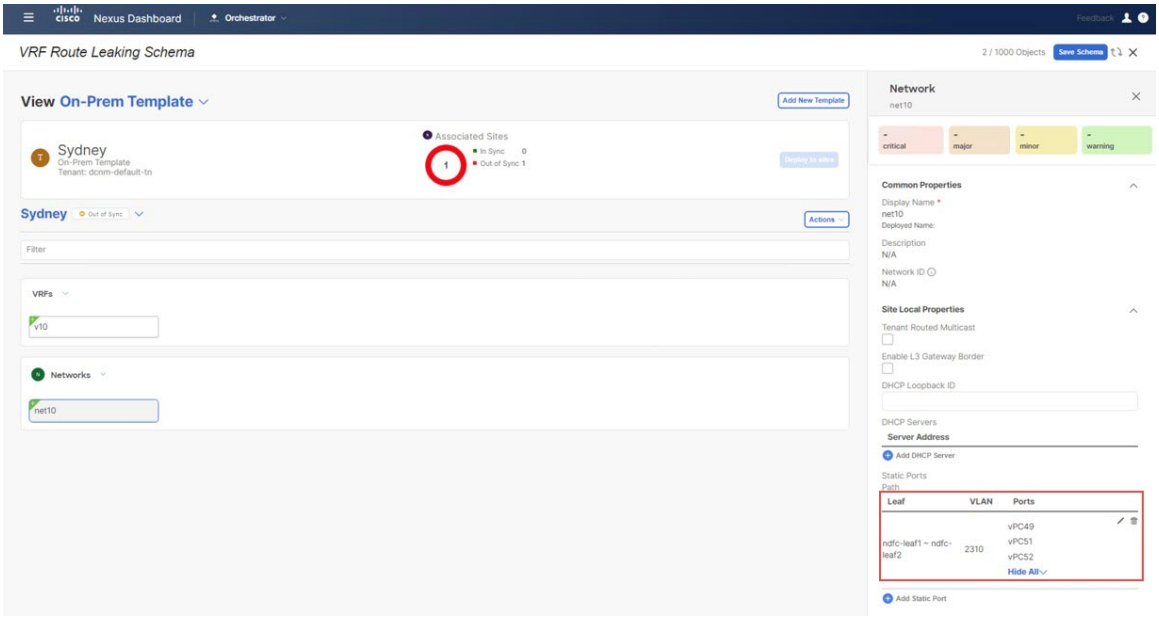

Step 26 Click the arrow next to the on-premises site, and from the drop-down menu, select **Template Properties**.

Step 27 Click **Deploy to Sites**.

Figure 175:

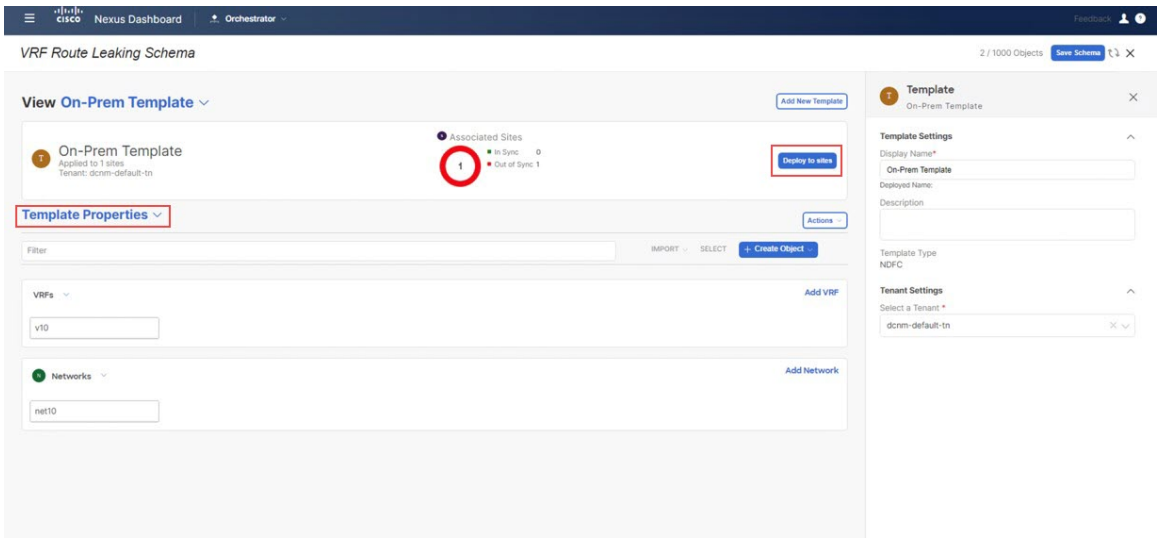

Step 28 Deploy `On-Prem Template` to the sites.

- Click **Deployment Plan** for additional verification.

  Click on the on-premises site to see the deployment plan for that specific site.

  Figure 176:

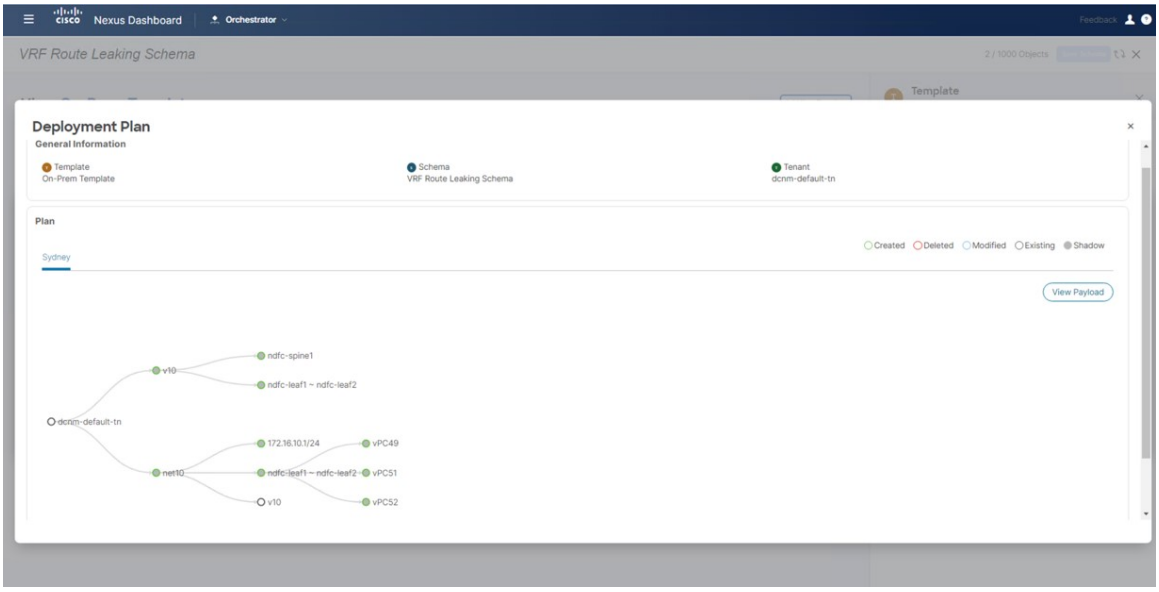

- Click **Deploy** to have NDO push the configurations to NDFC.

  This pushes the NDO configurations to NDFC.

Step 29 In NDFC, verify that the VRF was deployed successfully.

Figure 177:

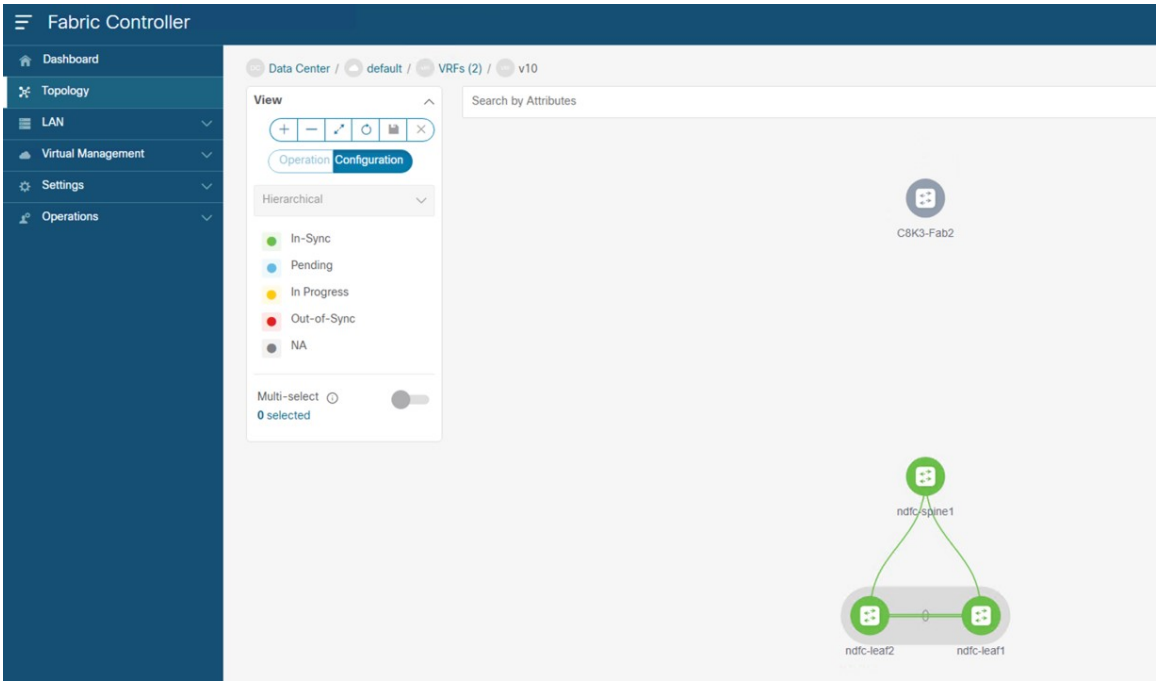

What to do next

Follow the procedures provided in Configure the Azure Site Template, on page 153.

Configure the Azure Site Template

In this section, you will configure the `Azure Template` that will be associated to the Azure site.

Before you begin

Follow the procedures provided in Configure the On-Premises Site Template, on page 145.

Step 1 Under the `VRF Route Leaking Schema` schema, click **Add New Template**.

Step 2 Choose the NDFC template.

Step 3 Enter a name in the **Display Name** field to create an NDFC-type template for the Azure site (for example, `Azure Template`).

Step 4 Select the `dcnm-default-tn` tenant in the **Select a Tenant** field to map the template to that tenant.

Figure 178:

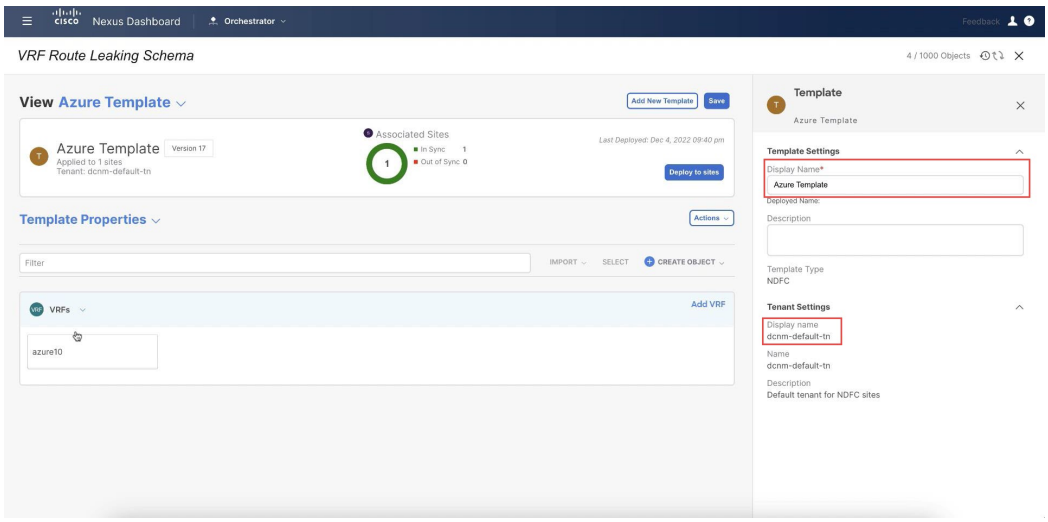

Step 5 Under **Template Properties**, click **Create Object** and choose **VRF** to create a VRF that will be used with the Azure site.

Figure 179:

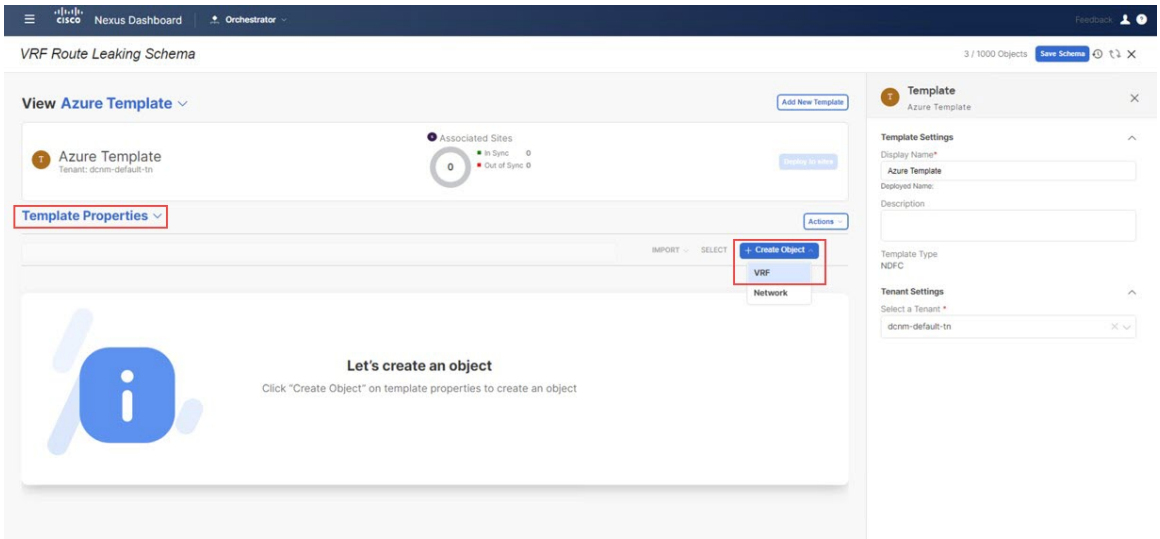

Step 6 Enter a name in the **Display Name** field for this VRF (for example, `azure10`).

Figure 180:

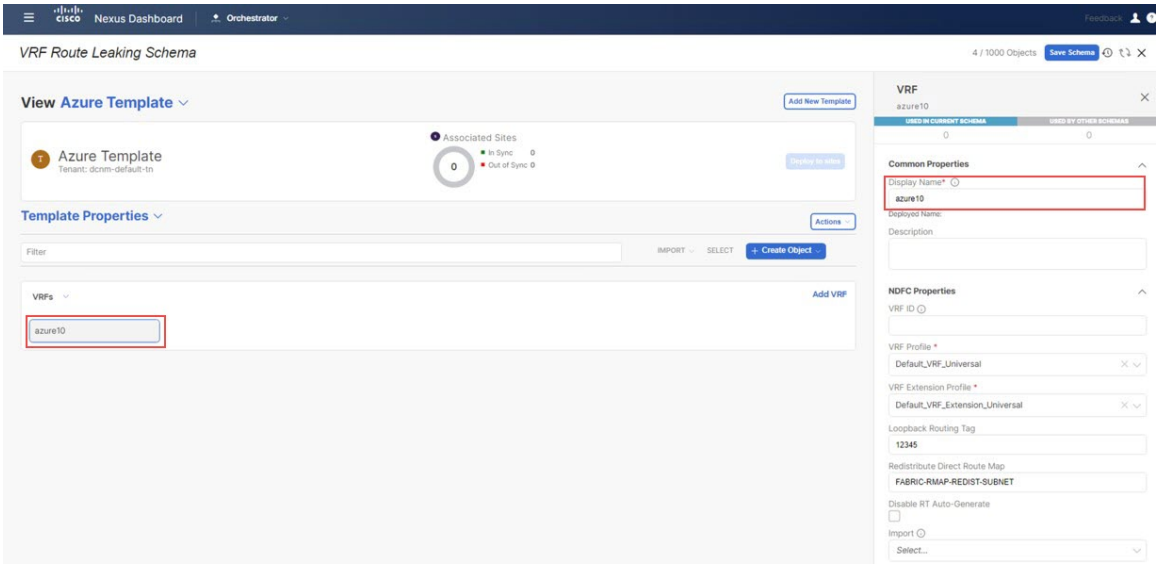

Step 7 In the **Template Properties** area, click **Actions** > **Sites Association**.

Figure 181:

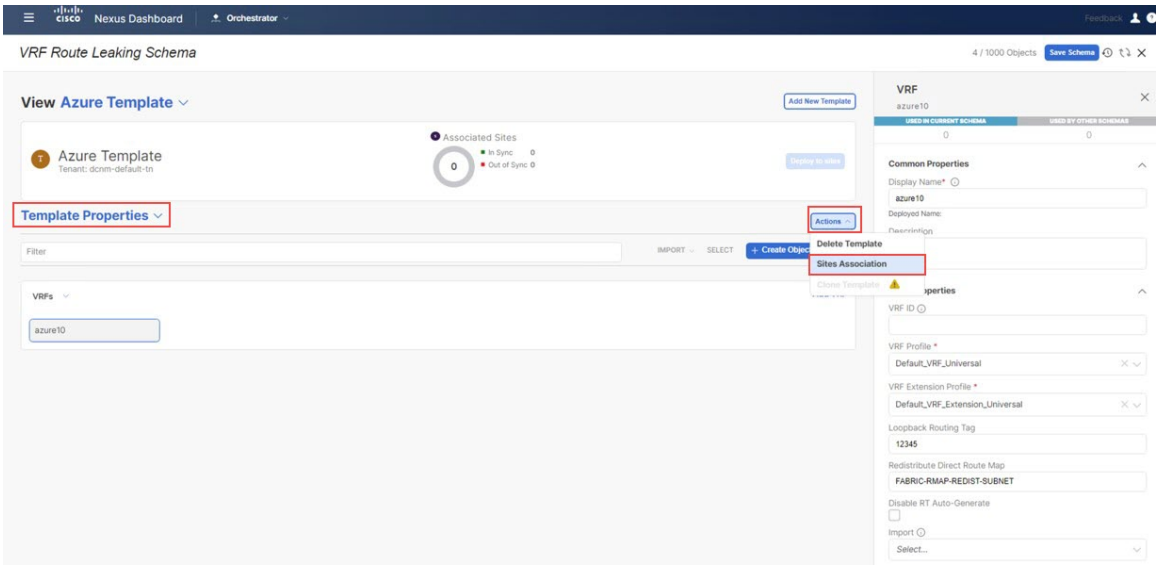

Step 8 Associate this template only to the Azure site, then click **Ok**.

Figure 182:

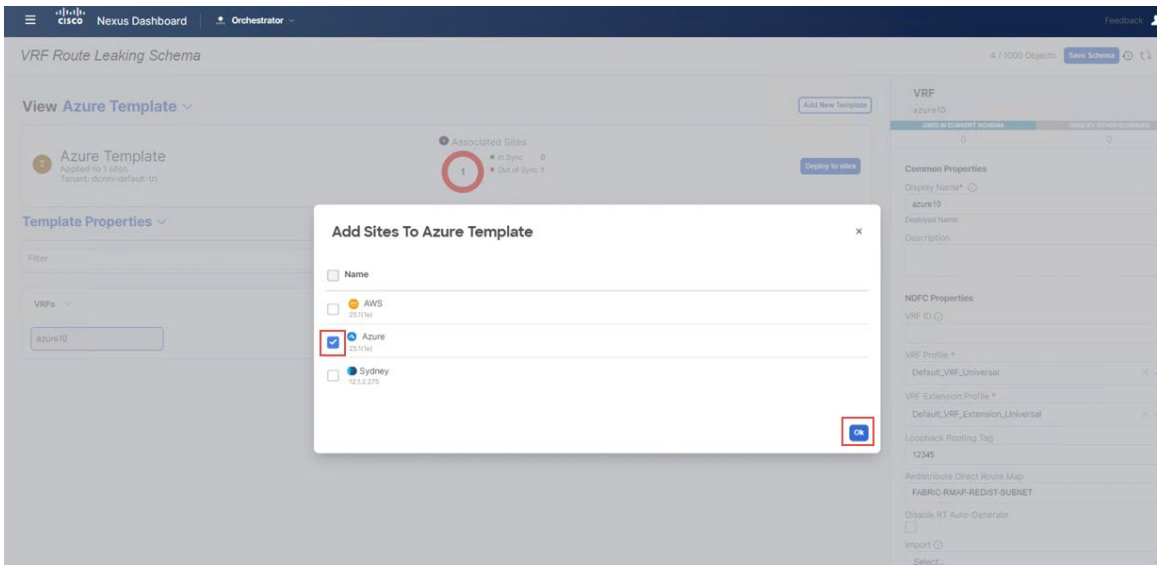

Step 9 Click the `azure10` VRF, then click **Add Region** to create the VNet in a selected region. The **Add Cloud Region CIDRs** window appears.

Step 10 In the **Region** field, choose the region where you want to create the VNet.

Step 11 In the **CIDR** field, click **Add CIDRs** and define a CIDR block for the VNet.

Step 12 Click **Add Subnet** to create the subnets, then click **Save**.

Figure 183:

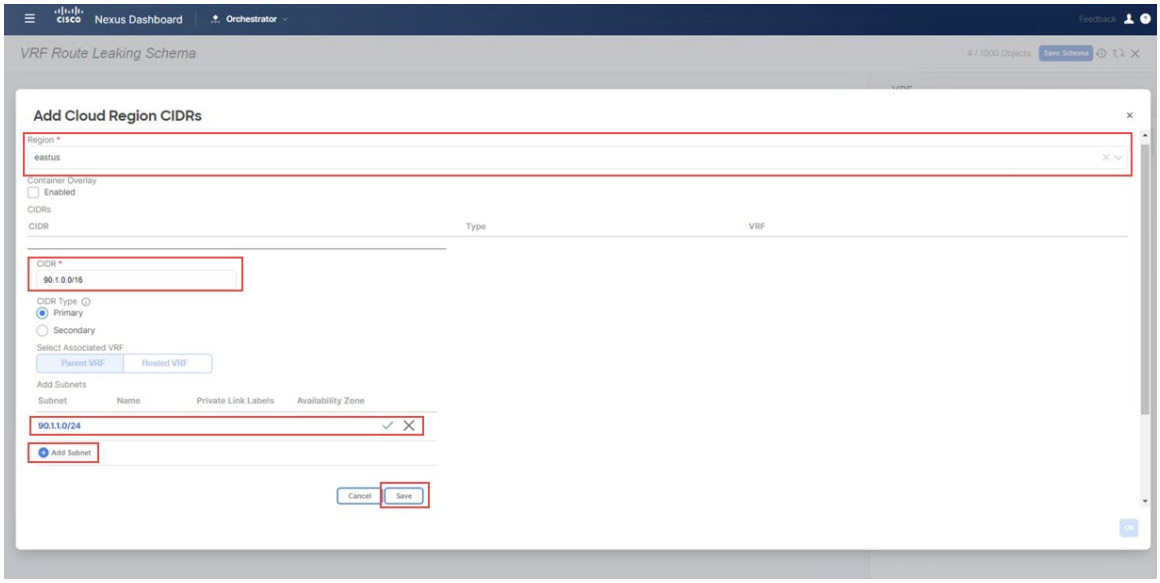

Step 13 Check the box under the **VNet Peering** field, then select the hub network that was created on the Cisco Cloud Network Controller for Azure.

Figure 184:

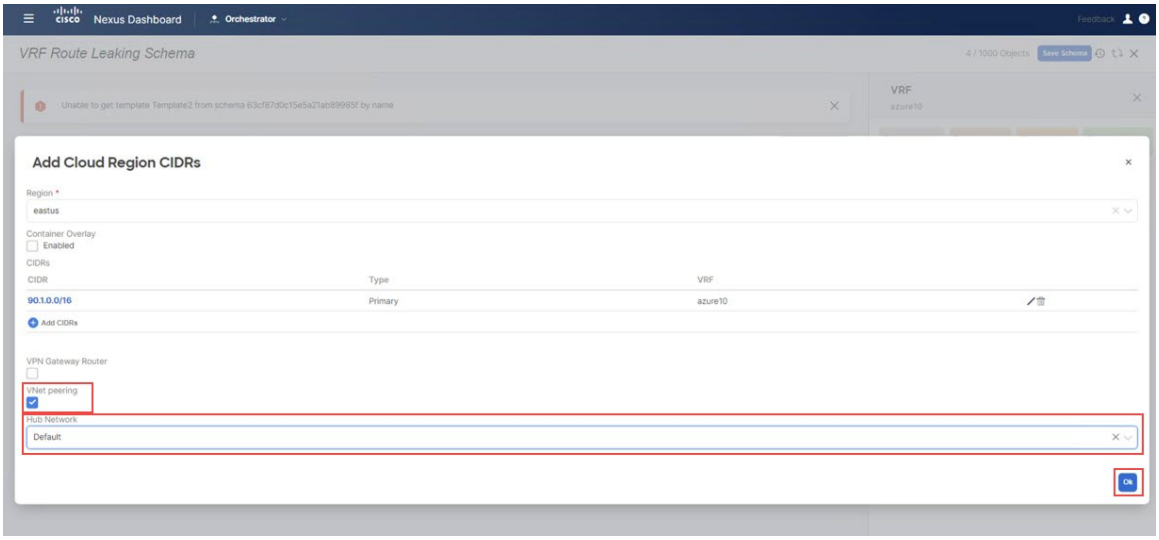

Step 14 Click **Ok**.
You are returned to the Azure template window.

Step 15 Click the arrow next to the Azure site, and from the drop-down menu, select **Template Properties**.

Step 16 Click **Deploy to Sites**.

Step 17 Deploy `Azure Template` to the sites.

- Click **Deployment Plan** for additional verification.

  Click on the Azure site to see the deployment plan for that specific site.

  Figure 185:

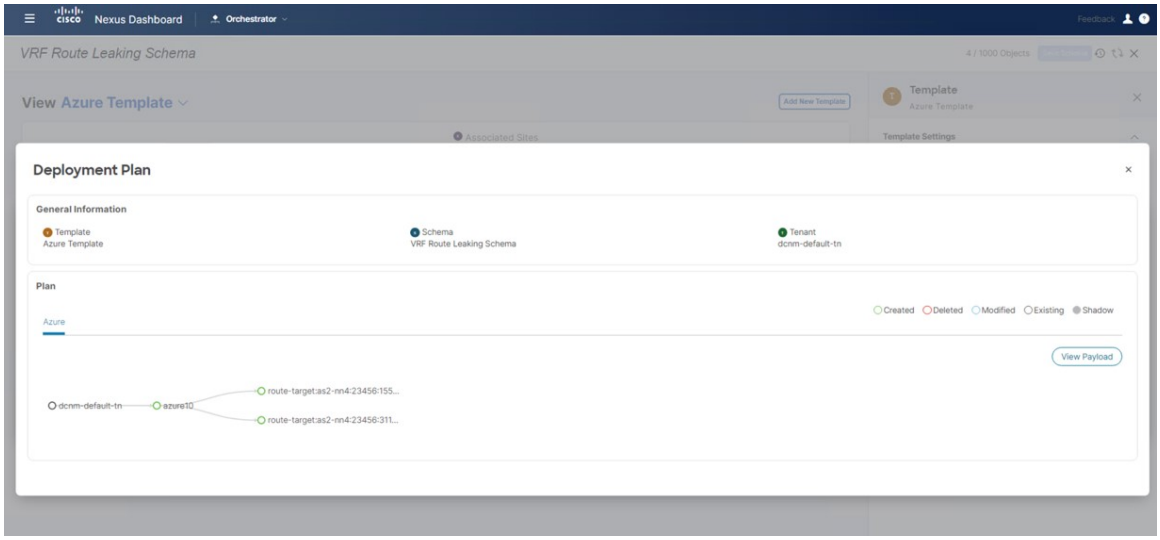

- Click **Deploy** to have NDO push the configurations to NDFC.

To verify that the configurations were pushed out correctly, connect to the Cloud Network Controller deployed on Azure and navigate to **Cloud Resources** > **Virtual Networks**, then click the `azure10` VNet and use the information in the Overview page for additional verifications:

Figure 186:

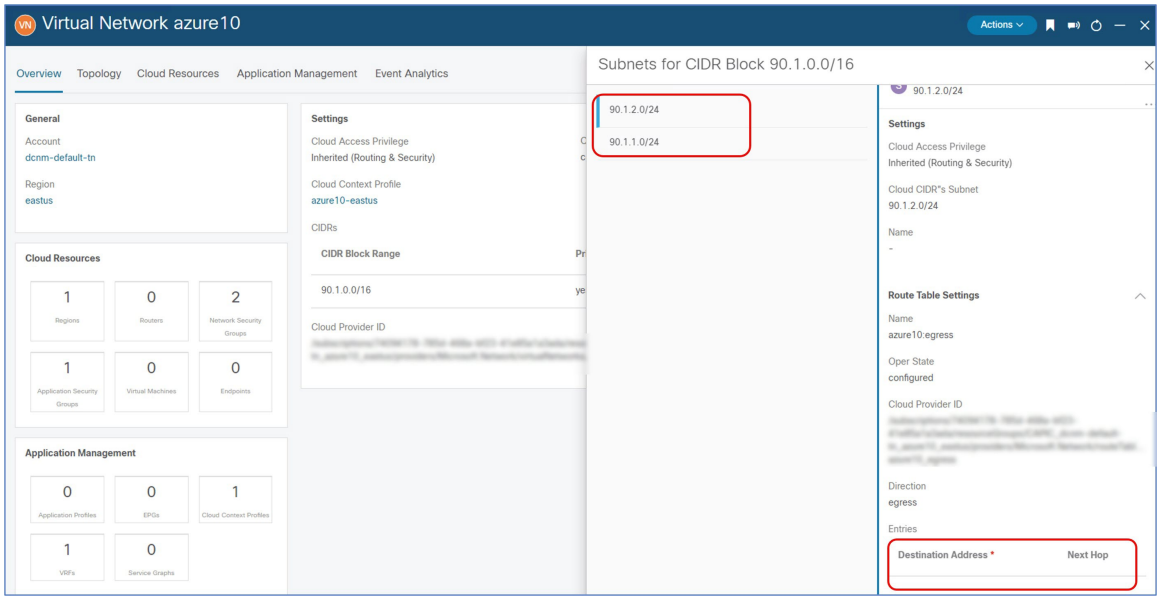

Note that there is no destination address configured at this point in the process, so the Azure site cannot talk to any other site yet at this point in the process. This destination address configuration will be pushed out after you have completed the route leaking procedure.

What to do next

Follow the procedures provided in Configure the AWS Site Template, on page 158.

Configure the AWS Site Template

In this section, you will configure the `AWS Template` that will be associated to the AWS site.

Before you begin

Follow the procedures provided in Configure the Azure Site Template, on page 153.

Step 1 Under the `VRF Route Leaking Schema` schema, click **Add New Template**.

Step 2 Choose the NDFC template.

Step 3 Enter a name in the **Display Name** field to create an NDFC-type template for the AWS site (for example, `AWS Template`).

Step 4 Select the `dcnm-default-tn` tenant in the **Select a Tenant** field to map the template to that tenant.

Figure 187:

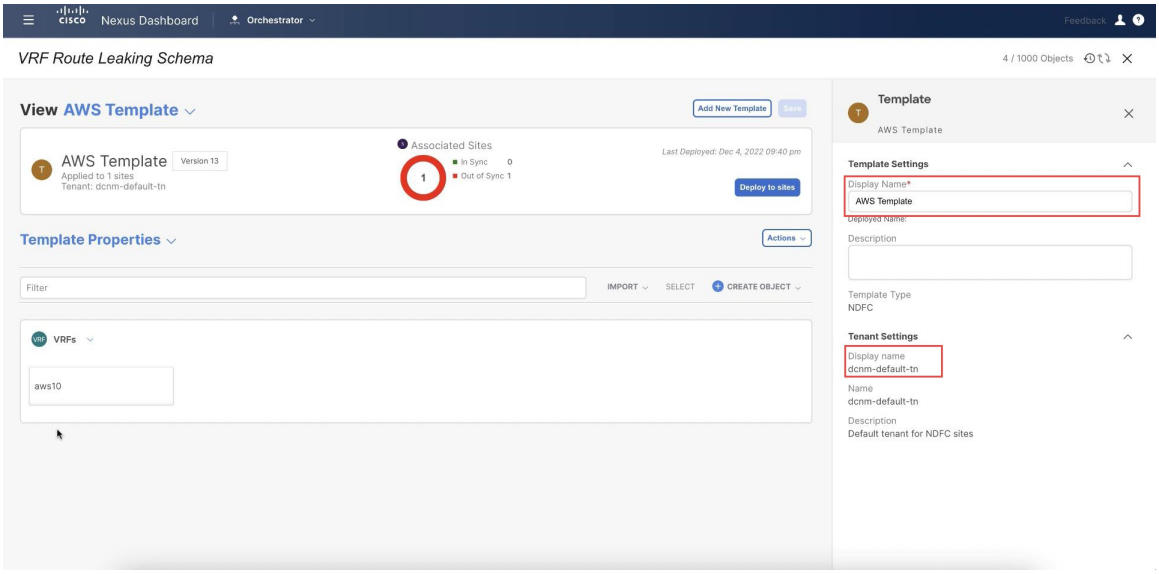

Step 5 Under **Template Properties**, click **Create Object** and choose **VRF** to create a VRF that will be used with the AWS site.

Step 6 Enter a name in the **Display Name** field for this VRF (for example, `aws10`).

Figure 188:

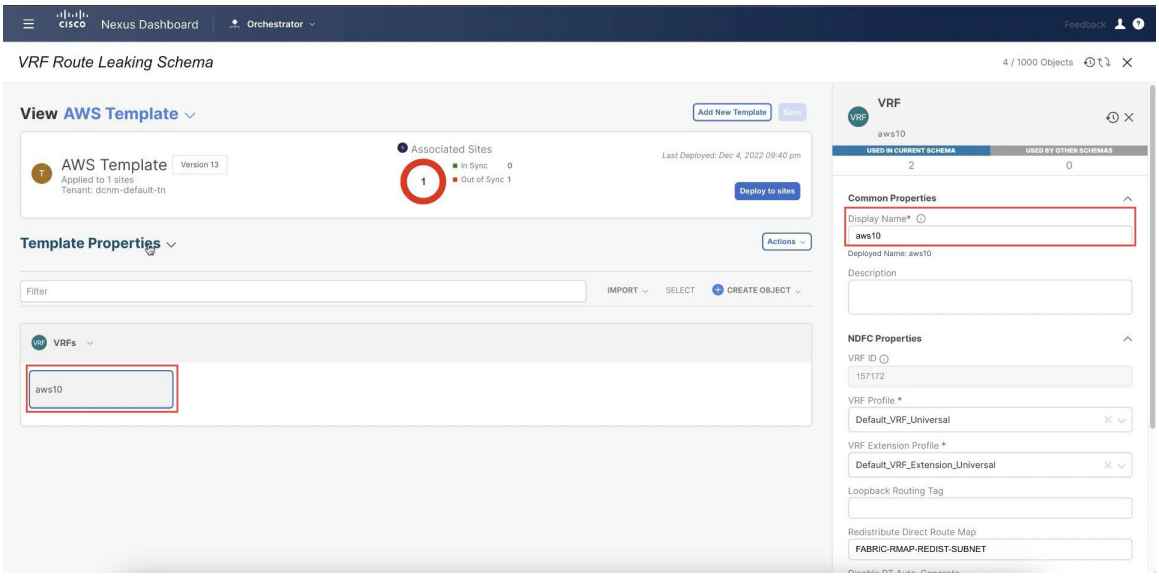

Step 7 In the **Template Properties** area, click **Actions** > **Sites Association**.

Step 8 Associate this template only to the AWS site, then click **Ok**.

Figure 189:

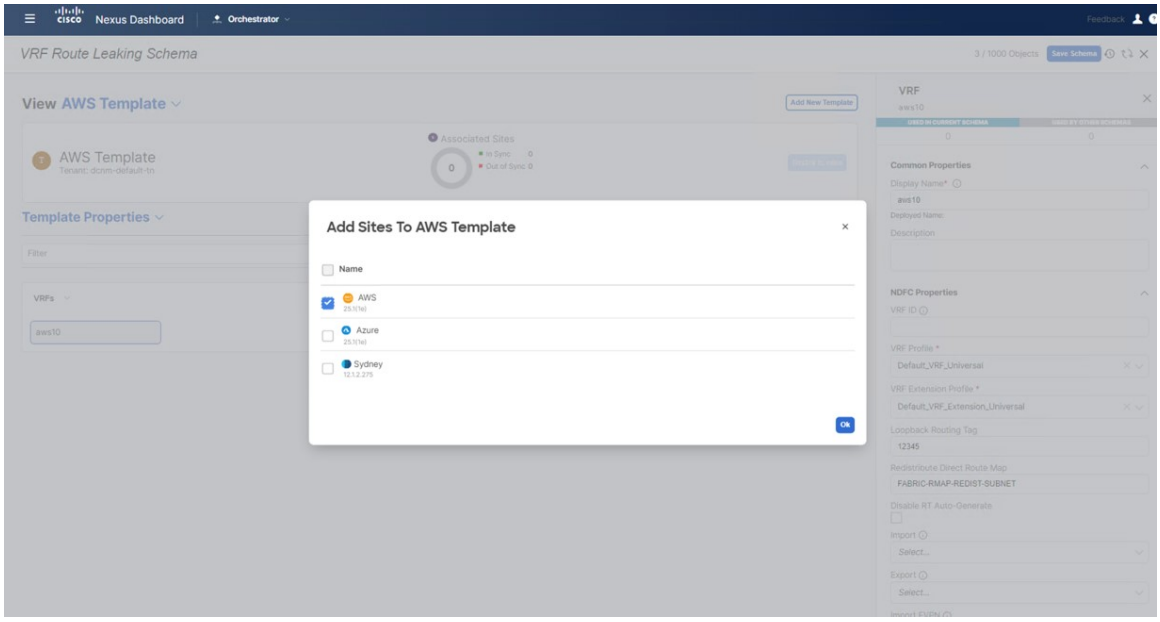

Step 9 Click the arrow next to **Template Properties**, and from the drop-down menu, select the AWS cloud site.

Step 10 Click the `aws10` VRF, then click **Add Region** to create the VPC in a selected region.
The **Add Cloud Region CIDRs** window appears.

Step 11 In the **Region** field, choose the region where you want to create the VPC.

Step 12 In the **CIDR** field, click **Add CIDRs** and define a CIDR block for the VPC.

Step 13 Click **Add Subnet** to create the subnets and map them to the availability zones, then click **Save**.

Figure 190:

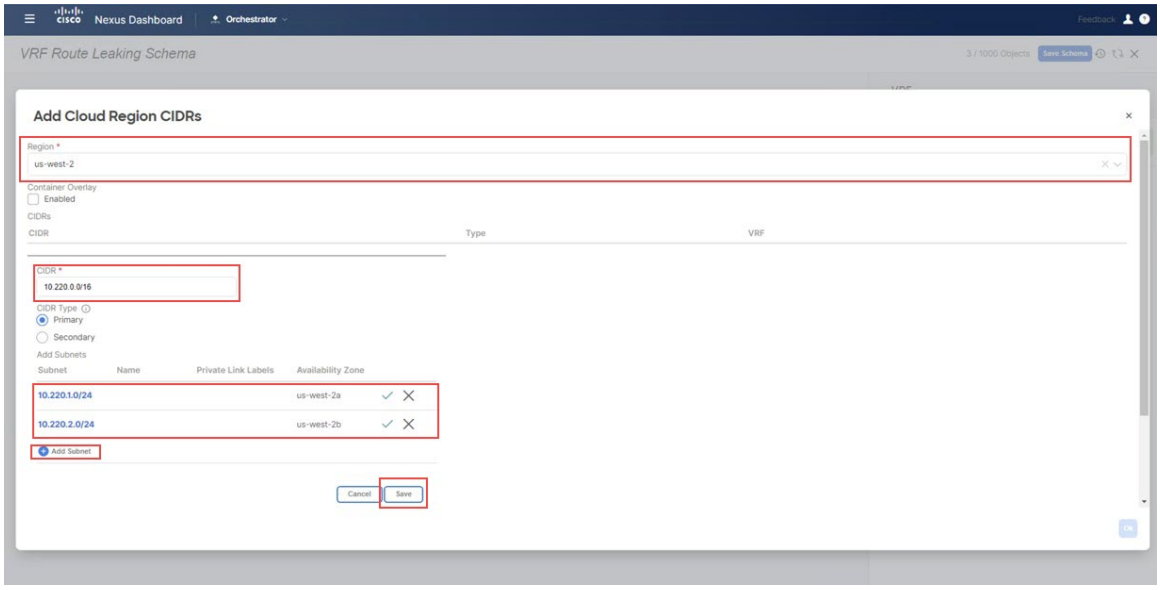

Step 14 Check the box under the **Hub Network** field, then select the hub network that was created on the Cisco Cloud Network Controller for AWS.

This allows the Cisco Cloud Network Controller to attach the subnets onto the transit gateway, which builds the connectivity from those subnets to the transit gateway, where the transit gateway already has the connectivity to the Cisco Catalyst 8000Vs in the cloud.

Step 15 In the **Subnets** field, map the subnets that will be used for the transit gateway.

It is best practice to have a dedicated subnet that will be used for the transit gateway.

Figure 191:

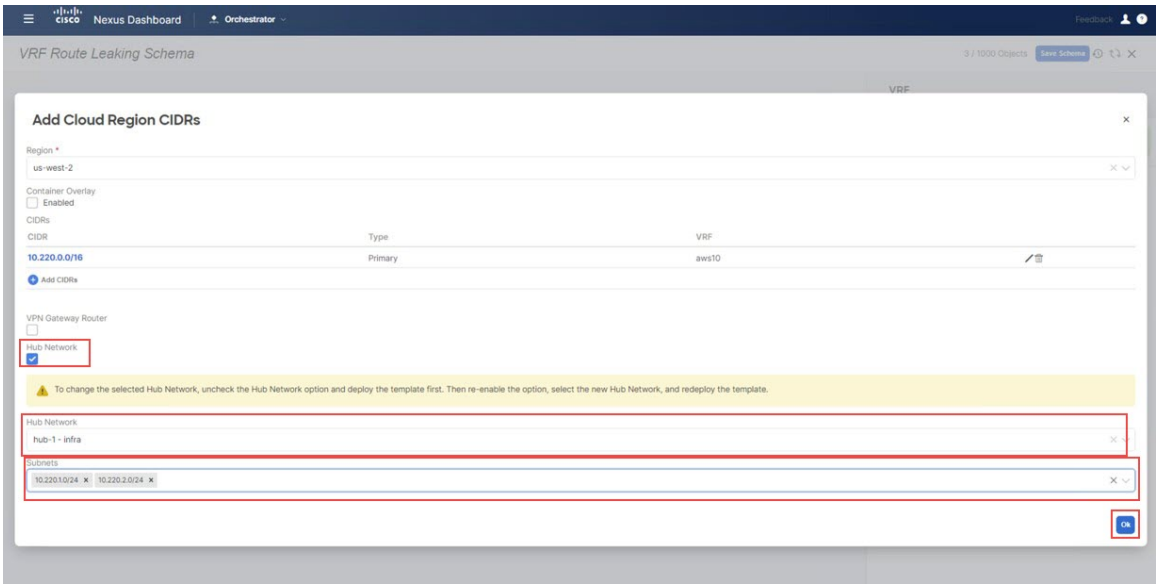

Step 16 Click **Ok**.
You are returned to the AWS template window.

Step 17 Click the arrow next to the AWS site, and from the drop-down menu, select **Template Properties**.

Step 18 Click **Deploy to Sites**.

Figure 192:

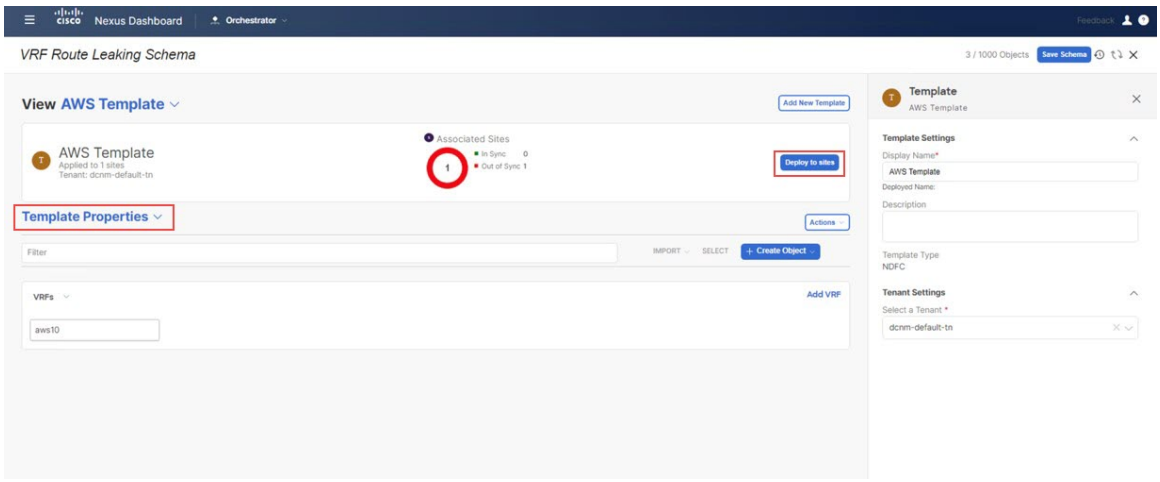

Step 19 Deploy `AWS Template` to the sites.

- Click **Deployment Plan** for additional verification.

  Click on the AWS site to see the deployment plan for that specific site.

  Figure 193:

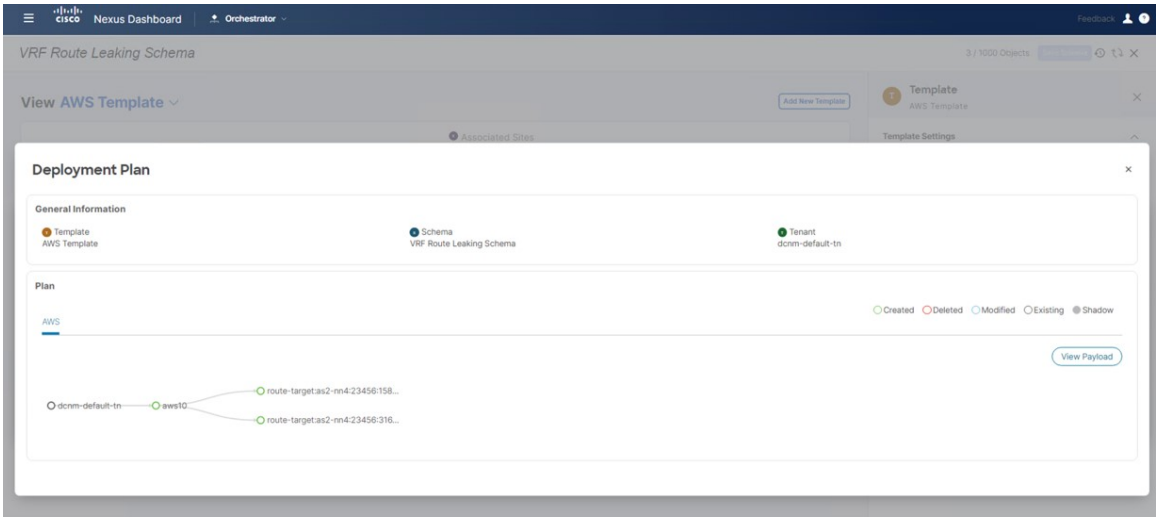

- Click **Deploy** to have NDO push the configurations to NDFC.

To verify that the configurations were pushed out correctly, connect to the Cloud Network Controller deployed on AWS and navigate to **Cloud Resources** > **VPCs**, then click the `aws10` VPC and use the information in the Overview page for additional verifications:

Figure 194:

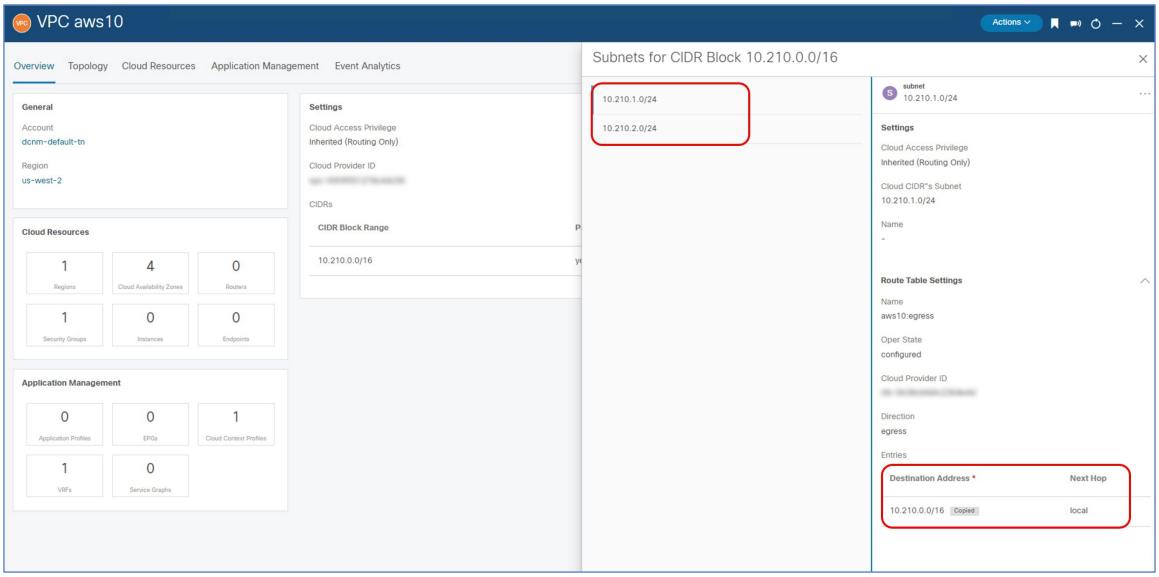

Note that there is a destination address configured at this point in the process for AWS, but this shows only that this AWS site can talk to itself; the AWS site cannot talk to any other site yet at this point in the process. The necessary destination address configuration that will allow the AWS site to talk to another site will be pushed out after you have completed the route leaking procedure.

What to do next

Configure route leaking using the procedures provided in Configure Route Leaking, on page 163.

Configure Route Leaking

Use the procedures in the following sections to configure the route leaking use case.

Configure Route Leak from Azure VRF to NDFC VRF

In this section, you will configure the route leak from the Azure VRF (`azure10`) to the NDFC VRF (`v10`).

Before you begin

Configure the necessary templates using the procedures provided in Configure the Necessary Templates, on page 145.

Step 1 Click the `Azure Template` that you configured earlier in these procedures and the dcnm-default-tn tenant.

Step 2 Click the `azure10` VRF that you configured earlier in these procedures.

Step 3 In the right pane, click **Add Leak Route**.

Figure 195:

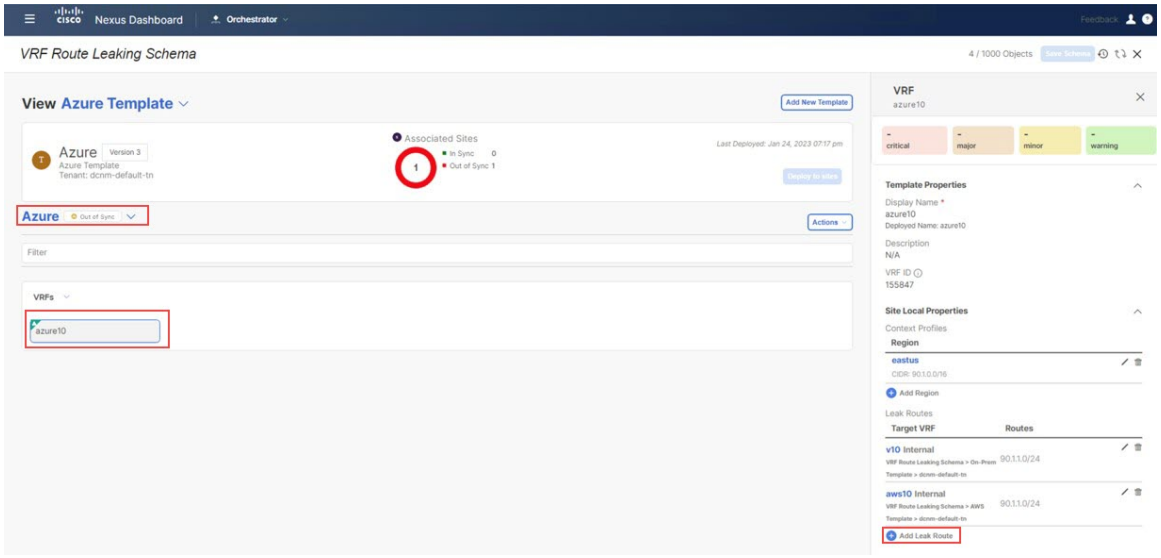

The **Add Leak Routes** window appears.

Step 4 In the **Add Leak Routes** window, click **Select a Target VRF**.

Figure 196:

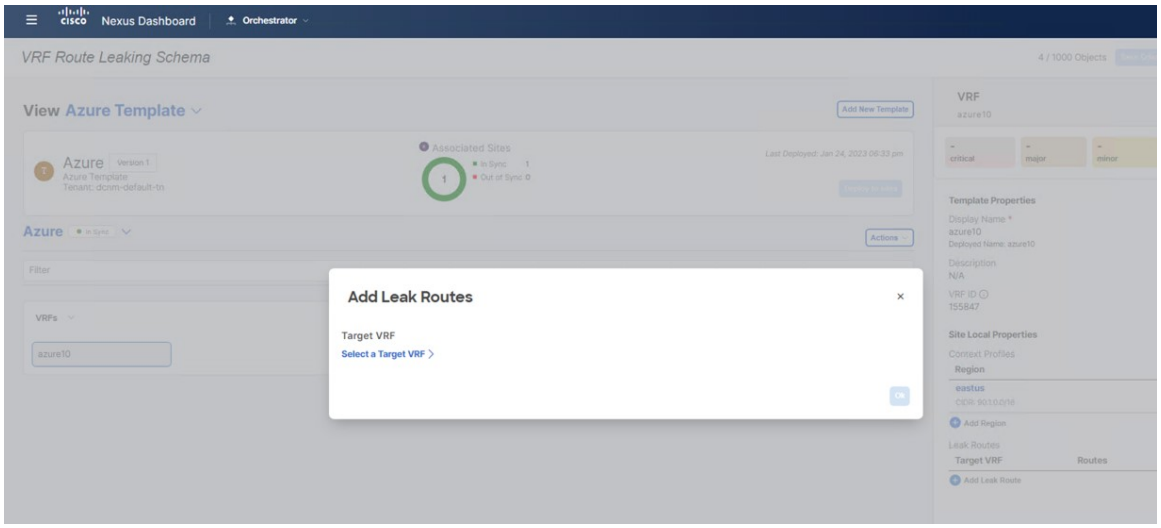

The **Select a Target VRF** window appears.

Step 5 In the **Select a Target VRF** page, select the NDFC VRF (`v10`) that you want to leak routes to, then click **Select**.

Figure 197:

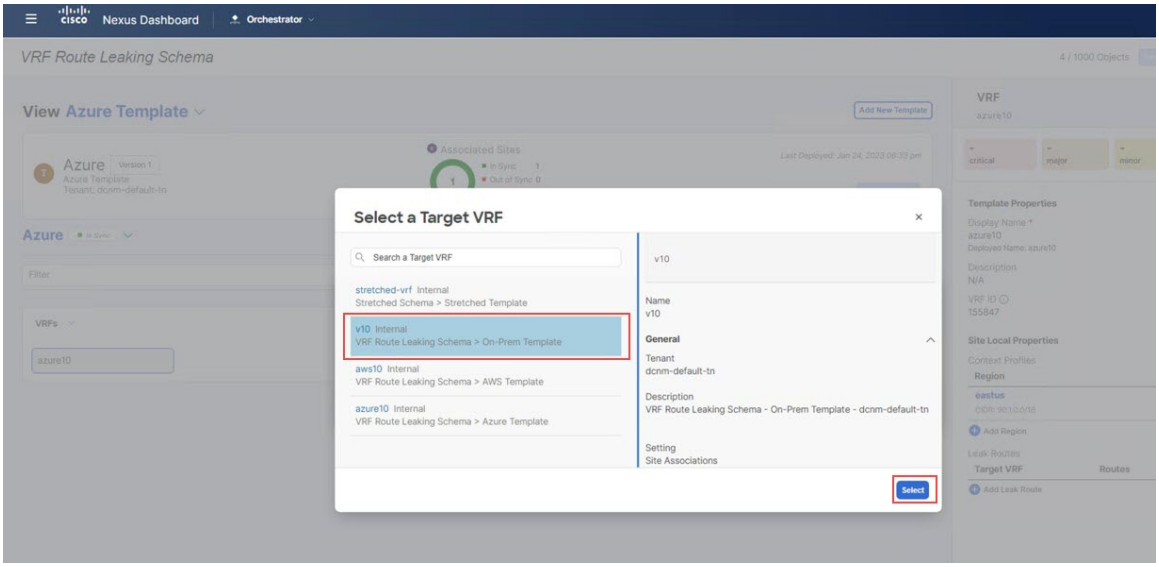

You are returned to the **Add Leak Routes** window.

Step 6 In the **Add Leak Routes** window, click **Add Subnet IP**, then add the Azure cloud subnets that you want to propagate to the on-premises site.

Note The **Add Subnet IP** option allows leaking of only selective subnets. Alternatively, you can use the **All Subnet IPs** option instead in the case where all the prefixes need to be leaked into a destination VRF.

Figure 198:

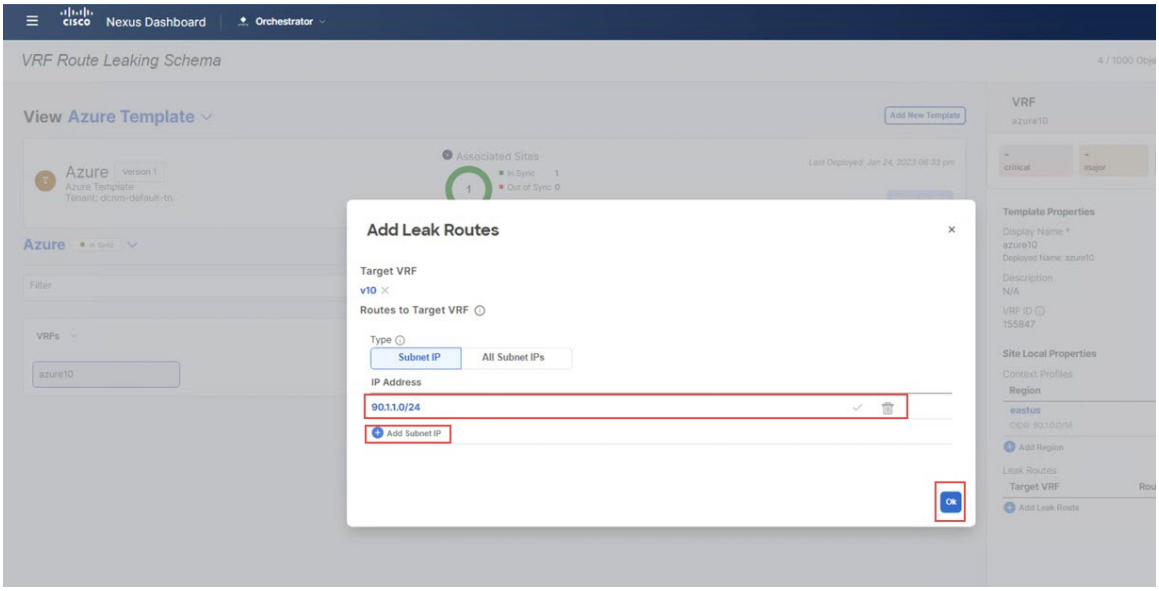

For this use case, you will use the `90.1.1.0/24` subnet.

Step 7 Click **Ok**.

You are returned to the `Azure Template` page, where you can see the configuration for this route leak from the Azure VRF to the NDFC VRF.

What to do next

Follow the procedures provided in Configure Route Leak from Azure VRF to AWS VRF, on page 166.

Configure Route Leak from Azure VRF to AWS VRF

In this section, you will configure the route leak from the Azure VRF (`azure10`) to the AWS VRF (`aws10`).

For these procedures, you will be going through the exact same procedures that you performed in Configure Route Leak from Azure VRF to NDFC VRF, on page 163, except in these procedures, you will be selecting a different target VRF (the AWS target VRF in these procedures).

Before you begin

Follow the procedures provided in Configure Route Leak from Azure VRF to NDFC VRF, on page 163.

Step 1 In the **Select a Target VRF** page, select the AWS VRF (`aws10`) that you want to leak routes to, then click **Select**.

Figure 199:

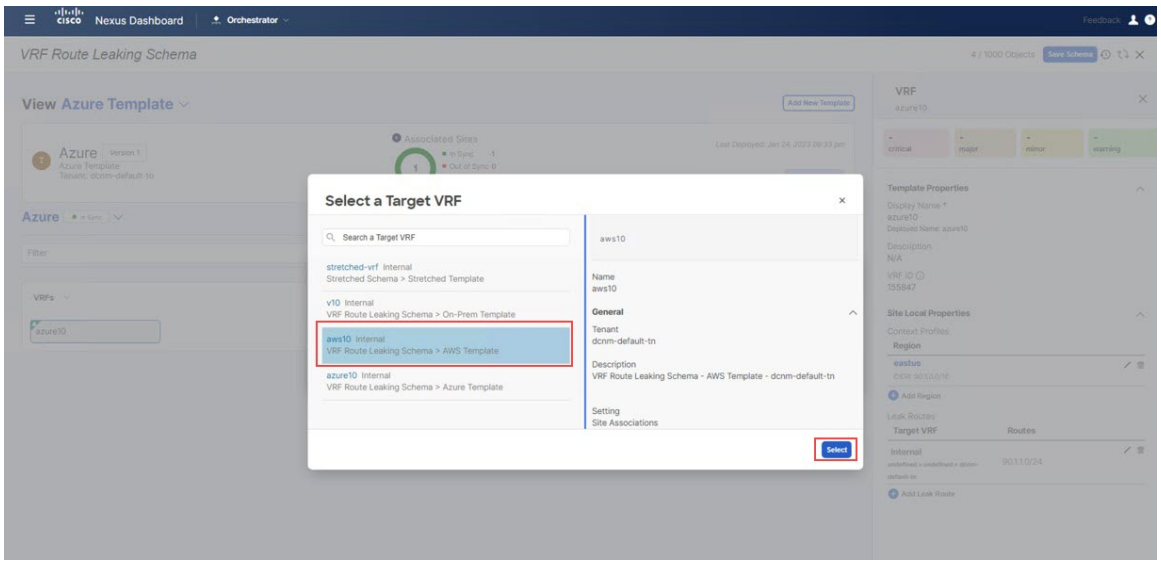

You are returned to the **Add Leak Routes** window.

Step 2 In the **Add Leak Routes** window, add the subnets that you want to propagate to the AWS cloud.

For this use case, you will use the `90.1.1.0/24` subnet. Therefore, you will click the dropdown menu and choose the `90.1.1.0/24` subnet.

Figure 200:

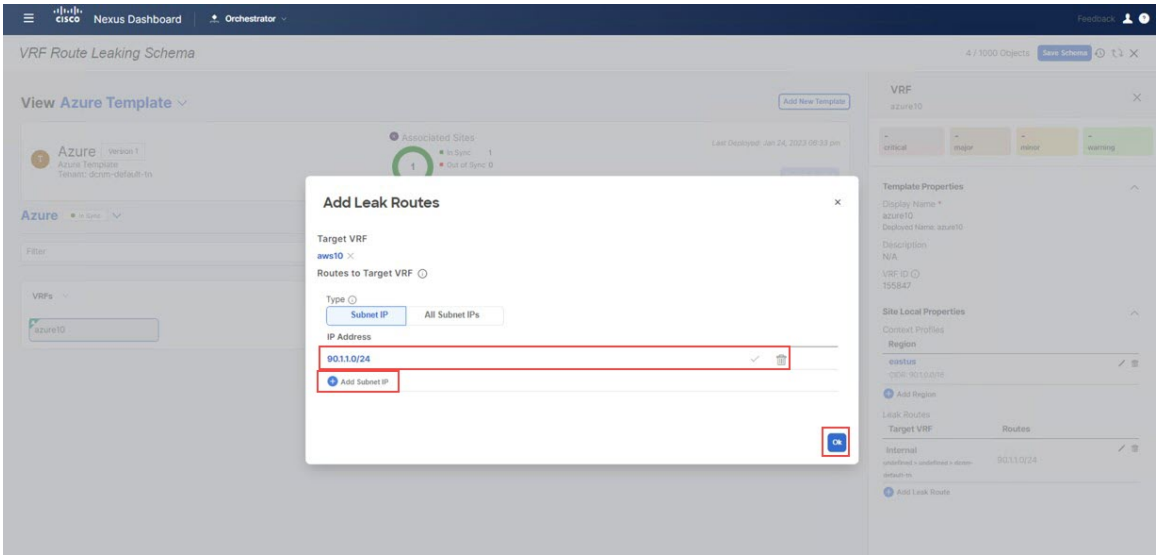

Step 3 Click **Ok**.

You are returned to the `Azure Template` page, where you can see the configuration for this route leak from the Azure VRF to the AWS VRF, as well as the route leak from the Azure VRF to the NDFC VRF that you configured in the previous set of steps.

Step 4 Click the arrow next to the Azure site, and from the drop-down menu, select **Template Properties**.

Step 5 Click **Deploy to sites**.

Figure 201:

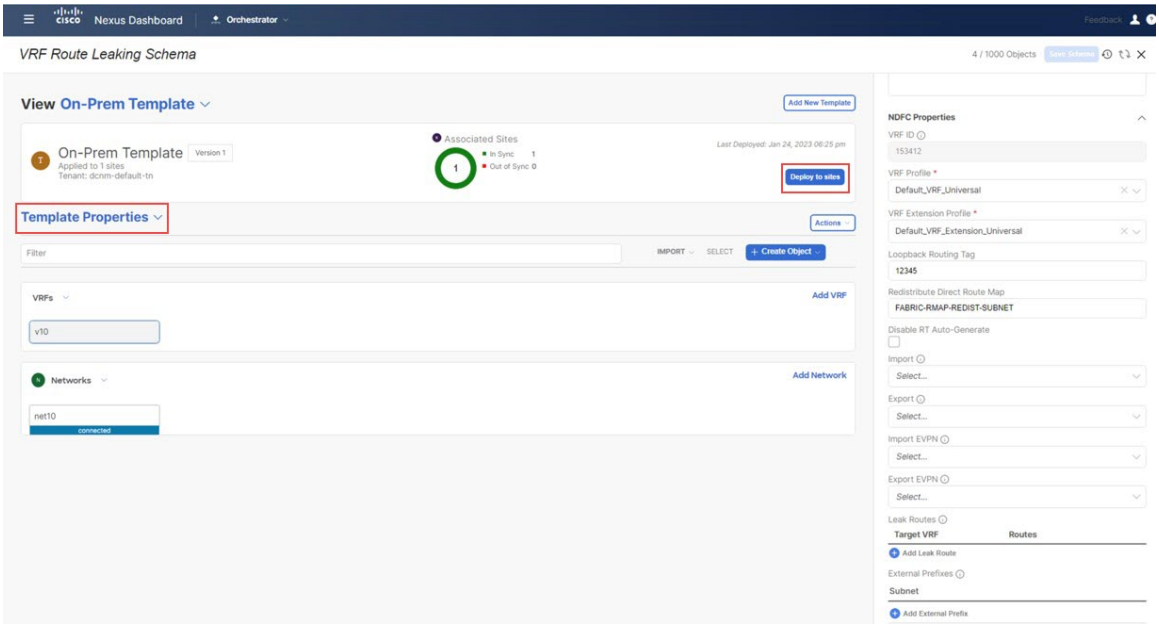

The **Deploy to sites** window appears, showing where the template will be deployed.

Step 6 Click **Deployment Plan** for additional verification, then click on a site to see the deployment plan for that specific site.

Step 7 Click **Deploy** to have NDO push the configurations to the site specific controllers.

Figure 202:

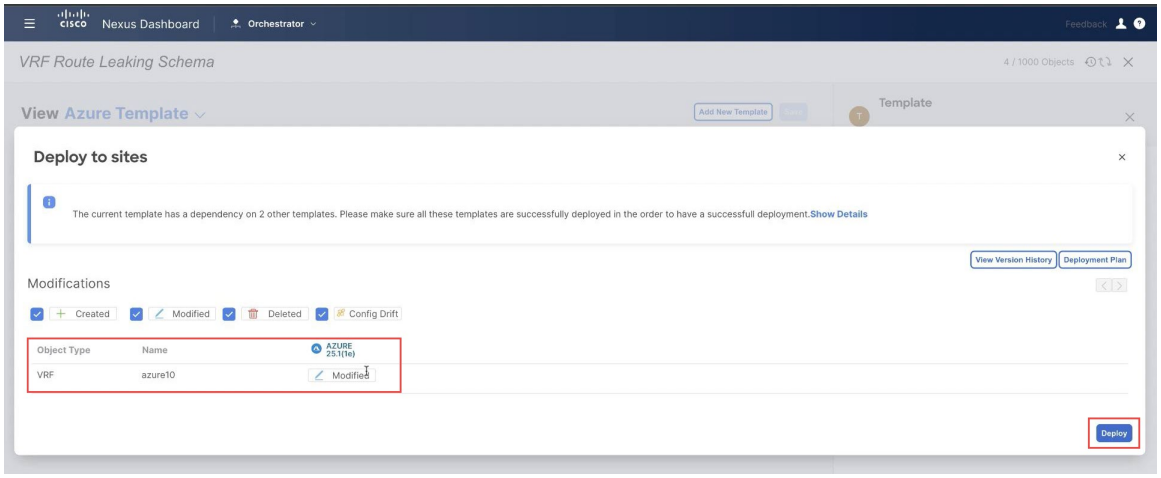

What to do next

Follow the procedures provided in Configure Route Leak from AWS VRF to NDFC VRF, on page 168.

Configure Route Leak from AWS VRF to NDFC VRF

In this section, you will configure the route leak from the AWS VRF (`aws10`) to the NDFC VRF (`v10`).

Before you begin

Follow the procedures provided in Configure Route Leak from Azure VRF to AWS VRF, on page 166.

Step 1 Click the `AWS Template` that you configured earlier in these procedures and the dcnm-default-tn tenant.

Step 2 Click the `aws10` VRF that you configured earlier in these procedures.

Step 3 In the right pane, click **Add Leak Route**.

Figure 203:

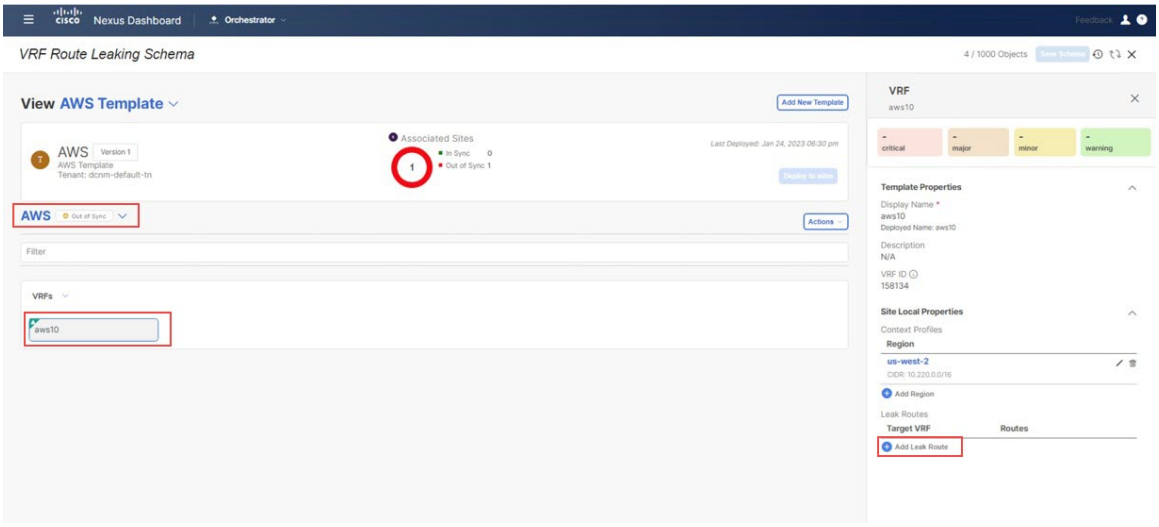

The **Add Leak Routes** window appears.

Step 4 In the **Add Leak Routes** window, click **Select a Target VRF**.
The **Select a Target VRF** window appears.

Step 5 In the **Select a Target VRF** window, select the NDFC VRF (`v10`) that you want to leak routes to, then click **Select**.
You are returned to the **Add Leak Routes** window.

Step 6 In the **Add Leak Routes** window, click **Add Subnet IP**, then add the AWS cloud subnets that you want to propagate to the on-premises site.

Note The **Add Subnet IP** option allows leaking of only selective subnets. Alternatively, you can use the **All Subnet IPs** option instead in the case where all the prefixes need to be leaked into a destination VRF.

Figure 204:

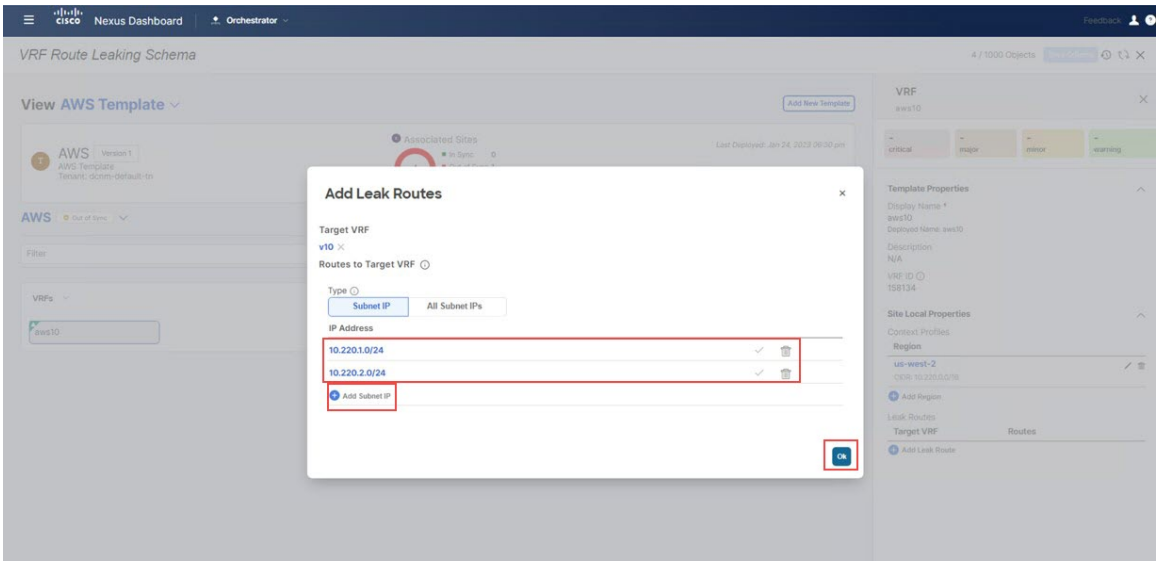

For this use case, you will use the following subnets:

- 10.220.1.0/24
- 10.220.2.0/24

Step 7 Click **Ok**.

You are returned to the `AWS Template` page, where you can see the configuration for this route leak from the AWS VRF to the NDFC VRF.

What to do next

Follow the procedures provided in Configure Route Leak from AWS VRF to Azure VRF, on page 170.

Configure Route Leak from AWS VRF to Azure VRF

In this section, you will configure the route leak from the AWS VRF (`aws10`) to the Azure VRF (`azure10`).

For these procedures, you will be going through the exact same procedures that you performed in Configure Route Leak from AWS VRF to NDFC VRF, on page 168, except in these procedures, you will be selecting a different target VRF (the Azure target VRF in these procedures).

Before you begin

Follow the procedures provided in Configure Route Leak from AWS VRF to NDFC VRF, on page 168.

Step 1 In the **Select a Target VRF** page, select the Azure VRF (`azure10`) that you want to leak routes to, then click **Select**. You are returned to the **Add Leak Routes** window.

Step 2 In the **Add Leak Routes** window, add the subnets that you want to propagate to the Azure cloud.

For this use case, you will use the following subnets:

- 10.220.1.0/24
- 10.220.2.0/24

Therefore, you will click the dropdown menu and choose those subnets.

Figure 205:

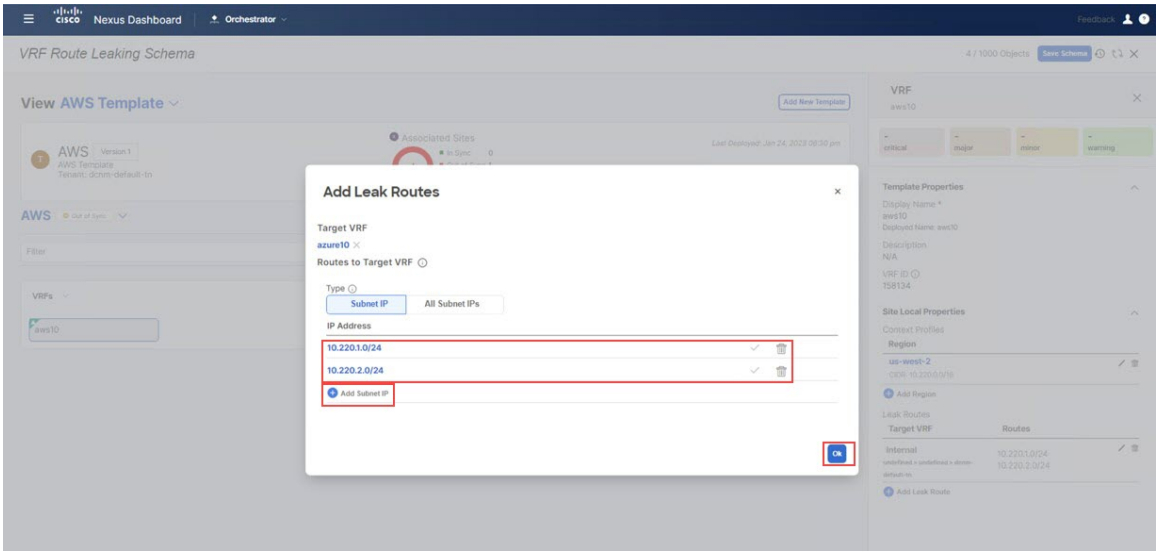

Step 3 Click **Ok**.

You are returned to the `AWS Template` page, where you can see the configuration for this route leak from the AWS VRF to the Azure VRF, as well as the route leak from the AWS VRF to the NDFC VRF that you configured in the previous set of steps.

Figure 206:

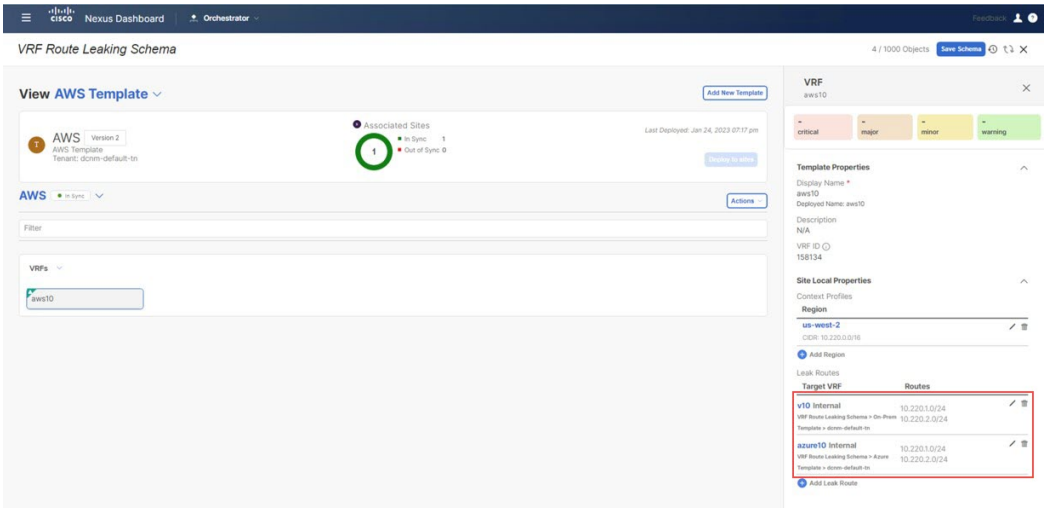

Step 4 Click the arrow next to the AWS site, and from the drop-down menu, select **Template Properties**.

Step 5 Click **Deploy to sites**.

The **Deploy to sites** window appears, showing where the template will be deployed.

Step 6 Click **Deployment Plan** for additional verification, then click on a site to see the deployment plan for that specific site.

Step 7 Click **Deploy** to have NDO push the configurations to the site specific controllers (NDFC and Cloud Network Controller).

What to do next

Follow the procedures provided in Configure Route Leak from NDFC VRF to AWS VRF, on page 172.

Configure Route Leak from NDFC VRF to AWS VRF

In this section, you will configure the route leak from the NDFC VRF (`v10`) to the AWS VRF (`aws10`).

Before you begin

Follow the procedures provided in Configure Route Leak from AWS VRF to Azure VRF, on page 170.

Step 1 Click the `On-Prem Template` that you configured earlier in these procedures and the dcnm-default-tn tenant.

Step 2 Click the `v10` VRF that you configured earlier in these procedures.

Step 3 In the right pane, click **Add Leak Route**.

Figure 207:

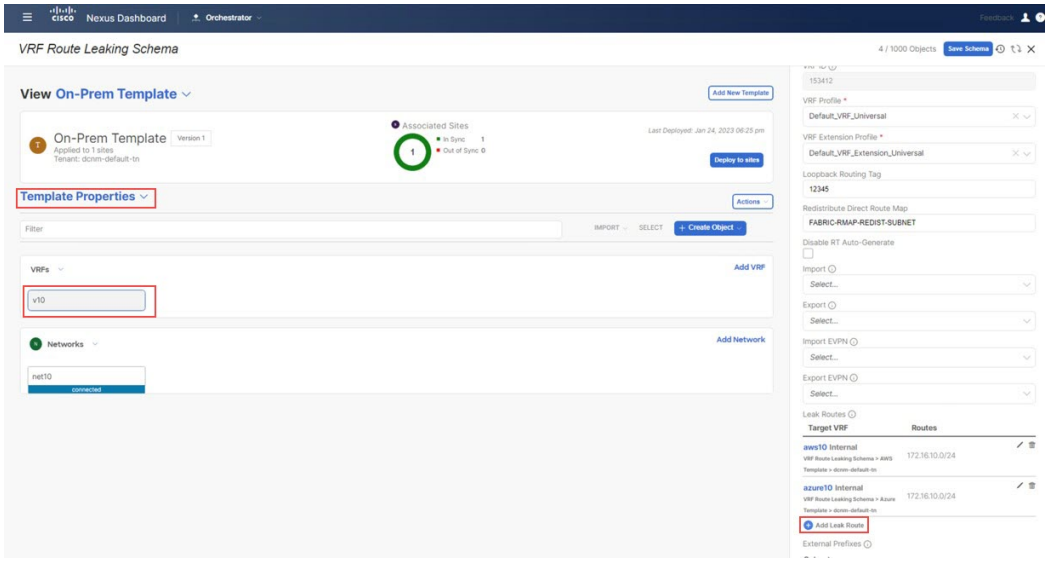

The **Add Leak Routes** window appears.

Step 4 In the **Add Leak Routes** window, click **Select a Target VRF**.
The **Select a Target VRF** window appears.

Step 5 In the **Select a Target VRF** window, select the AWS cloud site VRF (`aws10`) that you want to leak routes to, then click **Select**.
You are returned to the **Add Leak Routes** window.

Step 6 In the **Add Leak Routes** window, click **Add Subnet IP**, then add the AWS cloud subnets that you want to propagate to the on-premises site.

Note The **Add Subnet IP** option allows leaking of only selective subnets. Alternatively, you can use the **All Subnet IPs** option instead in the case where all the prefixes need to be leaked into a destination VRF.

For this use case, you will use the `172.16.10.0/24` subnet.

Figure 208:

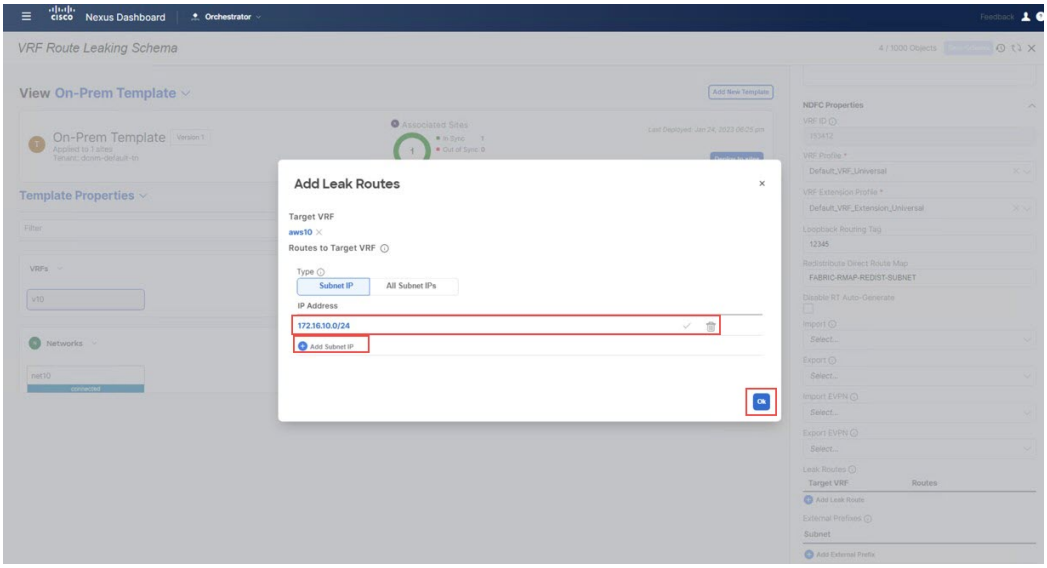

Step 7 Click **Ok**.

You are returned to the `On-Prem Template` page, where you can see the configuration for this route leak from the NDFC VRF to the AWS VRF.

What to do next

Follow the procedures provided in Configure Route Leak from NDFC VRF to Azure VRF, on page 173.

Configure Route Leak from NDFC VRF to Azure VRF

In this section, you will configure the route leak from the NDFC VRF (`v10`) to the Azure VRF (`azure10`).

For these procedures, you will be going through the exact same procedures that you performed in Configure Route Leak from NDFC VRF to AWS VRF, on page 172, except in these procedures, you will be selecting a different target VRF (the Azure target VRF in these procedures).

Before you begin

Follow the procedures provided in Configure Route Leak from NDFC VRF to AWS VRF, on page 172.

Step 1 In the **Select a Target VRF** window, select the Azure VRF (`azure10`) that you want to leak routes to, then click **Select**.

You are returned to the **Add Leak Routes** window.

Step 2 In the **Add Leak Routes** window, add the subnets that you want to propagate to the Azure cloud.

For this use case, you will use the `172.16.10.0/24` subnet. Therefore, you will click the dropdown menu and choose the `172.16.10.0/24` subnet.

Figure 209:

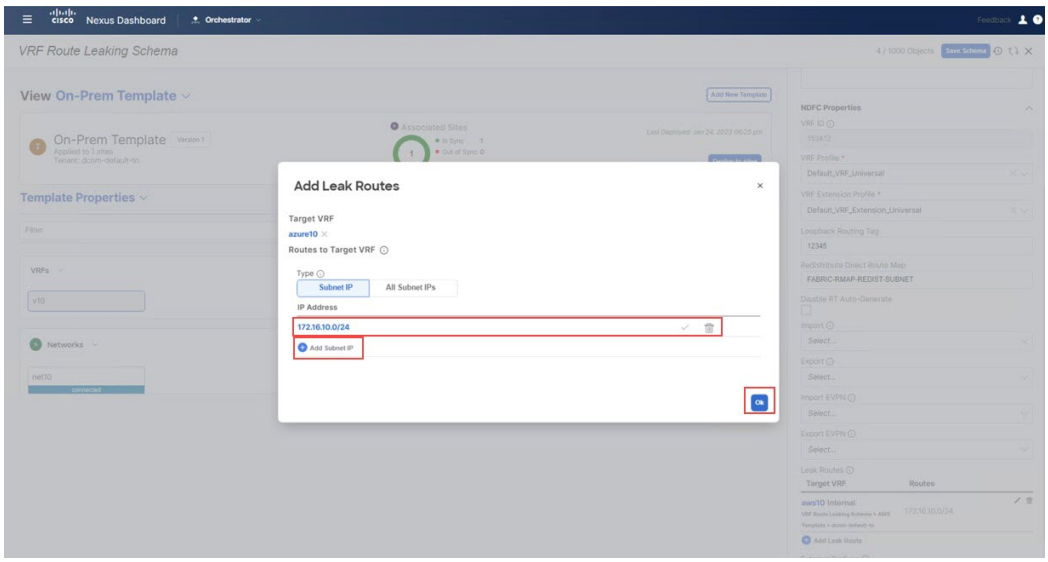

Step 3 Click **Ok**.

You are returned to the `On-Prem Template` page, where you can see the configuration for this route leak from the NDFC VRF to the Azure VRF, as well as the route leak from the NDFC VRF to the AWS VRF that you configured in the previous set of steps.

Step 4 Click the arrow next to the on-premises site, and from the drop-down menu, select **Template Properties**.

Step 5 Click **Deploy to sites**.

Figure 210:

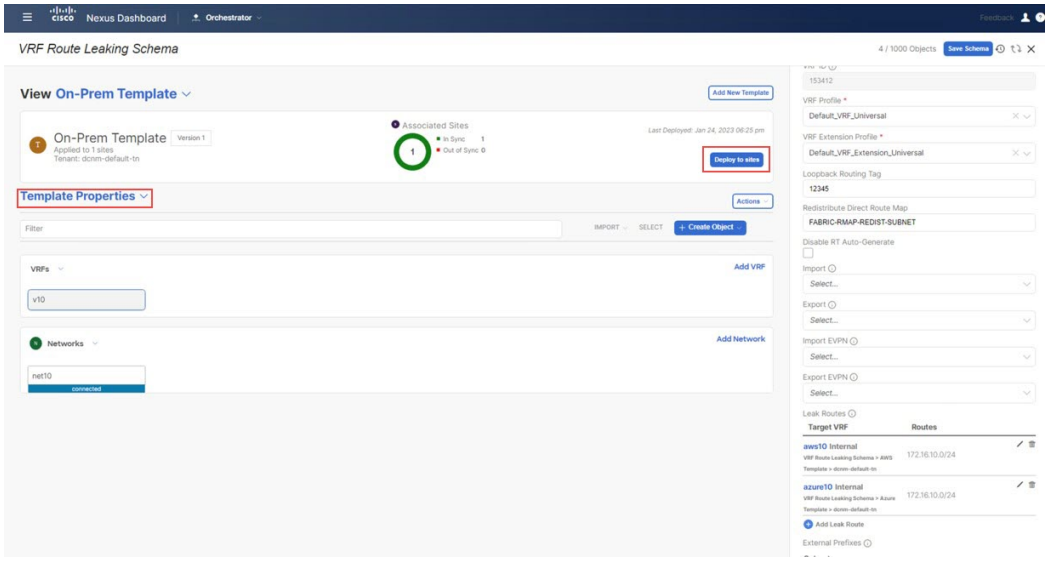

The **Deploy to sites** window appears, showing where the template will be deployed.

Step 6 Click **Deployment Plan** for additional verification, then click on a site to see the deployment plan for that specific site.

Figure 211:

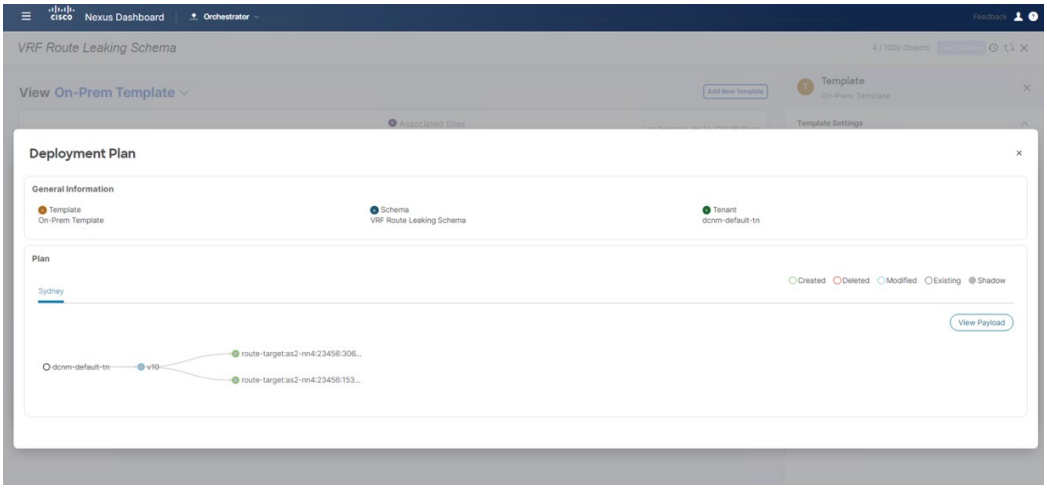

Step 7 Click **Deploy** to have NDO push the configurations to the site specific controllers (NDFC and Cloud Network Controller).

What to do next

Verify that the configurations were deployed successfully using the procedures provided in Verify the Configurations, on page 175.

Verify the Configurations

In this section, you will verify that the configurations were deployed successfully. Note that for each of these verification steps, the exact command that would be used specifically for the configurations in this use case are shown. Replace the appropriate variables in each command based on your configuration.

Before you begin

Follow the procedures provided in Configure Route Leak from NDFC VRF to Azure VRF, on page 173.

Step 1 Verify the configurations in NDO.

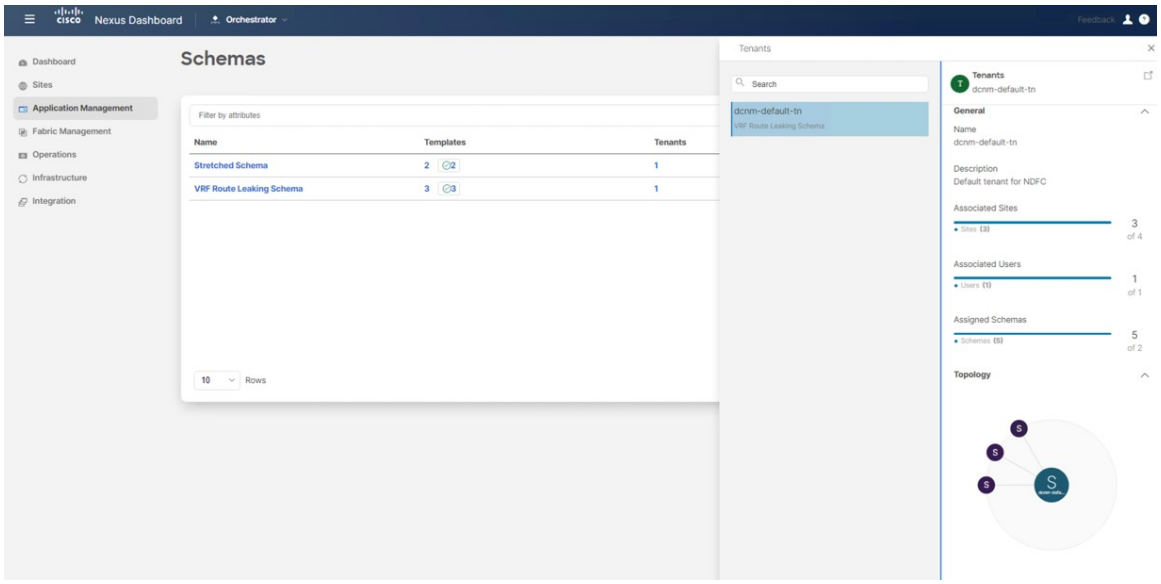
Nexus Dashboard
Orchestrator
Schemas
Dashboard
Sites
Application Management
Fabric Management
Operations
Infrastructure
Integration
Filter by attributes
Name
Templates
Tenants
Stretched Schema
2
2
1
VRF Route Leaking Schema
3
3
1
10
Rows
Tenants
Search
dcnm-default-tn
VRF Route Leaking Schema
Tenants
dcnm-default-tn
General
Name
dcnm-default-tn
Description
Default tenant for NDFC
Associated Sites
3
of 4
Associated Users
1
of 1
Assigned Schemas
5
of 2
Topology

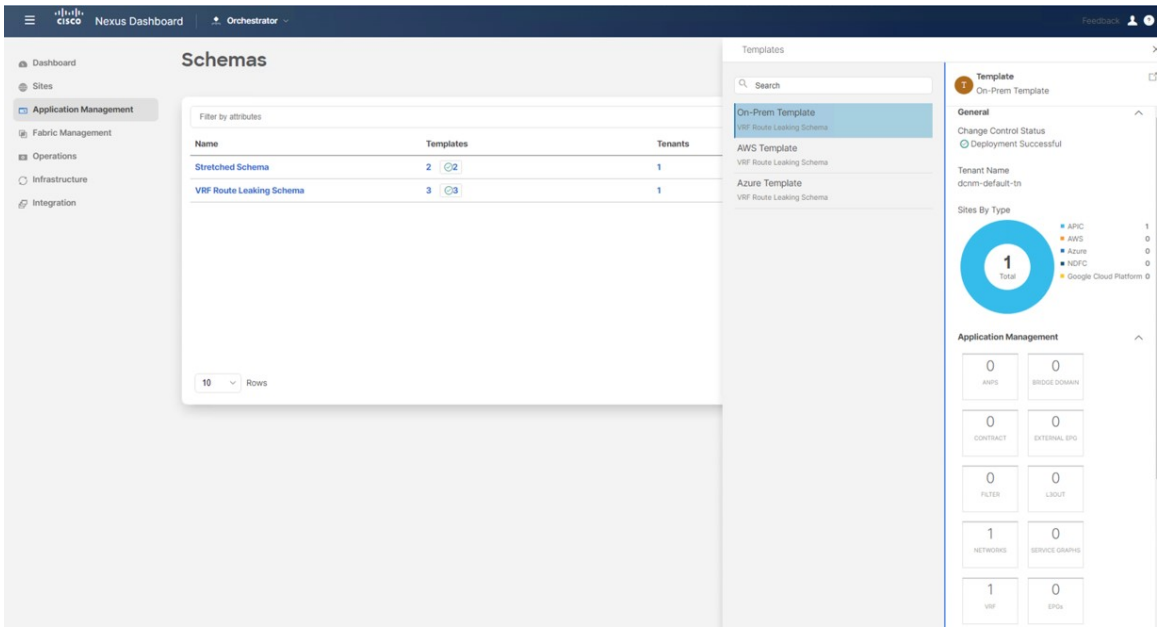
Nexus Dashboard
Orchestrator
Schemas
Dashboard
Sites
Application Management
Fabric Management
Operations
Infrastructure
Integration
Filter by attributes
Name
Templates
Tenants
Stretched Schema
2
2
1
VRF Route Leaking Schema
3
3
1
10
Rows
Templates
Search
On-Prem Template
VRF Route Leaking Schema
AWS Template
VRF Route Leaking Schema
Azure Template
VRF Route Leaking Schema
Template
On-Prem Template
General
Change Control Status
Deployment Successful
Tenant Name
dcnm-default-tn
Sites By Type
1
Total
APIC
AWS
Azure
NDFC
Google Cloud Platform
Application Management

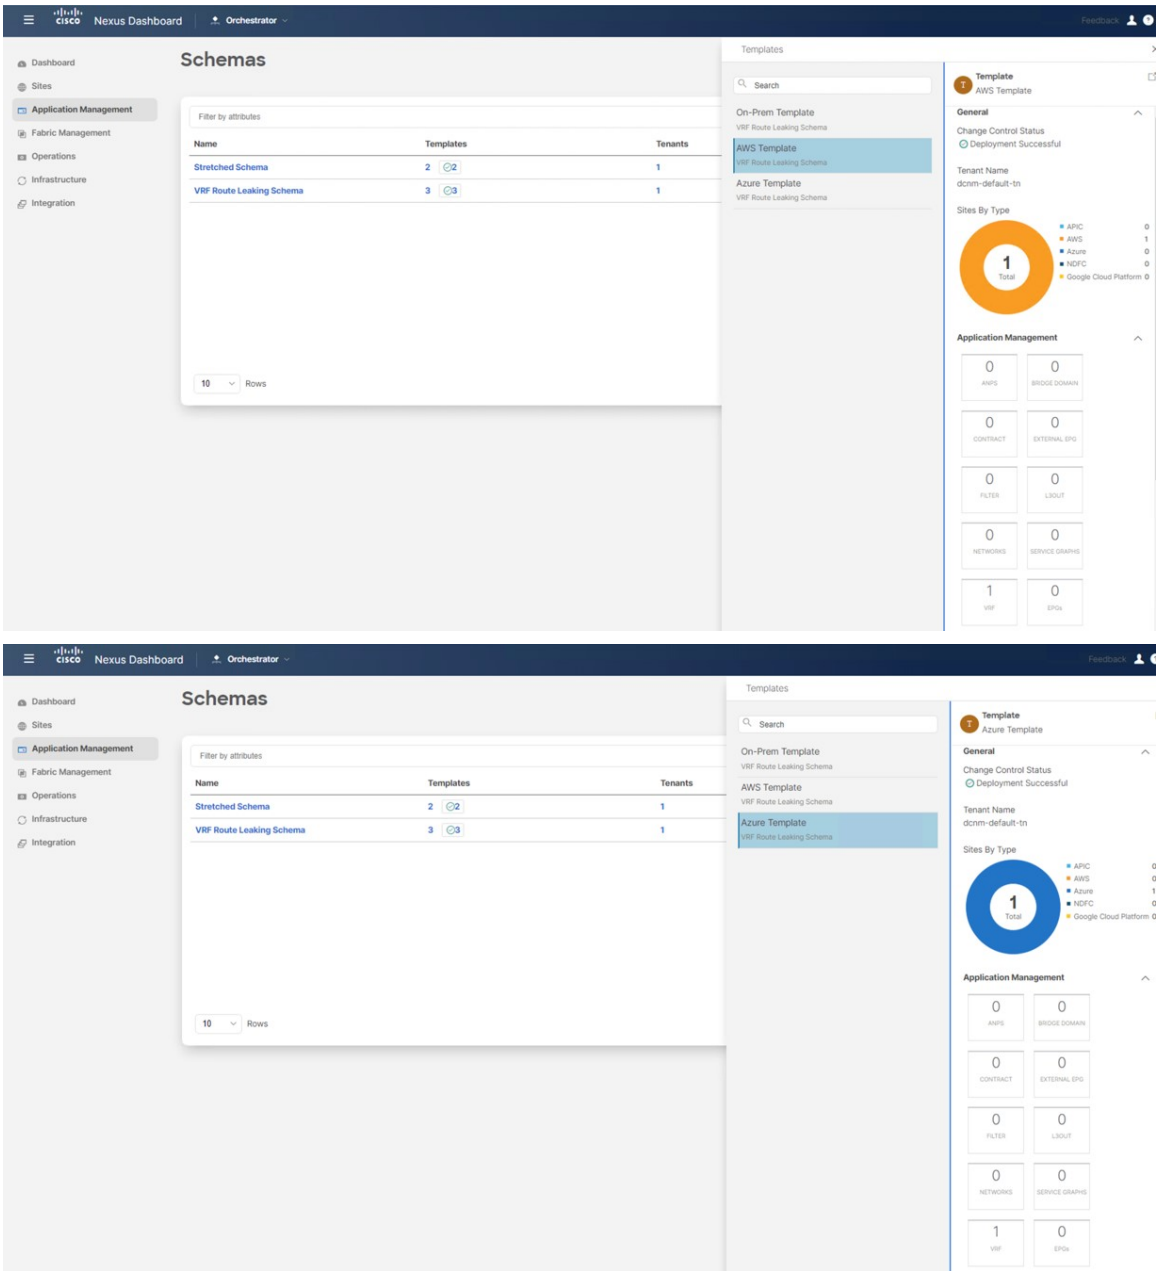

Step 2 Enter **sh ip route vrf v10** on the on-premises Border Gateway Spine device:

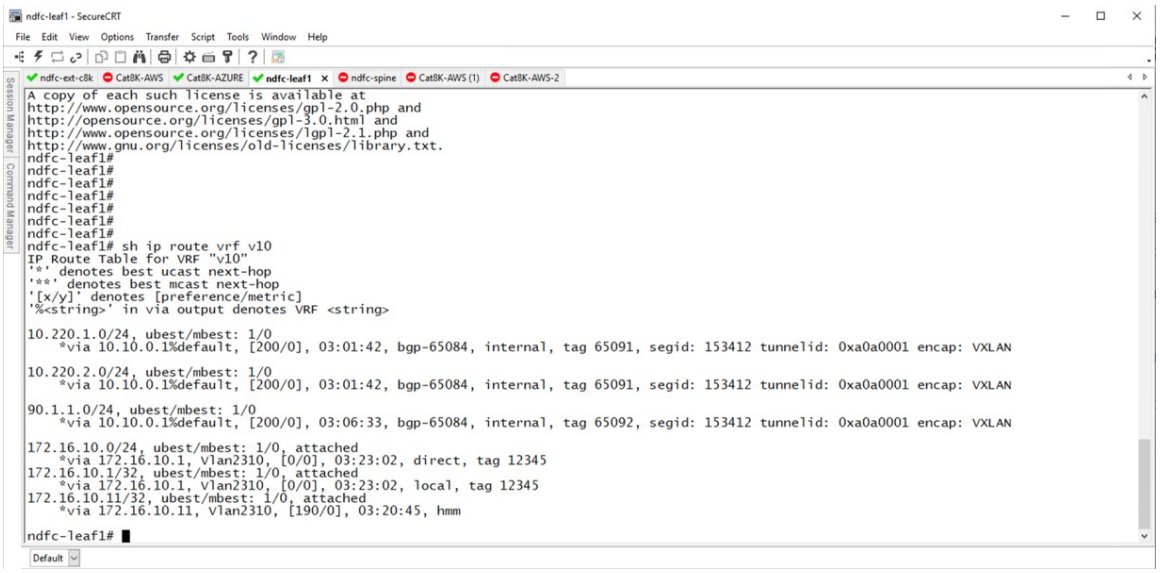

The routing table on the on-premises leaf switch shows that the reachable subnets are:

- **AWS**: 10.220.0.0/16
- **Azure**: 10.220.0.0/16

Step 3 Connect to the Cloud Network Controller deployed on AWS and navigate to **Application Management** > **VRFs**, and verify that you can see the Azure and NDFC VRFs.

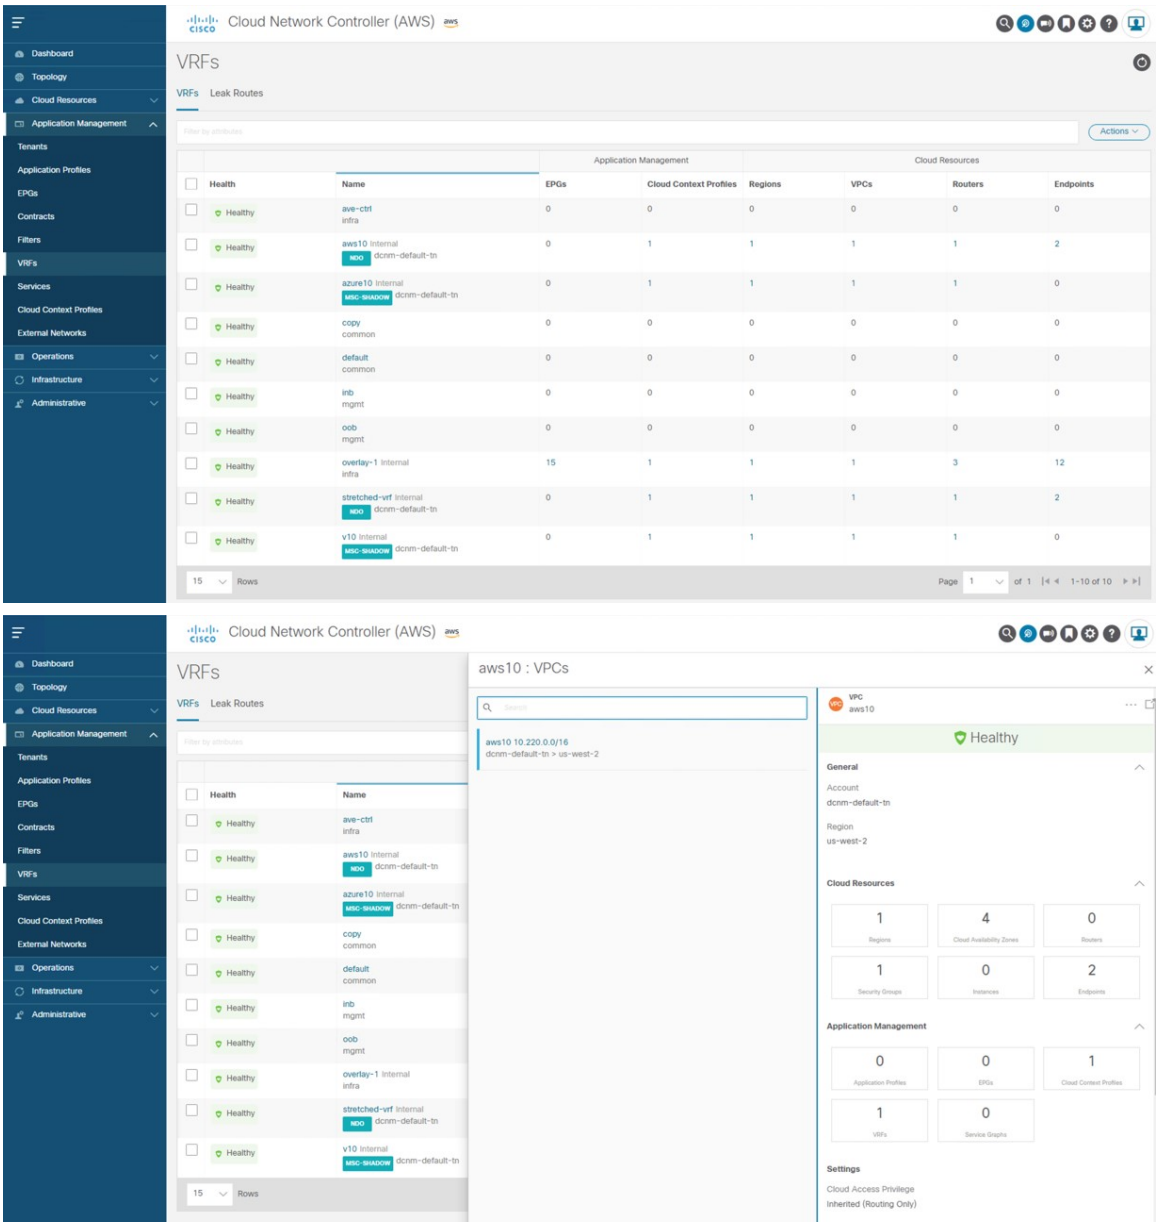

Step 4 Remaining in the Cloud Network Controller deployed on AWS, perform a verification on the route table view.

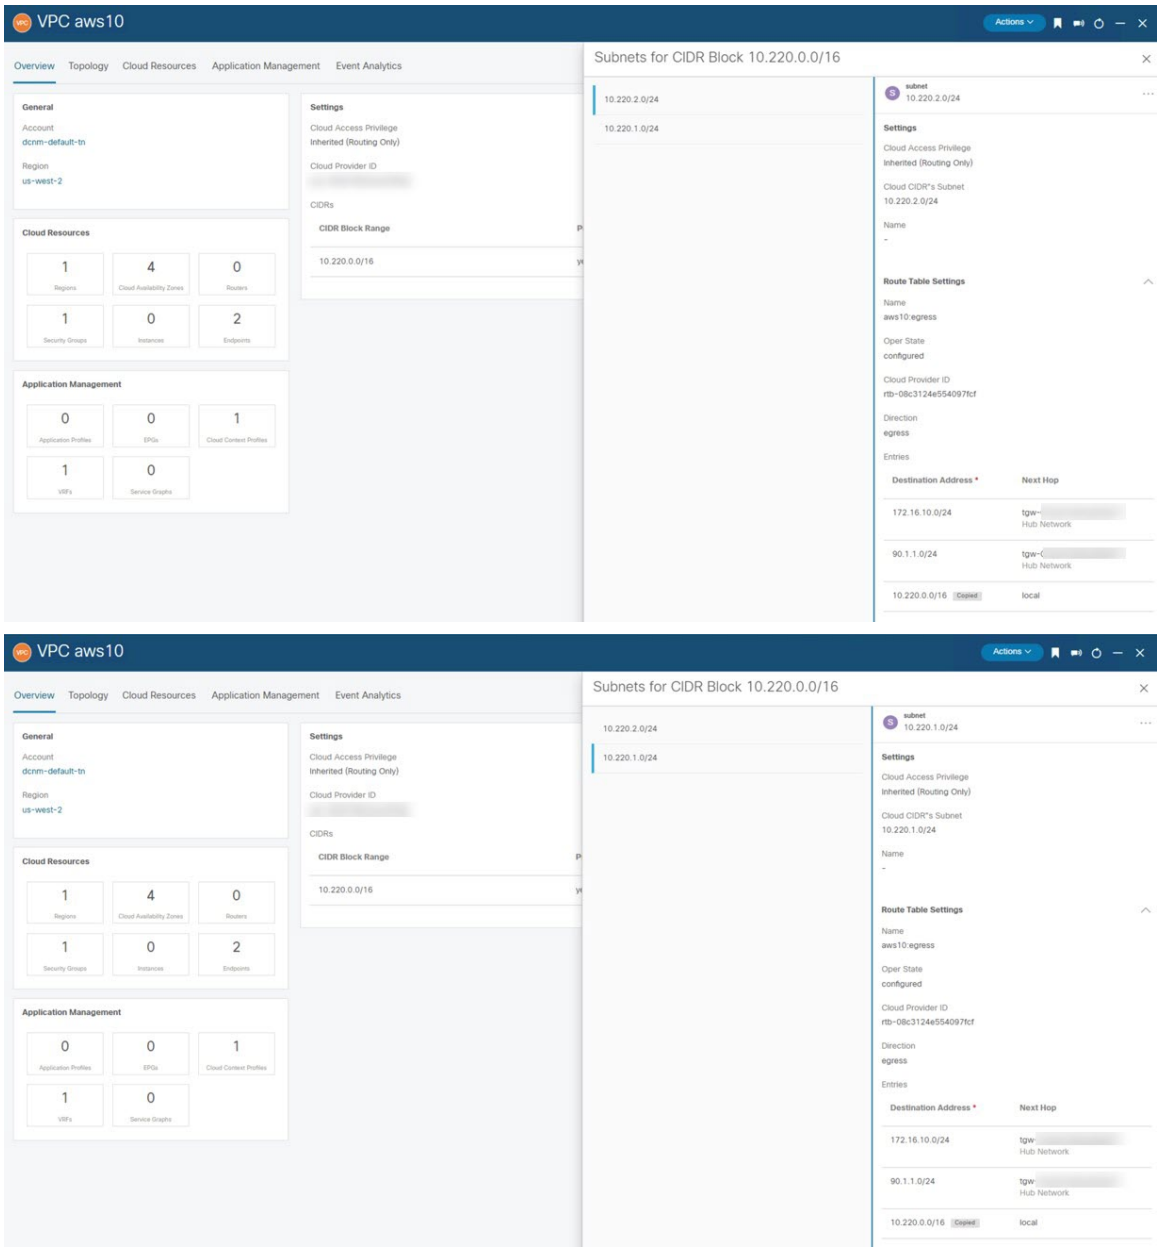

Step 5 In the AWS console, perform a verification on the route table view.

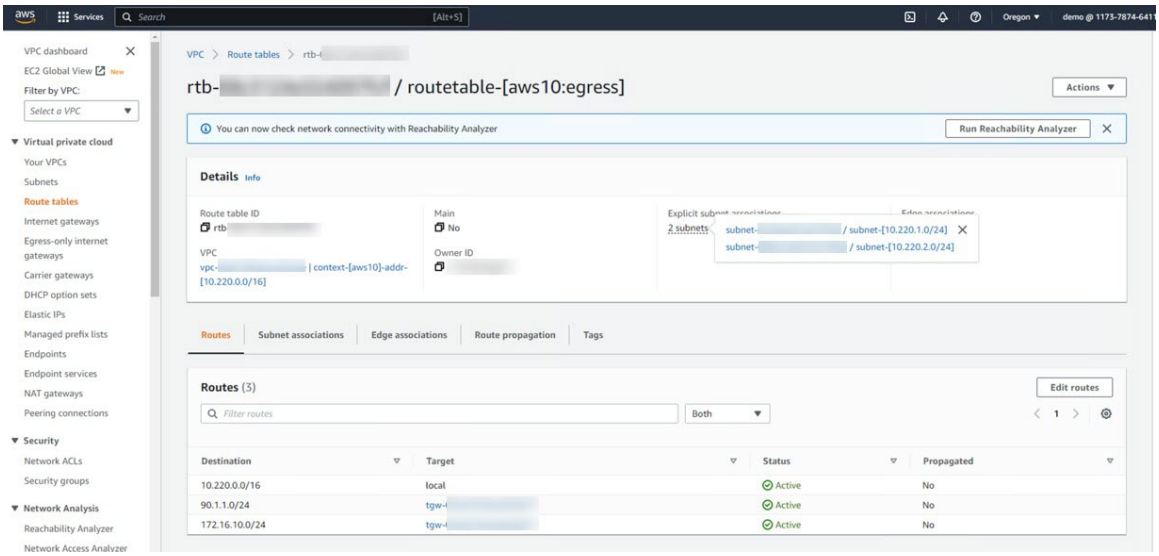

Step 6 Connect to the Cloud Network Controller deployed on Azure and navigate to **Application Management** > **VRFs**, and confirm that you can see the AWS and NDFC VRFs:

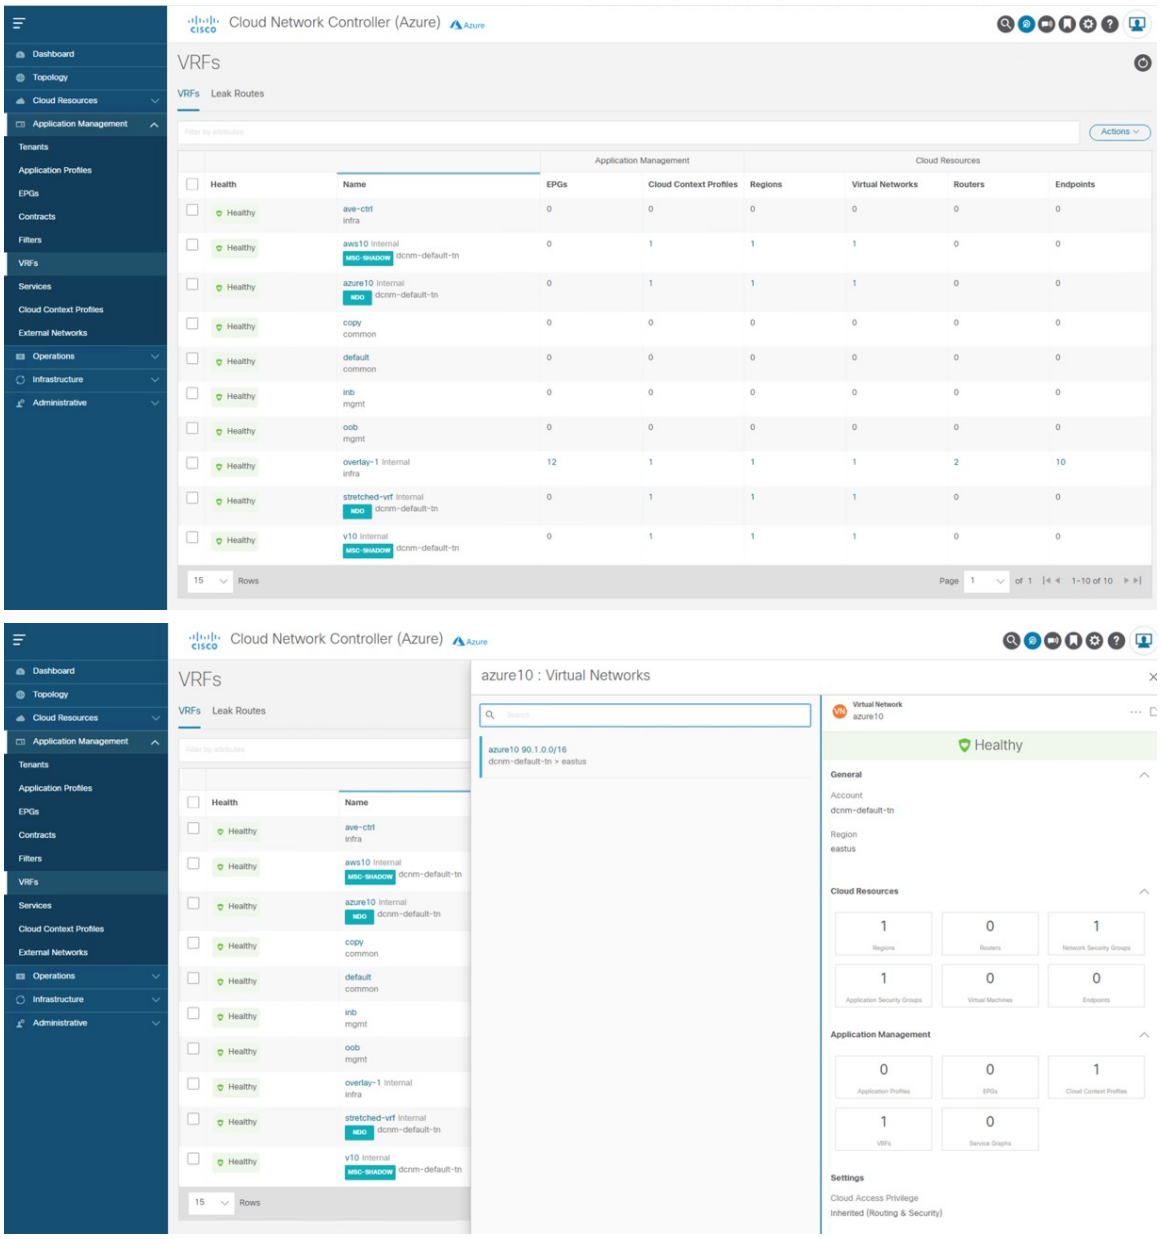

Step 7 Remaining in the Cloud Network Controller deployed on Azure, navigate to **Cloud Resources** > **Virtual Networks**, then click the `azure10` VNet and use the information in the Overview page for additional verifications.

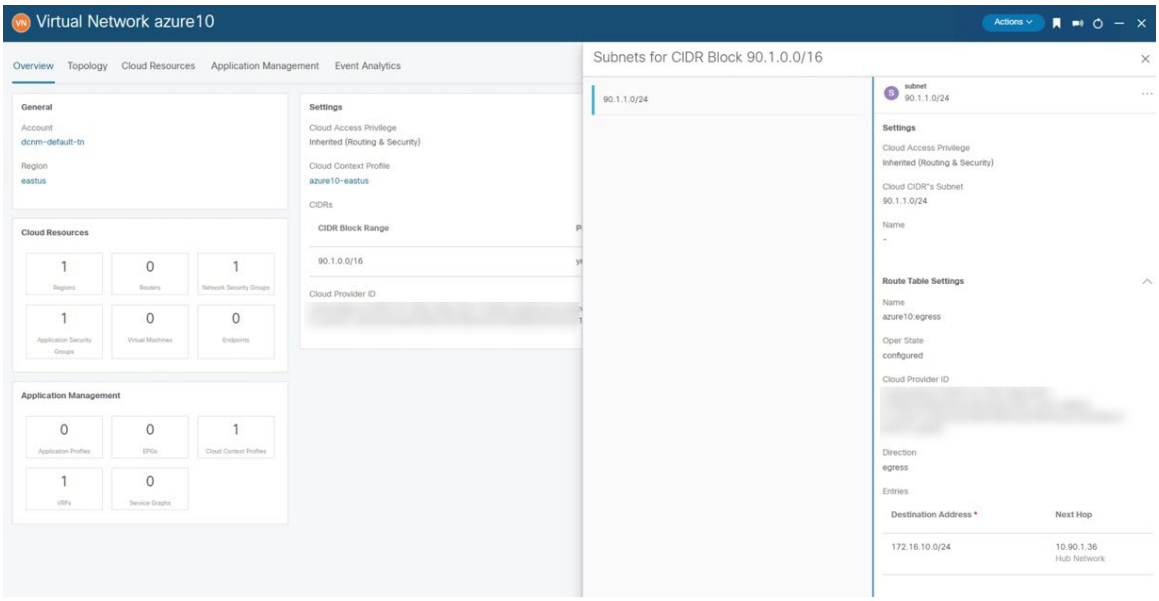

Step 8 In the Azure console, perform additional verifications.