Drive Consistent Policy and Operations Across ACI Sites with Cisco Nexus Dashboard

Publish Date: December 12, 2025

This guide for the preconfigured demonstration includes:

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

What's Next?

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

About This Demo

Cisco Nexus Dashboard Orchestrator provides consistent network and policy orchestration, scalability, and disaster recovery across multiple data centers through a single pane of glass while allowing the data center to go wherever the data is.

Nexus Dashboard Orchestrator allows you to interconnect separate Cisco Application Centric Infrastructure (Cisco ACI) sites, Cisco Cloud ACI sites, and Cisco Nexus Dashboard Fabric Controller (NDFC) sites, each managed by its own controller (APIC cluster, NDFC cluster, or Cloud APIC instances in a public cloud). The on-premises sites (ACI or NDFC with ND 4.2) can be extended to different public clouds for hybrid-cloud deployments while cloud-first installations can be extended to multicloud deployments without on-premises sites. In addition, the Nexus Dashboard Orchestrator can be deployed through the Cisco Nexus Dashboard, which provides a single automation platform to access the data center network's operational services and tools.

The single-pane network interconnect policy management and the consistent network workload and segmentation policy provided by the Nexus Dashboard Orchestrator allows monitoring the health of the interconnected fabrics, enforcement of segmentation and security policies, and performance of all tasks required to define tenant intersite policies in APIC, NDFC, and cAPIC sites.

Cisco Nexus Dashboard Orchestrator provides these main functions:

- Single pane of glass for administration and orchestration of multiple networking fabrics for both Cisco ACI and NDFC
- Automation of the configuration and management of intersite network interconnects across an IP backbone for both Cisco ACI and NDFC
- Consistent multitenant policy across multiple sites, which allows IP mobility, disaster recovery, and active/active use cases for data centers.
- Capability to map tenants, applications, and associated networks to specific availability domains within the Cisco Multi-Site architecture for both Cisco ACI and NDFC
- Hybrid cloud and multicloud orchestration supporting on-premises Cisco ACI sites and public cloud sites (AWS and Azure)
- Capability to have multicloud ACI deployments without on-premises sites.
- Ability to scale out sites and leaf switches based on resource growth.

Limitations

Certain features of Cisco ACI are outside the scope of this demonstration because the demonstration uses a simulated environment rather than a physical one:

- The simulator is unresponsive if left running for more than two weeks; you must then reboot it.
- The simulator returns to its initial state following a reboot.
- No traffic passes between the two simulators.
- Screen refresh may take slightly longer than expected.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

NOTE: Beginning with Cisco ACI release 6.1, the default Fabric Discovery Policy has been changed from **Permissive** to **Strict**. Due to a limitation of the simulator, onboarding nodes in **Strict** mode is not supported. If you wish to add in apic2 and apic3, then you must manually switch the Fabric Discovery Policy to **Permissive**. Instructions are located in the Appendix.

Customization Options

Customers are provided a dedicated lab. Users are free to customize any available service(s) to fit their needs.

Requirements

The table below outlines the requirements for this preconfigured demonstration.

Required	Optional
Laptop with a web browser (Google Chrome is preferred)	Cisco AnyConnect

About This Solution

This solution demonstrates how Cisco Nexus Dashboard orchestrates policy, connectivity, and segmentation across geographically distributed ACI fabrics. Through a sequence of scenarios, it illustrates the operational models and architectural patterns used to build, scale, and modernize multi-site environments. Each scenario highlights a distinct component of multi-site architecture—from foundational inter-fabric onboarding, to tenant provisioning, to advanced routing designs, application segmentation, disaster recovery, and migration workflows.

Scenario 1 – Initial Setup and Inter-Fabric Configuration: The first scenario defines the groundwork for all subsequent operations by establishing the Nexus Dashboard Multi-Site domain. The fabrics are onboarded, authenticated, and connected, giving NDO the authority to orchestrate policies across them. This scenario focuses on fabric discovery, registration, and establishing the ISN (Inter-Site Network) parameters that allow routing and policy exchange between sites.

Scenario 2 – User Provisioning and NDO Tenant Creation: This scenario introduces multi-site orchestration at the tenant level. A local user with Fabric Administrator privileges is created, and a greenfield tenant is defined entirely within NDO. The tenant, VRF, and associated objects are deployed simultaneously to both sites.The core idea is the centralized lifecycle management of a tenant, avoiding manual per-fabric configuration and ensuring policy uniformity across geographically distributed environments.

Scenario 3 – Layer 3-Only Communication Across Sites (Intra-Tenant, Intra-VRF): the solution demonstrates how to stretch only the Layer 3 constructs—the Tenant and VRF—while keeping all Layer 2 domains site-local. This design allows endpoints in different sites to communicate using routed inter-site paths, while preventing L2 broadcast domains from stretching unnecessarily.The objective is simple and highly practical: enable multi-site application communication without expanding the failure domain of Layer 2.

Scenario 4 – Layer 3-Only Communication Across Sites and VRFs (Shared Services): Building on Scenario 3, this design introduces shared services through a dedicated shared VRF, BD, and EPG. Applications in different VRFs and across different sites can reach centralized services such as DNS or DHCP. It cleanly demonstrates

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

cross-VRF, cross-site service insertion while still maintaining isolation and compliance with multi-tenant principles.

Scenario 5 – IP Mobility Across Sites (Stretched BD Without BUM Flooding): This use case addresses disaster recovery and workload relocation. A single BD and subnet span both sites, allowing endpoints to retain their IP address when moved between locations. By disabling BUM flooding, the design avoids classic pitfalls of stretched Layer 2 while still offering true IP mobility for virtual machines or container workloads that must migrate without renumbering.

Scenario 6 – Importing Brownfield Tenant Configuration: This scenario demonstrates how NDO can integrate an existing tenant from a brownfield ACI fabric into a new multi-site architecture. The existing EPGs, VRFs, and BDs are imported and orchestrated while preserving their current behavior. It shows the solution's capability to modernize and consolidate environments without service disruption, a frequent requirement in enterprise data center evolution.

Scenario 7 – Application-Centric Segmentation Using ESGs: Endpoint Security Groups (ESGs) introduce a higher-level segmentation model that reflects application architecture instead of physical or logical network boundaries. This scenario highlights how ESGs drastically reduce contract complexity and hardware resource consumption by collapsing many EPG-to-EPG relationships into a small number of ESG-to-ESG policies. It is a blueprint for scalable, application-aligned security policy in multi-site environments.

Scenario 8 – Migration of ACI Schema Policies From NDO 3.7 to NDO 4.4 via Configuration Restore: This final scenario covers inter-release migration using backup and restore. A configuration snapshot from NDO 3.7 is restored into a freshly deployed NDO 4.x cluster, where dependency ordering, object validation, and template restructuring are automatically handled. The model supports seamless transition from older design paradigms to the more structured, dependency-aware architecture introduced in NDO 4.x.

Topology

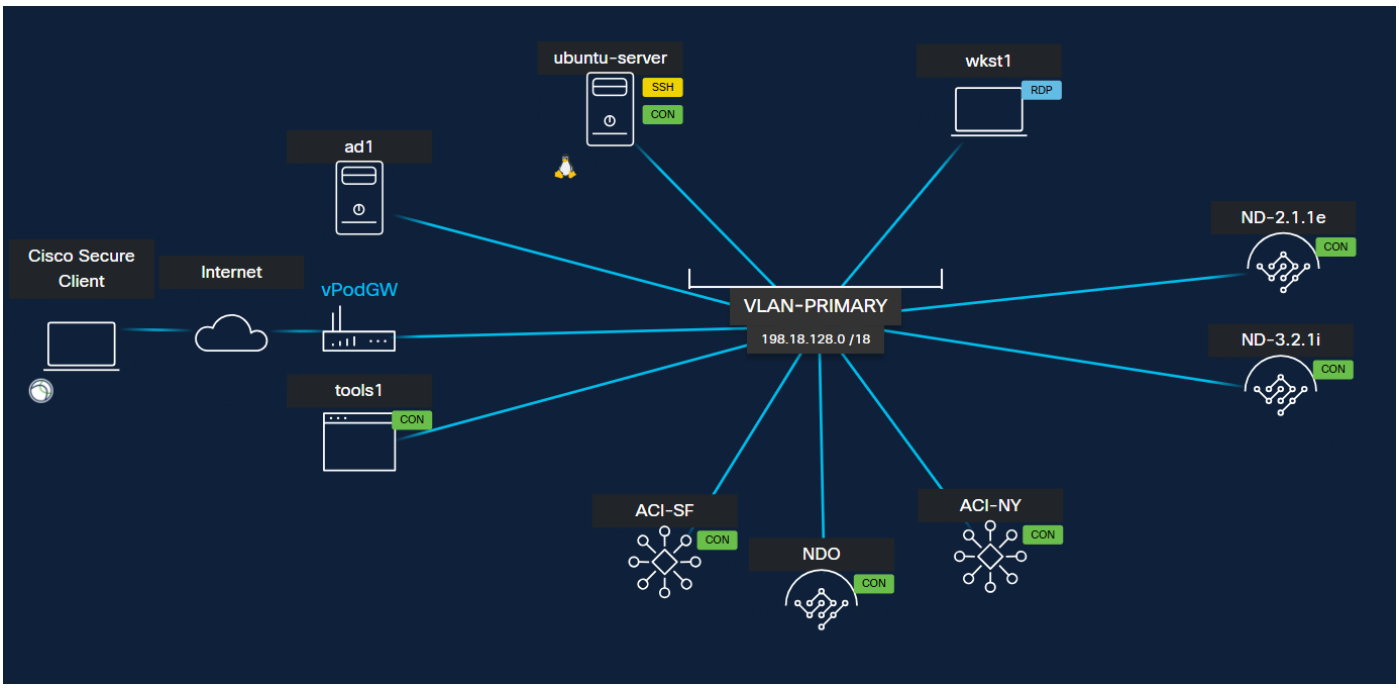

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Fabric Name	San_Francisco	New_York
Pod ID	1	1
AS	65001	65002
Spine 1 – EPVN RID	10.1.100.1	10.2.100.1
Spine 2 – EPVN RID	10.1.100.2	10.2.100.2
Spine 1 – Interface 5/32	10.1.0.1/31	10.2.0.1/31
Spine 2 – Interface 5/32	10.1.0.3/31	10.2.0.3/31
O-UTEP	10.1.100.100/32	10.2.100.200/32
O-MTEP	10.1.100.200/32	10.2.200.200/32
Infra TEP	10.0.0.0/12	10.0.0.0/12
Data Plane TEP	10.1.200.200/32	10.2.200.200/32

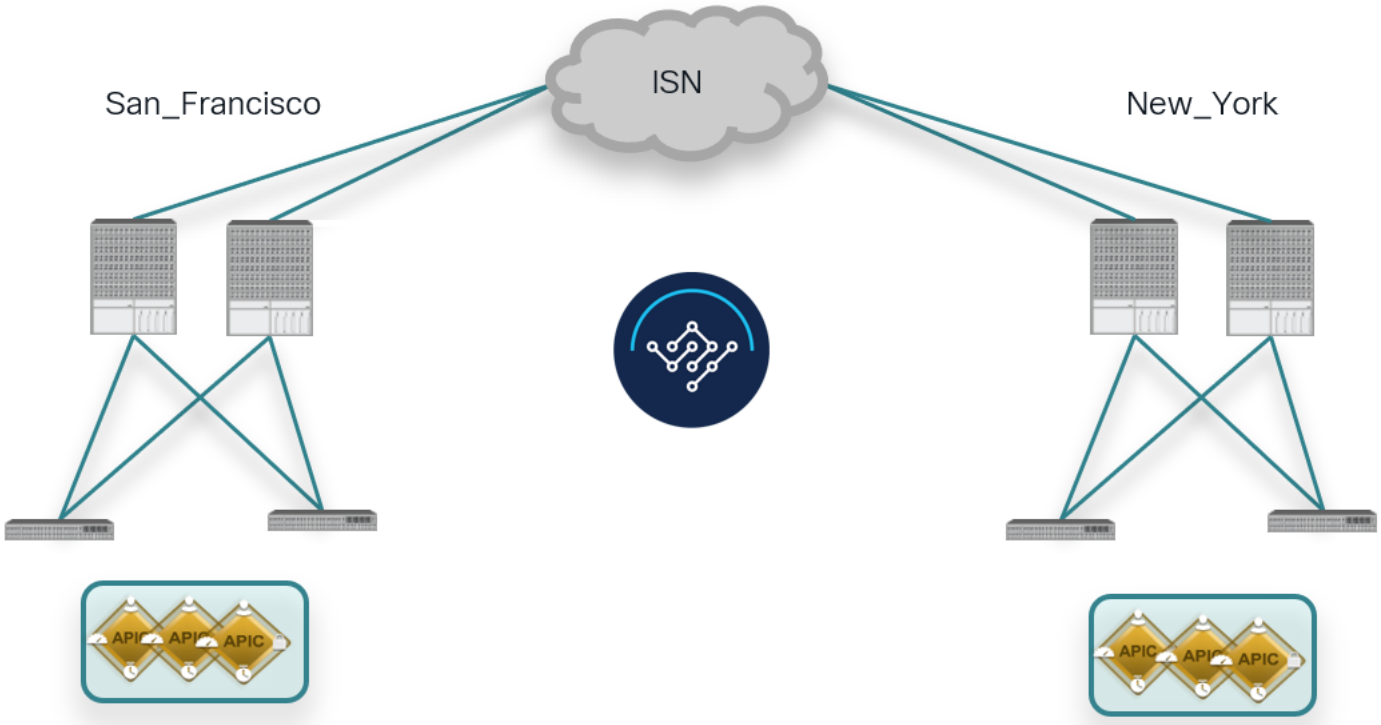

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Components

(Virtual Appliances)

- Cisco Nexus Dashboard 4.1.1g with Orchestrator enabled
- Cisco Nexus Dashboard 3.2.1i (NDO 4.4)
- Cisco Nexus Dashboard 2.1.1e (NDO 3.7)
- Two Cisco ACI simulator version 6.1.4h
- Ubuntu 20.04
- Windows 11

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Get Started

Follow the steps to schedule a session of the content and configure your presentation environment.

1. Schedule your dCloud session. When your session is available, go to **My Hub** > **Sessions** and click **View**.
   - It can take up to 50 minutes for your session to become available
2. Connect to the workstation using one of two connection methods:
   - Cisco AnyConnect VPN and your local RDP client (Best Quality)
   - Cisco dCloud Remote Desktop Client (Easiest)

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 1. Initial Setup and Inter-Fabric Configuration

This scenario defines the essential groundwork for all subsequent multi-site operations within this demonstration. It guides you through establishing the core Nexus Dashboard Multi-Site domain by onboarding, authenticating, and connecting your ACI fabrics. This process grants Cisco Nexus Dashboard Orchestrator (NDO) the authority to centrally orchestrate policies across your distributed environment.

Key steps in this scenario include fabric discovery, registration, and the crucial establishment of Inter-Site Network (ISN) parameters, which enable seamless routing and policy exchange between sites. This foundational setup is critical for building, scaling, and modernizing multi-site environments, preparing the topology for consistent policy and operations.

NOTE: Name the sites exactly as instructed, or you will be unable to complete the demo successfully.

Onboard ACI Sites to Nexus Dashboard

1. In the taskbar, open **Google Chrome**.
2. Click the **Nexus Dashboard** bookmark.
3. You should be presented with the ND login page. Enter the credentials to sign in:
    - **Username:** admin
    - **Password:** C1sco12345

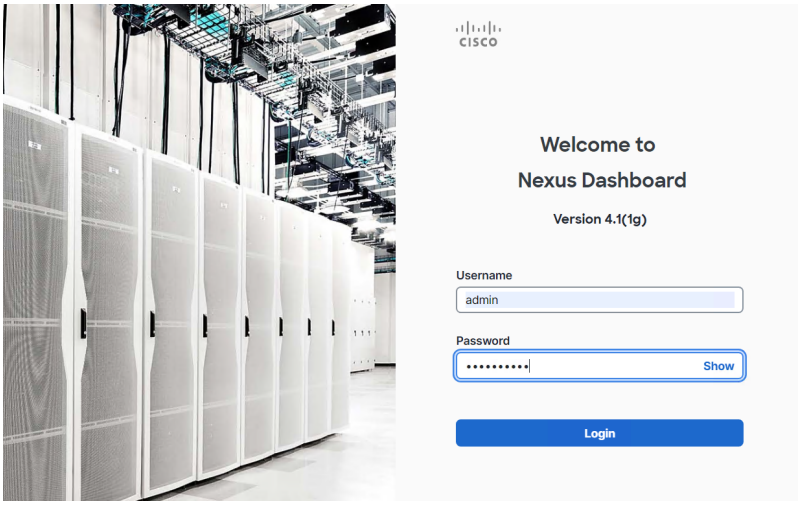

4. In Nexus Dashboard UI, navigate to the left-hand menu, and select **Admin > System Settings > Multi-cluster Connectivity**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

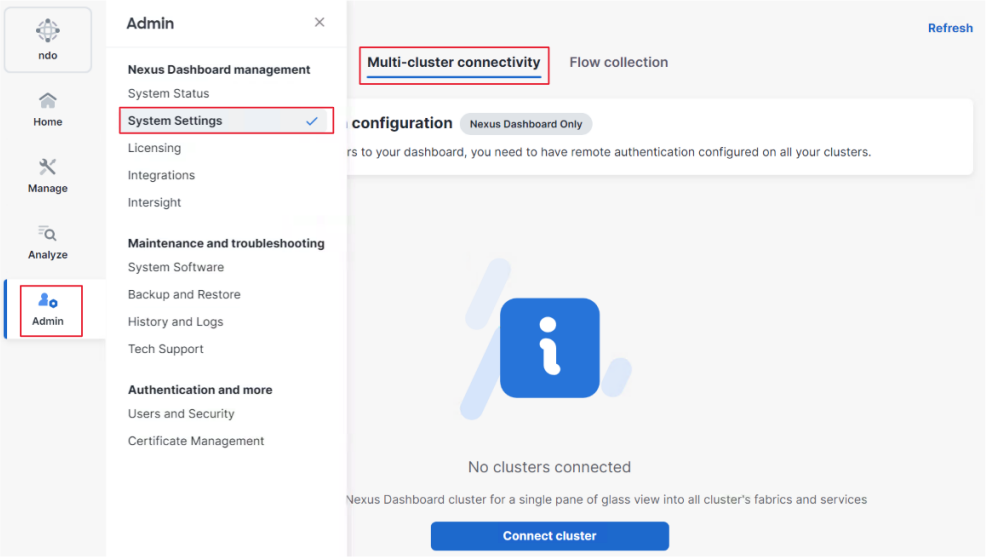

5. On the Multi-cluster Connectivity page, click **Connect Cluster**.
6. In the Select Type section, choose **ACI**, and then click **Next**.

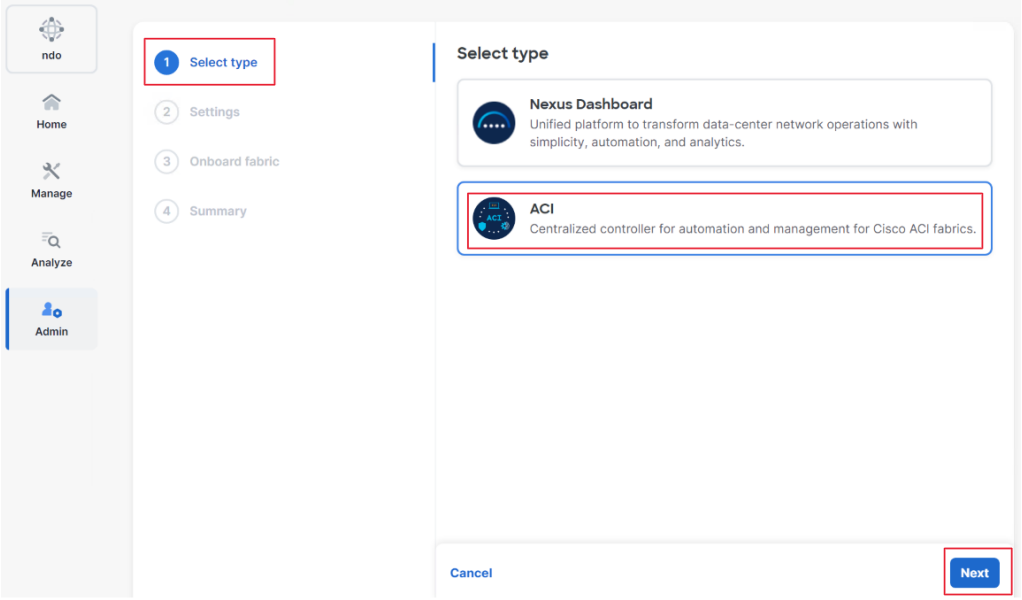

7. In the Settings section, add the following details:
   - **Hostname/IP address:** 198.18.133.200
   - **Username:** admin
   - **Password:** C1sco12345

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

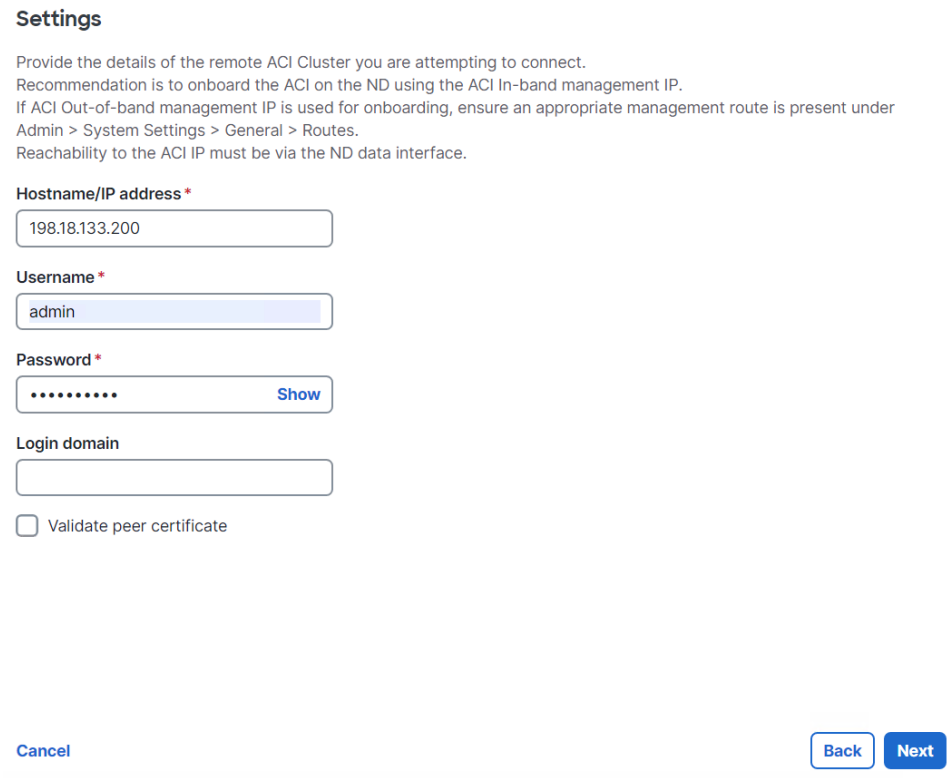

8. Click **Next.**
9. In the onboard fabric section add the following details:
    - Fabric name: **San_Francisco**
    - Location: San Francisco
    - License tier: **Premier**

NOTE: When the Premier license is selected, the telemetry option is enabled by default. Please ensure that you **deselect** this option.

- Enabled features: **Orchestration**

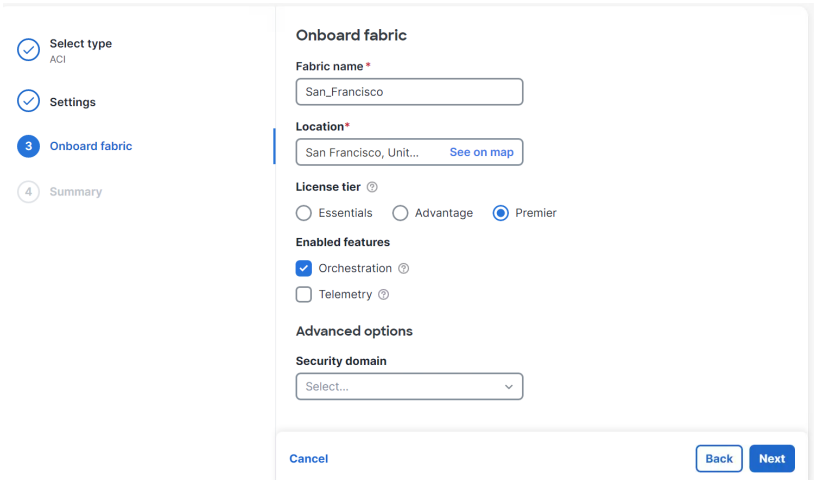

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

10. Click **Next.**

11. Review the configuration summary, then click **Connect** to onboard the site.

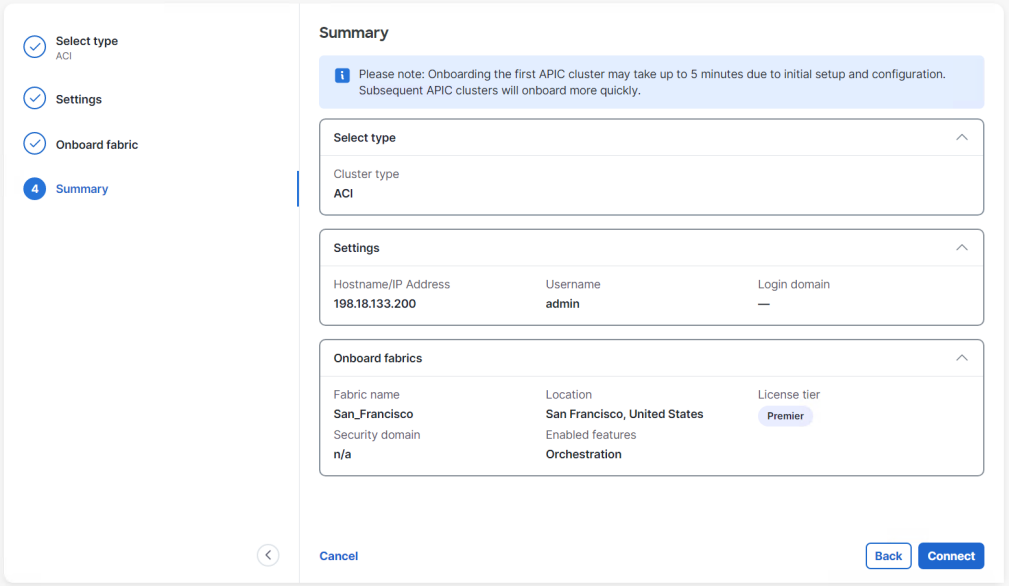

12. From the left-hand menu, navigate to **Admin** > **System Settings** > **Multi-cluster Connectivity**.

13. Verify that the **San_Francisco** site is listed, with a **Connectivity Status** of Up (green), along with its **APIC version** and **Hostname/IP address**.

14. Next, onboard the second site (**New_York**) by expanding the **Actions** menu and selecting **Connect Cluster**.

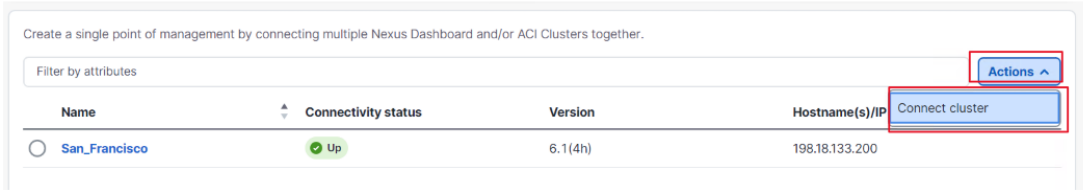

15. Repeat the same configuration steps as before, using the following IP address for the New_York site: **198.18.132.200**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

16. Navigate to **Admin** > **System Settings** > **Multicluster Connectivity and** verify that both the **San Francisco** and **New York** sites show a green “Up” status.

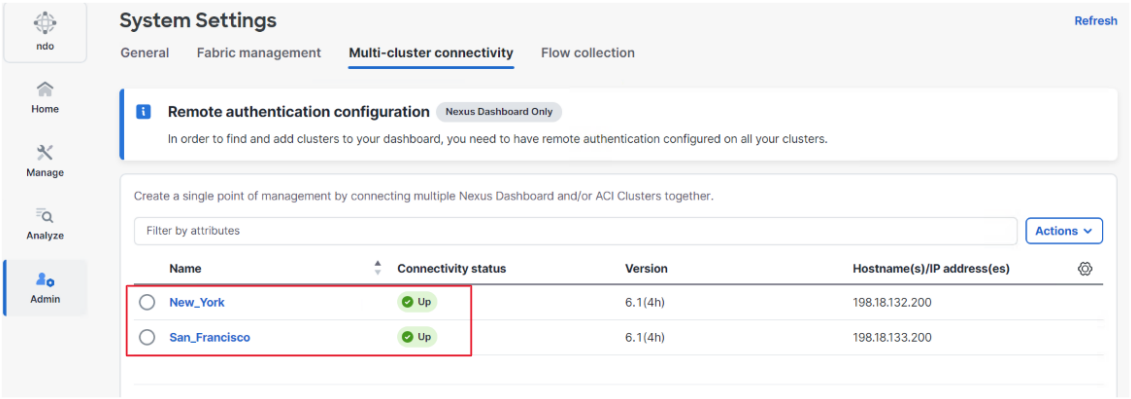

17. Navigate to **Manage** > **Orchestration** and confirm that there are two sites under **Fabrics with Associated Templates**.

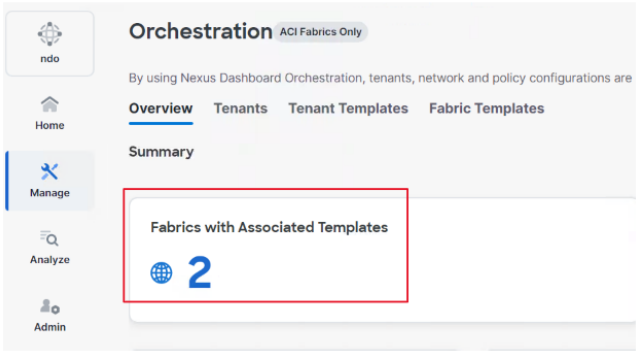

Site Visibility in Nexus Dashboard

After successfully onboarding the two sites, you can monitor the entire fabric’s health and configuration directly from the Nexus Dashboard Orchestrator (NDO) interface. The platform provides a unified operational view across all connected sites, allowing administrators to instantly assess status, configuration, and activity without logging into each APIC.

1. In the left-hand menu, navigate to **Manage** > **Fabrics**, where all onboarded ACI fabrics (New_York and San_Francisco) are listed, then double-click **New_Yo**rk to open its site-level dashboard.

Overview Tab

- From the Overview page, the New_York site’s key operational data is immediately visible.
- The General section displays essential information such as the fabric name, type, license tier, and software version, giving an instant summary of the site’s configuration.
- The Config-Sync Status shows In-Sync, confirming that the site configuration is fully synchronized with the Orchestrator.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

- On the right-hand side, the Recent Activity panel lists the latest administrative actions, such as enabling orchestration and connecting the APIC cluster, providing complete operational traceability.
- You can also see at a glance the number of switches, controllers, and interface states, helping you evaluate infrastructure status without navigating deeper into the APIC GUI.
- This page effectively acts as a dashboard within a dashboard, consolidating both configuration and operational data in a single pane of glass.

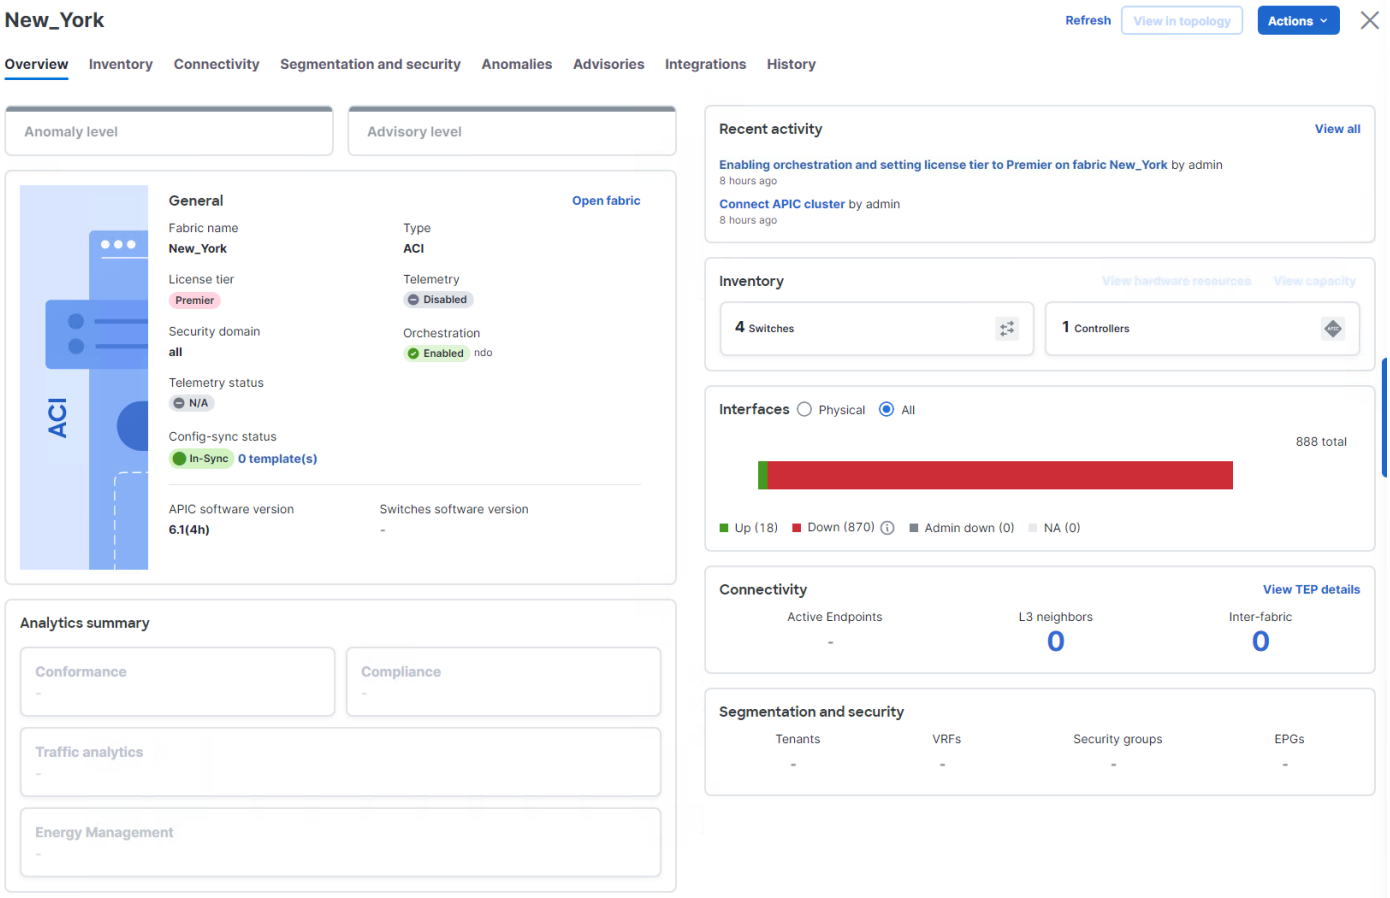

Inventory Tab

- The Inventory tab provides detailed visibility into all physical and logical components of the New York site.
- Under Switches, you can see the entire switch inventory: two spines and two leafs and showing an active status.
- The table lists serial numbers, TEP pools, and roles, which makes it easy to verify that the expected hardware components are online and properly integrated into the fabric.
- You can switch between Controllers and Switches views to get a complete hardware overview.
- This granular visibility lets administrators verify hardware health, software versions, and connectivity directly from the Nexus Dashboard—eliminating the need for manual logins to each site's APIC controller.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

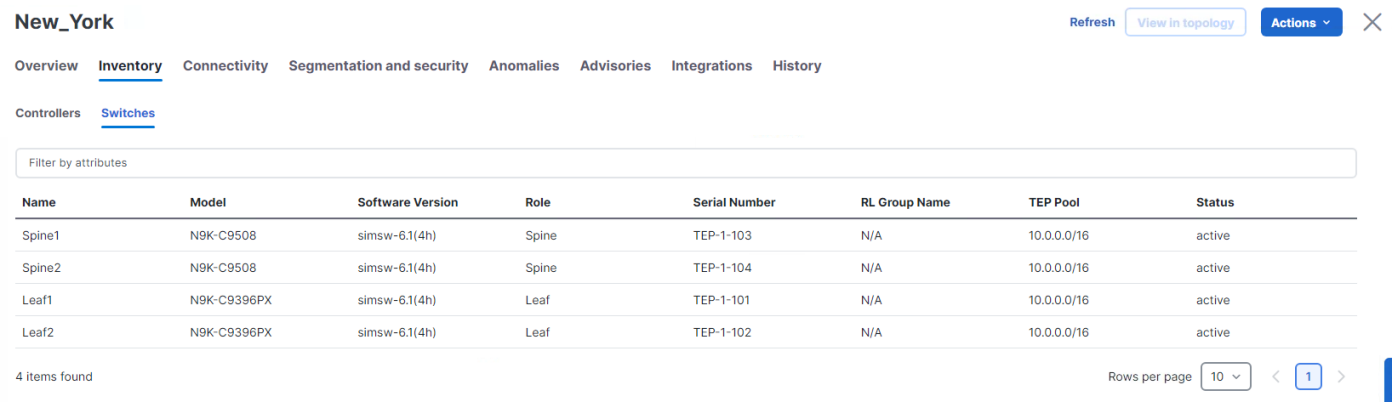

Connectivity Tab

- The Connectivity view offers deep operational insight into interface-level status across the entire fabric.
- The summary graphs display Interfaces by Operational State and Interfaces by Admin State, giving a quick health snapshot.
- In this example, 18 interfaces are operationally up, while 870 are down, which is typical in a lab or partially built environment.
- Each interface entry lists its node ID, operational state, and admin state, providing a complete picture of port-level behavior.
- This centralized visibility demonstrates the power of Nexus Dashboard as a single point of management and monitoring. Administrators can instantly detect issues, verify configurations, and track activity across multiple fabrics—all from one intuitive interface.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

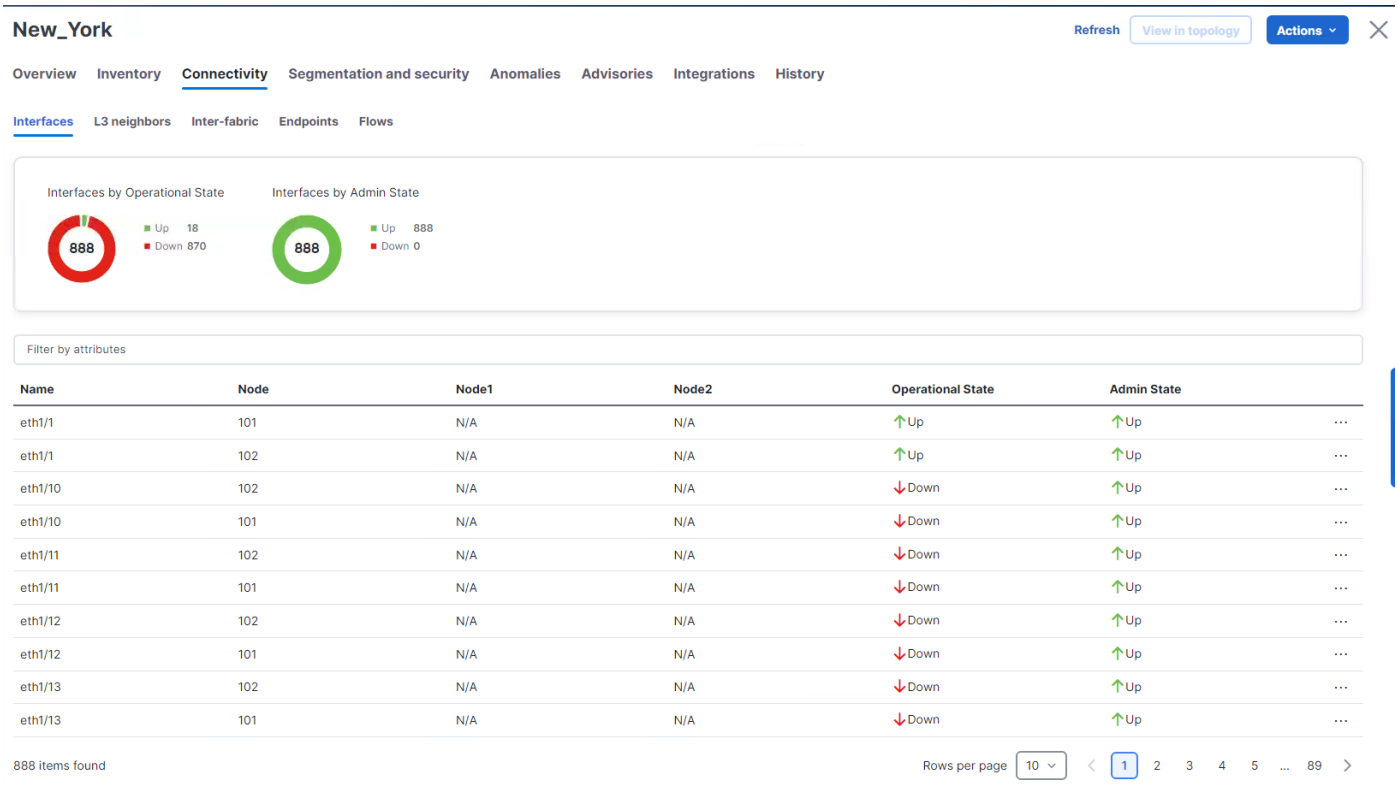

Inter-Fabric Connectivity

The Inter-Fabric Connectivity configuration establishes the foundational Layer 3 connectivity between ACI fabrics (sites) and is managed through the Nexus Dashboard Orchestrator (NDO).

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

NOTE: To bypass the remaining inter-fabric configuration tasks in this scenario, complete the following three steps.

Step 1: Execute the **FixMyDemo** Utility:

a. On the workstation desktop, launch the **FixMyDemo** script using the provided shortcut.

b. When the menu appears, select Option 3: "**Skip Intersite Infrastructure Configuration**".

Step 2: Collect Site details from Nexus Dashboard:

a. Log in to Nexus Dashboard.

b. In the upper-right corner, select ? icon and choose Help Center.

c. Under Deployment Planning section, click **API Reference: Swagger (In-product)**.

d. In the left navigation pane, click **Orchestration** and locate the **Site Information** section.

e. Expand GET /api/v2/sites/ and select Try it out.

f. **Execute** the request and record the returned "**id"**, "**siteId"**, and "**siteGroup"** values for **each site**. These values are required for preparing the inter-fabric configuration payload.

Step 3: Submit the Inter-Fabric Configuration Payload

a. Open the JSON file "inter-fabric" on the desktop.

b. Insert the appropriate values for every site, based on the data collected in step 2.

"To identify which site corresponds to each entry in the JSON file, reference the **bgpAsn** value. The **San_Francisco** site uses **65001**, and the **New_York** site uses **65002**"

c. Return to the **Site Information** section in Swagger and locate **POST /api/v2/sites/manage**.

d. Select **Try it out**, paste the completed JSON payload into the request body, and select **Execute**.

Upon completing these steps, Scenario 1 is fully bypassed and you can start from Scenario 2.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Configurations Managed from the Nexus Dashboard Orchestrator

The following parameters are defined and applied centrally from NDO to ensure consistent intersite connectivity and routing:

- Spine Interface Configuration: Interface selection, IP address assignment, and MTU settings for intersite links.
- Control Plane E-TEP Configuration: Defines overlay endpoints used for BGP EVPN control-plane peering between sites.
- Unicast Data Plane E-TEP Configuration: Establishes unicast transport endpoints for intersite traffic forwarding.
- Multicast Data Plane E-TEP Configuration: Configures multicast transport endpoints for intersite multicast forwarding.
- OSPF Area Configuration: Defines OSPF areas (area ID and area type) for spine-to-IP network connections.
- BGP Autonomous System Number (ASN): Assigns the BGP ASN for each fabric, aligning with the existing fabric-level ASN.
- BGP Community Configuration: Defines optional BGP community attributes for route tagging and policy control.
- External L3 Domain Selection: Associates with the external Layer 3 domain used for intersite routing.

Configurations Managed Locally from Each APIC

The following parameters are configured locally on the APIC at each site, as they are specific to the local fabric's topology and external connectivity policies:

- Access Policy Configuration for the External L3 Domain:
  Includes defining the spine switch profile, interface profile, interface policy group, attachable entity profile, and external Layer 3 domain.
- BGP Route Reflector Policy:
  Specifies the route reflector nodes and associated parameters for intra-fabric BGP peering.

The NDO automatically reads the BGP ASN and External L3 Domains from each site's APIC. Ensure these elements are correctly configured on the respective APICs before proceeding.

Configure BGP Route Reflector and L3 Domain

1. The Open Google Chrome shortcut on the workstation desktop, click the **APIC-SF** tab from the bookmark bar and log in:
   - **Username**: admin
   - **Password**: C1sco12345

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2. Navigate to **System** > **QuickStart**. Click Start **ACI Fabric Setup**.

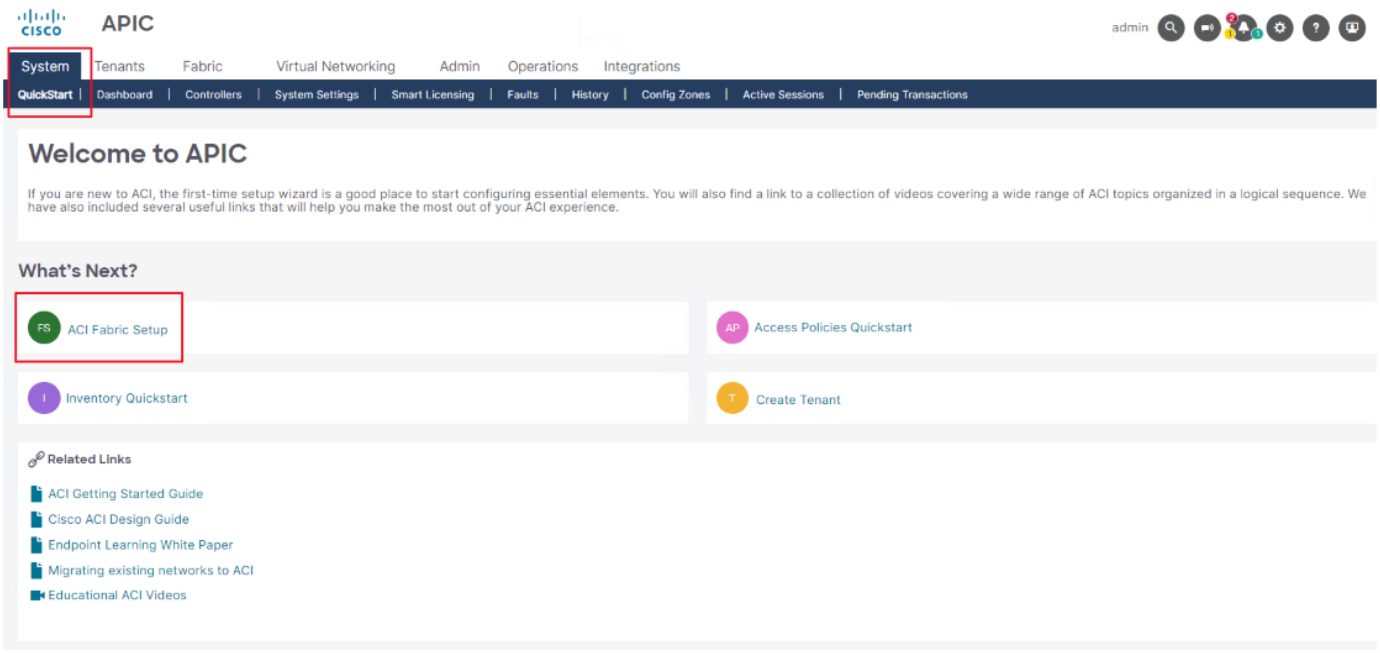

3. Click **Begin** under BGP setting.

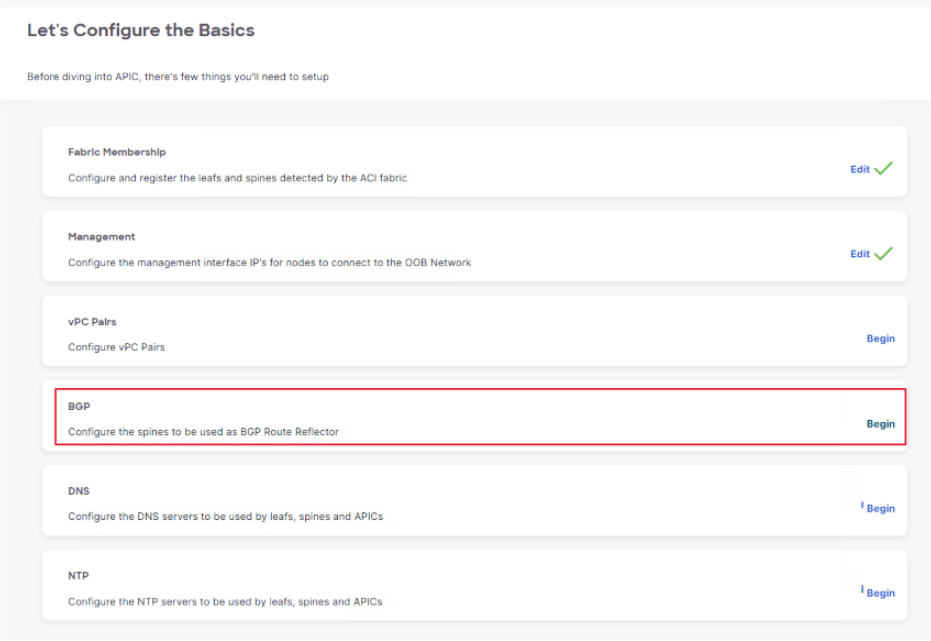

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

4. Enter **65001** in the Autonomous System Number field, select spine **104** and **103** and click **Save**.

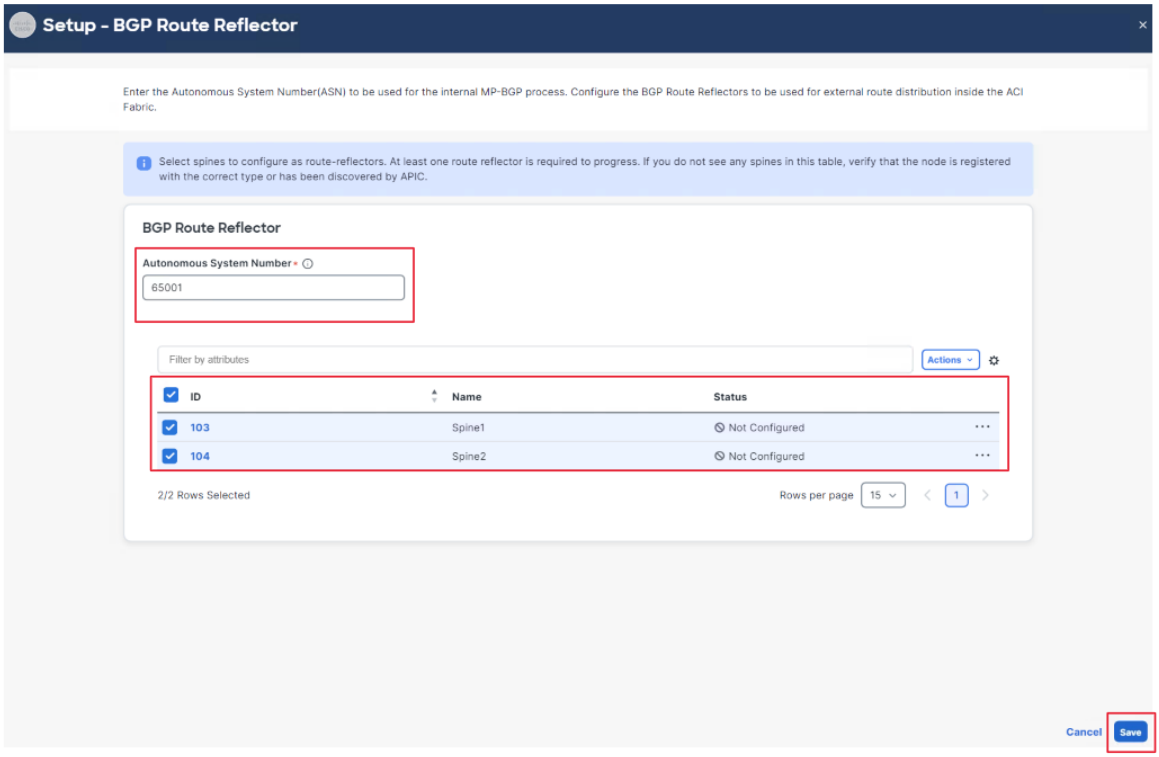

5. Close the Setup-Overview window on the right top corner

NOTE: Although additional configuration sections such as **DNS**, **NTP**, and **Intersight Proxy** are marked as mandatory on both sites, they are not relevant to this demonstration and can remain unconfigured.

6. Click **Fabric** > **Access Policies** in the top menu and expand **Physical and External Domains**.
7. Click **L3 Domains**.
8. Expand the **Tools** menu at the upper right and select **Create L3 Domains**.

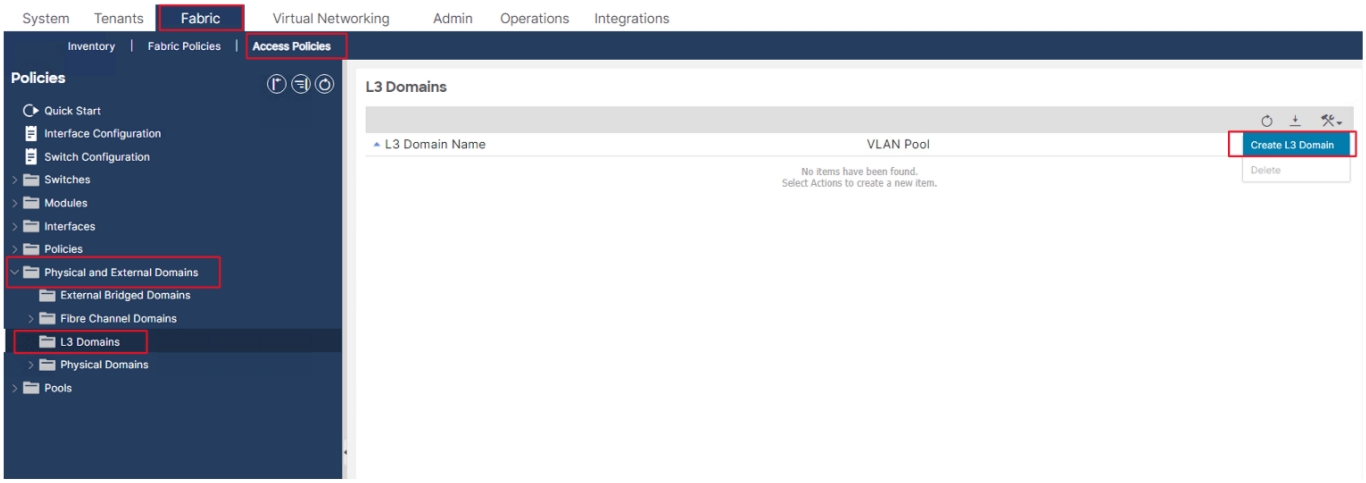

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

9. Enter **Multisite_External_L3_Domain** in the **Name** field and click **Submit**.

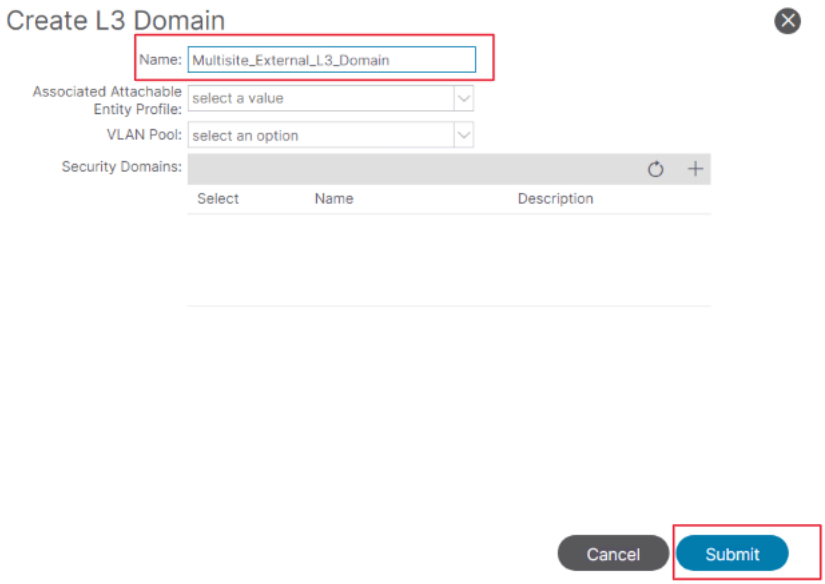

10. Now let's repeat the same steps to site 2. Go back to google Chrome and click **APIC-NY** tab from the bookmark bar and log in:
    - Username: **admin**
    - Password: **C1sco12345**
11. Navigate to **System** > **QuickStart** then, click Start **ACI Fabric Setup**.
12. Click **Begin** under BGP setting.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

13. Enter **65002** in the **Autonomous System Number** field, select spine **104** and **103** and click **Save**.

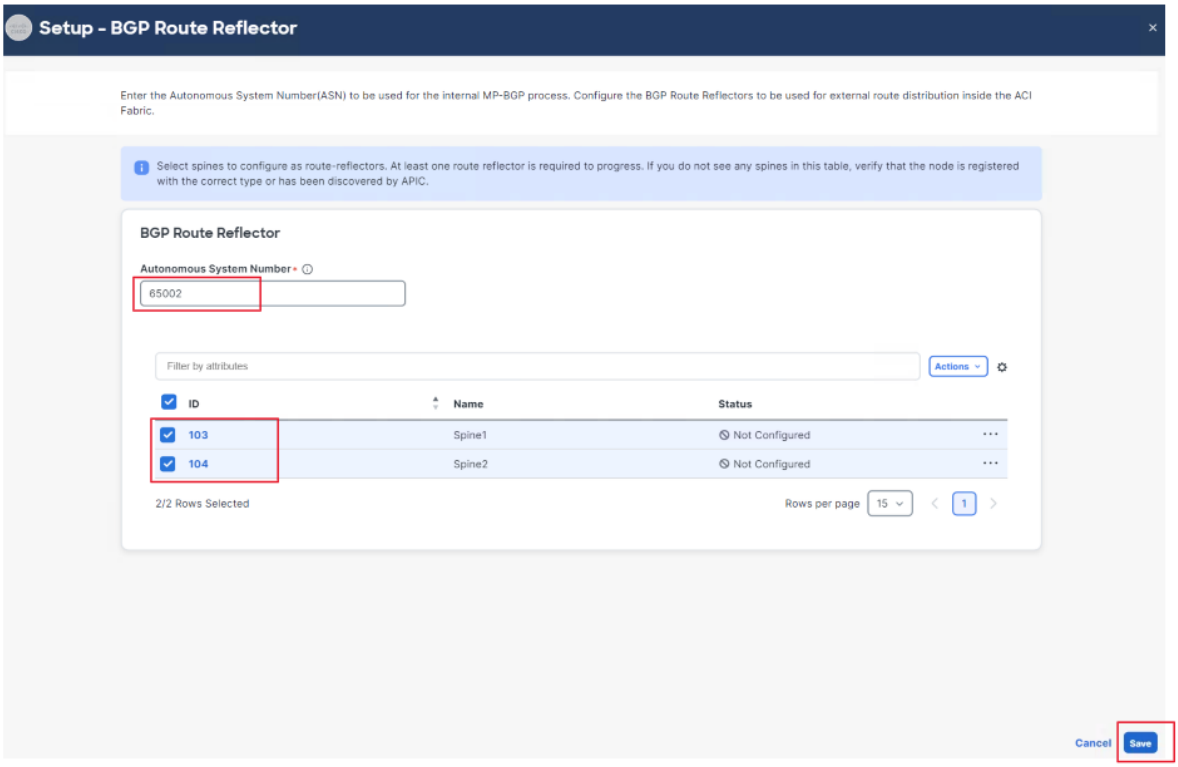

14. Close the Setup-Overview window on the right top corner.
15. Click **Fabric** > **Access Policies** in the top menu and expand **Physical and External Domains**.
16. Click **L3 Domains**.
17. Expand the Tools menu at the upper right and select **Create L3 Domains**.
18. Enter **Multisite_External_L3_Domain** in the Name field and click **Submit**.

Configure the Inter-Fabric Connectivity in NDO

1. The Open Google Chrome shortcut on the workstation desktop, click Nexus Dashboard tab from the bookmark bar and log in:
   - **Username**: admin
   - **Password**: C1sco12345
2. In the left-hand menu, navigate to **Manage > Orchestration**.
3. Click **Inter-Fabric Connectivity**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

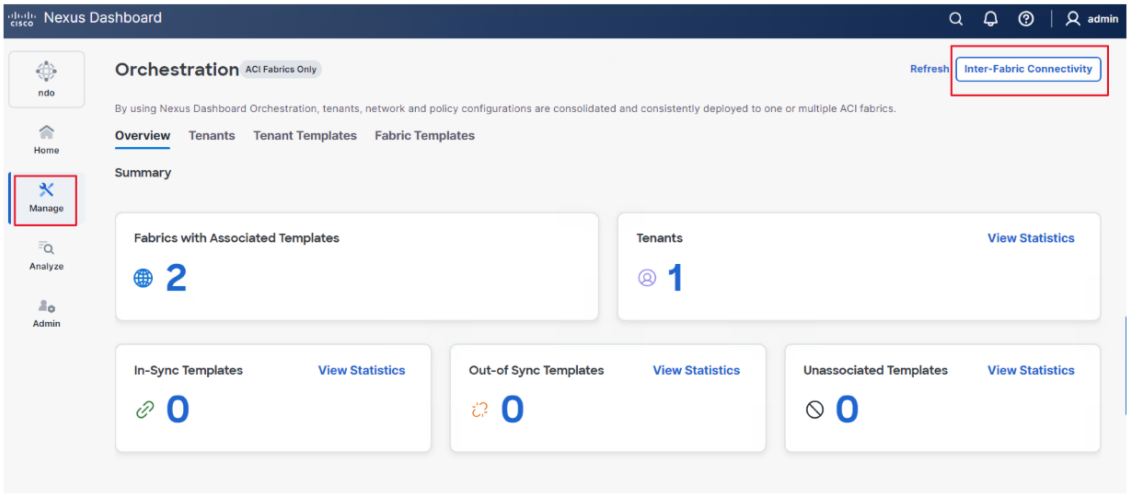

4. Now In the **Fabrics Configuration** step, verify that both onboarded sites show a **Deployed (green)** status, then expand the **General Settings** section to review the **BGP Configuration**, keep the default values unchanged, and click **Next** to proceed.

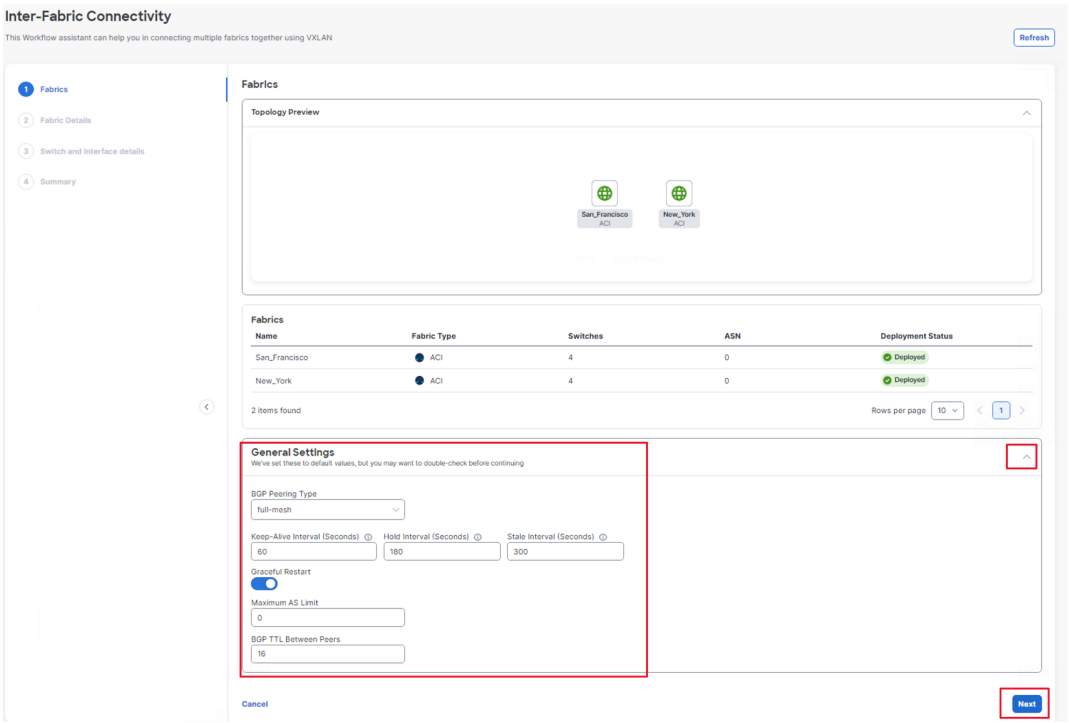

In the Fabric Details section, we will enable Multi-Fabric Connectivity and configure BGP, OSPF, External Routed Domain, Unicast TEP (UTEP), and Multicast TEP (MTEP). These settings establish the control-plane and data-plane connectivity required for intersite communication—BGP and OSPF handle route exchange between fabrics, while the UTEP and MTEP configurations define the transport endpoints for unicast and multicast traffic across sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5. Click **Edit** next to the **San_Francisco** site.

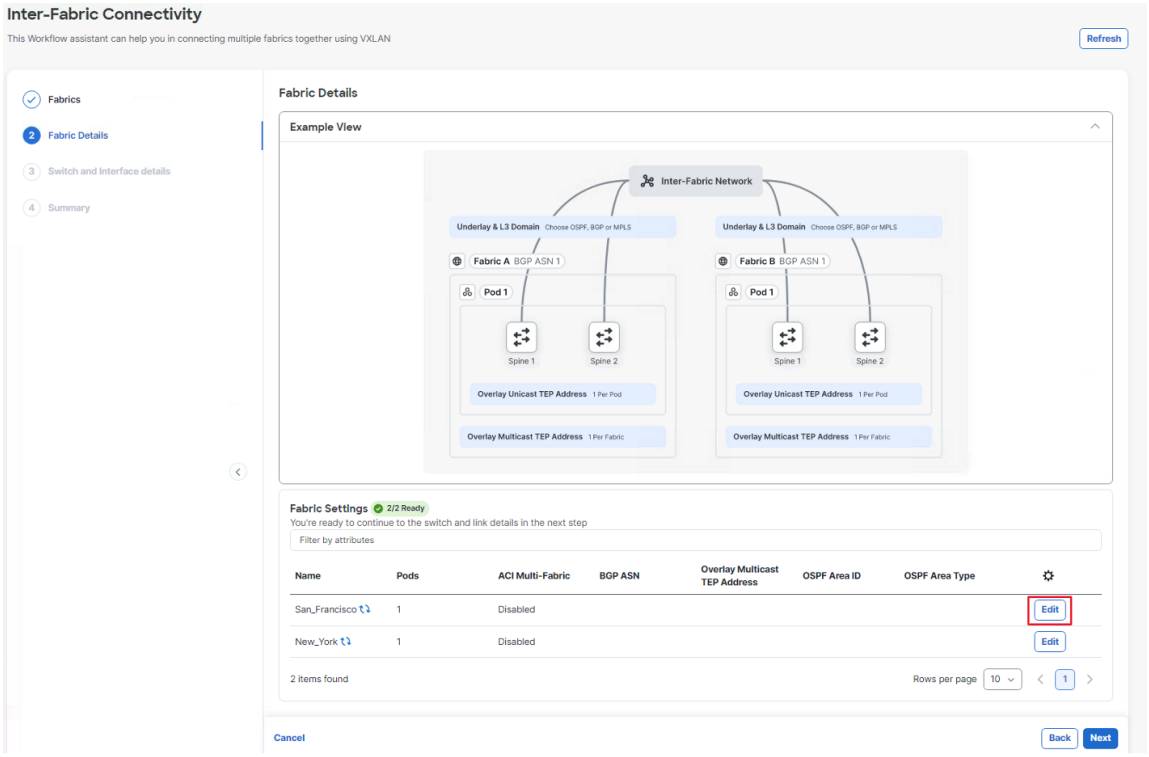

6. Click **Enable Multi-Fabric**, then complete the remaining fields as follows:
   - BGP ASN: **65001**
   - External Routed Domain: Choose **Multisite_External_L3_Domain**
   - Overlay Multicast TEP Address: **10.1.100.200**
   - Area ID: **0**
   - Area Type: **regular**
   - Network type: point-to-point
   - Expand pod-1, in Overlay Unicast TEP**: 10.1.100.100**
   - Click **save**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

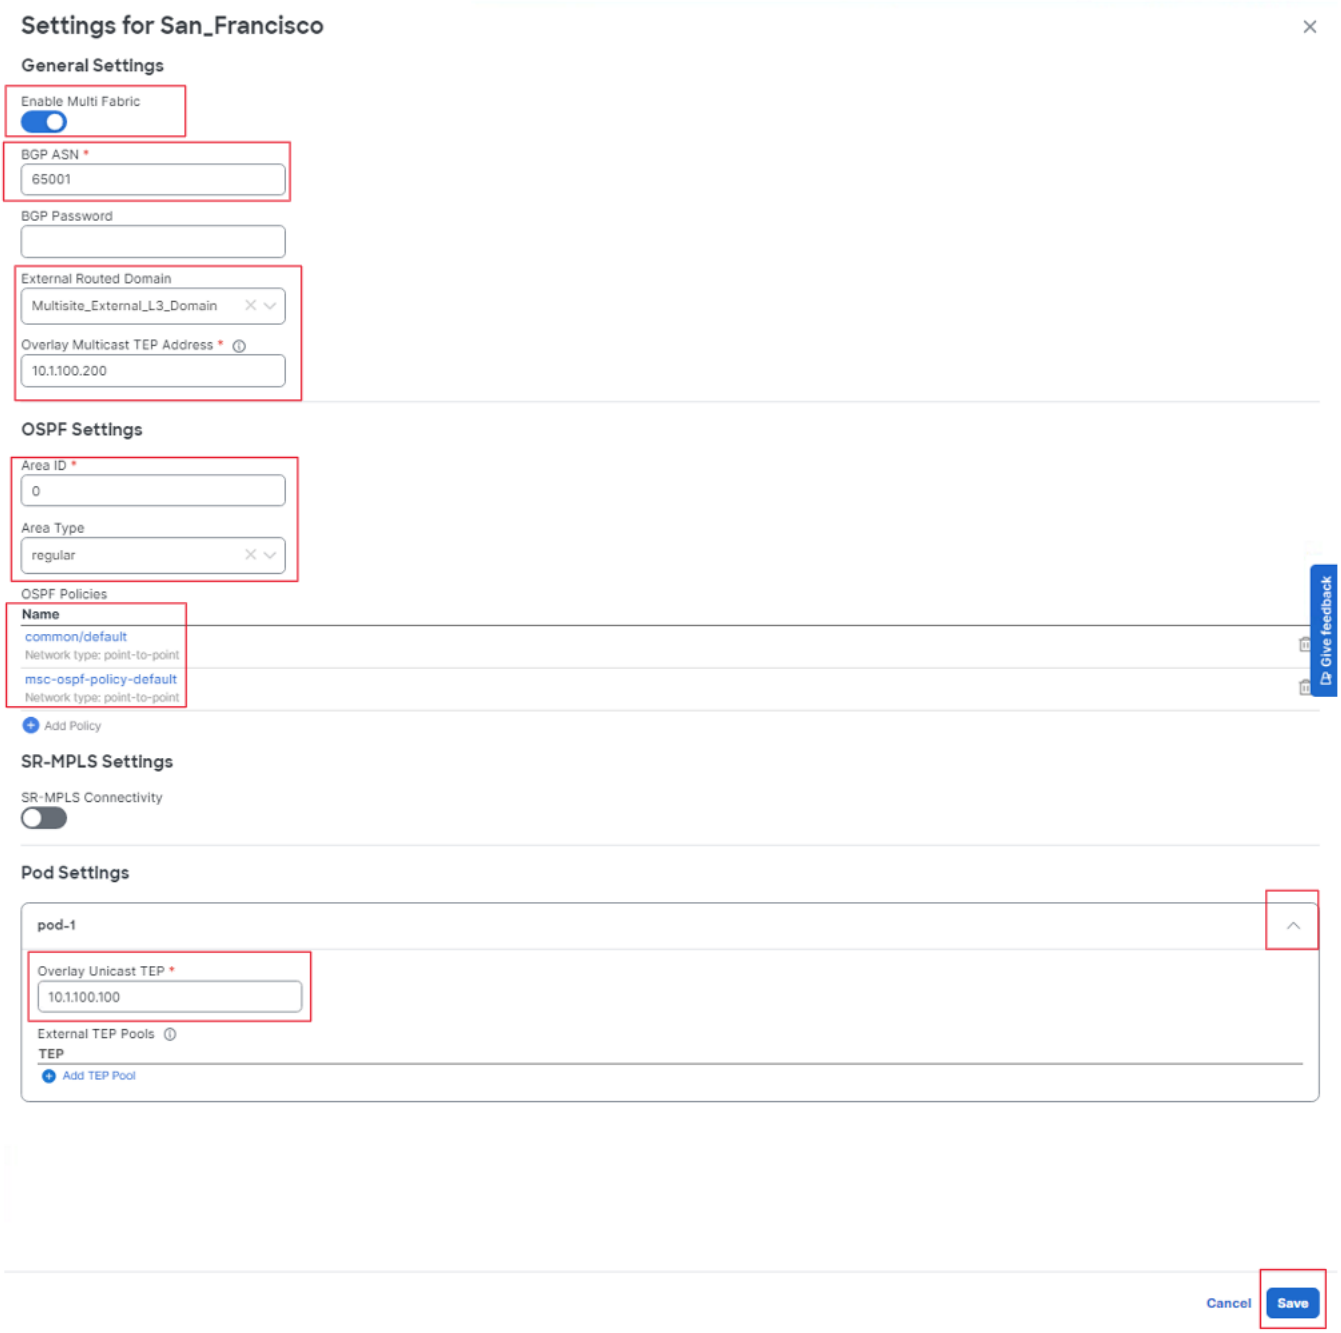

7. Click **Edit** next to the **New_York** site.
8. Click **Enable Multi-Fabric**, then complete the remaining fields as follows:
   - BGP ASN: **65002**
   - External Routed Domain: Choose **Multisite_External_L3_Domain**
   - Overlay Multicast TEP Address: **10.2.100.200**
   - Area ID: **0**
   - Area Type: **regular**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

- Network type: point-to-point
- Expand pod-1, in Overlay Unicast TEP**: 10.2.100.100**
- Click **save**

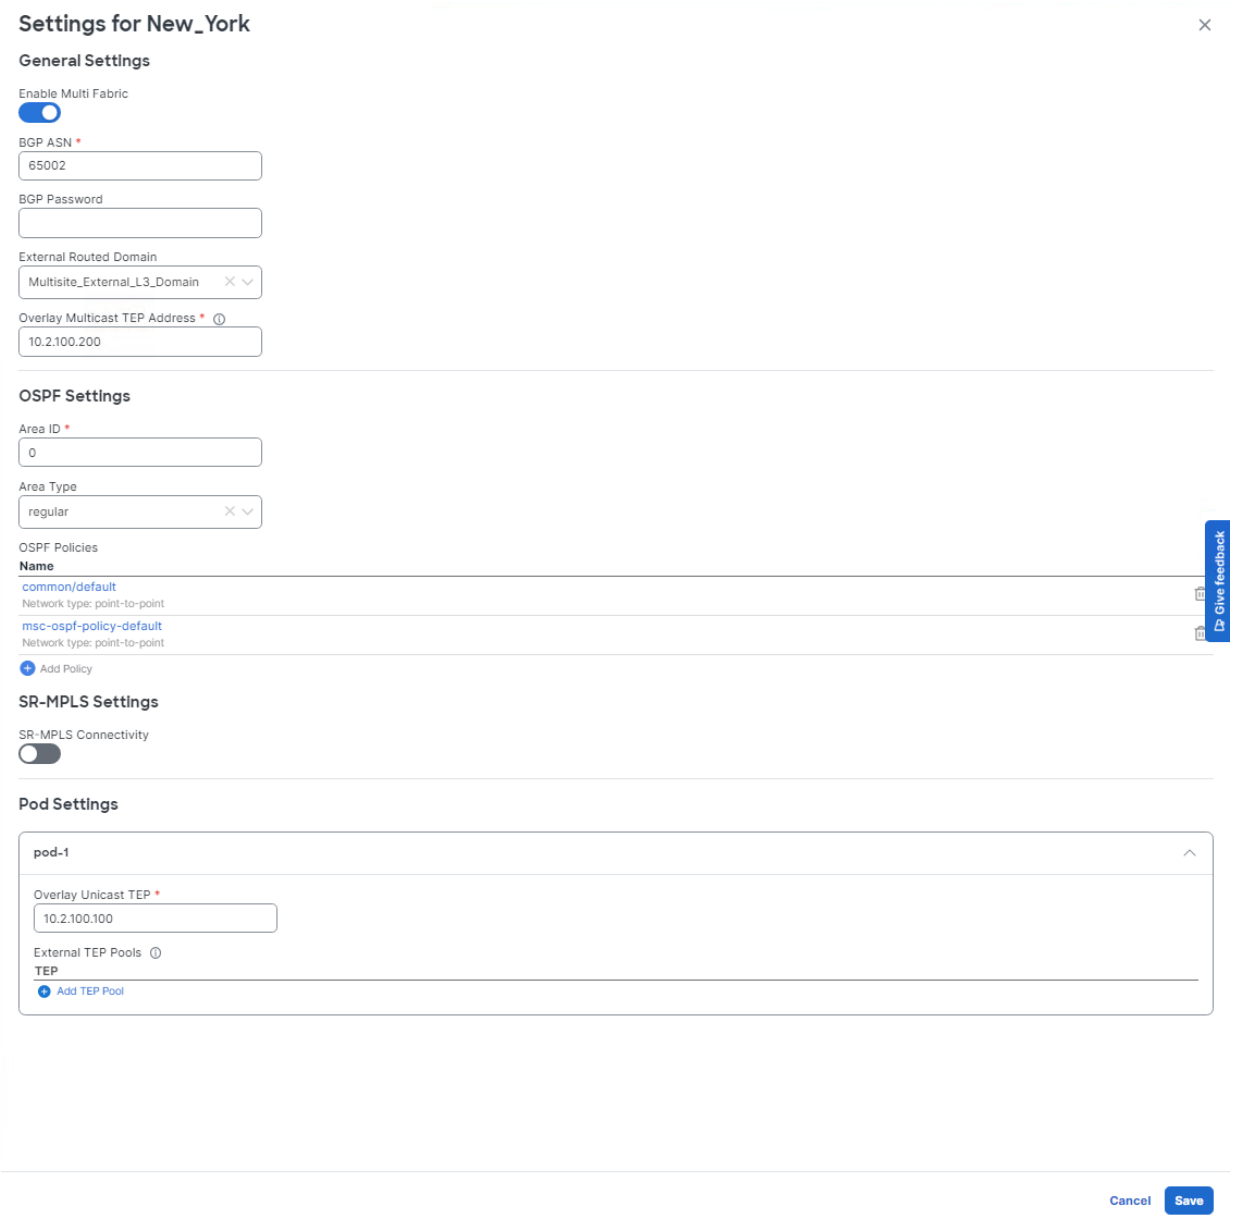

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

9. Review the **Fabric Settings**. Click **Next** to continue.

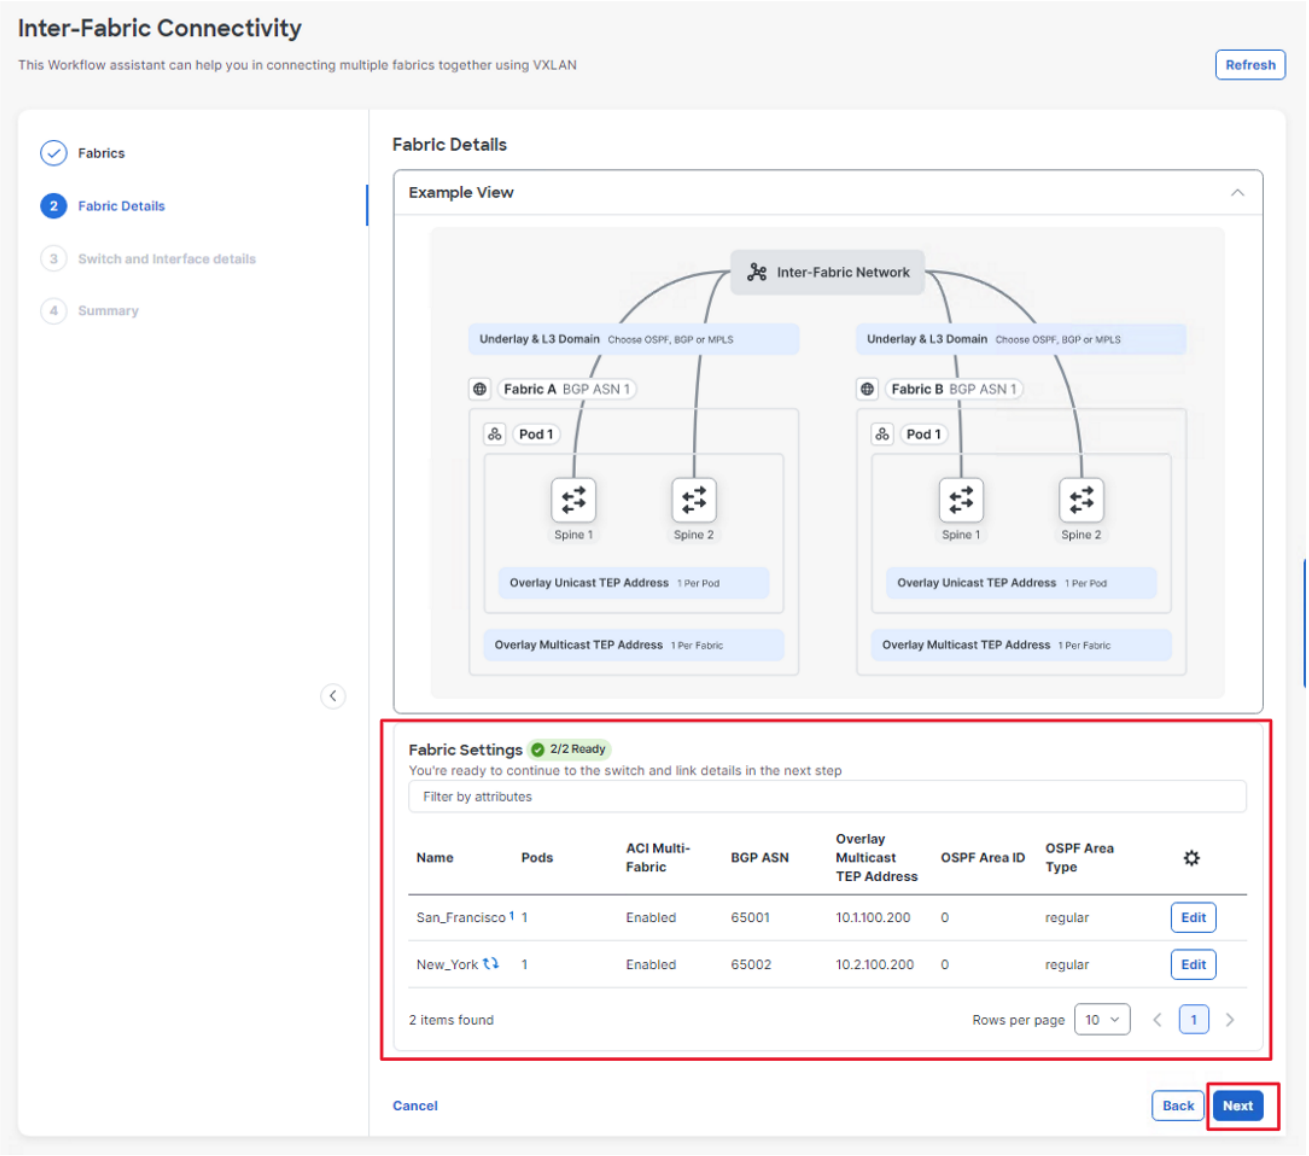

In this step, we will select the spine interfaces on each site that connect to the Inter-Site Network (ISN) devices. These interfaces are used for establishing the underlay transport between fabrics, enabling both control-plane (BGP EVPN) and data-plane (unicast and multicast) communication across sites.

In this demo scenario, we will configure one interface per spine for each site—All interface parameters (such as MTU, routing protocol, and BGP settings) will remain the same for both sites. The only variation will be the interface IP address, which must be unique per spine interface.

After interface selection and IP assignment, these links will serve as the transport path for all inter-fabric communication between the San Francisco and New York sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

10. Click **Add Link** at the bottom of the page.

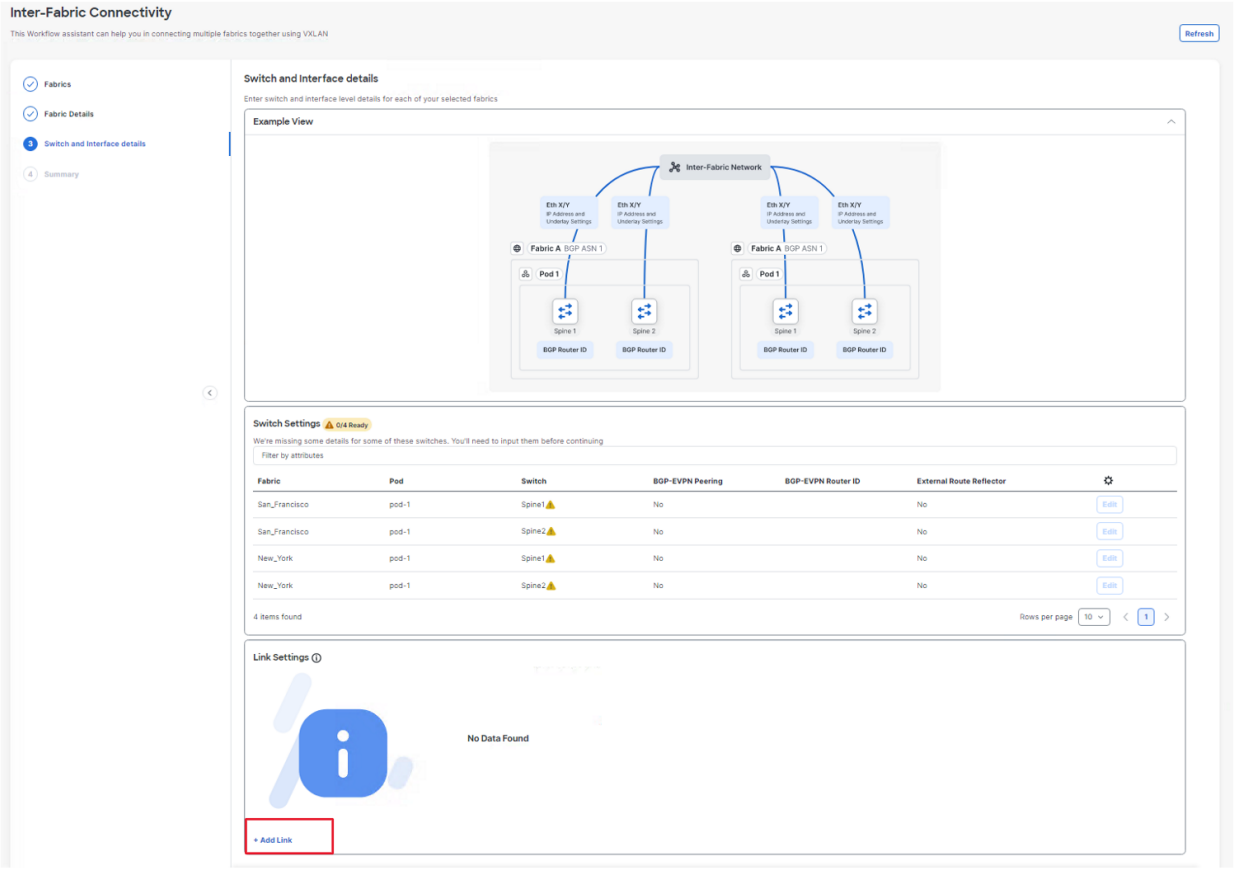

11. Click **Add Link** at the bottom of the page

12. Now, we will add four interfaces—one per spine in each site. All configuration parameters will remain identical, except for the **Interface IP Address** fields. Begin by adding the interfaces for the **San Francisco** site:

    - Fabric: San_Francisco
    - Switch: Spine1 (pod-1)
    - Interface: **5/32**
    - IP Address: **10.1.0.1/31**
    - MTU: **9216**
    - Enable OSPF: **Checked**
    - OSPF Policy: msc-ospf-policy-default

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

13. Repeat the same steps and parameters, changing only the IP address, for the following spines:

- Spine2 (San_Francisco): IP Address **10.1.0.3/31**
- Spine1 (**New_York**): IP Address **10.2.0.1/31**
- Spine2 (**New_York**): IP Address **10.2.0.3/31**

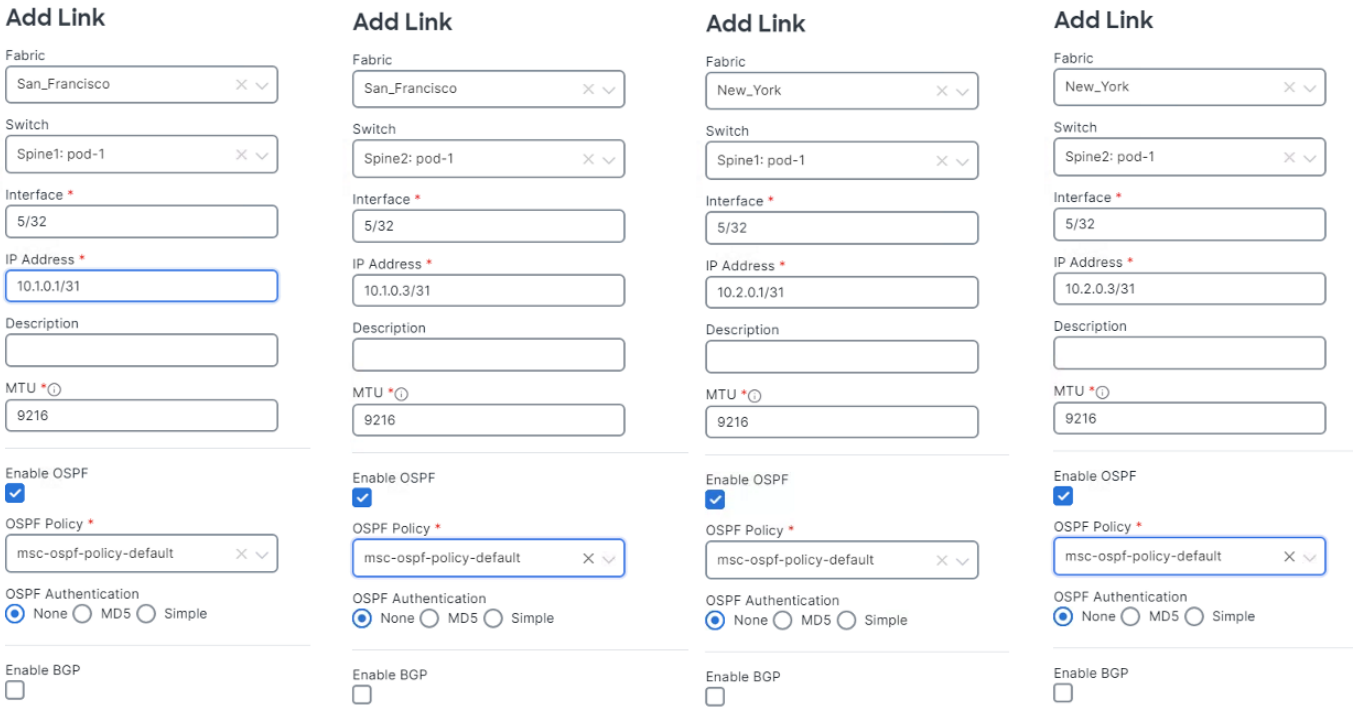

After configuring all interfaces, the next step is to enable them at the switch level and activate BGP EVPN peering. This ensures that each spine switch participates in the intersite control-plane exchange, allowing the fabrics to establish BGP sessions and advertise EVPN routes between the San Francisco and New York sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

14. In the same page, navigate to **Switch Settings**, then click **Edit** next to **San_Francisco – pod-1 – Spine1**

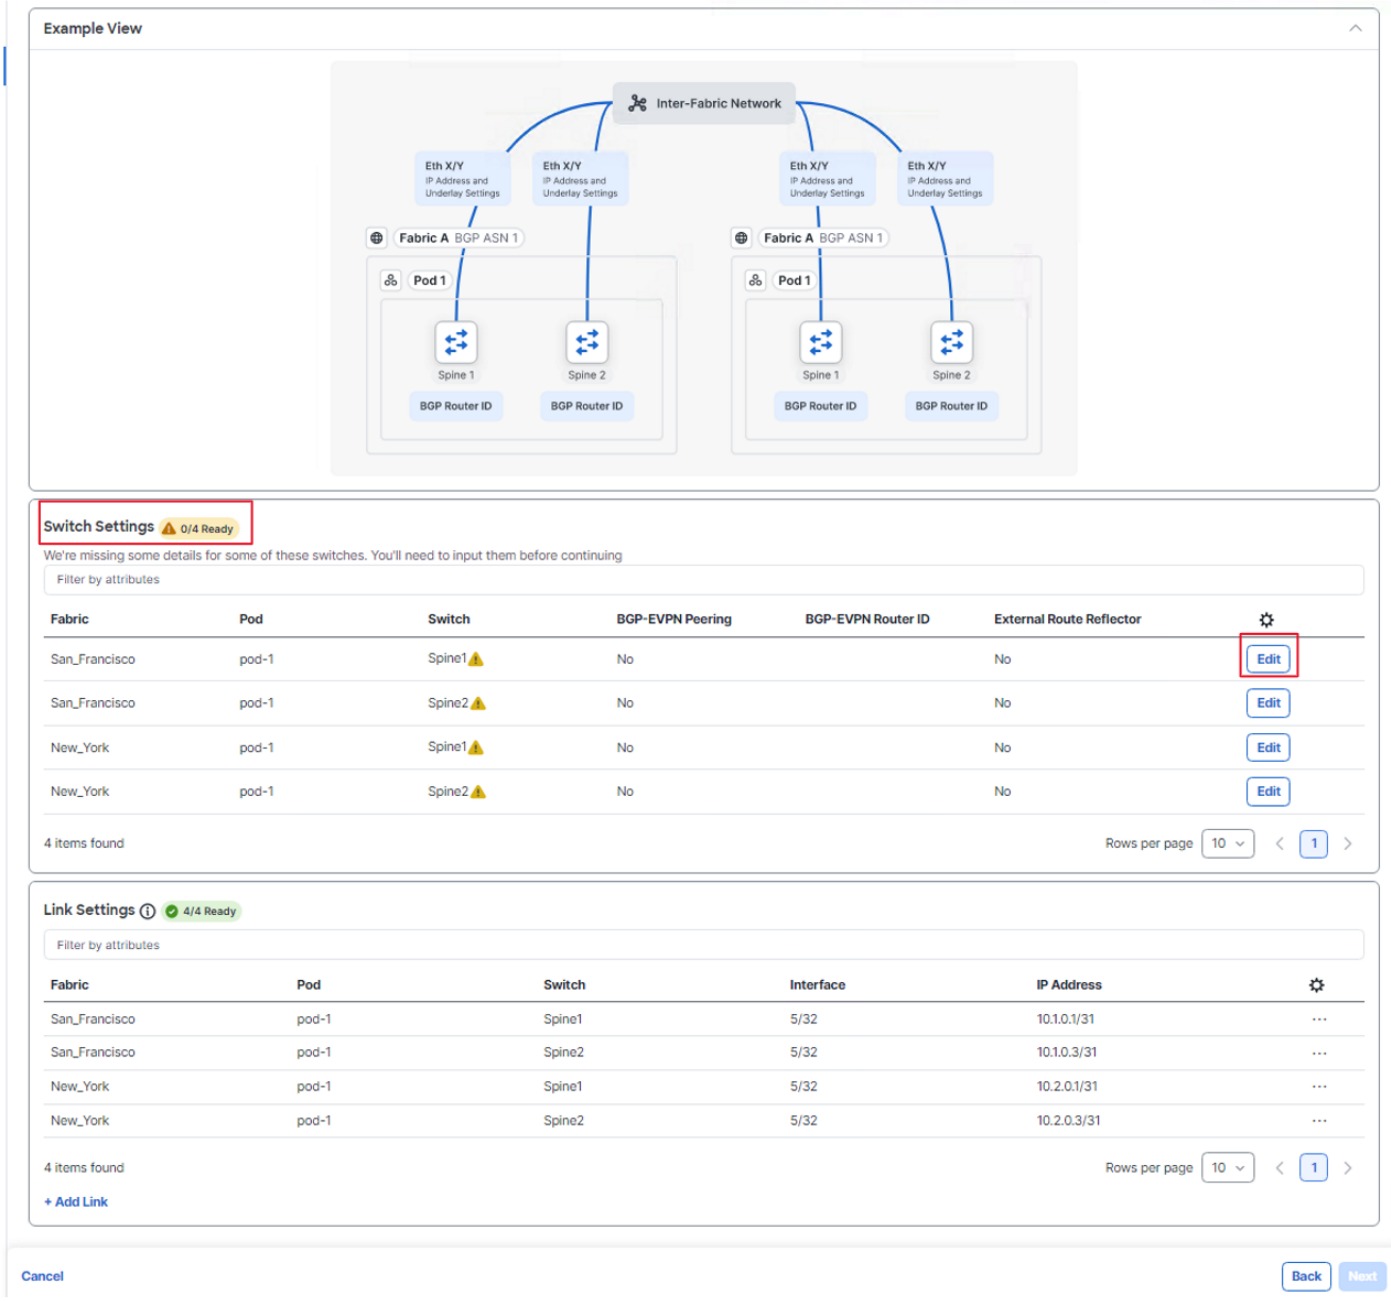

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

15. Click BGP-EVPN Peering to enable it, then enter 10.1.100.1 in the BGP-EVPN Router ID field.

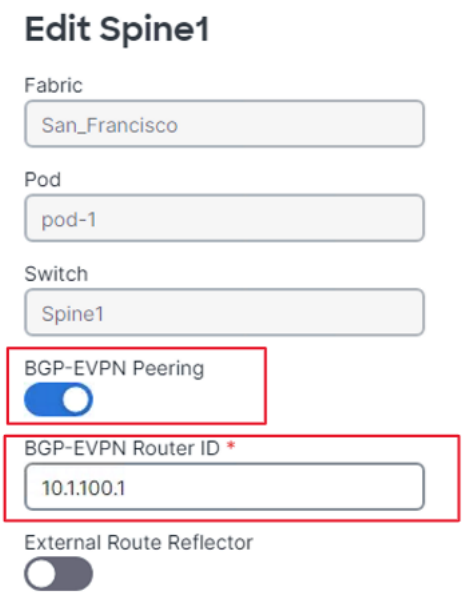

16. Repeat the same steps for the remaining three spines, updating only the **BGP-EVPN Router ID** field as follows:
    - San_Francisco – Spine2: **10.1.100.2**
    - New_York – Spine1: **10.2.100.1**
    - New_York – Spine2: **10.2.100.2**

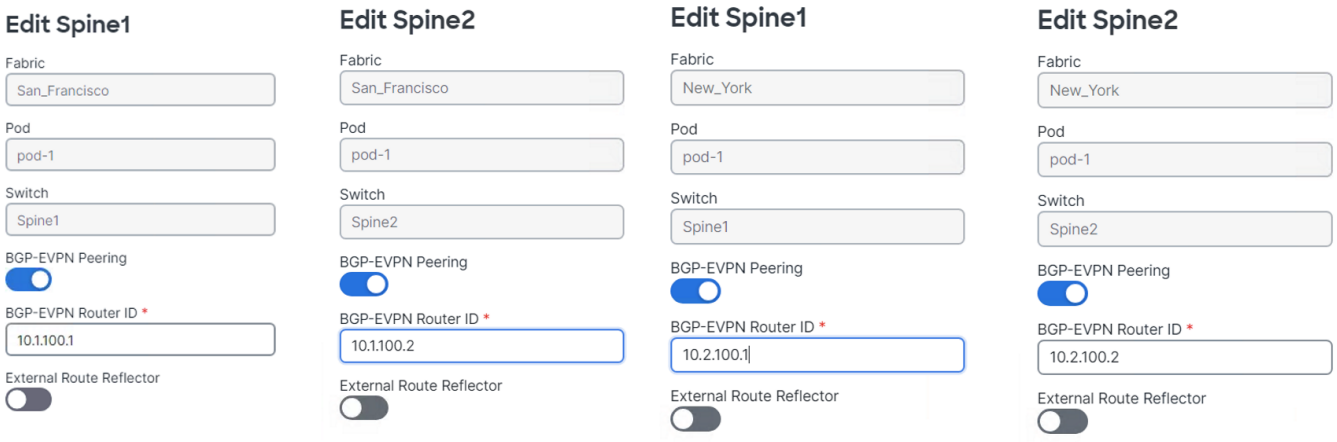

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

17. Now, verify that both **Link Settings** and **Switch Settings** show a **green (Ready)** status for all entries. Once confirmed, you will notice that the **Next** button becomes available. Click **Next** to proceed to the summary.

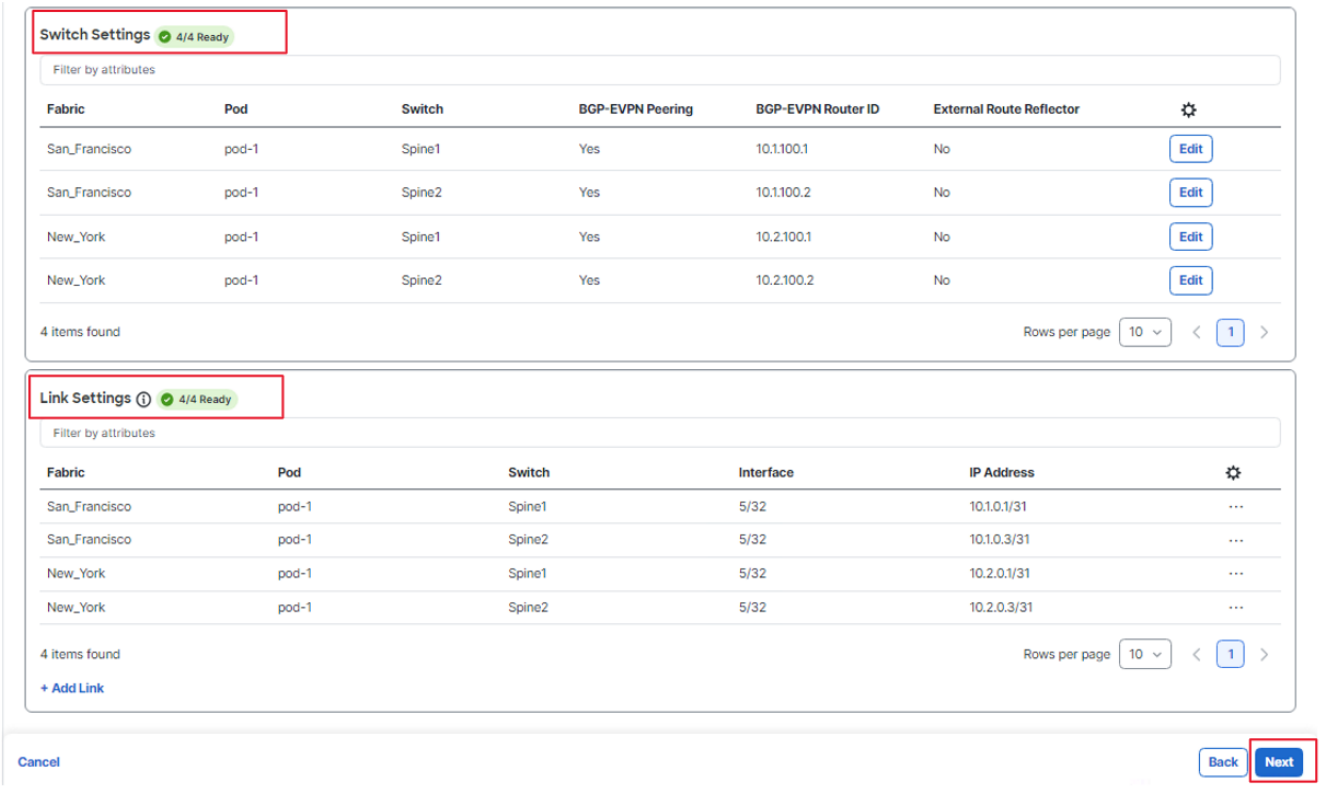

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

18. Review the **Summary** page to verify all configuration details, then click **Save** to proceed with deploying the configuration.

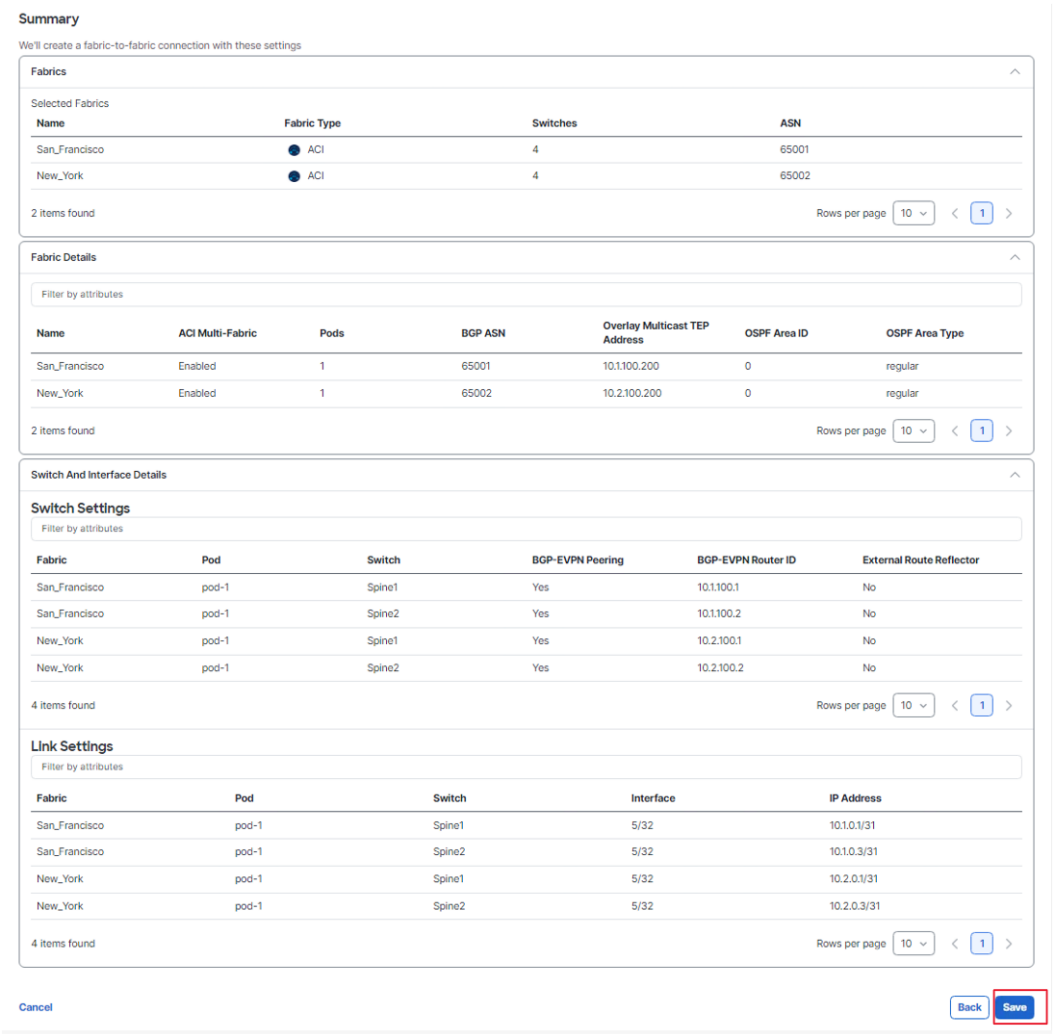

When routes are advertised to the physical spines from the intersite network (ISN), OSPF re-destribute those routes only to IS-IS. This redistribution happens when IS-IS is part of an existing tunnel endpoint (TEP) pool or part of a subnet that is configured under **Fabric Ext Connection Policies** in the APIC. Because of this behavior, spines that are not connected to the ISN cannot reach the vSpines. If your data plane TEP and vSpine router IDs are not part of the ACI site TEP pool, you must configure the Fabric Ext Connection Policies with those subnets.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

19. Add the Multipod Data Plane TEP configuration for San_Francisco and New_York sites as following:
    a. In the APIC window for the site being configured, click Tenants > infra > Policies > Protocol > Fabric Ext Connections Policies and click Fabric Ext Connection Policy.
    b. Enter **extended:as2-nn4:5:16** in the Community field.
    c. Click **Submit**.

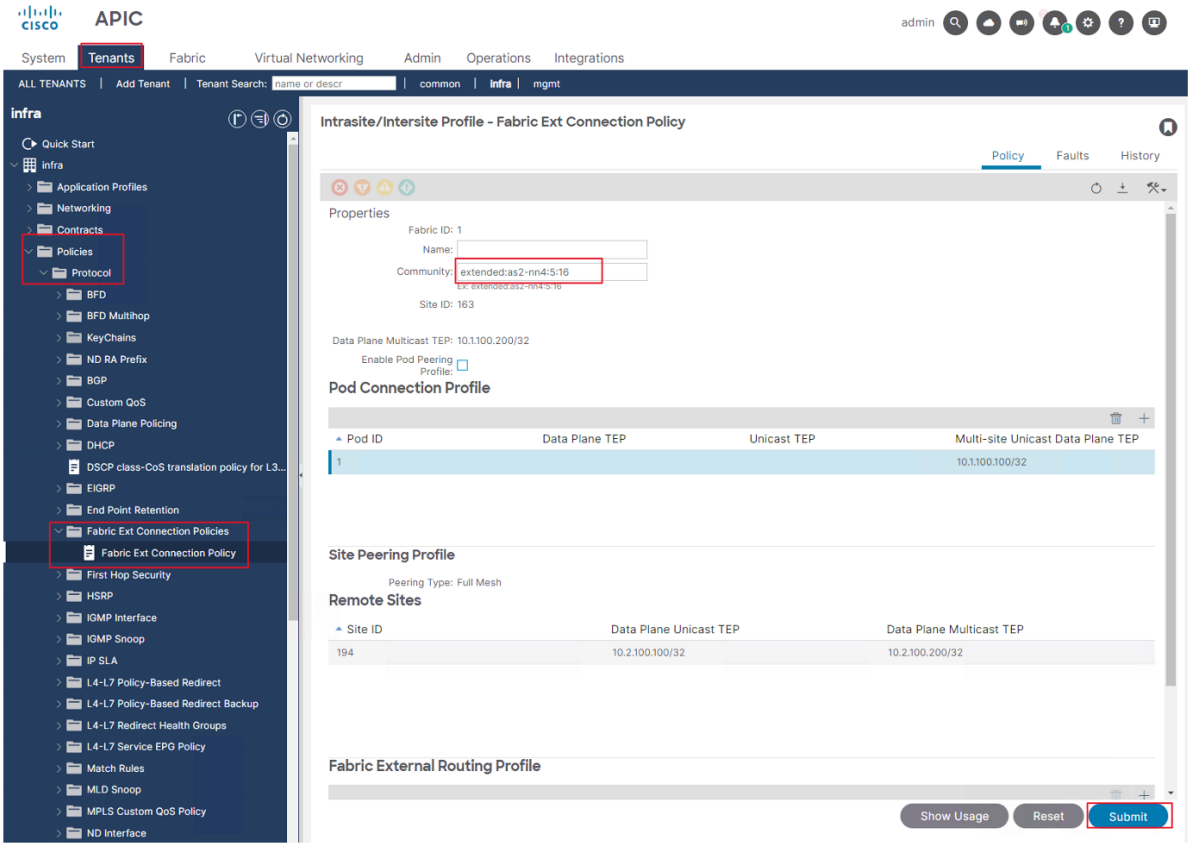

20. In the Pod Connection Profile, double-click **Pod ID 1**
21. Click the **(+)** icon to add a new subnet entry, then enter the **Data Plane TEP IP** of the site. After entering the value, click **Update**, and then **Close** the window to save the configuration.
    - San_Francsico Data Plane TEP IP: **10.1.200.200/32**
    - New_York Data Plane TEP IP: **10.2.200.200/32**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

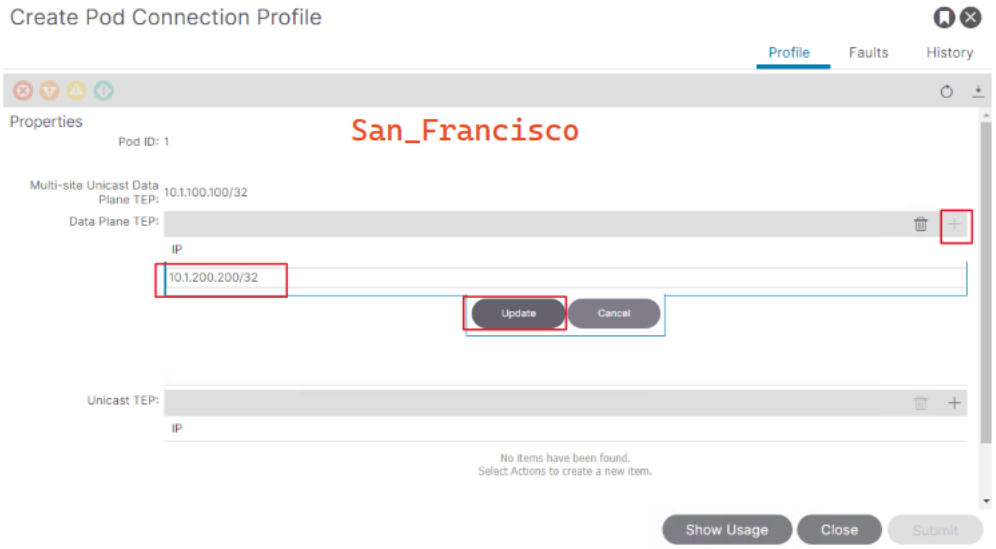

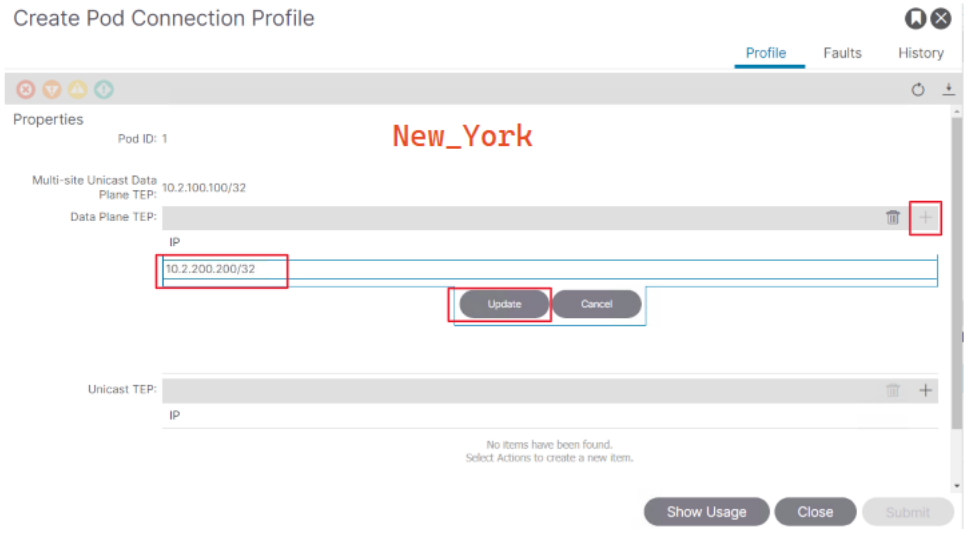

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

22. Navigate to **Tenant** > **infra** > **Networking** > **L3Outs** and verify that an L3Outs called **intersite** has been configured under the infra tenant. This indicates that Infra L3Out has been successfully configured.

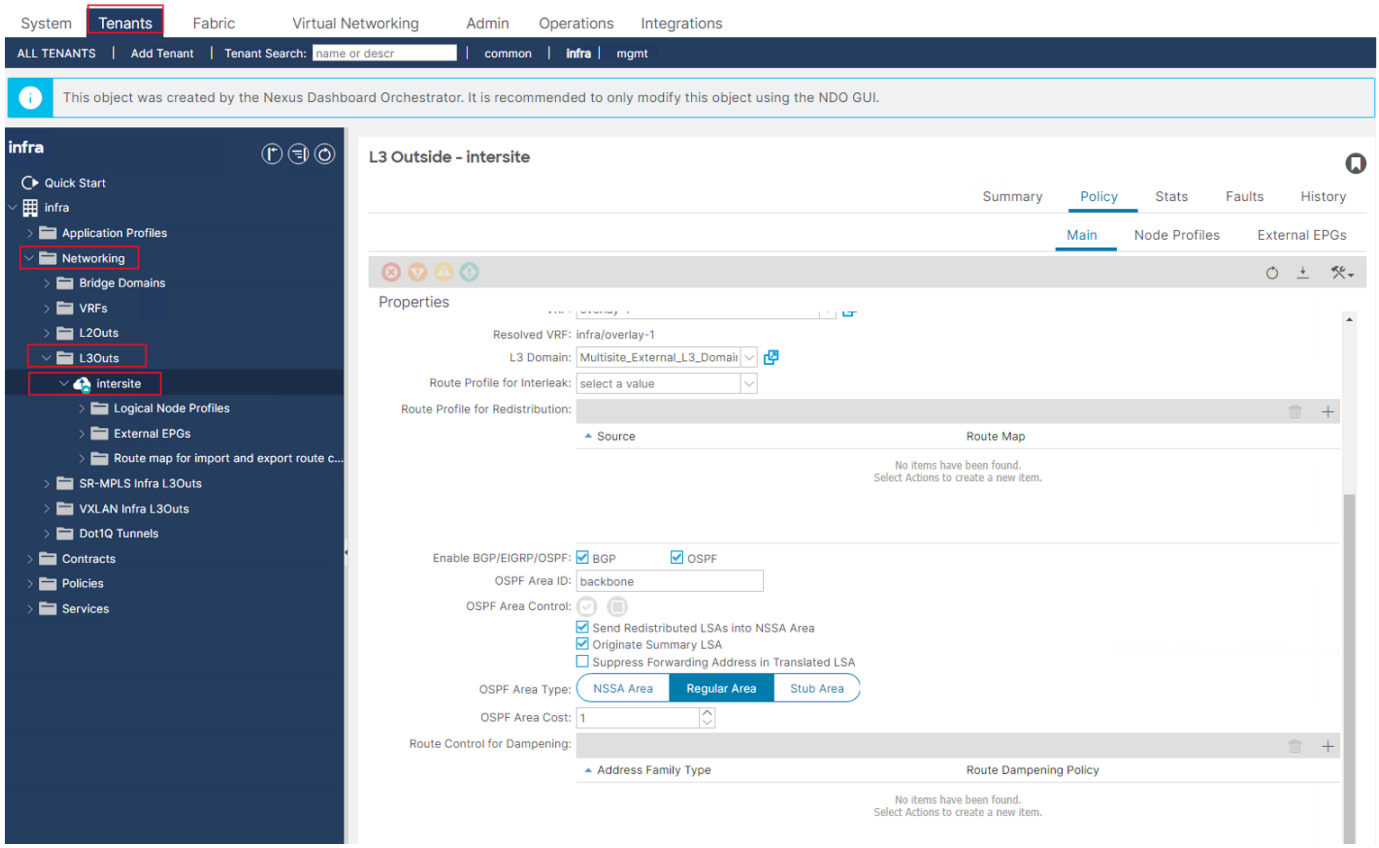

We have now successfully completed the **Inter-Fabric Connectivity** configuration, establishing the control-plane and data-plane links between the **San Francisco** and **New York** sites.

Value Proposition:

You have seen how Cisco Nexus Dashboard Orchestrator (NDO) simplifies the deployment of Inter-Fabric Connectivity by centralizing configuration and automation across multiple ACI sites. Through a unified interface, administrators can onboard fabrics, enable Multi-Fabric Connectivity, and configure Spine interfaces, BGP EVPN peering, OSPF underlay, and UTEP/MTEP endpoints in just a few steps.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 2. User Provisioning and NDO Tenant Creation

This scenario shifts focus to multi-site orchestration at the tenant level, demonstrating how Cisco Nexus Dashboard Orchestrator (NDO) streamlines the setup of new network environments. You will first create a local user with **Fabric Administrator** privileges, establishing the necessary access controls. Following this, a **greenfield tenant** will be defined entirely within NDO. This tenant, along with its associated VRF and other objects, will then be deployed simultaneously and consistently across both ACI sites.

The core objective here is to showcase NDO's capability for **centralized lifecycle management** of a tenant, eliminating the need for manual, per-fabric configuration and ensuring **policy uniformity** across geographically distributed environments from the outset.

Create a New User (**Optional**)

1. Navigate to **Admin > Users and Security**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

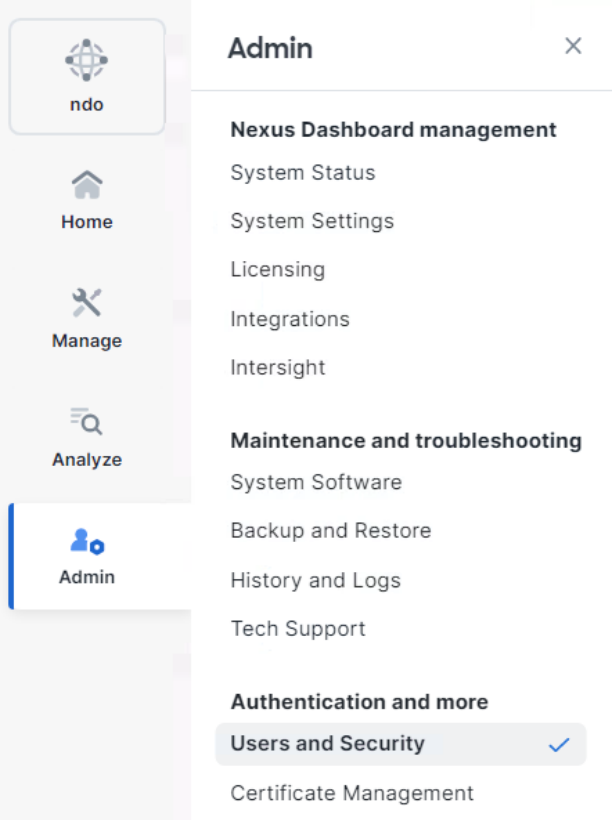

2. Navigate to **Users** > **Local**, expand the **Actions** menu, and click **Create Local User**.

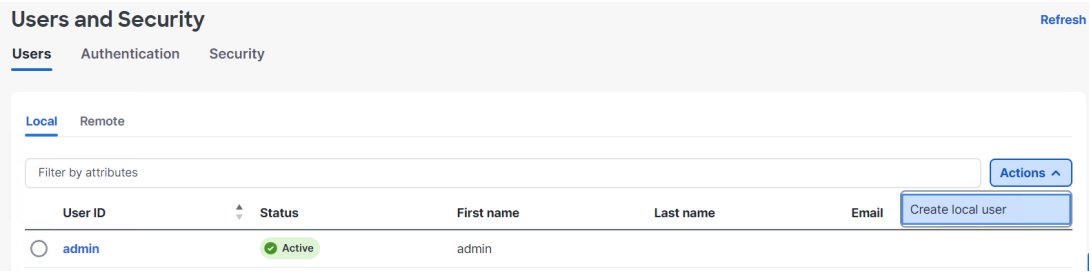

3. In the **Create Local User** configuration window, enter the following details:
   - User ID: **demouser**
   - Password: C1sco12345!
   - Confirm password: **C1sco12345!**
   - Email: demouser@corp.pseudoco.com
   - Security domain name: **All**
   - Role: choose **Fabric Administrator**
   - then click the check **(✓)** symbol to confirm the selection.
4. Click **Create**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Roles and permissions in Nexus Dashboard 4.1.1:

- **Approver**: When change control is enabled, users with this role can perform operations related to the approval or denial of submitted configuration changes.
- **Designer**: Users with this role can make changes to the configuration (the intent) but cannot deploy those changes to the network fabrics.
- **Fabric Administrator:** Users with this role have permissions for full fabric management, including creation of network policies, interface configurations and software upgrades. Users with this role and the all security domain access can also add and delete Nexus Dashboard and APIC clusters. See Connecting Clusters for more information.
- **Observer:** Primarily a read-only user.
- **Super administartor:** Users with this role have full access to all resources in every security domain. A user with this role can perform all operations on the Nexus Dashboard, including backup and restore processes.
- **Support Engineer:** Users with this role are able to perform tasks related to support, such as collecting tech support, creating backups, deploying templates, and general troubleshooting. When change control is enabled, users with this role can also deploy and revert policy changes that were approved. Users with this role are not able to make configuration changes.

Create an NDO Tenant

1. In Nexus Dashboard, navigate to **Manage** > **Orchestration** > **Tenants**.

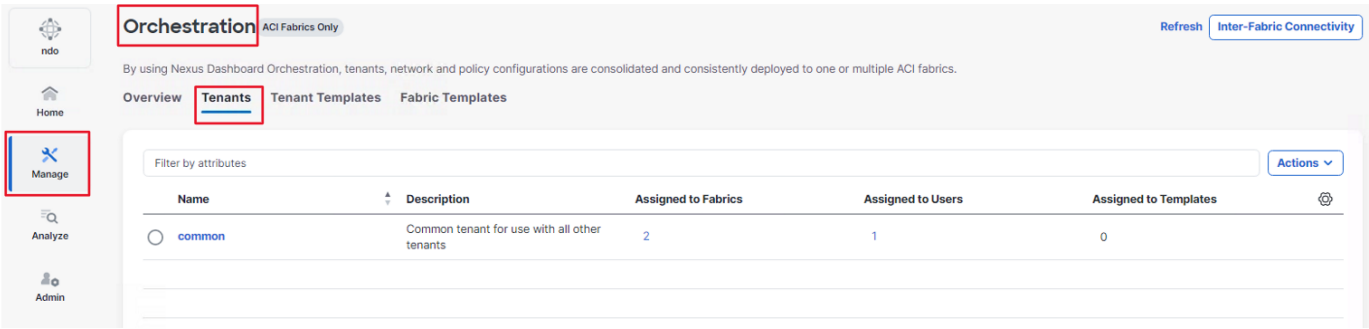

2. Expand the **Actions** menu on the right-hand side and select **Create Tenant**.

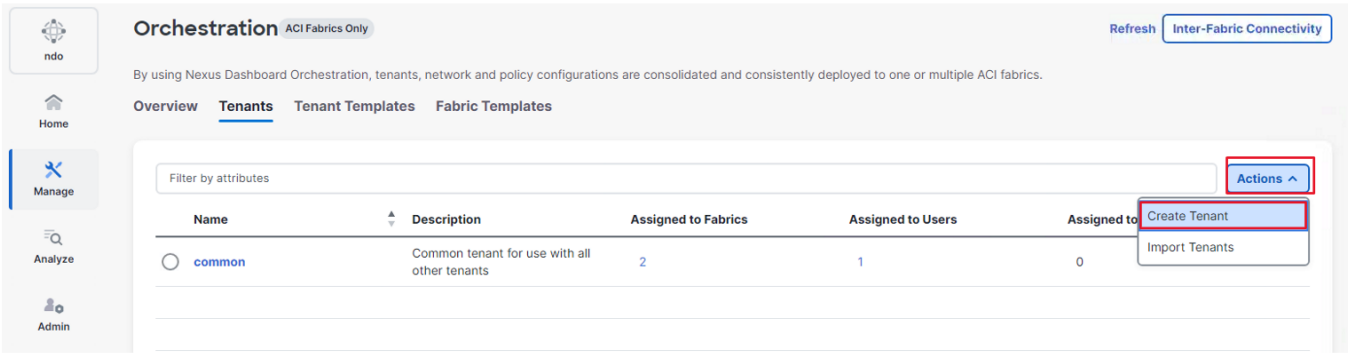

3. Enter **Pseudoco** in the Display Name field.
4. Select both **San_Francisco** and **New_York** and click **Save** to push the Tenant configuration to APIC.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

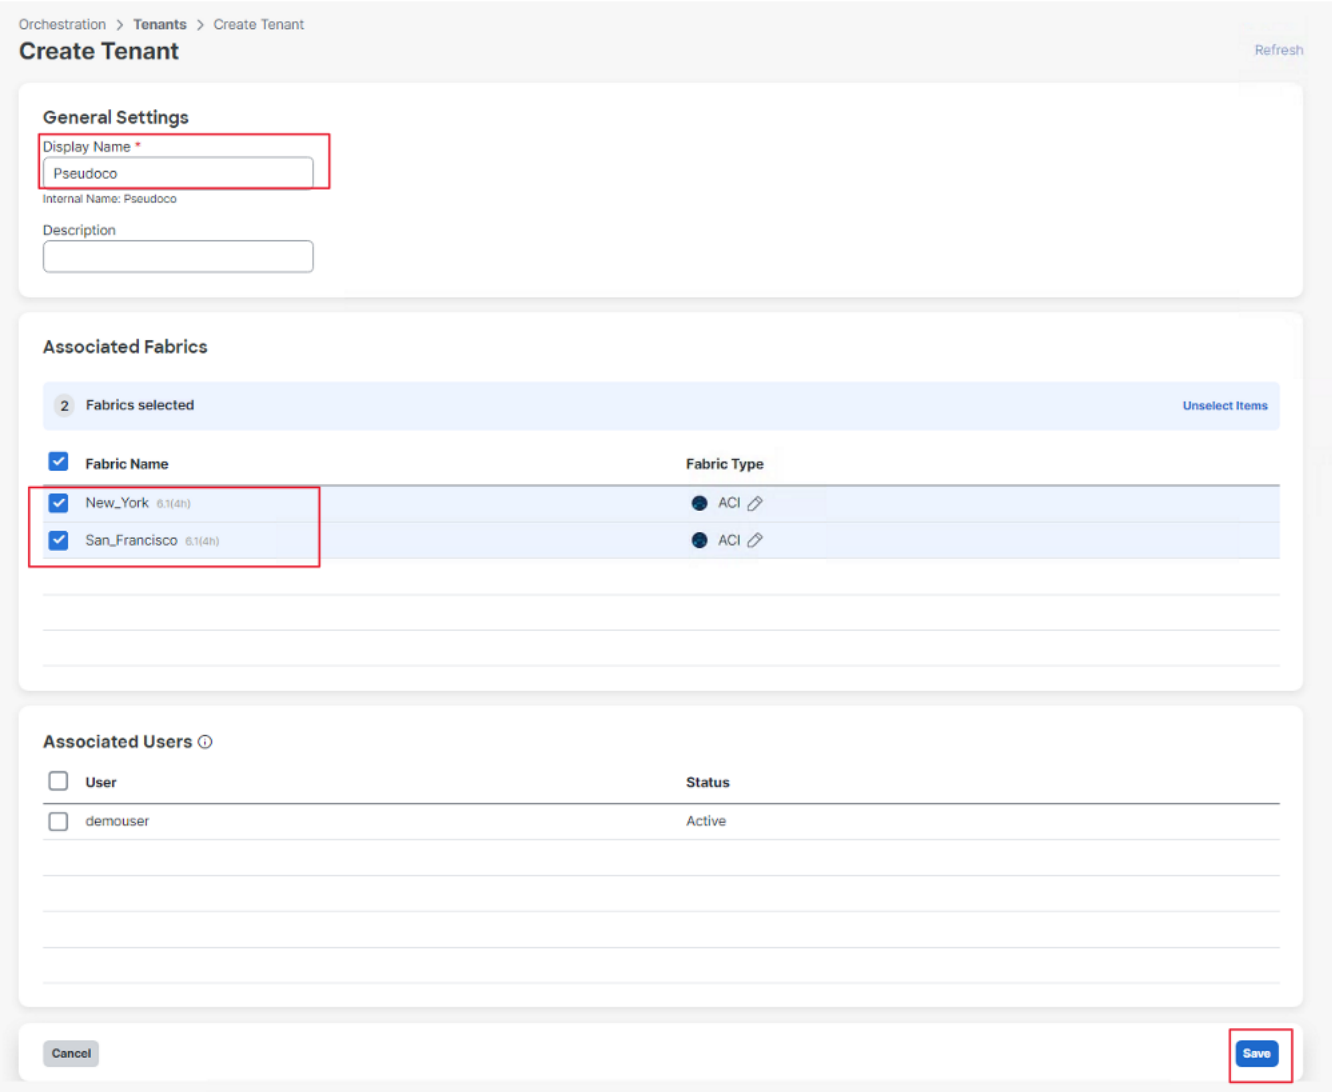

5. Return to the **San_Francisco** and **New_York** APICs. (If APIC is not open, click Google Chrome bookmark bar for each site and login with **admin/C1sco12345**). Click **Tenants** > ALL TENANTS in each window and verify that the **Pseudoco** tenant is created on both fabrics.

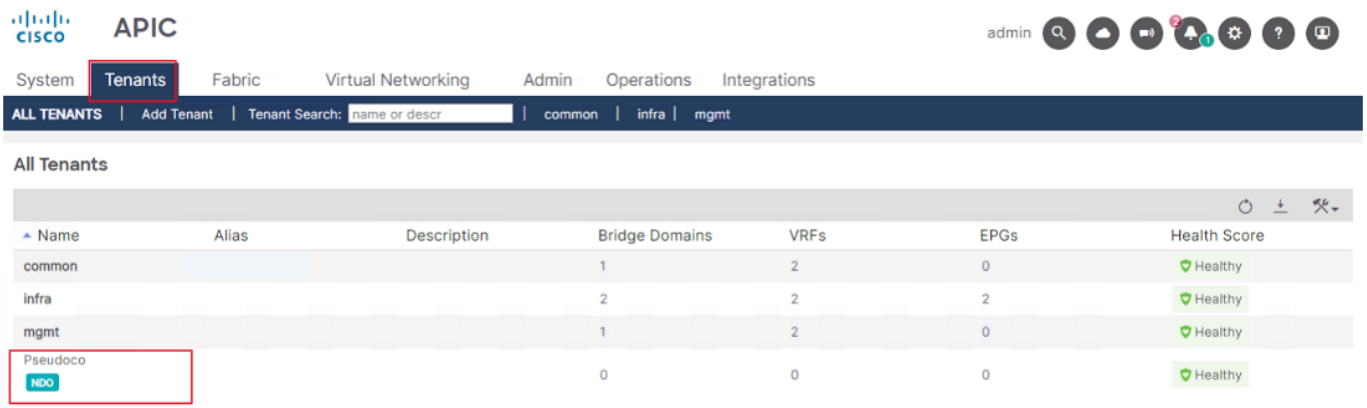

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5. In either **San_Francisco** and **New_York** APICs, double-click **Pseudoco** to proceed to the APIC window for Pseudoco. The tenant object includes the cloud symbol, indicating that this object has been configured from the NDO. The APIC GUI will also display a message to this effect.

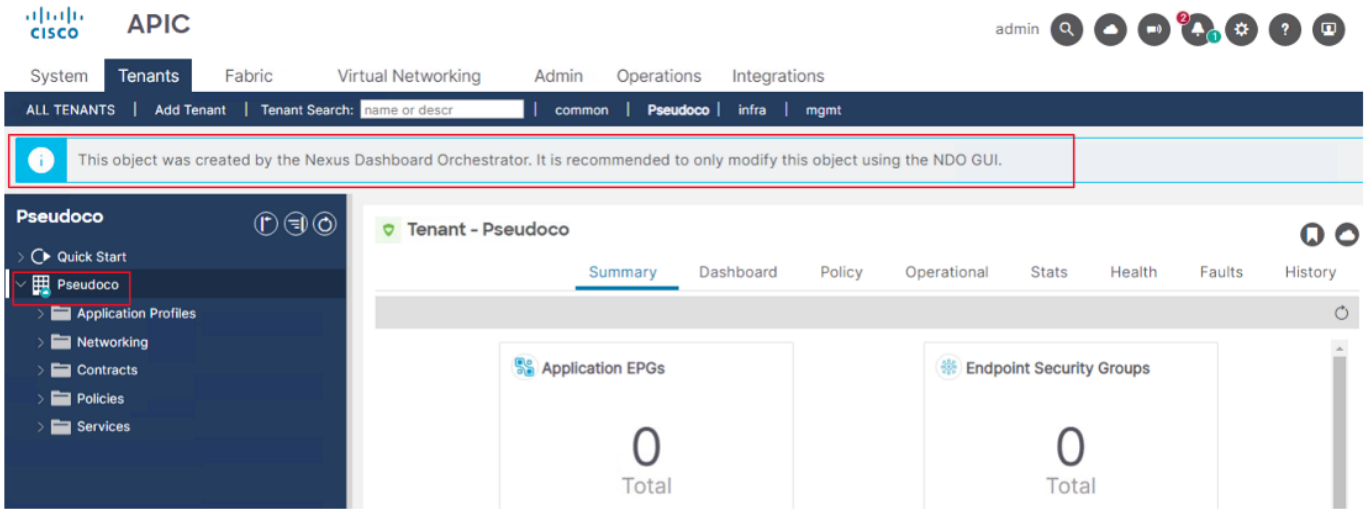

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 3. Layer 3-Only Communication across Sites (Intra Tenant and Intra VRF)

This scenario demonstrates inter-site Layer 3 connectivity within a single tenant and VRF context using Cisco Nexus Dashboard Orchestrator. In this deployment model, tenant and VRF constructs are stretched across both geographic sites (San Francisco and New York), enabling unified routing and policy enforcement. However, Bridge Domains (BDs), subnets, and Endpoint Groups (EPGs) remain site-local, meaning each site maintains its own independent Layer 2 forwarding domains and IP address space.

This architectural approach allows application workloads deployed within the same tenant and VRF to communicate seamlessly across sites via Layer 3 routing, while preserving site-specific broadcast domains. By keeping Bridge Domains local to each site, broadcast, unknown unicast, and multicast (BUM) traffic is contained within geographic boundaries, preventing unnecessary flooding across the inter-site network (ISN). This design optimizes This use case establishes routed communication between endpoints in different ACI sites that are part of the same tenant. The core principle is to create a unified Layer 3 routing domain that spans all sites, while keeping all Layer 2 broadcast domains and application endpoint groups strictly local to their respective physical sites.

The configuration is achieved by logically separating objects within Nexus Dashboard Orchestrator (NDO) into two categories:

Stretched Objects: These objects are defined once in a shared template and are deployed consistently across multiple sites. In this use case, the Tenant and VRF are stretched, ensuring a common policy and routing namespace across the entire Multi-Site domain.

Site-Local Objects: These objects are defined in site-specific templates and are only instantiated within a single designated site. For this scenario, the Bridge Domain (BD), Subnet, and Application EPGs are configured as site-local. This design is critical for isolating Layer 2 fault domains; issues like a broadcast storm in one site will not propagate across the WAN to other sites.

Communication between the site-local EPGs is enabled by a Contract which is also defined as a stretched object in the shared template. This allows NDO to orchestrate the necessary policy enforcement and route advertisements between sites.

What is a Schema and Template?

A Schema in Nexus Dashboard Orchestrator serves as a logical container and organizational structure for configuration templates. It acts as a repository that holds one or more templates, which in turn contain the actual policy objects (tenants, VRFs, bridge domains, EPGs, contracts, etc.) to be deployed across ACI sites. Within each schema, a Template is a child object that contains the actual configuration objects and policies to be deployed to one or more ACI sites. Templates define the scope and association of configuration elements, determining which sites will receive specific policy configurations.

This hierarchical structure exists to provide flexible and scalable multi-site orchestration. Schemas enable logical separation and organization of policies by tenant, application, or business unit, while templates within those schemas allow granular control over where specific configurations are deployed.

This architecture supports multiple deployment patterns: stretched templates can deploy identical objects across multiple sites for consistent policy enforcement and inter-site communication, while site-local templates contain configurations specific to individual locations, allowing independent addressing schemes and physical connectivity.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

By separating the logical organization (schema) from the deployment scope (template), Nexus Dashboard Orchestrator provides administrators with the flexibility to manage complex multi-site environments efficiently, maintain clear configuration boundaries, enable cross-template object references, and support both greenfield deployments and brownfield migrations where existing site-local configurations must coexist with newly stretched policies.

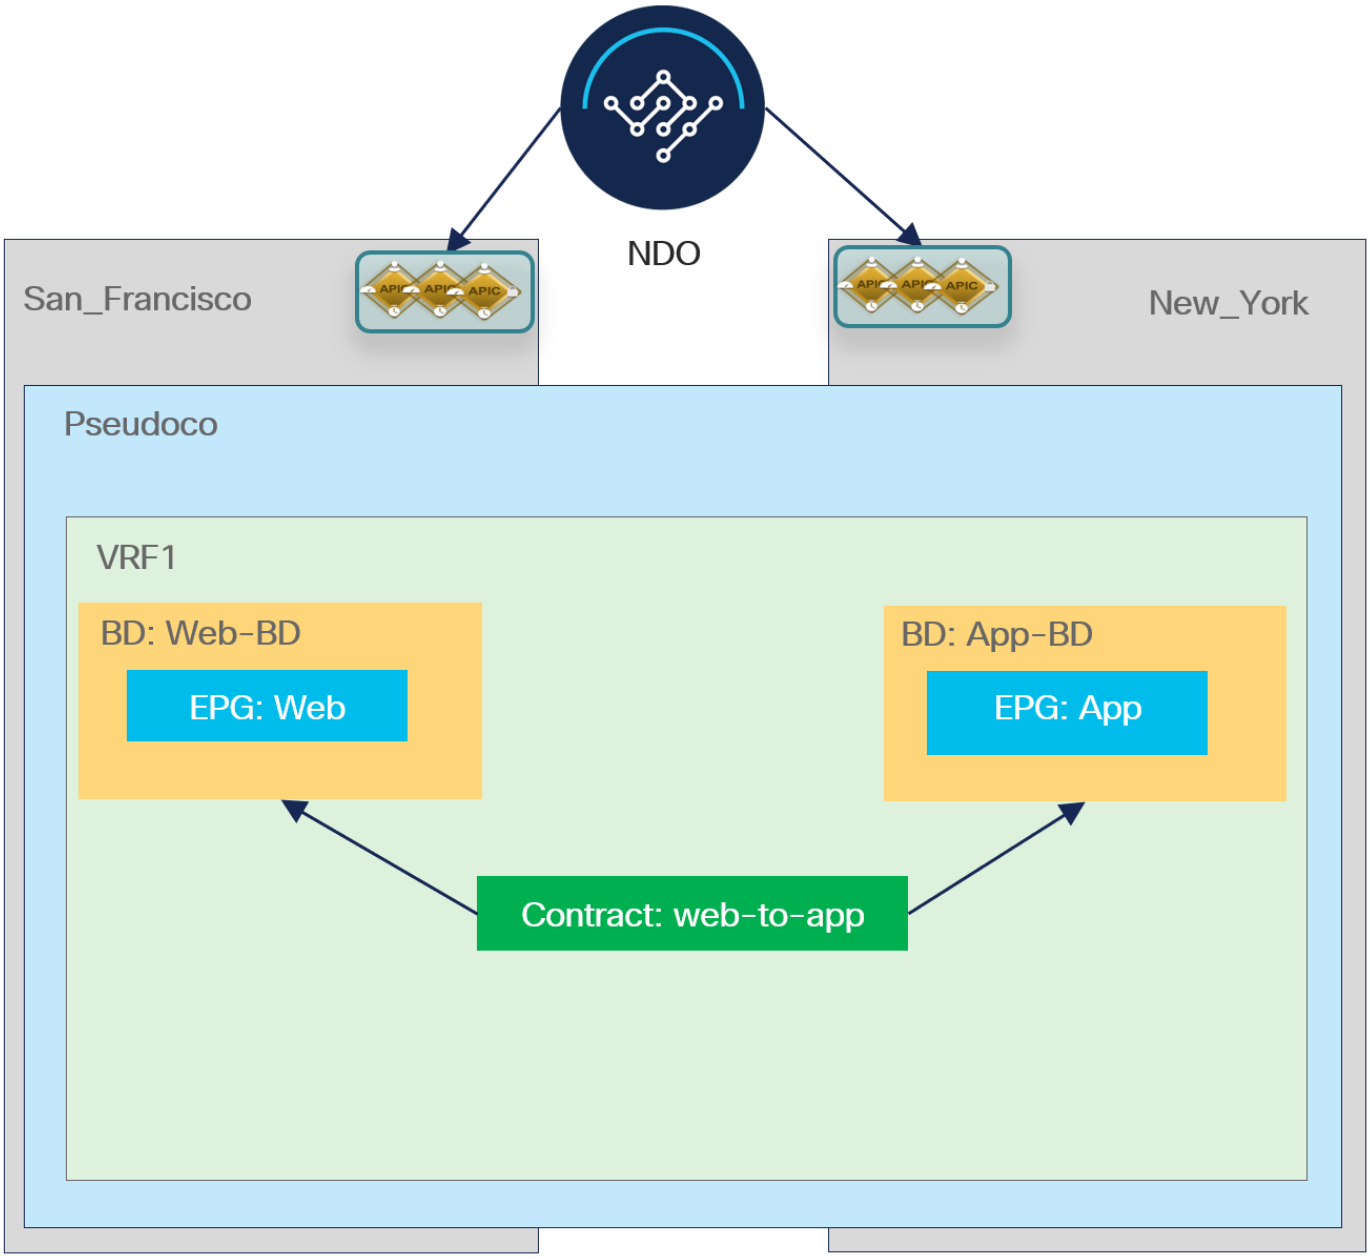

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

NOTE: After modifying a template, ensure that you click **Save Scheme** to retain the policy configuration.

1. Open a web browser and log in to Nexus Dashboard, then navigate to **Manage** > **Orchestration** > **Tenant Templates** > **Applications**, expand the **Actions** menu in the upper-right corner, and click **Create Schema**.

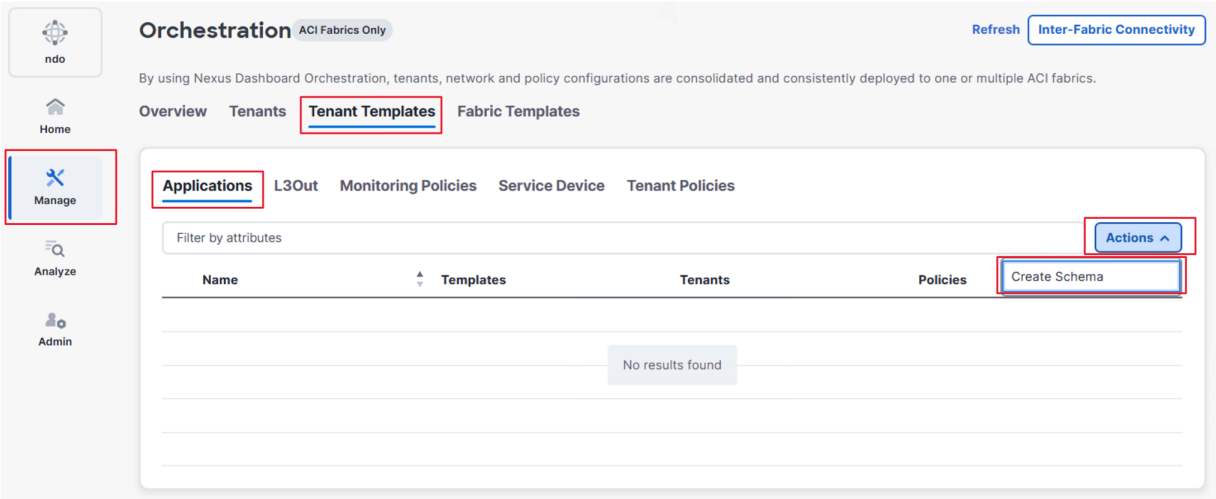

2. Enter **L3-stretch-schema** as the schema name in the Display Name field, then click **Add** to create the schema.

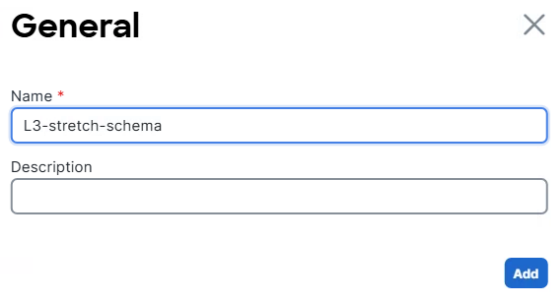

3. Click **Create New Template** to begin adding a template to the schema.

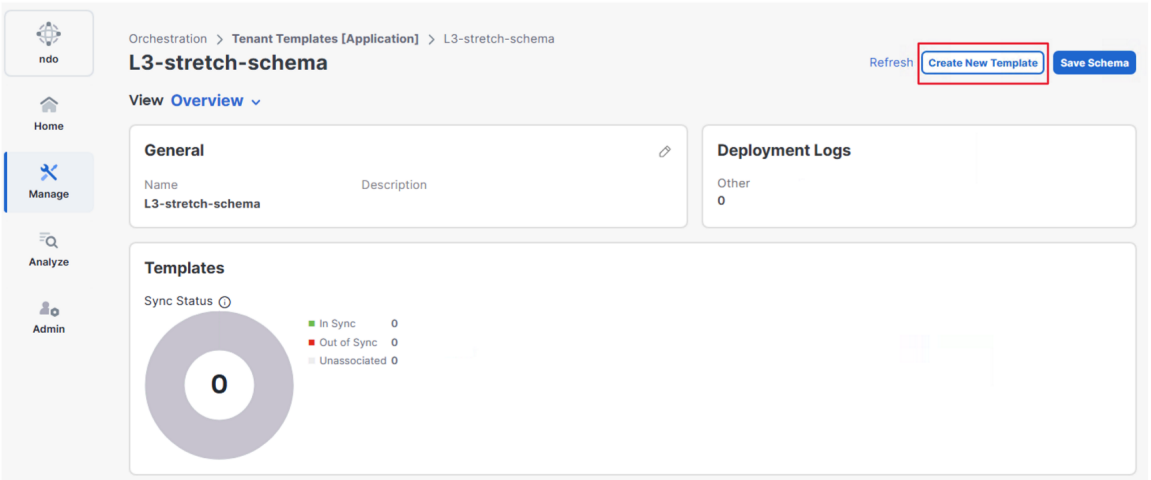

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

4. Enter **SF and NY Template** in the Display Name field, select **Pseudoco** from the Select a Tenant dropdown menu, ensure **Multi-Fabric** is selected in the Deployment Mode option, then click **Next**, and click **Continue to Template**.

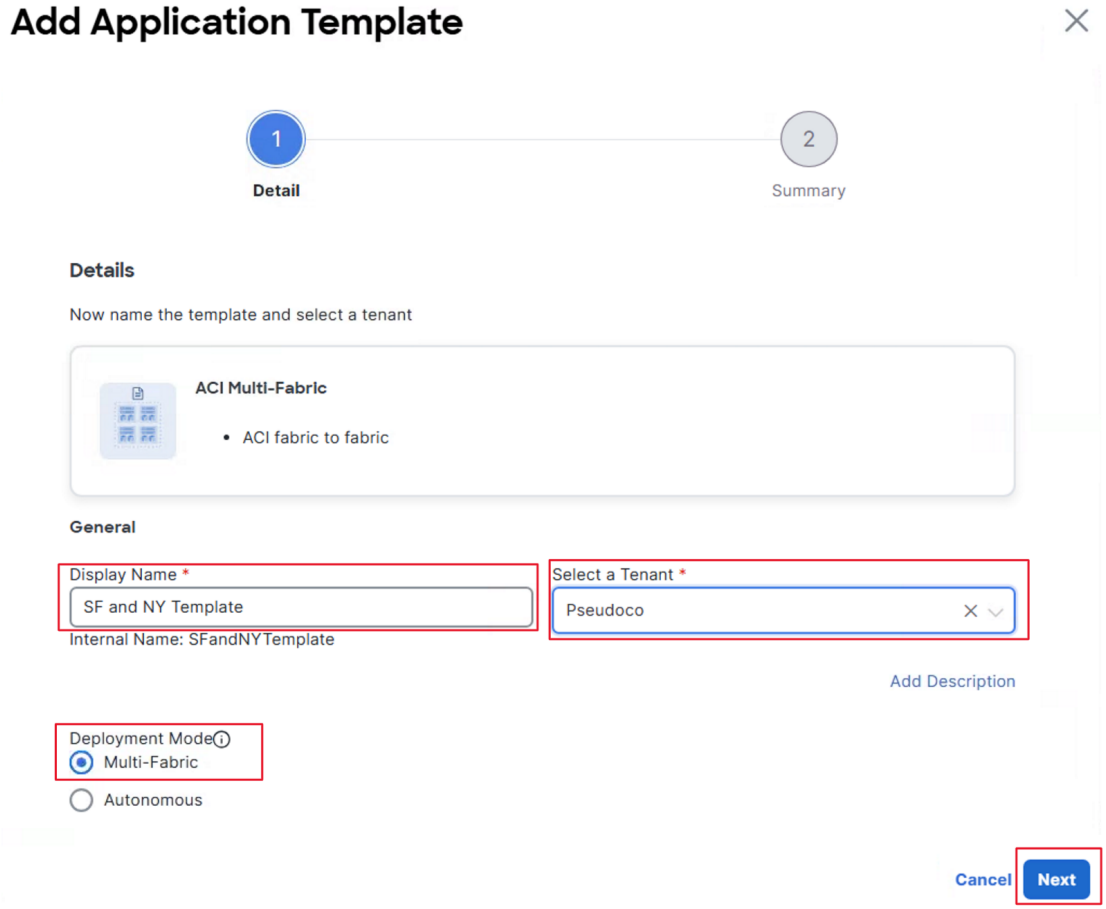

Multi-Fabric - The template can be associated with a single fabric (fabric-local policies) or multiple fabrics (stretched policies), and the option should be selected for Multi-Fabric Network (ISN) or VXLAN inter-fabric communication.

Autonomous - The template can be associated with one or more fabrics that are operated independently and are not connected via Inter-Fabric Network (no inter-fabric VXLAN communication).

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5. Expand **Create Object** and choose VRF.

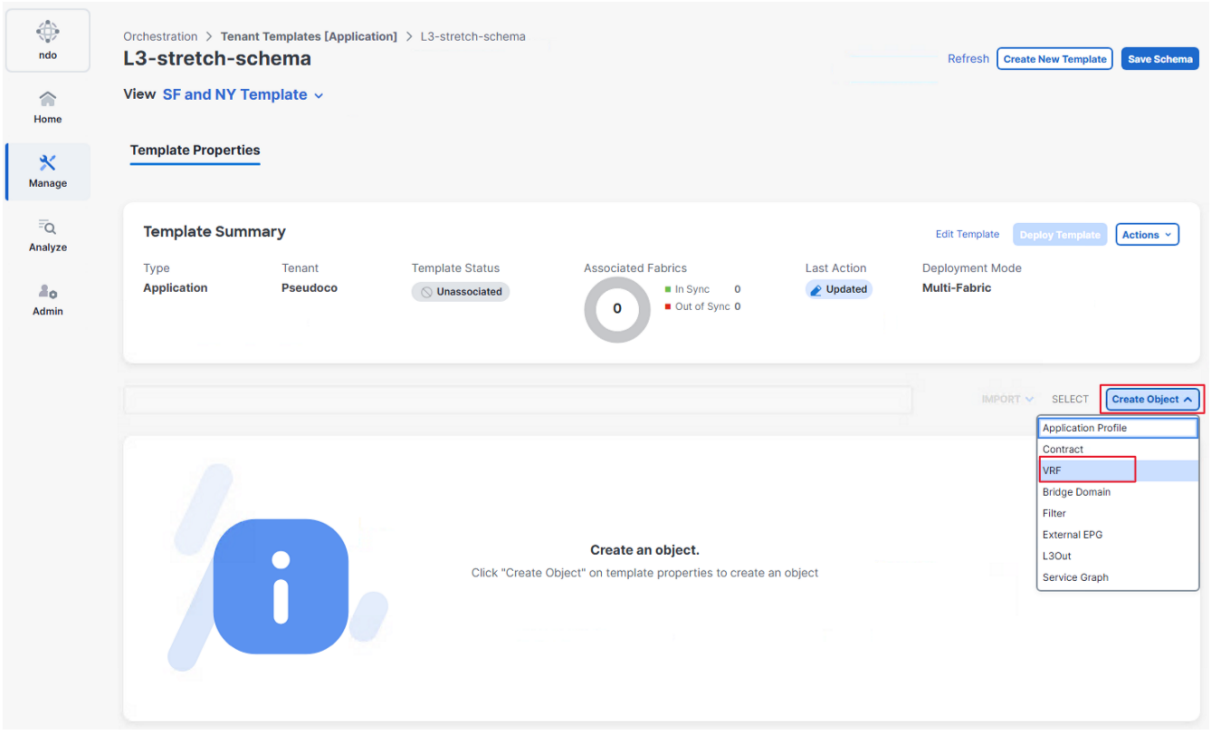

6. In the Display name field, enter **VRF1**, and click **OK**.

Because the objects in this template will be stretched across both sites, the template must be associated with both San_Francisco and New_York sites. Stretched objects are configuration elements (such as VRFs, Application Profiles, EPGs, and Contracts) that exist identically in multiple fabrics and enable inter-site communication. When a template is associated with multiple sites, any objects created within that template will be automatically deployed to all associated sites during template deployment, ensuring consistent policy configuration and enabling seamless workload mobility and communication across geographic locations.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

7. Expand **Actions** and click **Add/Remove Fabrics**.

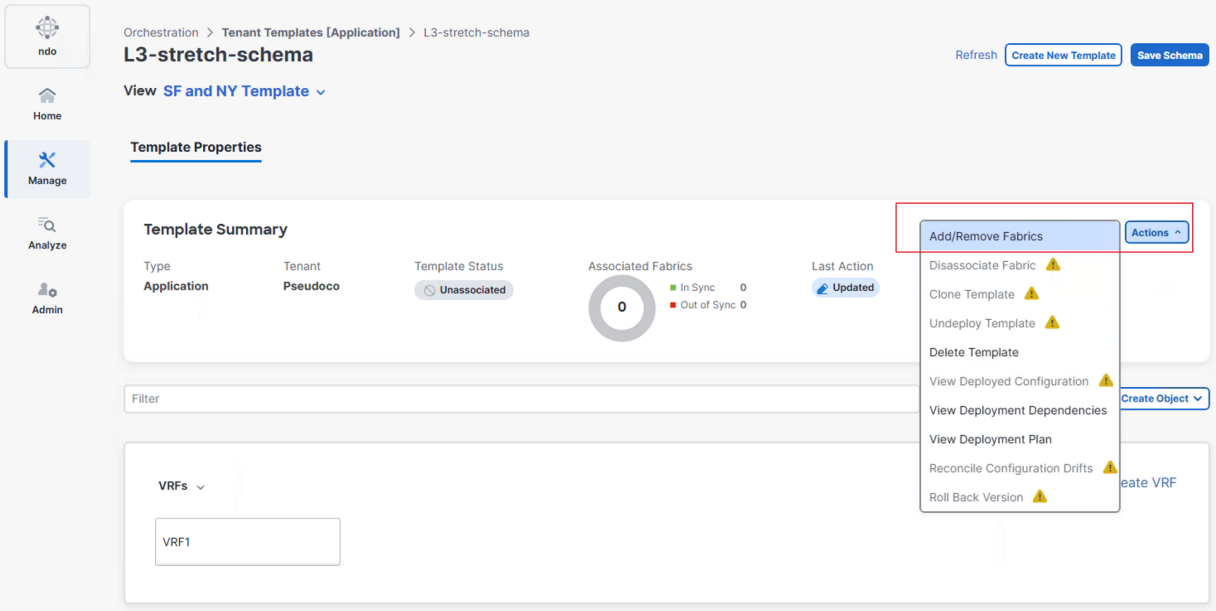

8. Select **New_York** and **San_Francisco**. Click **Ok**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

NOTE: When configuration is added to the NDO there is a Save Schema button and a Deploy to sites button. Saving the template configuration saves it to the NDO database but does not make any changes to the APICs. Only after selecting Deploy to sites is the configuration change pushed to the APICs. At this point in the configuration, we have added a template and created a VRF but have not saved nor deployed the configuration.

9. Click **Save Schema** to save the configuration to the NDO without deploying to APIC.

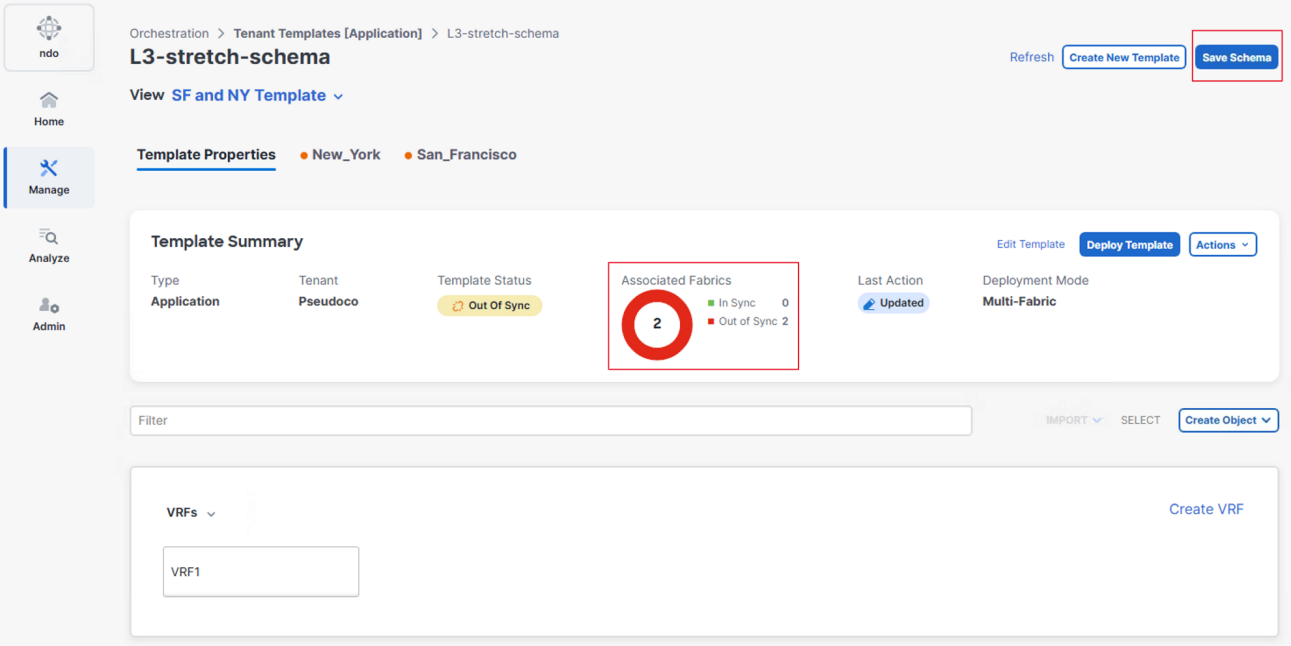

10. Click **Deploy Template** to deploy the configuration to the two sites. A window appears showing which changes will be deployed and to which sites, then click **Deploy**.

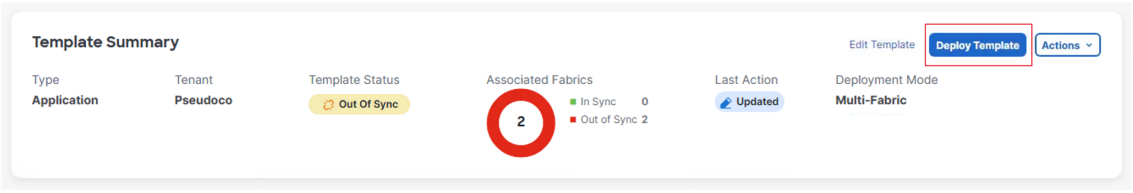

11. Another pop-up window will appear indicating Out-of-Sync Templates. Review the out-of-sync templates, then click **Deploy Out of Sync Templates** to proceed with deployment.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

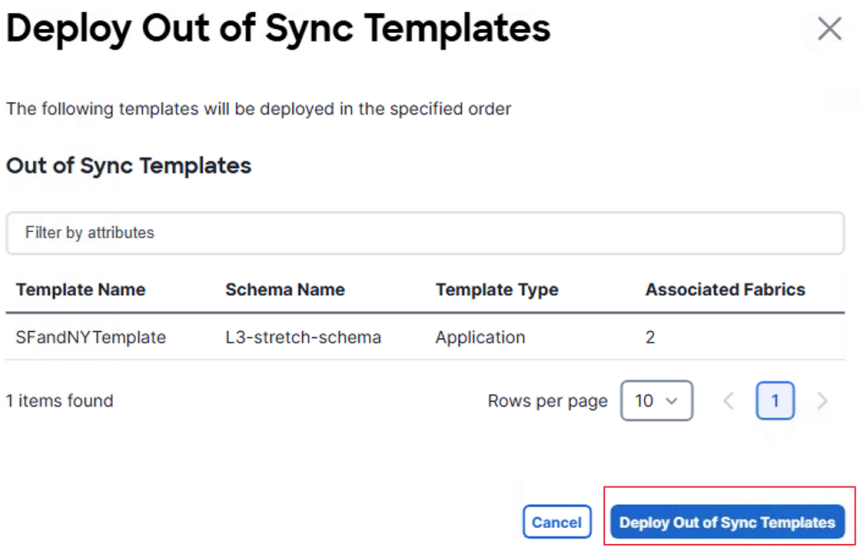

Now that we have configured the stretched VRF across the two sites, the next step is to create the site-local Application Profile, Bridge Domain (BD), and End Point Group (EPG).

12. Click on **Create New Template**, enter in the Display Name **SF Only**, Select **Pseudoco** as the **Tenant,** Click **Next** then click o**n Continue to template**.
13. Expand **Create Object** and select **Application Profile**.

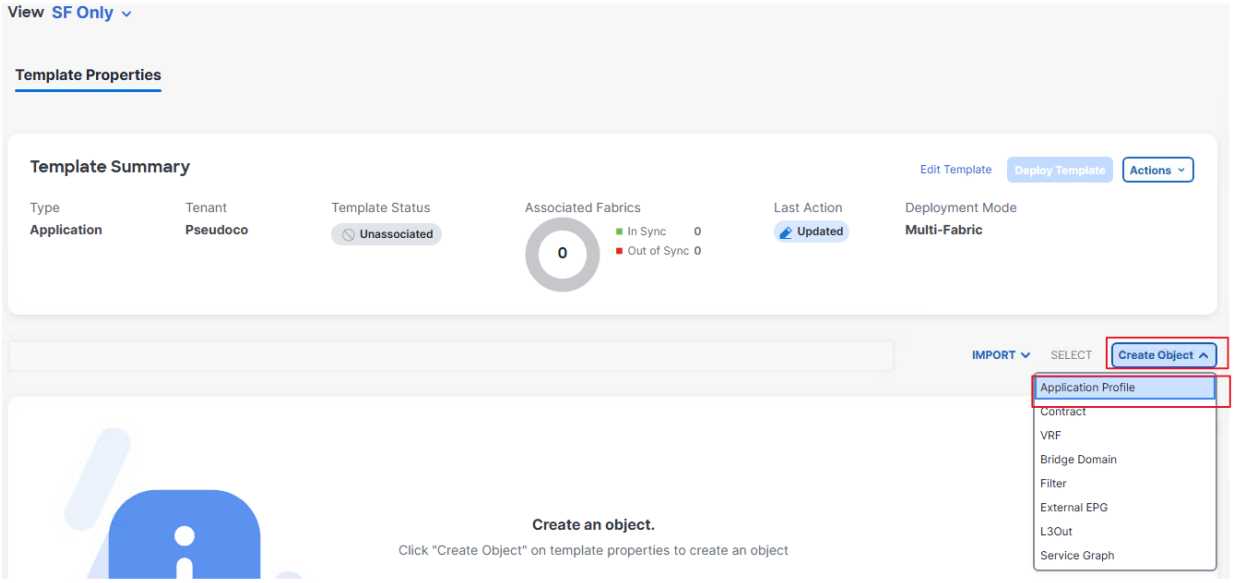

14. In the Display Name field, enter **Webapp**, and select **OK.**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

15. Expand **Create Object** and select **Bridge Domain** and configure the following:
    - Display Name field: **Web-BD**
    - Virtual Routing & Forwarding select: VRF1 (the VRF created in the SF and NY Template)
    - Uncheck L2 Stretch
    - Click **Ok**.

> When the L2STRETCH box is unchecked the option to add a BD subnet is removed. This is because the BD becomes a site local configuration. The site local configuration will be covered in a few more steps

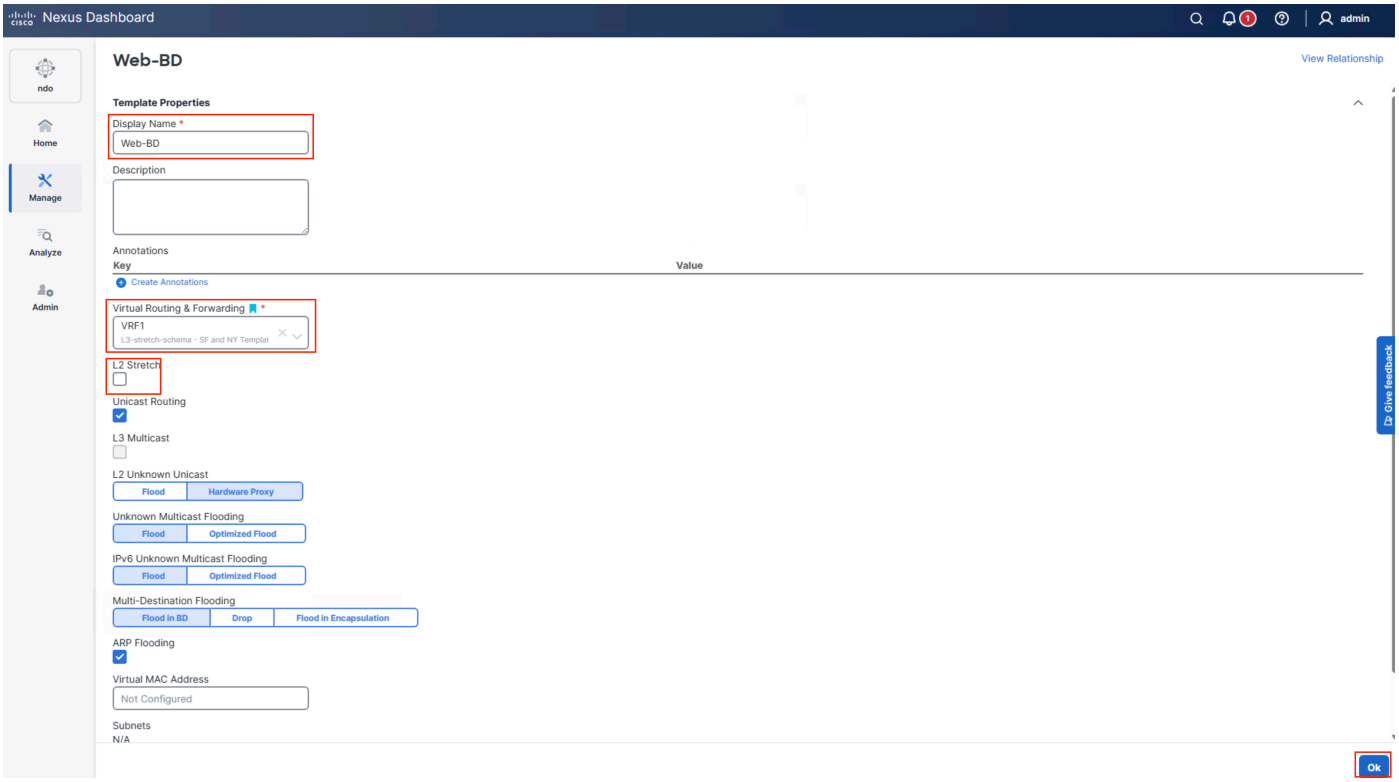

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

16. Click **Create Object > EPG**.

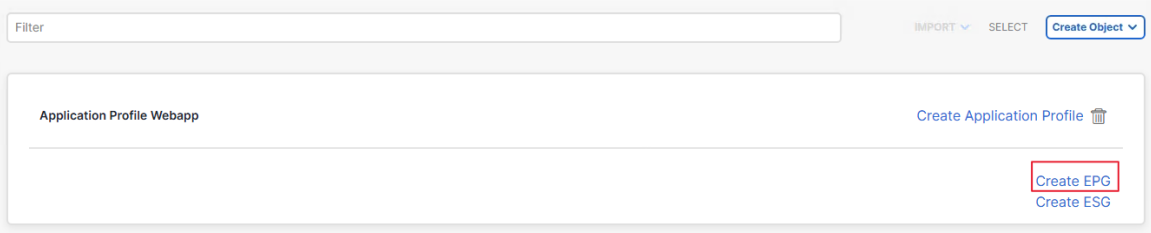

17. Enter **Web** in the Display Name field, and in Bridge Domain select **Web-BD**. Click **Ok**.

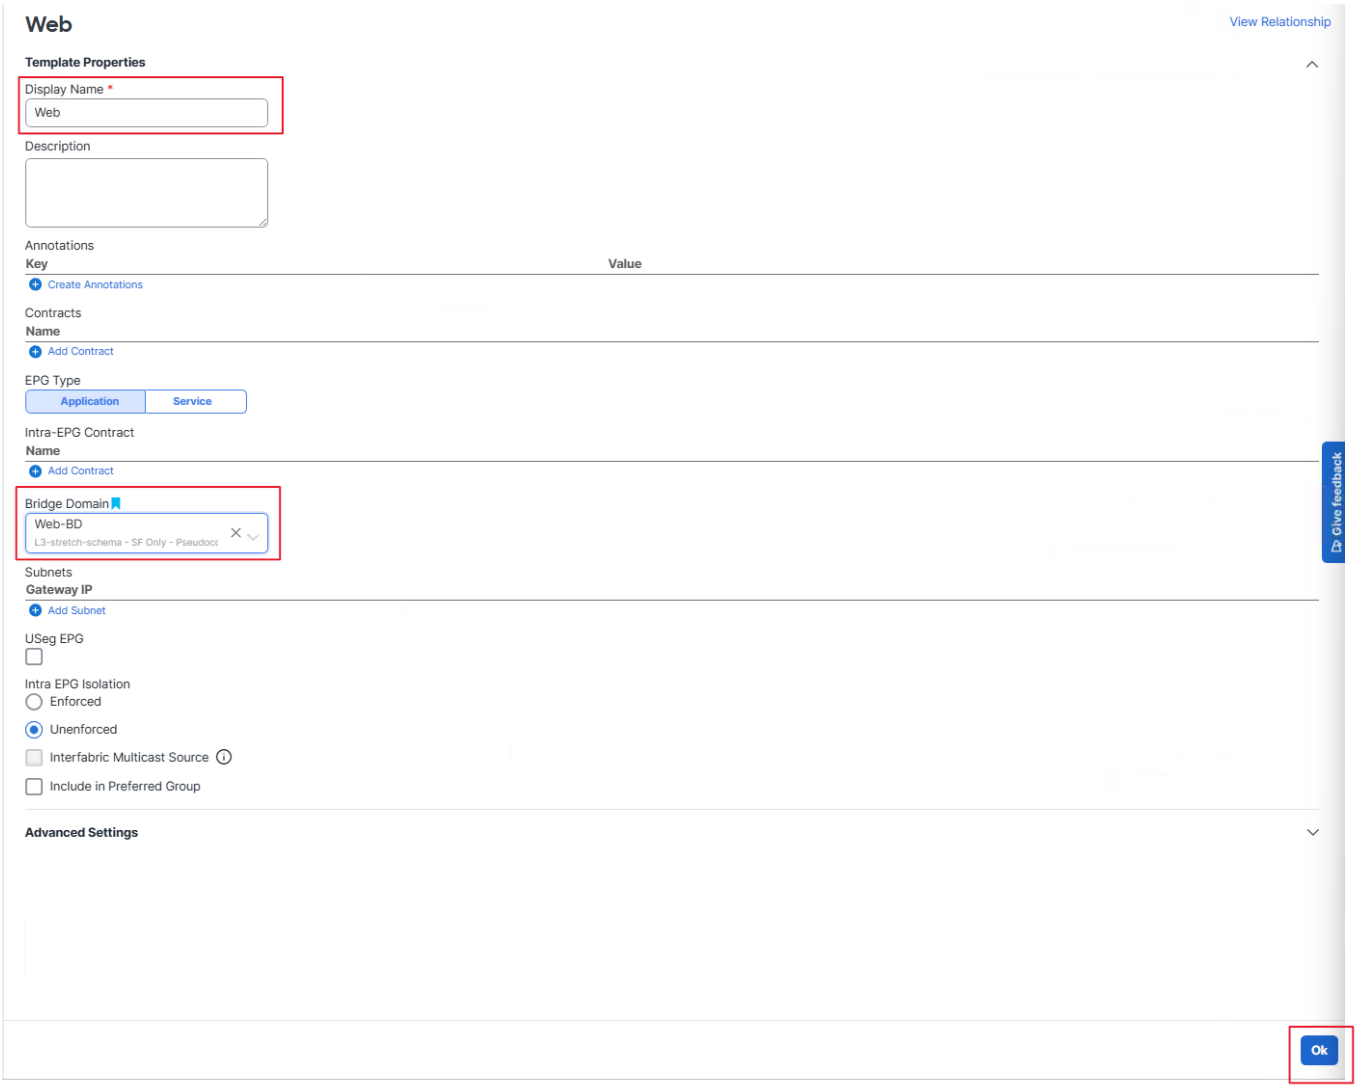

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

18. Expand **Actions** and select **Add/Remove Fabrics**. Choose **San_Francisco,** then click **Ok.**

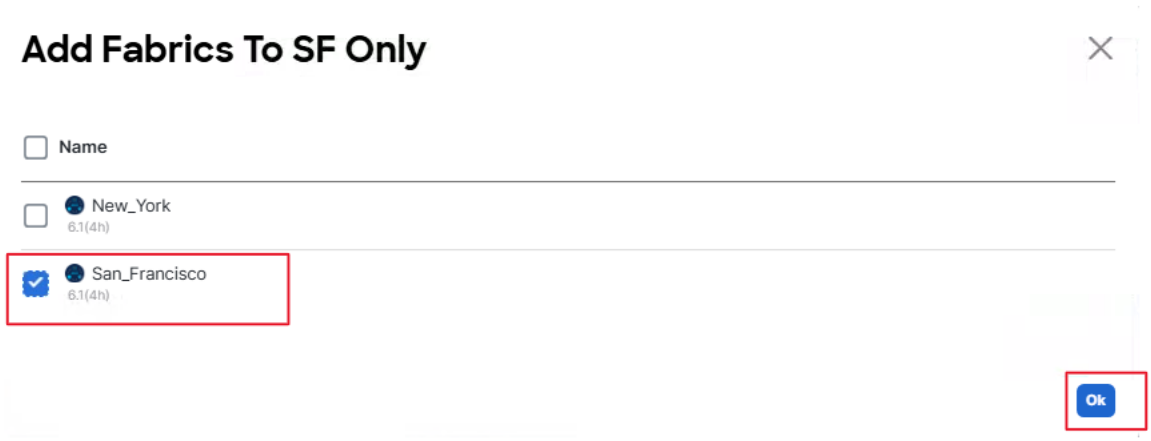

19. Click **Save Schema** to save all the objects that have been created.
20. Next, click **Deploy Template** to push the configuration to the **San_Francisco** site.

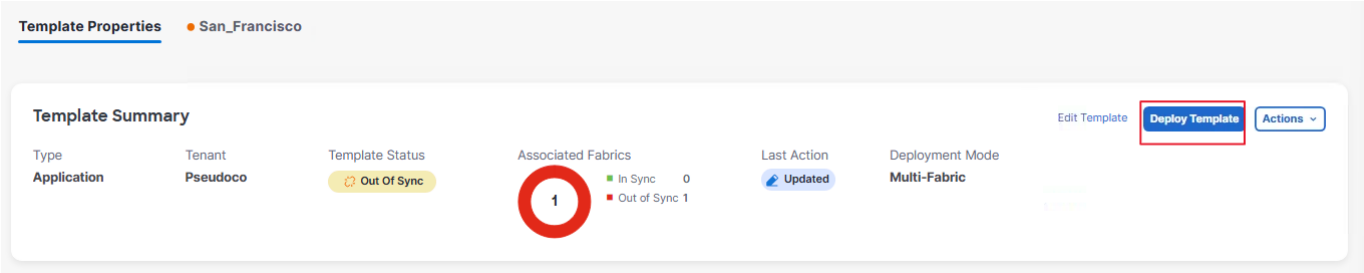

21. Review the Summary, then click on **Deploy Out of Sync Templates**.
22. Another pop-up window will appear indicating Out-of-Sync Templates. Review the out-of-sync templates, then click **Deploy Out of Sync** to proceed with deployment.

A dedicated tenant has now been configured for the **San Francisco** site, including its own **Bridge Domain (BD)**, **EPG**, and **Application Profile**. We will now perform a similar configuration for the **New York** site.

23. Repeat the previous section starting from **Step 12** to create an **NY Only** template and associate it with the **Pseudoco** tenant, making the following adjustments:
    - Create a **NY Only** template.
    - Application Profile Name: **Webapp**
    - Bridge Domain Name: **App-BD**
    - EPG Name: **App**
    - Add the NY Only template to **Site 2 (New_York)**
24. Click **Save Schema** to save all configurations.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

25. Click **Deploy Template** to push the configuration to Site 2 (New_York).

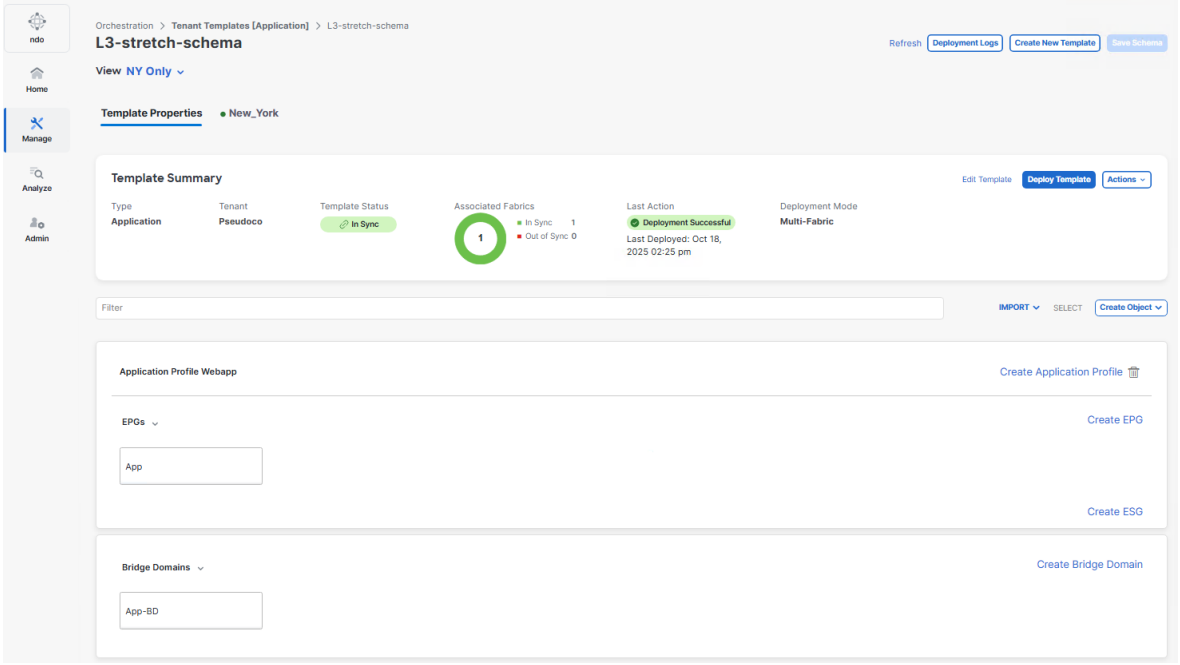

At this point, both sites have been configured with a tenant called Pseudoco and an application profile called Webapp. A Web EPG and BD has been configured in San Francisco and an App EPG and BD has been configured in New York. There is no communication at this time between sites, and the Web BD subnet in San Francisco is not known to New York and vice versa for the App BD in New York. A contract is required in order to allow communication between sites and to advertise endpoint IP address information between sites. Since this contract will be used by both EPGs we will configure it under the SF and NY Template.

26. Select **SF and NY Template** in the vertical menu.

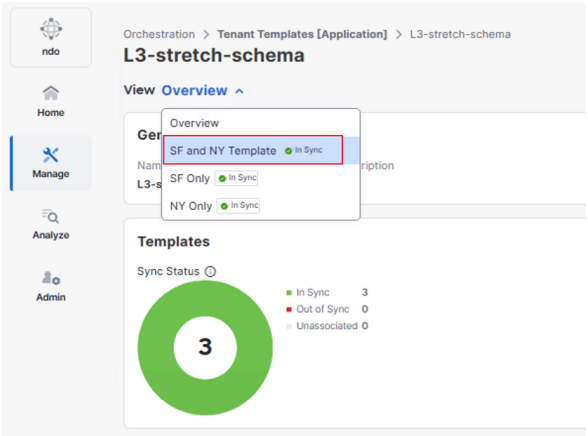

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

27. Expand **Create Object** and select **Contract**.

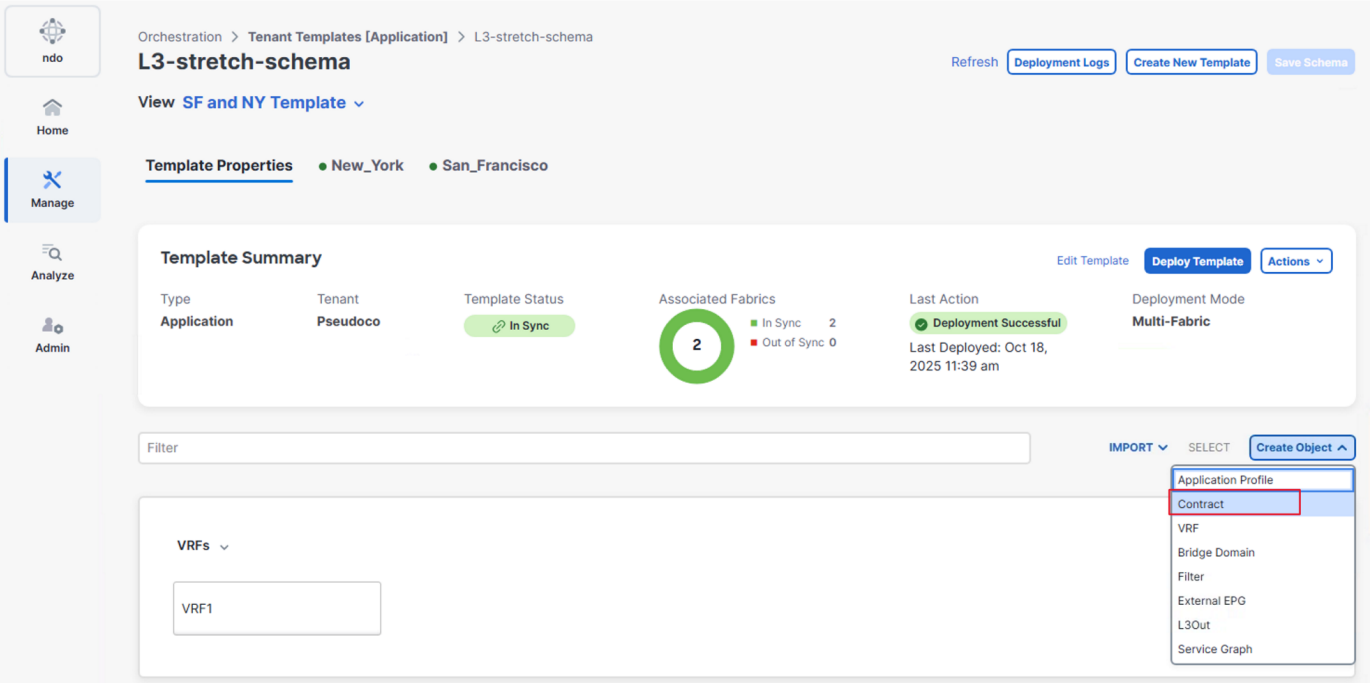

28. Enter **web-to-app** in the **Display Name** field. Click **Ok**.

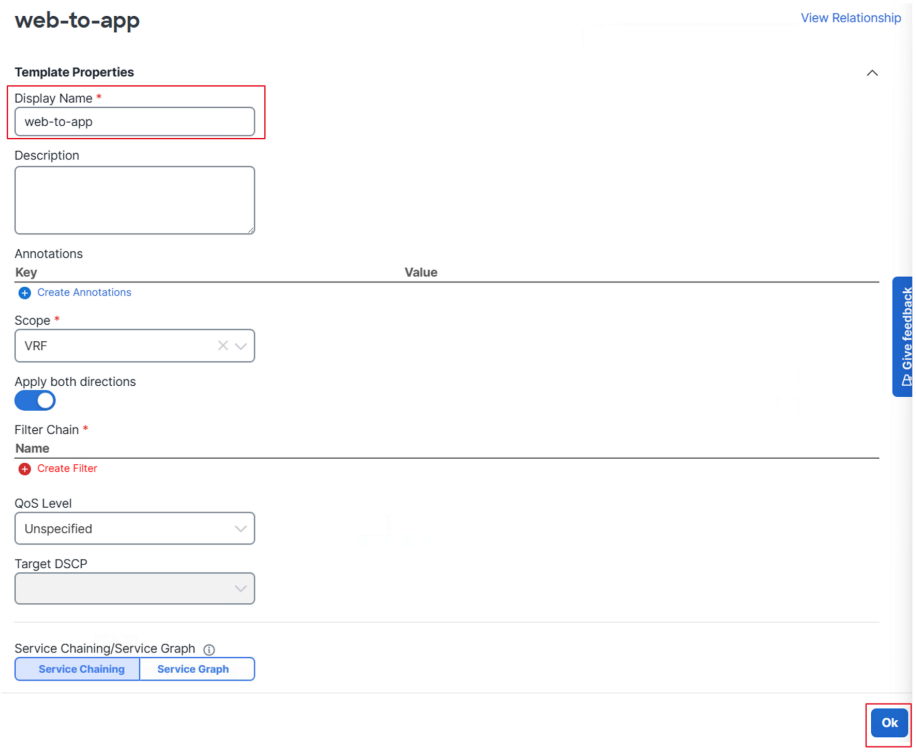

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

29. Expand **Create Object** and select **Filter**.

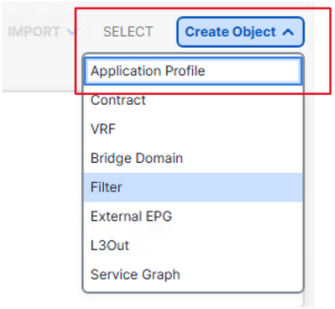

30. Enter **any** in the Display Name field. On the same page, click **Add Entry**, enter **any** in the name field, then click **save**, and finally click **Ok**.

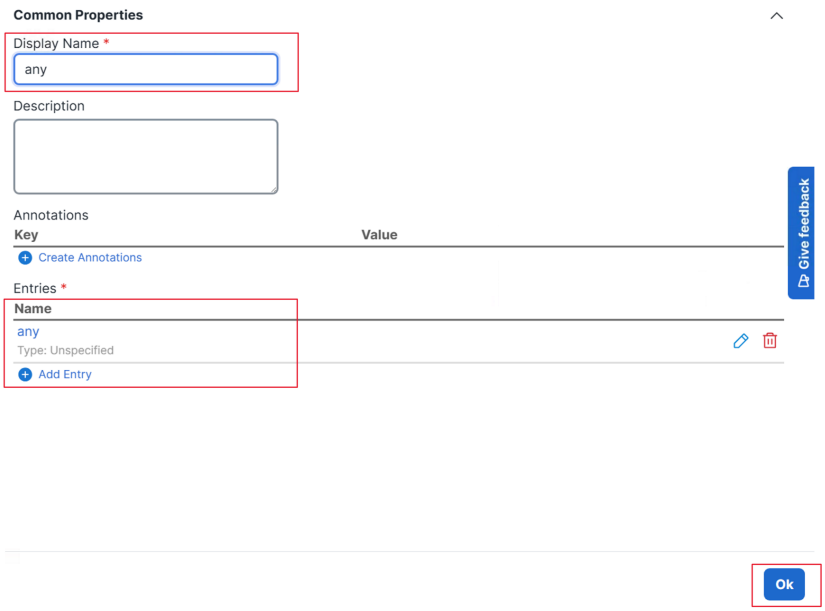

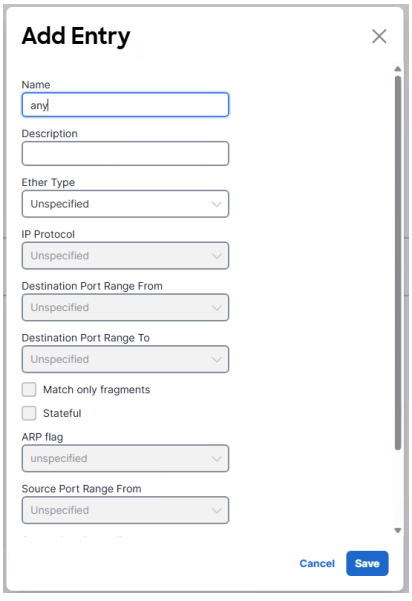

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

31. Next, add the filter that was just created to the contract. To do this, click on the **web-to-app** contract, then click **Create Filter**. In the name field, select **any** filter, leave all other options at their default values, and click **Save**.

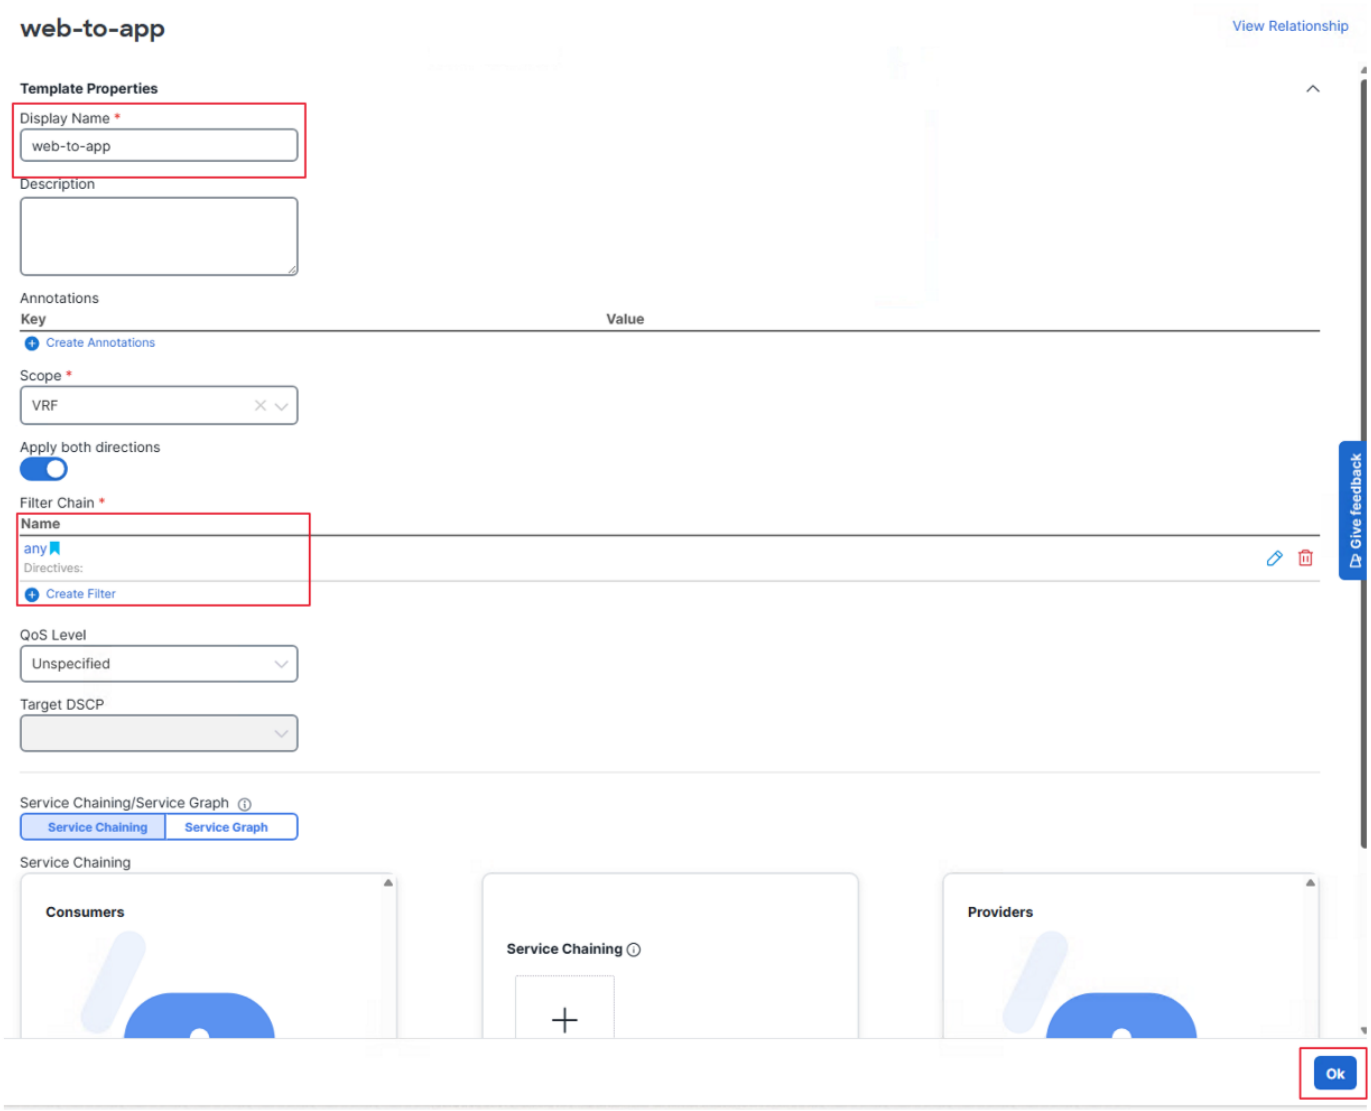

32. Click **Save Schema,** then click **Deploy Template > Deploy > Deploy Out of Sync Templates.**

Now you will add a **web-to-app** provider contract to the **Web EPG** in the **SF Only template**, and a **web-to-app** consumer contract to the **App EPG** in the **NY Only template**. This will enable the application tiers to communicate between sites as long as they are deployed in the same tenant.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

33. In the vertical menu, select **SF Only**.

34. Click **Web** epg.

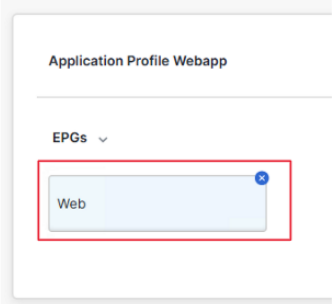

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

35. In the Contract section, click **Add Contract** to associate a contract with the EPG. From the Contract drop-down menu, select **web-to-app**. In type choose **provider**. Click **Save**. Click **OK**.

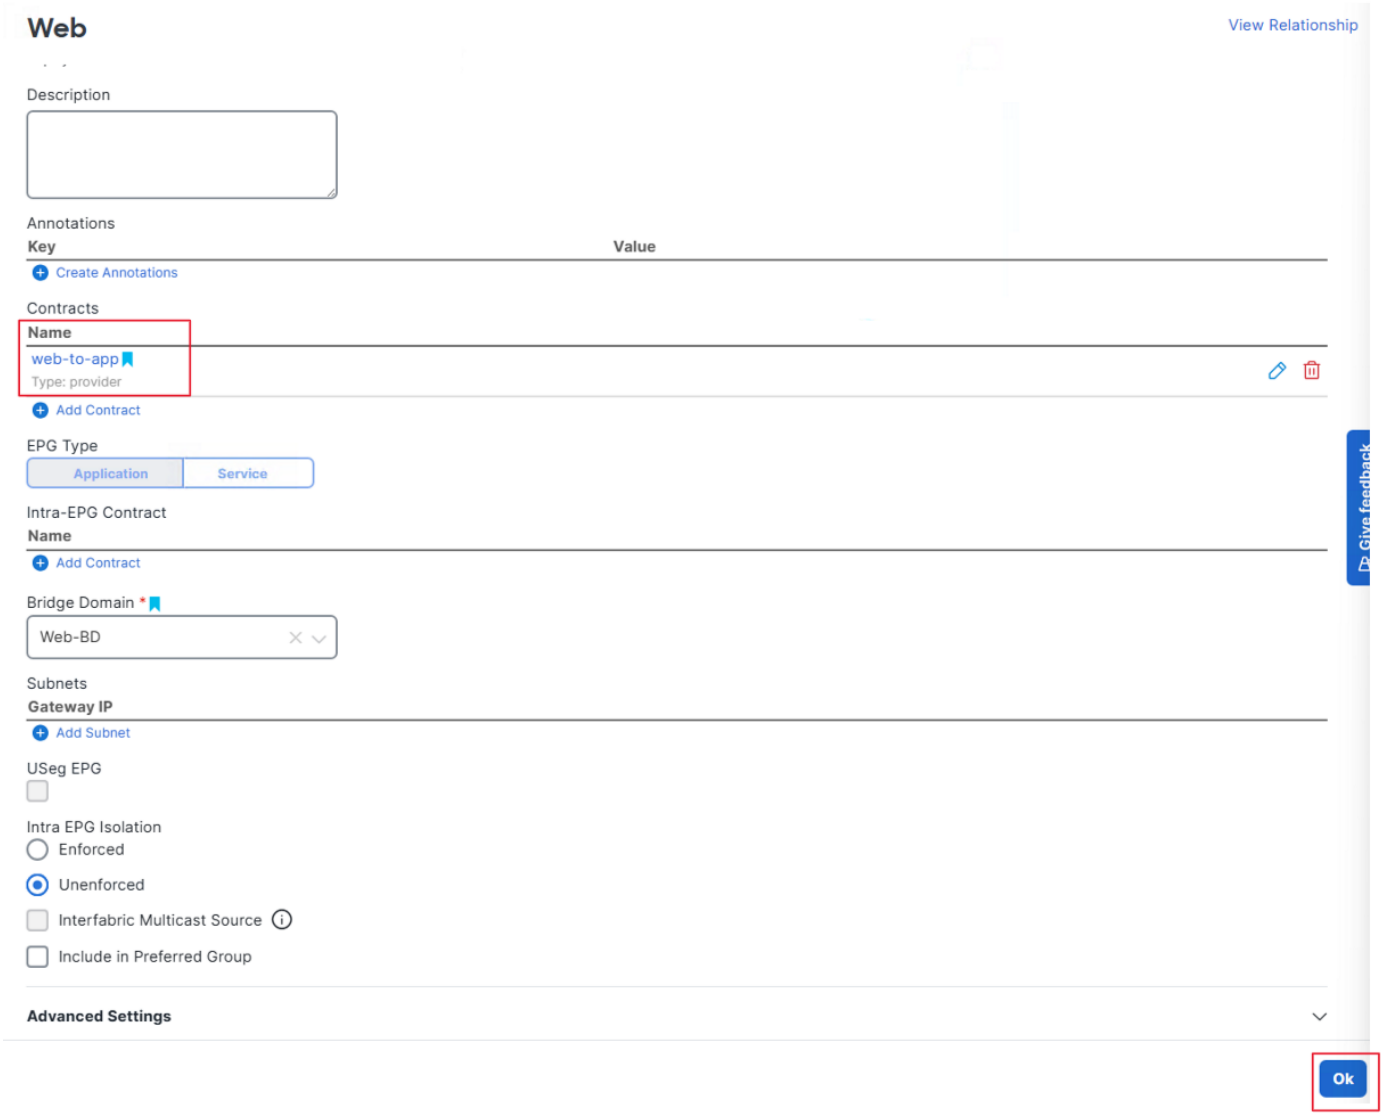

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

36. Click Save **Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.
37. In the vertical menu, select **NY Only**.

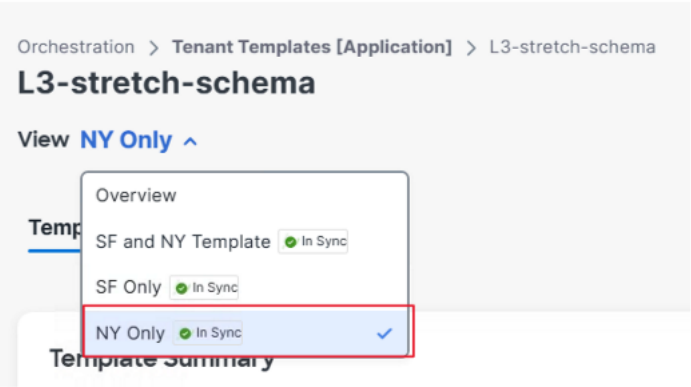

38. Click **App** epg, In the Contract section, click **Add Contract** to associate a contract with the EPG. From the Contract drop-down menu, select **web-to-app,** in type choose **consumer,** then click **Save**, and finally click **OK**.

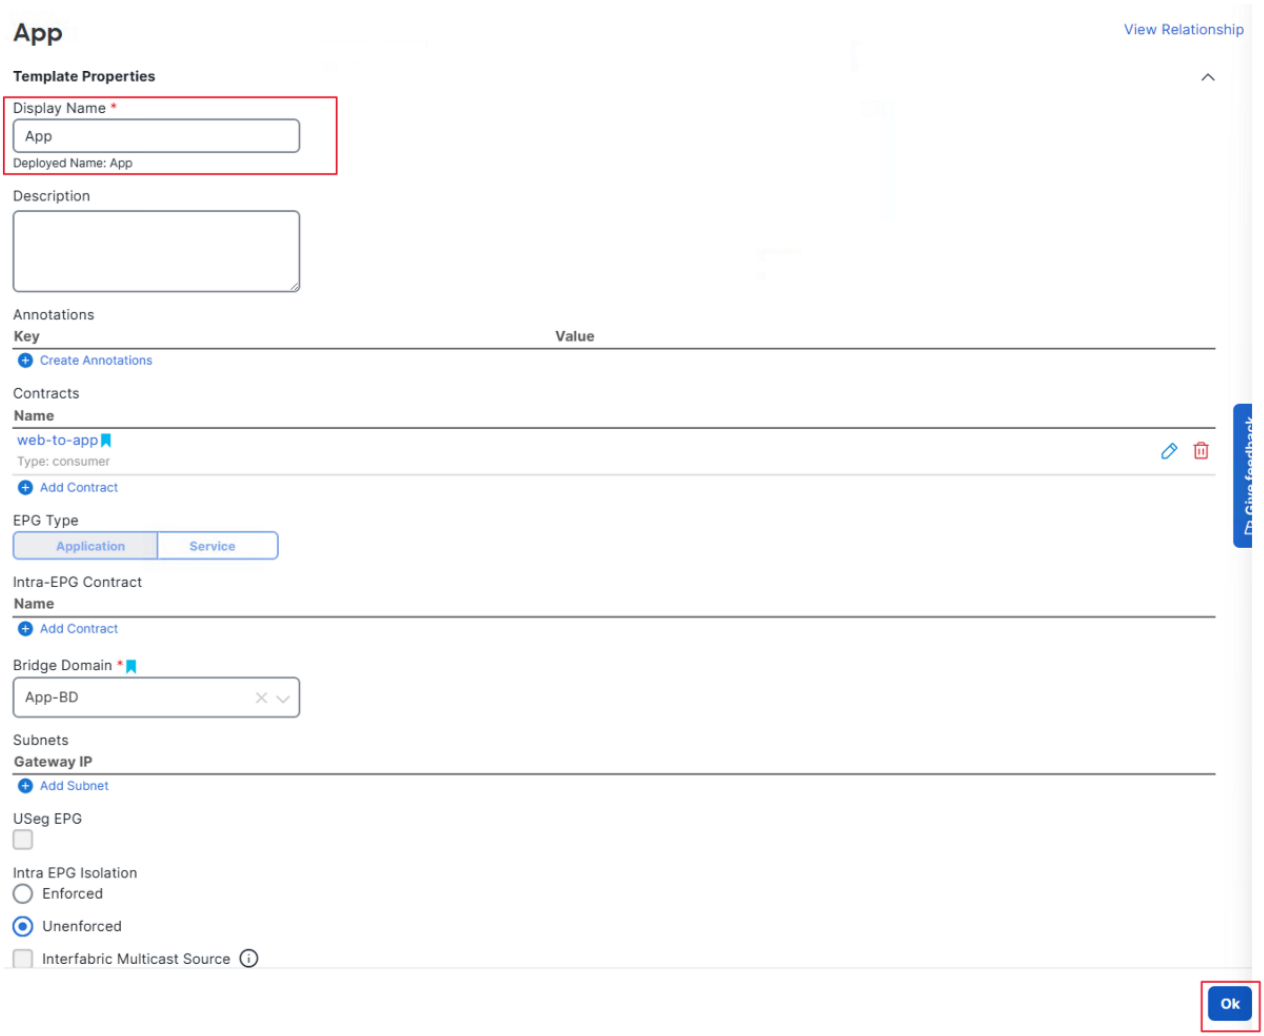

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

39. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.

Now verify all newly created objects, including the Application Profile (AP), Bridge Domain (BD), Endpoint Group (EPG), and Contract.

40. Open a web browser and either select the **APIC-SF** bookmark or enter the APIC San Francisco IP address **198.18.133.200** in the address bar. Navigate to **Tenants** > **Pseudoco**, and verify the following properties:
    - Application Profiles: **Webapp**
    - Application EPGs: **App**
    - Contracts: **web-to-app**
    - Contract type: **Consumed**
    - Bridge Domains: **Web-BD**
    - VRFs: **VRF1**

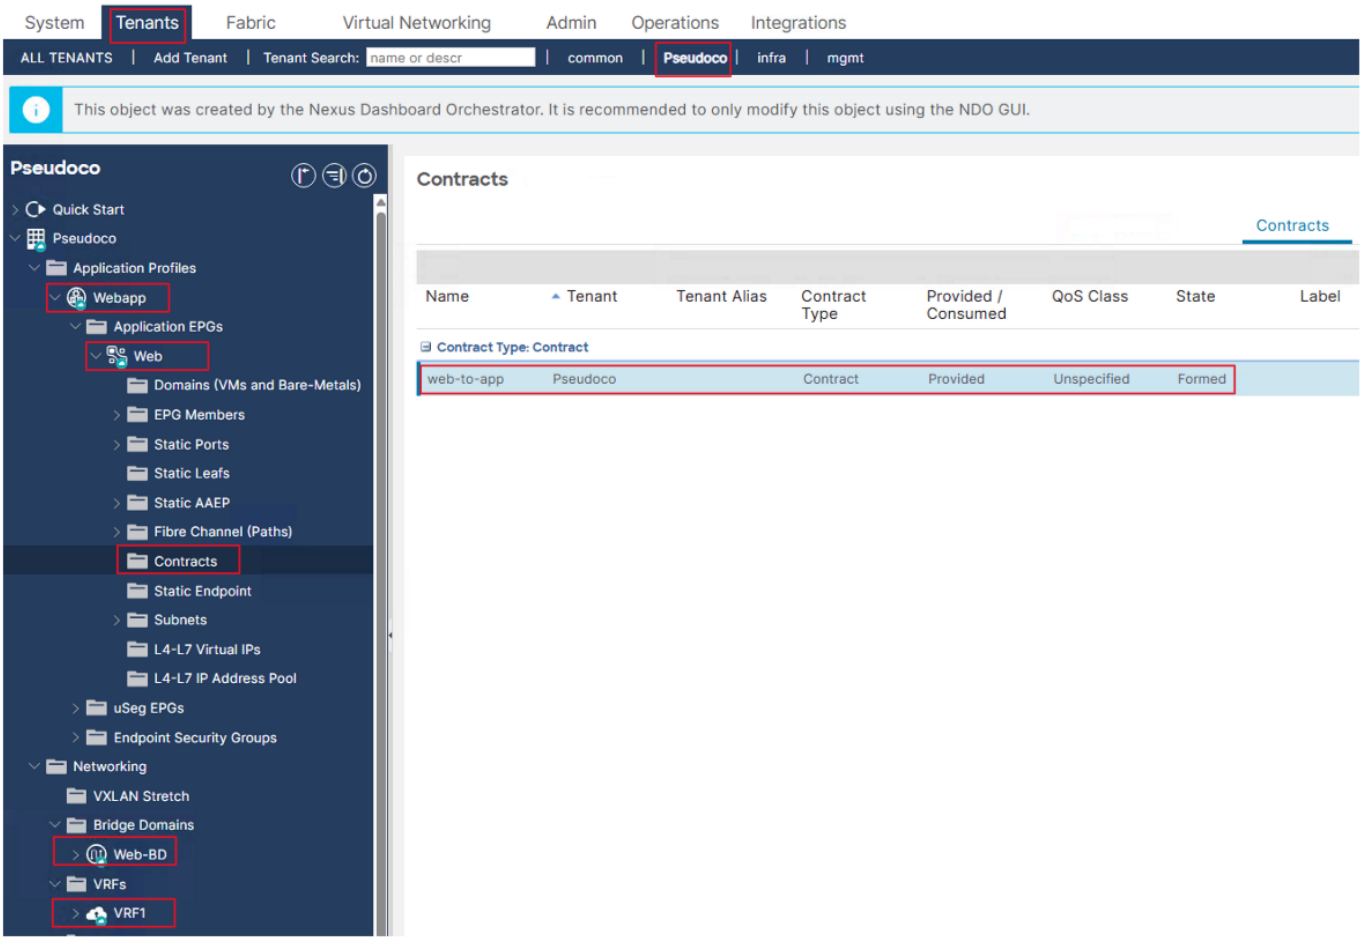

41. Repeat the same steps for **APIC-NY** (**198.18.132.200**).
    - Application Profiles: **Webapp**
    - Application EPGs: **App**
    - Contracts: **web-to-app**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

- Contract type: **Consumed**
- Bridge Domains: **App-BD**
- VRFs: **VRF1**

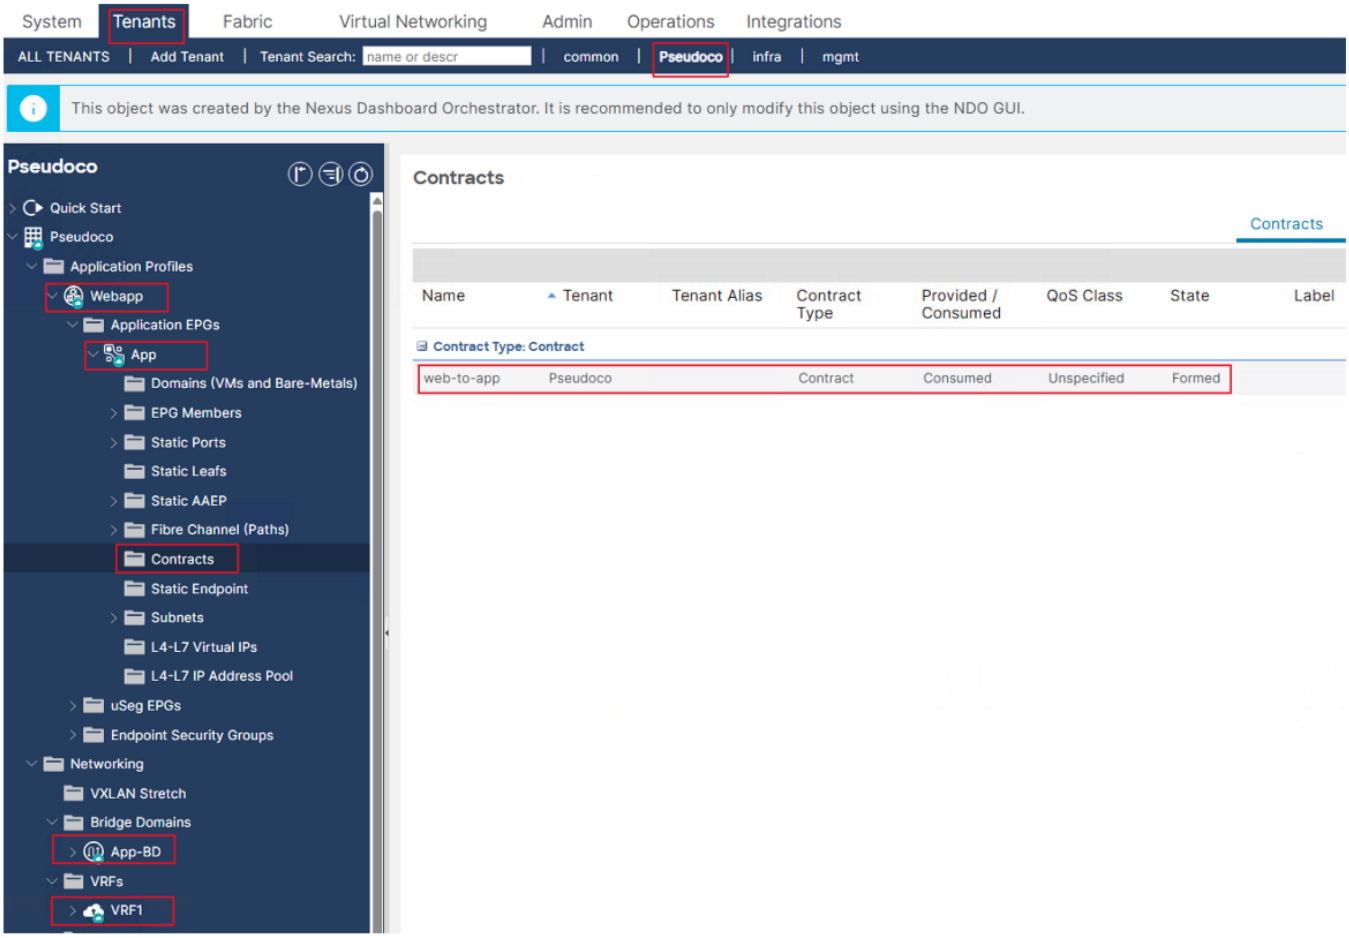

Value Proposition:

This scenario implements Layer 3–only intersite communication within a single tenant and VRF, enabling routed connectivity between endpoints located in different data centers without extending Layer 2 domains. The design leverages inter-VRF routing over the Cisco ACI Multi-Site architecture, ensuring efficient IP reachability, optimal path selection, and fault isolation across sites.

By using Layer 3 segmentation, this model minimizes broadcast traffic, enhances network stability, and provides scalable, policy-based routing between sites. It supports consistent policy enforcement, centralized control, and seamless integration with external networks, delivering a robust and operationally efficient multi-site architecture

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 4. Layer 3-Only Communication Across Sites and VRFs (Shared Services)

Building upon the foundational Layer 3 inter-site communication established in Scenario 3, this scenario introduces a critical capability for enterprise environments: **secure access to shared services** across multiple sites and different VRFs. It demonstrates how to provision centralized services—such as DNS, DHCP, or Active Directory—hosted in a single location, making them securely accessible to application endpoints distributed across various sites and even within different tenant VRFs.

This design is achieved by introducing a dedicated **Shared VRF, Bridge Domain (BD), and Endpoint Group (EPG)**. The objective is to enable controlled **inter-VRF communication** while meticulously maintaining tenant and application isolation. This approach ensures efficient resource utilization, consistent policy enforcement, and streamlined management, allowing your multi-site ACI fabric to provide essential services without compromising security or scalability.

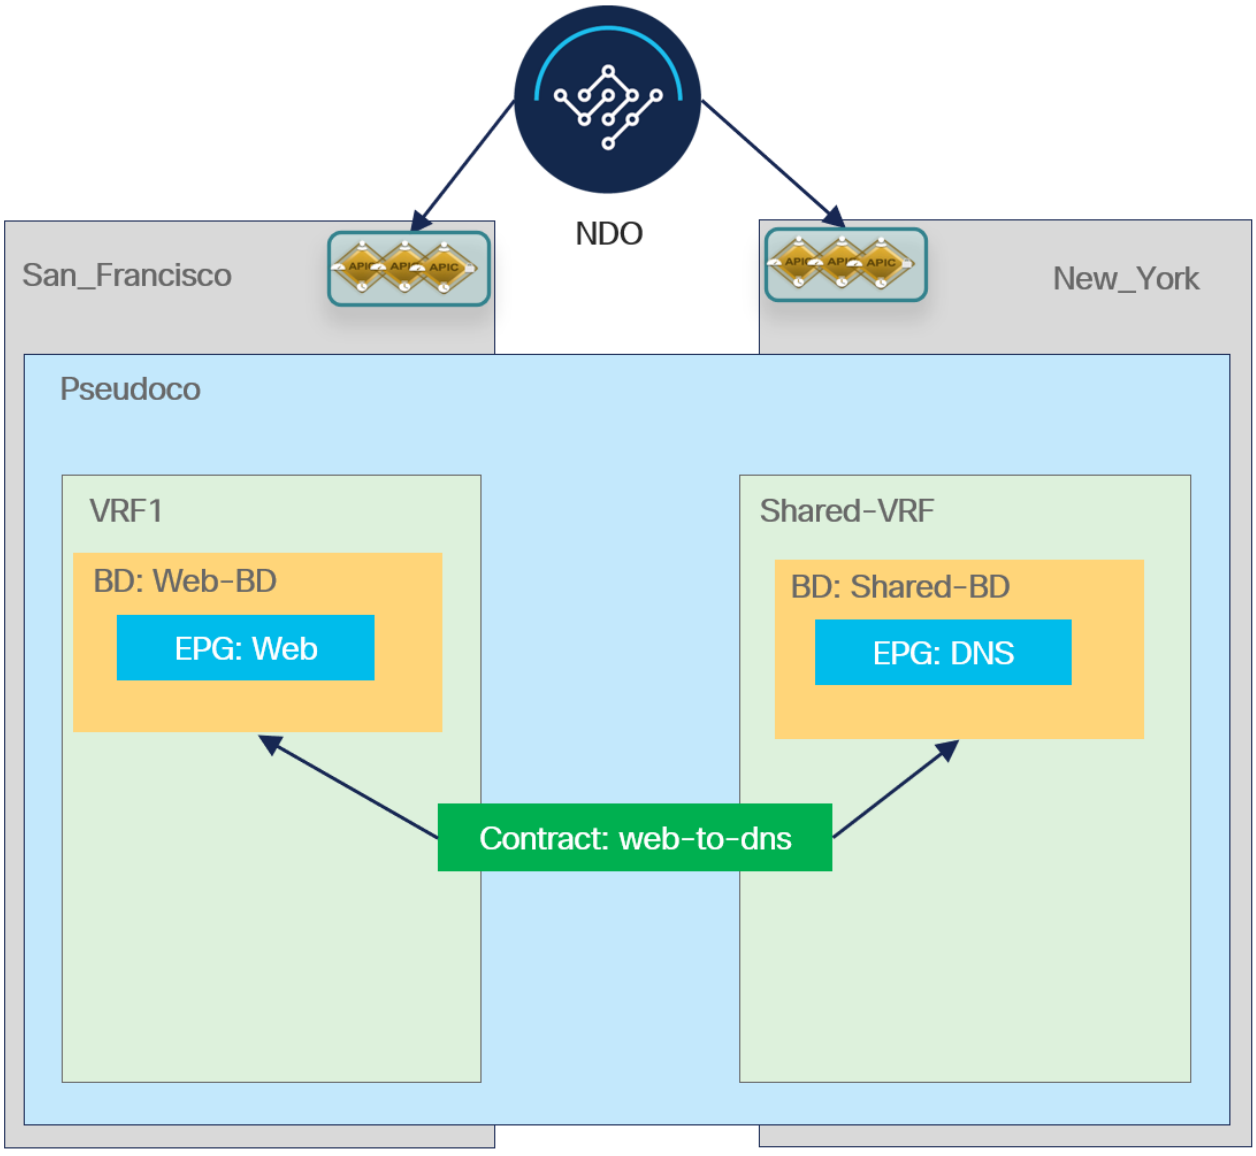

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

1. In the Nexus Dashboard window, navigate to **Manage** > **Orchestration** > **Tenant Templates** > **Applications**, then double click **L3-Stretch-schema**.
2. Expand the **View** menu, then select **NY Only.**

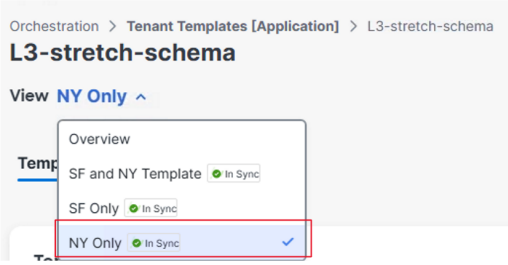

3. Expand **Create Object**, and select **VRF**, in the Display Name field, enter **Shared VRF**. Click **Ok.**
4. Expand **Create Object**, and select **Bridge Domain**, and enter the following properties:
   - Display Name: **Shared-BD**
   - Virtual Routing and Forwarding: **Shared VRF**
   - Subnets, add this Gateway IP: **10.10.1.254/24**, then enable (Shared between VRFs), click **Ok**.
   - Click **Ok** to save the Bridge Domain configuration**.**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

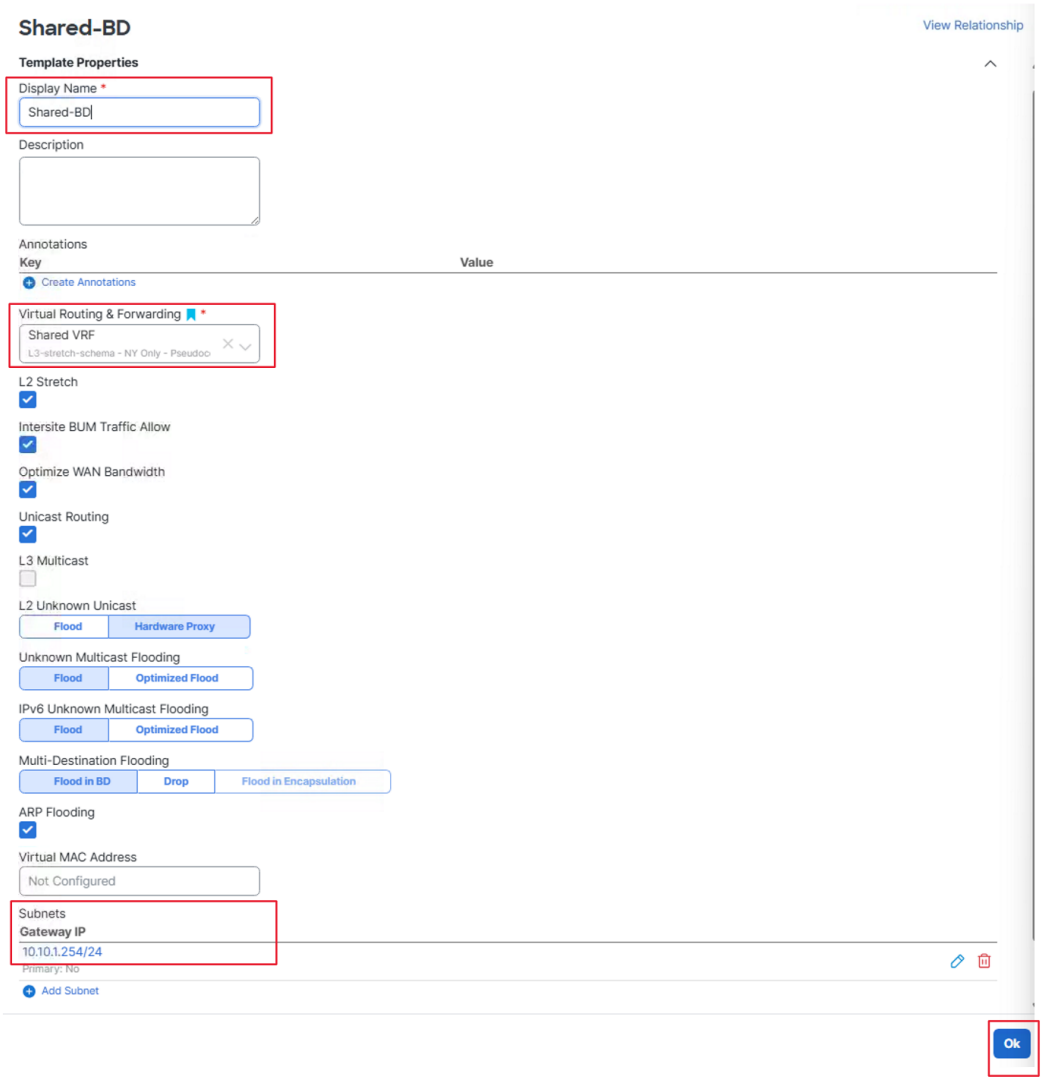

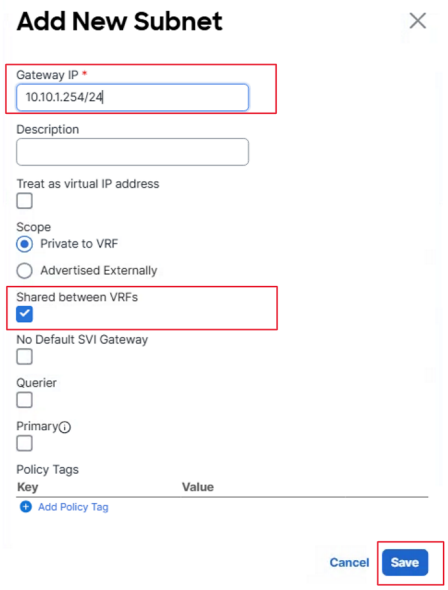

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5. In the **Application Profile Webapp** section, click **Create EPG**.

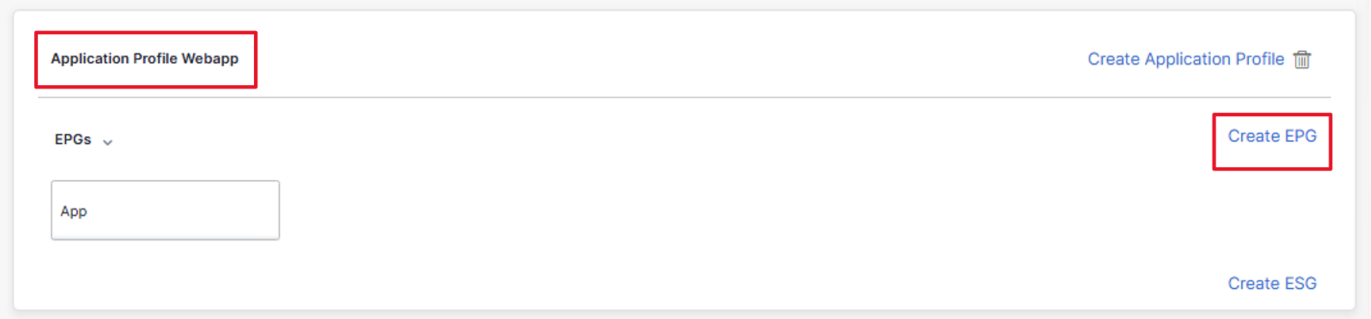

6. Enter **DNS** in the Display Name, select **Shared-BD** in the bridge domain field, and then Create a new subnet. In the DNS EPG work pane, click **+ Add Subnet**, enter the following attributes and click **OK** to save.
   - Gateway IP: **10.10.1.254/24**
   - Shared between VRFs: **checked**
   - No Default SVI Gateway: **checked**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

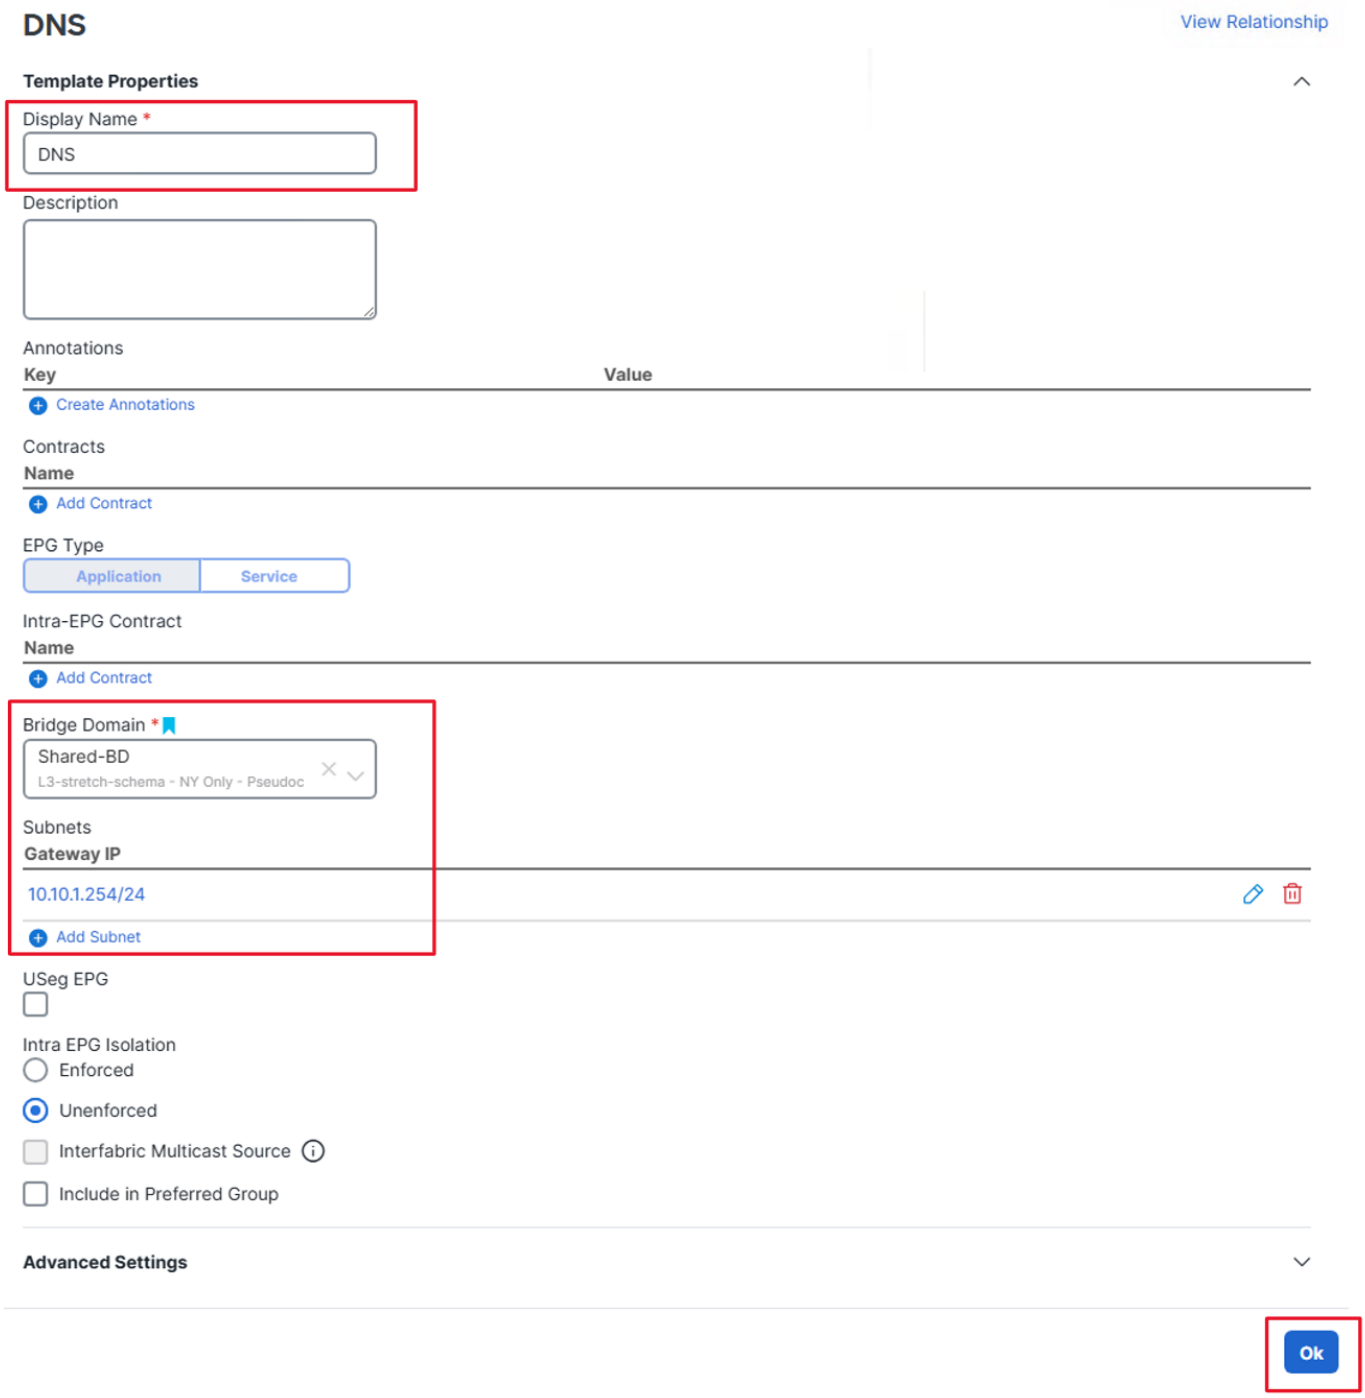

The purpose of defining the IP subnet information under the provider EPG is to enable the necessary VRF rout-leaking functions between the Shared VRF and the other VRFs accessing the shared services. The IP subnet that is configured at the BD provides the default gateway services, so it's important to select the No Default SVI Gateway flag.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Next, enable communication between **Web-BD** and **Shared-BD** by creating a **contract** in the stretched template (**SF and NY Template**). After creating the contract, add it to each **site-associated template** to establish communication between the corresponding Bridge Domains.

7. In the View menu, select **SF and NY Template**.

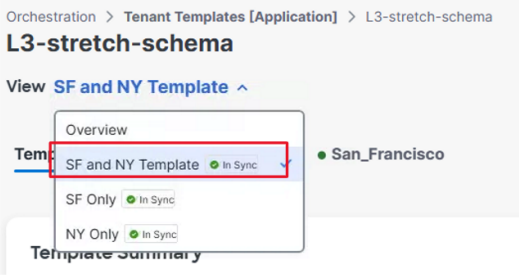

8. In the **Contract** section, click **Create Contract**, then enter the following attributes:
   - Display Name: **web-to-dns**
   - Scope: **Tenant**
   - Filter Chain: click **+ Add Filter** and then, select **any**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

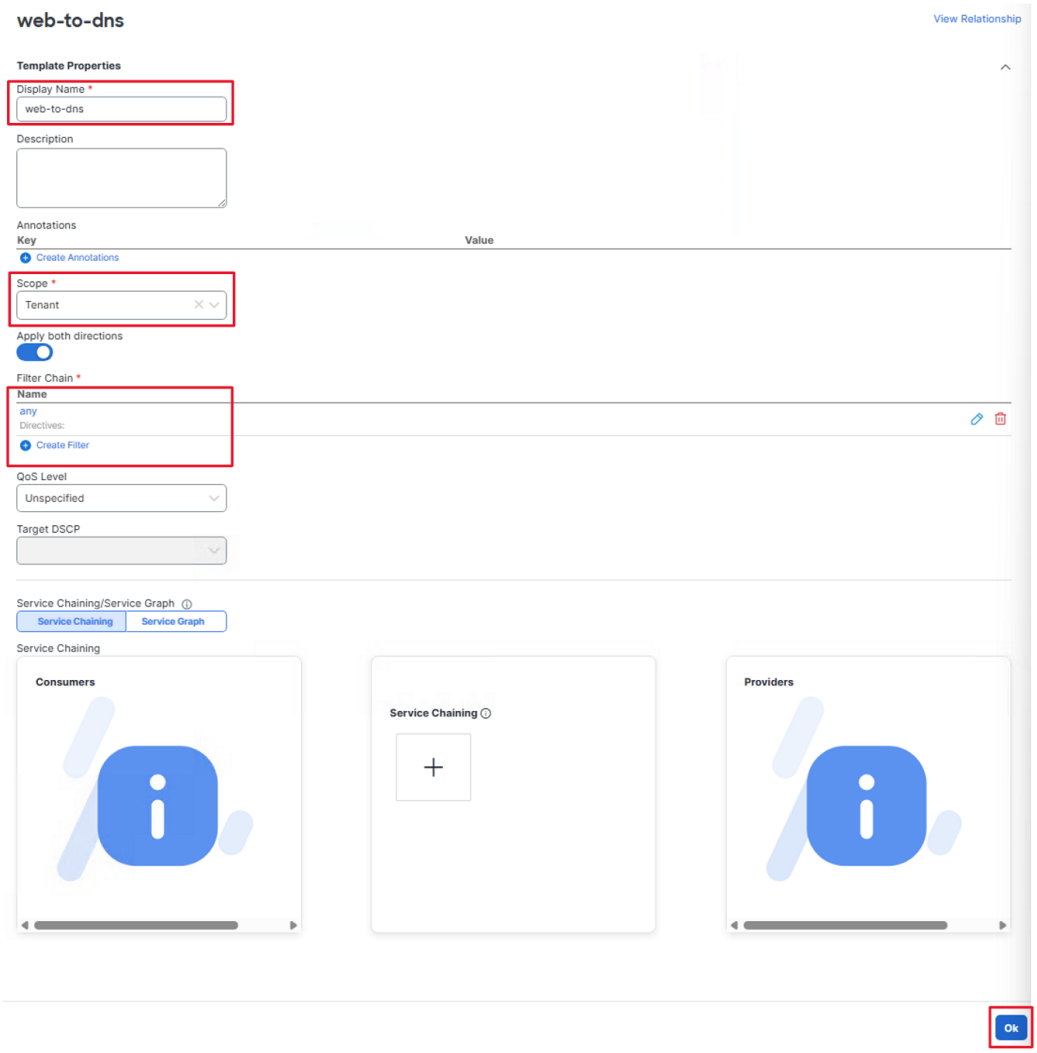

9. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.

Next, associate the contract with the appropriate EPGs by configuring the **DNS** EPG in the **New_York** site as the Provider, and the **Web** EPG in the **San_Francisco** site as the Consumer.

10. In the **View** menu, select **NY Only.** Under **EPG**, click **DNS** and then, click **+ Add Contracts**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

11. In the Add Contracts dialog, select **web-to-dns** then, in the Type field, select **provider** and then, click **OK**.

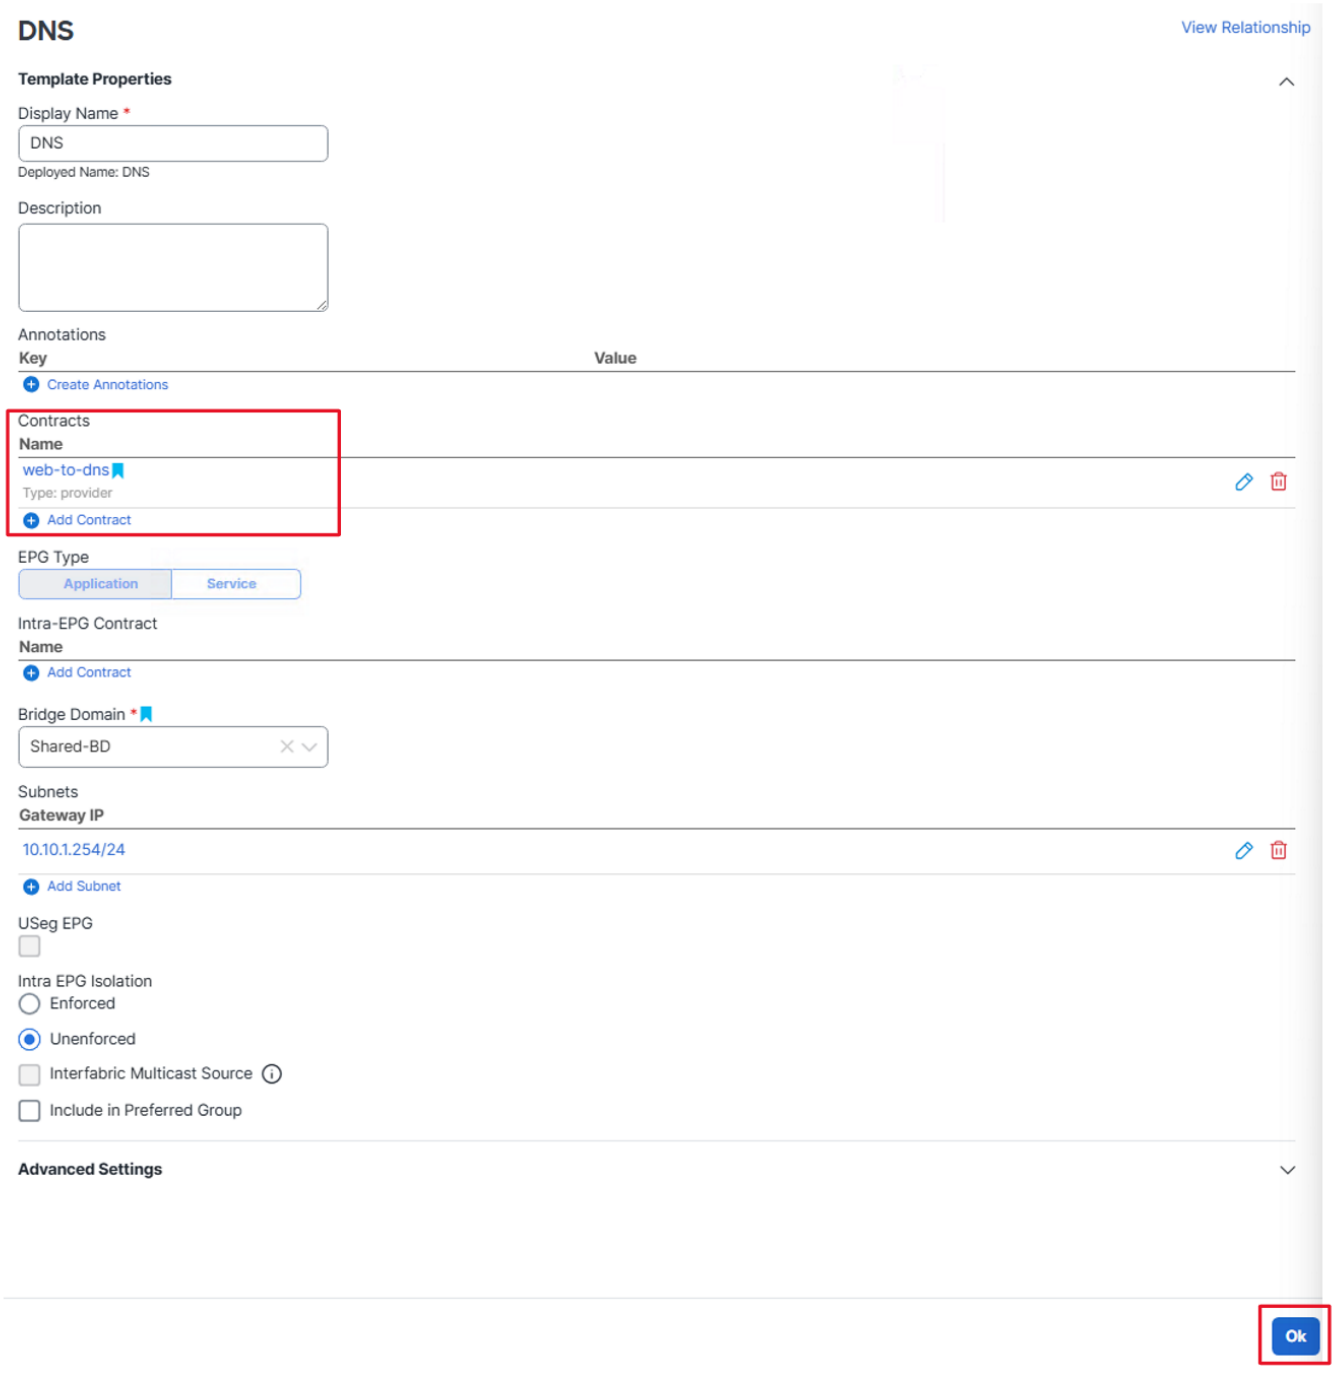

12. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.
13. In the View menu, select **SF Only**. Under EPG, click **Web** and then, click **+ Add Contracts**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

14. In the Add Contracts dialog, select **web-to-dns** then, in the Type field, select **consumer**. Click **OK**.

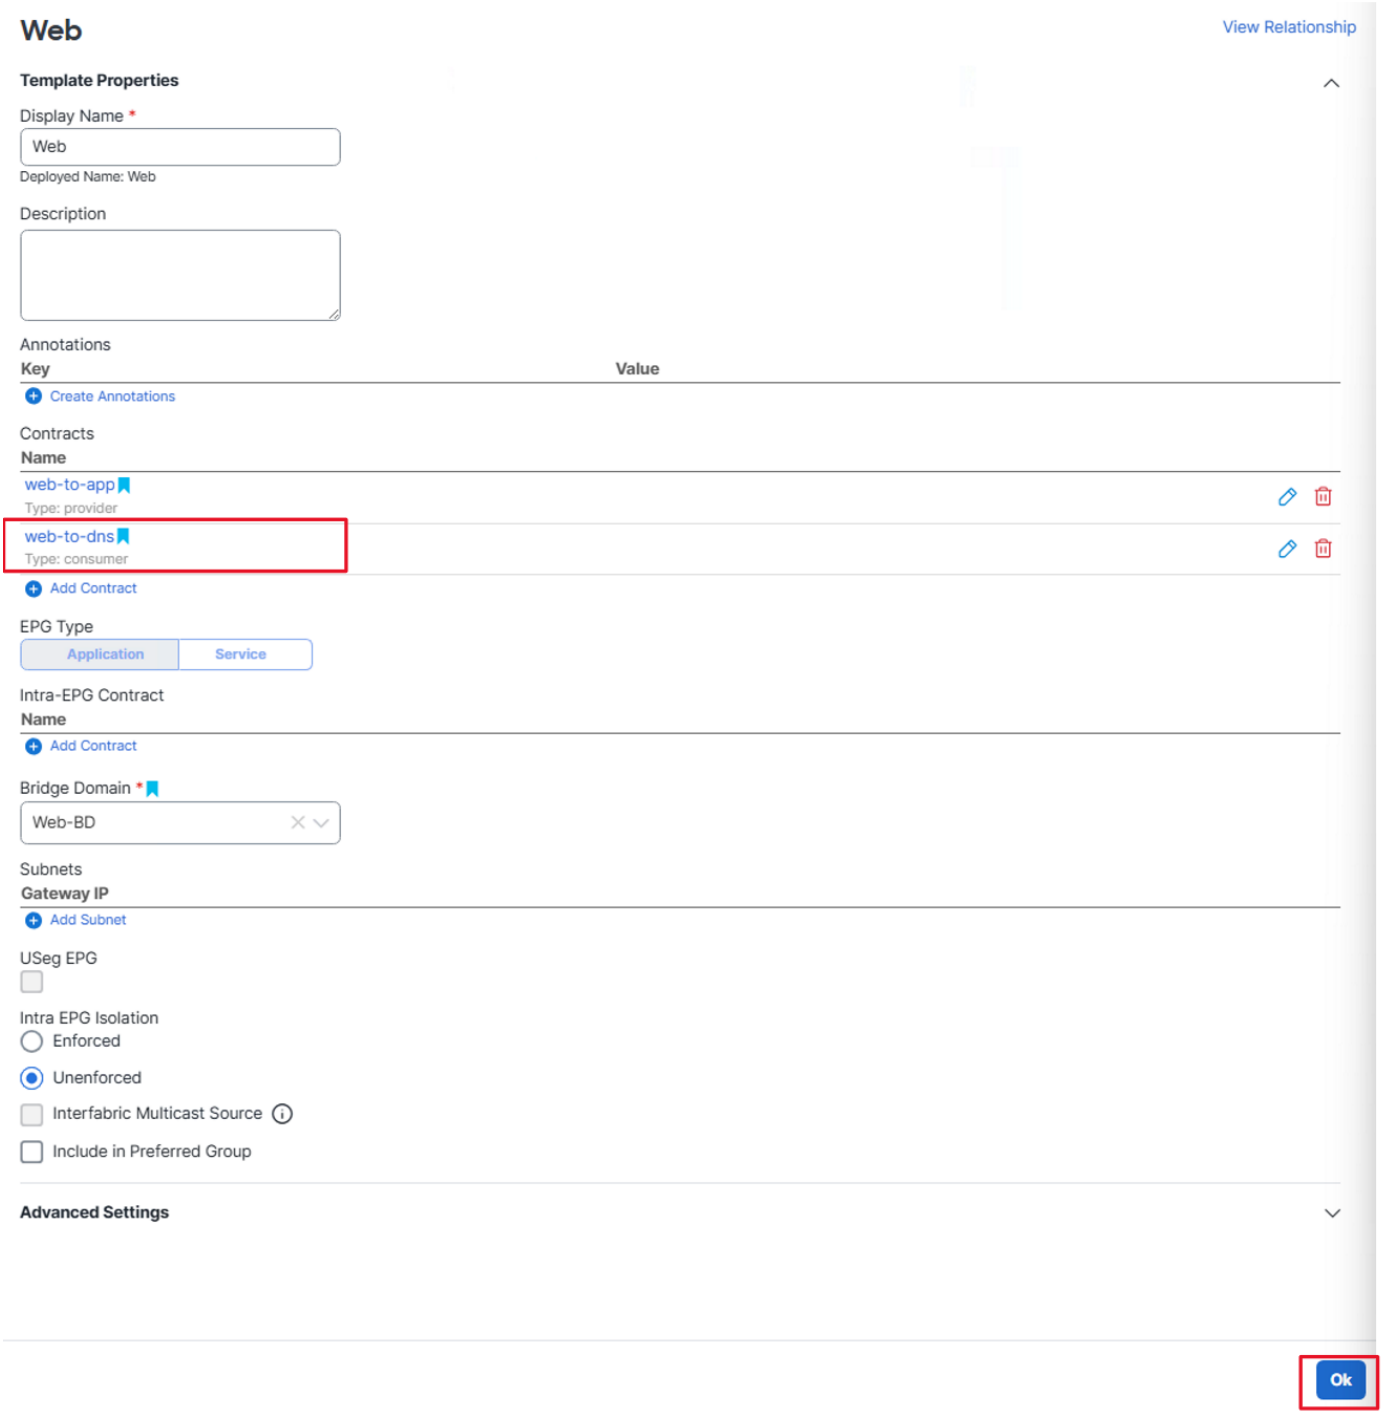

15. Click Save **Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

16. Verify the configuration on both sites. Log in to each site's **APIC**, then navigate to **Tenants > Pseudoco > Application Profiles > Webapp**, and click **Topology** to confirm the deployed objects and their relationships.

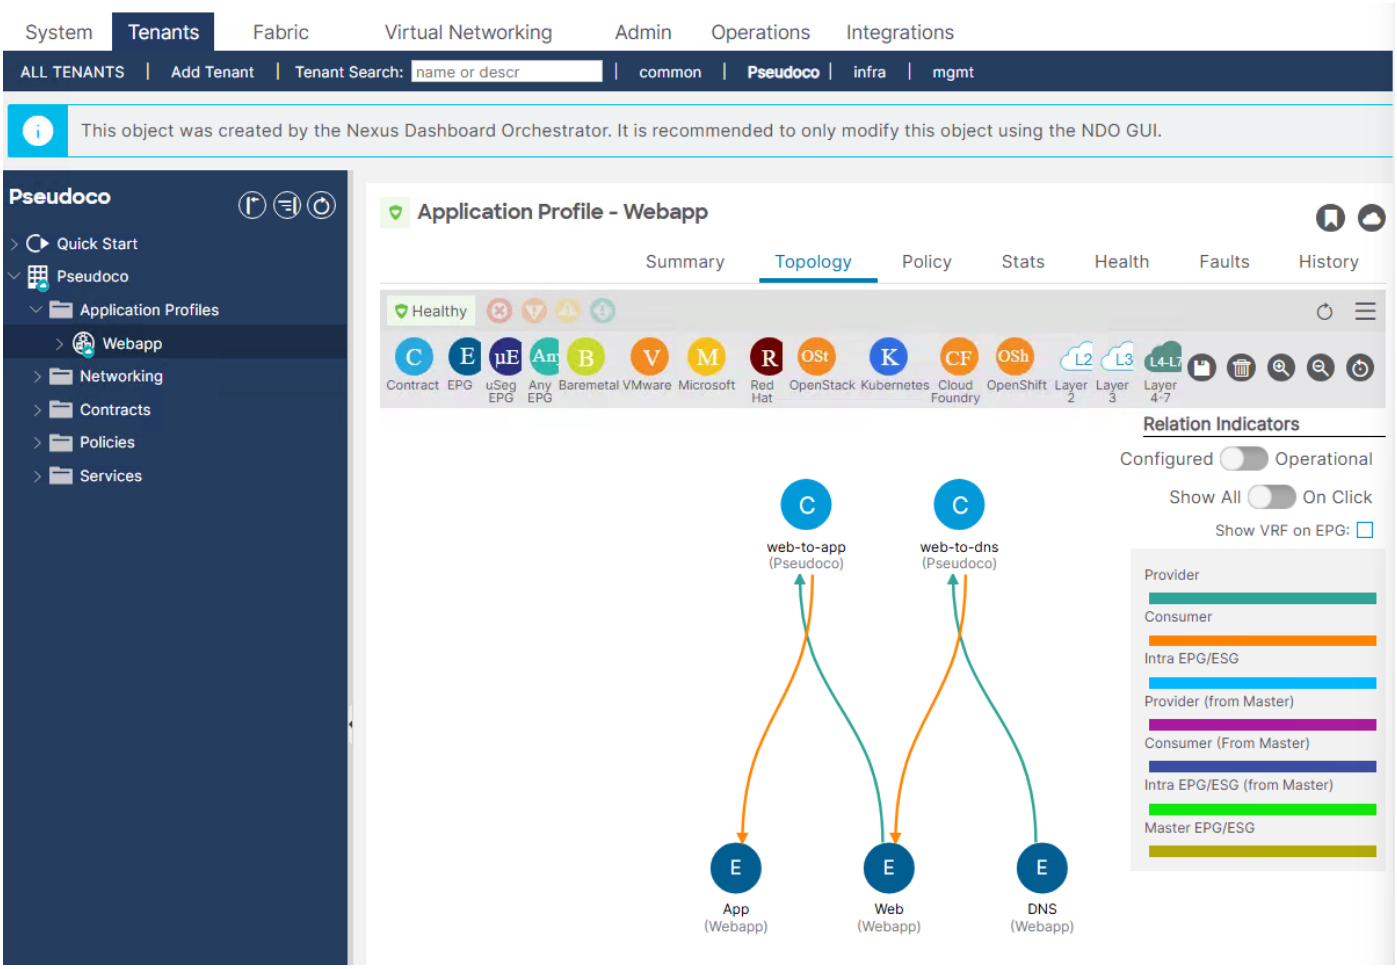

Value Proposition:

This scenario demonstrates **Layer 3 intersite and inter-VRF communication** within the Cisco ACI Multi-Site architecture, enabling shared services accessibility across geographically distributed data centers. By leveraging a **dedicated Shared VRF, Bridge Domain (BD), and Endpoint Group (EPG)**, centralized services such as DNS, DHCP, and authentication can be securely accessed by application tiers located in different sites and VRFs.

The design facilitates **controlled communication between isolated environments** while maintaining policy enforcement, routing efficiency, and operational scalability. Through centralized service hosting and contract-based communication, enterprises can reduce resource duplication, streamline management, and ensure consistent connectivity across multi-site deployments—all within a **secure, policy-driven Layer 3 framework**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 5. IP Mobility Across Sites (Stretched BD without BUM Flooding)

This crucial scenario addresses a key requirement for modern data centers: **seamless IP mobility** for workloads across geographically distributed sites, particularly vital for disaster recovery and workload relocation strategies. It demonstrates how to achieve this by implementing a **stretched Bridge Domain (BD) with a single subnet** spanning both the San Francisco and New York sites.

The innovation here lies in **disabling Broadcast, Unknown Unicast, and Multicast (BUM) flooding** across the inter-site network. This critical design choice prevents the traditional pitfalls of stretched Layer 2 domains, such as broadcast storms and inefficient traffic propagation, while still enabling endpoints to retain their IP addresses when moved between locations.

By isolating Layer 2 broadcast domains to their respective sites, this configuration ensures enhanced control, resiliency, and consistent Layer 3 reachability. You will deploy a new DB-BD under the Pseudoco tenant, showcasing a policy-driven approach to enable efficient IP mobility and seamless application recovery across data centers without the need for IP re-addressing.

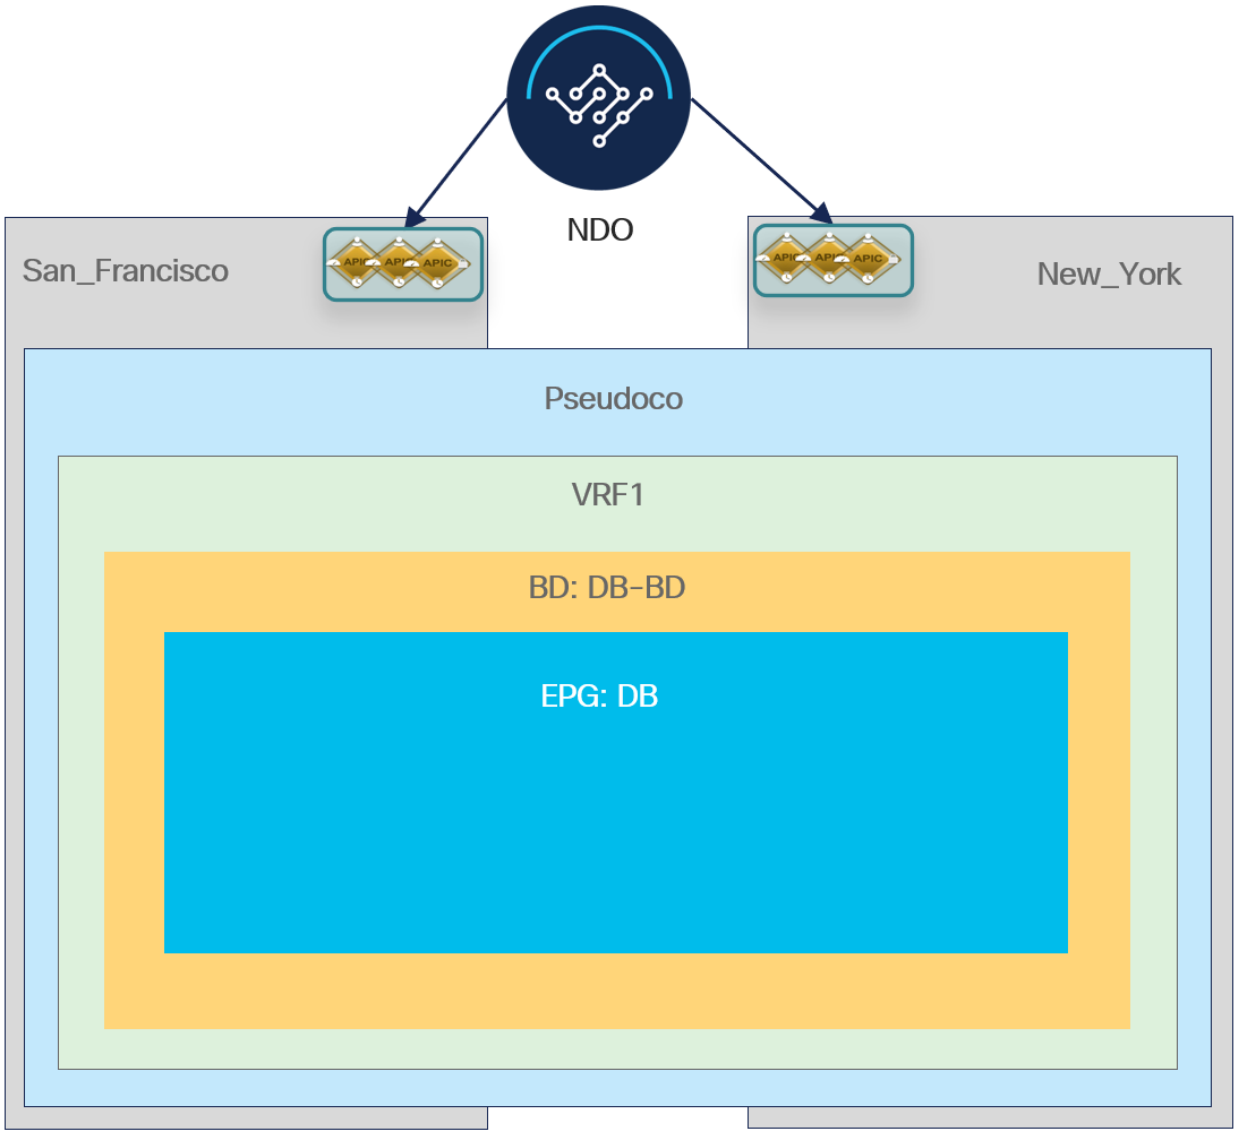

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

1. In the Nexus Dashboard window, navigate to **Manage** > **Orchestration** > **Tenant Templates** > **Applications** > **L3-Stretch-schema** > **SF and NY Template**.

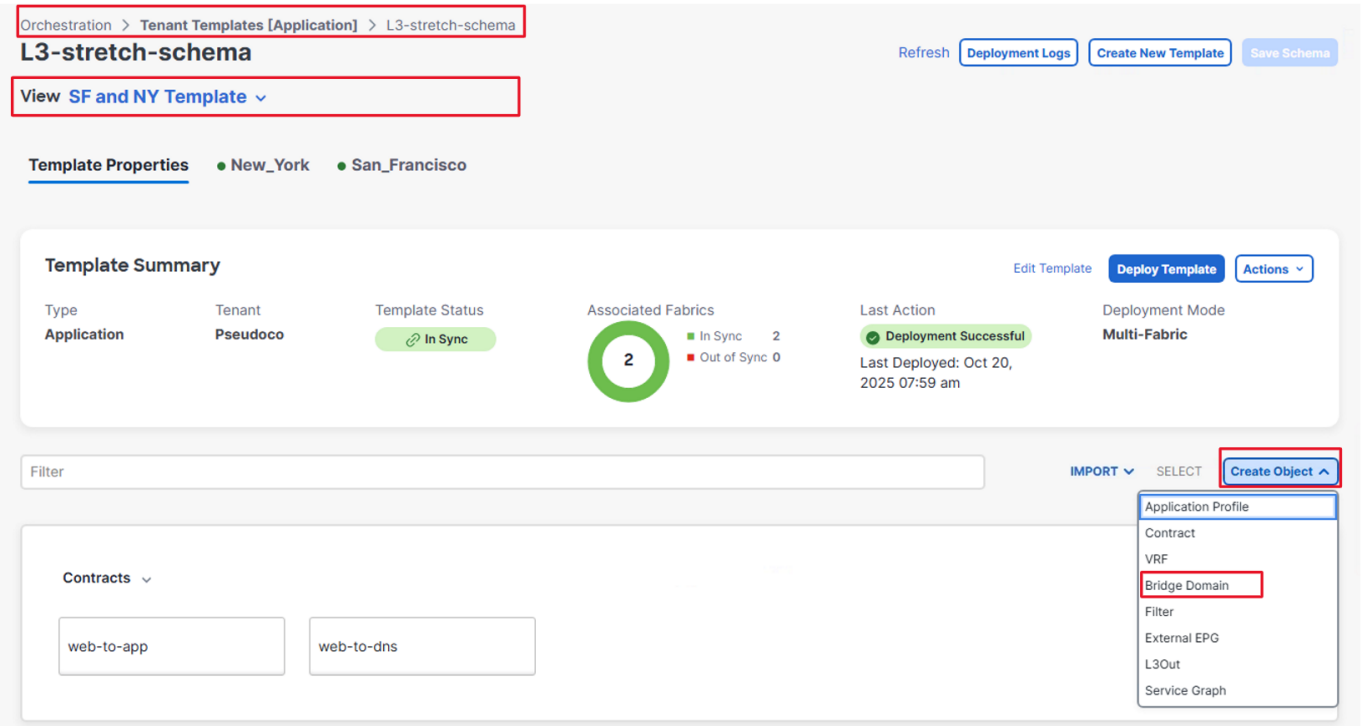

2. Expand **Create Object**, and select **Bridge Domain**, and configure the following, then **Ok**:
   - Display Name: **DB-BD**
   - Virtual Routing & Forwarding: **VRF1**
   - INTERSITE BUM TRAFFIC ALLOW: **Deselected** (Click Yes on the Warning dialog)
   - Add Subnet: **10.3.1.254/24**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

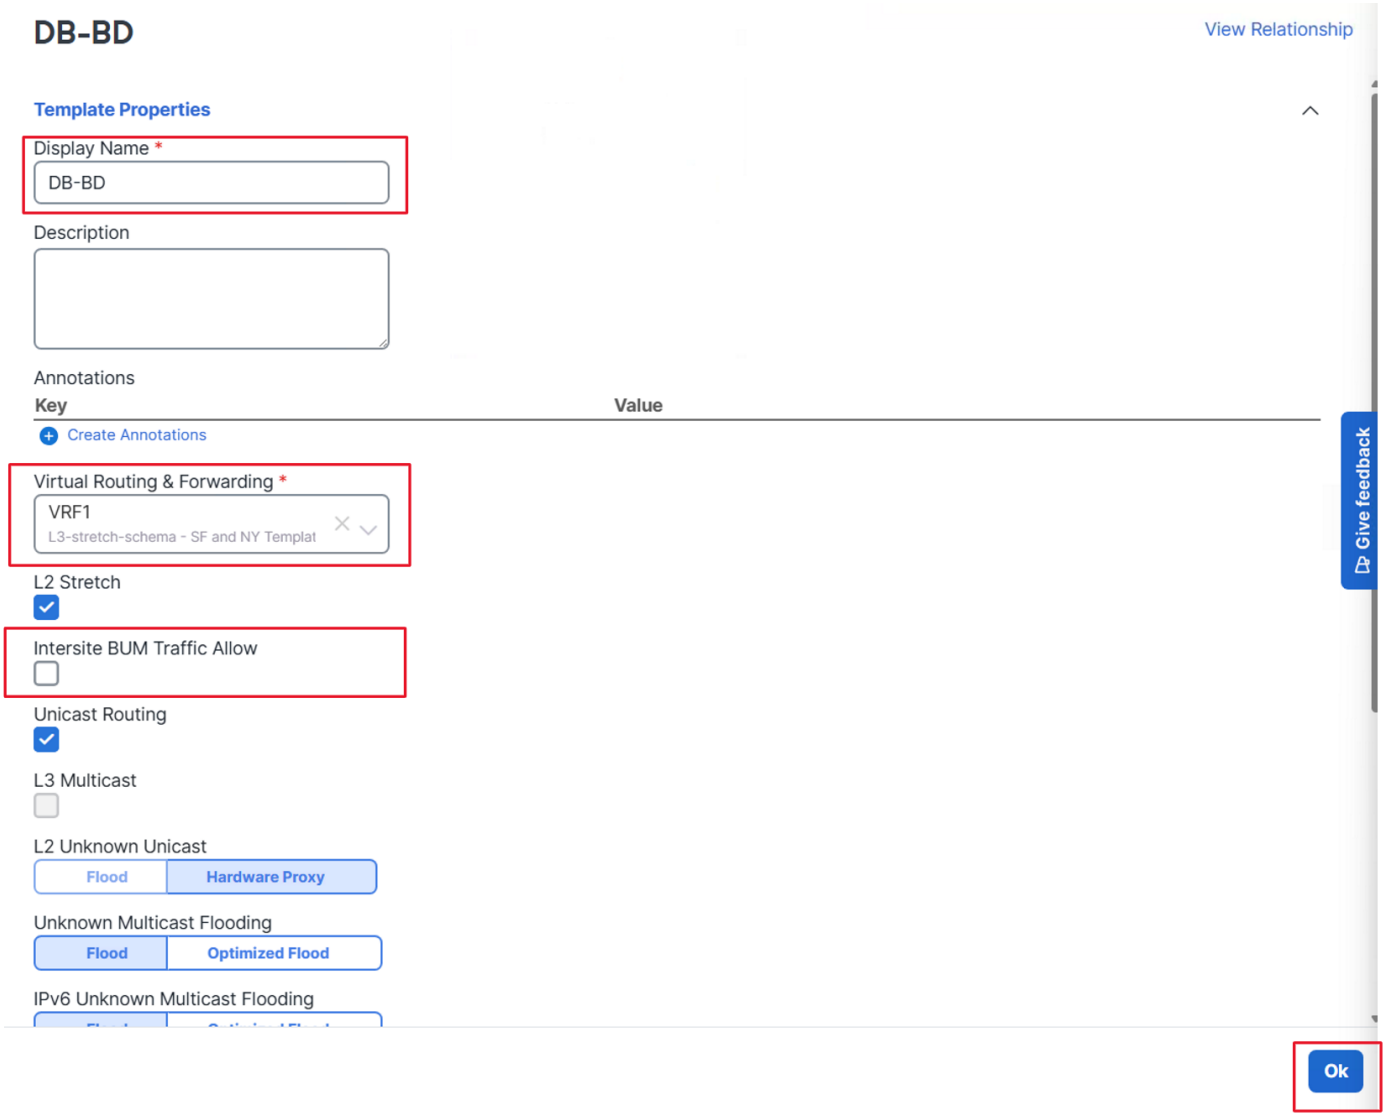

3. Click **Save Schema.**

> At this stage, a stretched Bridge Domain (**DB-BD**) and VRF (**VRF1**) have been created; however, no EPGs are currently shared across sites. As a result, the Webapp application profile has not yet been configured under the SF and NY Templates.
>
> The next step is to create stretched Application Profile (**Webapp**) and EPG (**DB**) to extend application connectivity across both sites.

4. Expand **Create Object,** and select **Application Profile**. In the Display Name, enter **Webapp**. Click **Ok.**

> The Application Profile name is case-sensitive, so be consistent with the name used previously.

5. Click **Save Schema,** then click **Create EPG**, in the Display Name enter **DB**, and in Bridge Domain select **DB-BD**. Click **Ok**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

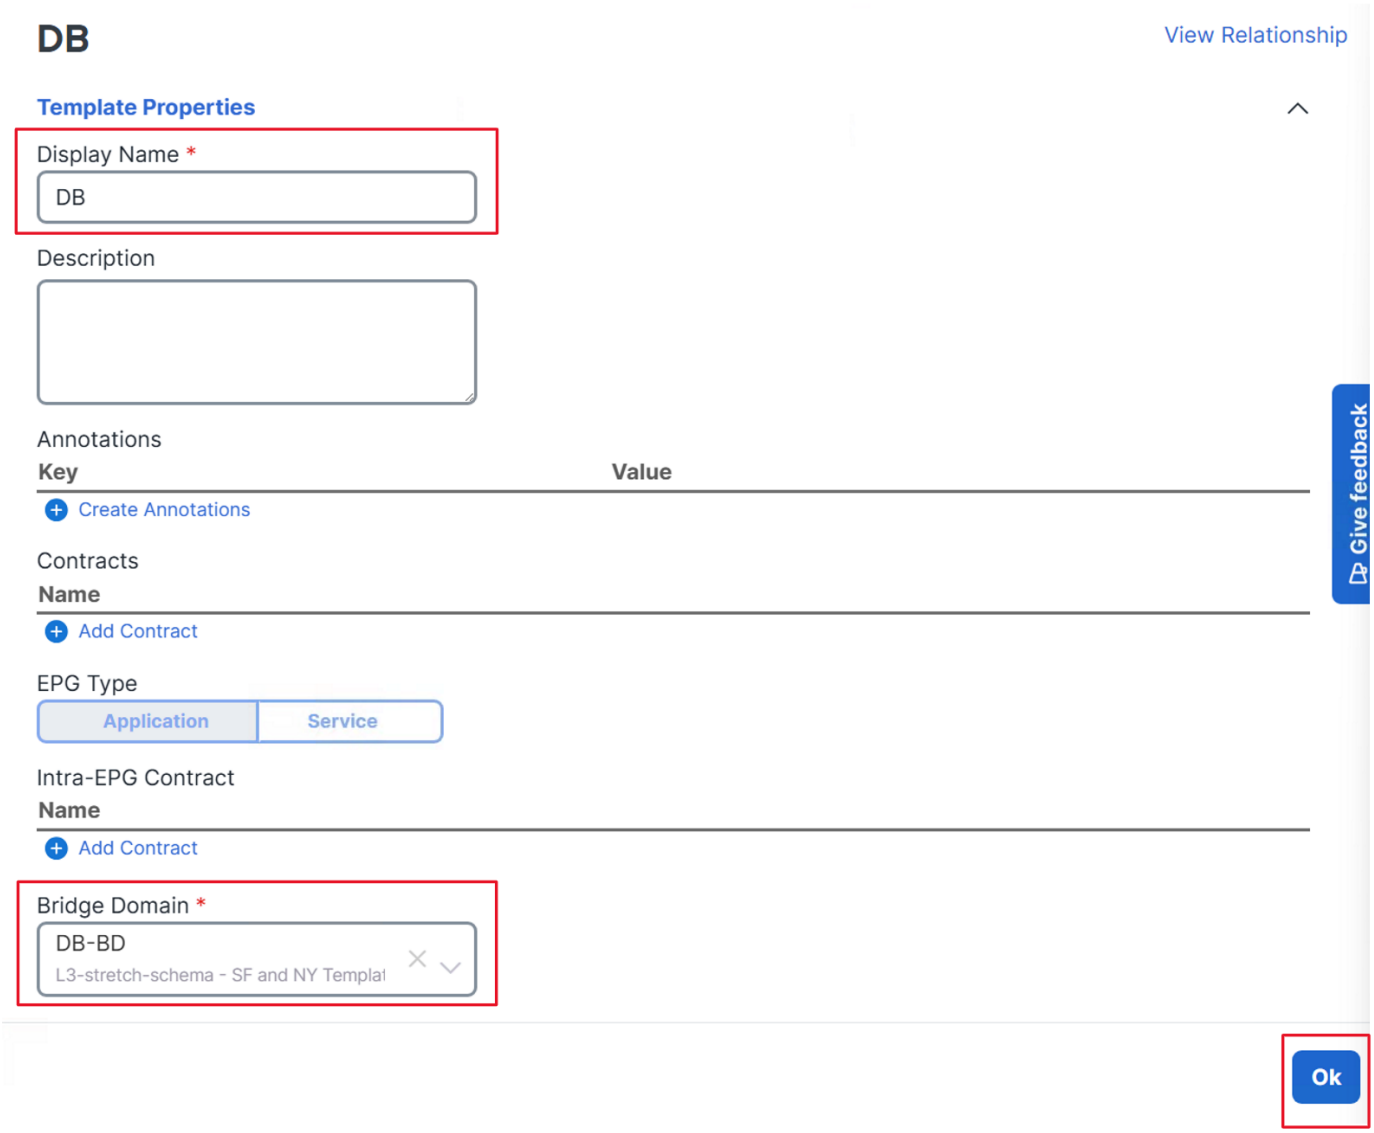

6. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.

Now verify the configuration on both sites to ensure that the stretched Bridge Domain (DB-BD), VRF (VRF1), and EPG (DB) have been properly deployed and synchronized across the San_Francisco and New_York sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

7. Log in to each site's APIC, then navigate to **Tenants** > **Pseudoco** > **Application Profiles** > **Webapp** > **Topology,** to verify the deployed objects and their inter-site connectivity.

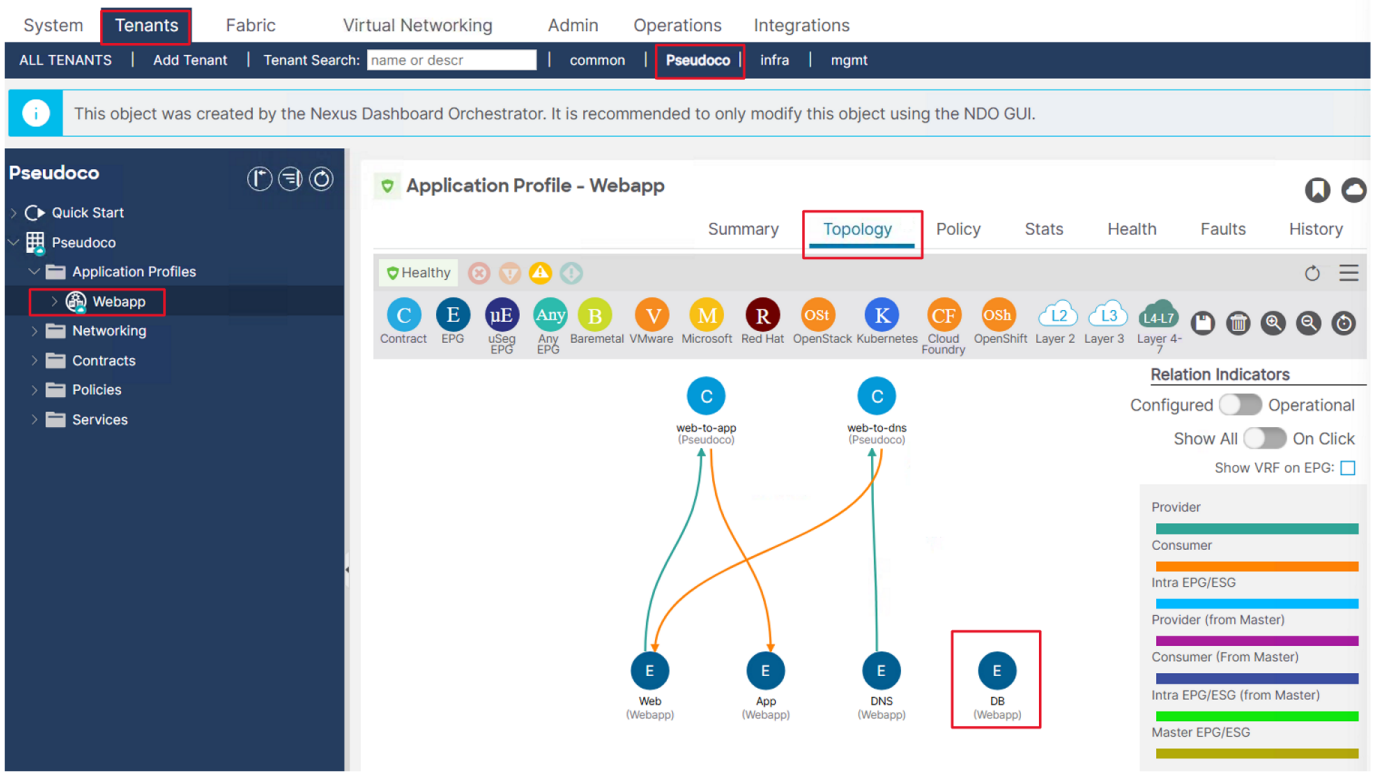

Value Proposition:
This scenario enables seamless IP mobility across sites using a stretched Bridge Domain while disabling BUM flooding to improve resiliency and control. It supports disaster recovery by allowing applications to be restored at a secondary site without IP changes, ensuring consistent Layer 3 reachability, fault isolation, and efficient multi-site operation.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 6. Importing Brownfield Tenant Configuration

This scenario directly addresses a common challenge in enterprise data center evolution: integrating existing, or "brownfield," ACI fabrics into a new multi-site architecture orchestrated by Cisco Nexus Dashboard. It demonstrates a powerful capability to seamlessly import an existing tenant's configuration — including Application Profiles, EPGs, Bridge Domains (BDs), and VRFs—from a brownfield site and then **extend these objects to a new, greenfield fabric.**

By leveraging this functionality, organizations can modernize and expand their ACI deployments **without service disruption.** This process ensures configuration consistency, policy alignment, and centralized orchestration across both the legacy and new environments. In this practical demonstration, the **San_Francisco fabric serves as the brownfield source**, while the **New_York fabric represents the new greenfield deployment**, showcasing a robust path for integrating established policies into a unified multi-site domain.

1. In the APIC for the San_Francisco (198.18.133.200) screen, click **Tenants** > **Brownfield**. This tenant has the following configuration:
   - VRF: **Brownfield-VRF**
   - Bridge Domain: **Brownfield-BD**
   - Application Profile: **AP**
   - Application EPG: **Brownfield-EPG**

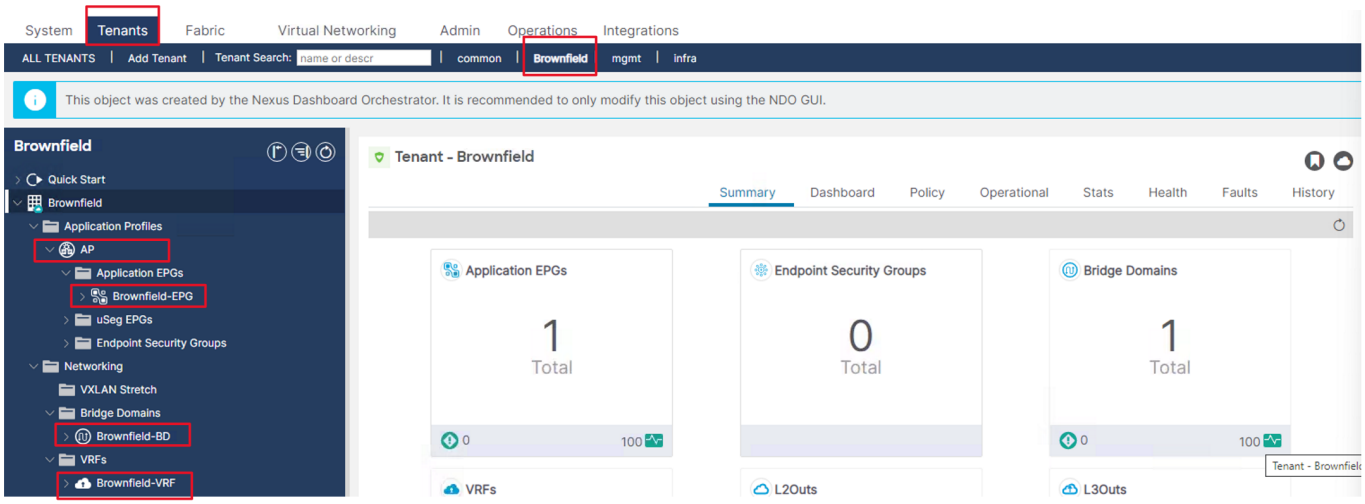

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2. Login to Nexus Dashboard and navigate to **Manage** > **Orchestration** > **Tenants > Actions > Create Tenant**.

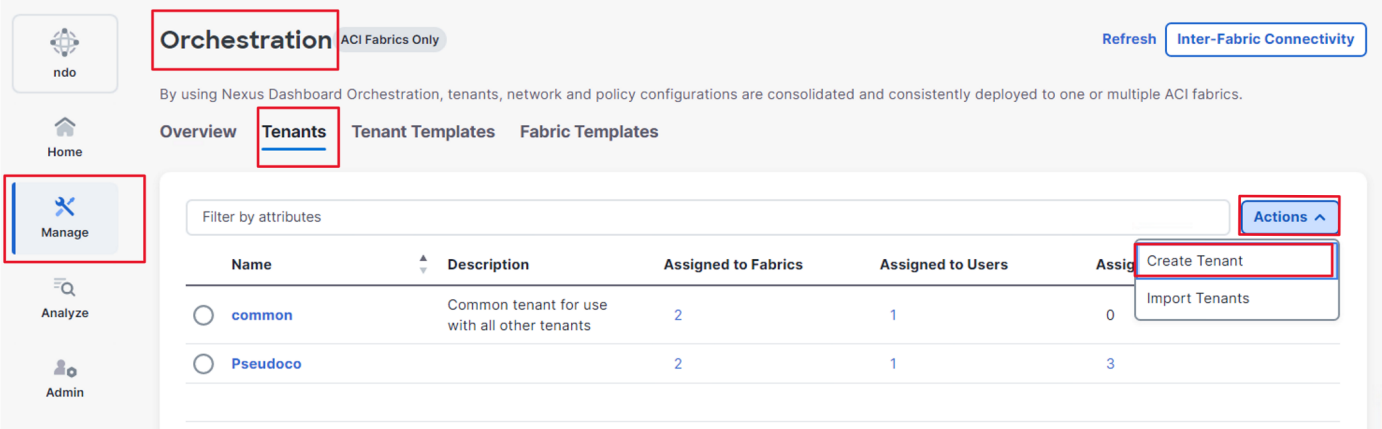

3. Enter the Display Name **Brownfield** and select both San_Francisco and New_York as the associated sites.

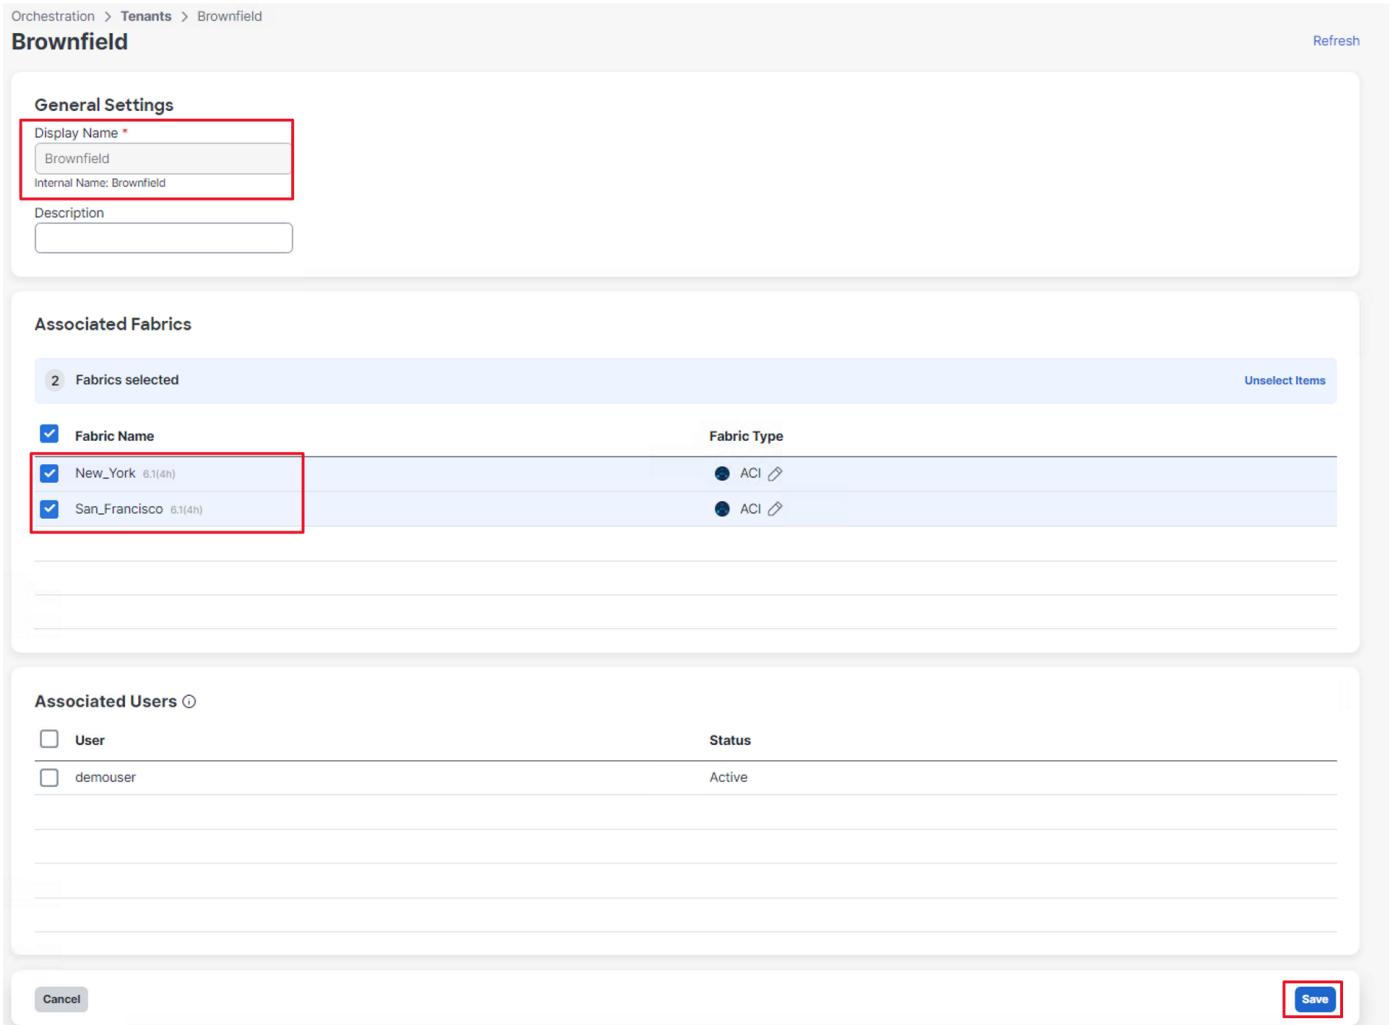

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

NOTE: The tenant name created in the Nexus Dashboard must exactly match the tenant name in the brownfield fabric from which the configuration will be imported. After creating the tenant, associate it with both sites—the brownfield site (source) and the greenfield site (destination)—since the configuration will be imported from one site and stretched to the other.

4. In the Nexus Dashboard window, navigate to **Manage** > **Orchestration** > **Tenant Templates** > **Applications > Actions > Create Schema.** In the Name display enter **Migration-Schema**. (This new schema will be used to perform the import of the configuration from San_Francisco into a new **Migration-Template**.)

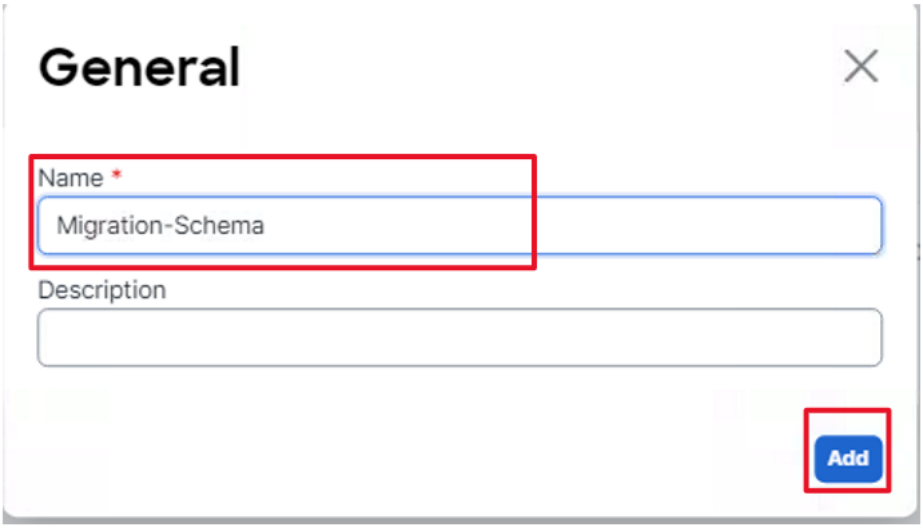

5. Click **Create New Template**.

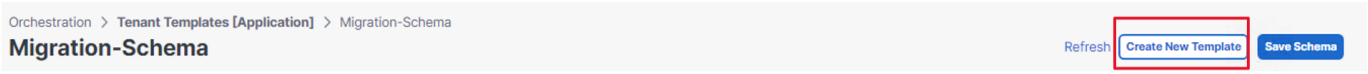

6. In the Display Name field, enter **Brownfield**. In **Select a Tenant** choose **Brownfield**. Click **Next**, followed by **Continue to Template**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

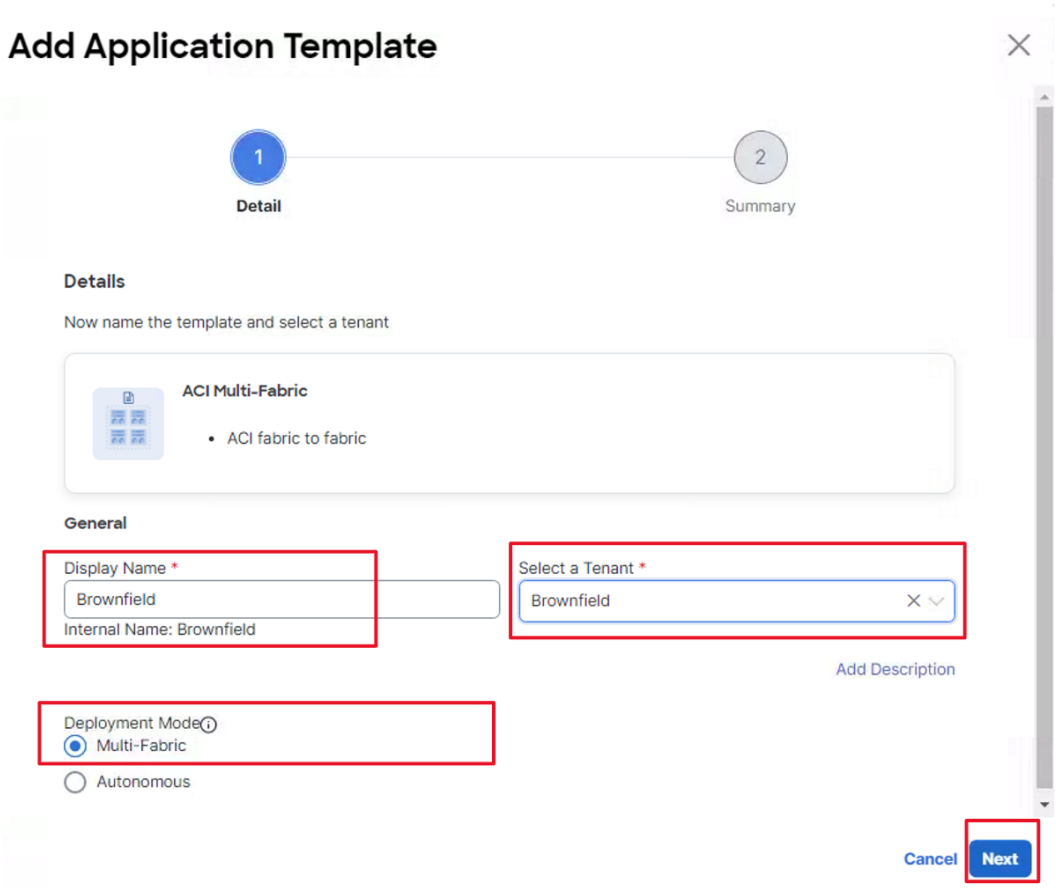

7. Click **Save Schema**.

NOTE: The Import button appears under Template Properties.

8. Click **Import** and select **San_Francisco** to import the **Brownfield** tenant configuration into the Orchestrator.

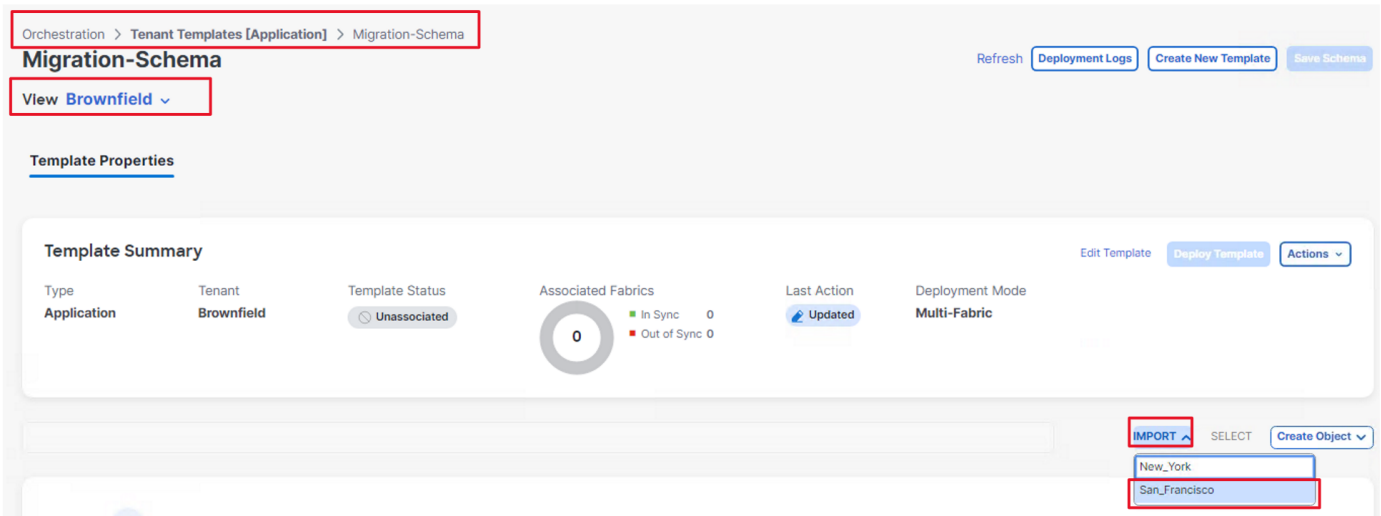

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

9. In the resulting window, select the Application Profile and select **AP**, then click **Import**.

NOTE: The Import Relations toggle automatically switches to ON to import all the objects associated with the Brownfield-AP application profile.

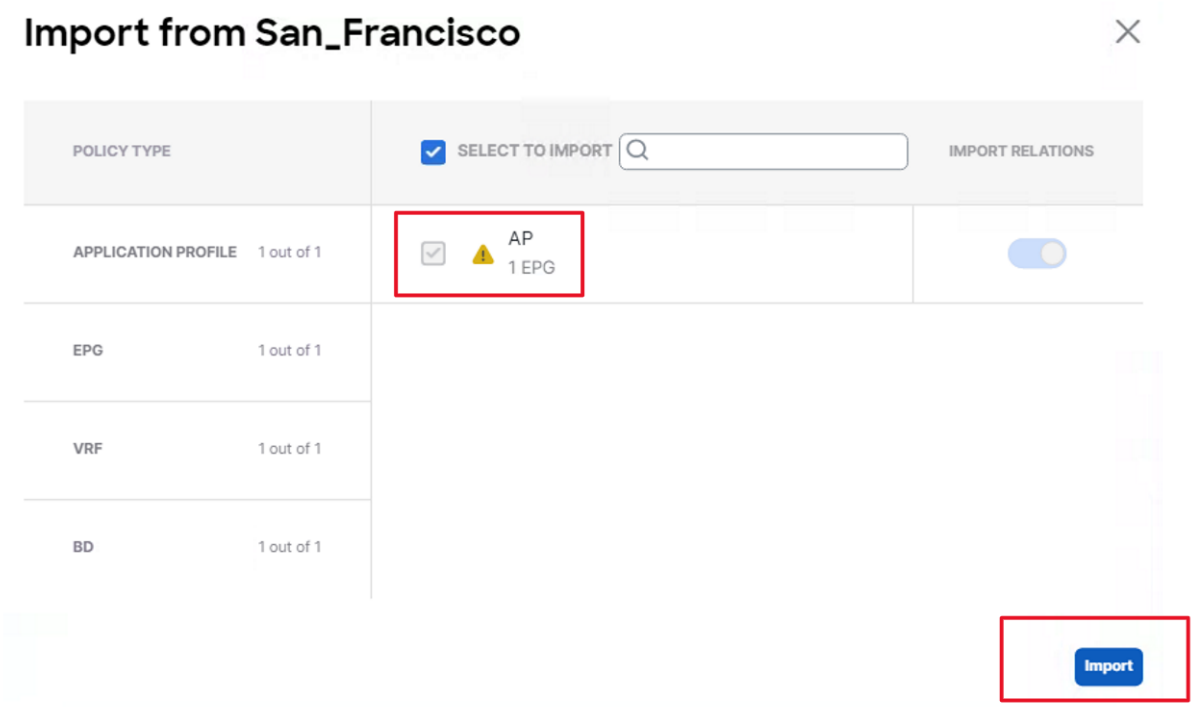

10. Click **Save Schema**.
11. Select the **Brownfield-BD** bridge domain to verify that the configuration isn't stretched (which is expected, since it was imported from a specific site).
12. Enable the **L2 Stretch** option to extend it to the **Greenfield** site and click **Yes** to acknowledge the warning message. Then, select the **Intersite BUM Traffic Allow** checkbox to permit BUM traffic between sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

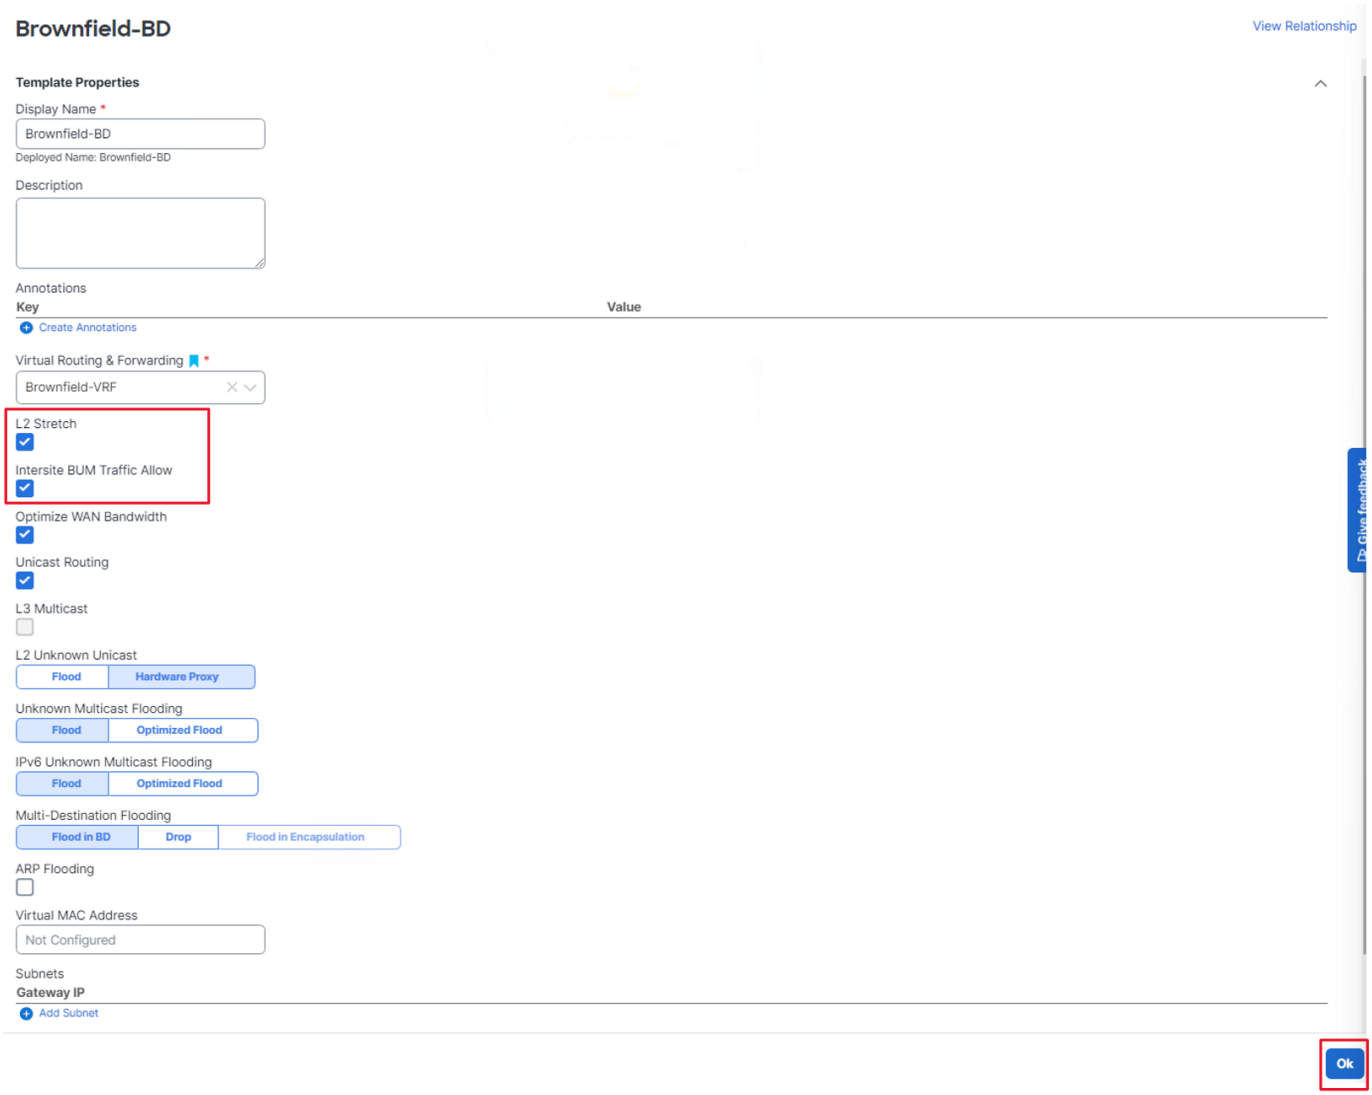

13. Now, associate the **Migration-Template** to the **New_York** site. Click **Actions**, then select **Add/Remove fabrics**, select **New_York**, and then click **OK**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

14. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.
15. Verify that the configuration is now correctly displayed in the **New_York** APIC controller (**198.18.132.200**). Navigate to **Tenants** > **Brownfield** to confirm that the imported objects have been successfully deployed to the Greenfield site (New_York).

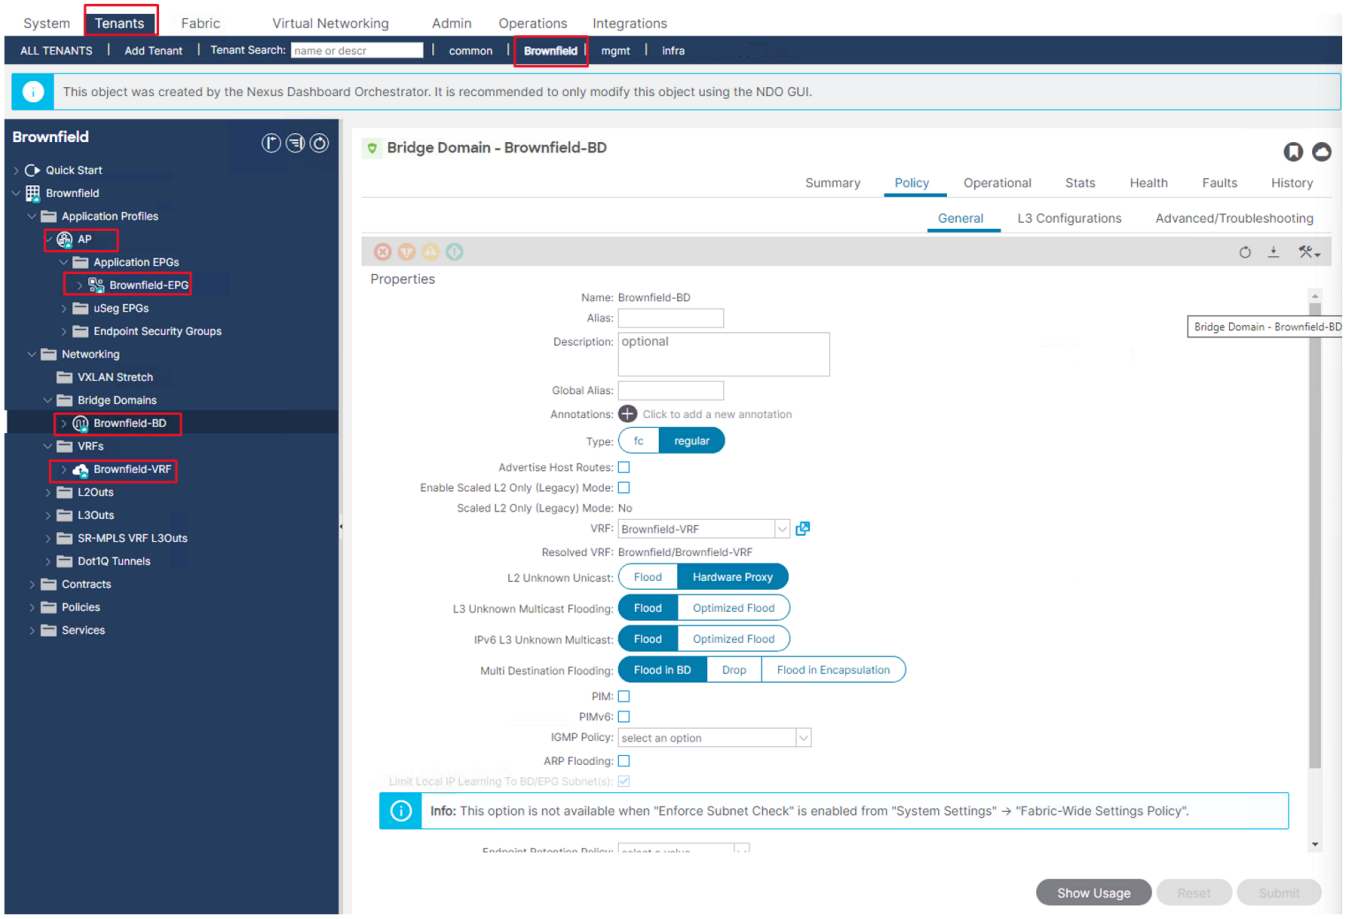

Value Proposition:
This scenario enables tenant-level configuration of import and extension between existing and new ACI fabrics using the Nexus Dashboard. It streamlines the integration of a brownfield deployment with a greenfield environment, allowing administrators to reuse and extend established policies, Application Profiles, EPGs, Bridge Domains, and VRFs without manual reconfiguration. By automating configuration of import and synchronization, this approach ensures operational continuity, uniform policy enforcement, and accelerated migration of workloads between fabrics. It provides a consistent and scalable method to onboard existing ACI deployments into a multi-site architecture, enhancing manageability and reducing migration risk.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 7. Application-Centric Segmentation with ESG

This scenario introduces **Endpoint Security Groups (ESGs)**, a transformative approach to security policy within Cisco ACI that moves beyond traditional network constructs. ESGs enable security policies to be defined around the functional roles of an application rather than individual EPGs, bridge domains, or physical boundaries.

In this demonstration, an application is represented by four distinct Endpoint Groups: **Web, App, DNS, and DHCP**. You will see how ESGs dramatically simplify segmentation by grouping these EPGs. For instance, the Web and App EPGs can be combined into a single "Frontend ESG," while DNS and DHCP form a "Shared Services ESG." This consolidation allows a single ESG-to-ESG contract to define the entire interaction between these tiers, significantly **reducing contract complexity** an1d **minimizing hardware resource consumption**.

Operating at the VRF level, ESGs form logical security domains whose membership can span multiple EPGs or even multiple sites, all governed by a single, consistent policy boundary. This approach provides a more scalable and application-aligned security model, ensuring that as your multi-site environment evolves, your security policies remain clear, efficient, and directly reflective of your application architecture.

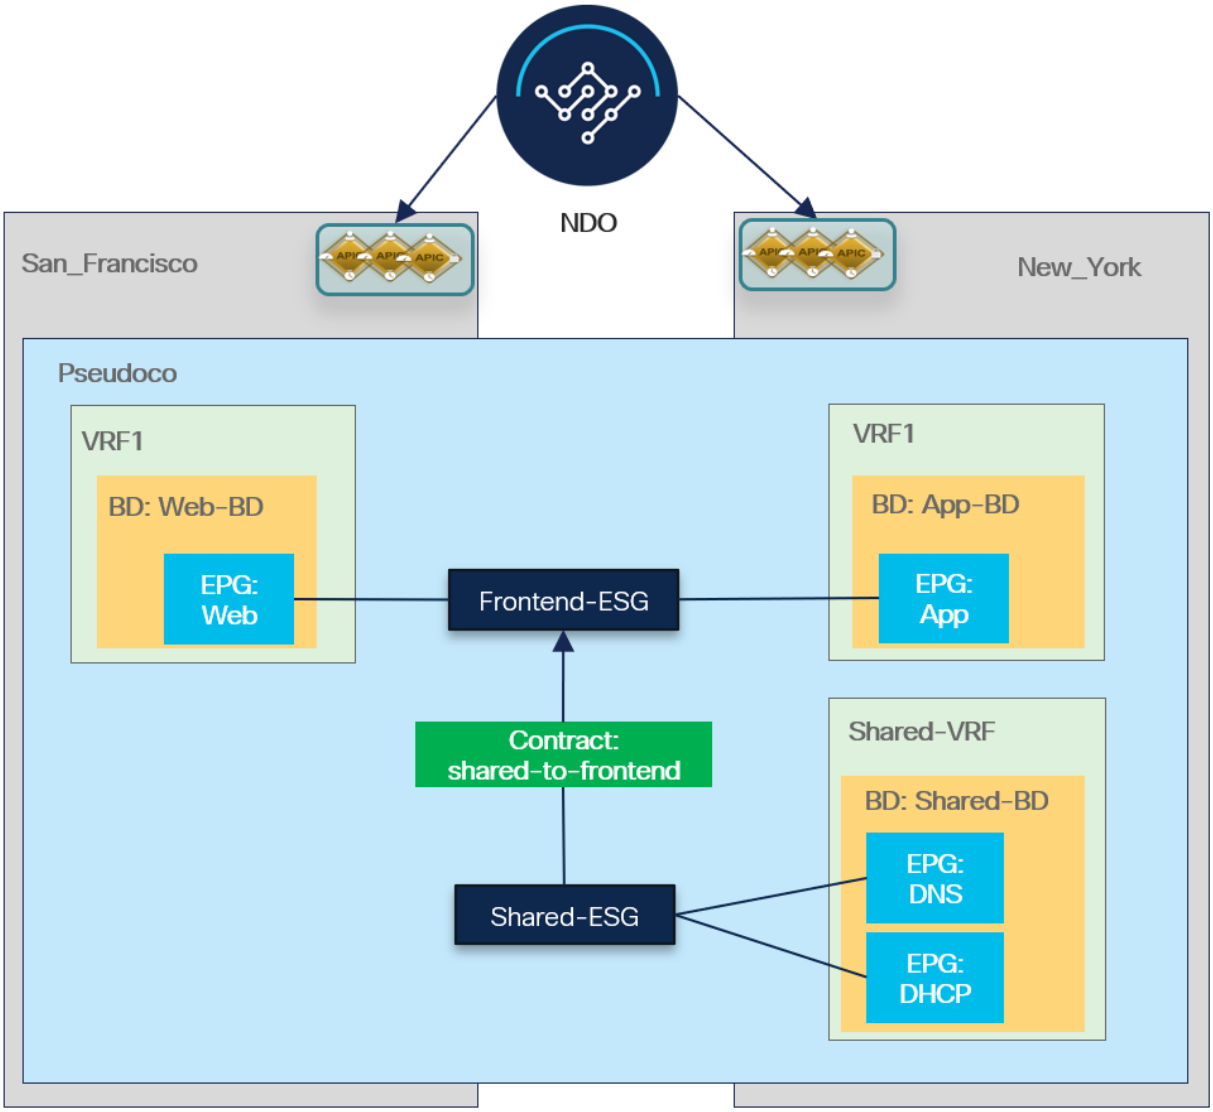

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

1. In the Nexus Dashboard window, navigate to **Manage** > **Orchestration** > **Tenant Templates** > **Applications**, then double click **L3-Stretch-schema**.
2. Expand the **View** menu, then select **NY Only**.
3. In the **Application Profile Webapp** section, click **Create EPG**.
4. Enter **DHCP** in the Display Name, select **Shared-BD** in the bridge domain field and click **OK**.
5. In the **Application Profile Webapp** section, click **Create ESG**.
6. Enter **Shared-ESG** in the Display Name, select **Shared VRF** in the VRF field.
7. In the EPG Selectors section, click **+** sign **Add selector**. Add the DNS and DHCP EPGs, and then click **OK**.

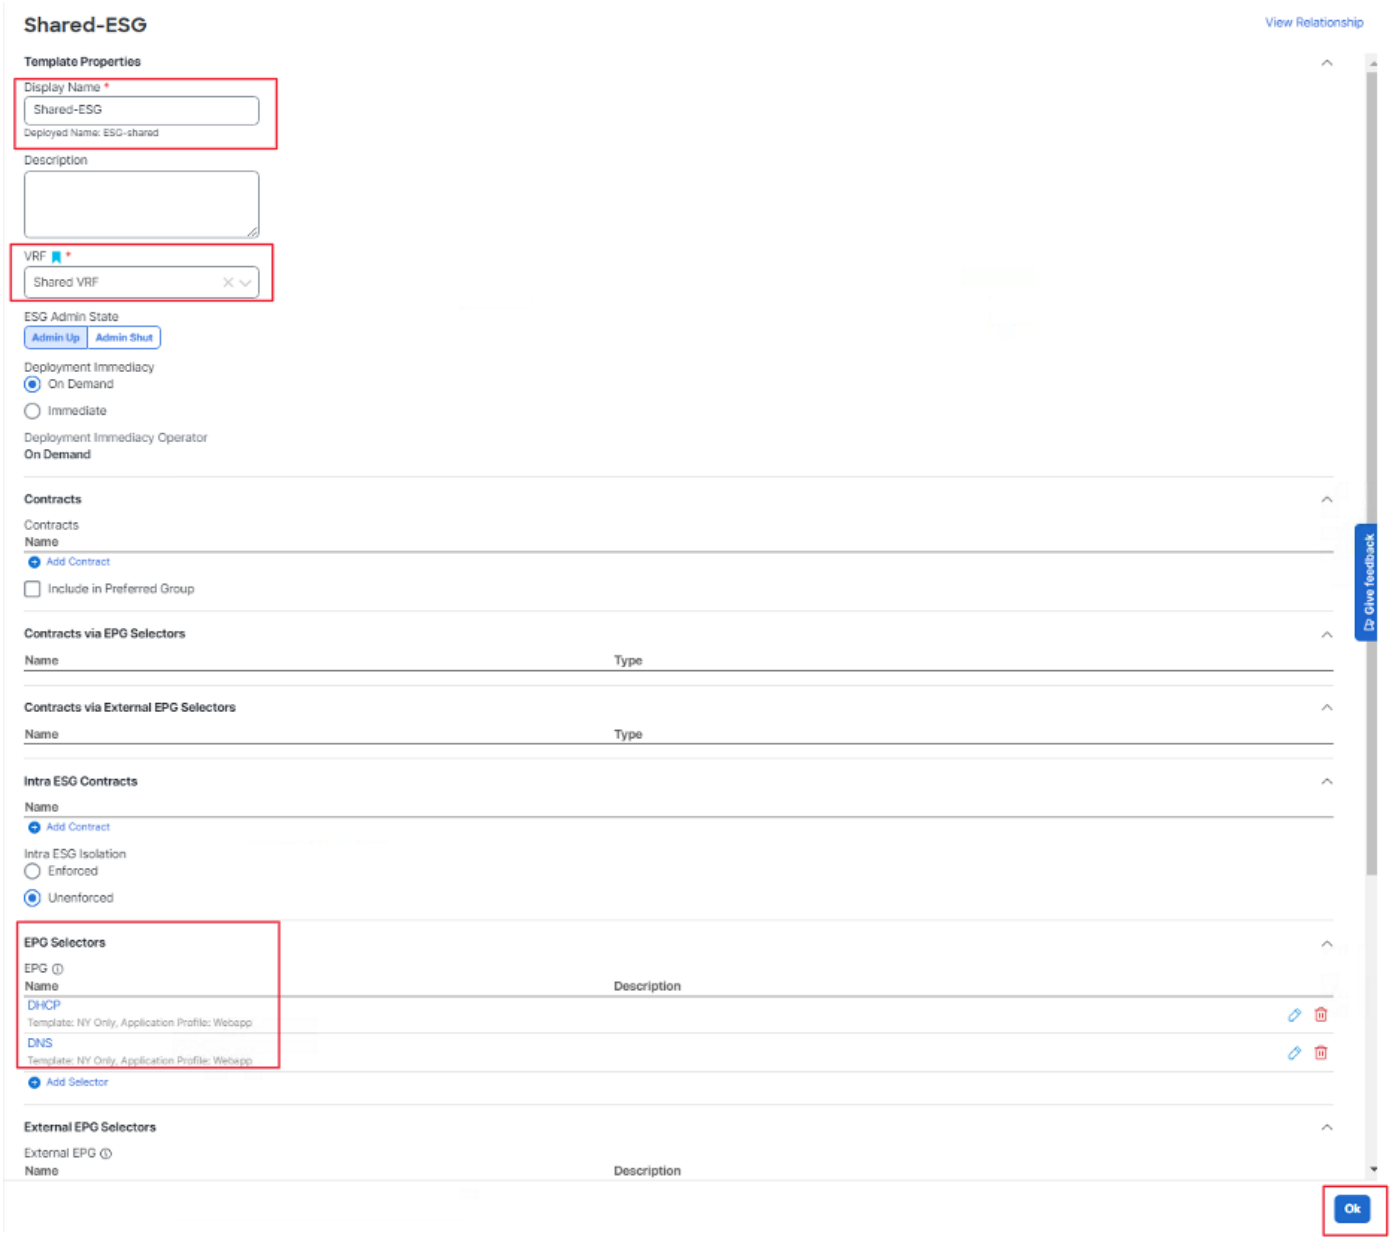

8. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.
9. In the **View** menu, select **SF and NY Template**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

10. In the **Application Profile Webapp** section, click **Create ESG**.
11. Enter **Frontend-ESG** in the Display Name, select **VRF1** in the VRF field.
12. In the **EPG Selectors** section, click **+** sign **Add selector**. Add the App and Web EPGs, and then click **OK**.

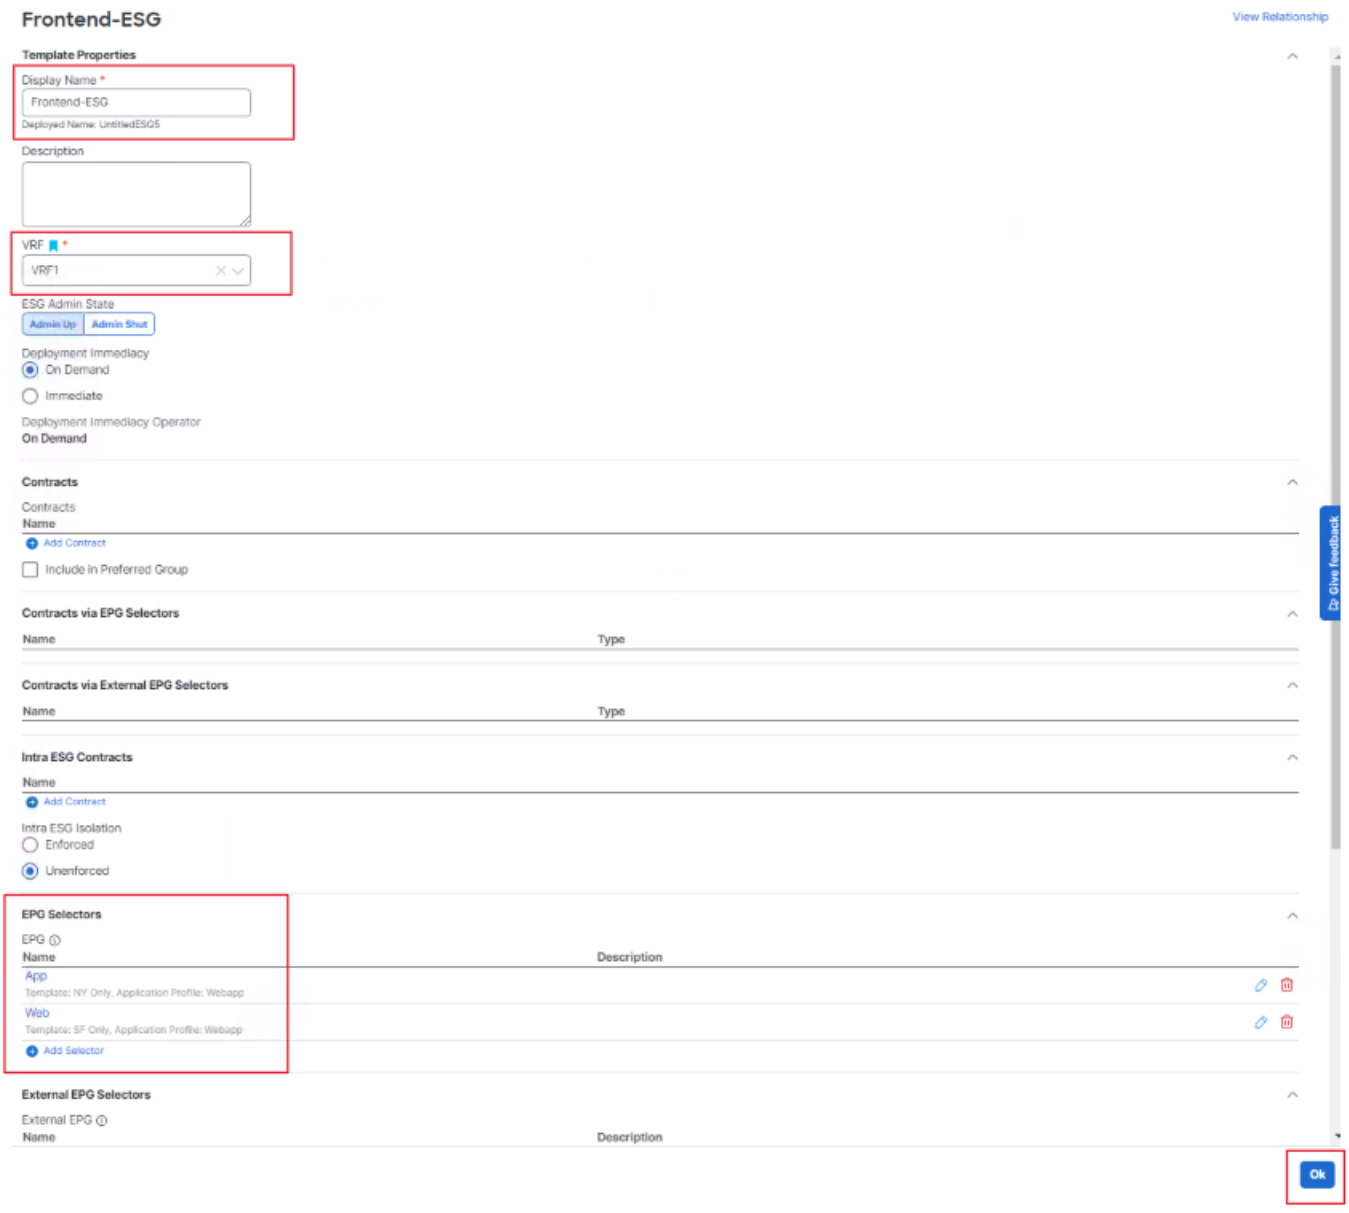

NOTE: After migrating the EPGs into the ESGs, their associated contracts are automatically inherited by the corresponding ESG. When you open any ESG—such as Frontend-ESG or Shared-ESG—and review the Contracts via EPG Selectors, you will see the previously created contracts, such as web-to-app and web-to-dns. At this stage, you can create a new contract named shared-to-frontend between Frontend-ESG and Shared-ESG, designating Shared-ESG as the Provider and Frontend-ESG as the Consumer. After the new contract is applied, the inherited contract can then be removed.

13. In **SF and NY Template**, navigate to **Contracts** section, click **Create Contract**.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

14. Enter **shared-to-frontend** as the Display Name, select **VRF** for the Scope, choose **any** in the Filter Chain, and then click **OK** to save.
15. In the **Application Profile Webapp** section, select **Frontend-ESG**, then go to the **Contracts** section.
16. Click **Add Contract**, choose **shared-to-frontend** and set the type to **Consumer**, then click **OK** to save.
17. Click **Save Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.
18. Expand the **View** menu, then select **NY Only.**
19. In the **Application Profile Webapp** section, select **Shared-ESG**, then go to the **Contracts** section.
20. Click **Add Contract**, choose **shared-to-frontend** and set the type to **Provider**, then click **OK** to save.
21. Click Save **Schema**, then click **Deploy Template** > **Deploy** > **Deploy Out of Sync Templates**.

NOTE: Now that we have migrated the frontend EPGs—such as **app** and **web**—into a single ESG, and the shared services—such as **DHCP** and **DNS**—into another ESG, we have also introduced a new contract to replace the inherited ones. Since this new contract is now in place, the inherited contracts associated with the EPGs can be removed. We can proceed by navigating each ESG and removing those inherited contracts. With that complete, we can move forward and verify the configuration on each ACI site.

Value Proposition:

Ths scenario introduces Endpoint Security Groups (ESGs) that enable security policies to be defined based on the functional roles within an application rather than on individual network constructs. ESGs operate at the VRF level, creating logical security domains that can span multiple Endpoint Groups (EPGs), bridge domains, or sites, all governed by a single consistent policy boundary. This approach simplifies segmentation, reduces configuration complexity, and aligns enforcement directly with application architecture. For example, grouping Web and App EPGs into a frontend ESG and DNS and DHCP into a shared-services ESG allows the entire application interaction to be managed with just one ESG-to-ESG contract. This reduces the number of required contracts and ACL entries significantly, lowering hardware resource consumption and operational overhead. As the application scales, new EPGs inherit policies automatically without additional configuration, ensuring scalable, clear, and consistent policy enforcement across all sites.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Scenario 8. Migration of ACI Schema Policies from NDO 3.7 to NDO 4.4 via Configuration Restore

This scenario addresses a critical operational task for network administrators: **migrating existing Cisco Nexus Dashboard Orchestrator (NDO) deployments to a newer release**, specifically demonstrating the transition from NDO 3.7 (run on Nexus Dashbooard 2.1) to NDO 4.4 (on Nexus Dashboard 3.2) using a robust backup and restore workflow. This method is particularly vital for deployments running older NDO versions (pre-3.3(1) where in-place upgrades are not supported) or for virtual Nexus Dashboard clusters where a fresh deployment offers a cleaner, less disruptive upgrade path with a straightforward rollback option.

NDO offers two primary upgrade approaches: an **in-place** upgrade for compatible existing clusters, or deploying a **new Nexus Dashboard cluster and restoring the configuration**. This scenario focuses on the latter, showcasing its broad compatibility for transitions from MSO 2.x/3.x and NDO 3.x/4.x releases to NDO 4.x.

A key aspect of this migration is understanding the evolution of NDO's policy model. Beginning with NDO 4.0(1), the system enforces a **more robust and structured policy architecture**. This includes strict dependency ordering (e.g., a Bridge Domain cannot deploy without its referenced VRF), disallowing circular dependencies, and automatically restructuring templates to comply with these new rules. This ensures greater consistency and predictability in policy deployment.

In this demonstration, you will experience the complete migration workflow:

- **Establishing a Baseline:** The process begins in an NDO 3.7 environment, where an ACI fabric is onboarded and a preconfigured development snapshot is restored to create a consistent baseline of schemas, templates, and ACI objects.
- **New NDO 4.x Deployment:** A new NDO 4.x cluster is then deployed, and the same ACI site is onboarded to this new instance.
- **Configuration Restore and Validation:** The configuration backup taken from NDO 3.7 is restored into the NDO 4.x environment. During this crucial step, the Orchestrator automatically examines imported objects, validates dependencies, resolves implicit stretching, and restructures any templates requiring modification for compliance with the newer release. A detailed validation report is presented, and the reconstructed schemas and templates are reviewed.
- **Decommissioning Legacy System:** Once the migration is verified, the legacy NDO 3.7 environment can be cleanly decommissioned by disassociating the ACI site, placing it in unmanaged mode, and removing it.

This scenario highlights how to achieve a seamless and controlled transition to NDO 4.x, preserving your existing policy structure while adapting it to the new platform's enhanced capabilities. It significantly reduces migration risk and operational disruption, ensuring uninterrupted operations during critical infrastructure upgrades.

1. Open a web browser and navigate to the bookmarks bar, click **NDO_3.7**. Alternatively, enter the IP address directly: https://198.18.134.200. In the login screen, enter the following:
   - Username: **admin**
   - Password: **C1sco12345**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2. After logging in, a “What’s New in Nexus Dashboard 2.1” information window is displayed. Click **Begin Setup**, then close it by clicking the X in the upper-right corner.
3. Navigate to the left-hand menu and click **Admin Console**.

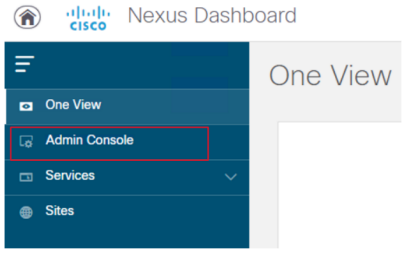

4. From the Admin Console navigation pane, select **Sites**.
5. In the upper-right corner, click **Add Site** to begin onboarding San Francisco ACI Site.

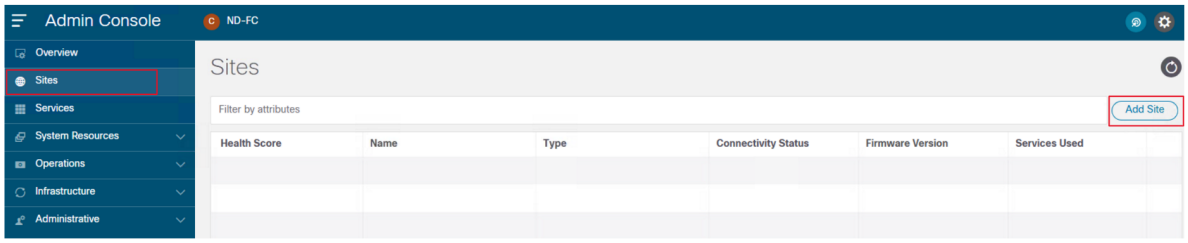

6. In the Select Type field, choose ACI and enter the following site information. Click **Save**:
   - Name: **SF**
   - IP Address: **198.18.133.200**
   - Username: **admin**
   - Password: **C1sco12345**

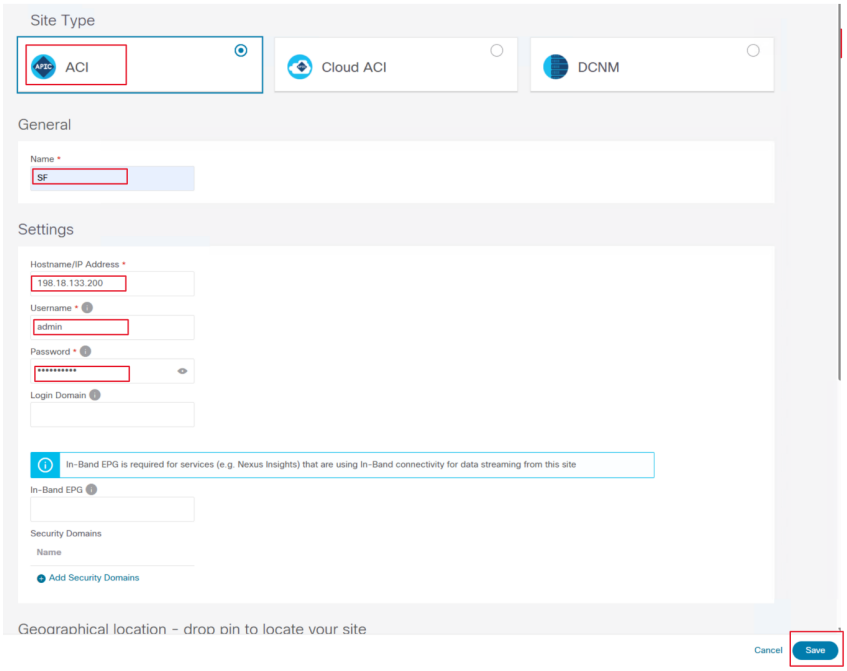

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

7. Wait a few seconds until the **Connectivity Status** indicator turns **UP** and **green**, confirming that the site has been successfully onboarded.
8. Navigate to **Services** from the left-hand menu, and under **Installed Services**, click the **Open** icon to launch the Nexus Dashboard Orchestrator interface.

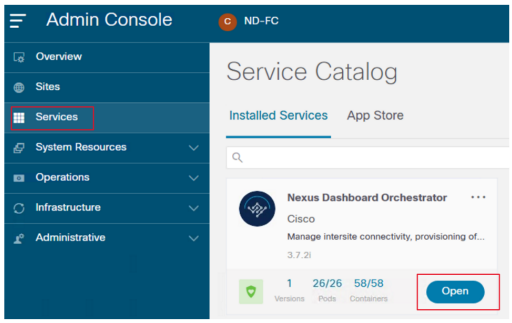

9. Expand **Infrastructure**, select **Sites**, locate the **SF** site, and change its **State** from **Unmanaged** to **Managed** to place the fabric under NDO policy control.

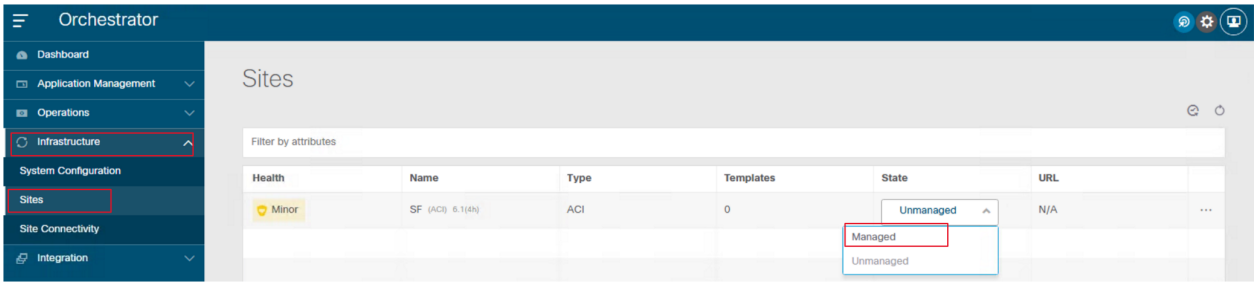

10. A new window will open prompting you to specify the Site ID; enter **1**, then click **Add**. Wait a few seconds until a green notification confirms that the site has been successfully switched to Managed mode.

NOTE: If an error is returned indicating an incorrect site ID, use the site ID value displayed in the error message.

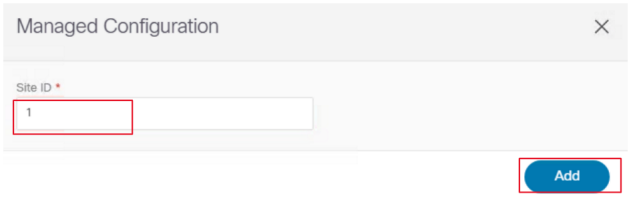

Instead of manually creating a new schema, templates, and ACI objects, you may choose to roll back to a pre-configured configuration snapshot that was previously generated for this environment. This snapshot already contains the required schema, templates, and policy objects. Alternatively, if preferred, you can proceed by building your own schema and templates from scratch.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

11. To restore the pre-configured environment, navigate to **Operations > Backup and Restore** from the left-hand menu. Locate the available backup in the list, click the **three-dot menu** on the right, and select **Roll Back to This Backup** to initiate the restoration.

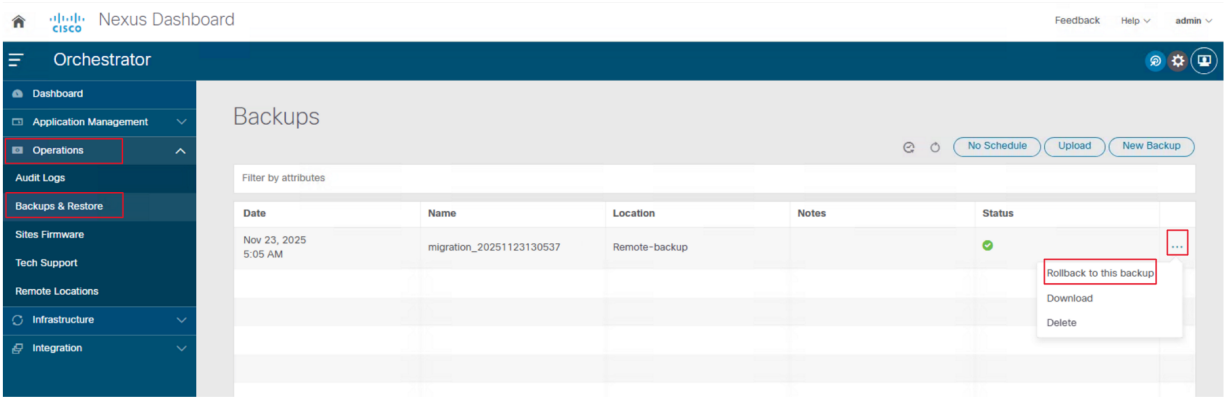

12. It will take **about 5 minutes** for the rollback process to complete. Once finished, a **green success notification** will appear confirming that the backup has been successfully restored. Click **OK** to proceed.

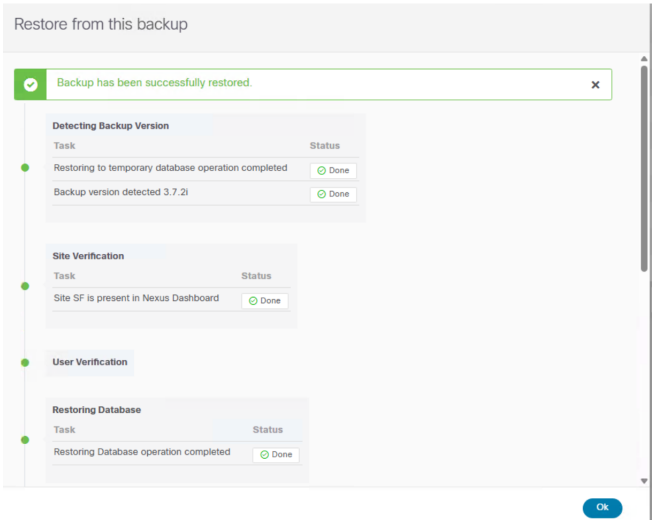

13. Expand **Application Management** from the left-hand menu and select **Schemas**. In the list of available schemas, locate **SCH_DEV** and double-click it to open the pre-configured configuration.

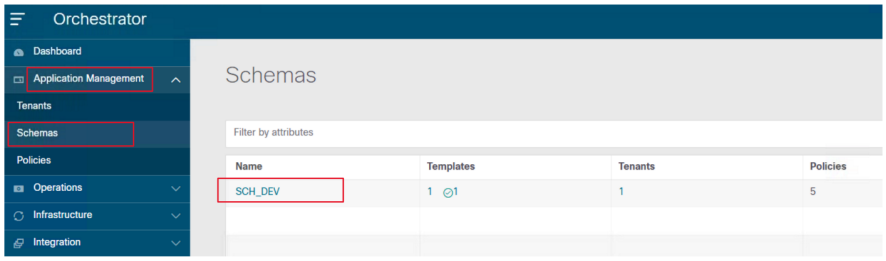

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

14. In the schema window, select the template **TMP_APP_DEV** from the left-hand panel and review the pre-configured ACI objects. When ready, click **Deploy to Sites**.

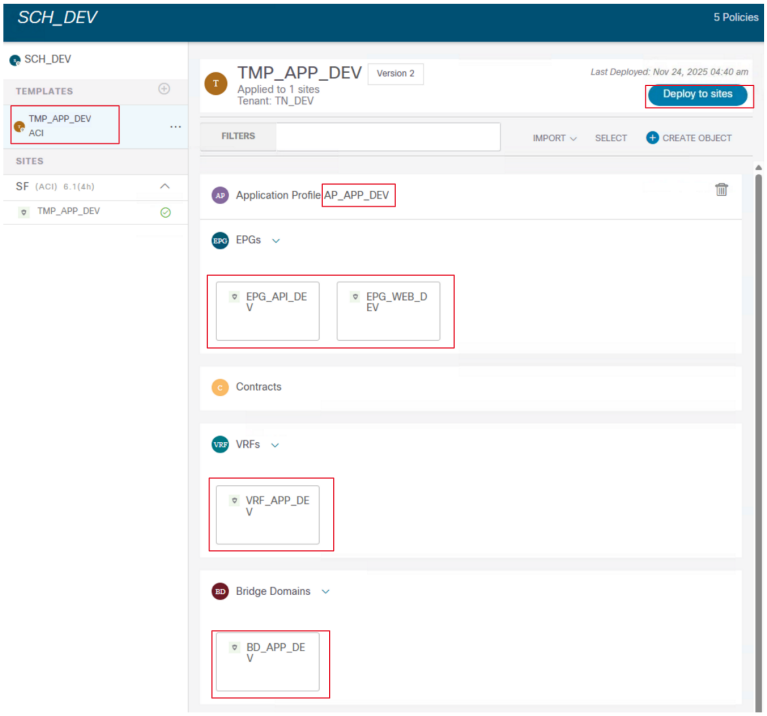

15. In the pop-up window, review the list of ACI objects that will be deployed to the **SF** site. Then click **Deploy** to push the configuration to the fabric.

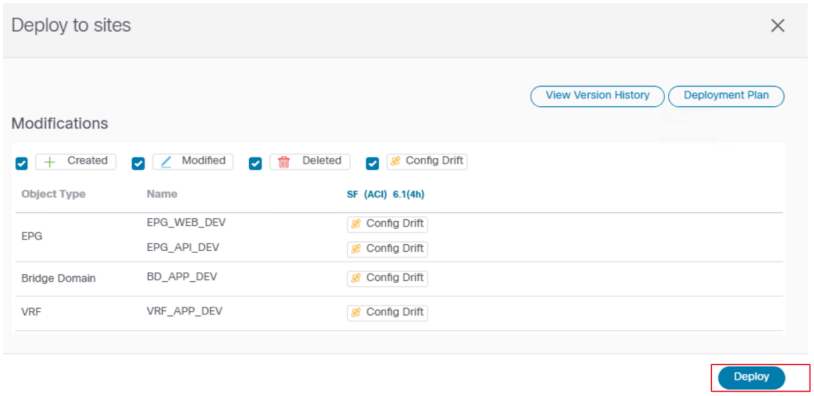

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

16. Log in to the **SF** ACI site (198.18.133.200) and navigate to the tenant **TN_DEV**. Review the tenant structure and confirm that all expected objects have been successfully created.

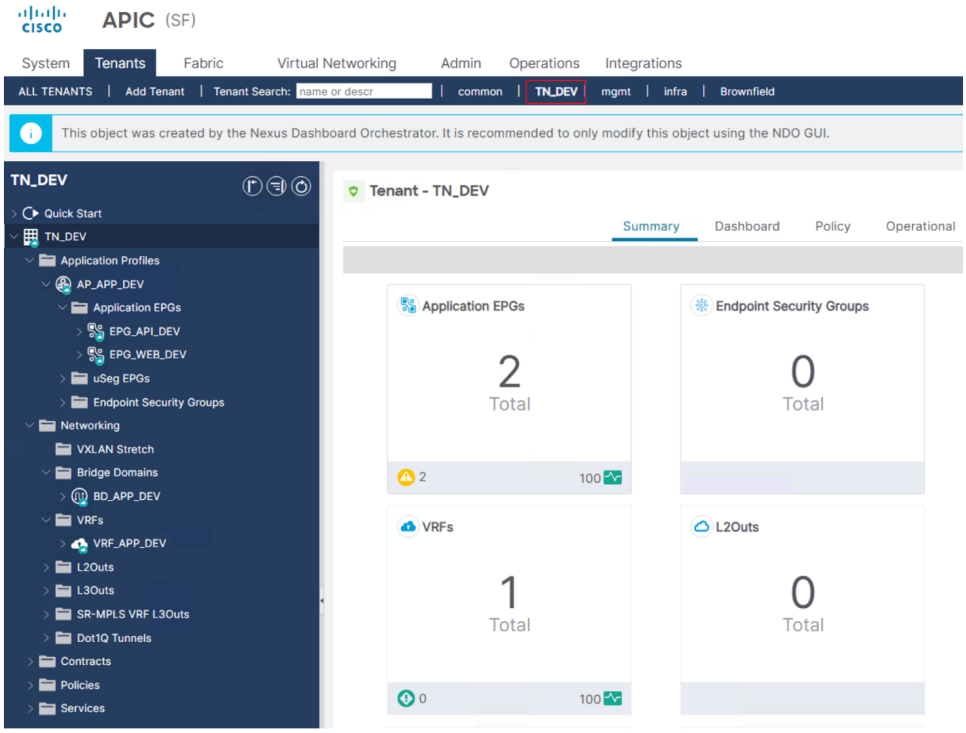

17. Download the NDO backup that will be used to migrate the existing configuration to NDO 4.4. To do this:
    a. Navigate to **Services > Orchestrator > Operations > Backups & Restore**.
    b. One backup job has already been run. Click the **three-dot** menu. Select **Download** to download the backup (migration-20251123130537.tar) to "C:\Users\demouser\Downloads".

NOTE: At this stage, we assume that you are working with an existing NDO 3.7 environment that already contains a development schema, templates, and associated ACI objects. The next steps will guide you through the process of migrating this configuration (from NDO 3.7 to NDO 4.4).

18. Open a web browser and navigate to the bookmarks bar, click **NDO_4.4**. Alternatively, enter the IP address directly: https://198.18.133.102 . In the login screen, enter the following:
    - Username: **admin**
    - Password: **C1sco12345**
19. Navigate to **Manage** > **Fabrics**, then click **Add Fabric,** and add the following configuration then click **Next:**
    - Hostname/IP Address: **198.18.133.200**
    - Username: **admin**
    - Password: **C1sco12345**

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

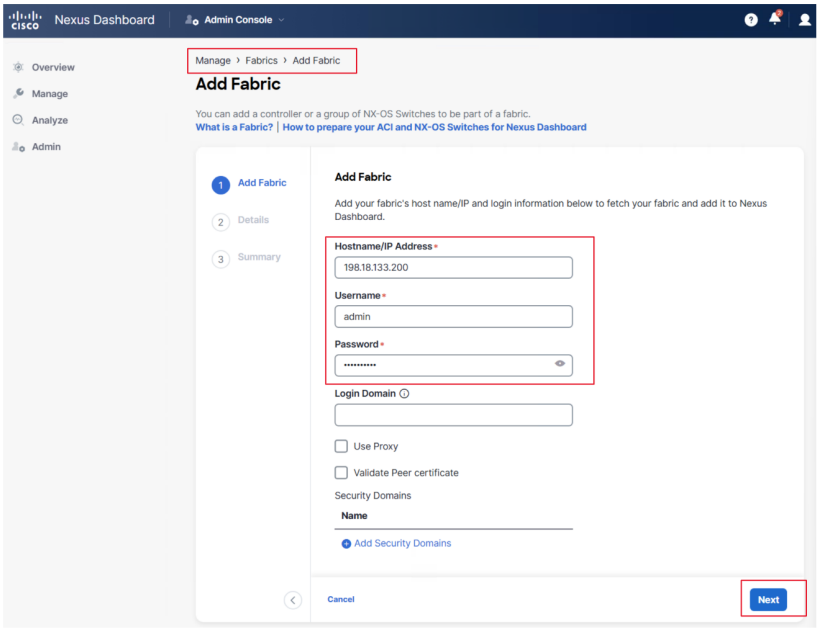

20. In the **Fabric Name** field, enter **SF**, select the Site Location, and click **Next** then **Save**.

21. Expand the **Admin Console** menu and choose **Orchestrator**.

22. Navigate to **Manage > Fabrics**, locate the **SF** site, and change its **State** from **Unmanaged** to **Managed**.

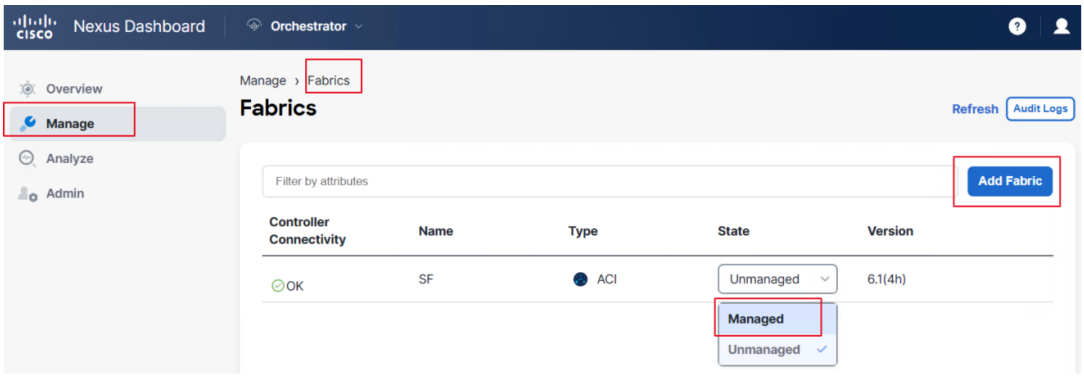

23. A new window will open prompting you to specify the Site ID; enter **1**, then click **Add**. Wait a few seconds until a green notification confirms that the site has been successfully switched to Managed mode.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

24. Navigate to **Admin** > **Configuration Import**, and click **Import backup File**.

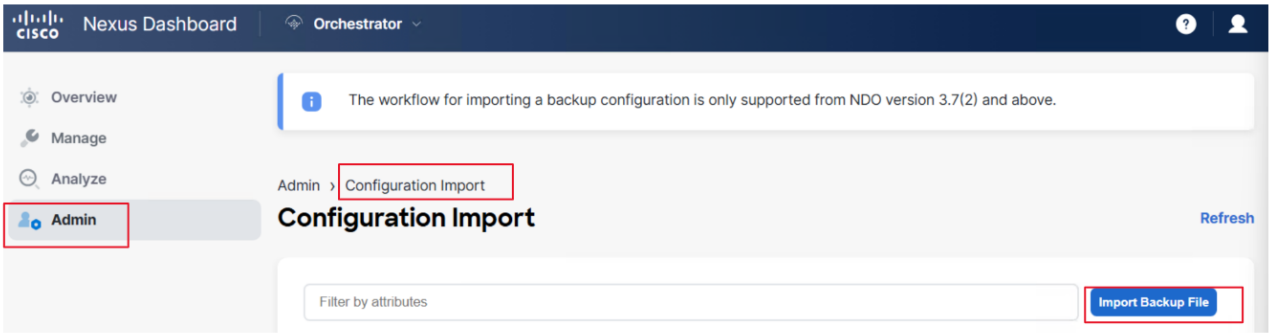

NOTE: The workflow for importing a backup configuration is only supported from NDO version 3.7(2) and above.

25. In the **Location** field, select **Local**, then click **Browse** to upload the NDO 3.7 backup file from the C:\Users\demouser\Downloads. Select “migration-20251123130537.tar”, then click **Restore** to begin importing the configuration into NDO 4.4.

26. A warning will appear indicating the backup is from a different NDO version and that templates will need redeployment after restore. When prompted click **Restore** to continue.

NOTE: The database rollback will now begin. This process may take around 7 minutes.

27. After the restore completes, review the displayed information—such as the backup version, status, and fabric verification details—then click **Restore Validation Required** to proceed.
28. A notification will appear indicating that NDO 4.0+ enforces template best-practice rules and has detected certain violations that will be automatically corrected during the restore. Review the message, then click **Restore and Continue** to proceed.

NOTE: Updating the NDO database is now in progress. This process may take around 5 minutes.

29. Once the database update is complete, a notification will appear. Click **OK** to continue.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

30. Go to **Manage > Tenant Templates > Applications**, then click on the schema **SCH_DEV**.

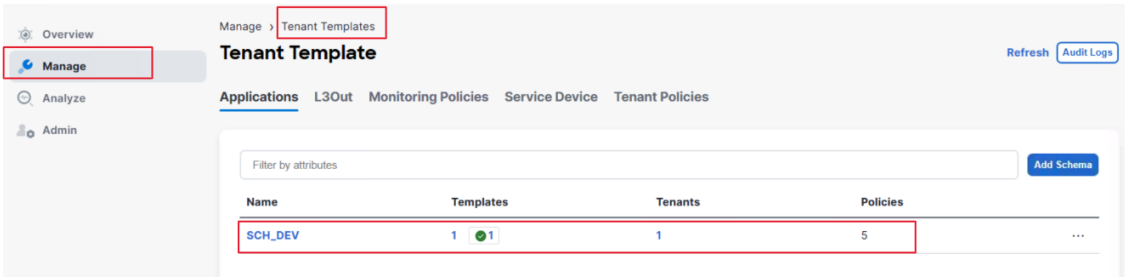

31. You can now review the schema, templates, and ACI objects within **SCH_DEV**. All configurations should match exactly as they existed in **NDO 3.7**.

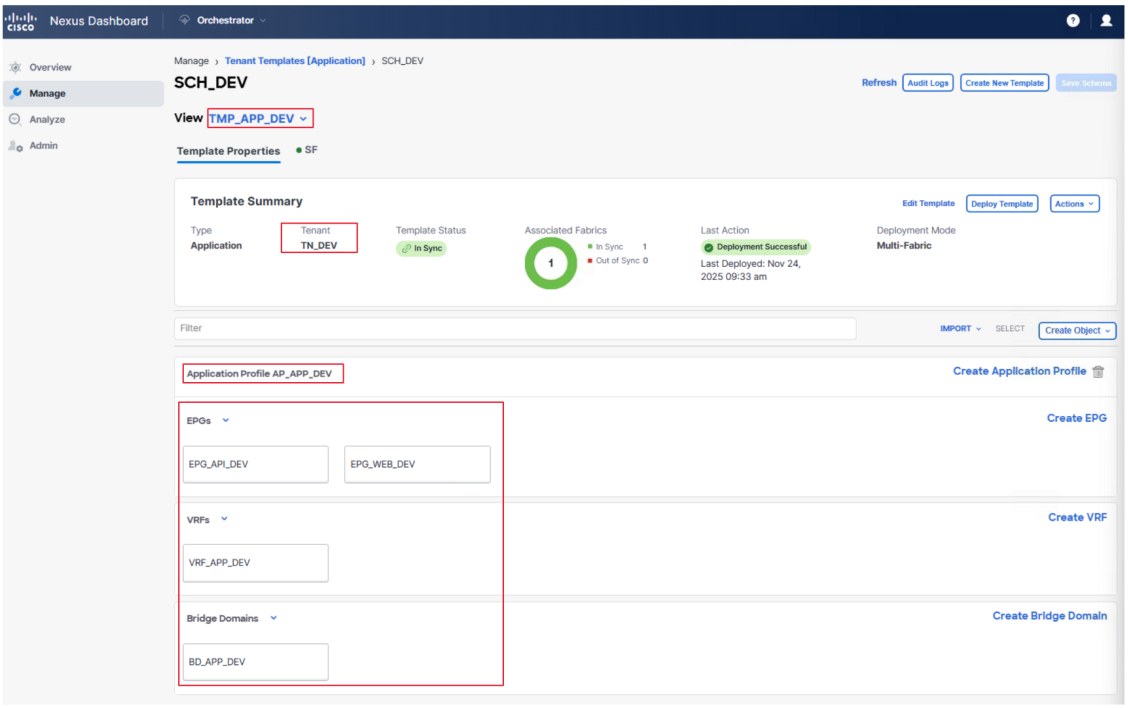

NOTE: With the migration to NDO 4.4 completed, you can now proceed with the decommissioning of the legacy NDO 3.7 environment. To retire the older orchestrator cleanly, first disassociate the site from any templates within NDO 3.7, then change the site's state from Managed to Unmanaged, and finally remove the site from the system. This ensures a proper and controlled decommissioning process.

Value Proposition:
This scenario covers the inter-release migration process using backup and restore functionality. A configuration snapshot from NDO 3.7 is restored into a newly deployed NDO 4.x cluster. The migration process automatically handles dependency ordering, object validation, and template restructuring, facilitating a seamless transition from older design paradigms to the more structured and dependency-aware architecture introduced in NDO 4.x. This model supports uninterrupted operations during upgrades, ensuring that existing configurations are preserved and adapted to the new platform's capabilities, thus reducing migration risk and operational disruption.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Key Takeaways

- **NDO centralizes and standardizes multi-site ACI operations:** It provides a single control system for provisioning, managing, and validating configuration across distributed fabrics.
- **Stretched and site-local templates define clear architectural boundaries:** They allow administrators to selectively share or isolate Layer 2 and Layer 3 constructs depending on application and operational requirements.
- **Layer 3-only multi-site designs reduce risk and improve scalability**: By keeping Layer 2 domains local, organizations avoid broadcast expansions while enabling seamless cross-site routing.
- **Shared-services VRF designs simplify cross-tenant or cross-VRF service consumption:** They enable centralization of common services without compromising segmentation.
- **IP mobility is achievable without traditional Layer 2 stretching pitfalls:** Stretched BDs with BUM flooding disabled allow disaster recovery and VM relocation with controlled propagation of failure domains.
- **Brownfield integration is fully supported:** Existing tenants can be imported into NDO, allowing modernization of legacy fabrics while preserving business-critical workloads.
- **ESGs provide dramatically more efficient and application-aligned segmentation:** They minimize the number of ACL entries, simplify policy management, and scale far more effectively than EPG-based segmentation alone.
- **Flexible Migration Paths:** Both upgrade methods—direct platform upgrade or Backup and Restore—enable seamless transitions to newer NDO releases while retaining existing policy structure.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Appendix A. Resetting APIC Simulator

This demo uses a ACI Simulator VM that contains the APICs, leaves, and spines. All of the configuration is temporary. When it is rebooted, the VM will take you back to the initial configuration steps. If you are having issues with the ACI Simulator, or perhaps you want to start over without scheduling a new session, you can simply reboot the VM.

1. In Cisco dCloud, click **My Hub > Sessions**, and then click **View** for the running demo.
2. You will be taken to Topology Builder, a service that hosts the virtual machines used in the demo. Find and click on the **ACI Simulator**

    c. In the left pane, expand **Control**. Click **Reset**. Using **Reset** instead of **Reboot** is acceptable, as the VM does not retain configuration on reboot.

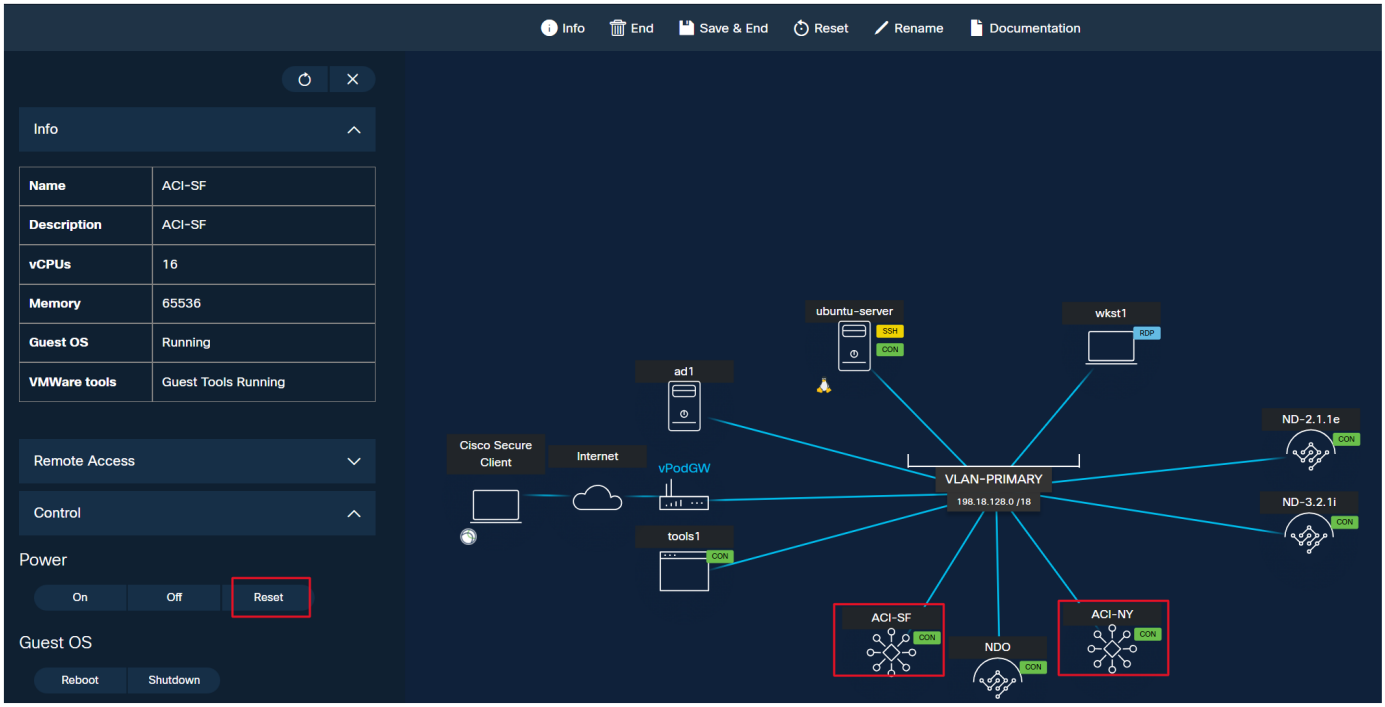

NOTE: It can take 20-25 minutes for the APIC UI to fully load into the login screen. If you see a "Cluster Bring-up" screen, please wait a few more minutes until the login screen loads.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Appendix B. Set Fabric to Permissive Mode

Starting in ACI 6.1, the default security mode has been changed from permissive to strict: https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/cisco-aci-releases-changes-in-behavior.html

Due to a limitation of the ACI simulator used in this demo, it is not possible to onboard **apic2** and **apic3** in strict mode. This means the default mode will be strict and only **apic1** will be available. The other two apics are already configured in the simulator, but with strict mode, they will not automatically be added to the APIC cluster.

If a user needs to access all 3 APICs for any use case beyond the scope of the demo, users can manually switch the mode from **Strict** to **Permissive** via the UI.

1. Sign into the **APIC** UI.
2. Navigate to **System > Controllers**
3. Using the left-pane, navigate to **Controllers > apic1 > Cluster as Seen by Node**.

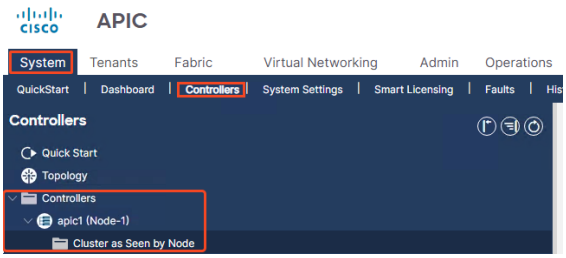

4. In the **Cluster Overview** pane, click **Edit**. You may notice that these panes on the **Cluster as Seen by Node** do not fully load. This is expected behavior due to a limitation of the simulator.
5. Select the **Permissive** radio button and click **OK**.

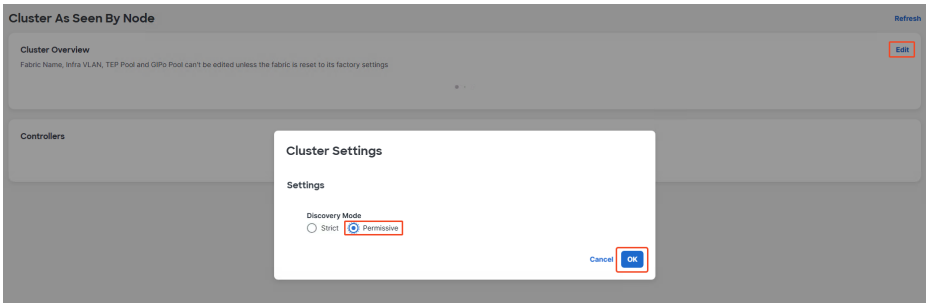

NOTE: This activity is not within the scope of this demo guide and is only included for informational purposes. This is not an activity supported by dCloud support.

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

What's Next?

Users are encouraged to further explore the UI to familiarize themselves with all features available.

Read more about Cisco Application Centric Infrastructure (ACI):
https://www.cisco.com/site/us/en/products/networking/cloud-networking/application-centric-infrastructure/index.html

Check out other dCloud demonstrations:

- https://www.ciscodcloud.com/go/dc-cloud
- https://www.ciscodcloud.com/go/ai

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2025 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.