aruba
a Hewlett Packard Enterprise company

# Validated Solution Guide

Aruba Solution TME

May 28, 2025

# Table of Contents

**ESP Data Center Deploy Guide** **5**
Document Conventions . . . 5

**Data Center Introduction** **6**
Overview . . . 6
EVPN-VXLAN Deployment Overview . . . 7
Layer 2 Two-Tier Deployment Overview . . . 7
Purpose of this Guide . . . 7
Customer Use Cases . . . 8

**Fabric Composer EVPN-VXLAN Fabric** **9**

**Initialize Fabric Components** **10**
Switch Installation . . . 10
Physical Cabling . . . 13
Out-of-Band Management . . . 14
Switch Initialization . . . 15
Download HPE Aruba Networking Fabric Composer . . . 17
Install HPE Aruba Networking Fabric Composer . . . 18
Download AMD Pensando Policy and Services Manager . . . 29
Install AMD Pensando Policy and Services Manager . . . 29
Configure the AMD PSM Cluster . . . 31

**EVPN-VXLAN Configuration** **36**
Physical Topology Overview . . . 36
HPE Aruba Networking Fabric Composer Process . . . 36
Plan the Deployment . . . 37
HPE Aruba Networking Fabric Composer Prerequisites . . . 41
Fabric Initialization . . . 41
Host Port Configuration . . . 120
Configure the Border Leaf . . . 139
Configure Overlay Test Loopbacks . . . 190
Configure Overlay IP Multicast . . . 210
VMWare vSphere Integration . . . 223

**Secure an Aruba EVPN Fabric** **230**
Overview . . . 230
Configure PSM Integration with Aruba Fabric Composer . . . 230
Configure Macro Firewall Policy . . . 235
VM Microsegmentation . . . 254
Verify Policy in PSM . . . 283

**AFC Multifabric Configuration** **288**
Second Fabric Guidance . . . 288
Configure Multifabric Underlay Connectivity . . . 290
Configure Overlay Control Plane . . . 304

**Aruba Central Two-Tier Data Center** **333**
Overview . . . 333
Planning the Deployment . . . 334

**Prepare Switches** **337**
Switch Installation . . . 337
Physical Cabling . . . 339
Management . . . 340

**Aruba Central Initial Configuration** **342**
Create Core Switch Group . . . 342
Configure Core Switch Group Settings . . . 345
Configure Group VLANs . . . 353
Clone Core Group for Access Switches . . . 354
Assign Switches to a Group . . . 355
Configure Switch Hostname . . . 356
Configure Data Center Site . . . 362
Verify Data Center Cabling . . . 366

**Two-Tier Core** **372**
Configure Core VSX ISL Interface . . . 372
Configure Routing VLAN . . . 375
Spanning Tree . . . 377
Enter MultiEdit Configuration . . . 378
Configure Core Switch VSX . . . 380
Configure Core Switch MC-LAGs . . . 381
Configure Routing Services . . . 382
Two-Tier Multicast . . . 387
Verify Operational State . . . 388

**Two-Tier Server Access** **396**
Configure Access VSX ISL Interface . . . 396
Spanning Tree . . . 399
Enter MultiEdit Configuration . . . 400
Configure Access Switch VSX Pairs . . . 401
Configure Access to Core MC-LAGs . . . 402
Configure Access Switch to Host MC-LAGS . . . 402
Configure Multicast . . . 404
Configure Physical Port Speeds . . . 405
Verify Configuration . . . 408

**Ansible Two-Tier Data Center** **413**
Overview . . . 413
Ansible Project Prerequisites . . . 413
Ansible Project Structure . . . 414
Two-Tier Data Center Topology . . . 414

**Ansible Two-Tier Inventory** **417**
YAML Inventory File . . . 417
Planning Inventory File Values . . . 422

Inventory Variable Reference . . . . . 424

**Ansible Two-Tier Template** **433**
Overview . . . . . 433
Template Syntax . . . . . 433

**Run the Two-Tier Playbook** **437**

**Aruba Validated Hardware and Software** **438**
EVPN-VXLAN Fabric . . . . . 438
Layer 2 Two-Tier . . . . . 438

# ESP Data Center Deploy Guide

This guide provides IT professionals with prescriptive steps to deploy a Data Center network using the following products, as outlined in the Design document:

- Aruba Fabric Composer
- Aruba CX 10000 Series
- Aruba CX 8300 Series
- Aruba CX 8400 Series
- VMWare vSphere

## Document Conventions

**Bold** text indicates a command, navigational path, or a user interface element.

Examples:

- the **show vsx status** command
- Go to **Configurations > Routing > VRF**

*Italic* text indicates important terminology or a value that appears in a field on the user interface.

Examples:

- A *VXLAN Tunnel Endpoint (VTEP)* function within leaf switches manages the origination and termination of point-to-point tunnels forming an overlay network.
- **Name**: *DB_NET_PROD_DC1*

Shaded blocks indicate variables for which you should substitute a value appropriate for your environment.

Example:

- BGP router id: `10.0.5.1`

**NOTE:**

For the most up-to-date information on ESP Data Center solutions, please refer to the Validated Solution Guide Program

# Data Center Introduction

The HPE Aruba Networking data center uses technology and tools to transform the data center into a modern, agile, services delivery platform. Organizations of any size, distributed or centralized, can benefit from streamlined perfomance and improved network cost-effectiveness using HPE Aruba Networking data center products.

## Overview

The HPE Aruba Networking AOS-CX operating system simplifies overall operations and maintenance using a common switch operating system across the campus, branch, and data center. The system can be managed in the cloud or on-premises and supports a comprehensive application programming interface (API). AOS-CX employs robust artificial intelligence functions that continually analyze and realign network flow to ensure that the system operates seamlessly in accordance with network management best practice, without requiring manual IT staff intervention.

The use of converged Ethernet has changed the way hosts access storage within the modern data center. Dedicated storage area networks are no longer required. Lossless Ethernet and bandwidth management protocols ensure timely reads and writes using a traditional IP LAN. The combined cost savings and operational simplicity are driving a major conversion to converged Ethernet.

At the same time, network topologies have become virtualized. Although virtualization delivers the flexibility required to meet the changing data center requirements, it can present complexity and challenges with implementation and management. An HPE Aruba Networking data center addresses these challenges by automating installation and implementation of the Aruba AOS-CX operating system, with features such as automated device group configuration, Zero Touch Provisioning, scheduled configuration backups, dashboard-ready network performance metrics, and built-in alerts for critical network functions.

Before designing a new or transformed data center, it is important to consider the organization’s current and projected strategy for hosting and accessing applications from the cloud. Determine the applications that will remain on-premises so you can establish a data center with ample throughput and storage to meet requirements.

The HPE Aruba Networking CX 93xx, 83xx, CX 84xx, and CX 10000 switching platforms provide a best-in-class suite of products featuring a variety of high-throughput port configurations, industry-leading operating system modularity, real-time analytics, and “always up” performance.

This guide explores deploying HPE Aruba Networking switches to create a modern EVPN-VXLAN solution and a traditional Layer 2 two-tier architecture.

## EVPN-VXLAN Deployment Overview

An EVPN-VXLAN architecture accommodates growth and provides network flexibility using a Layer 3 spine-and-leaf underlay with a software-defined fabric overlay. Spine switches provide connectivity between leaves, while data center hosts are attached to leaf switches. The EVPN-VXLAN fabric overlay simultaneously supports Layer 3 segmentation and Layer 2 adjacency between hosts anywhere in the data center using standards-based protocols. Reachability information between hosts is shared using BGP's L2VPN-EVPN address family. VXLAN encapsulation is used to forward traffic between overlay hosts using the Layer 3 underlay as a transport service.

The HPE Aruba Networking EVPN-VXLAN architecture provides the following benefits:

- A fault tolerant design that accommodates hardware failures at multiple levels.
- Easy incremental east-west capacity expansion by adding switches at the spine layer.
- Programmatic Layer 2 VLAN reachability across the data center.
- Programmatic expansion of Layer 3 segments supporting data center multitenancy.
- Inline policy enforcement using the Aruba CX 10000 switch.
- Microsegmentation of attached hypervisor VMs and containers.
- Switch upgrades without a service outage.
- Orchestrated configuration, management, and operations using HPE Aruba Networking Fabric Composer.

## Layer 2 Two-Tier Deployment Overview

A Layer 2 two-tier architecture supports a resilient, high capacity data center design without the need for specialized knowledge in overlay protocols. Multi-chassis link aggregations (MC-LAGs) provide multiple Layer 2 redundant data paths between a collapsed data center core and access switches, and between access switches and their attached data center hosts.

The HPE Aruba Networking Layer 2 two-tier architecture provides the following benefits:

- A fault tolerant design that can accommodate hardware failures at multiple levels.
- Layer 2 VLAN reachability across the data center.
- Inline policy enforcement using the Aruba CX 10000 switch.
- Microsegmentation of attached hypervisor VMs and containers.
- Switch upgrades without a service outage.
- Simplified configuration, management, and operations using Aruba Central cloud-based controls.

## Purpose of this Guide

This guide describes HPE Aruba Networking data center implementation procedures, with reference for architectural options and associated hardware and software components. It delivers best-practice recommendations for the following deployment models:

- A next generation spine-and-leaf data center fabric using VXLAN and BGP EVPN to take advantage of the orchestration capabilities of **HPE Aruba Networking Fabric Composer**.
- A traditional Layer two-tier data center configured via HPE Aruba Networking **Central**.
- An Ansible automated configuration via AOS-CX API for a traditional Layer 2 two-tier data center.

Refer to volume one of this VSG for additional design guidance: HPE Aruba Networking Data Center Design

This guide assumes the reader has an equivalent knowledge of an Aruba Certified Switching Associate.

### Audience

This guide is written for IT professionals who need to deploy an Aruba Data Center Network. These IT professionals serve in a variety of roles:

- Systems engineers who require a standard set of procedures to implement network solutions
- Project managers who create statements of work for Aruba implementations
- Aruba partners who sell technology or create implementation documentation.

## Customer Use Cases

Data center networks change rapidly. The most pressing challenge is maintaining operational stability and visibility for users while moving or upgrading computing and storage resources. In addition, data center teams must continue to support the rapid pace of DevOps environments and meet growing requirements to connect directly and continue operations within the public cloud infrastructure.

Within a rapidly changing landscape, it is critical that network and system engineers responsible for meeting data requirements have efficient tools to streamline and automate complex infrastructure configurations.

# Fabric Composer EVPN-VXLAN Fabric

HPE Aruba Networking Fabric Composer automates initial provisioning, ongoing configuration, and management of an EVPN-VXLAN data denter fabric. Fabric Composer's Guided Setup process configures baseline switch features, underlay addressing and routing, and overlay control and data plane components.

Fabric Composer integration with VMware vCenter provides visibility, policy automation, and vCenter configuration of DVS and PVLAN policy components.

Fabric Composer enables flexible management of AMD Pensando's Policy and Services Manager (PSM) firewall policy for the CX 10000 and access control lists available on all switch models. The API-based integration with PSM and AOS-CX switches allows Fabric Composer to provide a single pane of glass for all network-based policy management.

Fabric Composer provides a wizard-based microsegmentation policy builder. The end result of its PSM, switch, and vCenter policy component automation is a model that enables systems administrators to dynamically add a VM to network policy enforcement in their own independent workflows using VM tags.

Additionally, Fabric Composer provides ongoing visibility into fabric components. Alerts and troubleshooting tools allow quick identification and resolution to failures in the EVPN-VXLAN fabric.

# Initialize Fabric Components

The first step for deploying a data center is the physical installation of the switches and computing hosts.

## Switch Installation

Verify the airflow configuration for the products to be installed to ensure that they support the cooling design for the data center. If required, an optional air duct kit is available for Aruba data center top-of-rack (ToR) switches to redirect hot air away from servers inside the rack.

Before installing switches, download the Aruba Installation Guide for the specific models. Review the Installation Guide before installing and deploying the switches.  Carefully review requirements for power, cooling, and mounting to ensure that the data center environment is outfitted adequately for safe, secure operations .

**Step 1** Open a web browser, navigate to the HPE Networking Support Portal, and login with using appropriate credentials.

**Step 2** On the landing page, click **Software and Documents Search** panel.

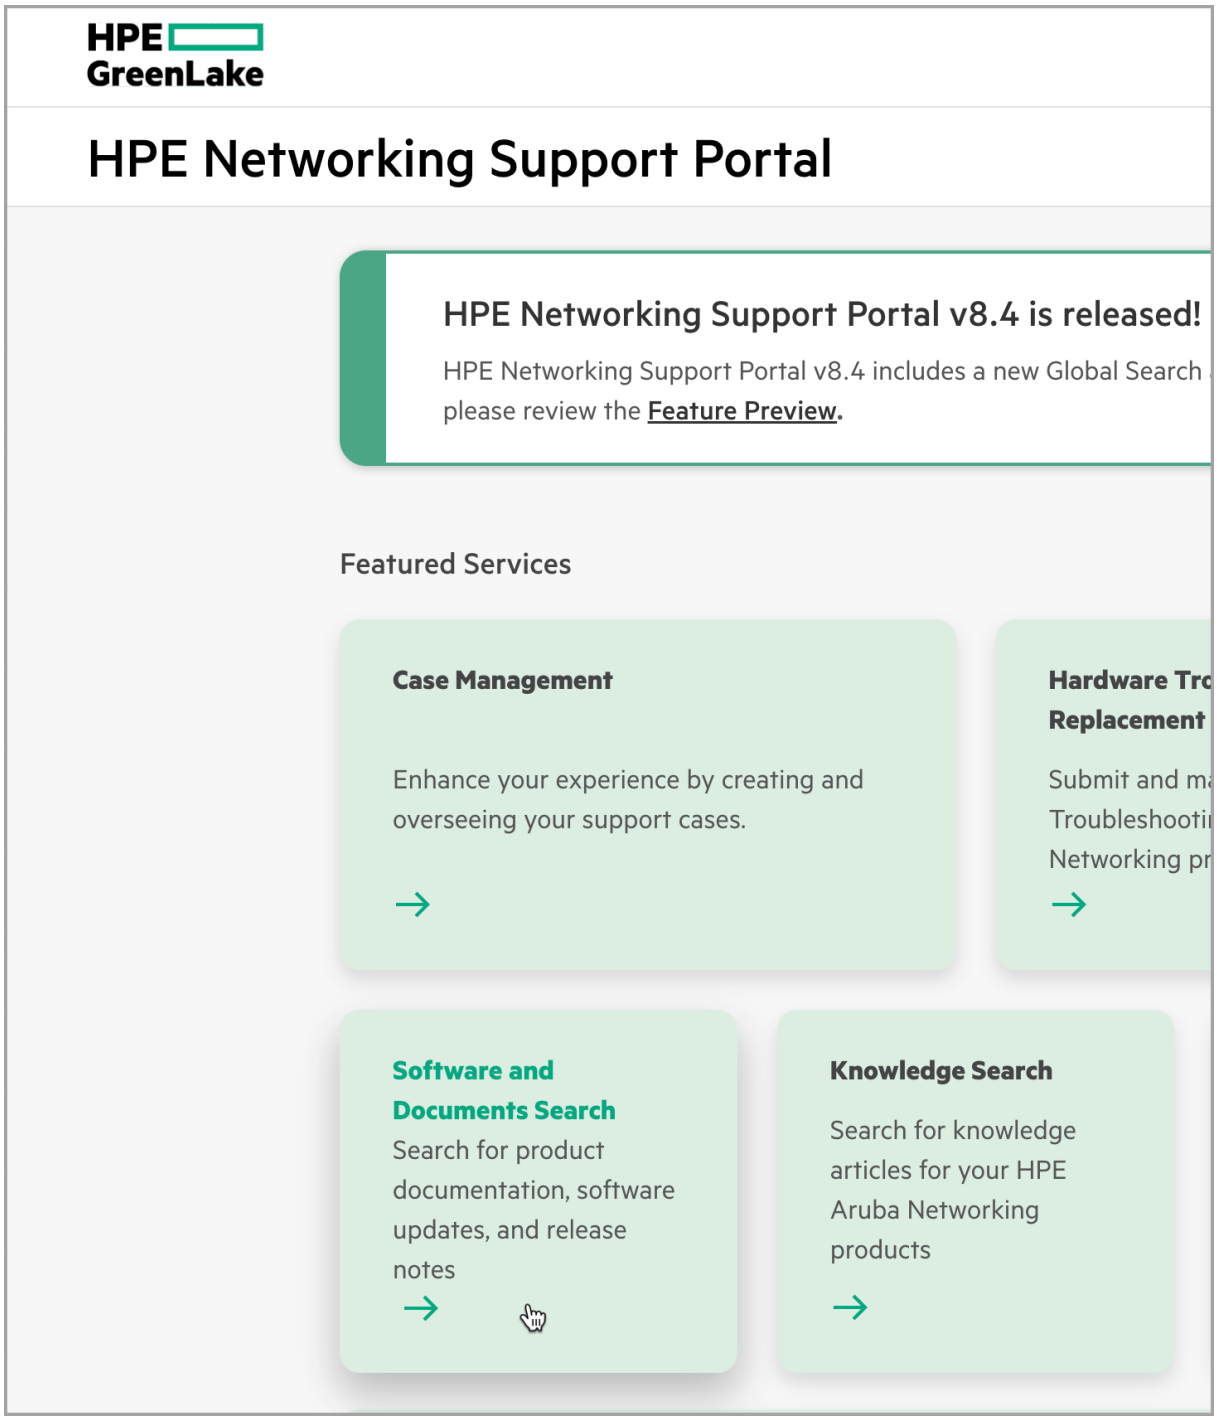


**Step 3** On the **Software & Documents** page, select the following filters. - **File Type:** *Document* - **Product:** *Aruba Switches* - **File Category:** *Installation Guide*

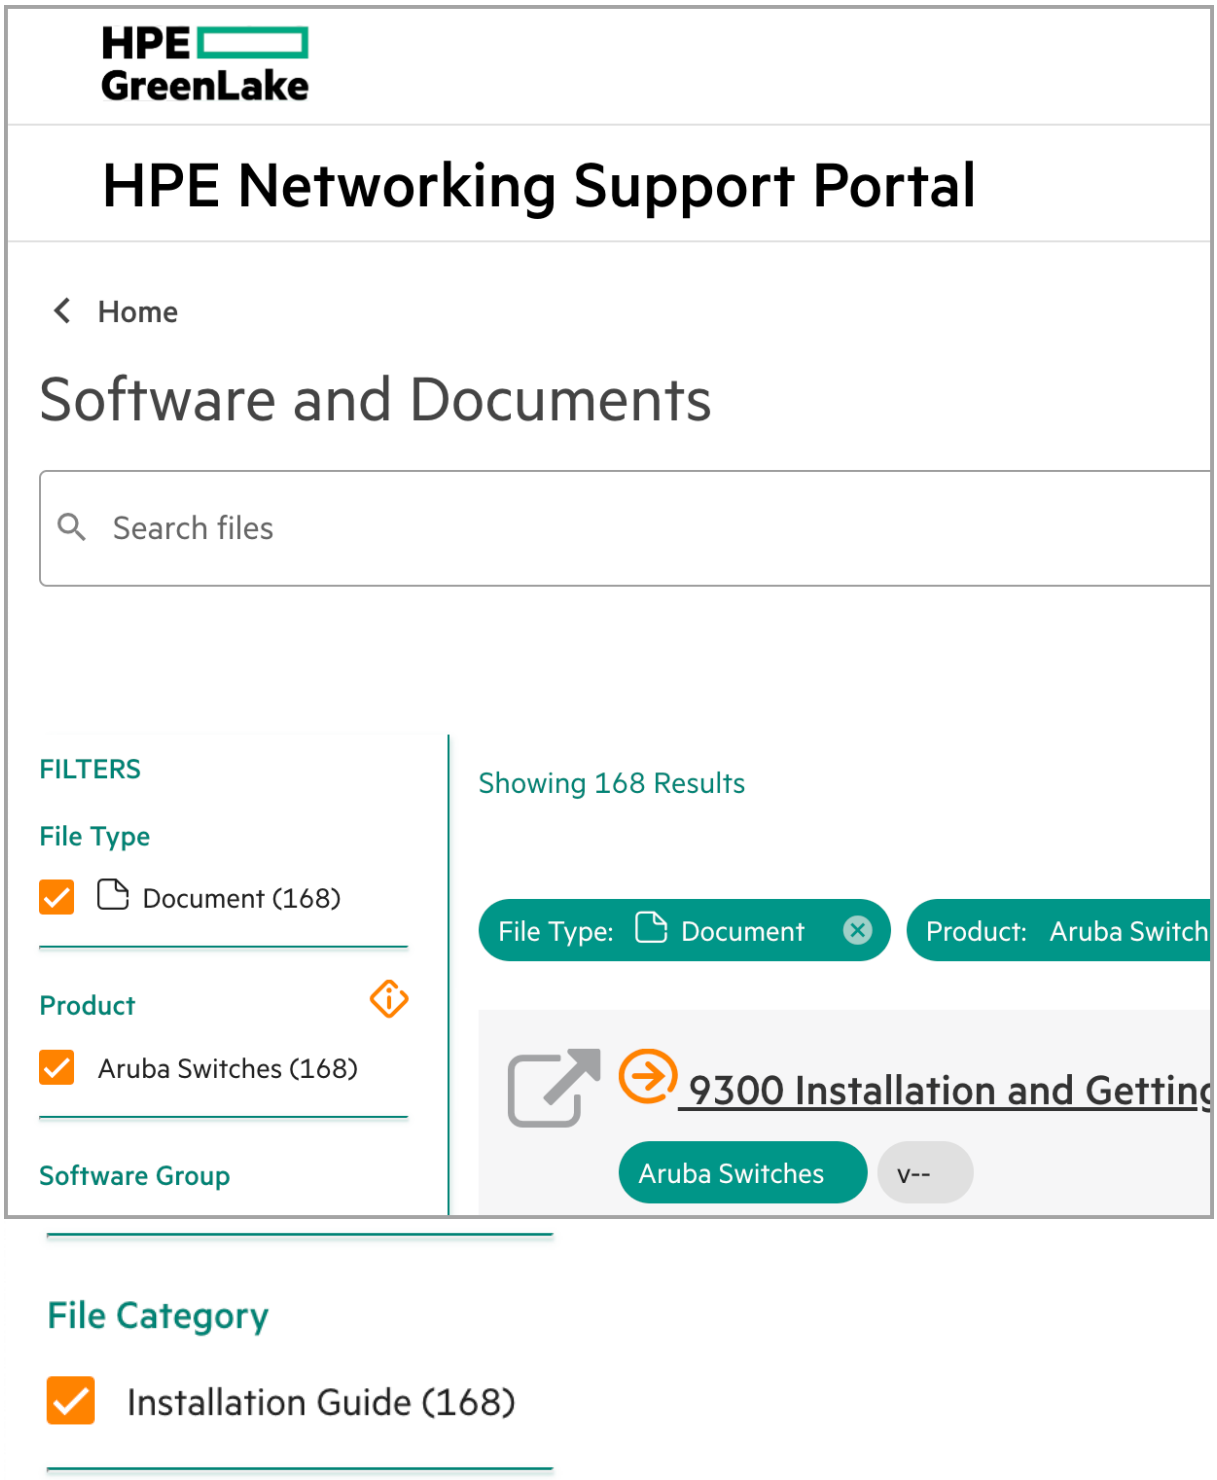


**Step 4** Download the Installation Guide version for the switch model to be installed.

**Step 5** Complete the physical installation of switches in the racks.

**NOTE:**

Spine switches can be installed centrally, in middle-of-row or end-of-row locations depending on cabling requirements and space availability. The key consideration is cable distance and the types of media used between leaf and spine switches. Leaf switches should be installed top-of-rack (ToR) in high-density environments or middle-of-row in low-density environments.

# Physical Cabling

Consistent port selection across racks and in the spine switches increases the ease of configuration management, monitoring, reporting, and troubleshooting tasks in the data center.

Breakout cables are numbered consistently with their split port designation on the switch.

Document all connections.

Ensure that distance limitations are observed for your preferred host connection media and between switches.

## Top of Rack Cabling

The illustrations below show the port configuration on two types of 48-port ToR switches. Redundant ToR switch pairs must be the same model.

Ports on an Aruba CX 8325-48Y8C:


**Figure 1:** 8325 ToR switch

Ports on an Aruba CX 10000-48Y6C:

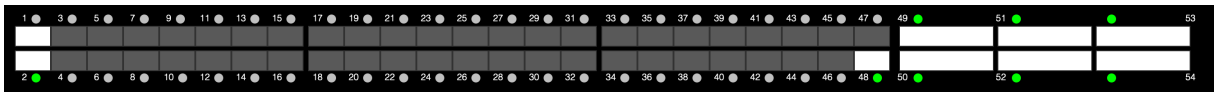


**Figure 2:** 8325 ToR switch

In a redundant ToR configuration, the first two uplink ports should be allocated to interconnect redundant peers (ports 49-50 on 8325-48Y8C and 10000-48Y6C switches), which provides physical link redundancy and sufficient bandwidth to accommodate a spine uplink failure on one of the switches.

Two links between redundant peers are sufficient for most deployments, unless the design may result in high traffic utilization of the inter-switch links under normal operating conditions, such as when many hosts in a rack are single-homed to only one of the redundant switches.

Additional uplink ports should be allocated to connect spine switches (ports 51-56 on an 8325-48Y8C and ports 51-54 on a 10000-48Y6C).

The highest numbered non-uplink port should be reserved as the VSX keepalive link between a ToR redundant pair.

**NOTE:**

VSX automation in HPE Aruba Networking Fabric Composer requires a dedicated physical port or a loopback address for the VSX keepalive interface. The recommended configuration is a dedicated port.

Determine a consistent number of leaf-to-spine links required on each ToR to achieve the desired oversubscription ratio. The number of spine switches is equal to the number of per ToR links required.

Follow a similar approach when using lower density ToR designs. Before deploying ToR configurations that require server connectivity at multiple speeds, review the switch guide to determine if adjacent ports are affected.

Configuration steps for changing port speeds are covered later in this guide. Refer to the Data Center Reference Architecture section for guidance on port speed groups on different hardware platforms.

## Spine-to-Leaf Cabling

The illustration below shows the port configuration on an 8325 32-port spine switch.

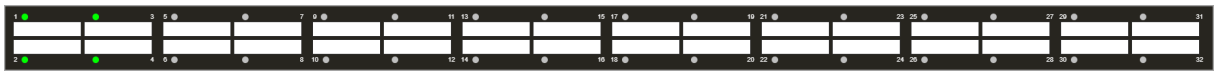

**Figure 3:** Spine switch

In a dual ToR configuration, a spine switch must be connected to each switch in the redundant ToR pair in each rack. A 32-port spine switch supports up to 16 racks in this design. Use the same port number on each spine switch to connect to the same leaf switch to simplify switch management and documentation. For example, assign port 1 of each spine switch to connect to the same leaf switch.

## Border Leaf Cabling

In a VXLAN spine-and-leaf design, a pair of leaf switches serves as the single entry and exit point to the data center. This is called the border leaf, but it does not require dedication to only border leaf functions. It may provide services leaf functions and, in some cases, provide connectivity to directly attached data center workloads. Cabling the border leaf can vary among deployments, depending on how the external network is connected and if services such as firewalls and load balancers are connected.

After all switches are physically installed with appropriate power and networking connections, continue to the next procedure.

# Out-of-Band Management

The use of a dedicated management LAN for the data center is strongly recommended.

A dedicated management LAN on separate physical infrastructure ensures reliable connectivity to data center infrastructure for automation, orchestration, and traditional management access. The management LAN provides connectivity to HPE Aruba Networking Fabric Composer, Aruba NetEdit, and AMD Pensando Policy and Services Manager (PSM) applications. Ensure that the host infrastructure needed for those applications also can be connected to the management LAN or is reachable from the management LAN.

Deploy management LAN switches top-of-rack with switch and host management ports connected. Plan for an IP subnet with enough capacity to support all management addresses in the data center. DNS and NTP services for the fabric should be reachable from the out-of-band management network.

Configuration steps for the management LAN are not covered in this guide.

## Switch Initialization

Go to the HPE Networking Support Portal and download the AOS-CX Fundamentals Guide for the version of the operating system you plan to run using the steps noted above for “Switch Installation.”

**NOTE:**

Refer to the operating system release notes and consult with an HPE Aruba Networking SE or TAC team member for assistance with determining and selecting the version.

The “Initial Configuration” section of each Fundamentals Guide presents detailed instructions for connecting to the switch console port. After connecting to the console port, follow the steps below.

**Step 1** Enable power to the switch by connecting power cables to the switch power supplies.

**Step 2** Login with the username **admin** and an empty password.

**Step 3** Enter a new password for the admin account.

**NOTE:**

The “Initial Configuration” section of the Fundamentals Guide provides detailed instructions for logging into the switch the first time.

**Step 4** Confirm that all CX 10000 switches in the fabric are running an AOS-CX version compatible with the Fabric Composer and PSM versions of a deployment. Table 2 in the Fabric Composer’s Pensando PSM & AOS-CX 10000 Software Selection Guidance document provides a matrix for compatibility. This guide uses the following versions of firmware and software:

- **AOS-CX**: *10.13.1050*
- **HPE Aruba Networking Fabric Composer**: *7.0.5*
- **PSM:** *1.80.1-T-6*

**Step 5** Confirm that all other switches are running AOS-CX 10.10 long-term stability release or AOS-CX 10.13+ for compatibility with Fabric Composer 7.0.5 used in this guide.

**Step 6** If the switch was previously configured, reset it to the factory default configuration. Fabric Composer requires a factory default configuration for orchestration during the fabric configuration process.

```
8325# erase all zeroize
This will securely erase all customer data and reset the switch
to factory defaults. This will initiate a reboot and render the
switch unavailable until the zeroization is complete.
This should take several minutes to one hour to complete.
Continue (y/n)? y
```

**Step 7** Configure 6300M VSF stacks using the Aruba AOS-CX VSF Guide.

**NOTE:**

VSF stacks should be configured on 6300 switches before making any other configuration changes after zeroization.

**Step 8** Configure switch hostnames.

```
hostname RSVDC-FB1-LF1-1
```

**NOTE:**

It is important to use a canonical naming scheme to easily identify the function of each switch. The hostname scheme above uses <*physical location*>-<*fabric identifier*>-<*role and unique VSX pair identifier*>-<*VSX pair member id*> to identify the correct fabric and role when using Fabric Composer. When using this scheme for switches that are not in a VSX pair, the number in the role field is sufficient for unique identification (i.e., RSVDC-FB1-SP1).

**Step 9** Configure the Switch Management Interface. By default, the management interface uses DHCP for its configuration. DHCP reservations can be used to assign a consistent IP address, default gateway, and nameserver. Static IP configuration eliminates dependence on DHCP service availability.

```
interface mgmt
    no shutdown
    ip static 172.16.116.101/24
    default-gateway 172.16.116.1
    nameserver 172.16.1.98
```

**NOTE:**

Based on the existing IP address management process, determine a subnet to be used for the management LAN, where out-of-band (OOB) management ports on your switches are connected. Aruba Fabric Composer must be reachable from this network. The “Initial Configuration” section of the Fundamentals Guide provides detailed instructions for configuring the management interface.

**Step 10** When spines use breakout cabling, configure split ports with the appropriate number of child interfaces and connection speeds, then confirm the operational port change.

```
interface 1/1/1-1/1/3
    split 2 100g
```

```
RSVDC-FB1-SP1(config)# interface 1/1/1-1/1/3
RSVDC-FB1-SP1(config-if-<1/1/1-1/1/3>)# split 2 100g
This command will disable the specified port, clear its configuration,
and split it into multiple interfaces.

Continue (y/n)? y
RSVDC-FB1-SP1(config-if-<1/1/1-1/1/3>)#
```

**NOTE:**

Typically, a spine uses a consistent split port strategy. An interface range is used to assign the same split configuration to multiple ports. The **confirm** parameter in the split configuration statement disables the operational warning. For example, **split 2 100g confirm**. Split interfaces also can be configured in HPE Aruba Networking Fabric Composer.

## Download HPE Aruba Networking Fabric Composer

**Step 1** Navigate to the HPE Networking Support Portal.

**Step 2** Click the **Software and Documents** pane.

**Step 3** In the **File Type** filter, select **Software**.

**Step 4** In the **Product** filter, select “Aruba Fabric Composer”, and click **Apply**.

**Step 5** In the search results, select the appropriate OVA version and download it to your computer. This guide uses Fabric Composer 7.0.5.

**Step 6** In the **File Type** filter, uncheck **Software**, then select **Documents**.

**Step 7** Type *release notes* in the **Search Files** bar.

**Step 8** Click on the HPE Aruba Networking Fabric Composer release notes for the version of software downloaded. The download link on the resulting page forwards the browser to Fabric Composer’s online help, install guide, and compatibility matrix. Review the installation considerations in the **Install Guide** to ensure that adequate host resources are available.

**NOTE:**

HPE Aruba Networking Fabric Composer is provided in ISO format for installation using other hypervisors. High availability Fabric Composer clusters are only supported when using an ISO image.

# Install HPE Aruba Networking Fabric Composer

Install Fabric Composer using the best process for your organization. The following process installs the Fabric Composer OVA using VMware vCenter.

**Step 1** In the **Hosts and Clusters** tab, right click on the location to install Fabric Composer and select **Deploy OVF Template...** to launch the installation wizard.

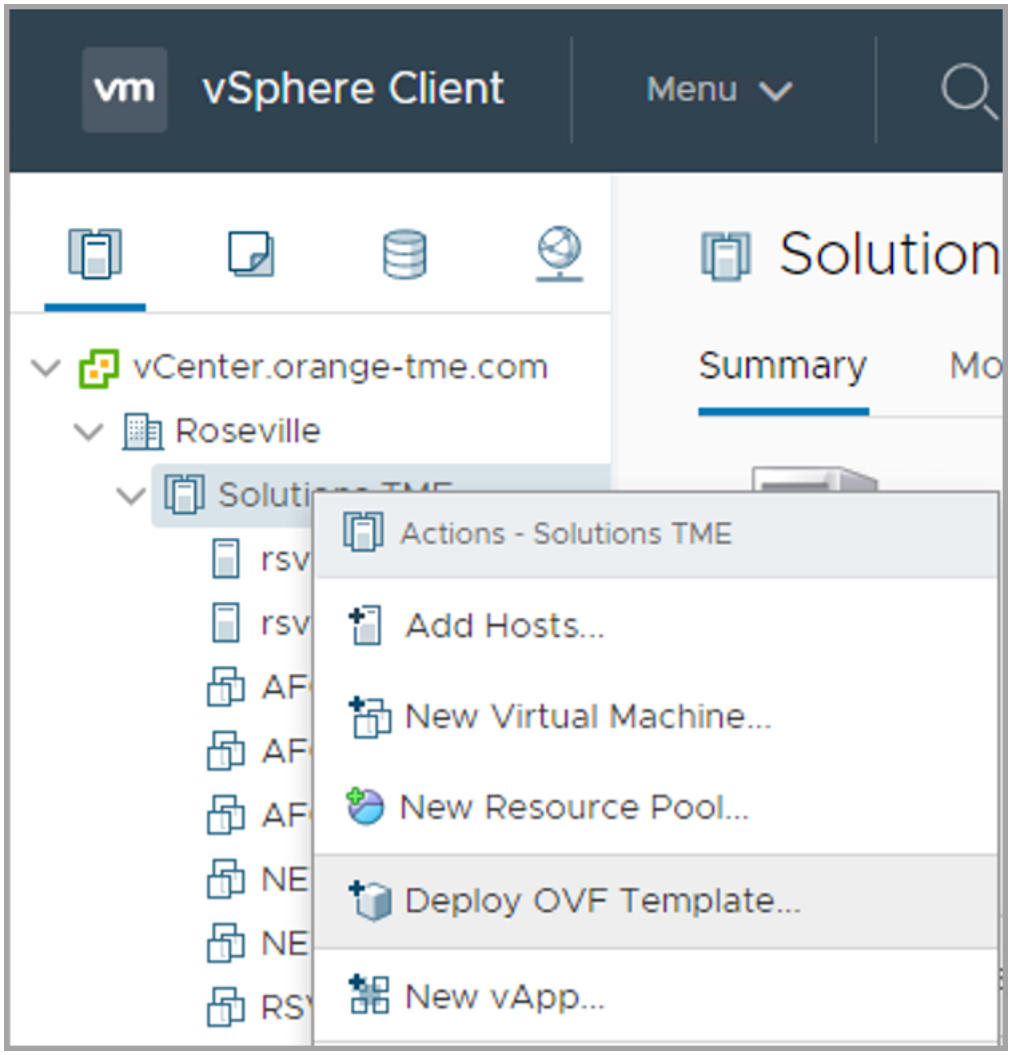


**Step 2** On the **Select an OVF template** page, click **Local file**, choose the downloaded AFC OVA file, and click **NEXT**.

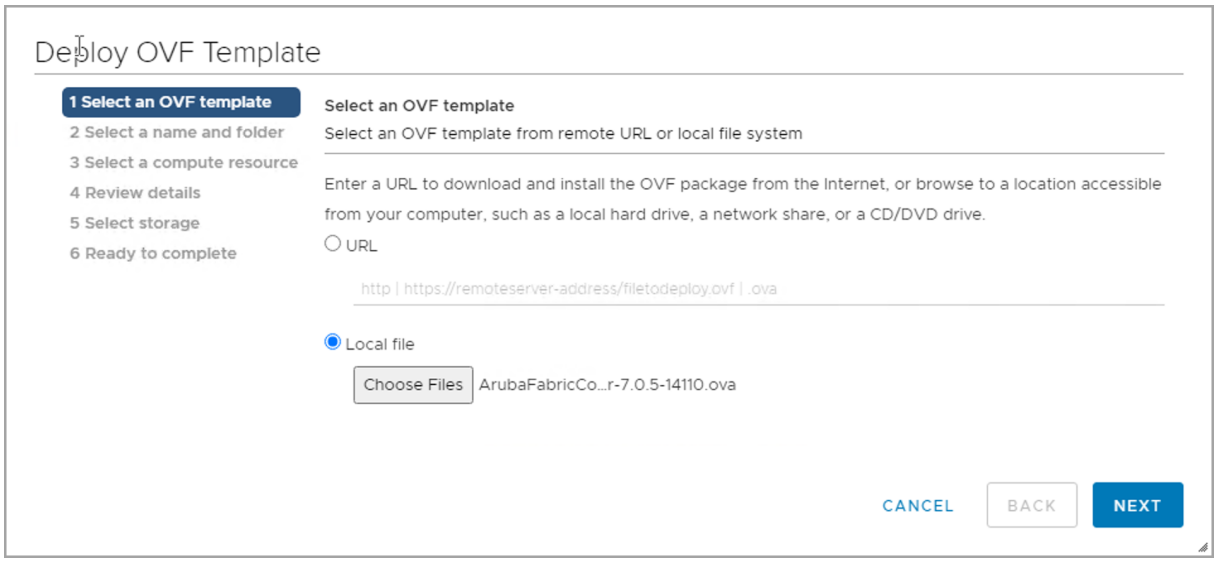


**Step 3** On the **Select a name and folder** page, enter a virtual machine name, select a target folder for the installation, and click **NEXT**.

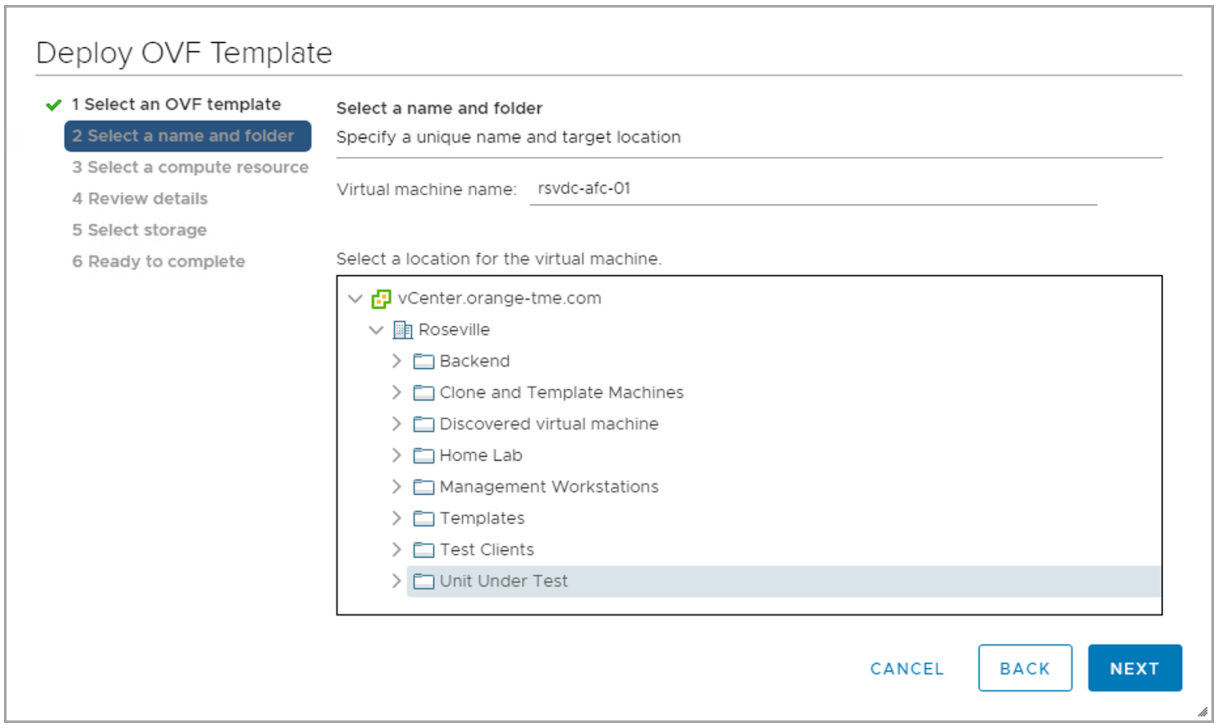


**Step 4** On the **Select a compute resource** page, select a cluster or cluster member and click **NEXT**.

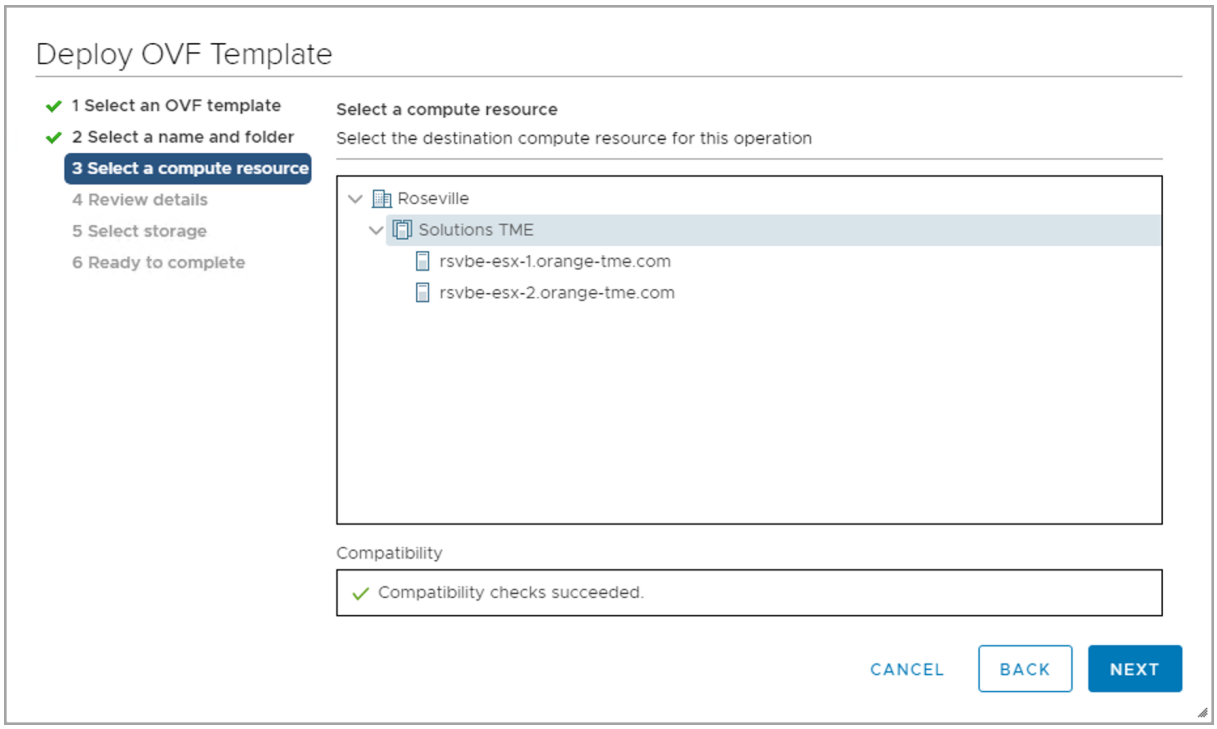


**Step 5** On the **Review details** page, read the information presented and click **NEXT**.

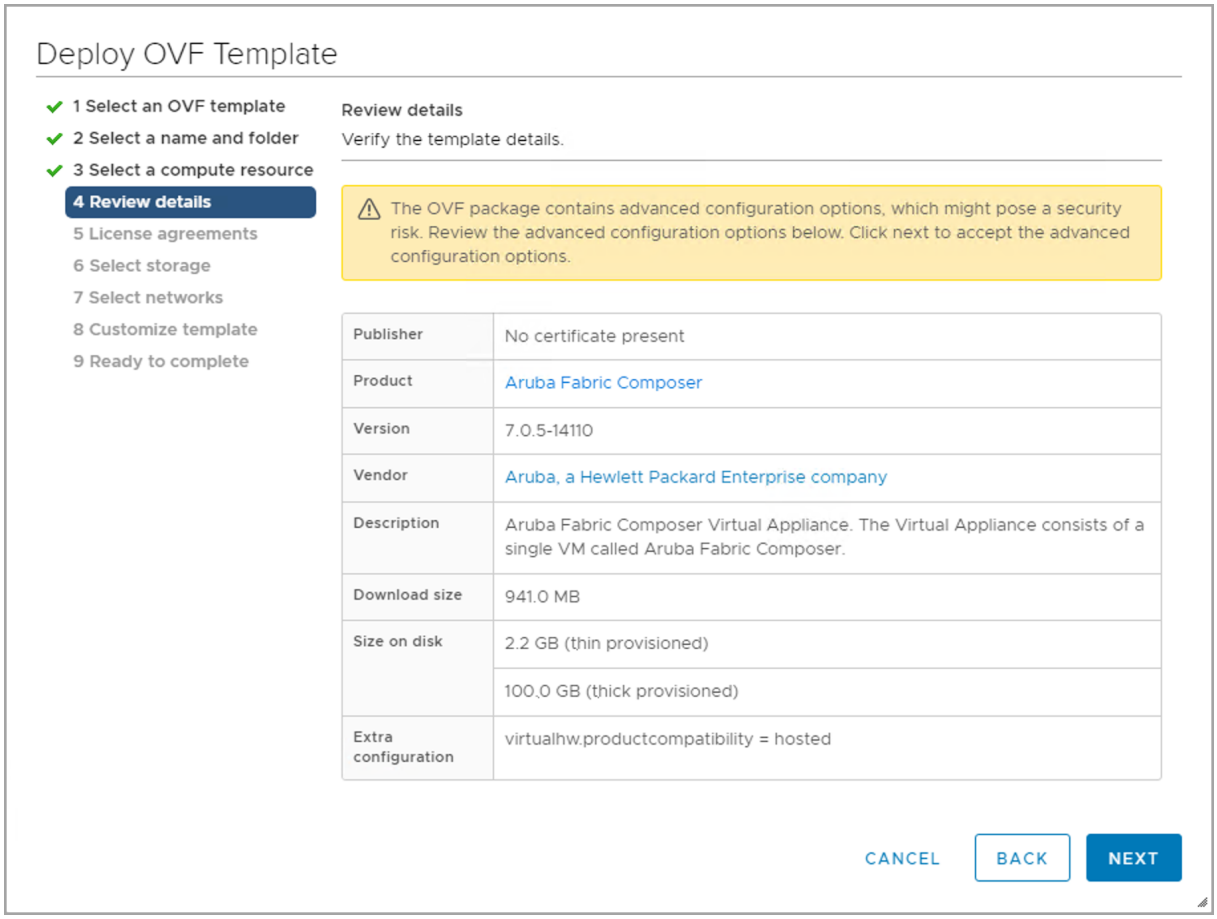


**Step 6** On the **License agreements** page, read the license agreement, select **I accept all license agreements**, and then click **NEXT**.

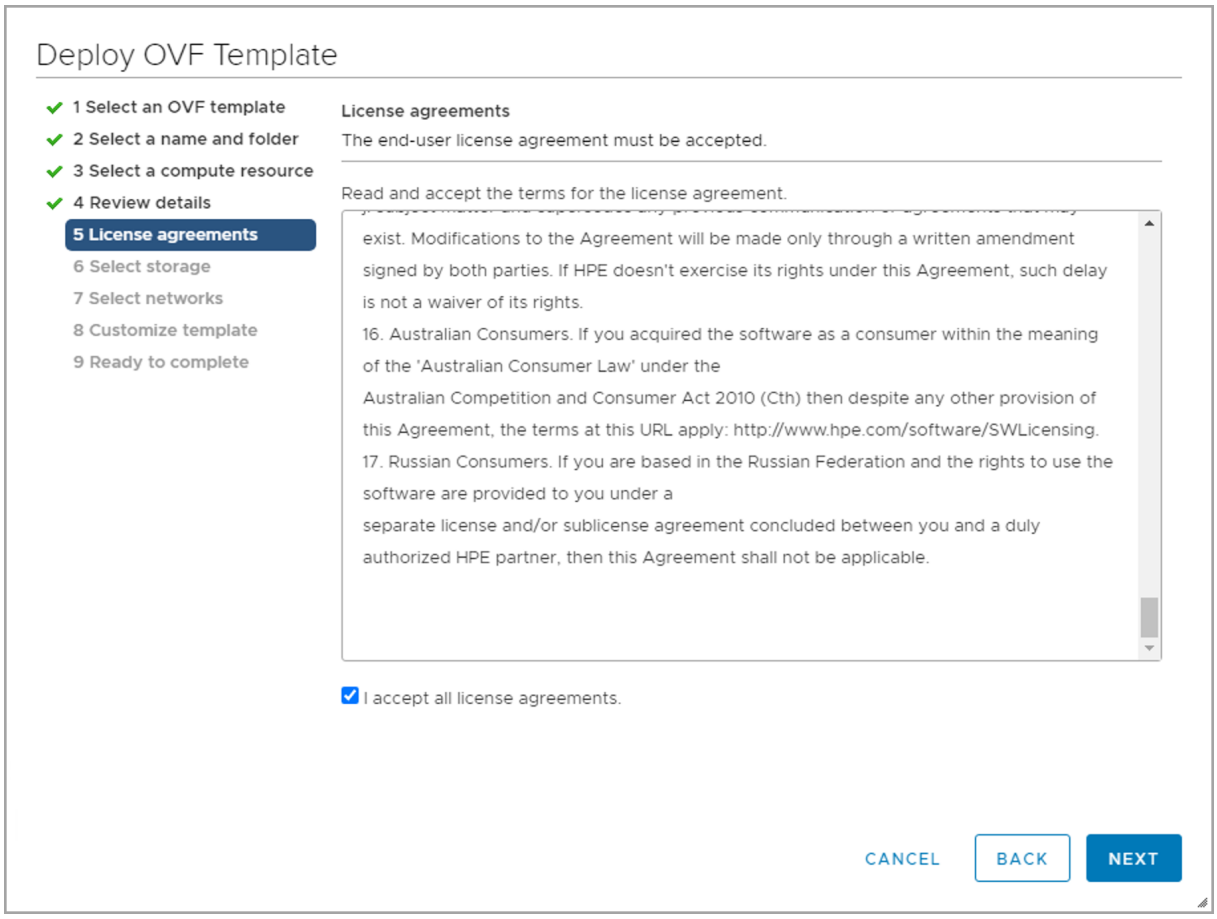


**Step 7** On the **Select storage** page, select the preferred provisioning method and storag volume, then click **NEXT**.

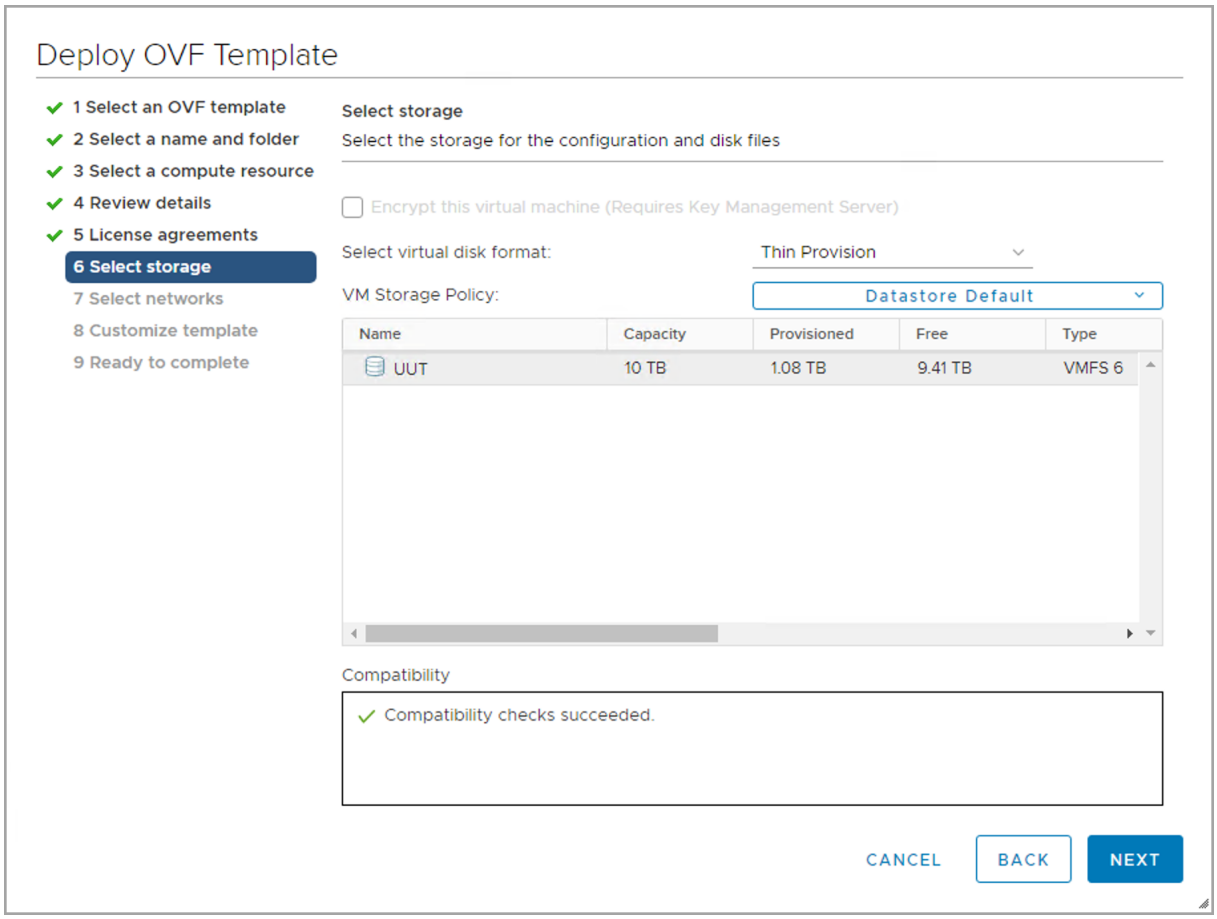


**Step 8** On the **Select networks** page, select a VM Network with connectivity to the data center out-of-band netowkr and click **NEXT**.

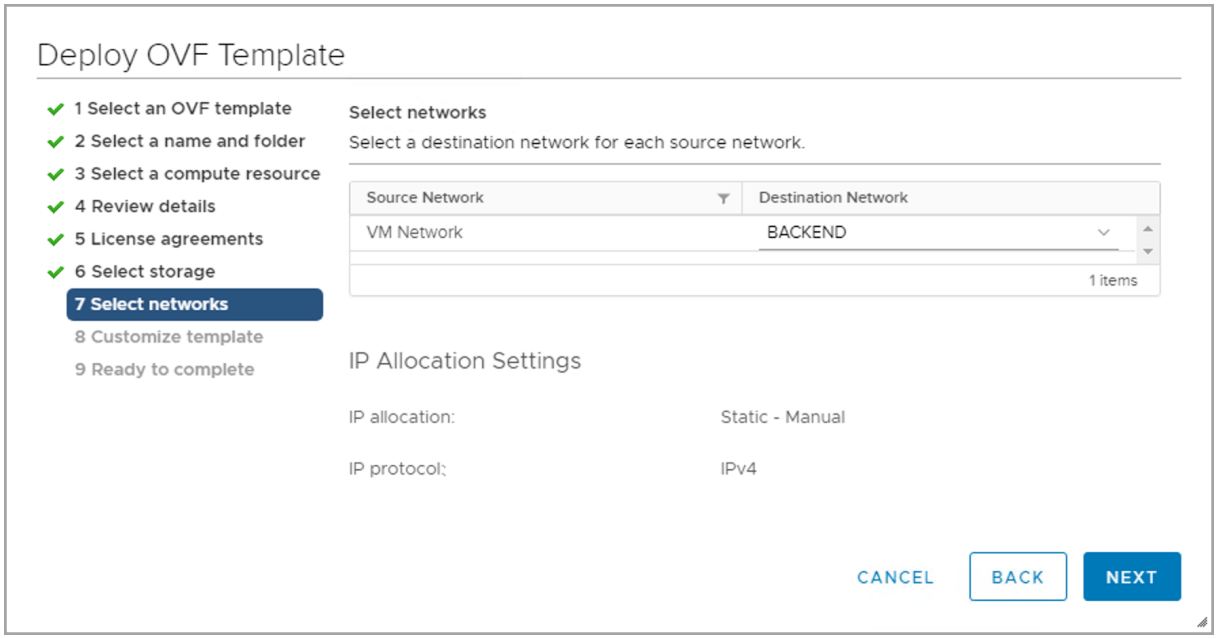


**Step 9** On the **Customize template** page, enter values for the following fields and click **Next**.

*(A) Network - General settings* - **(1) Hostname**: *rsvdc-afc-01* - **(2) Domain Name**: *example.local* - **(3) Primary NTP Server**: *172.16.1.99* - **(4) Secondary NTP Server**: *172.16.1.98*

*(B) Network - Static IP settings* - **(1) IP Address**: *172.16.1.50* - **(2) Network Mask**: *255.255.255.0* - **(3) Default Gateway**: *172.16.1.1* - **(4) Primary Name Server**: *172.16.1.99* - **(5) Secondary Name Server**: *172.16.1.98*

*(D) Linux Password* - **Password**: *<password&gt* - **Confirm Password**: *<password>*

**NOTE:**

Check **Use DHCP** when dynamic addressing is preferred over static IP assignment.

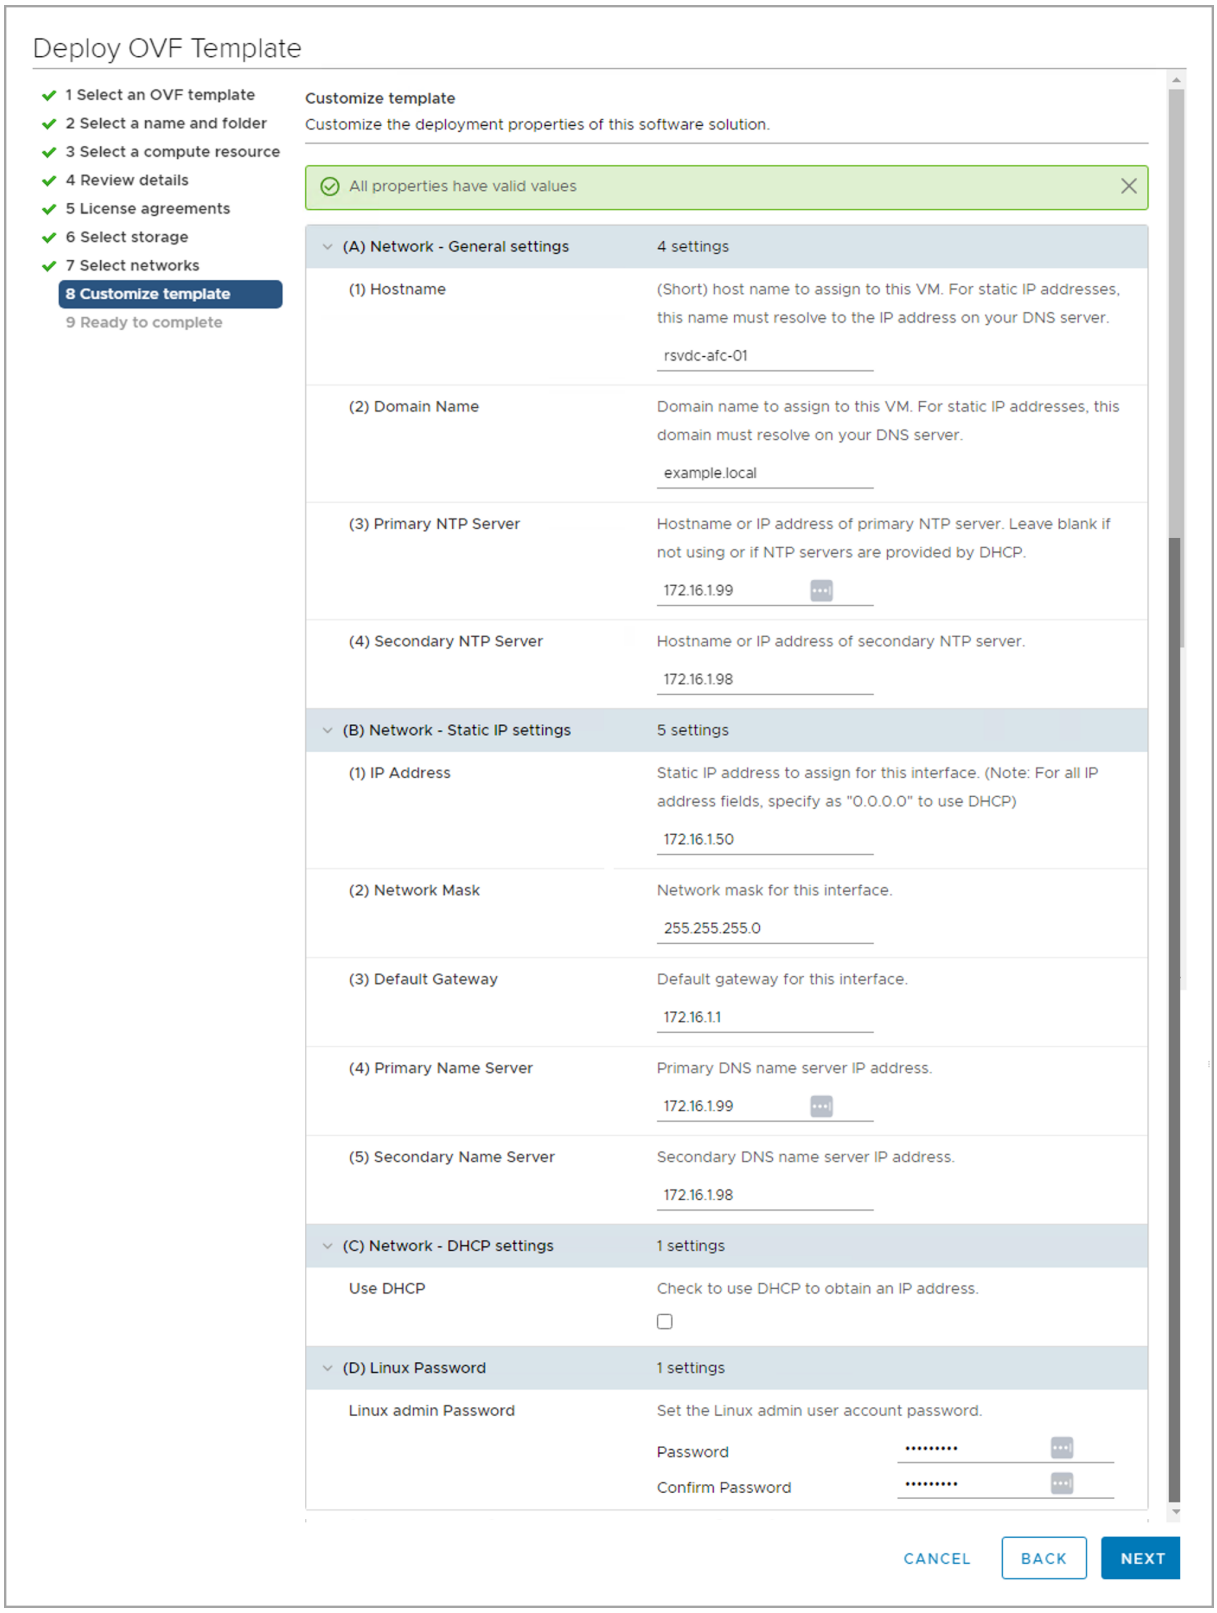


**Step 10** On the **Ready to complete** page, click **FINISH**.

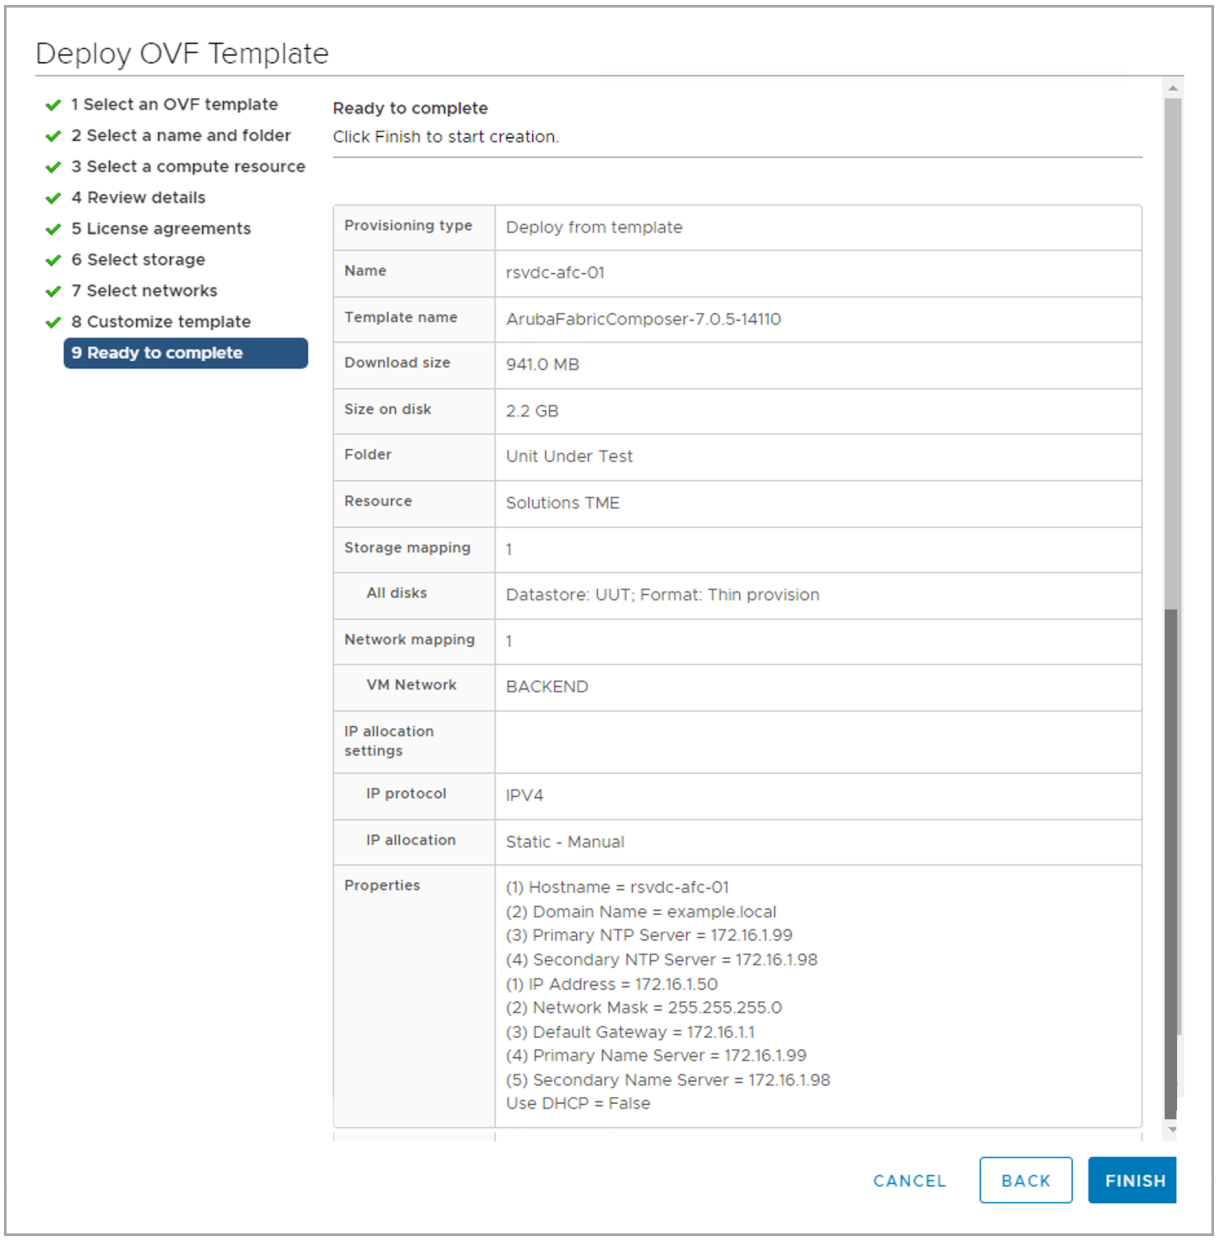


**Step 10** Open a web browser and connect to Fabric Composer at the previously configured IP address.

**NOTE:**

The software version is not displayed and login is not allowed while the system is initializing.

**Step 12** On the **Fabric Composer** page, enter the following default credentials, and click **LOGIN**.

- **Username:** *admin*
- **Password:** *aruba*


**Step 13** Enter the current and new password and click **APPLY**.

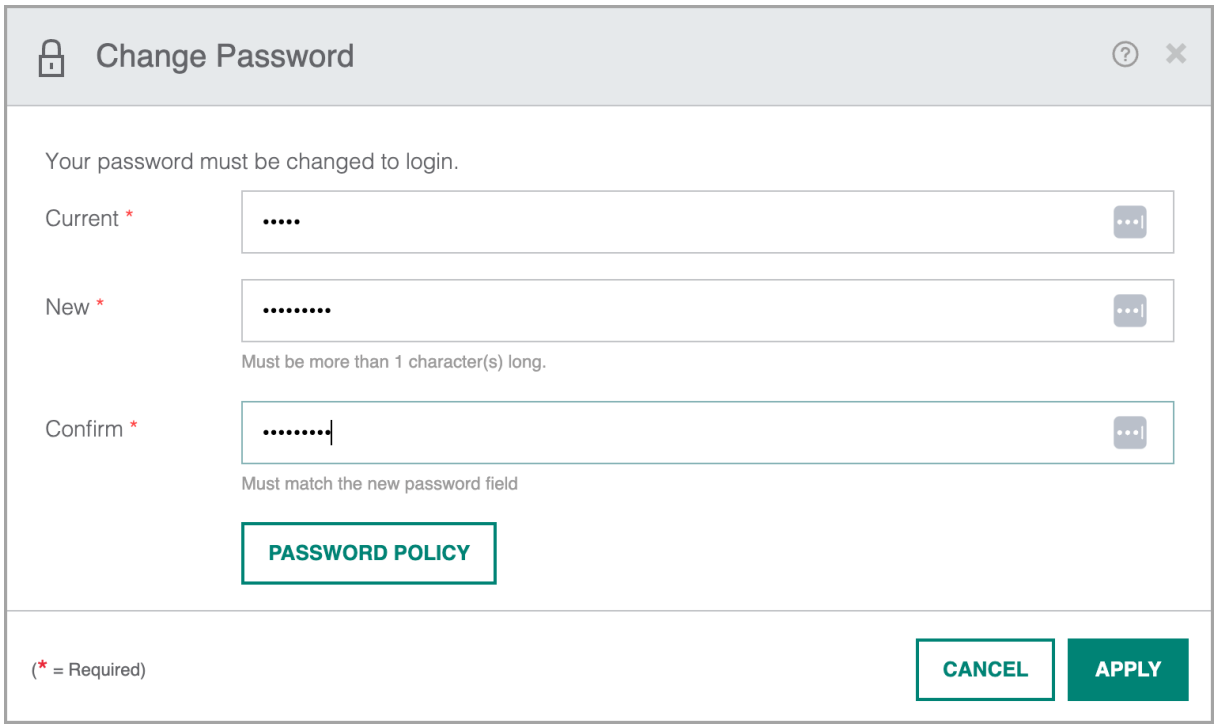


## Add HPE Aruba Networking Fabric Composer Licenses

**Step 1** On the **Maintenance** menu, select **Licenses**.

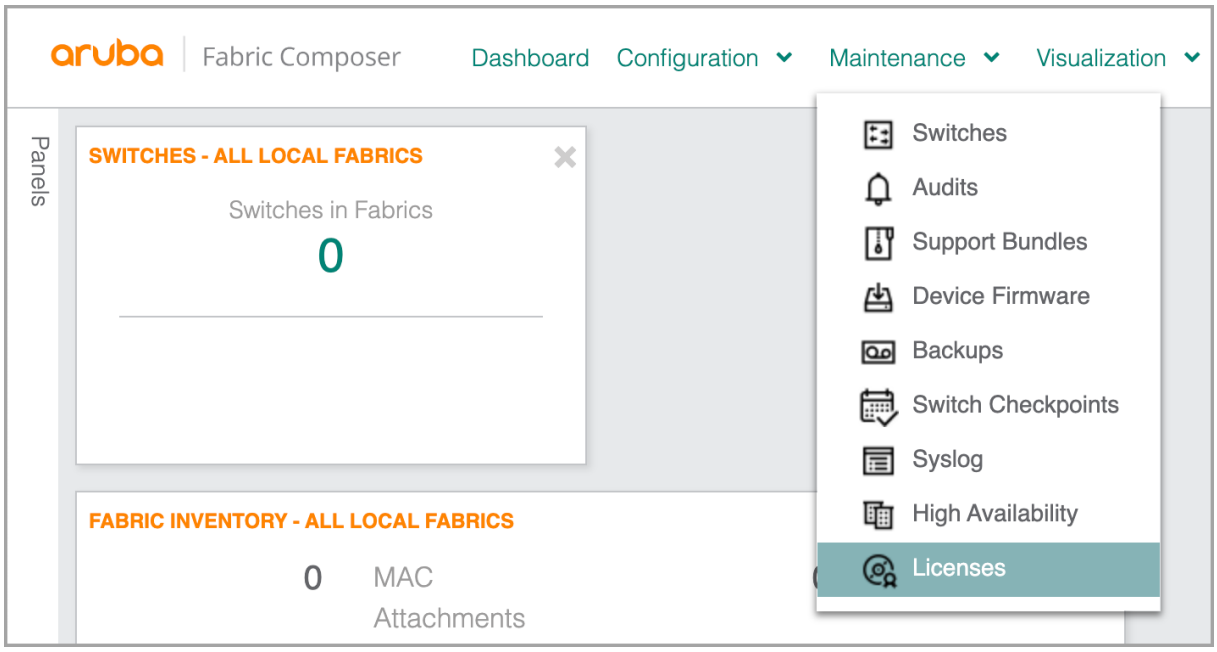


**Step 2** On the **ACTIONS** menu in the **Maintenance/Licenses** pane, select **ADD**.

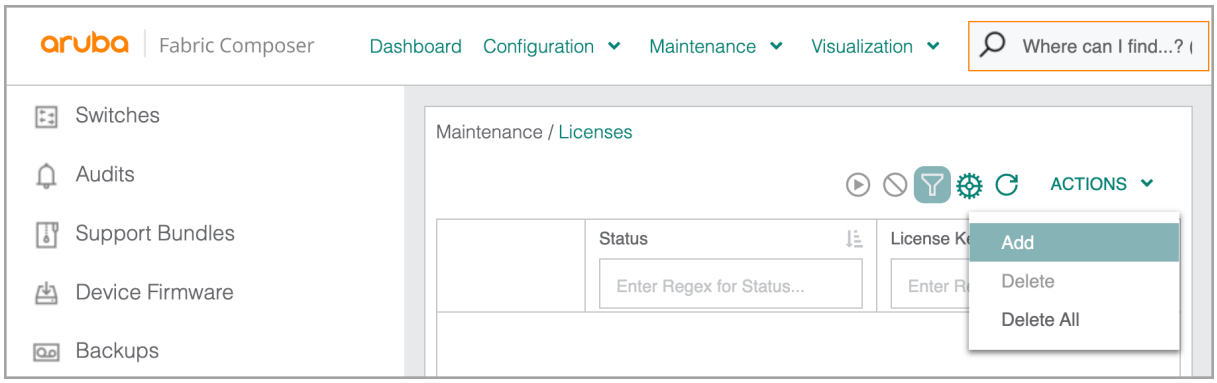


**Step 3** On the **License** page, paste the JSON license string in the **License** field and click **APPLY**.

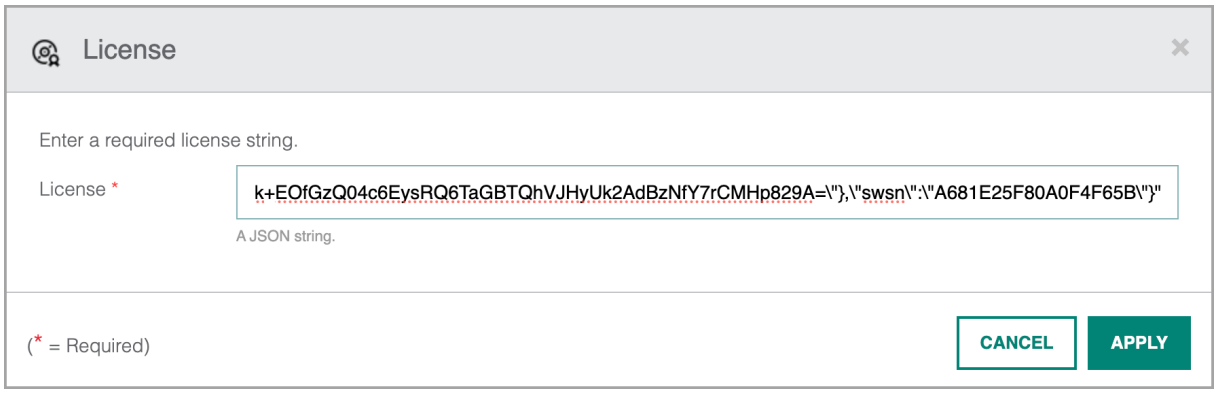


**Step 4** Review the installed license to verify that the **Start Date**, **End Date**, **Quantity**, and **Tier** values display as expected.

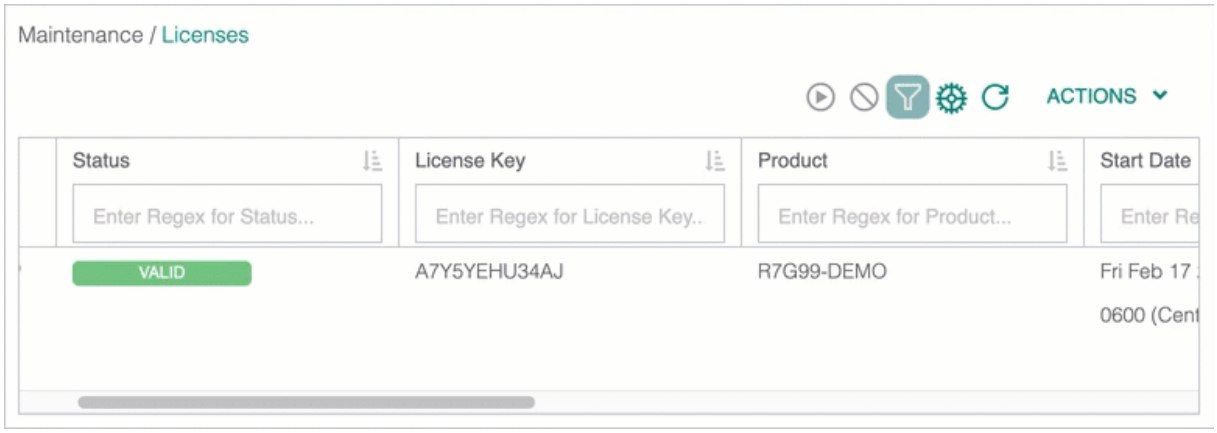


**NOTE:**

Fabric Composer manages two tiers of switches (Tier 3 and Tier 4). The datasheet for each switch model identifies the license tier required.

### Install HPE Aruba Networking Fabric Composer for High Availability

Refer to the *HPE Aruba Networking Fabric Composer Installation Guide* available on the HPE Networking Support Portal. In the “Installing High Availability for HPE ANW Fabric Composer using ISO” section, review the installation requirements and ensure that adequate host resources are available. Follow the steps provided to deploy the HA cluster.

## Download AMD Pensando Policy and Services Manager

When using the firewall capabilities of the CX 10000 switch in a data center, AMD Pensando Policy and Services Manager (PSM) VMs must be installed on a network that is accessible by Fabric Composer and switch management interfaces.

**Step 1** Navigate to https://asp.arubanetworks.com/.

**Step 2** On the menu at the top of the page, select **Software & Documents**.

**Step 3** In the **Search Files** field at the top, type *Pensando*.

**Step 4** In the search results, select the latest OVA version and download it to your computer.

## Install AMD Pensando Policy and Services Manager

In the Aruba Support Portal search results, find the *Pensando Policy and Services Manager for Aruba CX 10000: User Guide*. Review the “PSM Installation” section and ensure that adequate host resources are available. PSM requires a minimum of three VM instances for a production deployment.

**Step 1** Select the OVA file using the **Deploy OVF Template** workflow within vCenter and click **NEXT**.

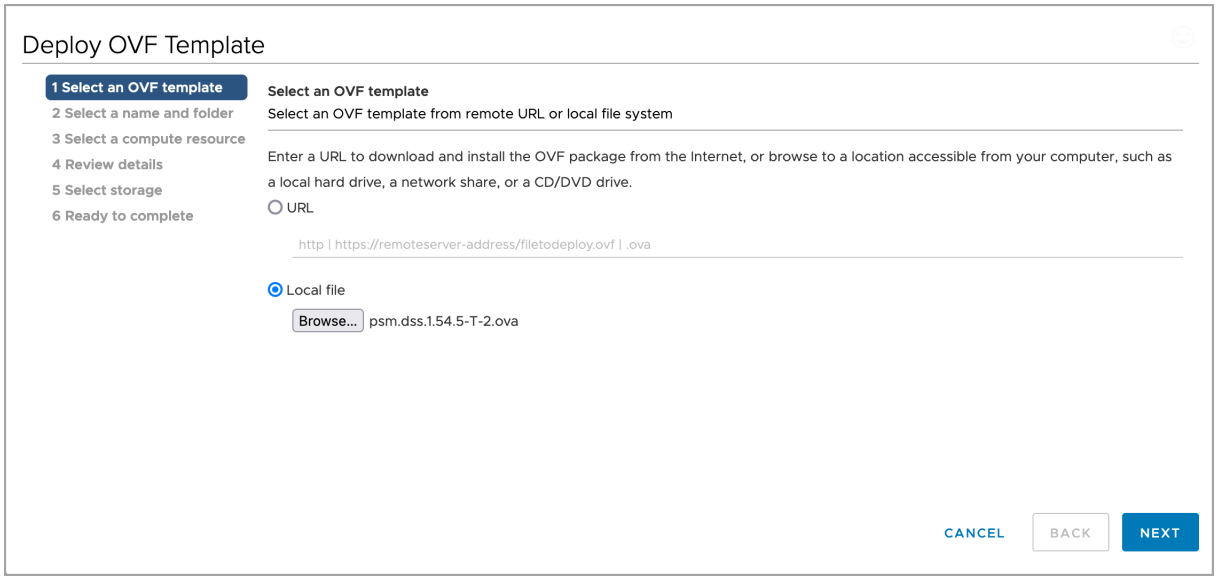


**Step 2** Choose the appropriate options in **Select a compute resource** and proceed through **Review details**.

**Step 3** On the **Configuration** page, click the radio button for **Production** and click **NEXT**.

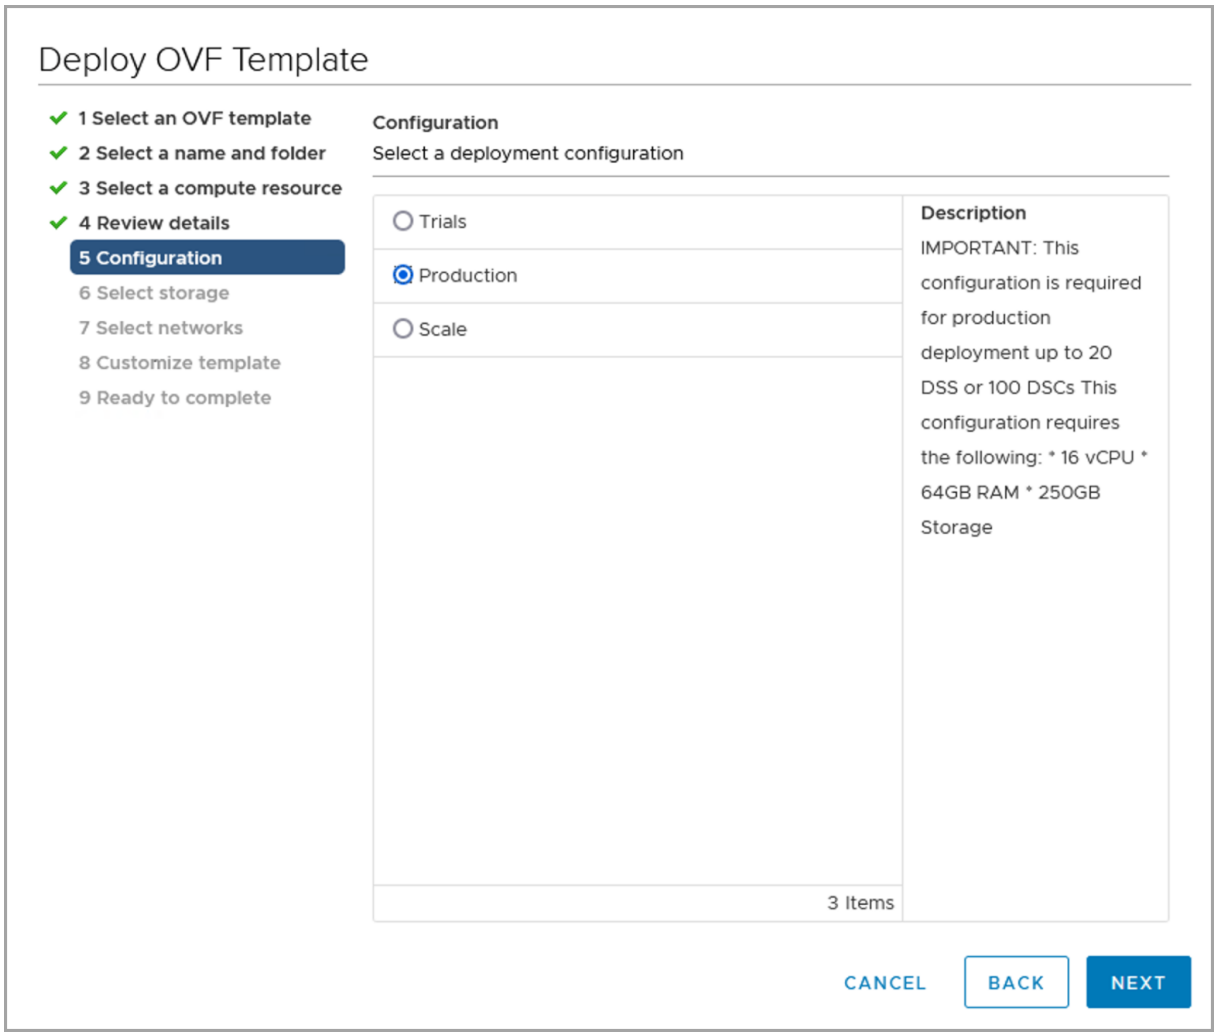


**Step 4** Proceed with selecting the appropriate storage and network resources for the deployment.

**Step 5** Complete the **Customize template** form using the example below.

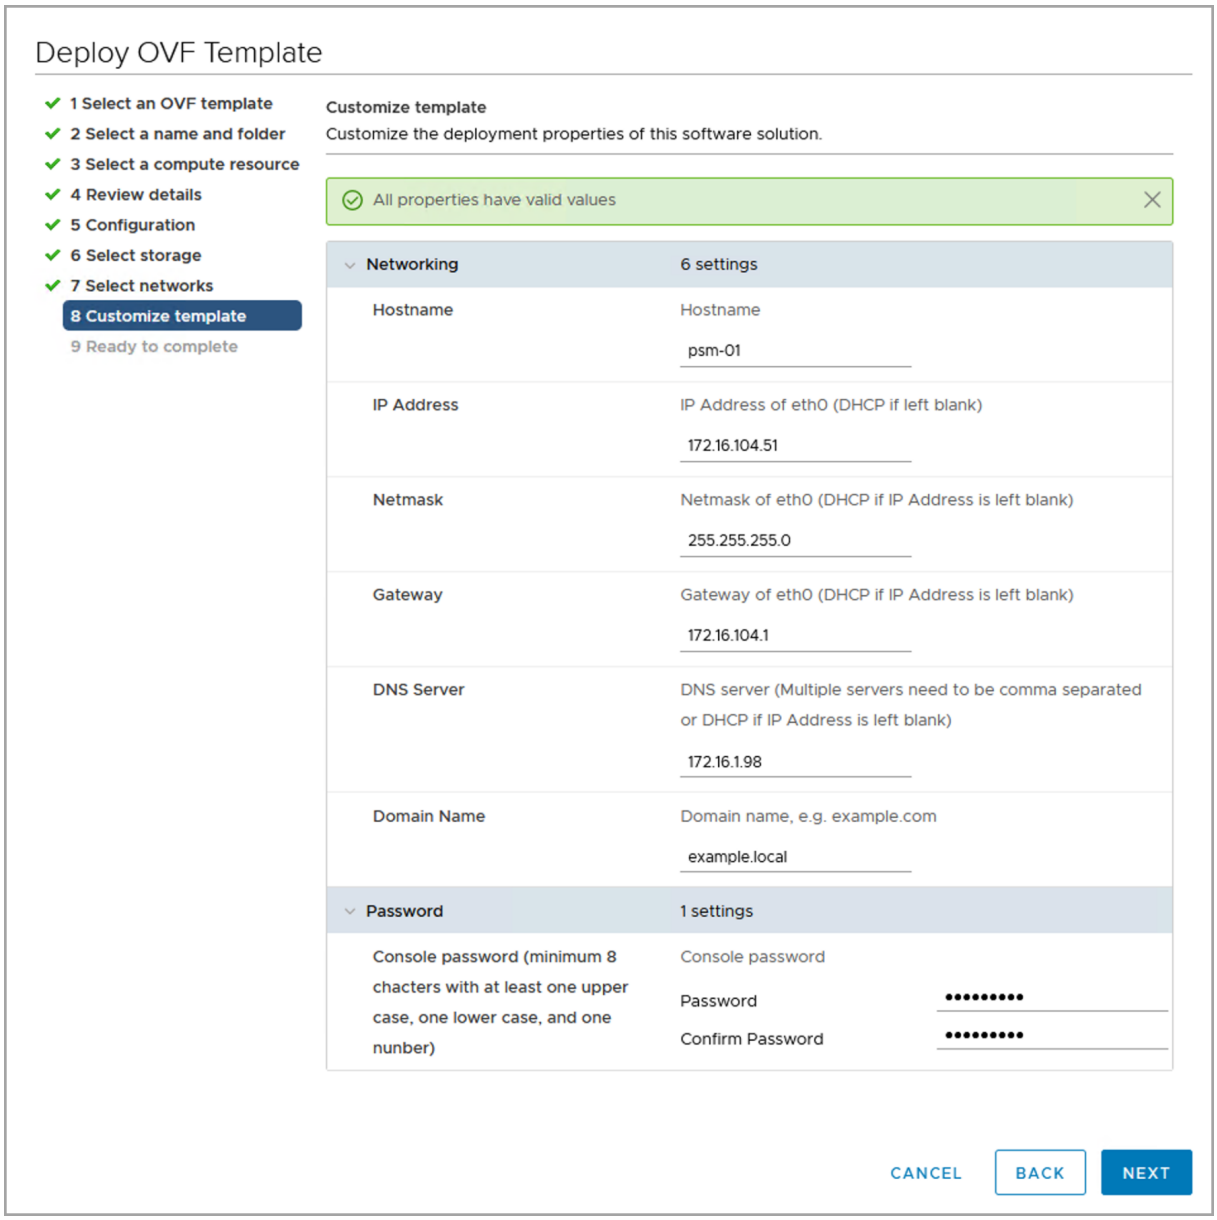


**Step 6** Complete the VM creation workflow.

**Step 7** Create additional PSM VMs as needed.

**NOTE:**

Additional VMs can be created by importing the OVA again or by cloning the first VM as a template as described in the “Installing OVA on ESXi” section of the *Pensando Policy and Services Manager for Aruba CX 10000: User Guide.*

## Configure the AMD PSM Cluster

**Step 1** In vCenter, login to one of the Penando PSM VM consoles.

- **Username:** *root*

- **Password:** < *Specified during VM creation process* >

**Step 2** At the VM console, bootstrap the PSM cluster with the bootstrap_PSM.py utility using the following command-line switch/value pairs followed by a space-delimited list of IP addresses for all cluster members.

- **-enablerouting:** < *No value required* >
- **-distributed_services_switch:** < *No value required* >
- **-autoadmit:** False
- **-clustername:** < *User supplied cluster name* >
- **-domain:** < *Domain name* >
- **-ntpservers:** < *Comma-separated list of NTP servers* >

```
bootstrap_PSM.py -enablerouting -distributed_services_switch -autoadmit False -
    clustername FB1_PSM -domain example.local -ntpservers 172.16.1.98,172.16.1.99
    172.16.104.51 172.16.104.52 172.16.104.53
```

**NOTE:**

The **-autoadmit** command line switch is set to *True* by default. This automatically enables any Distributed Services Switch to join PSM. When a strict admission policy to PSM is required, set this command line switch to *False*.

**Step 3** When prompted, read and accept the End User License Agreement.

**Step 4** Verify that the PSM cluster bootstrap completes successfully.

```
2023-02-16 23:30:55.815164: * PSM bootstrap completed successfully
2023-02-16 23:30:55.815989: * you may access PSM at https://172.16.104.51
2023-02-16 23:30:55.816779: * Note: For backup and disaster recovery, please fetch and store PSM sec
urity token using the "psmctl" command.
2023-02-16 23:30:55.818379: ** To generate and store PSM token which can be used to access all DSE's
 -
2023-02-16 23:30:55.819859: ** /usr/pensando/bin/psmctl get node-token --psm-ip localhost --psm-port
 443 --audience "*" --token-output ~/dse-tok
[root@psm-01 ~]#
```

**Step 5** On the VM console, enter the following to generate a PSM security token.

```
/usr/pensando/bin/psmctl get node-token --psm-ip localhost --psm-port 443 --
    audience "*" --token-output ~/dse-tok
```

**NOTE:**

The token can be used for disaster recovery and backup purposes. Store it with other sensitive network credentials.

**Step 6** When prompted, enter the following default credentials:

- **User name:** *admin*
- **Password:** *Pensando0$*

**Step 7** Open a web browser and connect to PSM at one of the configured VM IP addresses.

**Step 8** On the **AMD Pensando** login page, enter the following default credentials and click **SIGN IN**.

- **Username:** *admin*
- **Password:** *Pensando0$*


**Step 9** Go to **System > Cluster** and verify that each PSM VM is listed under **Nodes** in the **Cluster Detail** pane with the following values.

- **Quorum:** *true*
- **Phase:** *Joined*

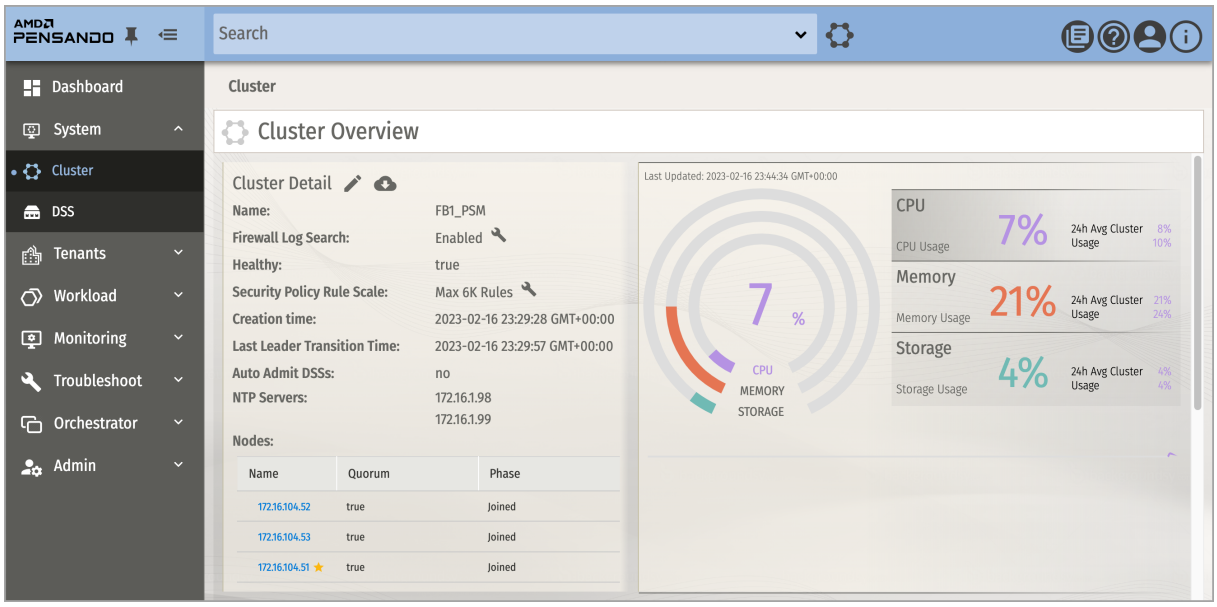


**Step 10** Go to **Admin > User Management**, mouse-over the **admin** user, and click the **Change password** icon.

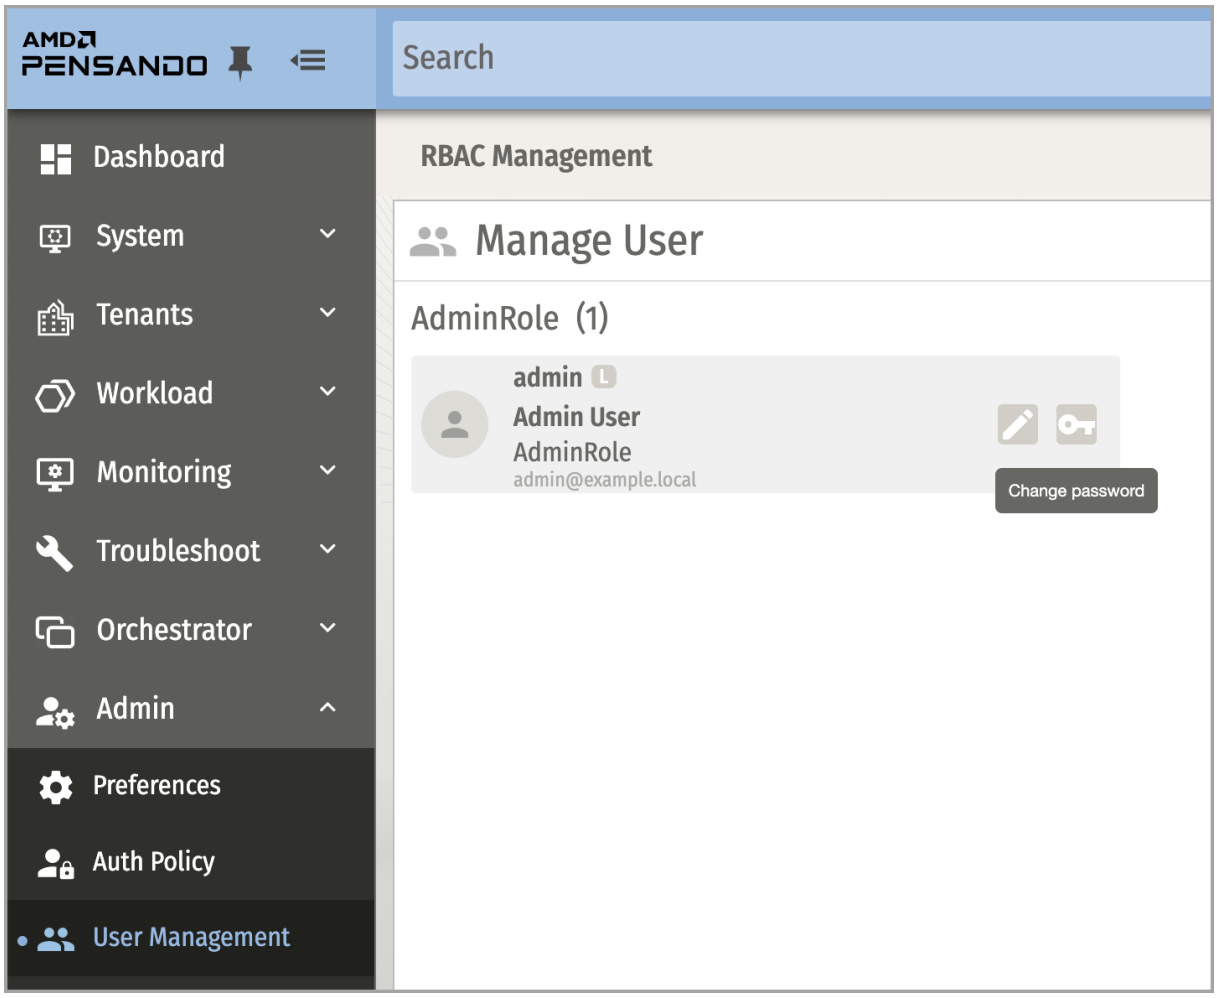


**Step 11** Enter the old and new passwords and click **Save changes**.

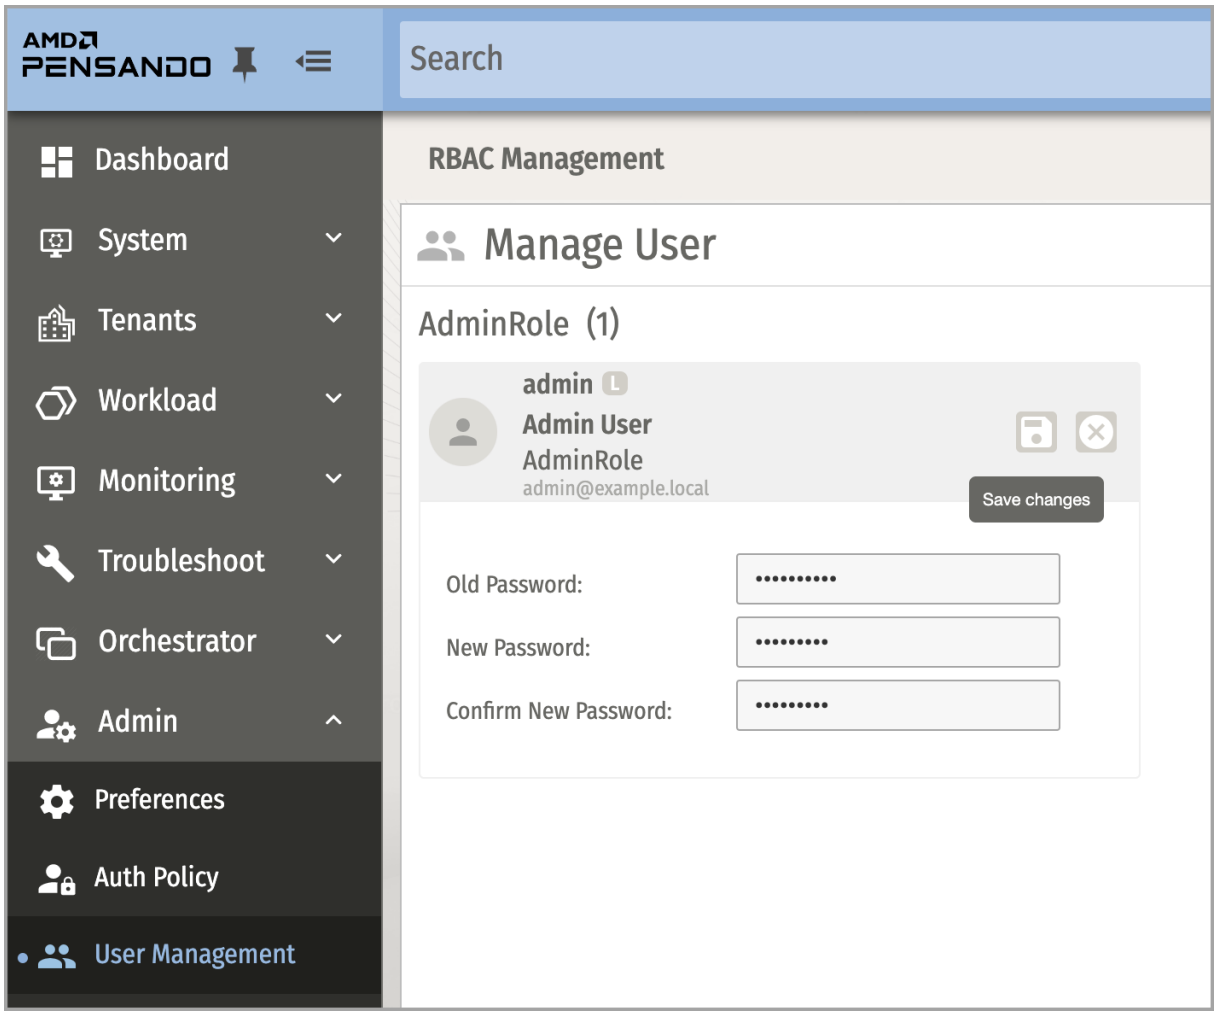


**NOTE:**

Changing the password on one VM updates all cluster members.

# EVPN-VXLAN Configuration

Configuring an HPE Aruba Networking data center fabric is best performed using the HPE Aruba Networking Fabric Composer guided setup process. Fabric Composer automates switch provisioning, underlay link and routing configuration, overlay configuration, and integration with VMware vCenter.

## Physical Topology Overview

The diagram below illustrates the physical links and hardware that comprise the primary data center in this guide. Fabric Composer is used to configure a routed underlay and EVPN-VXLAN overlay for the topology.

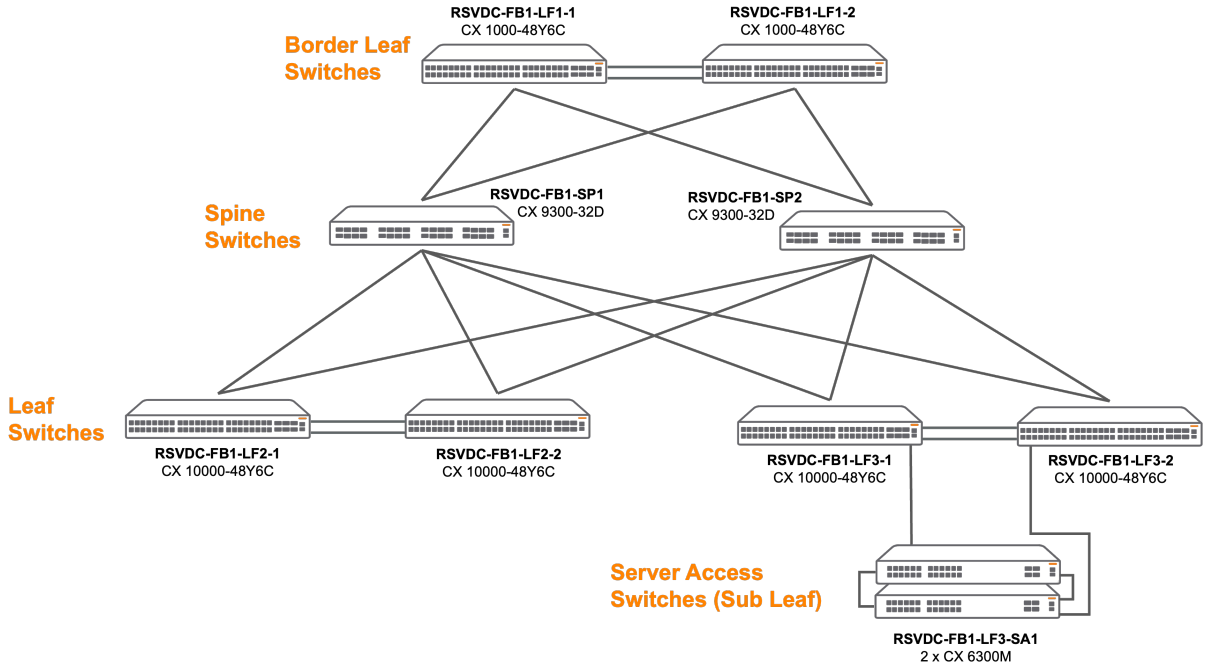


## HPE Aruba Networking Fabric Composer Process

Fabric Composer's Guided Setup automates configuration following these steps:

- **Switch discovery:** Discover and inventory data center switches in Fabric Composer.
- **Fabric Composer fabric creation:** Define the logical construct that identifies a fabric within Fabric Composer.
- **Switch assignment:** Assign roles to fabric switches.
- **NTP and DNS configuration:** Assign NTP and DNS servers to fabric switches.
- **VSX configuration:** Create VSX-redundant ToR leaf pairs.
- **Leaf/Spine configuration:** Assign IP addresses to leaf/spine links.
- **Underlay configuration:** Establish OSPF underlay to support the EVPN-VXLAN overlay data plane and control plane.

- **Overlay configuration:** Establish BGP peerings to enable the EVPN overlay control plane and VXLAN tunnel endpoints for overlay data plane.
- **EVPN configuration**: Establish Layer 2 EVPN mapping of VLANs to VXLAN Network Identifiers (VNIs).

When the Guided Setup is complete, additional configuration details for host onboarding, external fabric connectivity, testing, and multicast are required:

- Layer 3 services within overlays.
- Multi-chassis LACP LAG configuration for host connectivity.
- Routing between the data center and campus.
- Overlay loopbacks for testing reachability to directly connected hosts and resources both inside and external to the fabric.
- PIM-SM and IGMP to support overlay multicast services.

For additional details on the Guided Setup steps, refer to the “Guided Setup” section of the *HPE Aruba Networking Fabric Composer User Guide*.

# Plan the Deployment

Before starting the guided setup, plan ahead and develop a naming convention and address scheme with values that can accommodate the current deployment size and leave room for growth. Using a consistent approach in the physical and logical configurations improves the management and troubleshooting characteristics of the fabric.

This section provides sample values and rationale. Adjust the values and formats as needed to accommodate the current and projected sizes of the fabric effectively.

## Naming Conventions

Fabric Composer supports the execution of operations on a single switch or on a selected group of switches.

Establish a switch naming convention that indicates the switch type, role, and location to simplify identification and increase efficiency when operating production-scale fabrics. Configure switch names before importing them into Fabric Composer.

Example values used in this guide:

| Switch Name | Fabric Role | Description |
|---|---|---|
| RSVDC-FB1-SP1 | Spine | Fabric #1, Spine #1 |
| RSVDC-FB1-LF1-1 | Leaf | Fabric #1, VSX Leaf Pair #1, Member #1 |

| Switch Name | Fabric Role | Description |
|---|---|---|
| RSVDC-FB1-LF1-2 | Leaf | Fabric #1, VSX Leaf Pair #1, Member #2 |
| RSVDC-FB1-LF3-SA1 | Server Access (Sub Leaf) | Fabric #1, VSF Server Access Stack #1 attached to Leaf Pair #3 |

**NOTE:**

VSF stacks used in the server access role contain two or more switches. The stack operates as a single logical switch with a single control plane. It is not possible to differentiate between stack members using a unique hostname.

The Guided Setup prompts for a **Name Prefix** on some steps. Name prefixes are logical names used within Fabric Composer. Choose a descriptive name to make it easy to monitor, edit, and execute operations. The procedures below include examples of effective names that can be used.

## Underlay Connectivity and Addressing

Point-to-point connections between spine-and-leaf switches are discovered and configured automatically for IP connectivity using /31 subnets within a single network range. Fabric Composer supports addressing up to 128 links inside a fabric using a /24 subnet mask. The maximum number of links on a fabric is determined by the aggregate port count of the spine switches.

Another network range is provided to create: * A /32 loopback address on each switch, used as the router ID for OSPF and BGP. * A /31 transit VLAN between ToR switch pairs to ensure data plane continuity in case of host link failure. * A /31 point-to-point interface between ToR switch pairs to transmit keep-alive messages for VSX peer loss detection.

Fabric Composer creates each of these subnet types automatically from a single network range provided during the VSX setup process. If VSX is not used, the network range is provided during the underlay configuration process.

Example values used in this guide are:

| Purpose | Description | Example |
|---|---|---|
| Leaf-Spine IP address block | An IPv4 address block used to create /31, point-to-point layer 3 links between leaf and spine switches. | 10.255.0.0/23 |

| Purpose | Description | Example |
|---|---|---|
| Routed loopback, VSX transit VLAN, and VSX Keep-Alive Interface IP address block | An IPv4 address block used to allocate unique loopback addresses (/32) for each switch, for VSX keep-alive point-to-point connection (/31) and also used as a transit-routed VLAN between redundant ToRs (/31) | 10.250.0.0/23 |

## Overlay Connectivity and Addressing

The overlay network is created using VXLAN tunnels established between Virtual Tunnel Endpoints (VTEPs) within the leaf switches in the fabric. Loopback addresses assigned to establish route peerings are unique per switch and cannot be used as a VTEP IP when using VSX. A single logical VTEP per rack is defined by creating a dedicated /32 loopback interface common to both ToR peer switches. The interfaces are assigned automatically from a single subnet scope provided during the overlay guided setup.

| Purpose | Description | Example |
|---|---|---|
| VTEP IP address block | An IPv4 address block used to allocate VXLAN tunnel endpoint (VTEP) loopback addresses (/32) for each ToR switch pair | 10.250.2.0/24 |

A Virtual Network Identifier (VNI) is a numerical value that identifies network segments within the fabric's overlay topology. The VNI is carried in the VXLAN header to enable switches in the fabric to identify the overlay to which a frame belongs and apply the correct policy to it.

When configuring the overlay topology, a Layer 3 VNI represents the routed component of the overlay. Each Layer 3 VNI maps to a VRF. A Layer 2 VNI represents the bridged component of the overlay. Each Layer 2 VNI maps to a VLAN ID. Multiple Layer 2 VNIs can be associated to a single VRF.

Plan your VNI numbering scheme in advance to ensure that values do not overlap. Example values used in this guide are:

| VNI Type | Description | Example |
|---|---|---|
| L2 VNI | VLAN ID + 10,000 | VLAN100 == L2 VNI 10100, VLAN200 == L2VNI 10200 |
| L3 VNI | Overlay # + 100,000 | Overlay1 == L3 VNI 100001, Overlay2 == L3 VNI 100002 |

Internal BGP (iBGP) is used to share overlay reachability information between leaf switches. Layer 3 and Layer 2 information associated to a local switch's VNIs is advertised with its associated VTEP to other members of the fabric. Two of the spines operate as BGP route reflectors. All leaf switches are clients of the two route reflectors.

A unique IP loopback IP address is assigned to each overlay VRF for testing and troubleshooting.

Overlay VLAN switched virutal interface (SVI) and Active Gateway IP address assignments are not unique on leaf switches. A ping test to a directly attached host is not supported when the SVI/AG IP is the traffic source, because the ping response may be sent to the switch that did not originate the ping. Using a unique IP source from a loopback IP allows the attached switch to source a ping test to the host.

Additionally, the SVI/AG IP address for each VLAN is present on all leaf switches in a fabric. A unique loopback IP provides a source IP address for testing reachability within the fabric and to external hosts. It also provides a unique destination IP to test individual switch reachability within an overlay.

Plan an IP block per overlay VRF large enough that a unique IP address can be assigned to each leaf switch in the fabric. A single maskable block allows summarizing route advertisements to external networks.

| Purpose | Description | Example |
|---|---|---|
| VRF 1 IP address block | An IPv4 address block used to assign loopback IPs for reachability testing in VRF 1 | 10.250.4.0/24 |
| VRF 2 IP address block | An IPv4 address block used to assign loopback IPs for reachability testing in VRF 2 | 10.250.5.0/24 |

Each VSX pair requires an overlay transit VLAN to share routed reachability of the loopback IP addresses assigned from the IP block above. Assign a single maskable block of IP addresses for each overlay VRF, where /31 blocks can be assigned to the transit VLANs.

| Purpose | Description | Example |
|---|---|---|
| VRF 1 IP address block | An IPv4 address block used to assign loopback IPs for reachability testing in VRF 1 | 10.255.4.0/24 |
| VRF 2 IP address block | An IPv4 address block used to assign loopback IPs for reachability testing in VRF 2 | 10.255.5.0/24 |

## MAC Address Best Practice

A Locally Administered Address (LAA) should be used when Fabric Composer requires entry of a MAC address for the switch virtual MAC, a VSX system-MAC, or an Active Gateway MAC for a distributed SVI. An LAA is a MAC in one of the four formats shown below:

```
x2-xx-xx-xx-xx-xx
x6-xx-xx-xx-xx-xx
xA-xx-xx-xx-xx-xx
xE-xx-xx-xx-xx-xx
```

The *x* positions can contain any valid hex value. For more details on the LAA format, see the IEEE tutorial guide.

An active gateway IP distributes the same gateway IP across all leaf switches in a fabric to support gateway redundancy and VM movement across racks. An active gateway MAC associates a virtual MAC address with an active gateway IP. Only a small number of unique virtual MAC assignments can be configured per switch. The same active gateway MAC address should be reused for each active gateway IP assignment.

## HPE Aruba Networking Fabric Composer Prerequisites

The following items must be configured before building a Fabric Composer-based fabric.

- **Physically cable all switches in the topology.** VSX pairs, VSF stacks, and leaf-spine links must be connected fully to support Fabric Composer's automation.
- **Configure VSF stacking for server access switches.** When optional server access switches are present, VSF auto-stacking must be configured when the switches are at their default configuration. VSF configuration guidance is available in the HPE Networking Support Portal. Enable split detection after the stack is formed.
- **Assign management interface IP addresses.** A DHCP scope using MAC address reservations for each switch can be used in place of manual IP address assignment. When using DCHP, MAC address reservations ensure that each switch is assigned a consistent IP address.
- **Assign switch hostnames**. Assigning unique hostnames using a naming convention helps administrators identify a switch and its role quickly during setup and future troubleshooting.

## Fabric Initialization

Configuring an HPE Aruba Networking data center fabric using a spine-and-leaf topology is best performed using the Fabric Composer guided setup process. To return to guided setup at any time, simply select it in the menu bar at the top right of the Fabric Composer user interface.

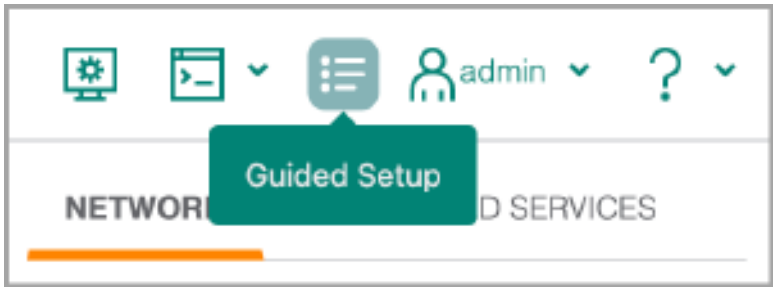

## Discover Switches on the Network

The first procedure adds switches to the Fabric Composer device inventory. An orderly naming convention for switch host names should be implemented before continuing with this procedure in order to simplify switch selection in the following steps.

**Step 1** On the **Guided Setup** menu, select **SWITCHES**.

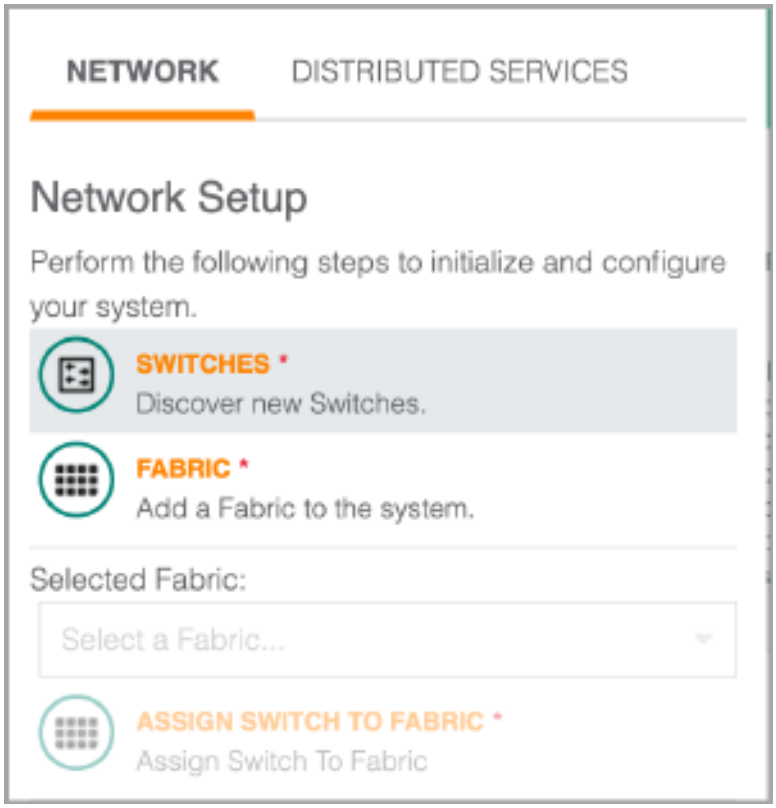


**Step 2** In the **Discover Switches** window, enter the following switch information and click **APPLY**.

- **Switches:** < *OOBM IP addresses for fabric switches* >
- **admin Switch Password:** < *password created during switch initialization* >
- **admin Switch Password:** < *password created during switch initialization* >
- **Service Account Password:** < *new password for the afc_admin account* >
- **Confirm Service Account Password:** < *new password for the afc_admin account* >

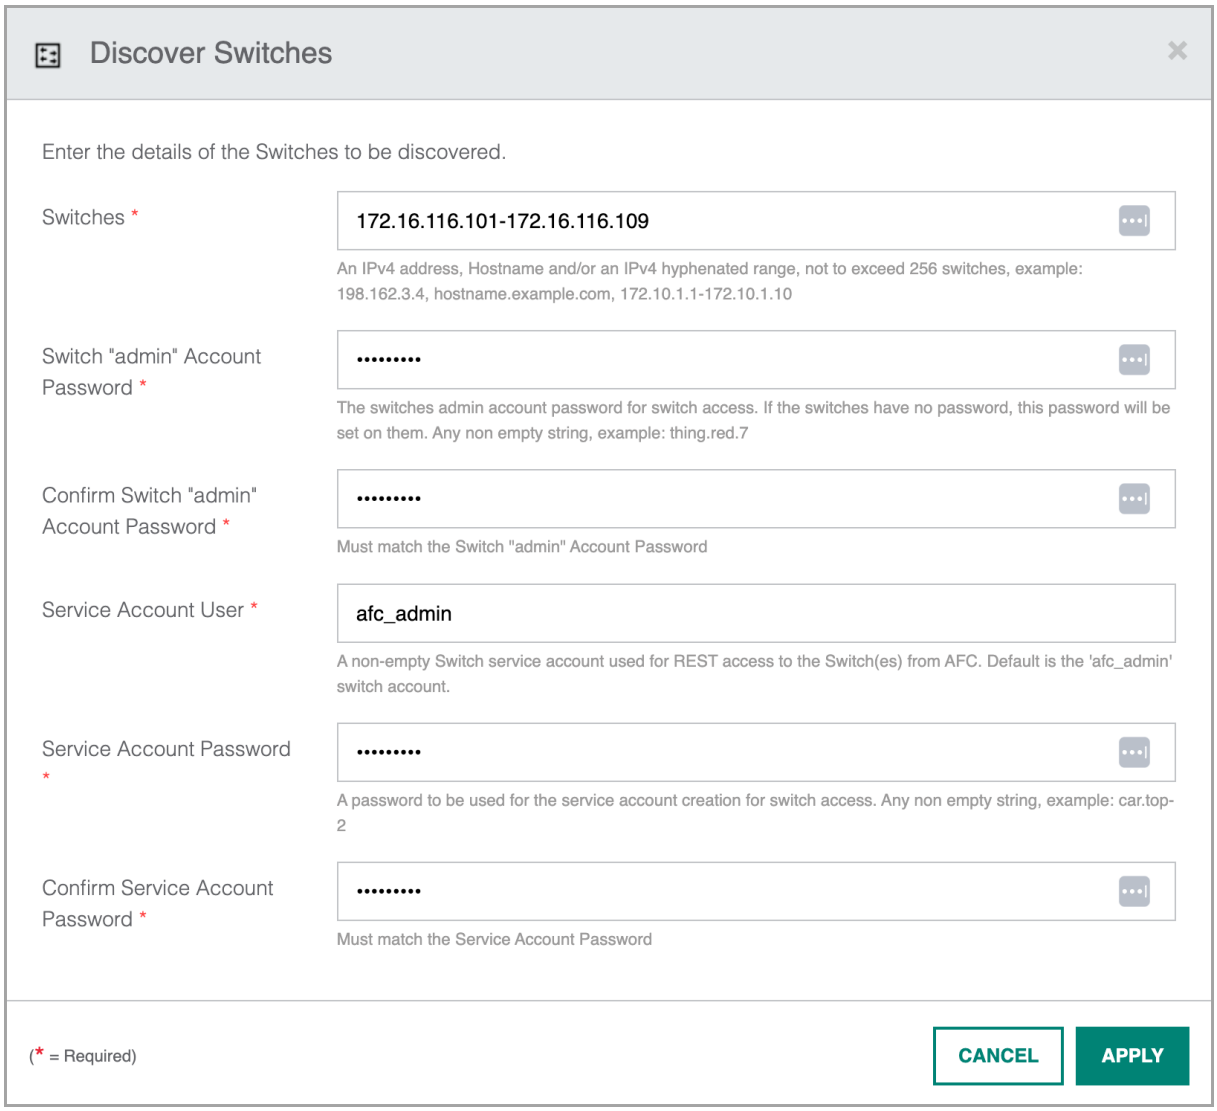


**NOTE:**

Switch IP addresses can be entered in a comma-separated list or in one or more ranges. If the IP addresses provided include devices not supported by Fabric Composer or switches with different credentials, a “Discovery Partially Successful” warning message appears after the import.This step creates a new afc_admin account on all the switches for API access from Fabric Composer.

**Step 3** Review the list of imported switches in the **Maintenance > Switches** window and verify that the health status of each switch is **HEALTHY, BUT...** Hovering over the health status value of an individual switch provides additional details.

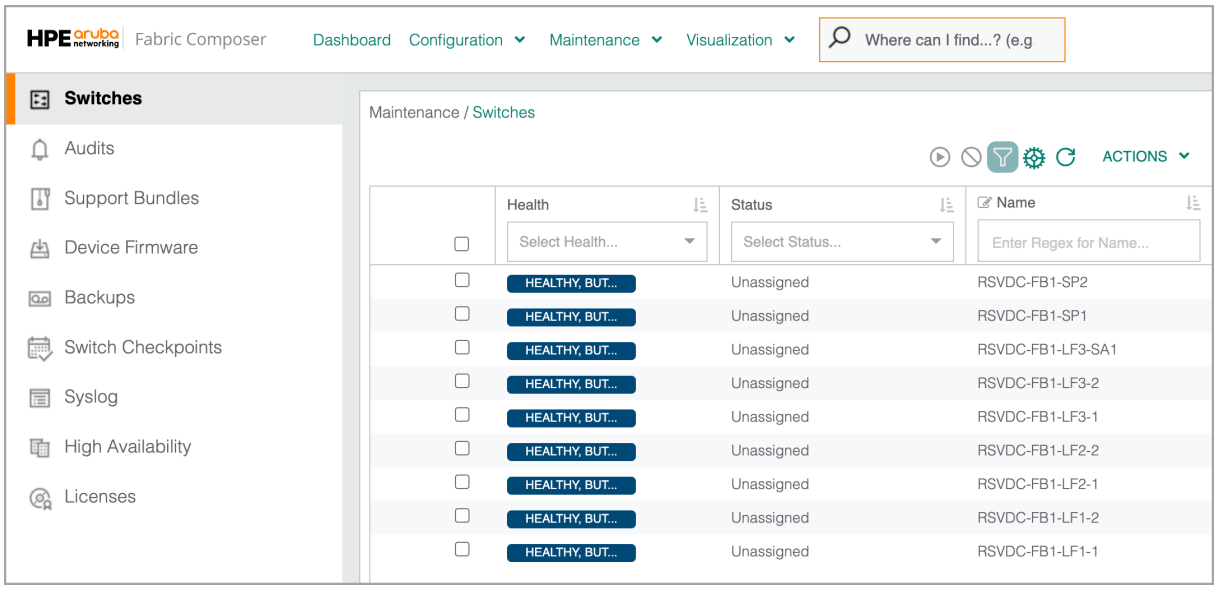


## Create a Fabric

A fabric container is created in Fabric Composer for collective configuration of a group of switches. The fabric name is internal to Fabric Composer operations and is not tied to configuration elements on a switch. Fabric Composer supports the configuration of spine and leaf, Layer 2 two-tier, and management networks. All topologies assign switches to a Fabric Composer internal fabric for configuration and management.

The fabric in this guide is used to implement a spine-and-leaf routed network with an EVPN-VXLAN overlay.

**Step 1** On the **Guided Setup** menu, select **FABRIC**.

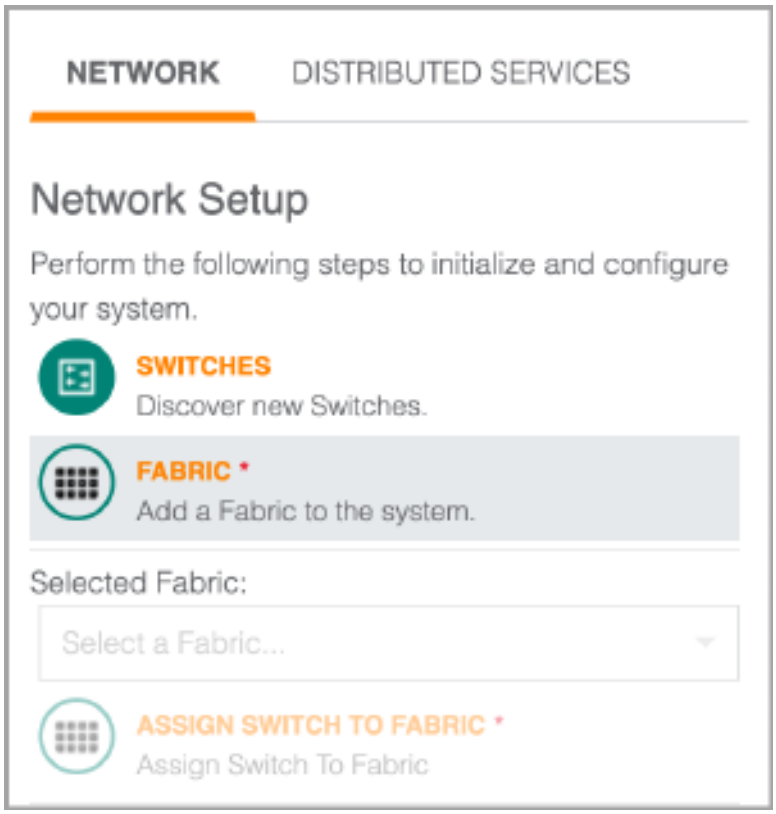


**Step 2** Define a unique logical name, set the **Type** to *Data*, specify a time zone, and click **APPLY**.

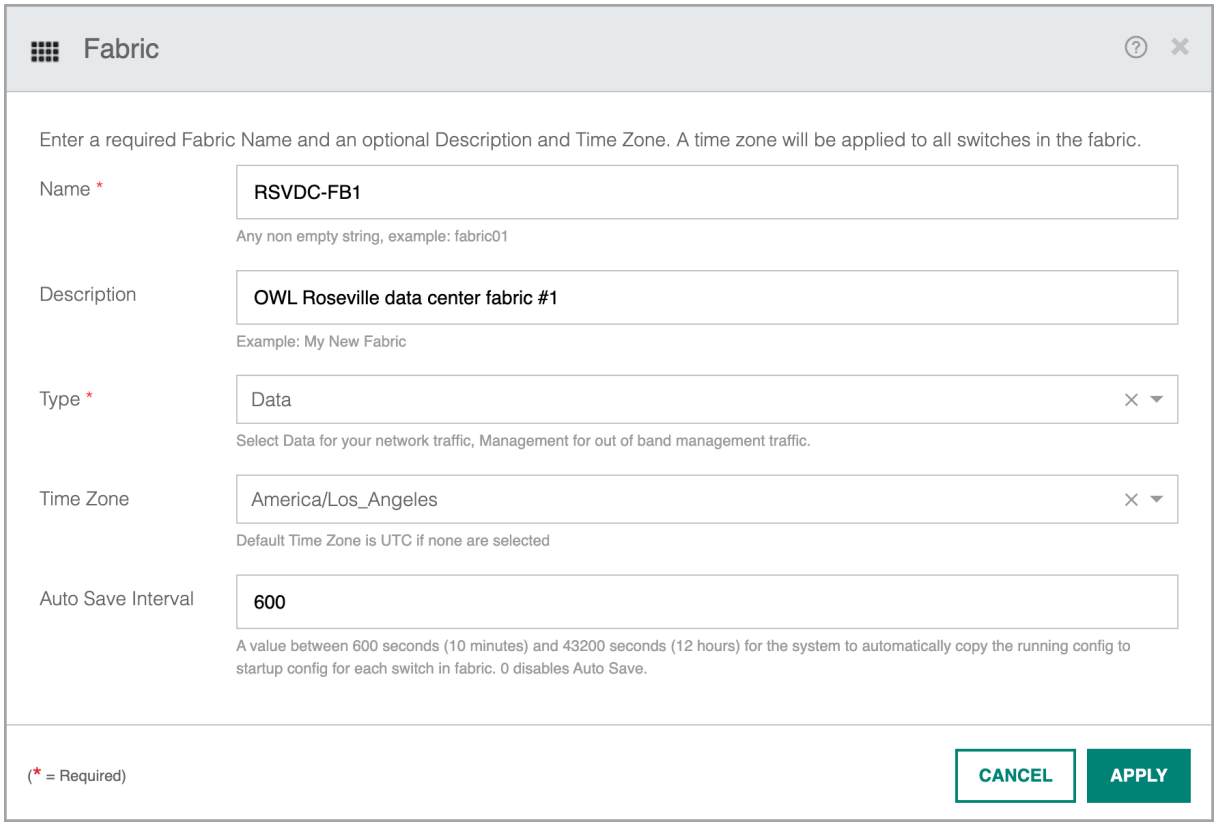

## Assign Switches to the Fabric

Switches must be added to a fabric before they can be configured. When adding a switch to a fabric, a role is declared. In the following steps, begin by adding spine switches. Leaf switches can then be added more easily as a group.

**Step 1** On the **Guided Setup** menu, verify that the fabric created in the previous step appears under **Selected Fabric** and click **ASSIGN SWITCH TO FABRIC**.

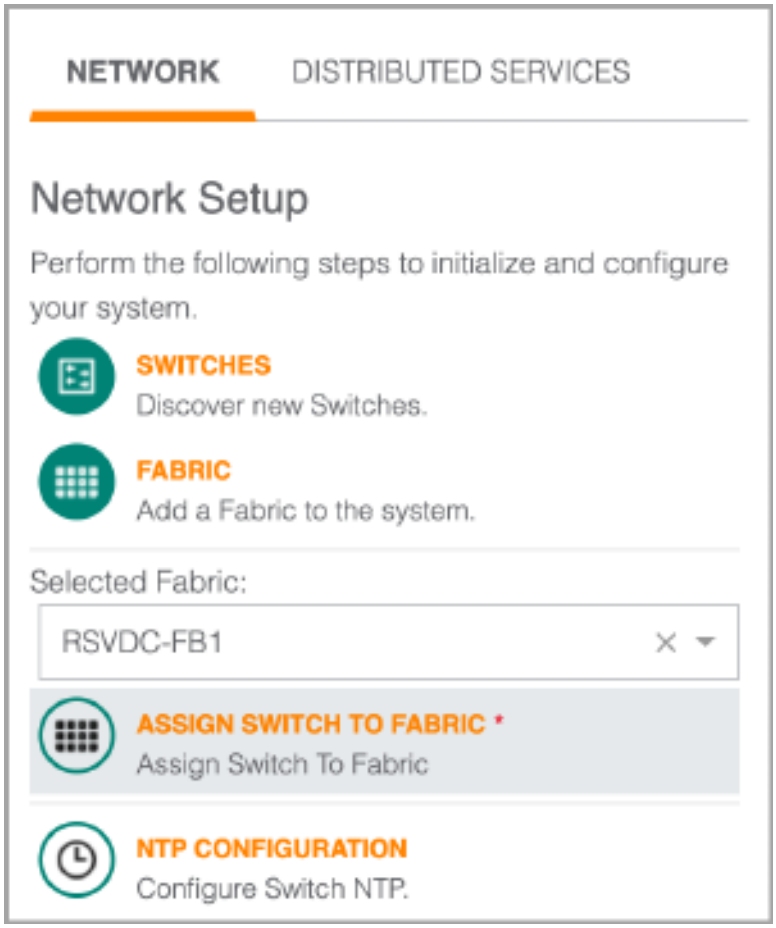


**Step 2** Assign switches to the fabric grouped by role. Assign the following values for spine switches, then click **ADD**.

- **Fabric:** *RSVDC-FB1*
- **Switches:** *< All spine switches >*
- **Role:** *Spine*
- **Force LLDP Discovery:** *checked*
- **Initialize Ports:** *checked*
- **Exclude this switch from association to any Distributed Services Manager:** *unchecked*

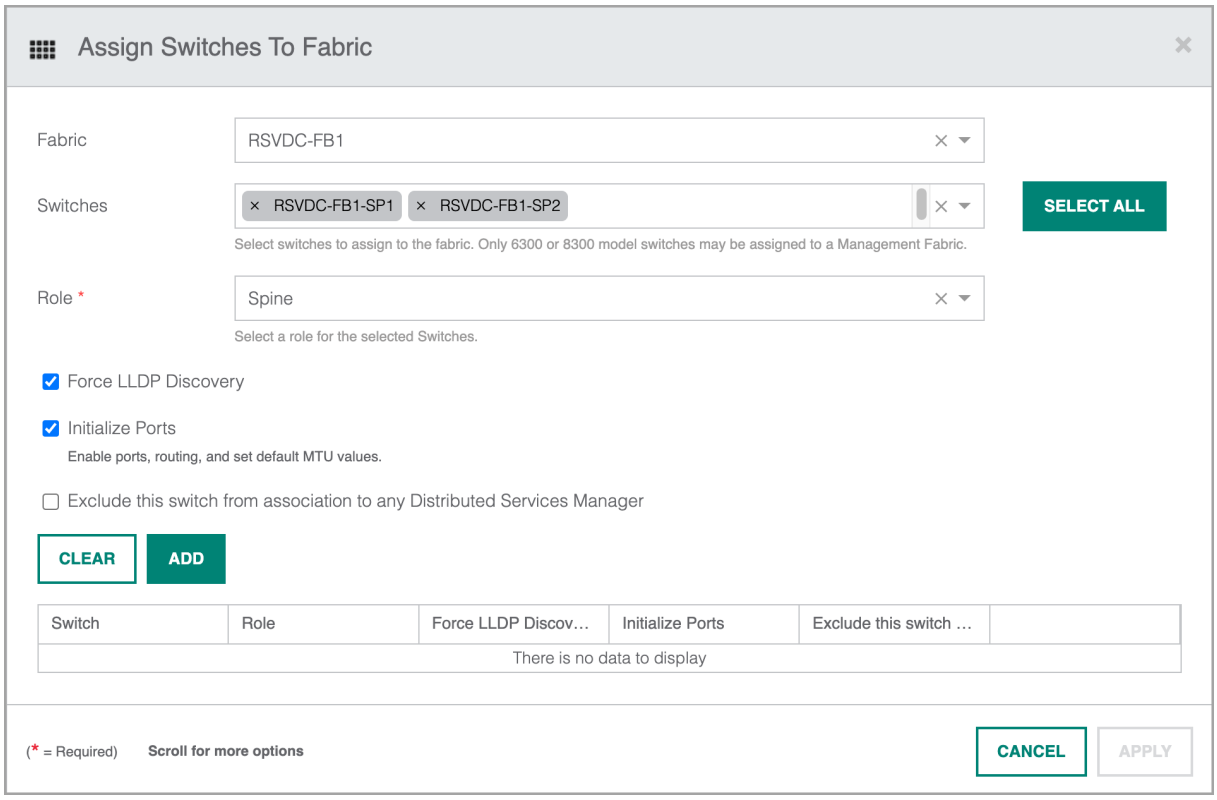


**NOTE:**

Checking **Initialize Ports** enables all switch ports for use in LLDP neighbor discovery. Split port configuration is performed in the previous *Switch Initialization* procedure to allow proper port initialization by Fabric Composer. The MTU of the physical ports also is adjusted to 9198 in order to support jumbo frames that allow VXLAN encapsulation overhead. Checking **Force LLDP Discovery** prompts Fabric Composer to use LLDP neighbor information to discover link topology between spine-and-leaf switches and ToR VSX pairs dynamically.

**Step 3** Repeat the steps above for VSF server access switch stacks. Verify that all server access switch stacks are listed with the **Sub Leaf** role selected and click **ADD**.

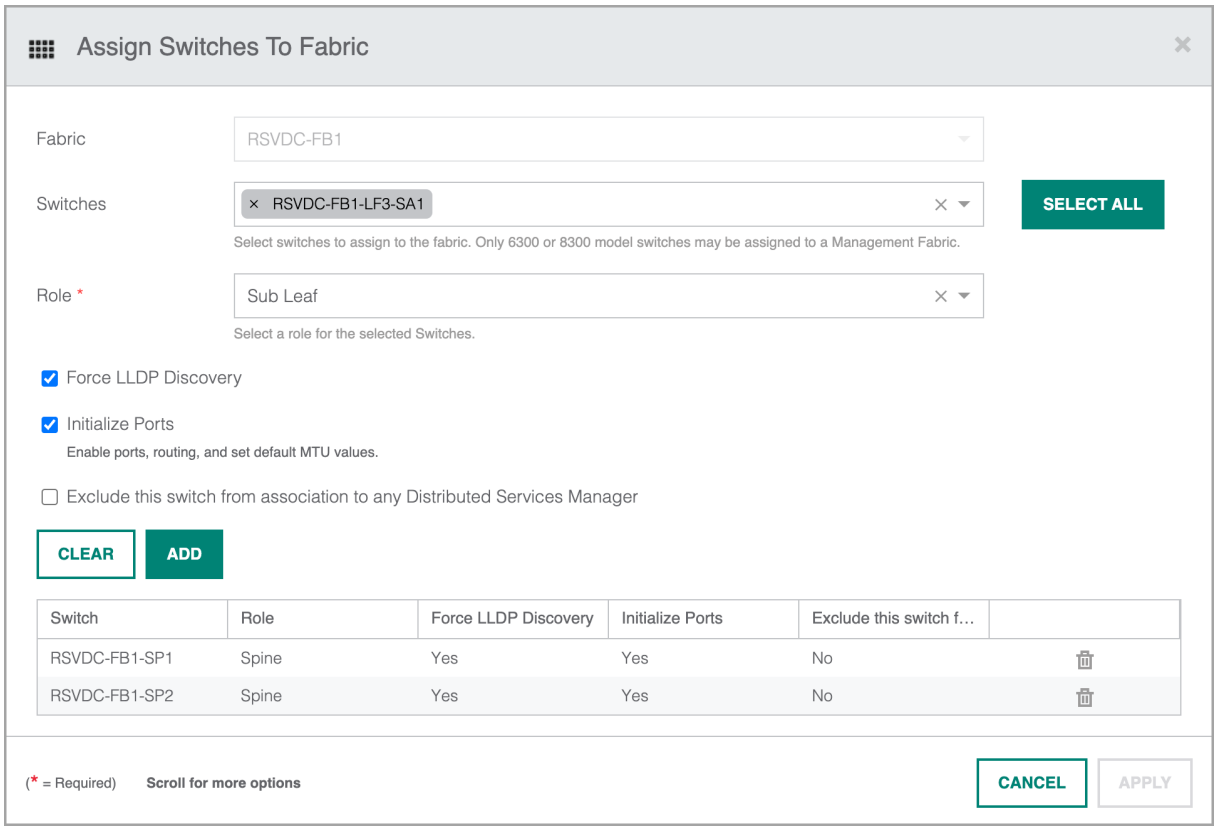


**NOTE:**

This step is optional. It is required only when server access switches are present in the topology. Each VSF switch stack has a single entry that represents all switch members of the stack. This example implementation contains a single VSF stack.

**Step 4** Repeat the previous step for border leaf switches with the **Border Leaf** role selected and click **ADD**.

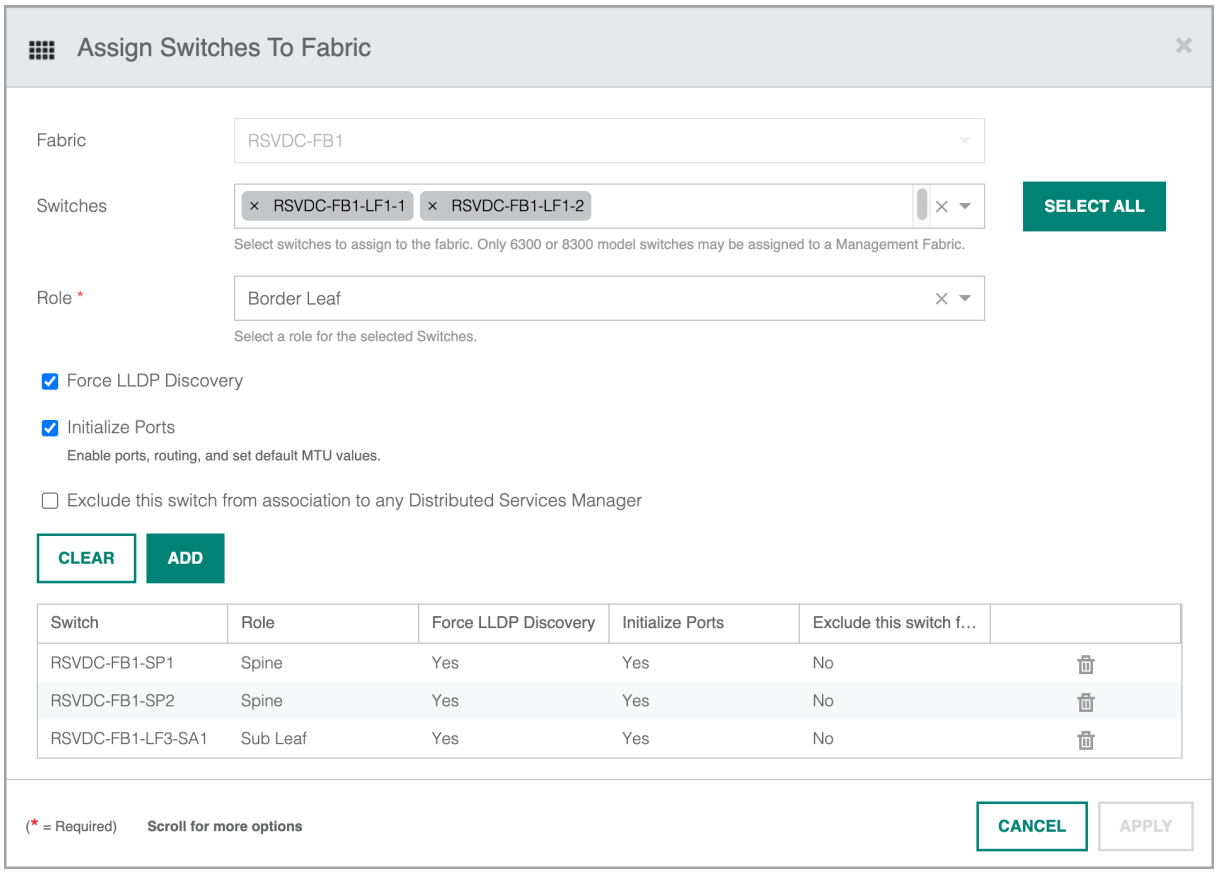


**Step 5** Repeat the previous step for the remaining leaf switches with the **Leaf** role selected and click **ADD**.

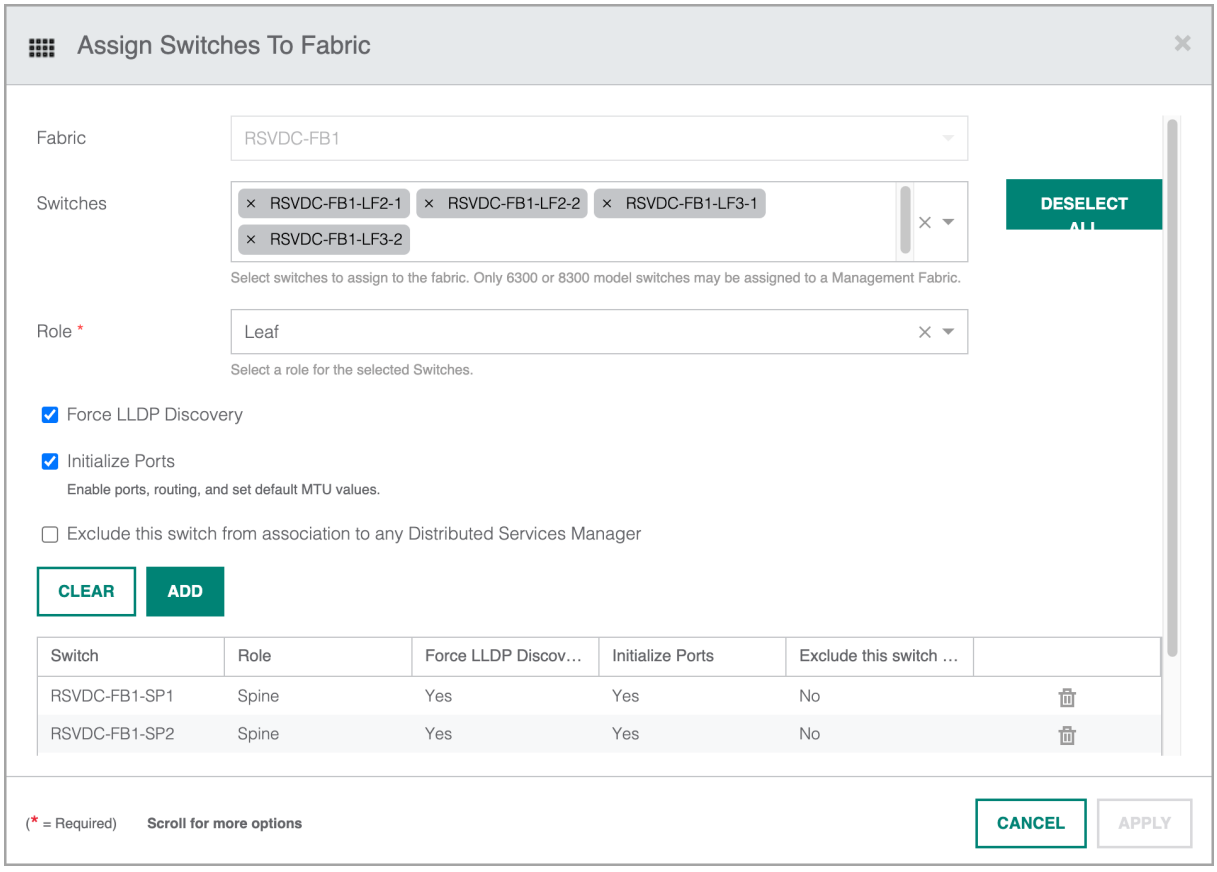

**NOTE:**

Leaf switches typically comprise the majority of switches in a fabric. Use **SELECT ALL** to catch all remaining leaf switches, when switch assignments containing smaller sets of switches are assigned first.

**Step 6** Scroll through the list of switches to verify role assignments and ensure successful configuration of the fabric. After adding all switches to the fabric with the correct role, click **APPLY**.

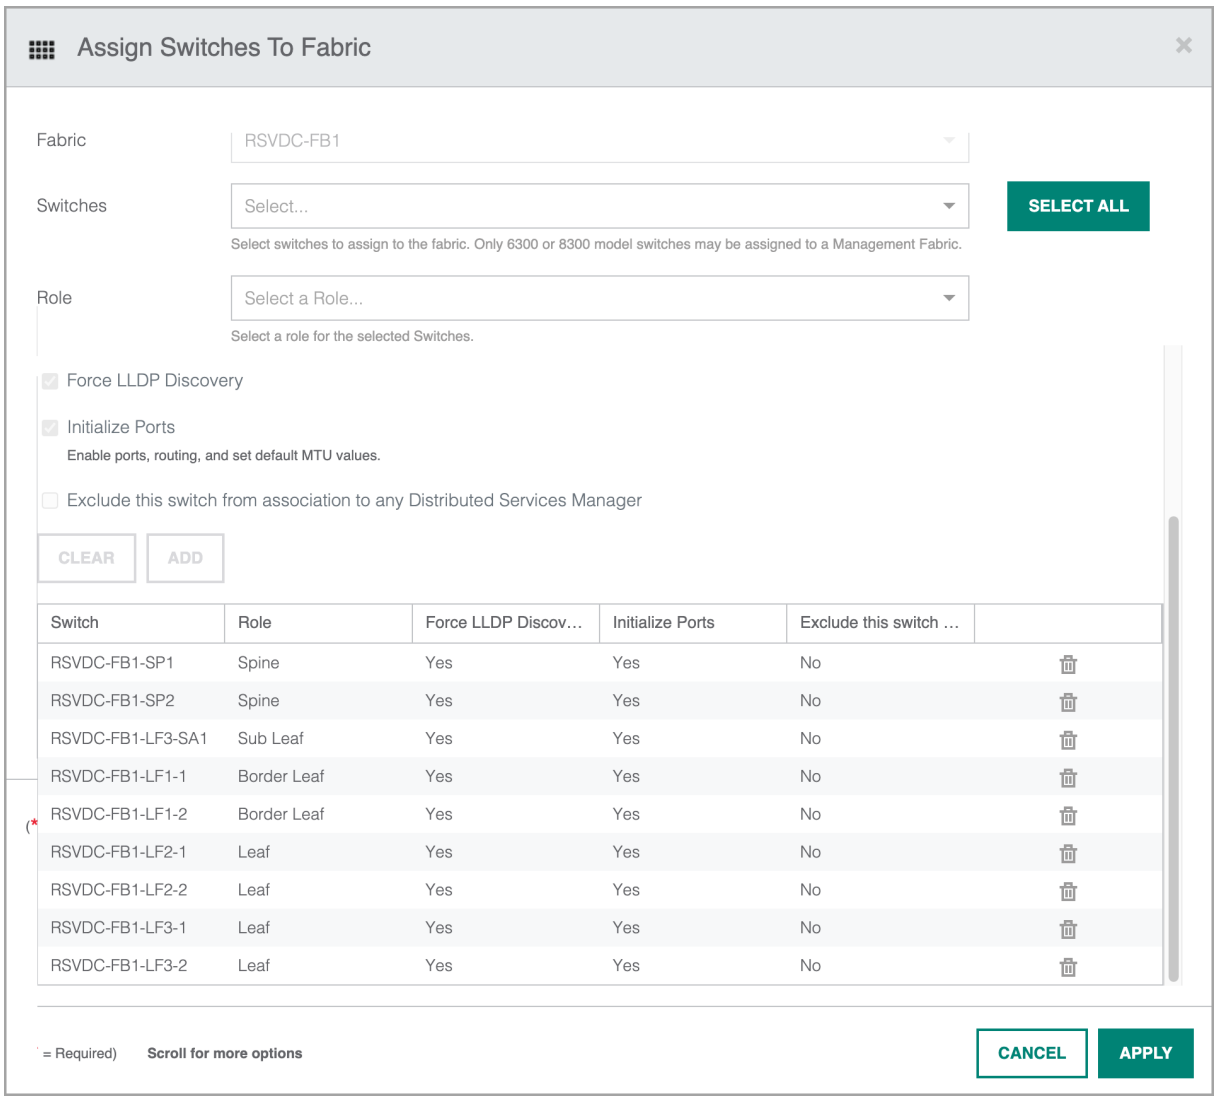


**Step 7 Guided Setup** displays the list of switches in the **Maintenance > Switches** window. Switch status should sync in a few seconds. Verify that all switches in the fabric are listed as **HEALTHY** in green.

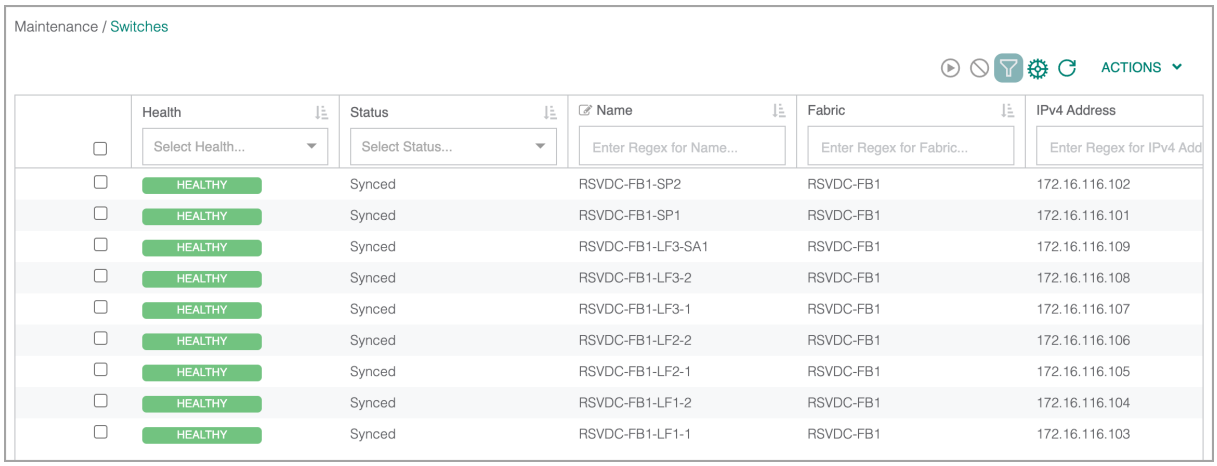


## Configure Switch Profile

The switch profile optimizes hardware resources for a switch's role in the network. Most switches are assigned a leaf role by default. The following procedure assigns the spine profile to the spine switches in the network.

**NOTE:**

When using IPsec or NAT on a CX10000 border leaf, the border leaf switches profile must be changed to **spine**. East-west policy enforcement for hosts attached to the border leaf is not supported after this change.

**Step 1** On the **Maintenance > Switches** page, click a checkbox to select one of the spine switches.

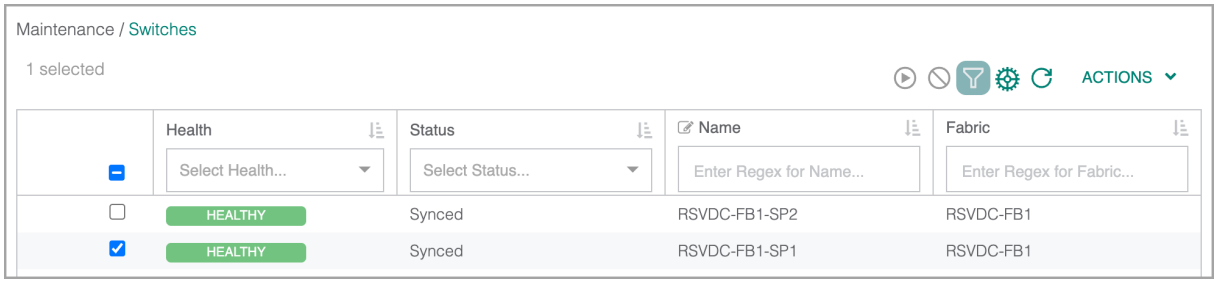


**Step 2** Click the **ACTIONS** menu on the right, and select **Change Profile**.

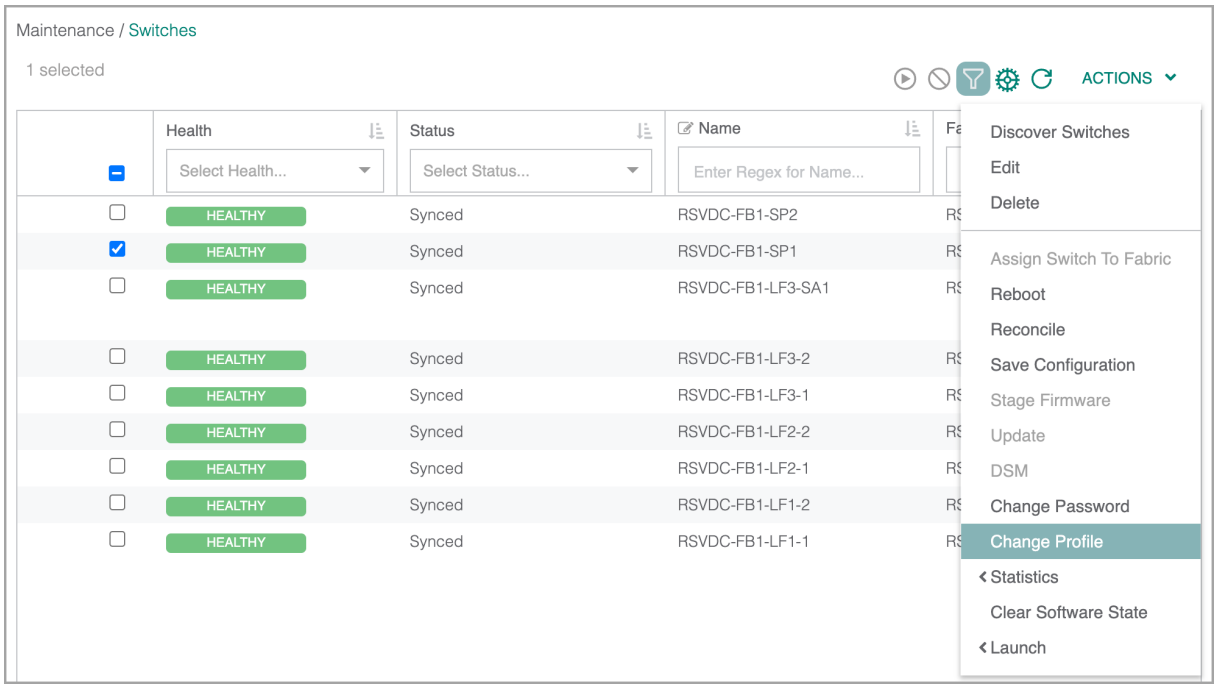


**Step 3** Select **Spine** in the **New Profile** field dropdown, check **Reboot switch after changing profile**, and click **Apply**.

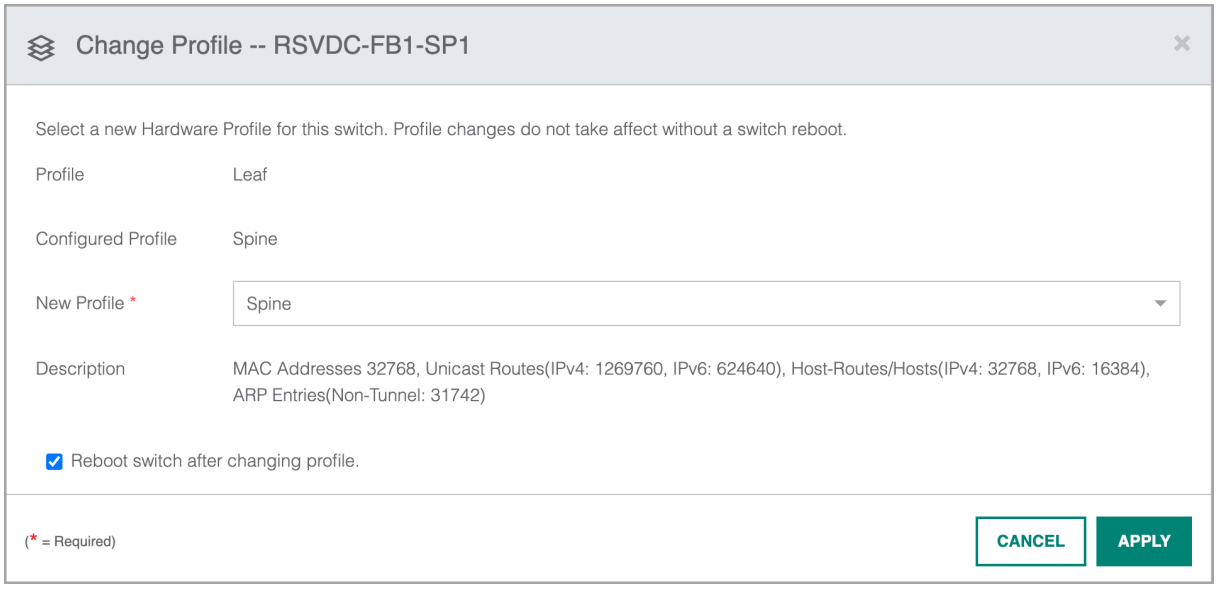


**NOTE:**

When selecting **Spine** in the **New Profile** field, the **Configured Profile** value changes dynamically from **Leaf** to **Spine**.

**Step 4** Repeat the procedure for each spine switch.

## Configure Infrastructure Split Ports

This process is necessary only when using links between fabric switches that require split port operation. The most common case is using a CX 9300 in the spine role to increase rack capacity of a fabric. In this sample deployment, CX 9300-32D spine ports are set to operate in 2 x 100 Gbps mode.

**Step 1** On the **Configuration** menu, select **Ports > Ports**.

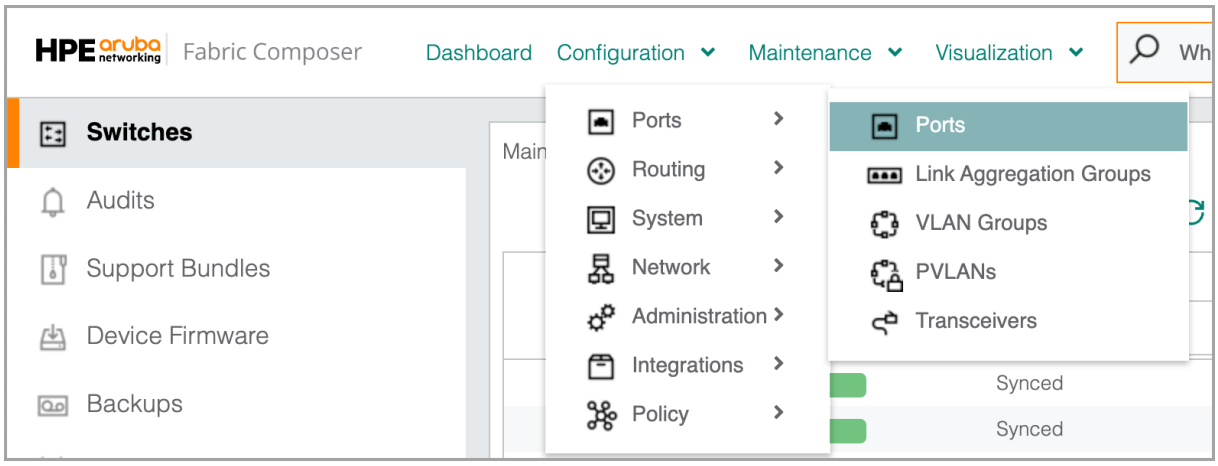


**Step 2** On the **Ports** page, select both spine switches in the **Switch** field.

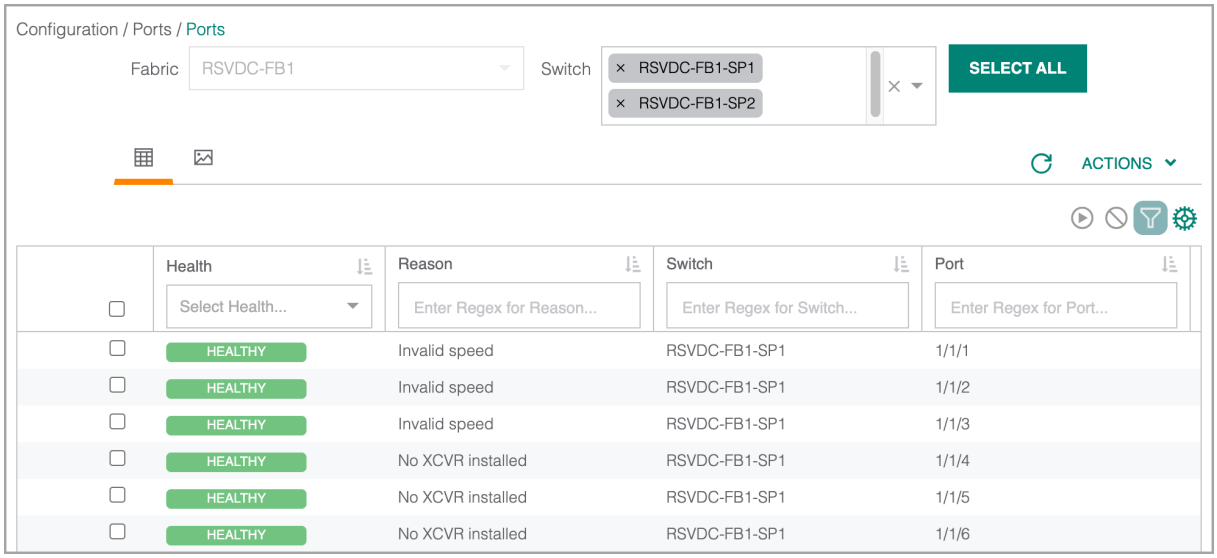


**NOTE:**

Typing a value in the **Switch** field filters selectable switch names to those containing that value in their name. Following the naming convention in this guide, only the spine switches are displayed for selection by typing **sp** in the **Switch** field.

**Step 3** Filter displayed ports by entering **Invalid** in the regex field below the **Reason** column heading and click the **Apply Table Filters** icon.

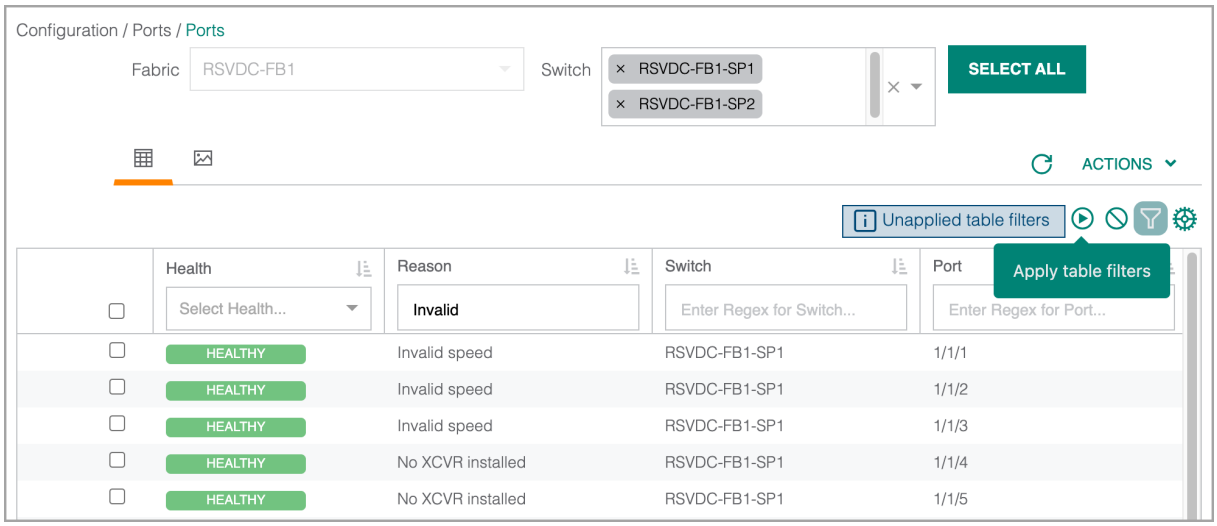


**NOTE:**

**Invalid speed** is displayed in the **Reason** column when there is a mismatch between a physical port's configured operation and an attached Active Optical Cable's (AOC's) physical split configuration. No error message is displayed when using a standard 400 Gbps transceiver before defining split port operation.

**Step 4** Click the box at the top of the selection column to select all the displayed ports on both spine switches that match the search criteria.

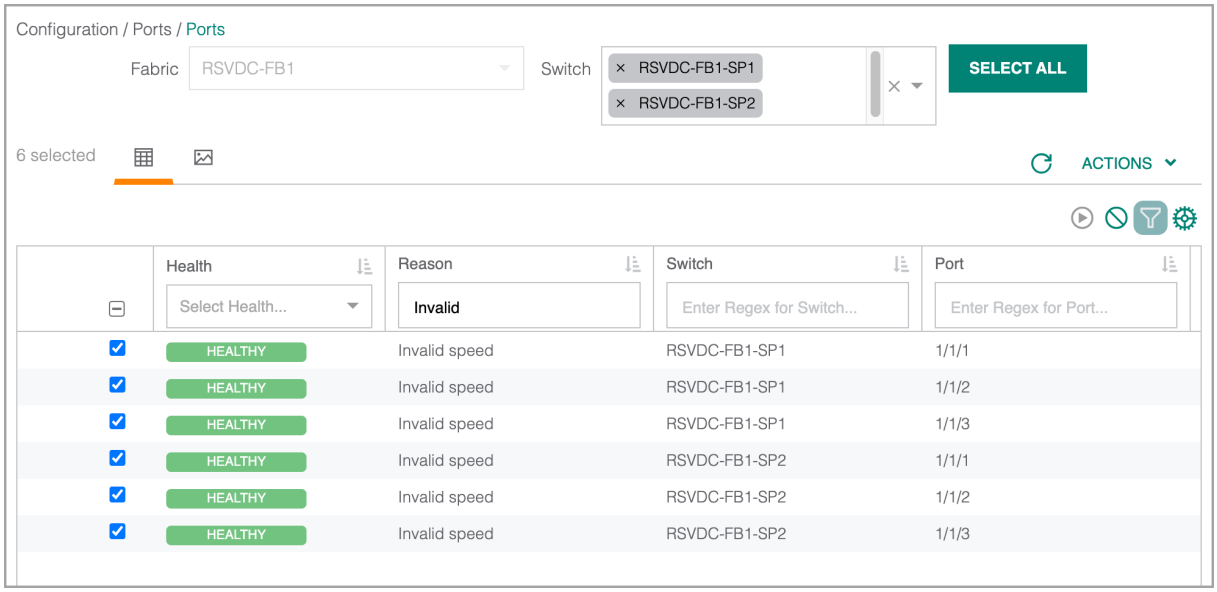


**Step 5** On the **ACTIONS** menu, select **QSFP Transform > Split > 2x 100**.

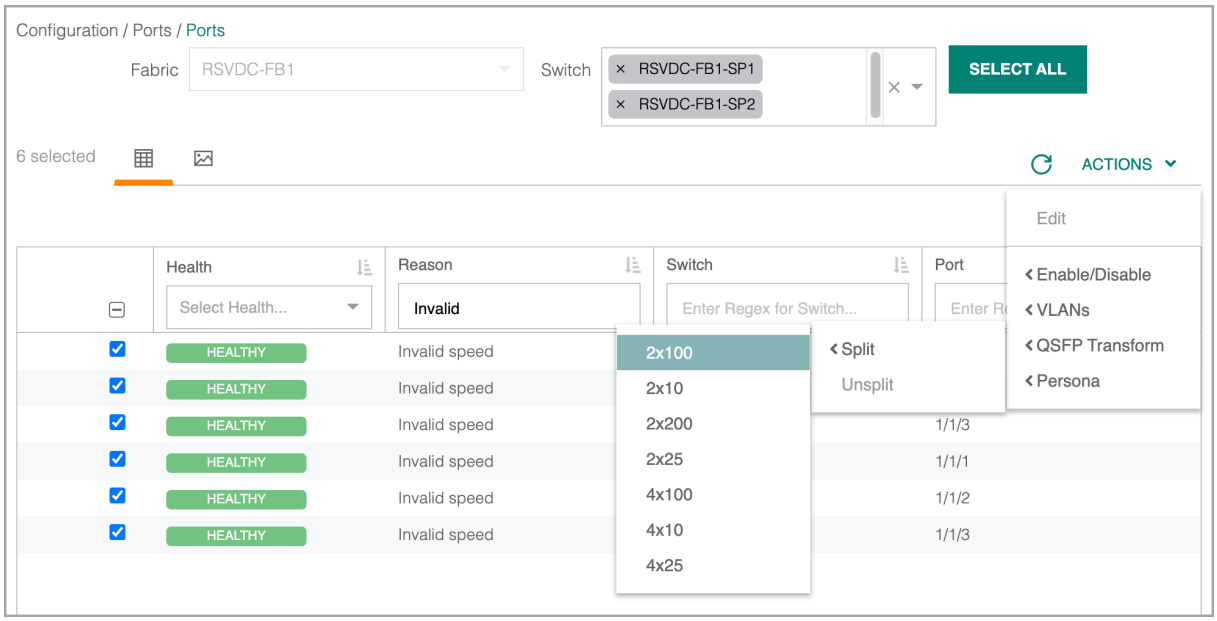


**Step 6** When prompted to confirm the split operation, click **OK**.

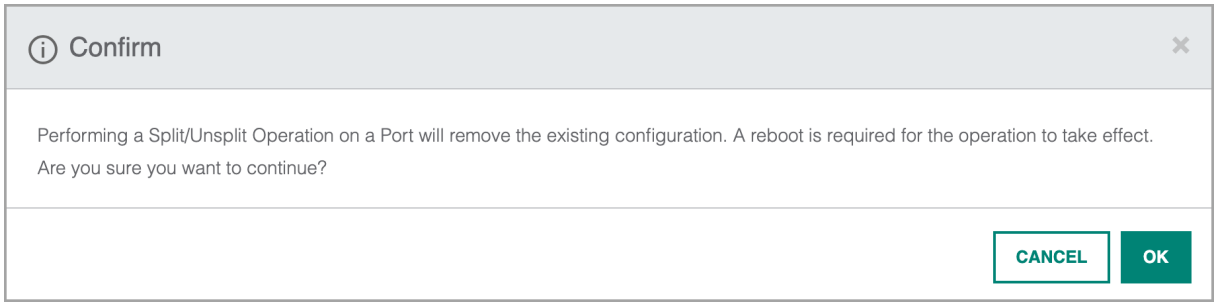


**NOTE:**

The split ports are enabled by Fabric Composer for use in LLDP neighbor discovery, and the MTU of the split ports is adjusted to 9198 to support jumbo frames for VXLAN encapsulation.The **Confirm** prompt indicates that a reboot is required, but a reboot is not required to enable split ports.

## Configure NTP for the Fabric

Modern networks require accurate, synchronized time. The NTP wizard is used to enter NTP server hosts and associate them with all fabric switches. The NTP servers must be reachable from the data center management LAN. The Fabric Composer CLI Command Processor shows the time synchronization status of each switch. At the completion of this procedure, the date and time are synchronized between the data center switches and the NTP servers.

**Step 1** On the **Guided Setup** menu, select **NTP CONFIGURATION**.

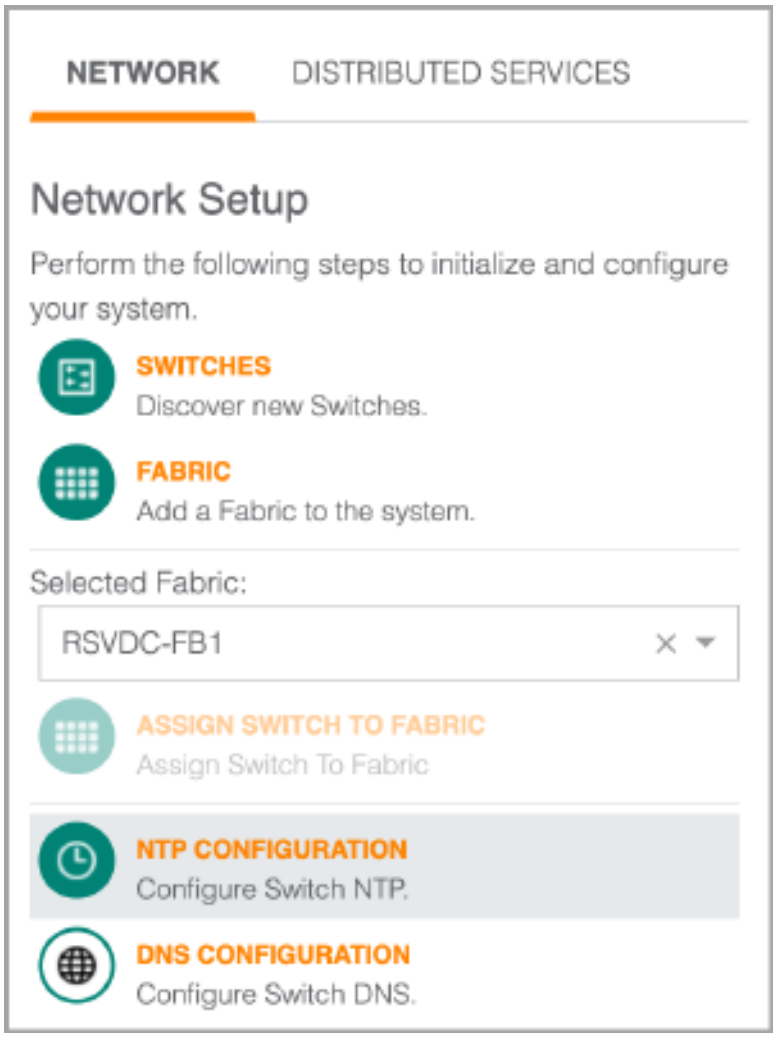


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

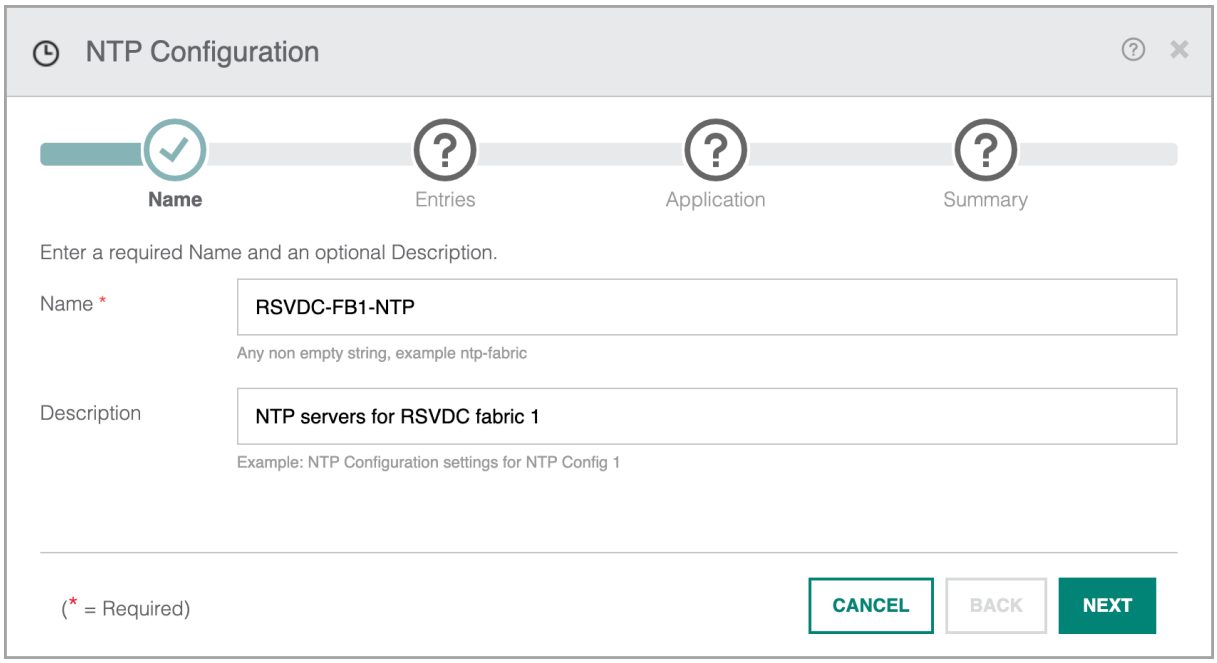


**Step 3** On the **Entries** page, enter a valid hostname or IP address and optional NTP authentication information, then click **ADD**.

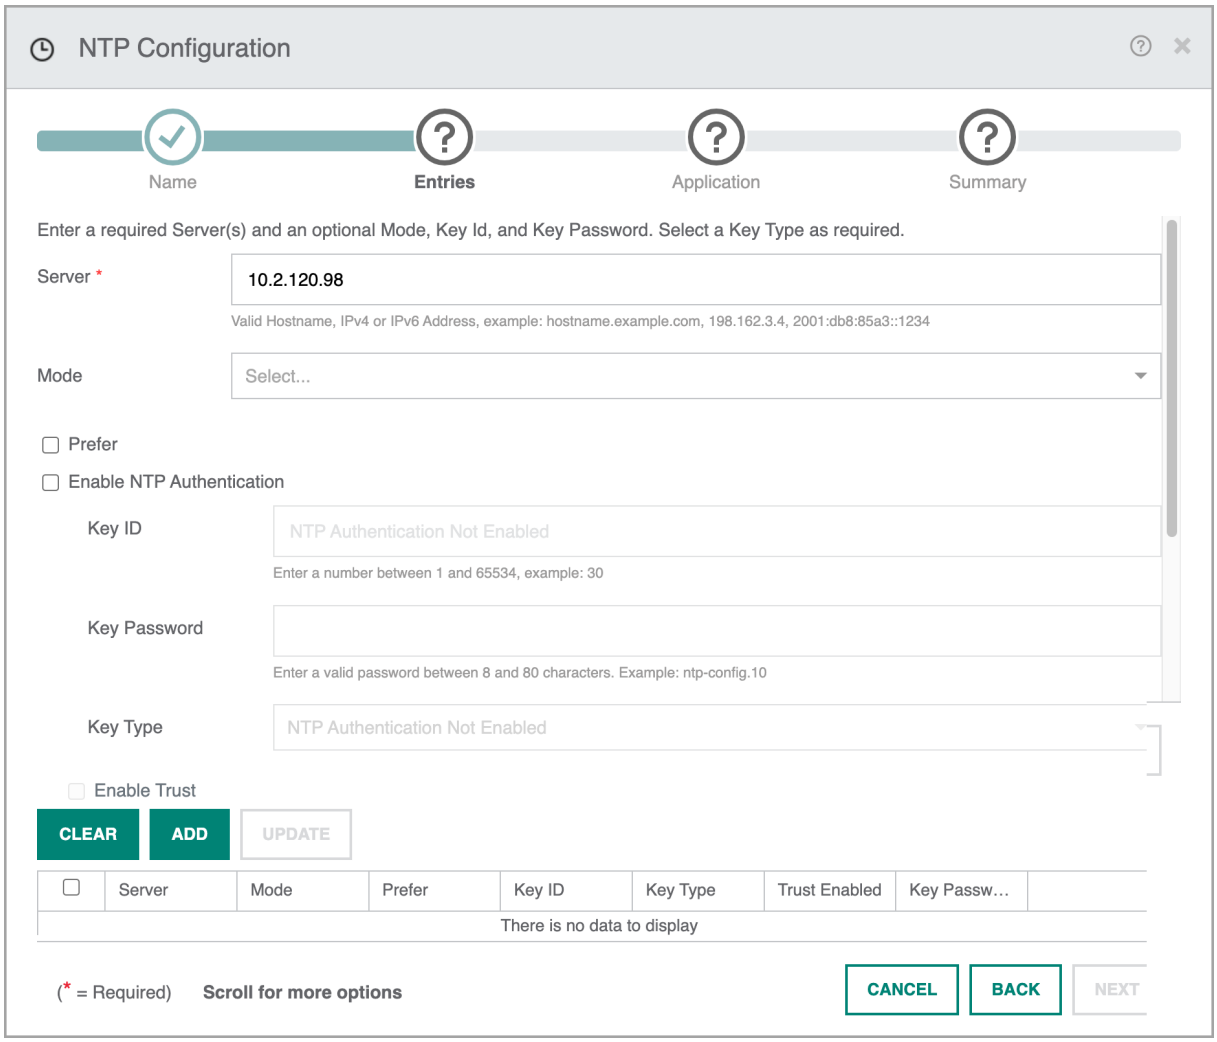


**Step 4** Repeat the step above for each NTP server in the environment.

**Step 5** After all NTP servers have been added, click **NEXT**.

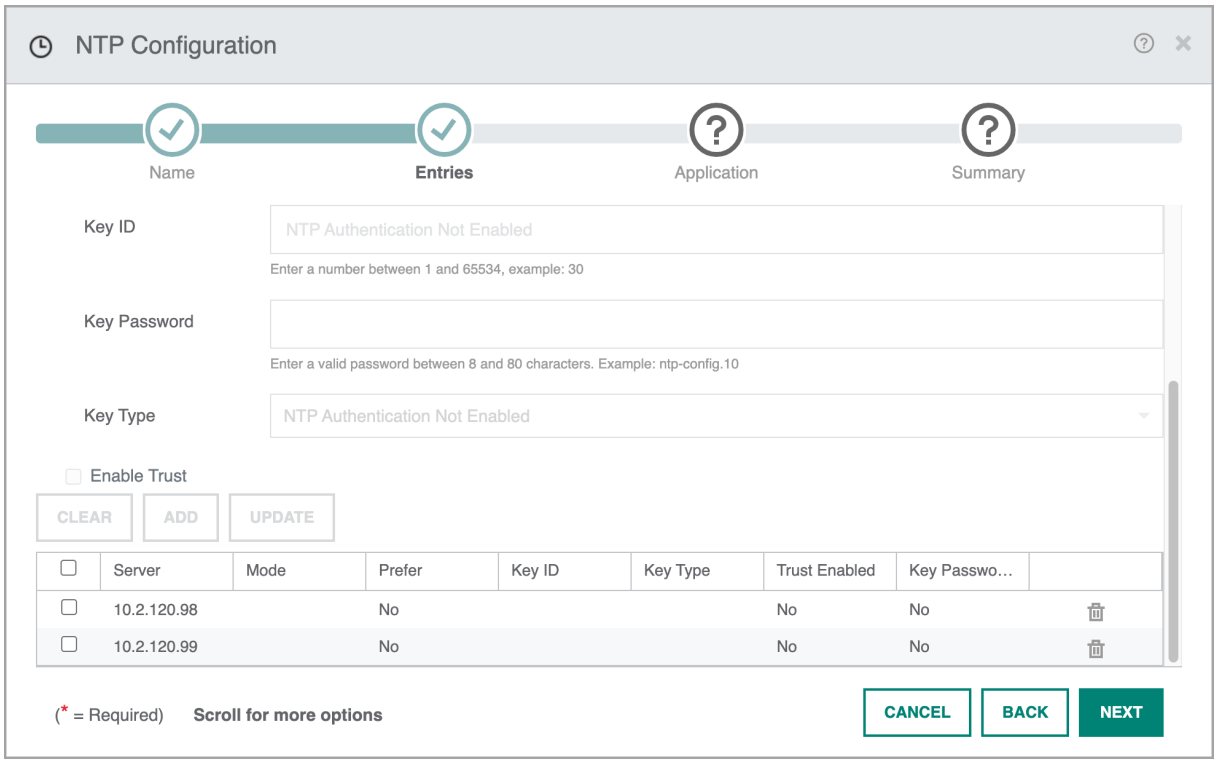


**Step 6** On the **Application** page, select the name of the fabric in the **Fabric** field and click **NEXT**.

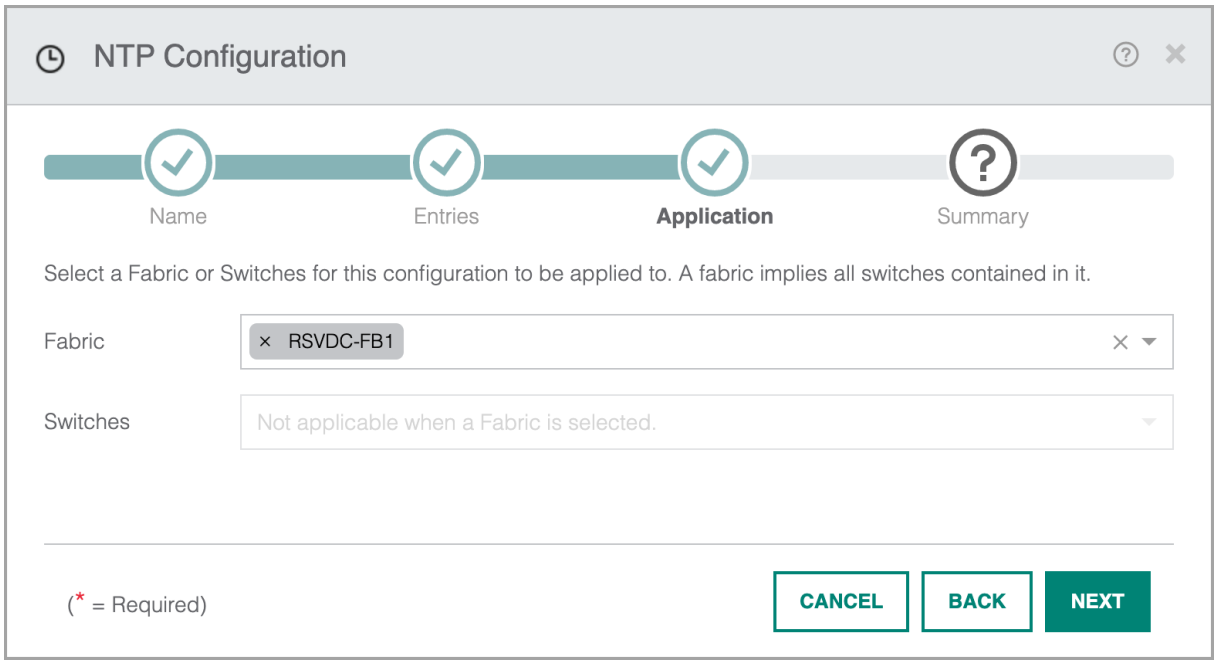


**Step 7** On the **Summary** page, verify that the information is entered correctly and click **APPLY**.

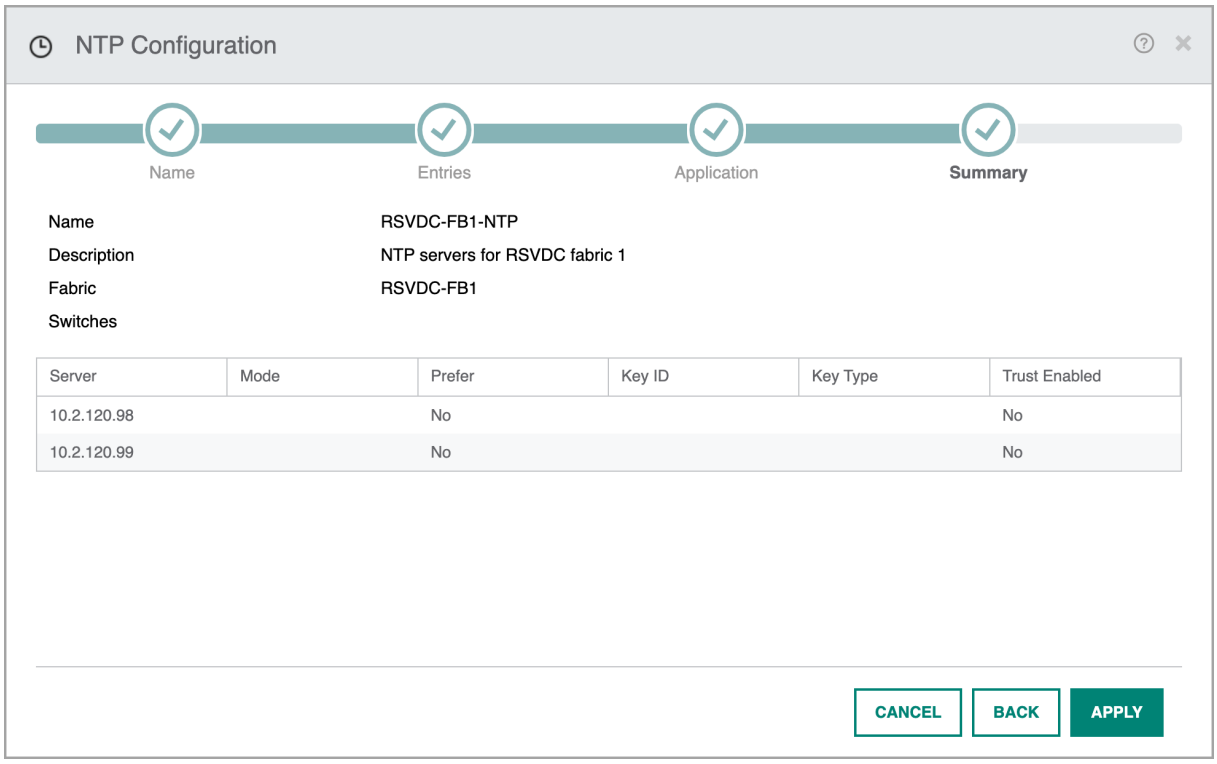


**Step 8** On the **Configuration > Network > NTP** page, click the radio button for the NTP config applied to an individual switch, click the **ACTIONS** menu on the right, and click **Delete**.

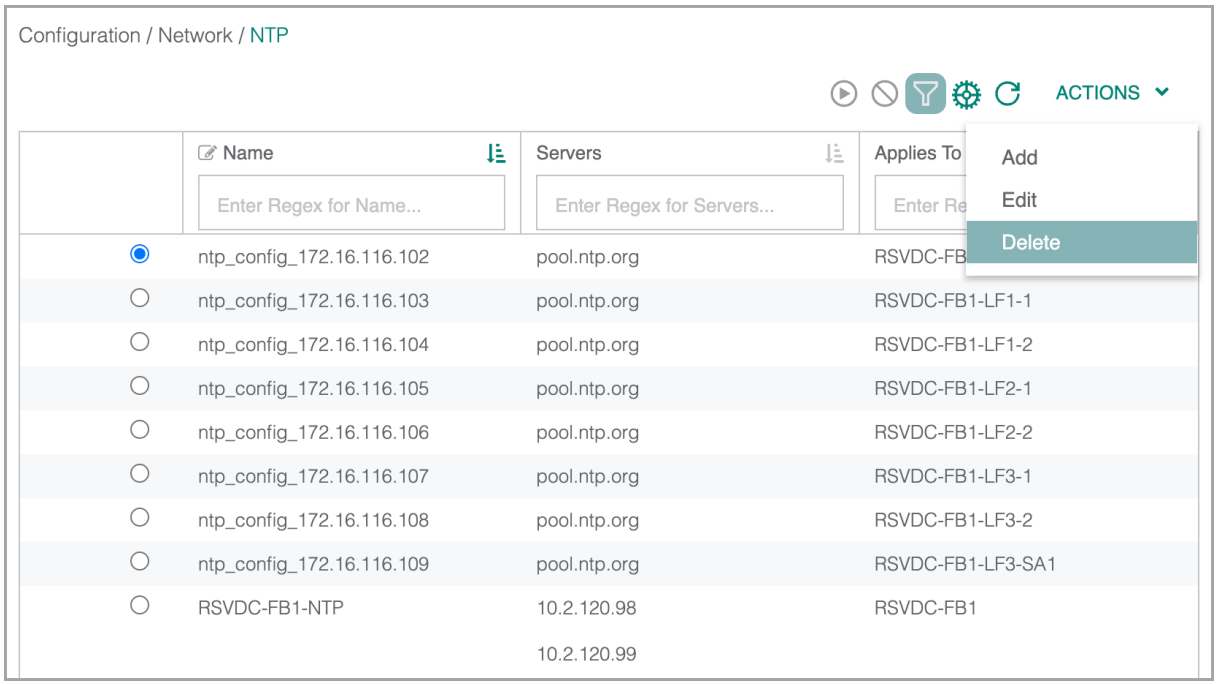

**NOTE:**

Fabric Composer dynamically creates switch level objects that reconcile configuration performed by an administrator directly on the switch. A switch level configuration object has a higher precedence than Fabric Composer objects defined at the fabric level. At this time, the default NTP config is reconciled in a switch level configuration object. In this case, it is necessary to delete switch level NTP configuration objects to apply the fabric level config. If per-switch reconciled config is not present, omit steps 8, 9, and 10.

**Step 9** In the **Delete** confirmation window, click **OK**.

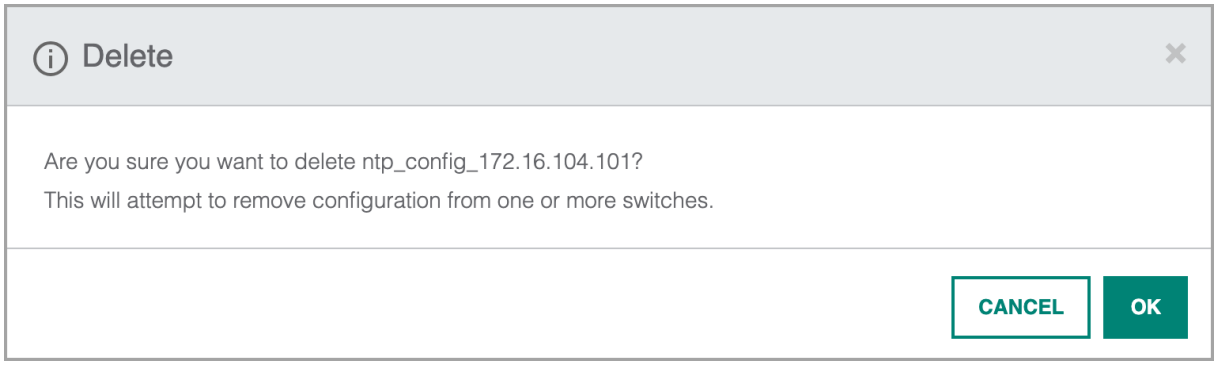


**Step 10** Repeat steps 8 and 9 to remove reconciled NTP configuration for all switches.

**Step 11** In the menu bar at the top right of the Fabric Composer display, click the **CLI Commands** icon and select **Show Commands**.

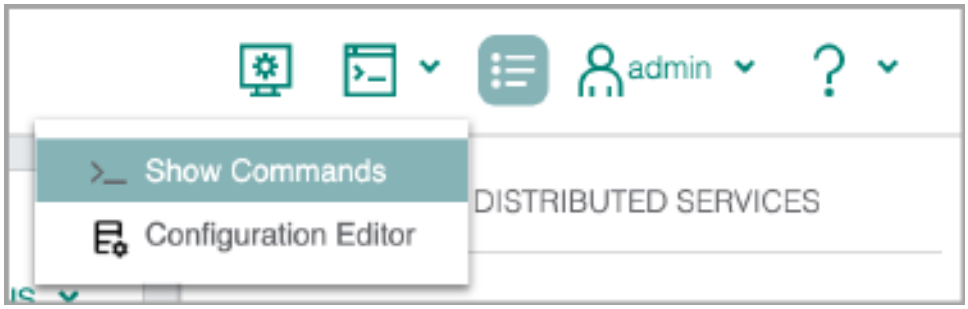


**Step 12** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Fabrics:** *RSVDC-FB1*
- **Commands:** *show ntp status*

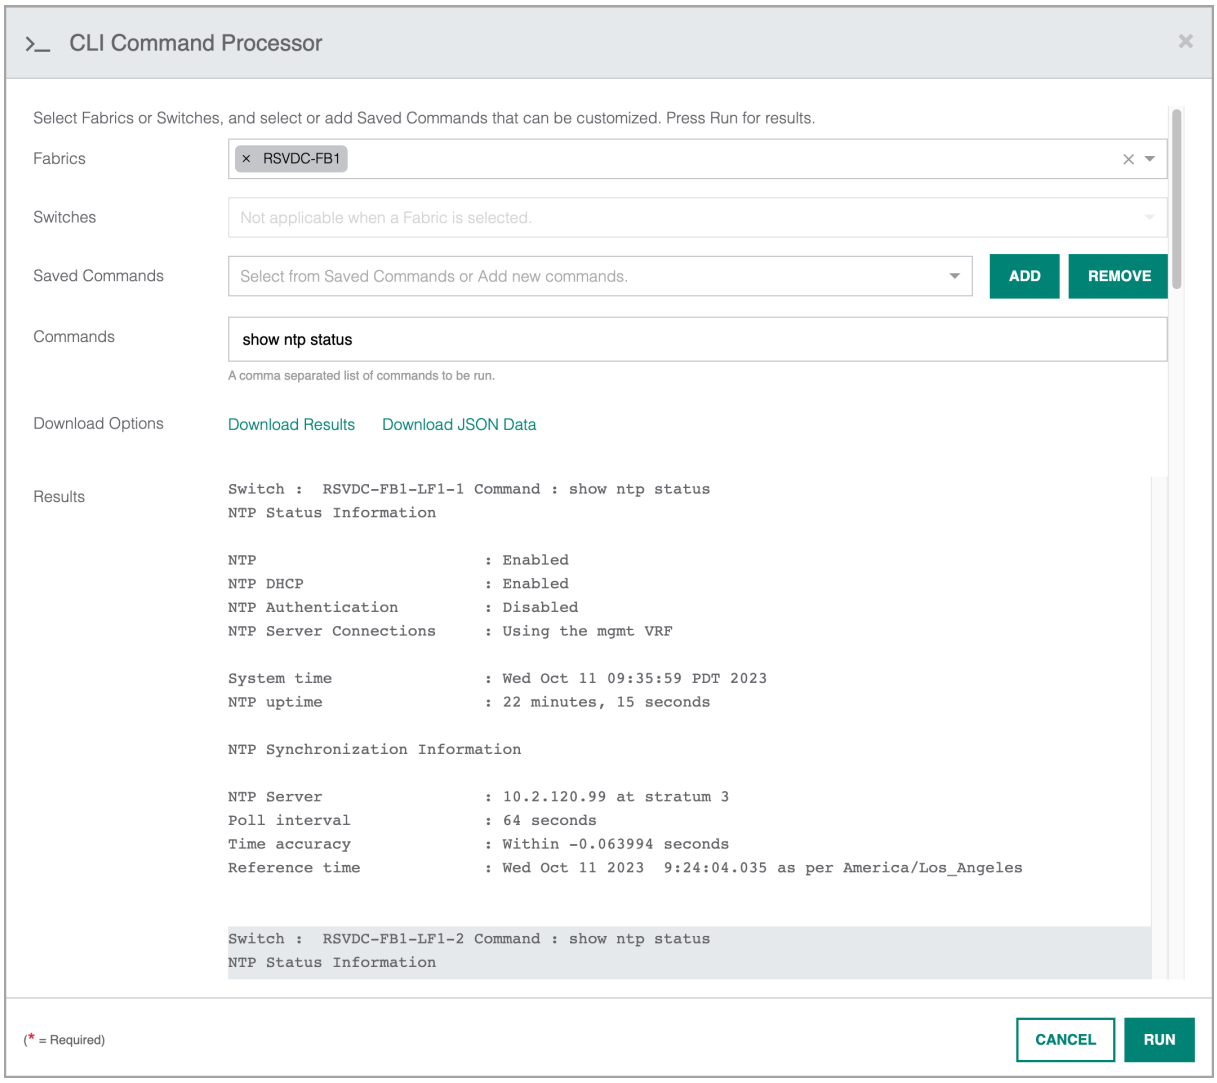


:::info Multiple commands are supported in the **Commands** field in a comma-separated list. CLI commands can be saved for future reuse by clicking the **ADD** button. When typing a command in the **Saved Commands** field, preconfigured and saved commands appear in a list. Select a command in the list to add it to the **Commands** field.

:::

**Step 13** Verify that the output for each switch displays an NTP server IP address with stratum level, poll interval, and time accuracy information.

**NOTE:**

NTP synchronization can take several minutes to complete. If a hostname was used instead of an IP address, complete the next step to configure DNS for the fabric before NTP verification.

## Configure DNS for the Fabric

Use the DNS wizard to enter DNS host details and associate them with all fabric switches. The DNS servers must be reachable from the data center management LAN.

At the completion of this procedure, the data center switches can resolve DNS hostnames to IP addresses.

**Step 1** On the **Guided Setup** menu, select **DNS CONFIGURATION**.

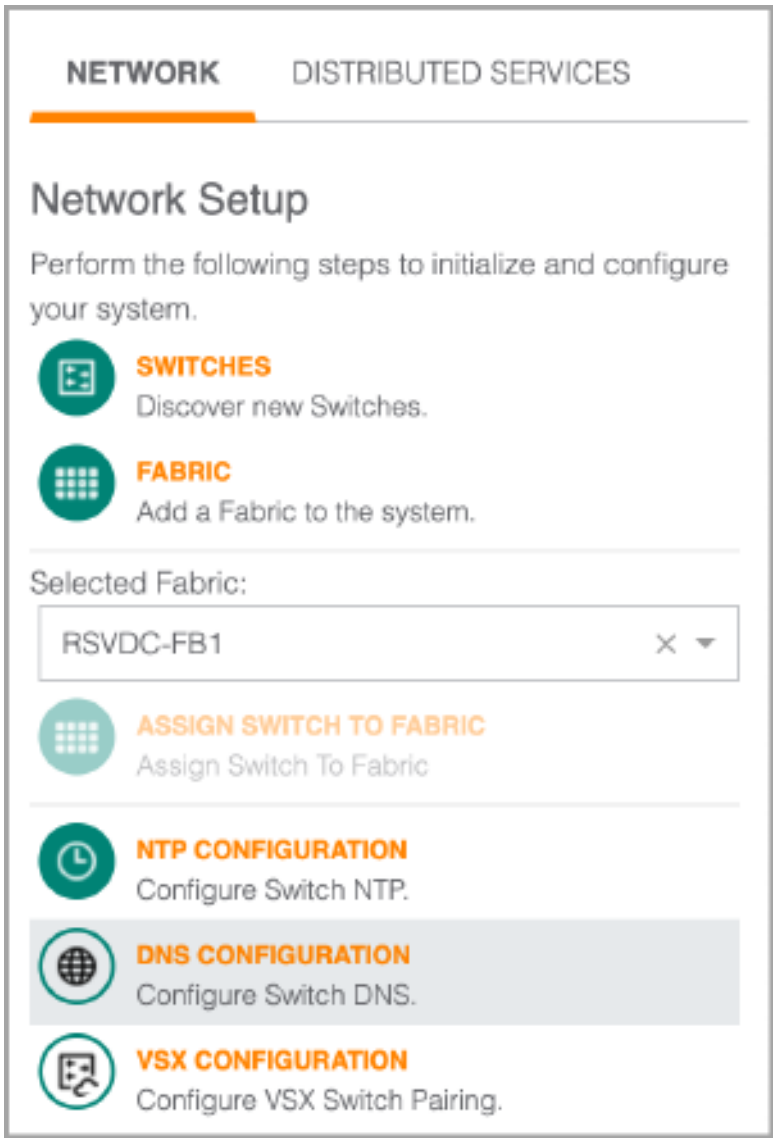


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

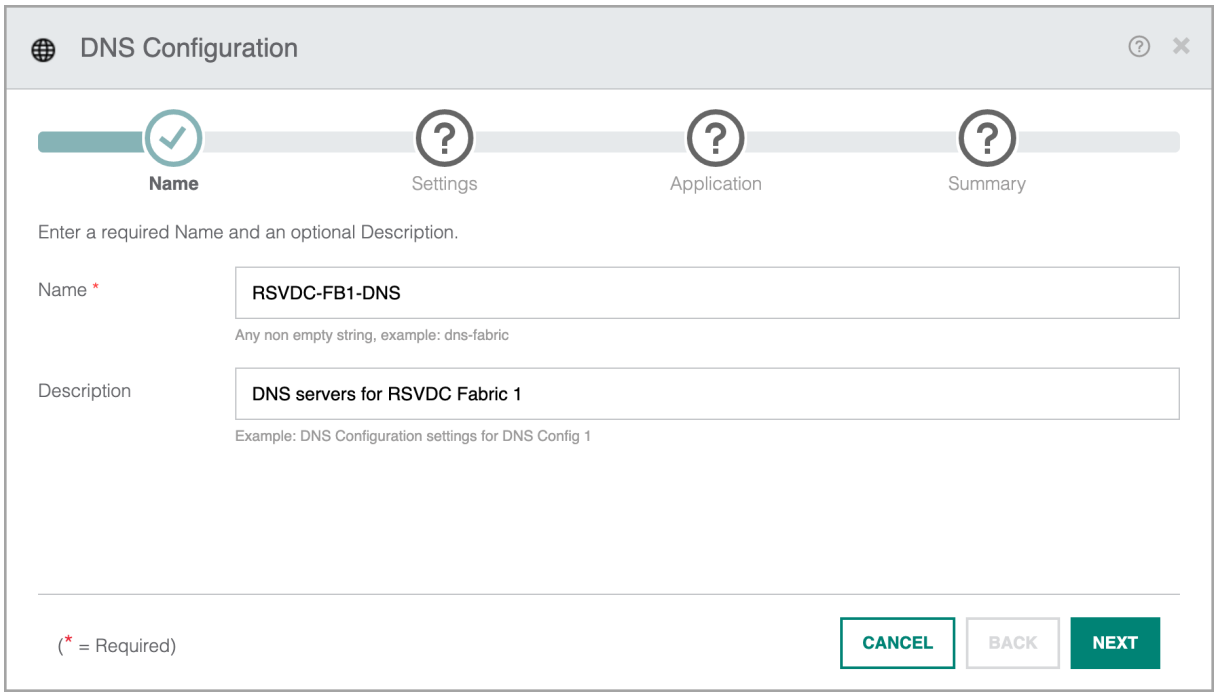


**Step 3** On the **Settings** page, enter the **Domain Name**. Enter a valid DNS server IP address in the **Name Servers** field. Press the TAB or ENTER key to complete the server entry.

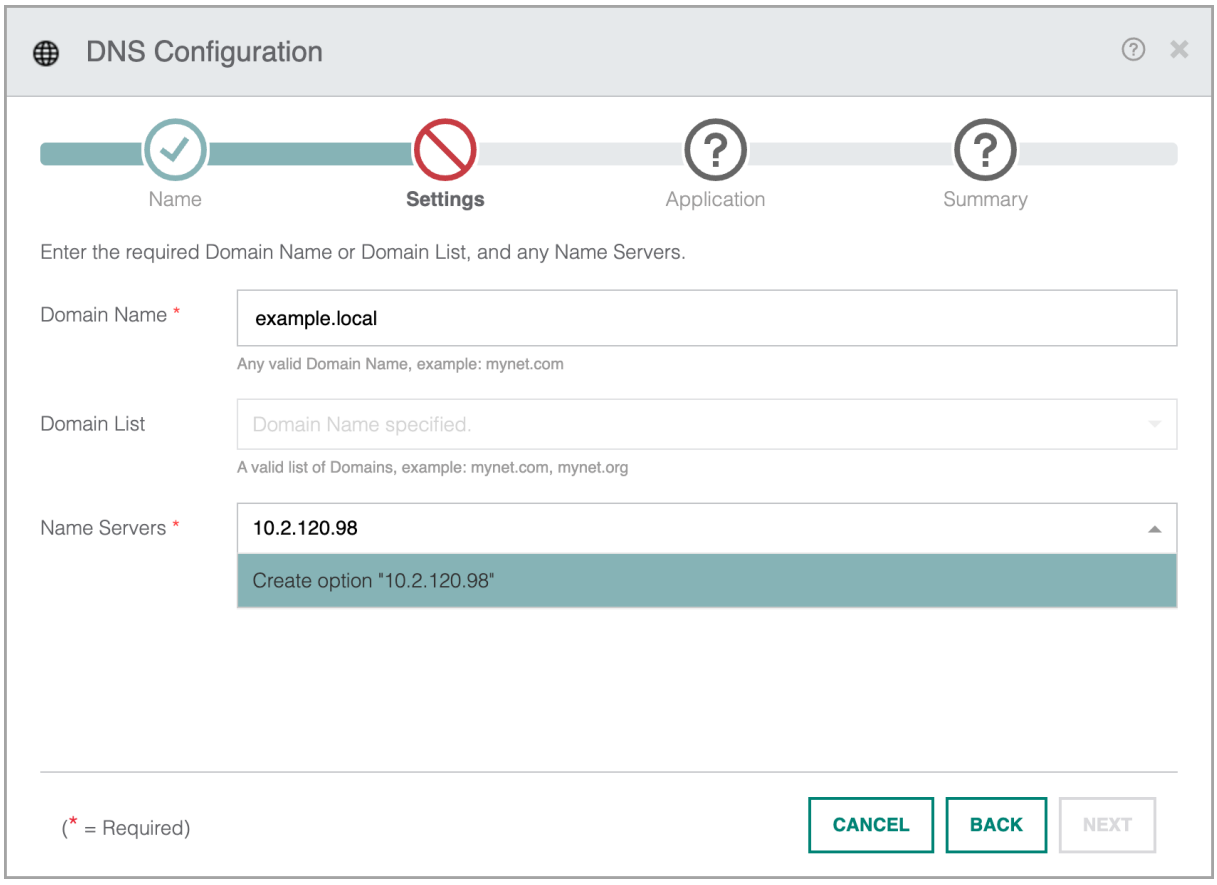


**Step 4** Create additional entries as needed. After all required DNS servers are entered, click **NEXT**.

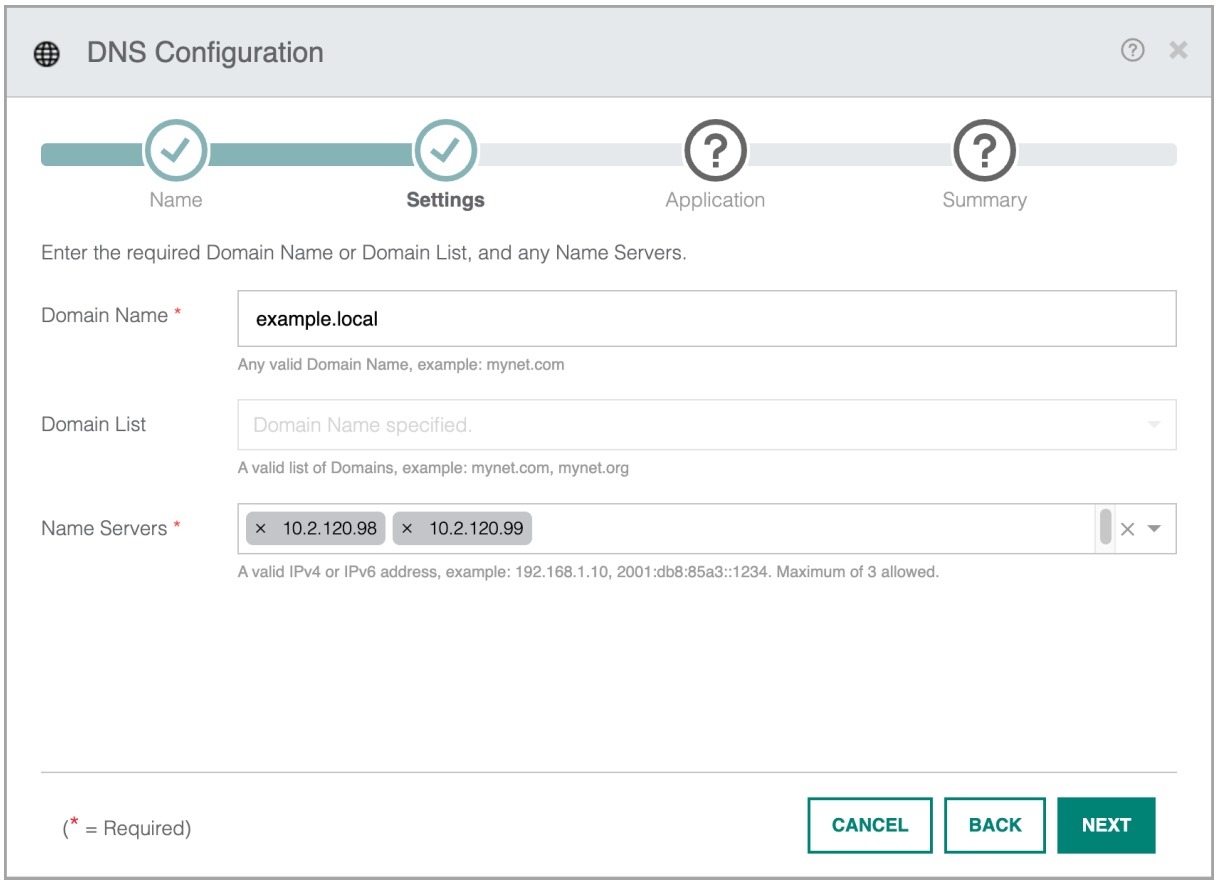


**Step 5** On the **Application** page, select the name of the fabric in the **Fabrics** field and click **NEXT**.

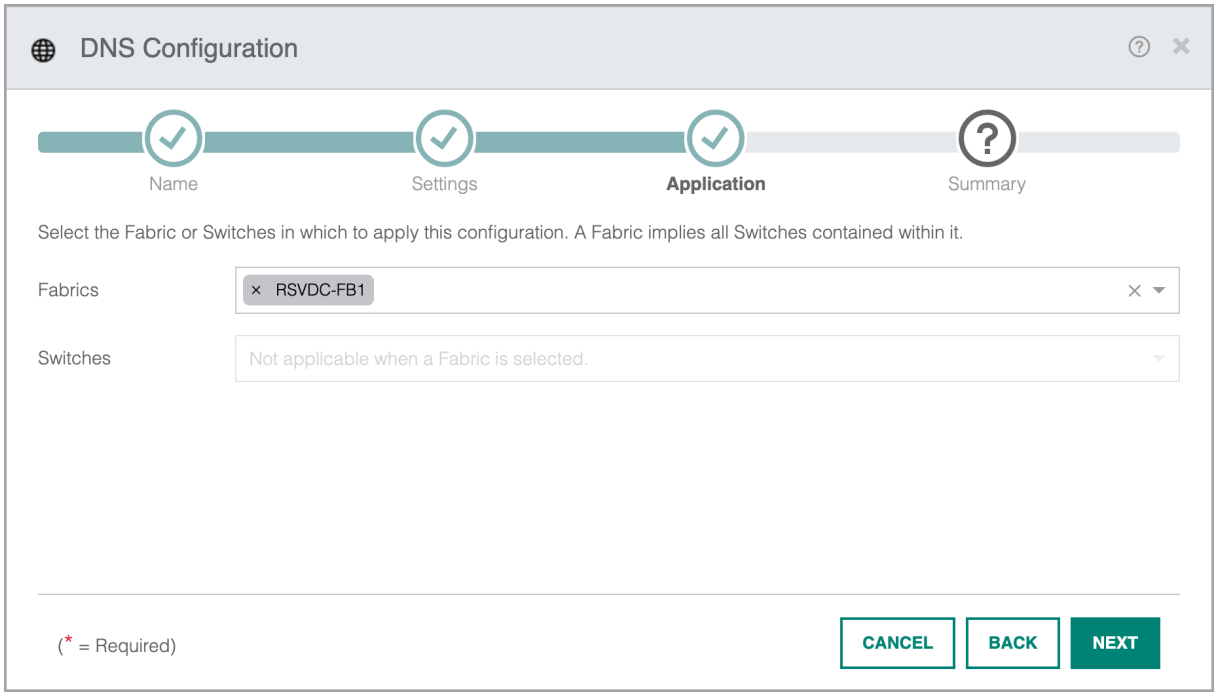


**Step 6** On the **Summary** page, verify that the information is entered correctly and click **APPLY**.

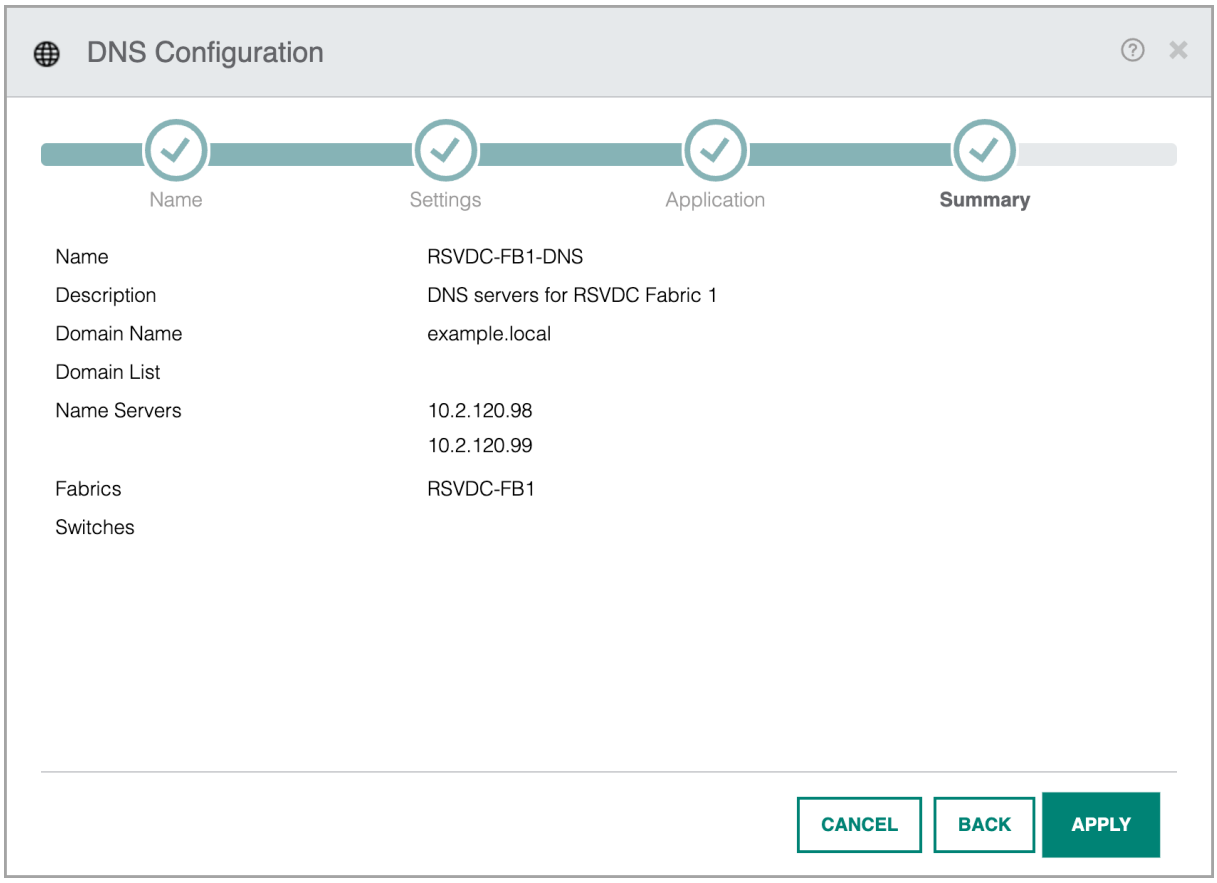


## Configure VSX on Leaf Switches

VSX enables a pair of ToR leaf switches to appear as a single logical switch to downstream hosts using multi-chassis link aggregation. VSX improves host availability in case of switch failure or maintenance downtime. Fabric Composer automatically identifies VSX switch pairs and configures them with the values supplied in the VSX wizard. Resource Pool wizards create IP and MAC address objects. The Fabric Composer CLI Command Processor verifies VSX operational status.

The diagram below highlights leaf and border leaf VSX pairs created in this procedure.

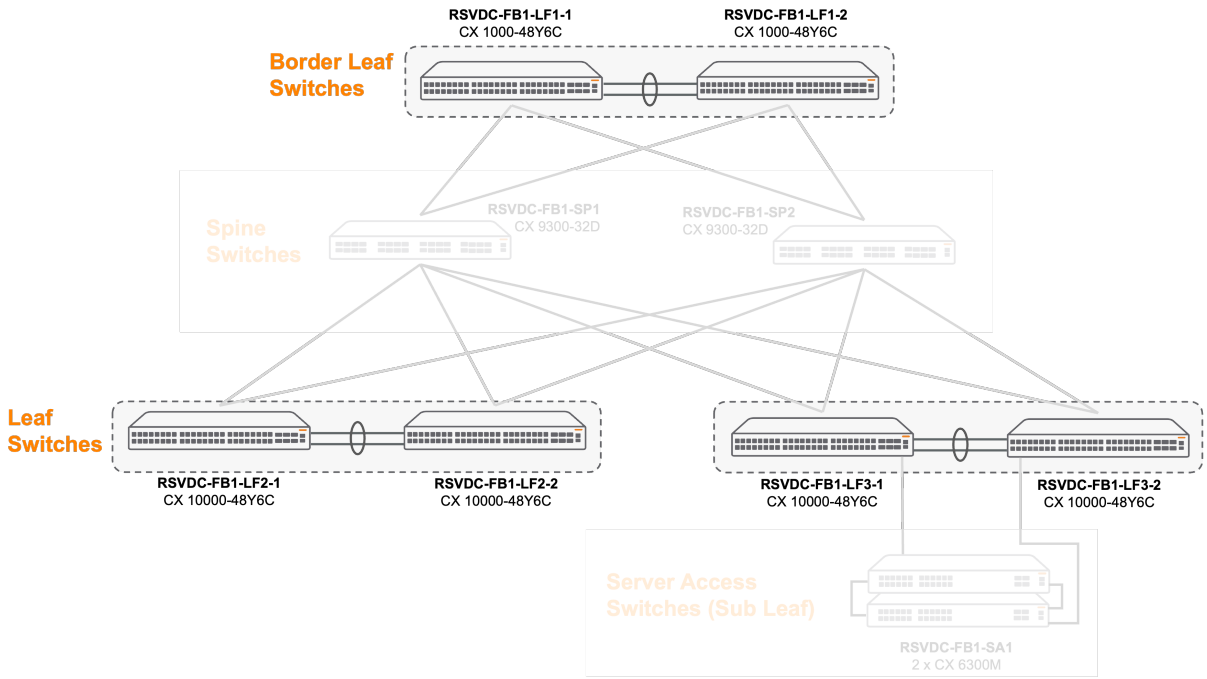


**NOTE:**

Use of a non-uplink port for keep-alive messages between VSX peers is recommended to maximize fabric capacity.

**Step 1** On the **Guided Setup** menu, select **VSX CONFIGURATION**.

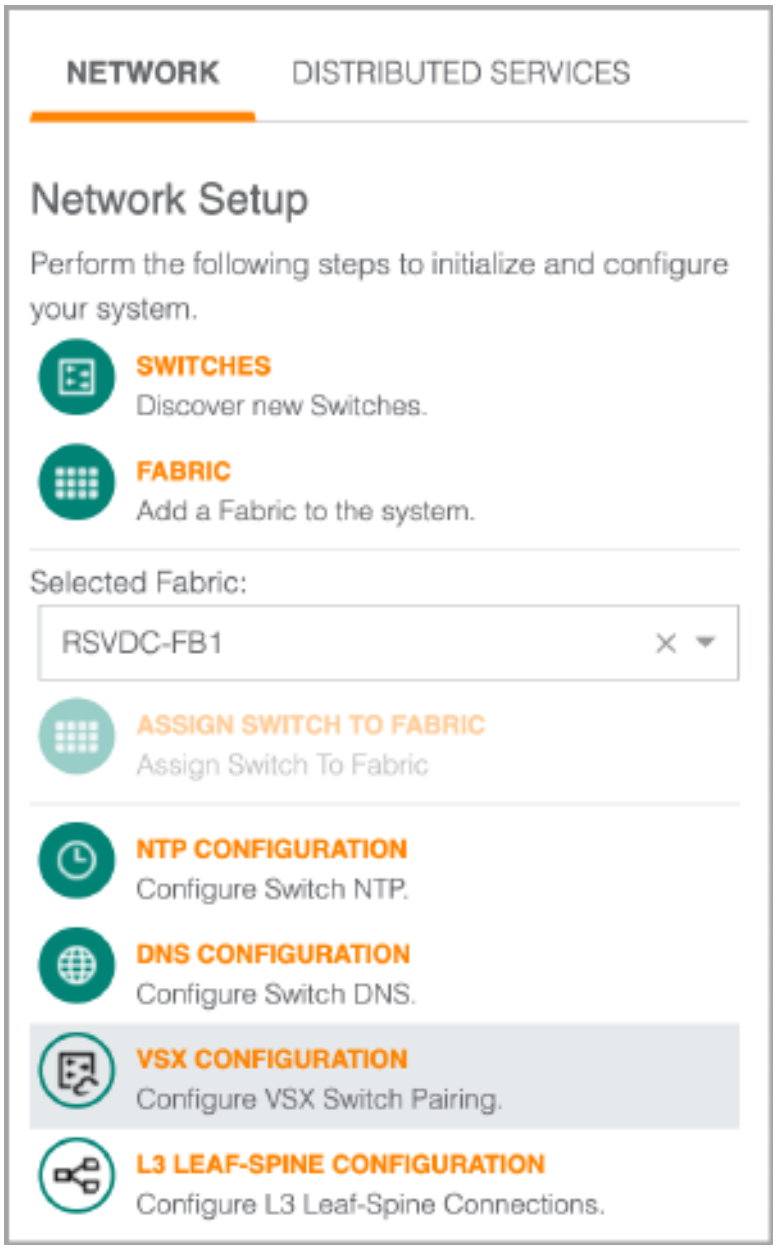


**Step 2** On the **Create Mode** page, leave **Automatically generate VSX Pairs** selected and click **NEXT**.

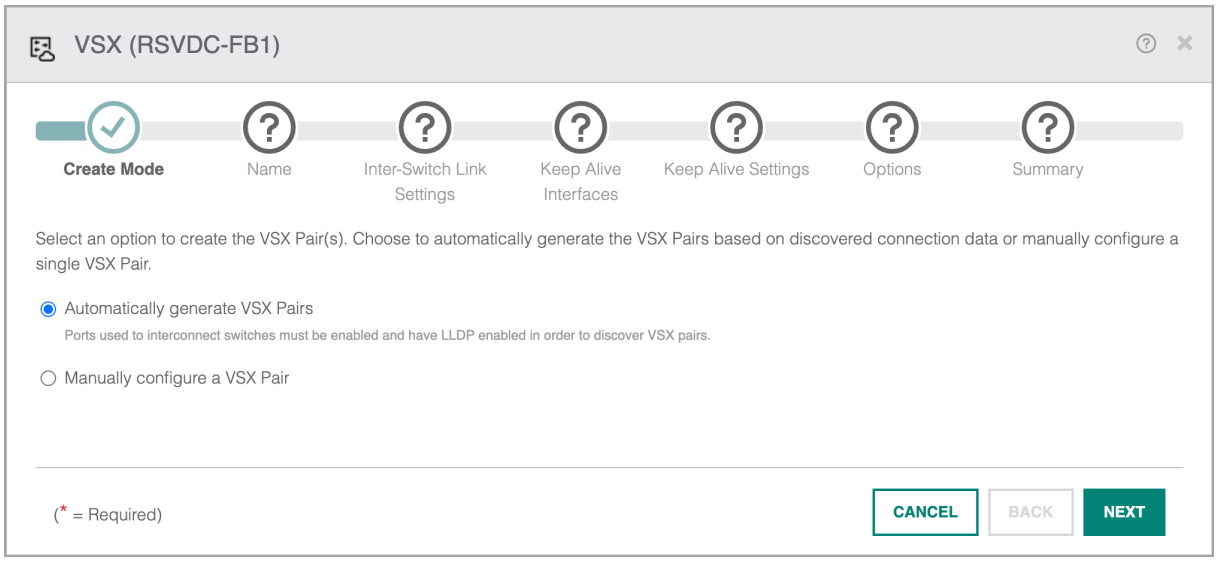


**Step 3** On the **Name** page, enter a **Name Prefix** and **Description**, then click **NEXT**.

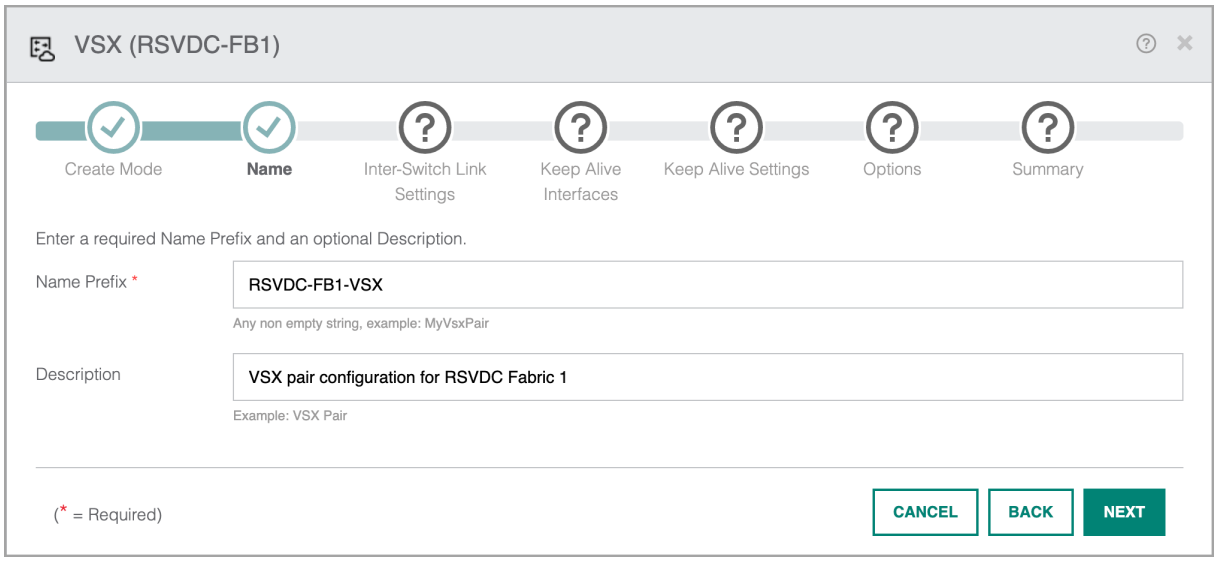


**Step 4** On the **Inter-Switch Link Settings** page, leave the default values and click **NEXT**.

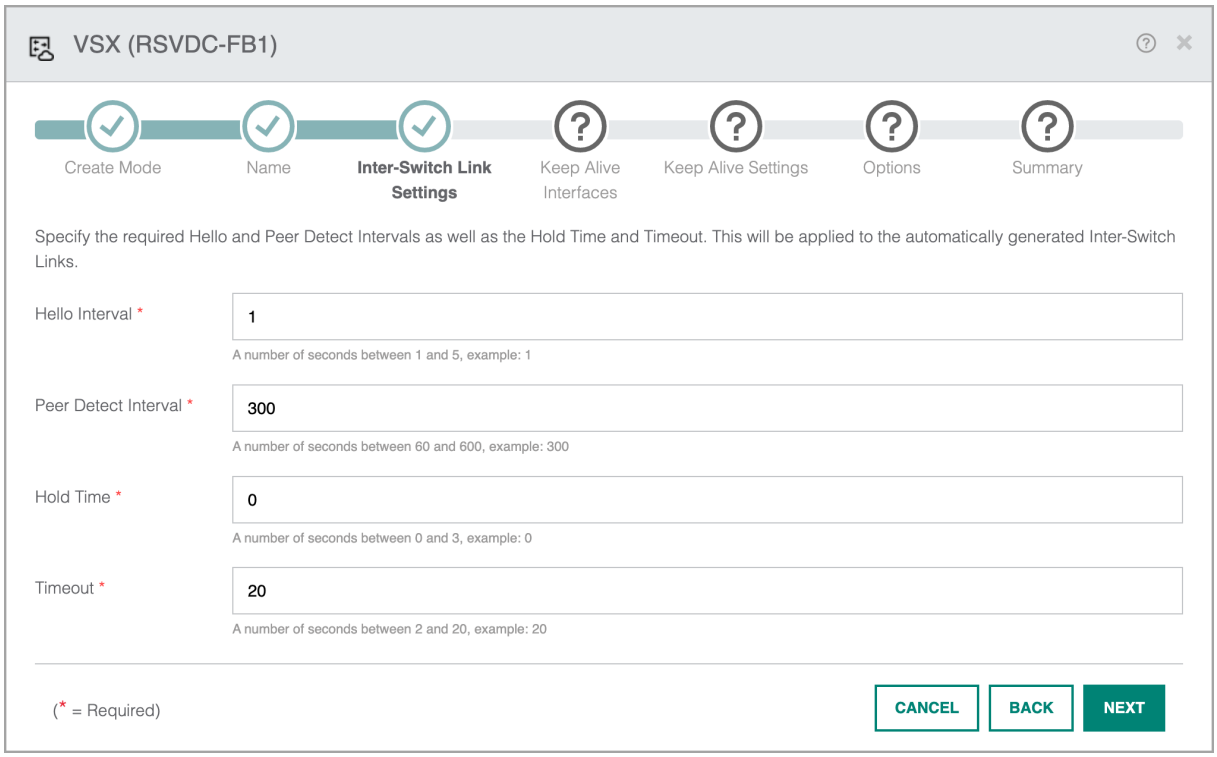


**Step 5** On the **Keep Alive Interfaces** page, select *Point-to-Point* as the **Interface Mode**. Click **ADD** to launch the **Resource Pool** wizard.

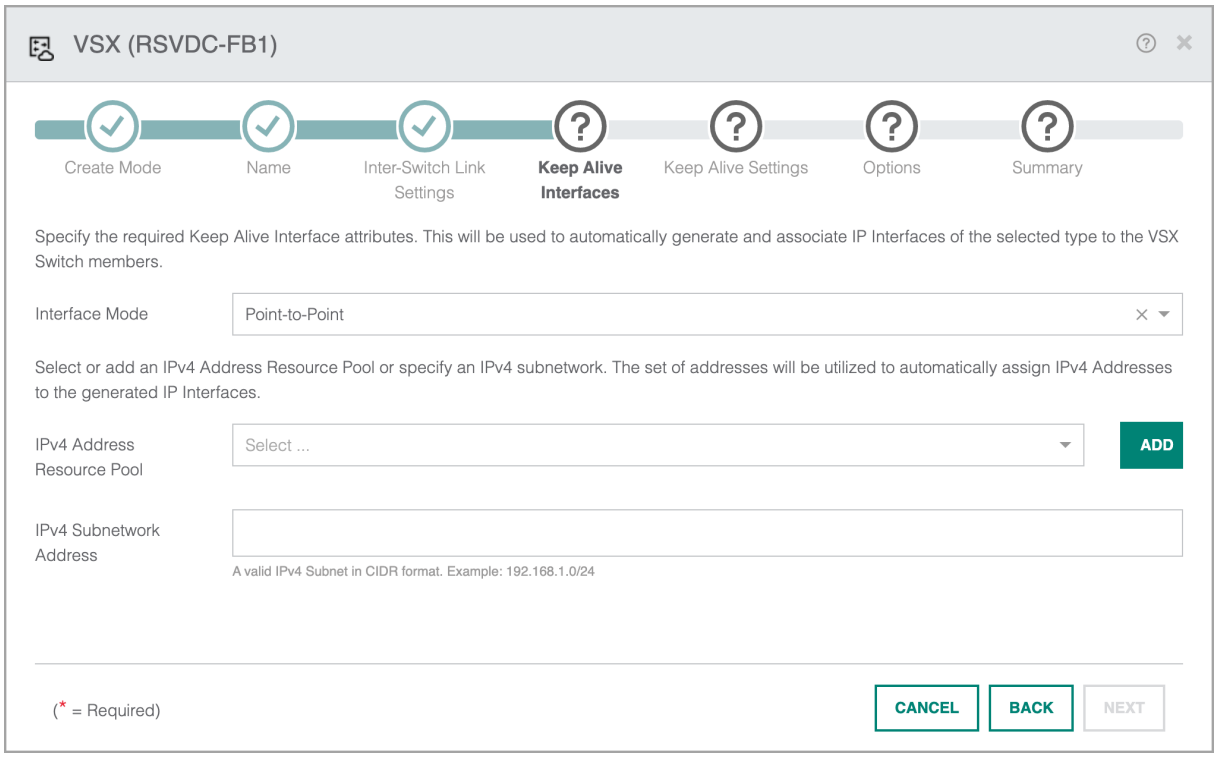

**NOTE:**

The **Resource Pool** wizard is launched in this step to create an object representing the IPv4 address range used for underlay loopback interfaces on all switches, VSX keep-alive interfaces, and routed transit VLAN interfaces on VSX pairs. A resource pool is a reusable object that ensures consistency and reduces errors when adding switches to the fabric in the future.

**Step 6 Resource Pool** wizard: On the **Name** page, enter a **Name** and **Description** for the IPv4 address pool, then click **NEXT**.

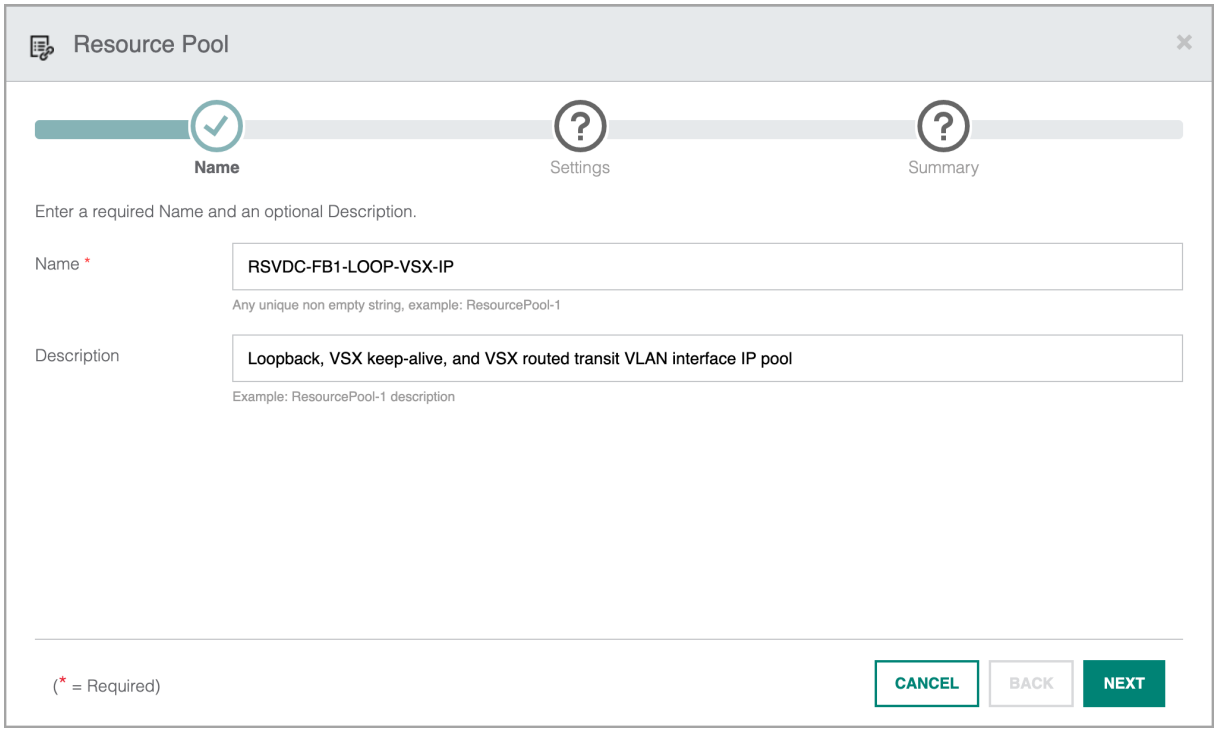


**Step 7 Resource Pool** wizard: On the **Settings** page, enter an IPv4 address block in the **Resource Pool** field and click **NEXT**.

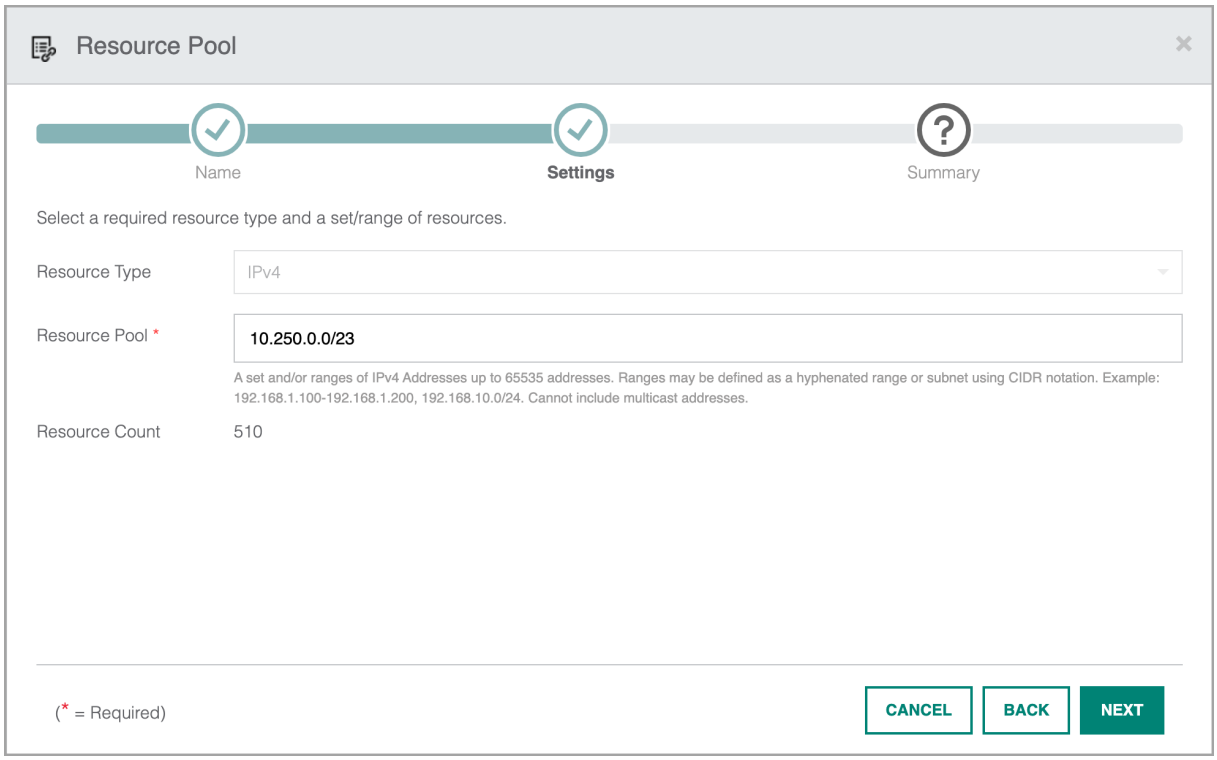


**NOTE:**

This IPv4 address block is used to allocate IP addresses to loopback interfaces (/32) for all fabric switches, VSX keep-alive point-to-point interfaces (/31), and routed transit VLAN interfaces on VSX pairs (/31). Use a block large enough to support addressing these interfaces across the entire fabric.

**Step 8 Resource Pool** wizard: On the **Summary** page, verify the IP address pool information and click **APPLY**. The **Resource Pool** wizard closes and returns to the main **VSX Configuration** workflow.

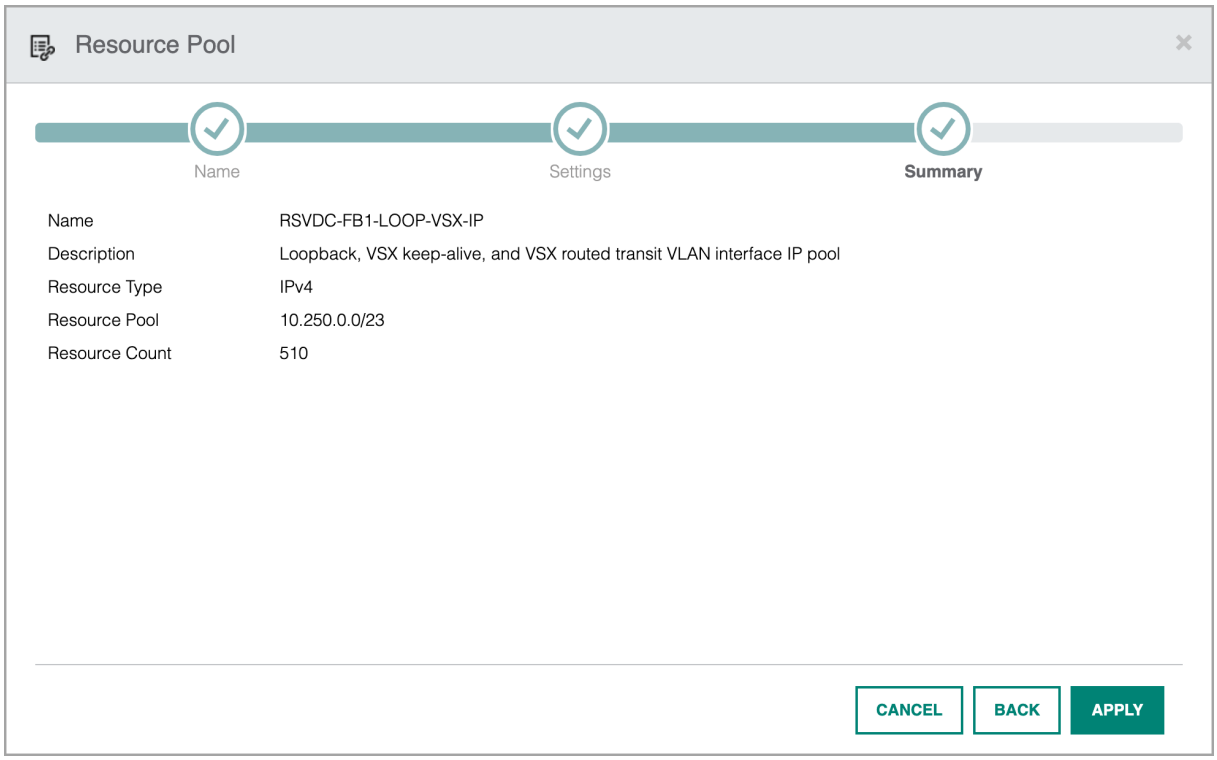


**Step 9** On the **Keep Alive Interfaces** page, verify that the new **IPv4 Address Resource Pool** is selected and click **NEXT**.

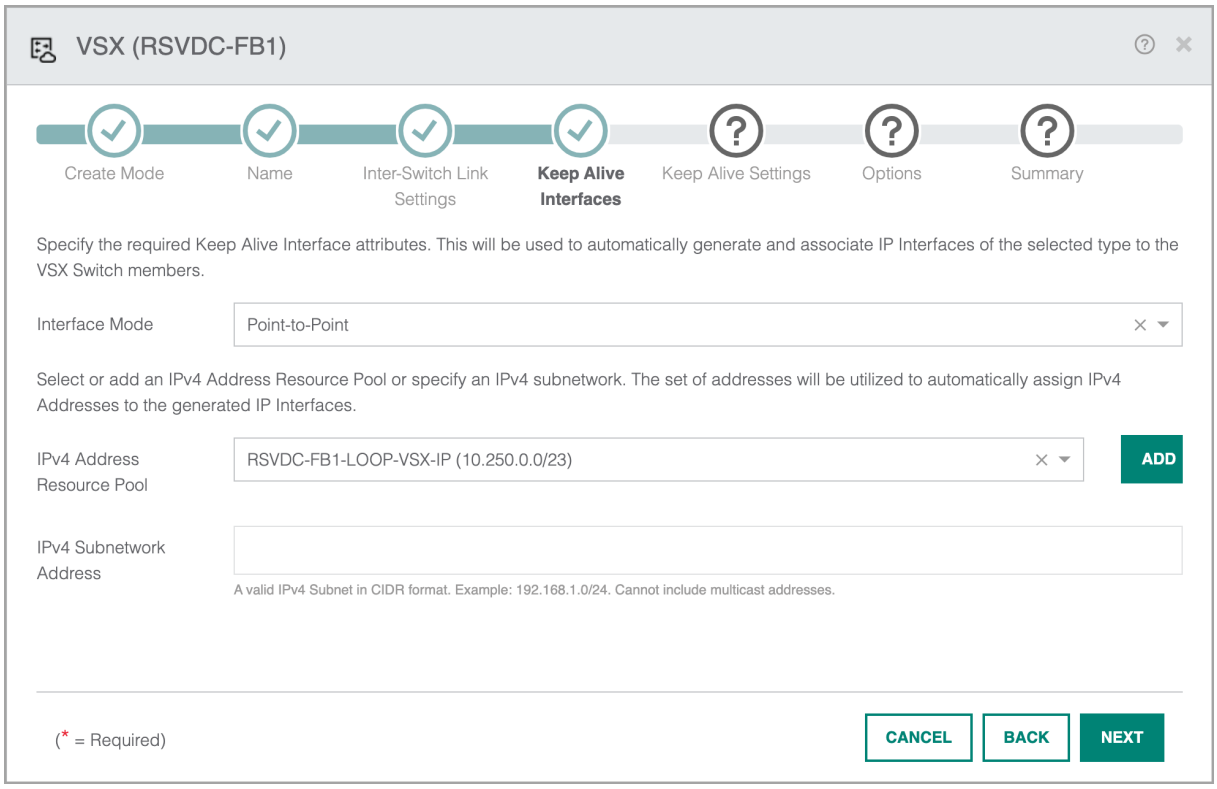


**Step 10** On the **Keep Alive Settings** page, leave the default values and click **NEXT**.

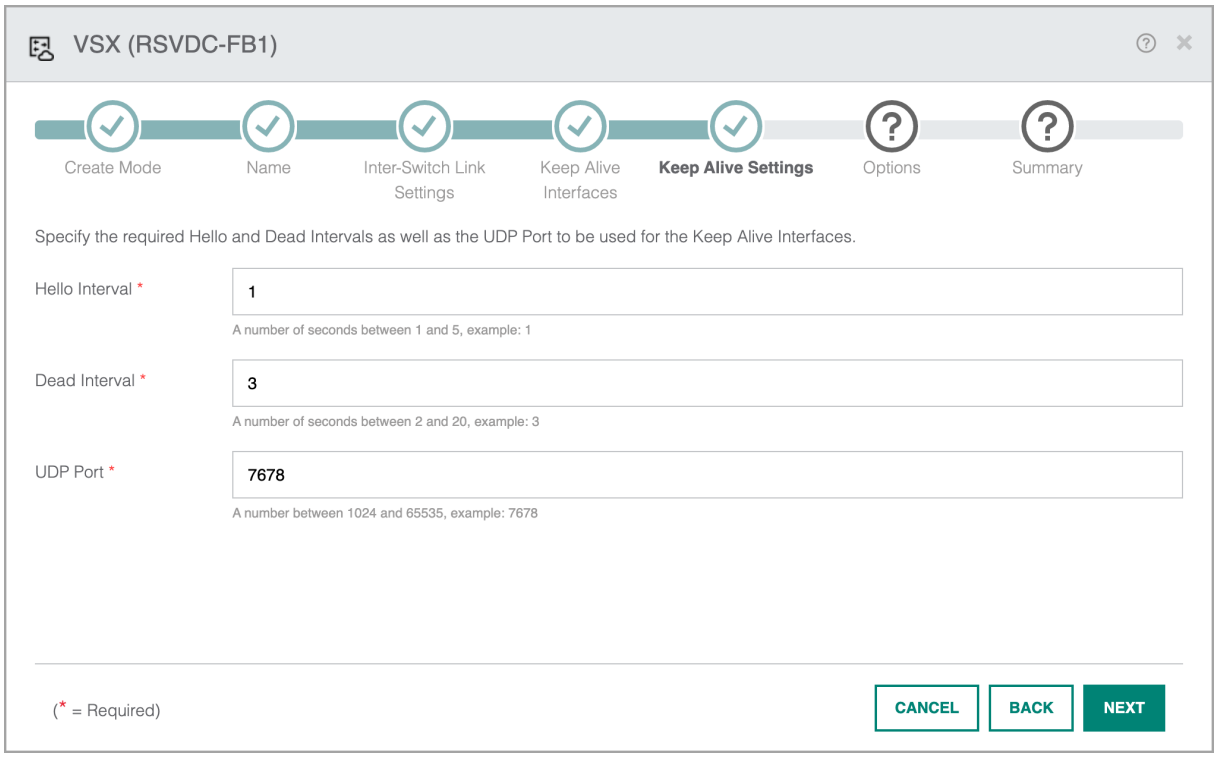


**Step 11** On the **Options** page, enter the value *600* for the **Linkup Delay Timer** field. Click **ADD** to launch the **Resource Pool** wizard.

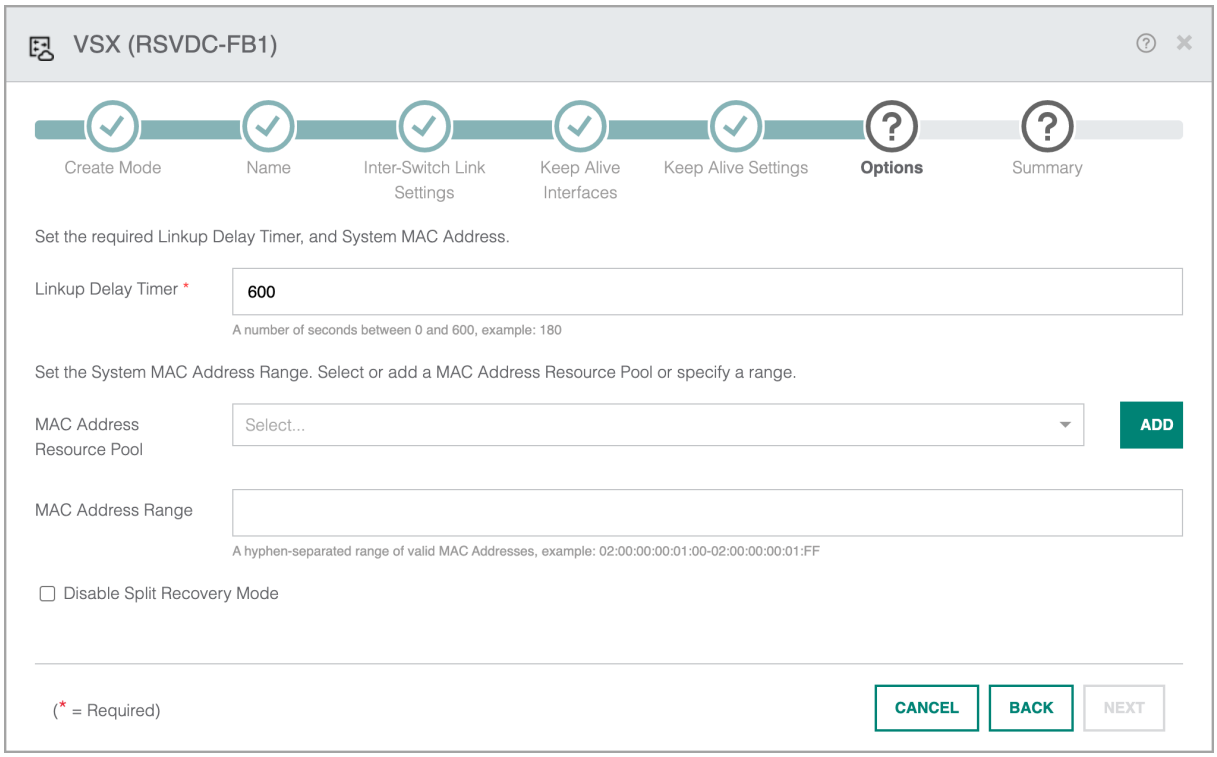


**NOTE:**

It is recommended to set a 600-second **Linkup Delay Timer** value on CX 10000 switches using firewall policy to ensure that policy and state have synchronized before forwarding traffic attached on a multi-chassis LAG.

**Step 12 Resource Pool** wizard: On the **Name** page, enter a **Name** and **Description** for the system MAC address pool. Click **NEXT**.

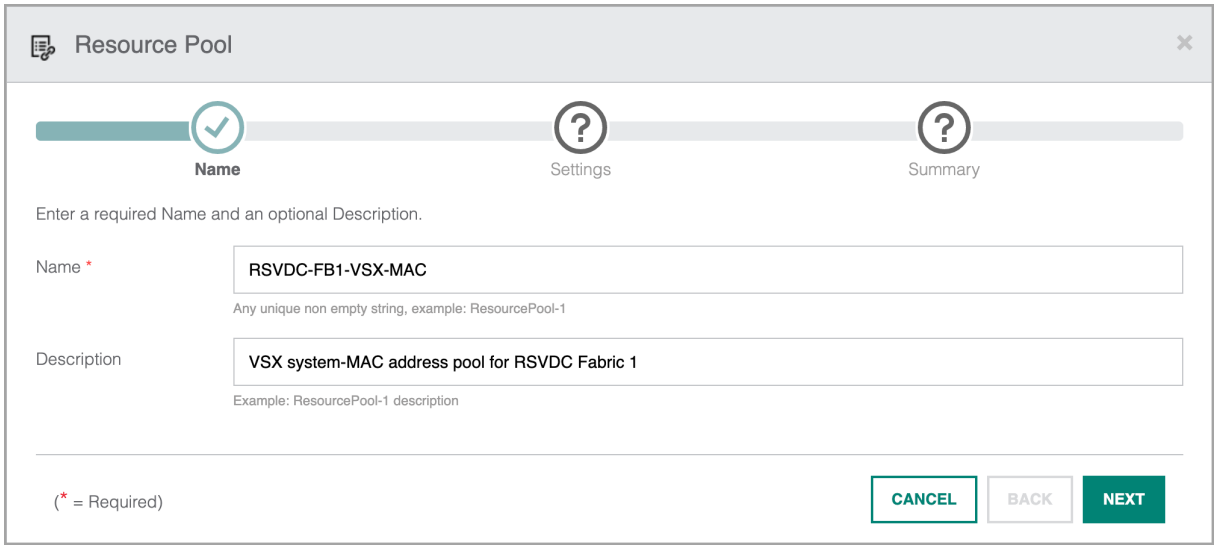


**Step 13 Resource Pool** wizard: On the **Settings** page, enter a MAC address range to be used for the VSX system MAC addresses, then click **NEXT**.

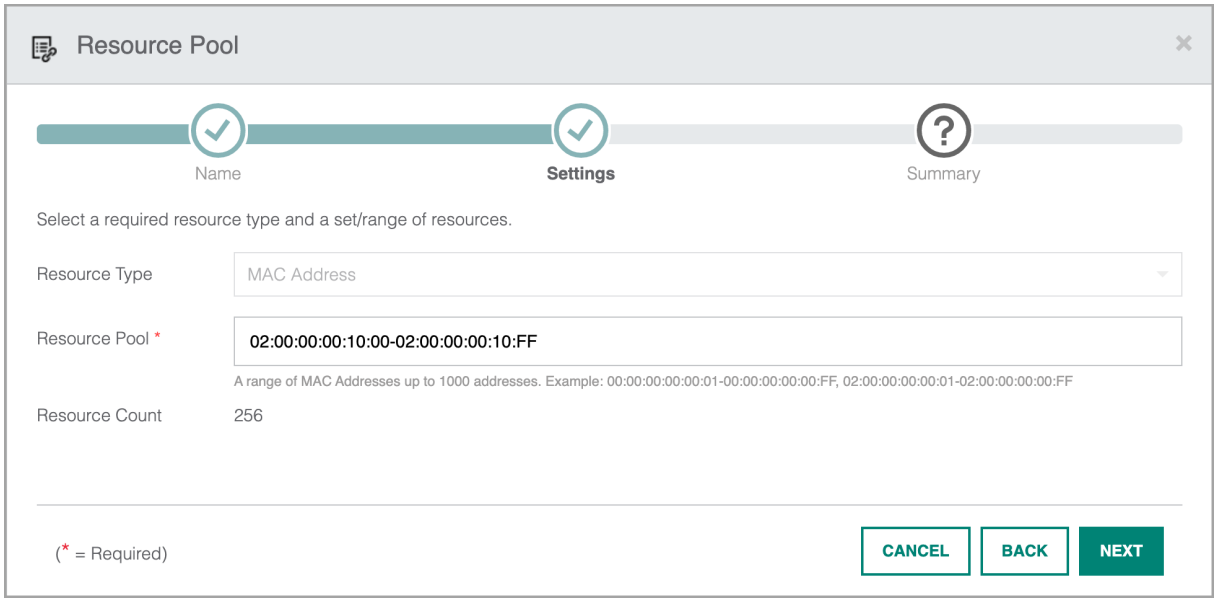


**Step 14 Resource Pool** wizard: On the **Summary** page, verify the system MAC address pool information and click **APPLY**. The **Resource Pool** wizard closes and returns to the main **VSX Configuration** workflow.

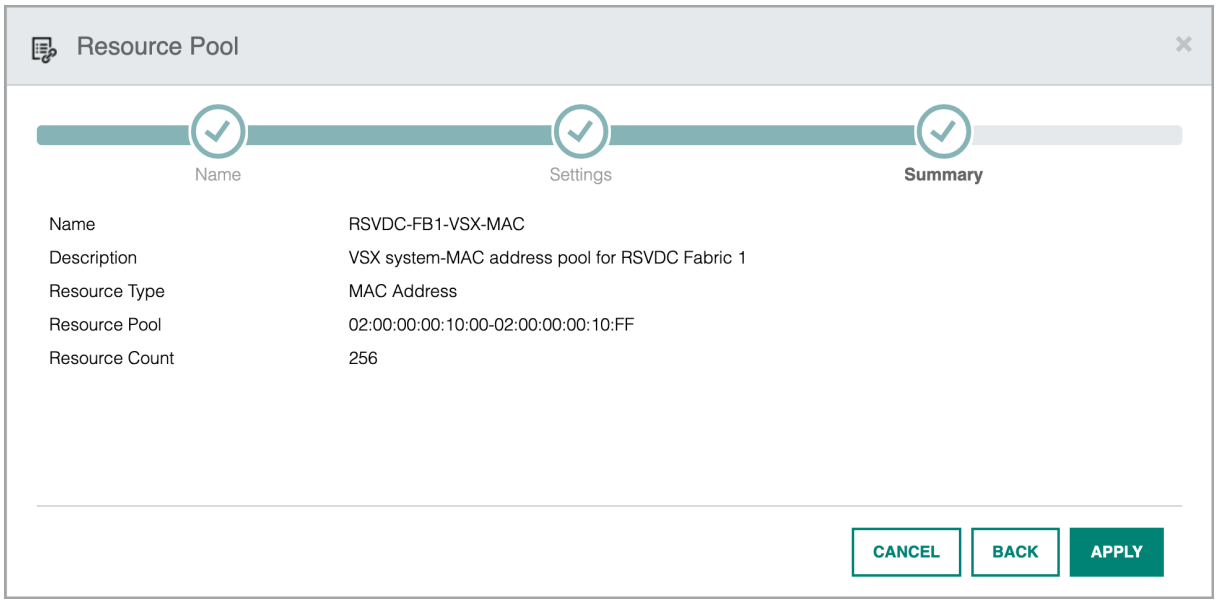


**Step 15** On the **Options** page, verify that the new **MAC Address Resource Pool** is selected and click **NEXT**.

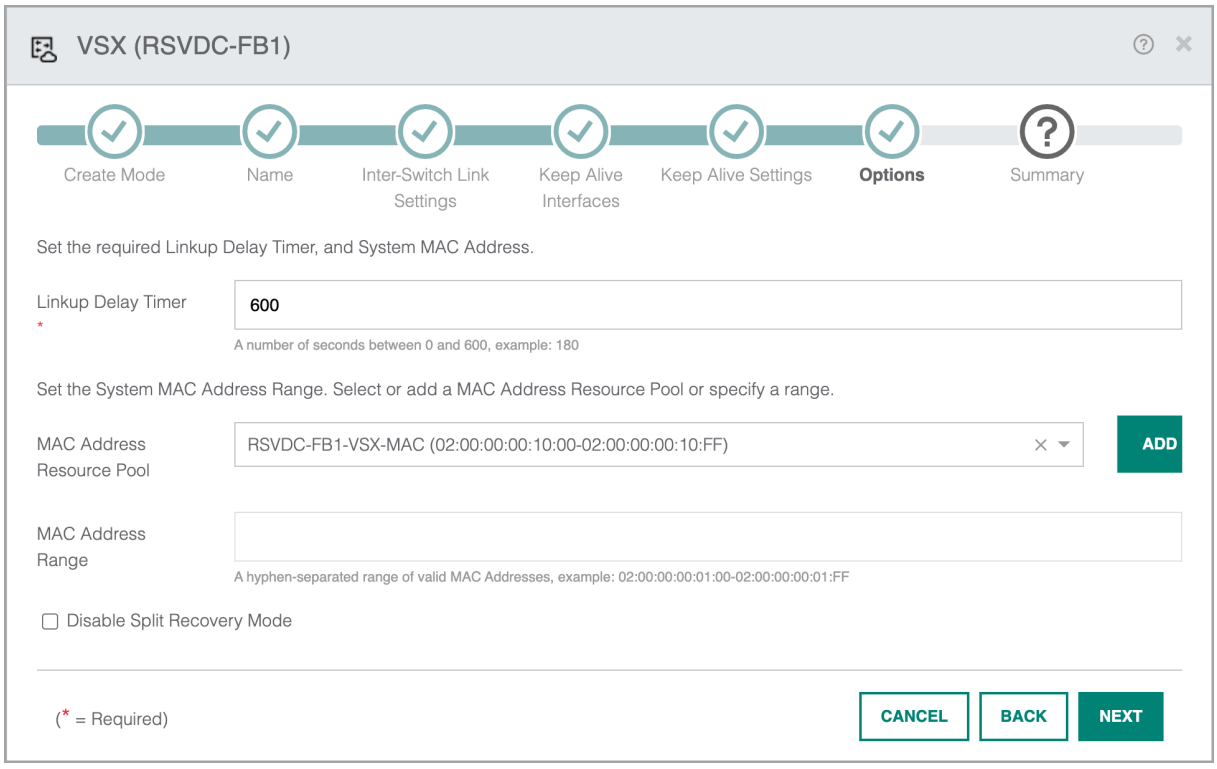


**Step 16** On the **Summary** page, verify the complete set of VSX settings and click **APPLY**.

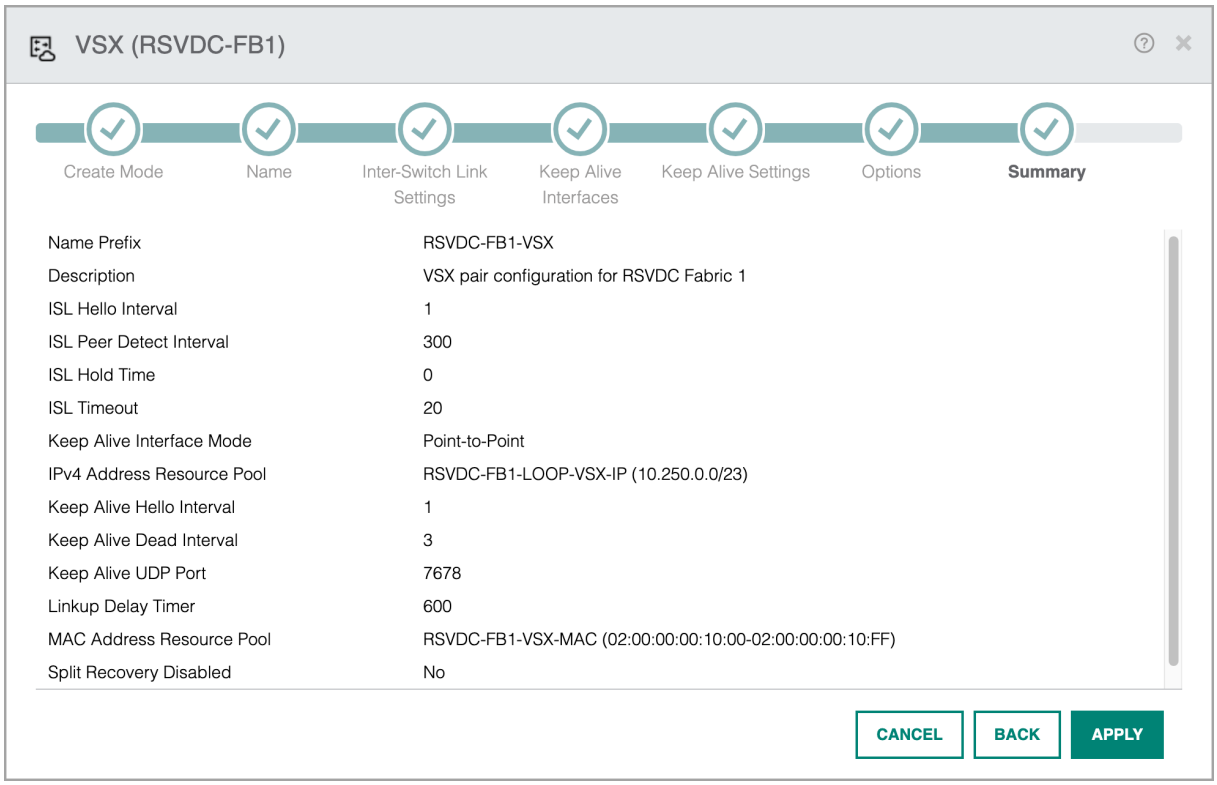


**Step 17 Guided Setup** displays the list of VSX pairs in the **Configuration / Network / VSX** window. Review the information to verify that the VSX pairs created are consistent with physical cabling.

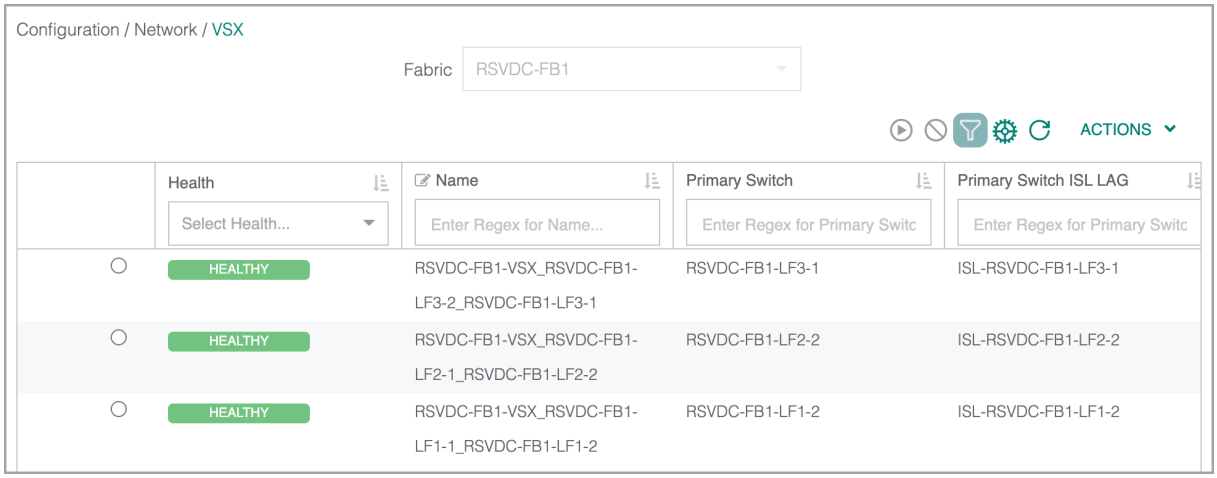


**NOTE:**

VSX **Health** status in Fabric Composer can update slowly. Click the **Refresh** button in the upper right of the **Configuration / Network / VSX** window to refresh the switch status manually.

**Step 18** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**

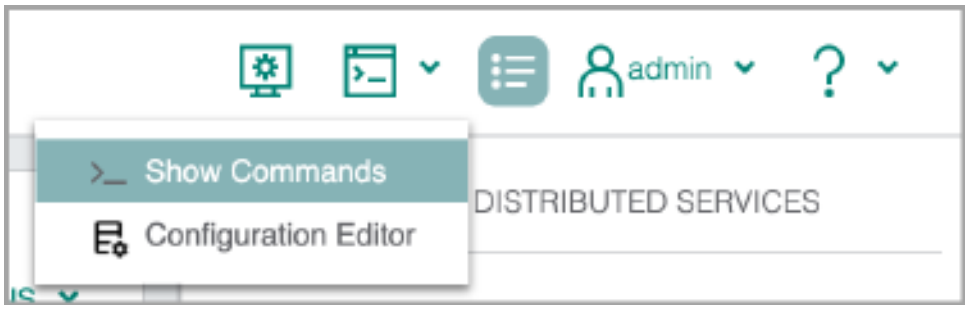


**Step 19** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *All leaf switches* >
- **Commands:** *show vsx status*

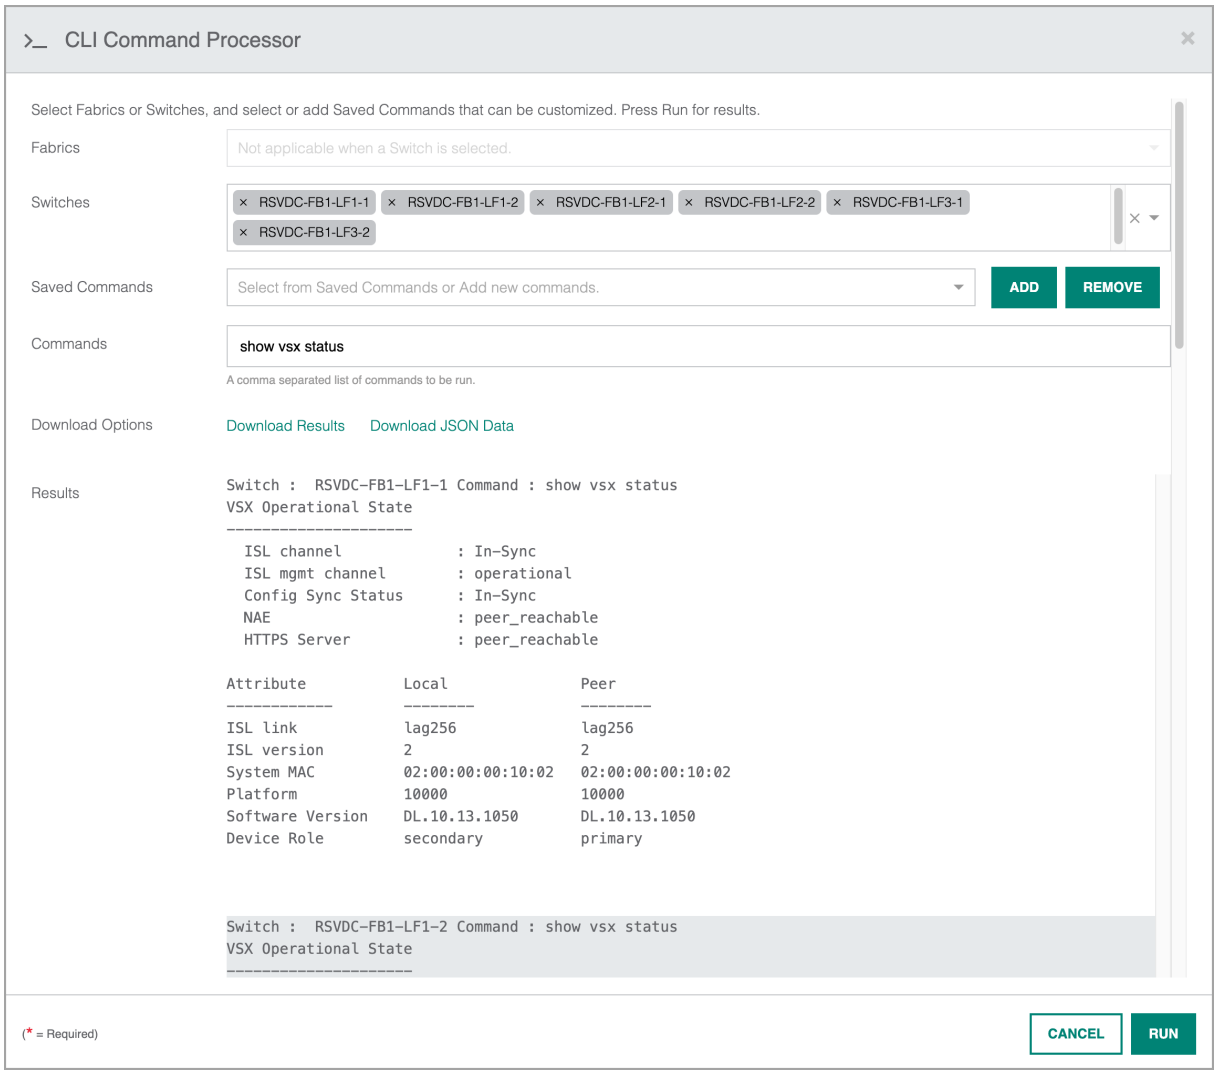


**Step 20** Verify that each switch has both **Local** and **Peer** information populated with the following values:

- **ISL channel:** *In-Sync*
- **ISL mgmt channel:** *operational*
- **Config Sync Status:** *In-Sync*
- **NAE:** *peer_reachable*

- **HTTPS Server:** *peer_reachable*

## Configure Layer 3 Leaf-to-Spine Connections

Fabric Composer automatically identifies leaf-to-spine connections and configures them with the values supplied in the **Leaf-Spine** wizards. A resource pool is created to assign IP addresses to routed leaf and spine interfaces using /31 subnets. At the completion of this procedure, IP addresses are assigned to all interfaces required to support deployment of the OSPF fabric underlay.

**Step 1** On the **Guided Setup** menu, select **L3 LEAF-SPINE CONFIGURATION** to start the **Leaf-Spine** workflow.

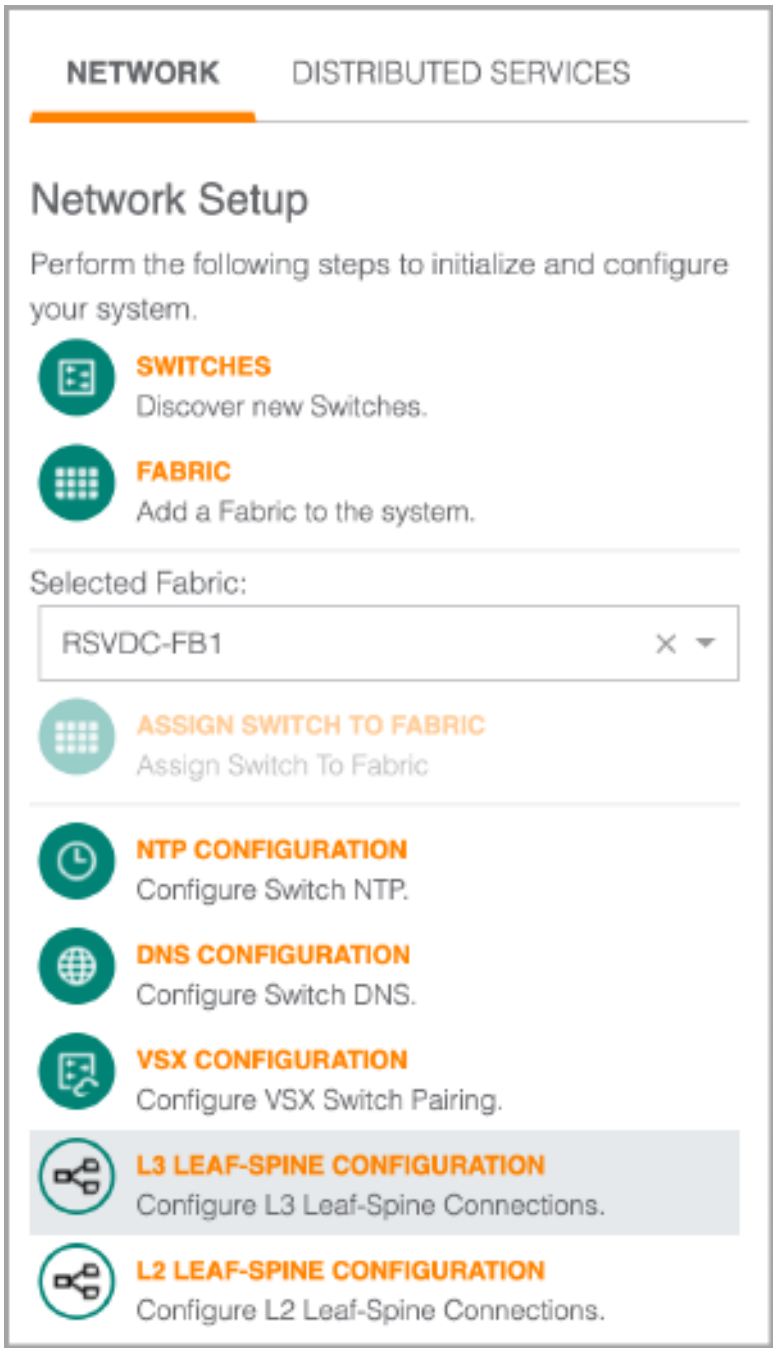

**Step 2** On the **Create Mode** page, leave **Automatically generate Leaf-Spine Pairs** selected and click **NEXT**.

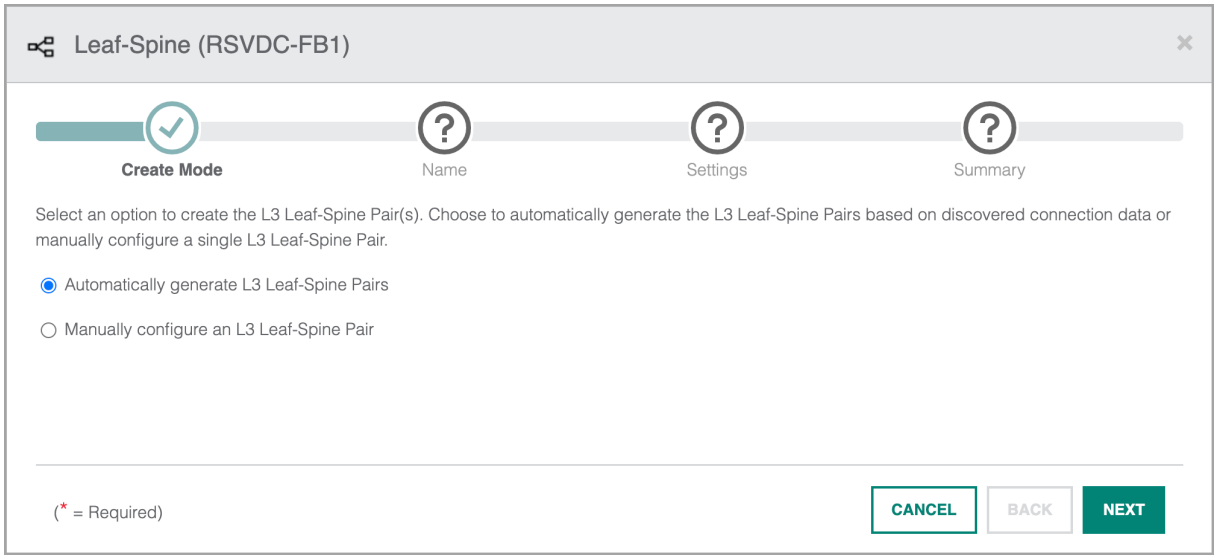


**Step 3** On the **Name** page, enter a **Name Prefix** and **Description**, then click **NEXT**.

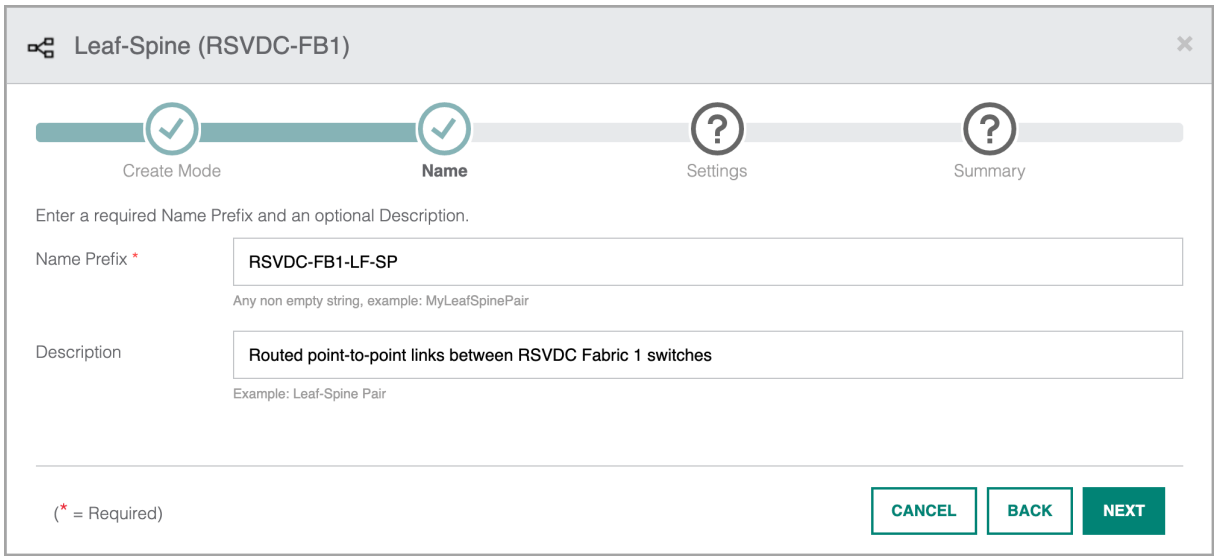


**Step 4** On the **Settings** page, click **ADD** to launch the **Resource Pool** wizard.

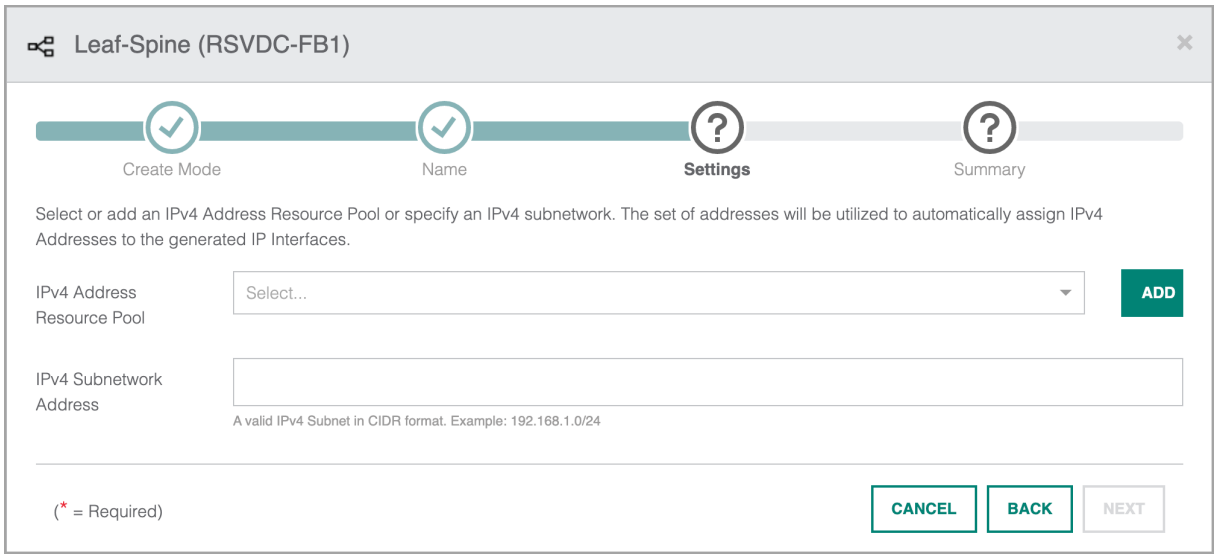


**Step 5 Resource Pool** wizard: On the **Name** page, enter a **Name** and **Description** for the IPv4 address pool, then click **NEXT**.

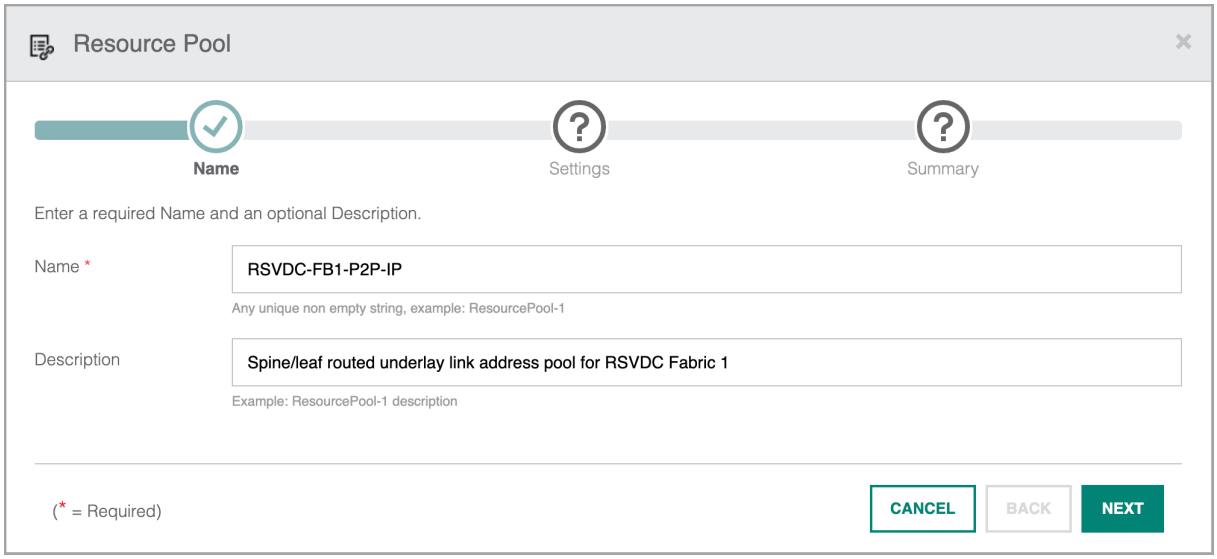


**Step 6 Resource Pool** wizard: On the **Settings** page, enter an IPv4 address block in the **Resource Pool** field and click **NEXT**.

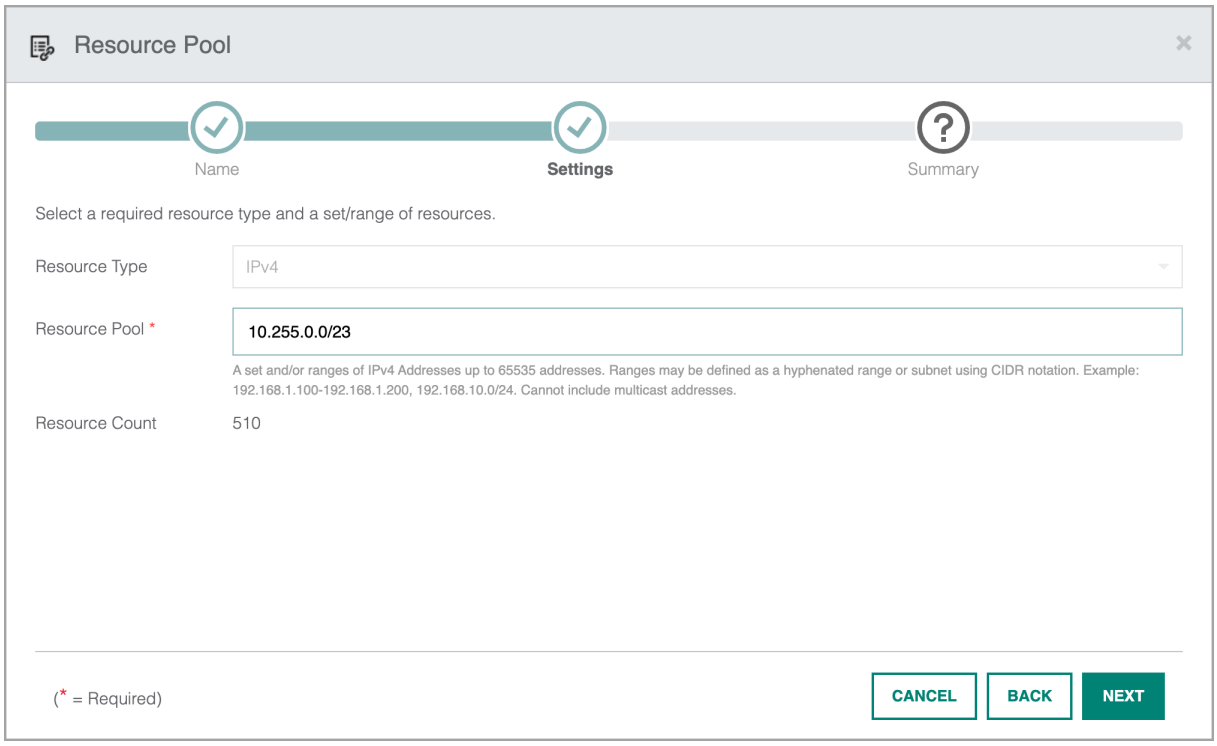


**NOTE:**

Use a subnet distinct from other subnets used in the overlay networks. The assigned subnet is used to configure routed ports between fabric switches. Use a block large enough to accommodate anticipated fabric growth.

**Step 7 Resource Pool** wizard: On the **Summary** page, verify the IP address pool information and click **APPLY**. The **Resource Pool** wizard closes and returns to the main **Leaf-Spine Configuration** workflow.

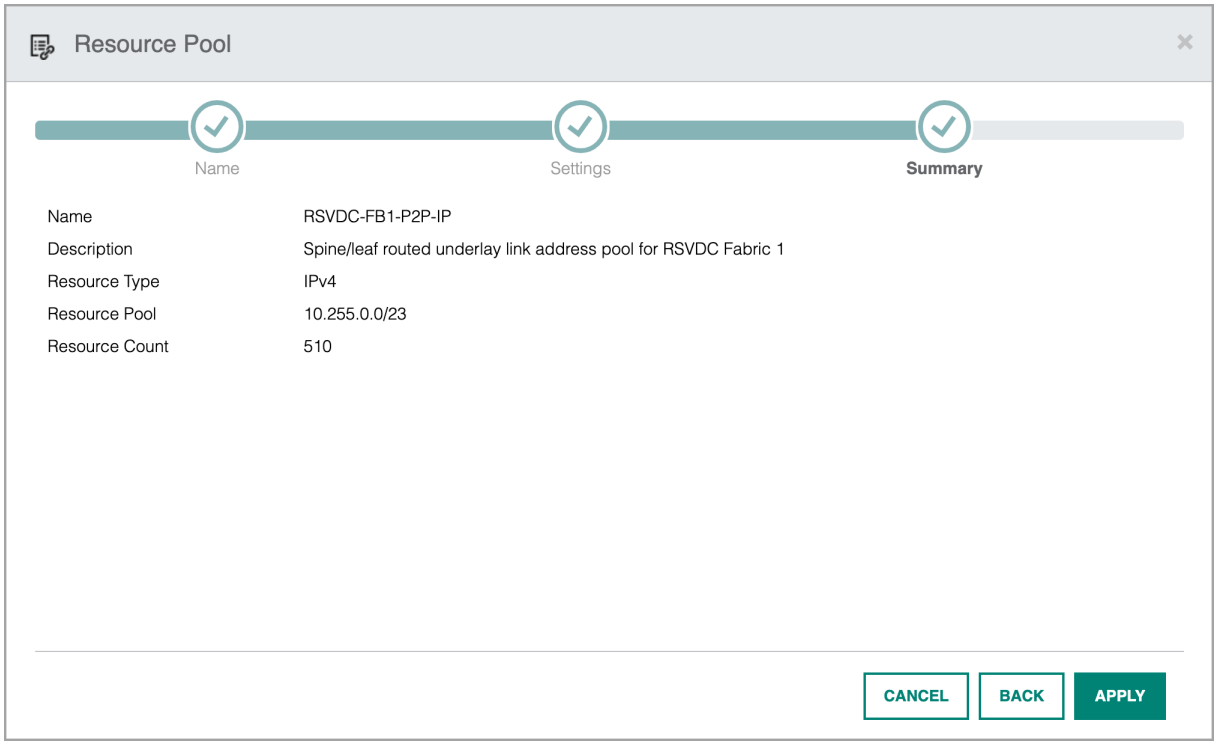

**Step 8** On the **Settings** page, verify that the new **IPv4 Address Resource Pool** is selected and click **NEXT**.

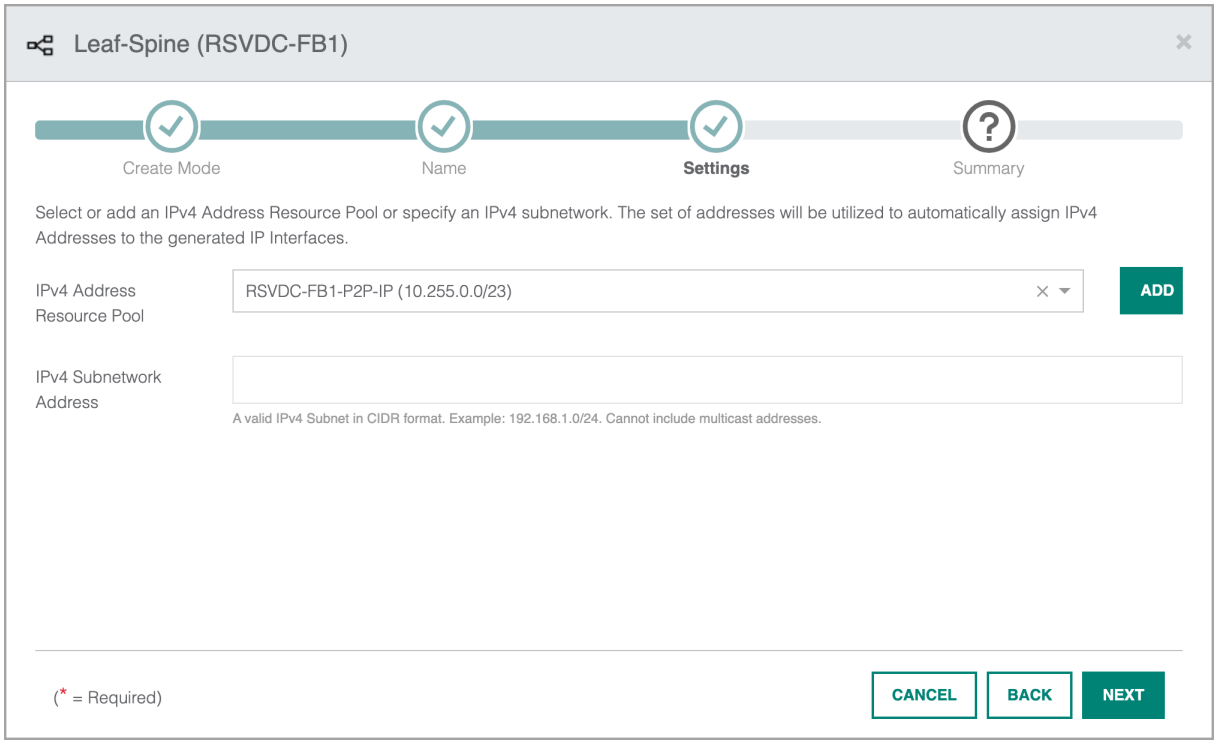


**Step 9** On the **Summary** page, verify that the information is correct and click **APPLY**.

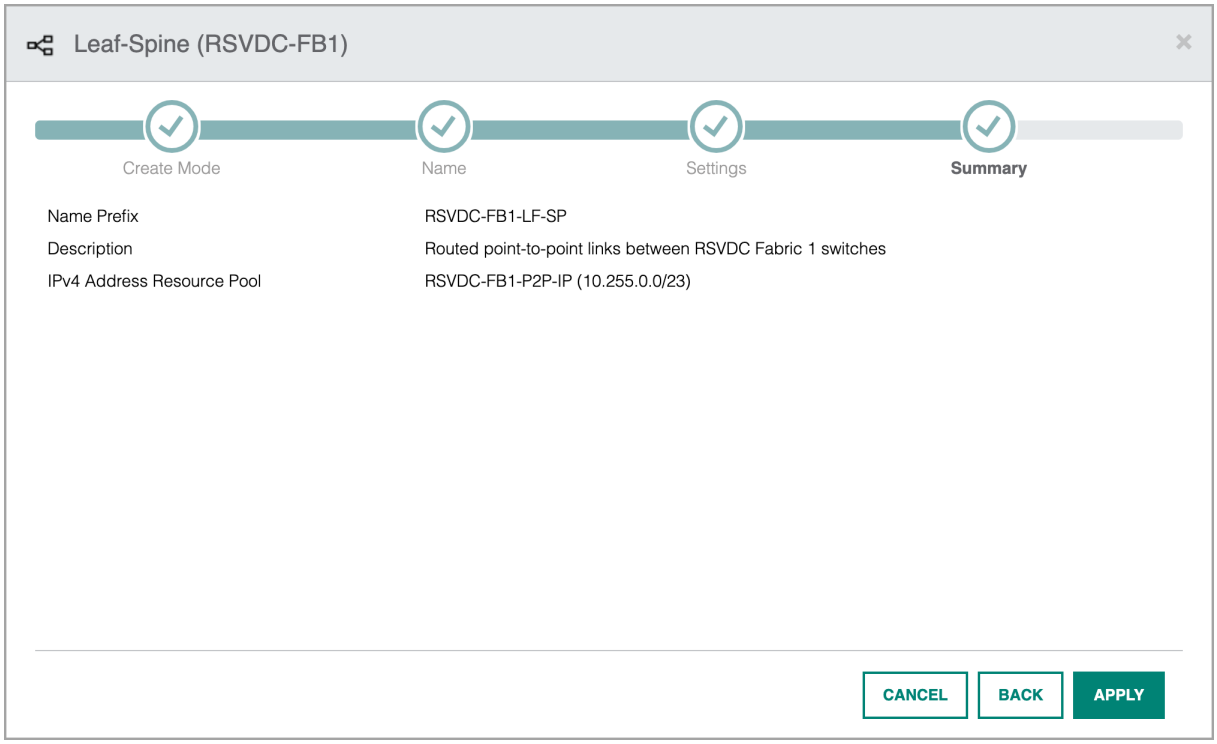


**Step 10 Guided Setup** displays the list of leaf-to-spine links in the **Configuration/Network/Leaf-Spine** window. Review the information to verify that the leaf-spine links created are consistent with physical cabling.

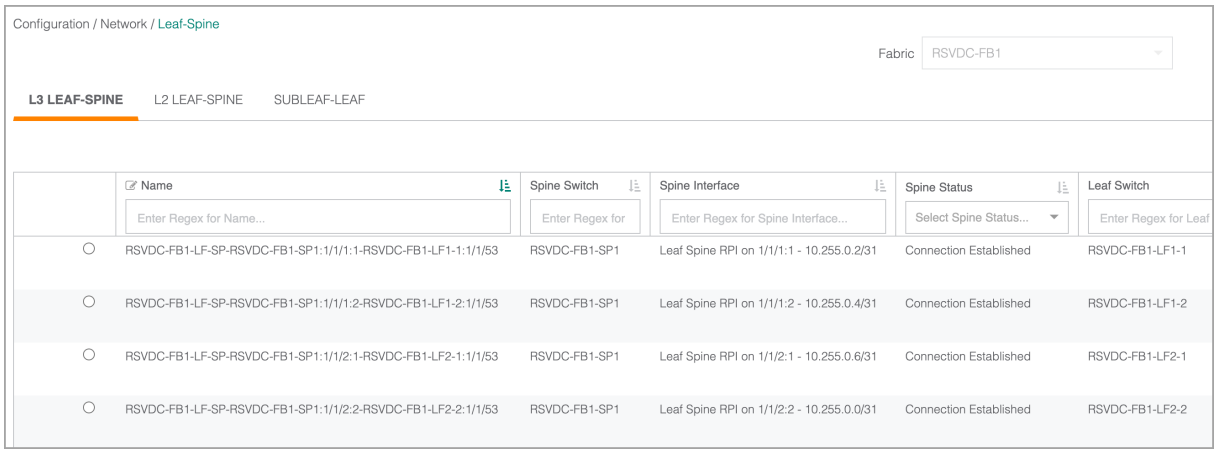


## Configure Server Access Switch Links

Fabric Composer refers to server access switches as subleaf switches. Compute and storage hosts are typically attached directly to leaf switches. Server access switches are primarily used to achieve two objectives: they provide a transition strategy to connect existing server infrastructure into an EVPN-VXLAN fabric, and they provide an economical strategy to support a large number of 1 Gbps connected hosts. Server access switches extend Layer 2 services from the leaf, but do not participate directly in underlay routing or overlay virtualization mechanisms.

The following procedure establishes an MC-LAG between a VSX leaf pair and a downstream VSF server access switch stack. The LAGs defined on both sets of switches are 802.1Q trunks that allow all VLANs.

The diagram below highlights the server access MC-LAG created in this procedure.

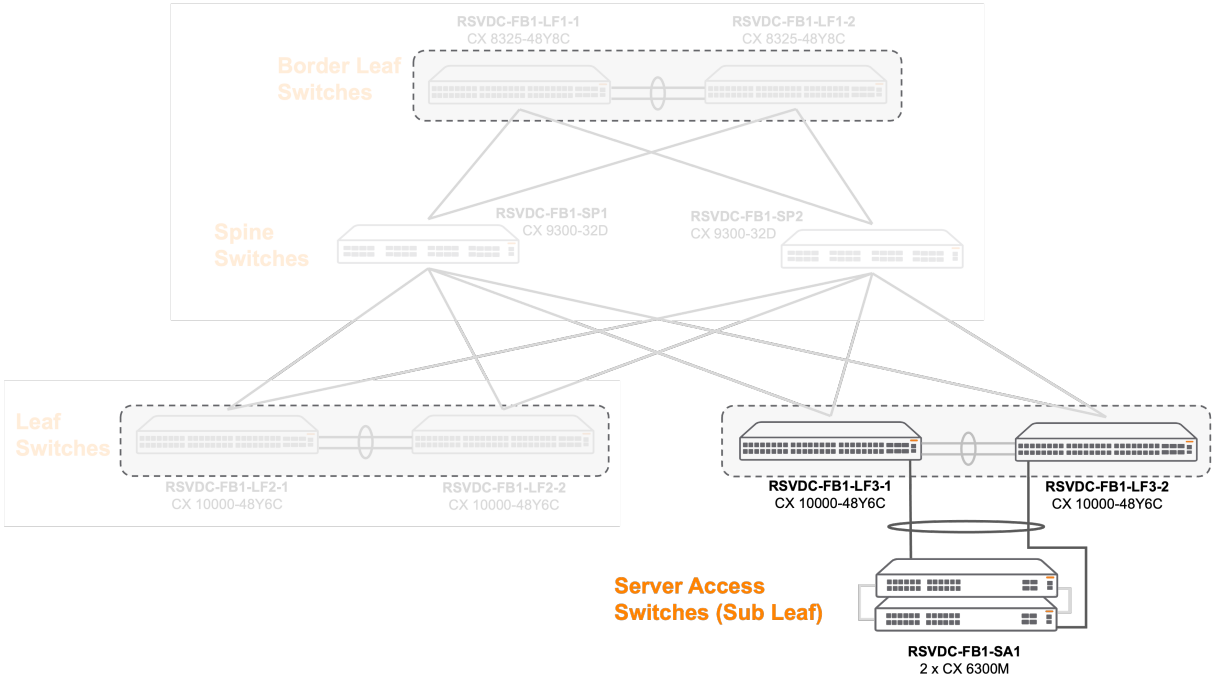


**Step 1** On the **Configuration > Network > Leaf-Spine** page, click **SUBLEAF-LEAF**.

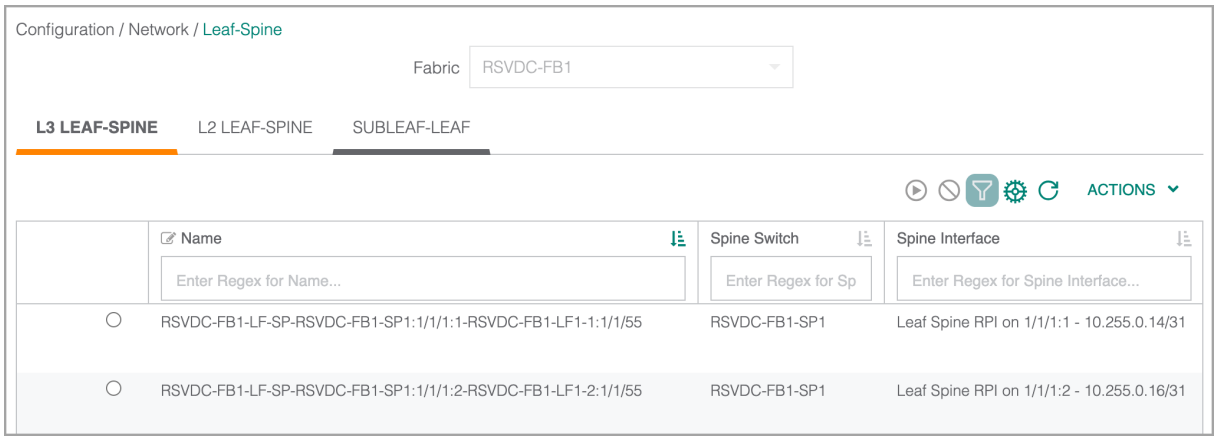


**Step 2** On the **ACTIONS** menu, select **Add**.

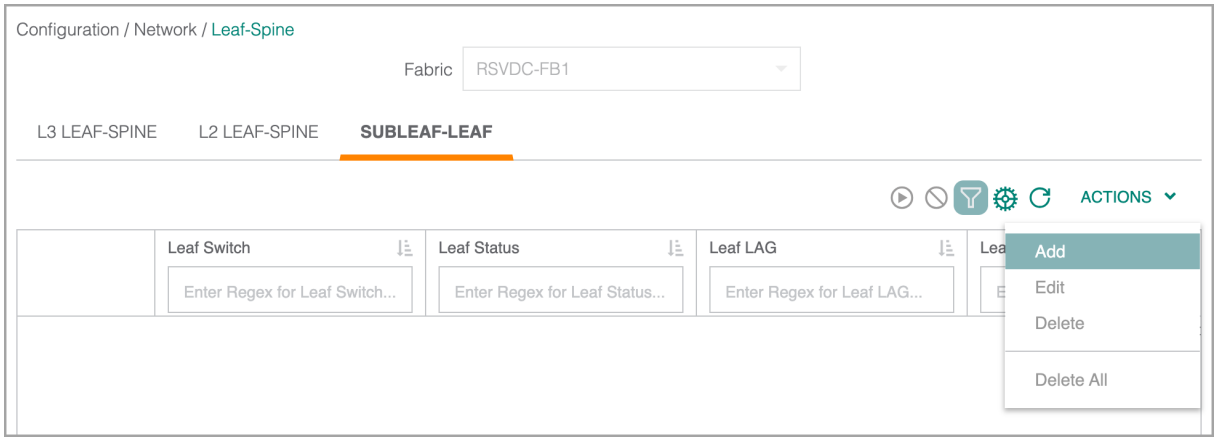


**Step 3** When prompted to continue, click **OK**.

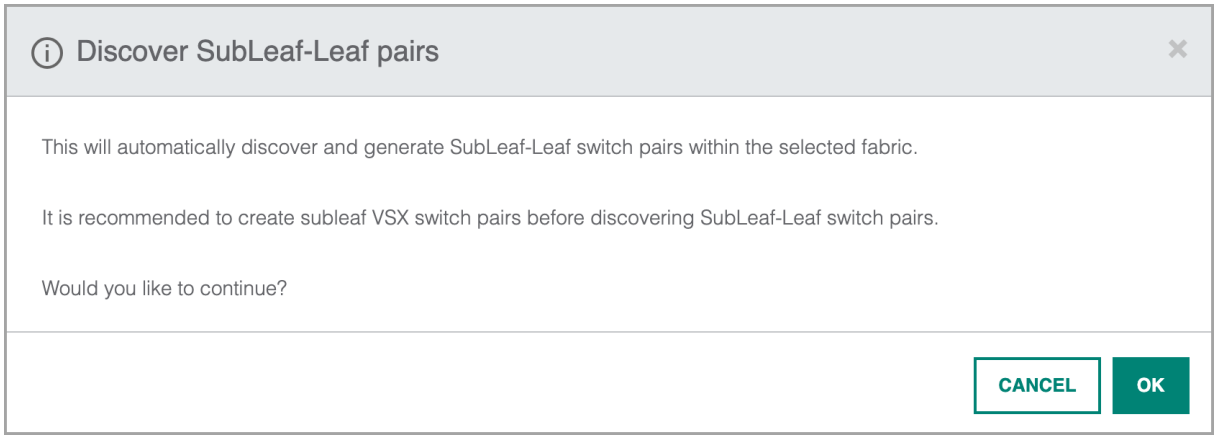


**Step 4** Review the leaf and server access MC-LAG information. Verify that the values in the **Leaf LAG Status** and **SubLeaf LAG Status** columns are *up*.

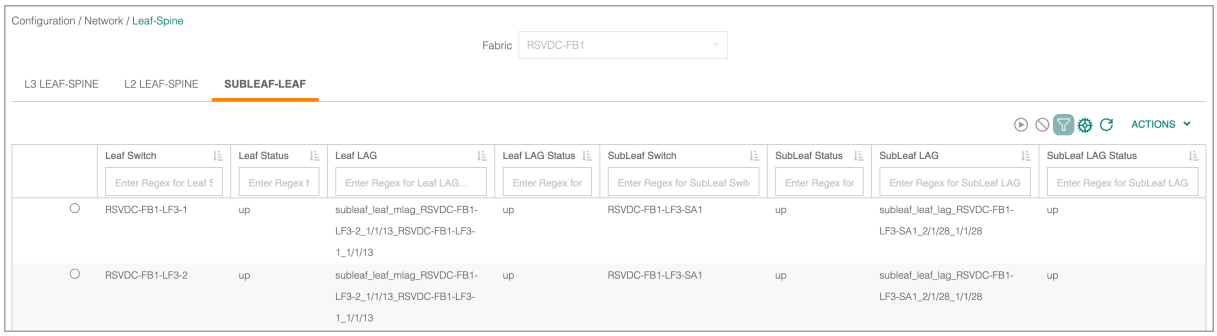


**NOTE:**

The status field values may take a few minutes to populate and may require a screen refresh.

## Configure Underlay Network Routing

The HPE Aruba Networking data center spine-and-leaf design uses OSPF as the underlay routing protocol. The Fabric Composer **Underlay Configuration** wizard creates a transit VLAN between redundant ToRs to support routing adjacency, assigns IP addresses to loopback and transit VLAN interfaces, and creates underlay OSPF configuration. OSPF shares the loopback0 IP addresses for later use in establishing overlay routing. The Fabric Composer CLI Command Processor verifies OSPF adjacencies.

At the completion of this procedure, a functional underlay for the data center fabric is complete. The diagram below illustrates the assigned loopback IP addresses and the links where OSPF adjacencies are formed between leaf and spine switches.

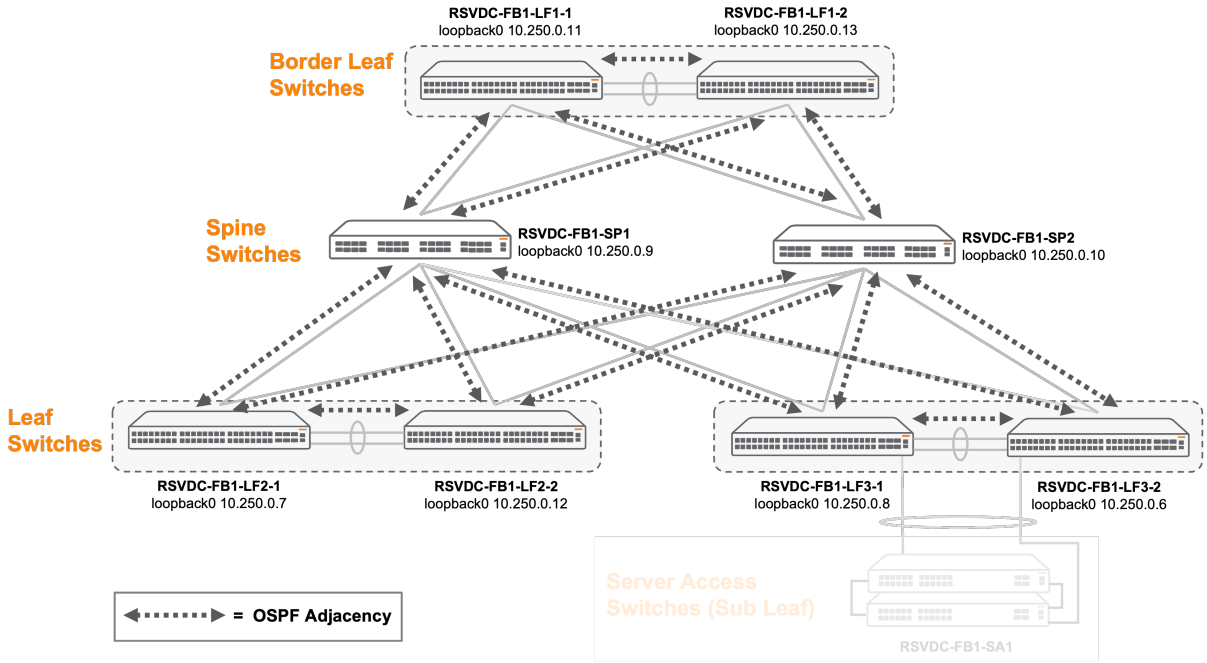


**Step 1** On the **Guided Setup** menu, select **UNDERLAYS** to start the **Underlay Configuration** workflow.

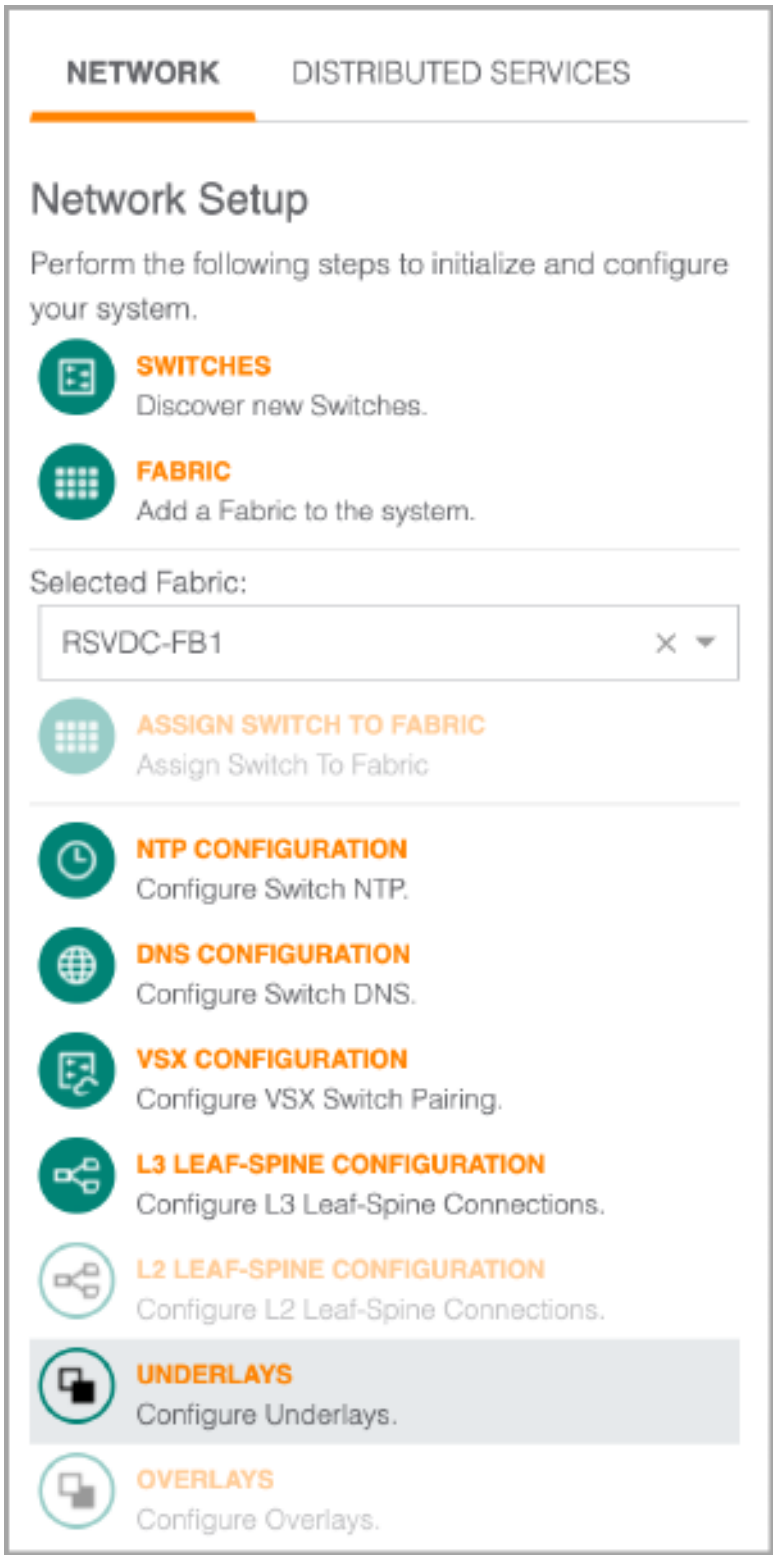


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

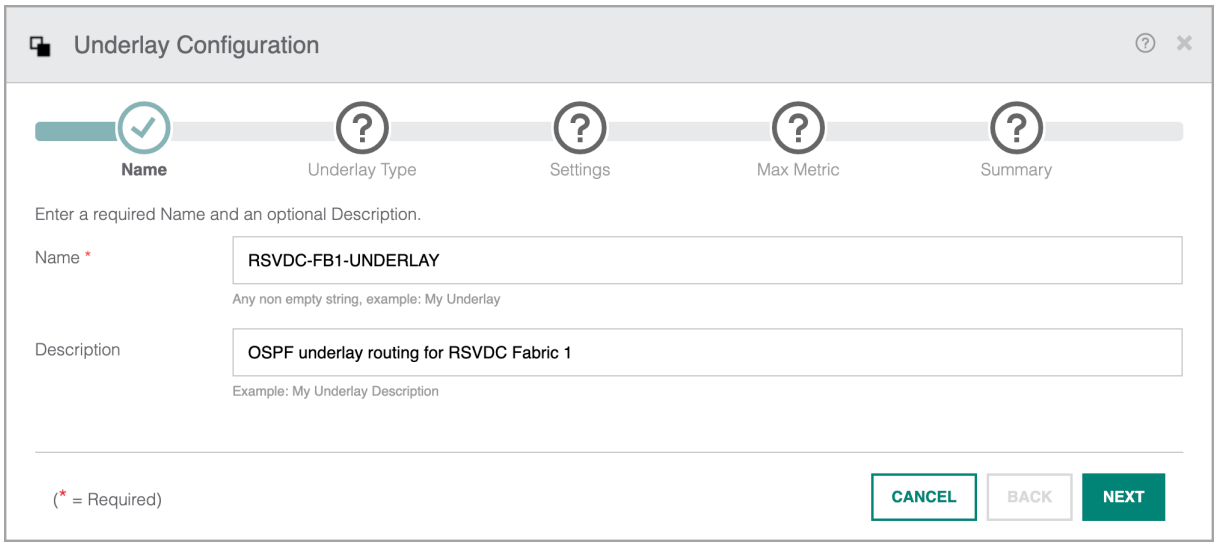


**Step 3** On the **Underlay Type** page, leave the default **OSPF** selection and click **NEXT**.

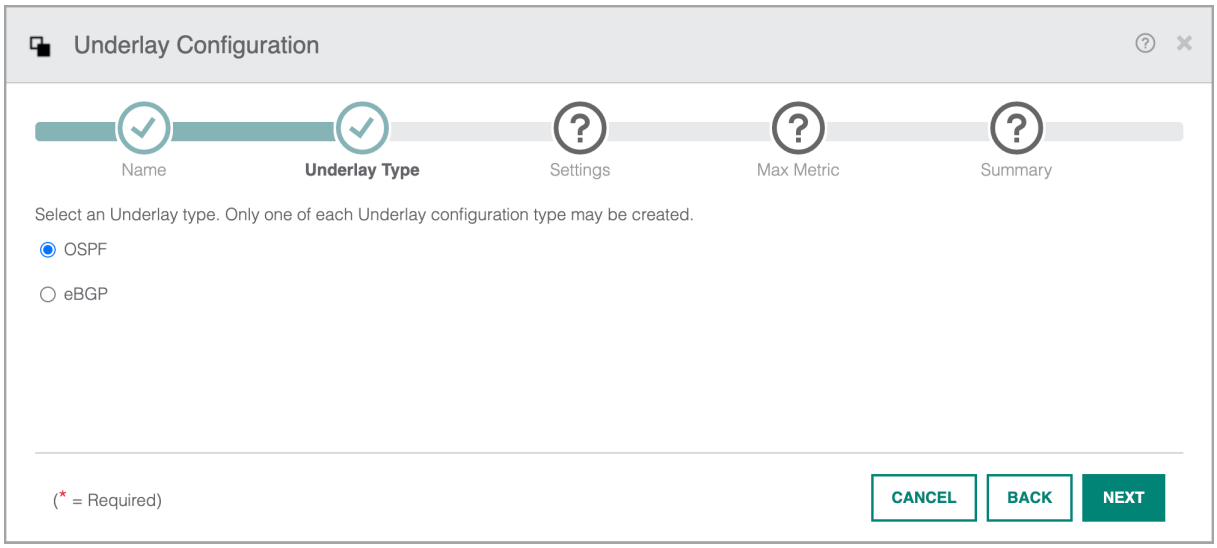


**Step 4** On the **Settings** page, set the **Transit VLAN** to *3999*. Leave other settings at their defaults and click **NEXT**.

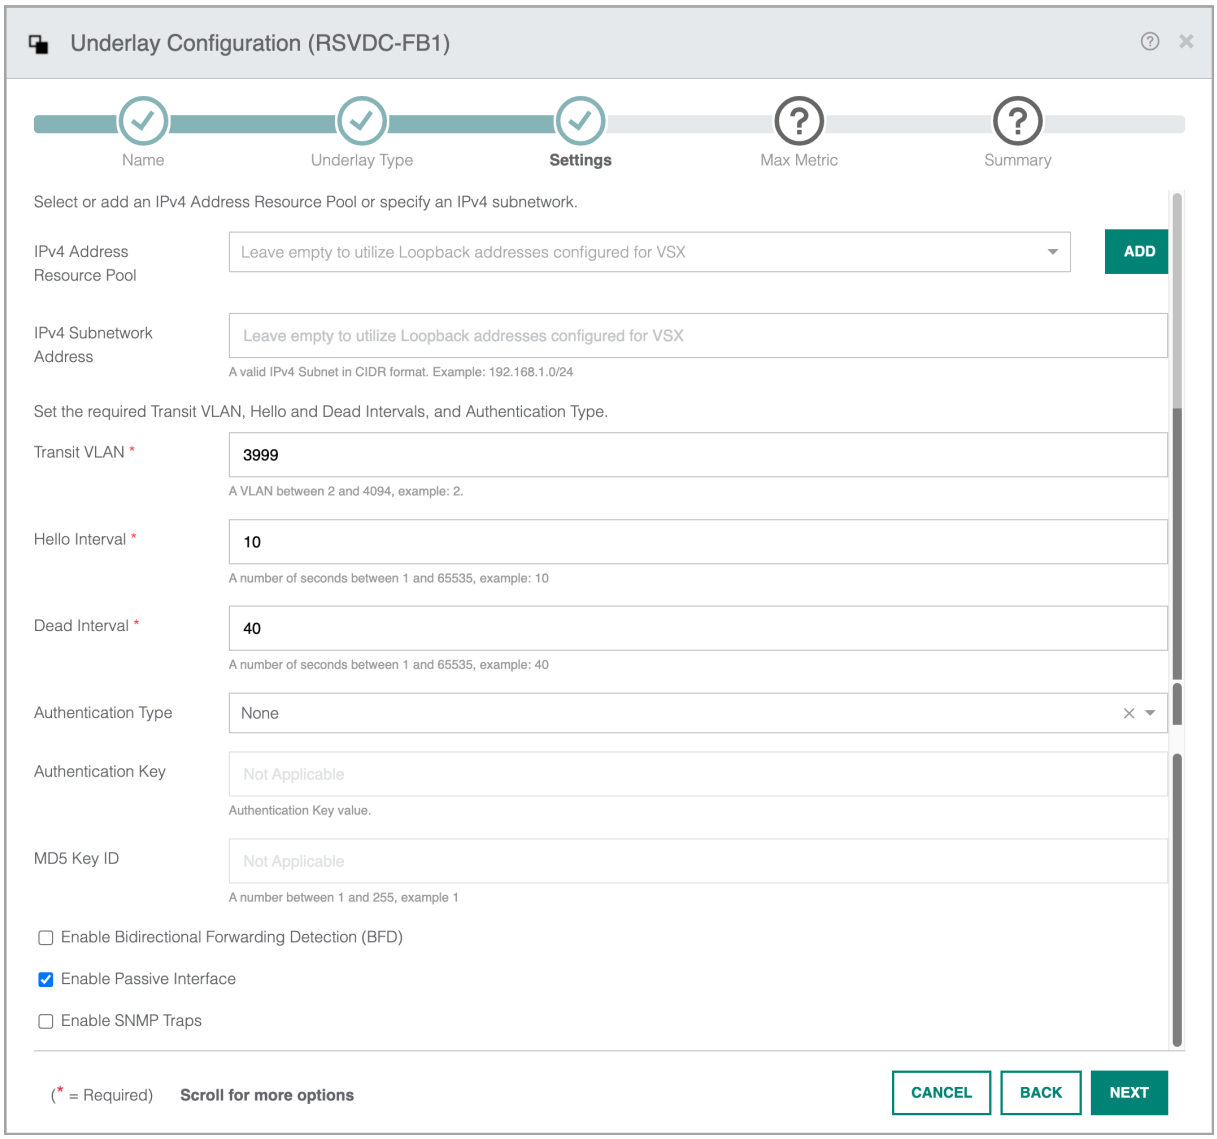


**NOTE:**

Enter a VLAN ID that cannot be confused easily with other VLANs within the network.

**Step 5** On the **Max Metric** page, enter the value *600* in the **On Startup** field. Leave other settings at their defaults and click **NEXT**.

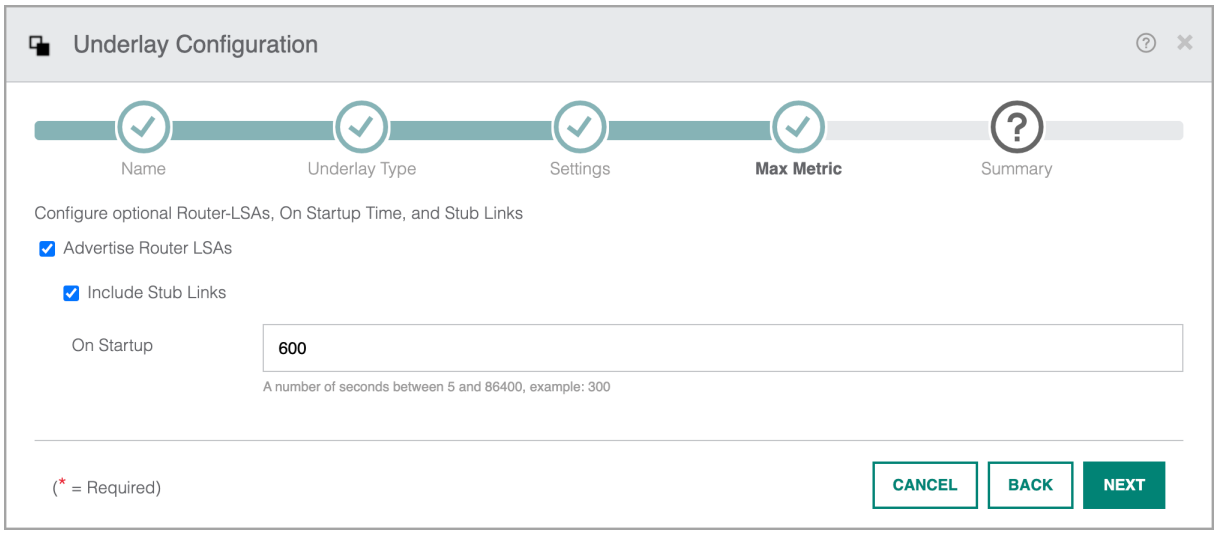


**NOTE:**

It is recommended to set a 600-second OSPF **On Startup** max metric value for CX 10000 switches using firewall policy in a VSX pair to ensure that policy and state have synchronized before fabric traffic is forwarded to the switch VTEP. The same value is applied to all switches in this sample fabric.

**Step 6** On the **Summary** page, verify that the information is entered correctly and click **APPLY** to create the OSPF configuration.

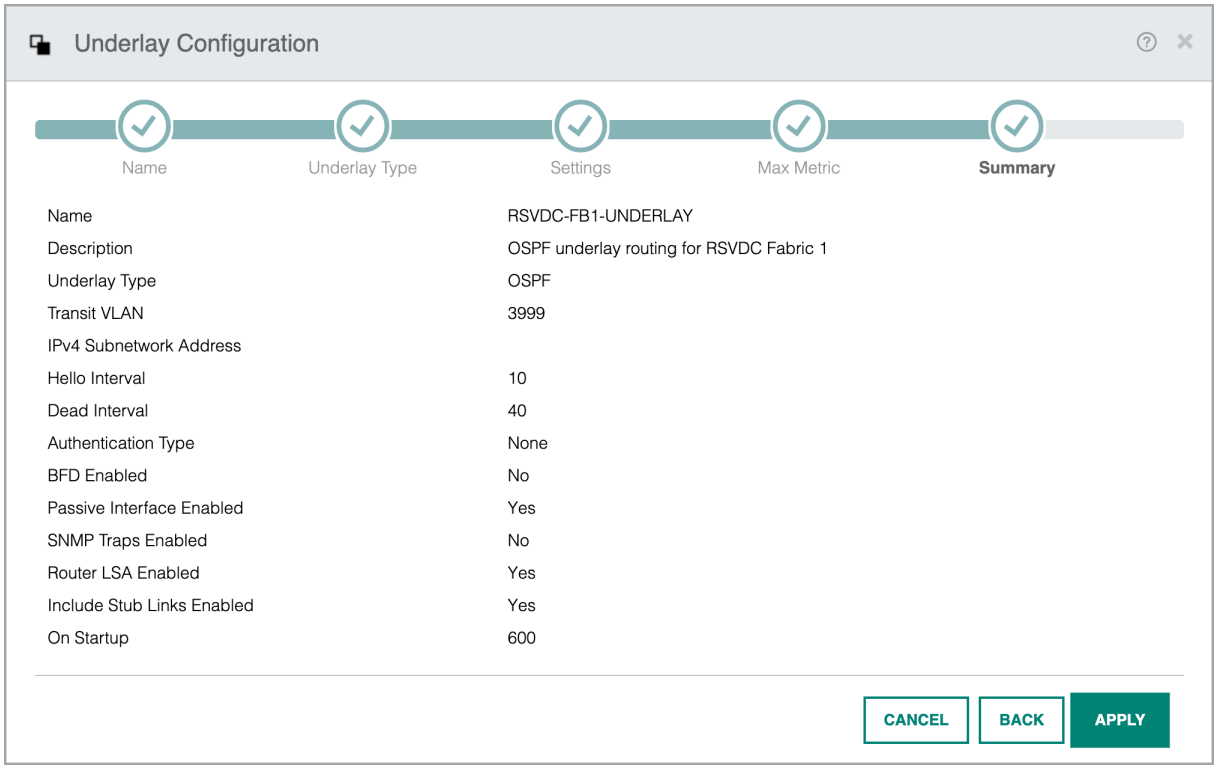


**Step 7** In the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 8** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Fabrics:** *RSVDC-FB1*
- **Commands:** *show ip ospf neighbors*

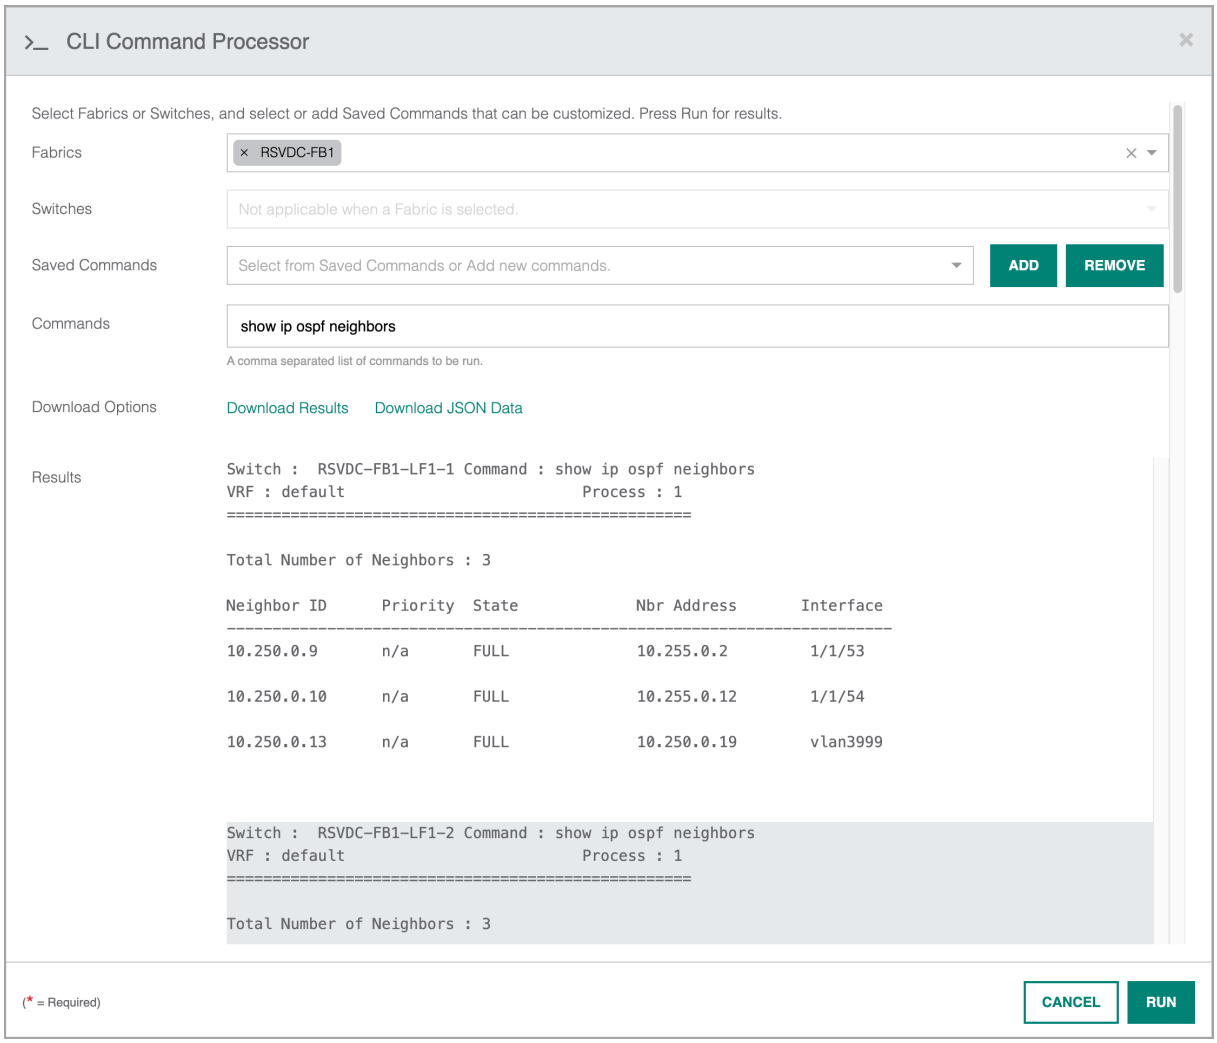


**Step 9** Verify that each spine switch shows an OSPF neighbor adjacency in the “FULL” state for all leaf switches. Verify that all leaf VSX pairs show an OSPF neighbor adjacency in the “FULL” state between themselves over the routed transit VLAN in addition to an adjacency in the “FULL” state with each spine.

## Configure Overlay Network Routing

The HPE Aruba Networking data center uses iBGP as the control plane for the fabric overlay within a single fabric. BGP provides a mechanism to build VXLAN tunnels dynamically and share host reachability across the fabric using the L2VPN EVPN address family. VTEP interfaces are the VXLAN encapsulation and decapsulation points for traffic entering and exiting the overlay. VSX leaf pairs share the same anycast VTEP IP address.

Use the Fabric Composer **Overlay Configuration** wizard to implement iBGP peerings using a private ASN and to establish VXLAN VTEPs. VTEP IP addresses are assigned as a switch loopback using a resource pool. iBGP neighbor relationships are verified using the Fabric Composer CLI Command Processor.

The diagram below illustrates the iBGP L2VPN EVPN address family peerings established using loopback interfaces between leaf switches and the two spines operating as iBGP route reflectors.

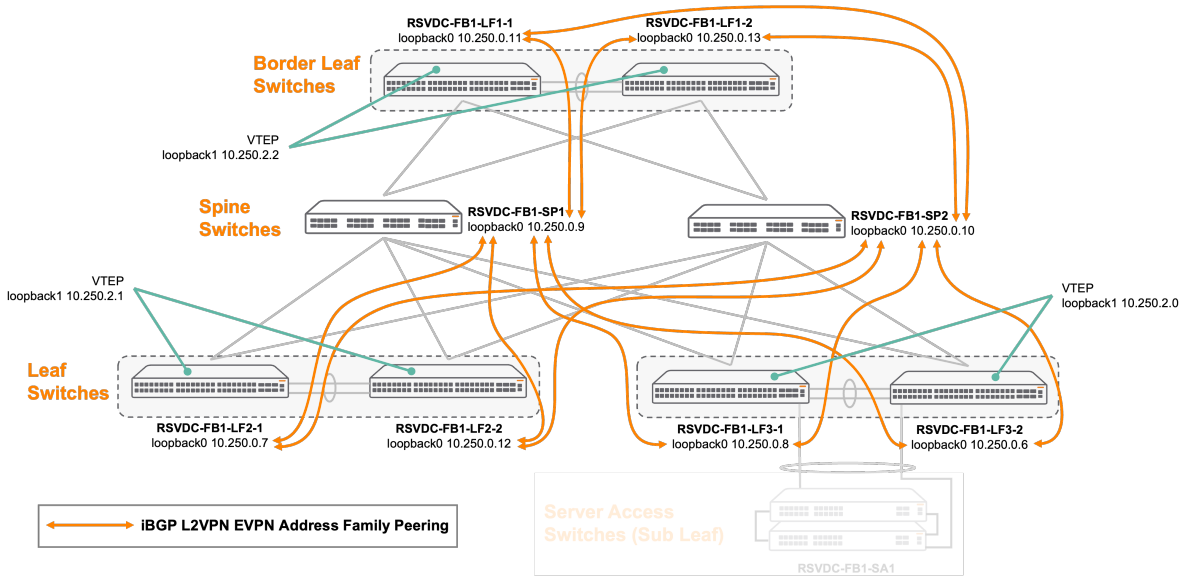


**Step 1** From the **Guided Setup** menu, select **OVERLAYS** to start the **Overlay Configuration** workflow.

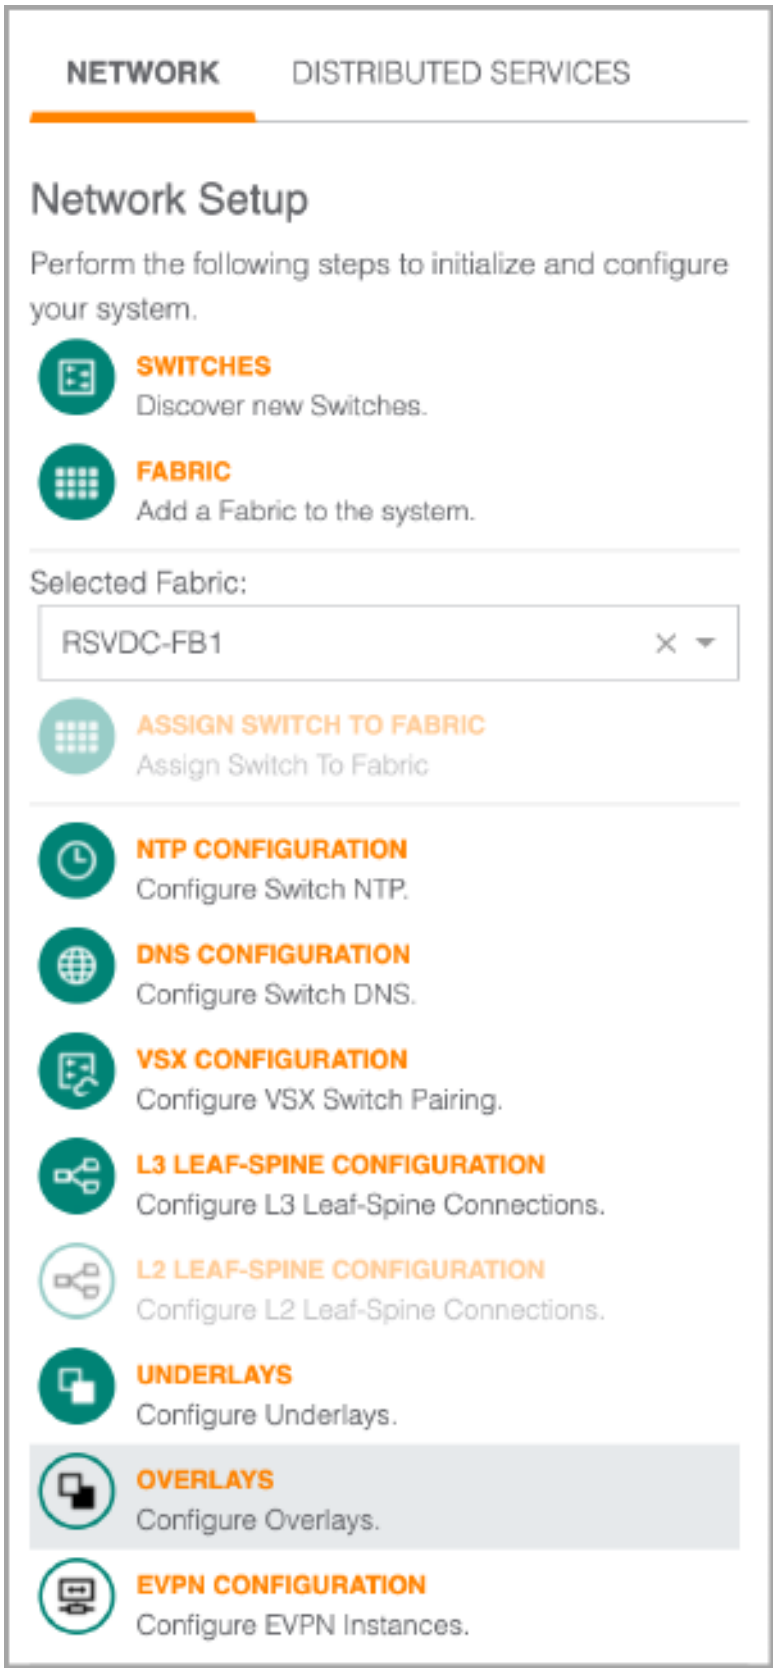


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

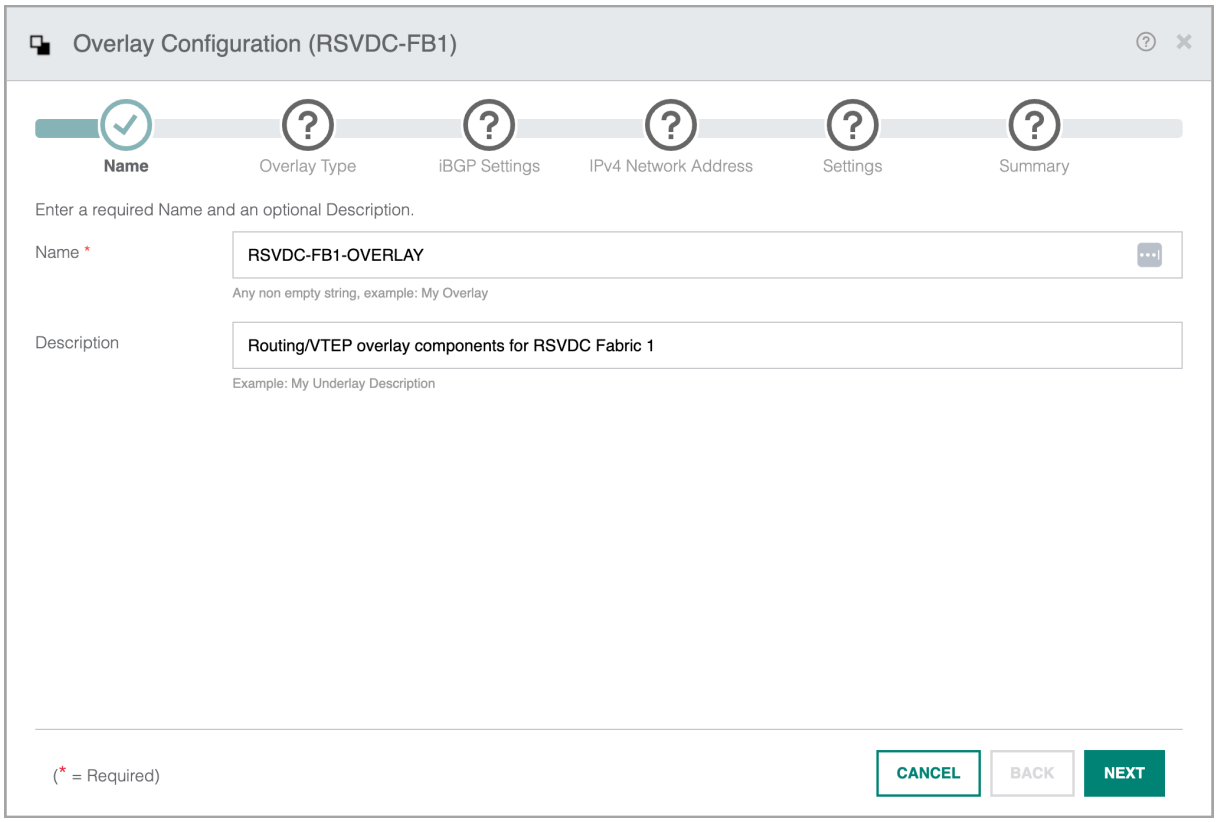


**Step 3** On the **Overlay Type** page, leave **iBGP** selected and click **NEXT**.

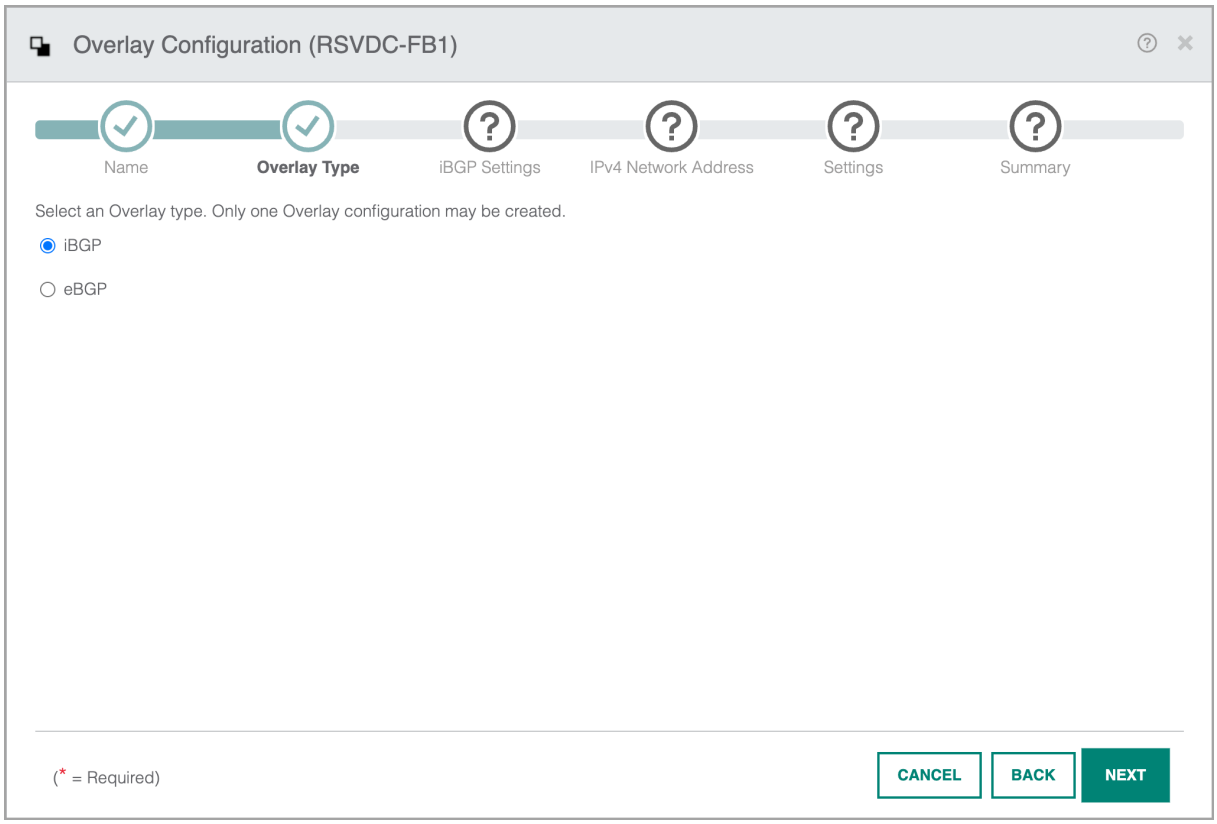


**Step 4** On the iBGP Settings page, enter the following settings, then click **NEXT**.

- **Spine-Leaf ASN:** *65001*
- **Route Reflector Servers:** < *Select two spine switches* >

- **Leaf Group Name:** *RSVDC-FB1-LF*
- **Spine Group Name:** *RSVDC-FB1-RR*

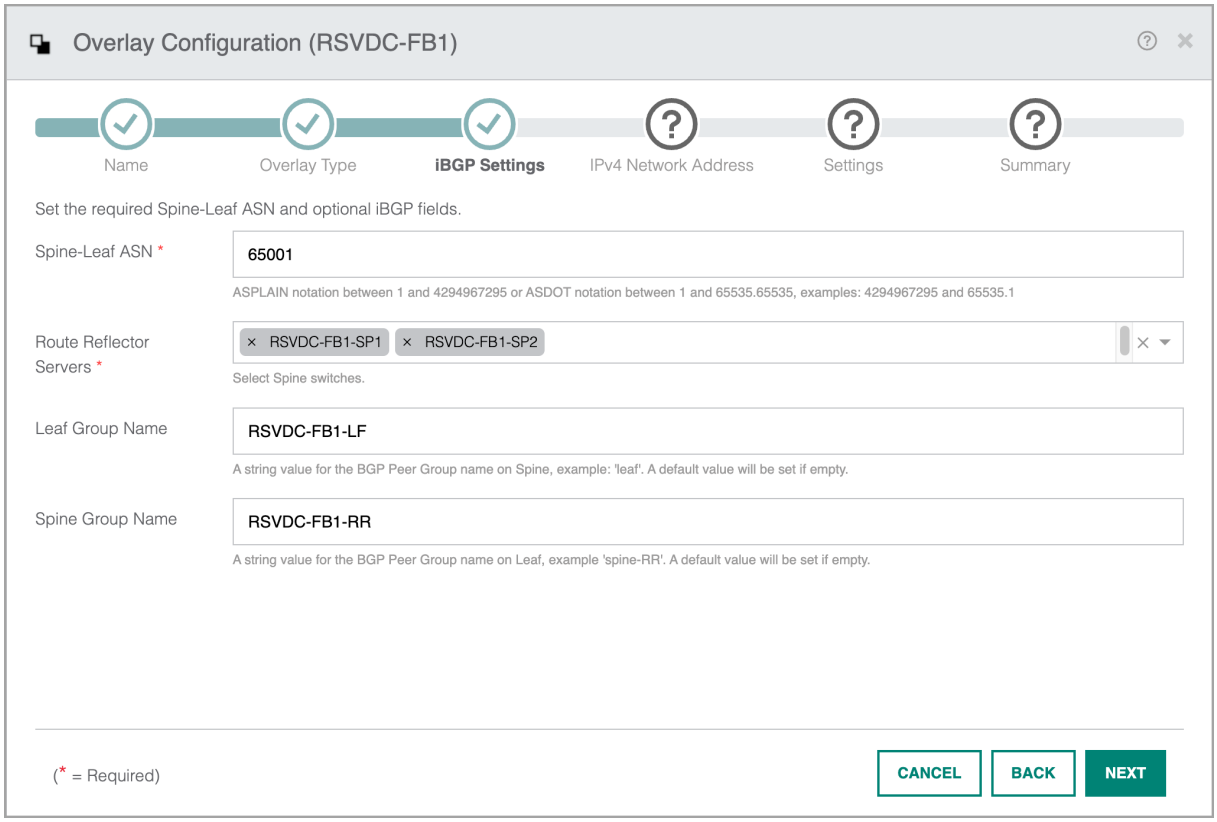


**NOTE:**

Use a 2-byte ASN in the private range of 64512-65534 for an easy-to-read switch configuration. A 4-byte ASN is supported.

**Step 5** On the **IPv4 Network Address** page, click **ADD** to launch the **Resource Pool** wizard.

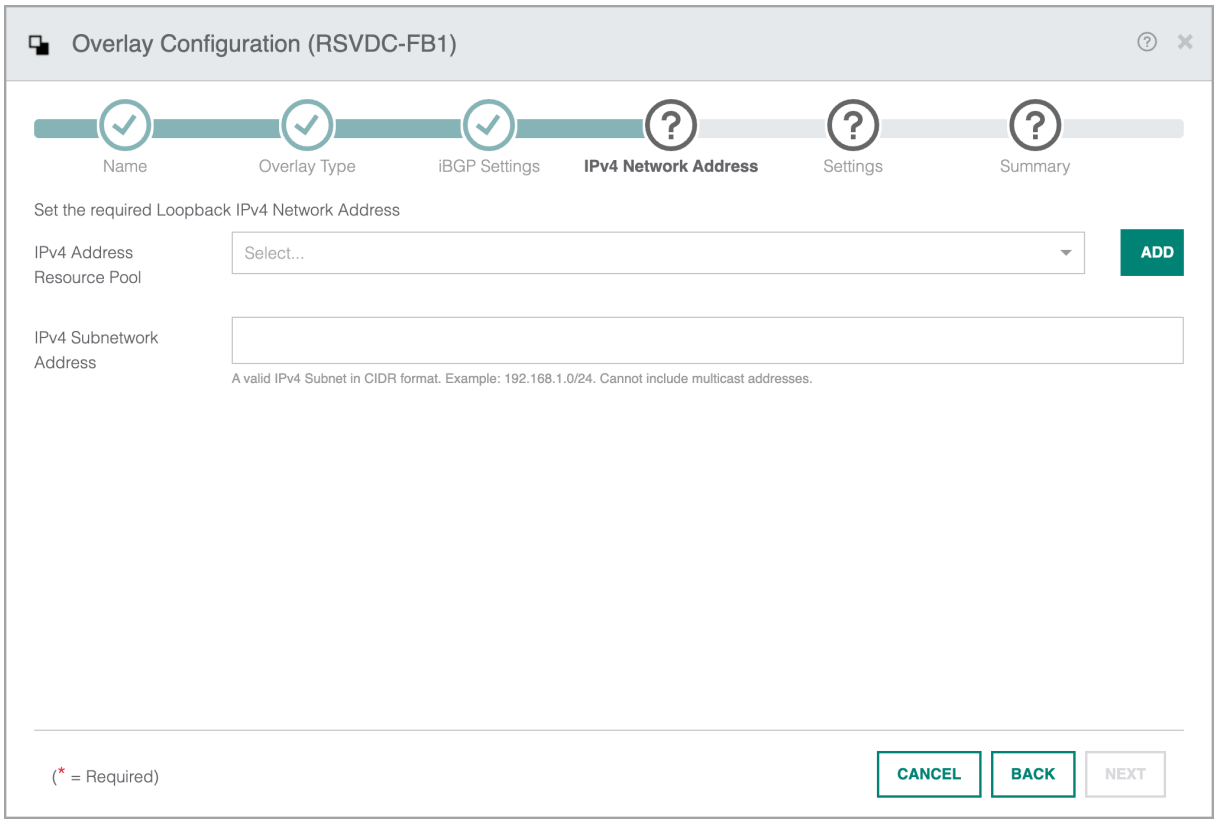


**Step 6 Resource Pool** wizard: On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

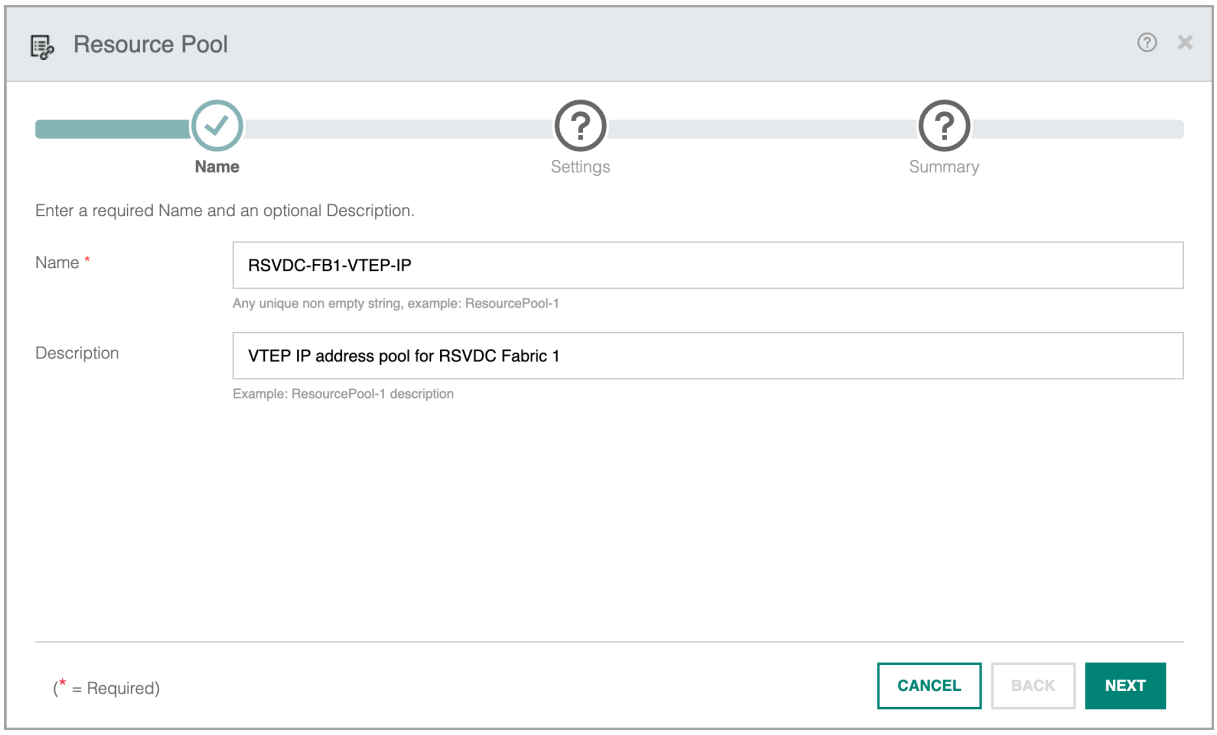


**Step 7 Resource Pool** wizard: On the **Settings** page, enter an IPv4 address block in the **Resource Type** field and click **NEXT**.

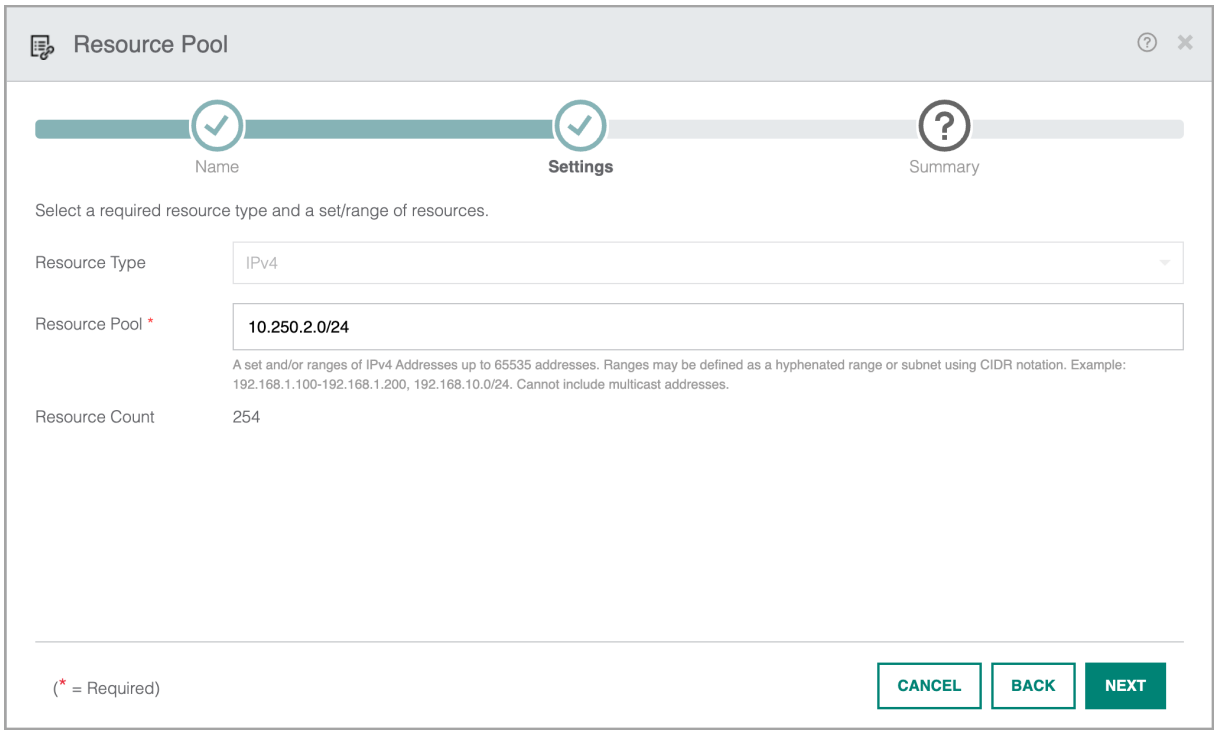


**NOTE:**

This IPv4 address block is used to configure loopback addresses on all leaf switches for VXLAN VTEPs. Each member of a VSX leaf pair uses the same IP loopback address.

**Step 8 Resource Pool** wizard: On the **Summary** page, verify the VTEP IP address pool information and click **APPLY**. The **Resource Pool** wizard closes and returns to the main **Overlay Configuration** workflow.

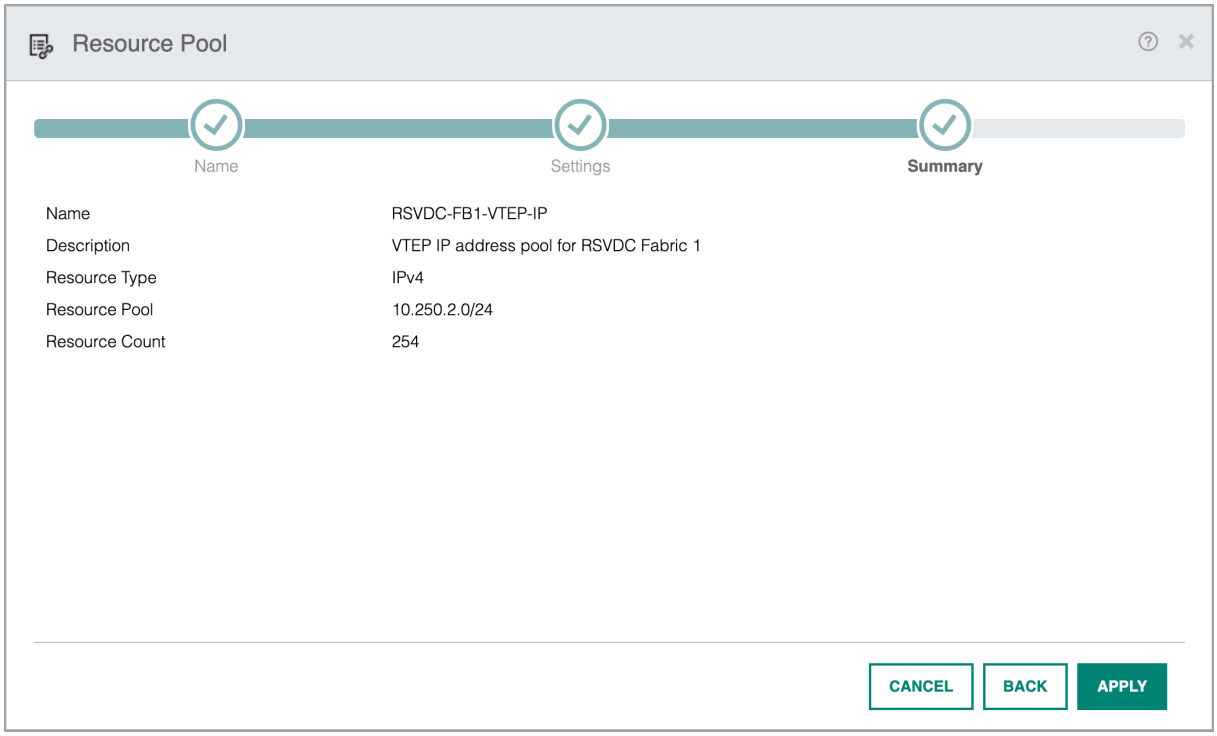

**Step 9** On the **IPv4 Network Address** page, verify that the new **IPv4 Address Resource Pool** is selected and click **NEXT**.

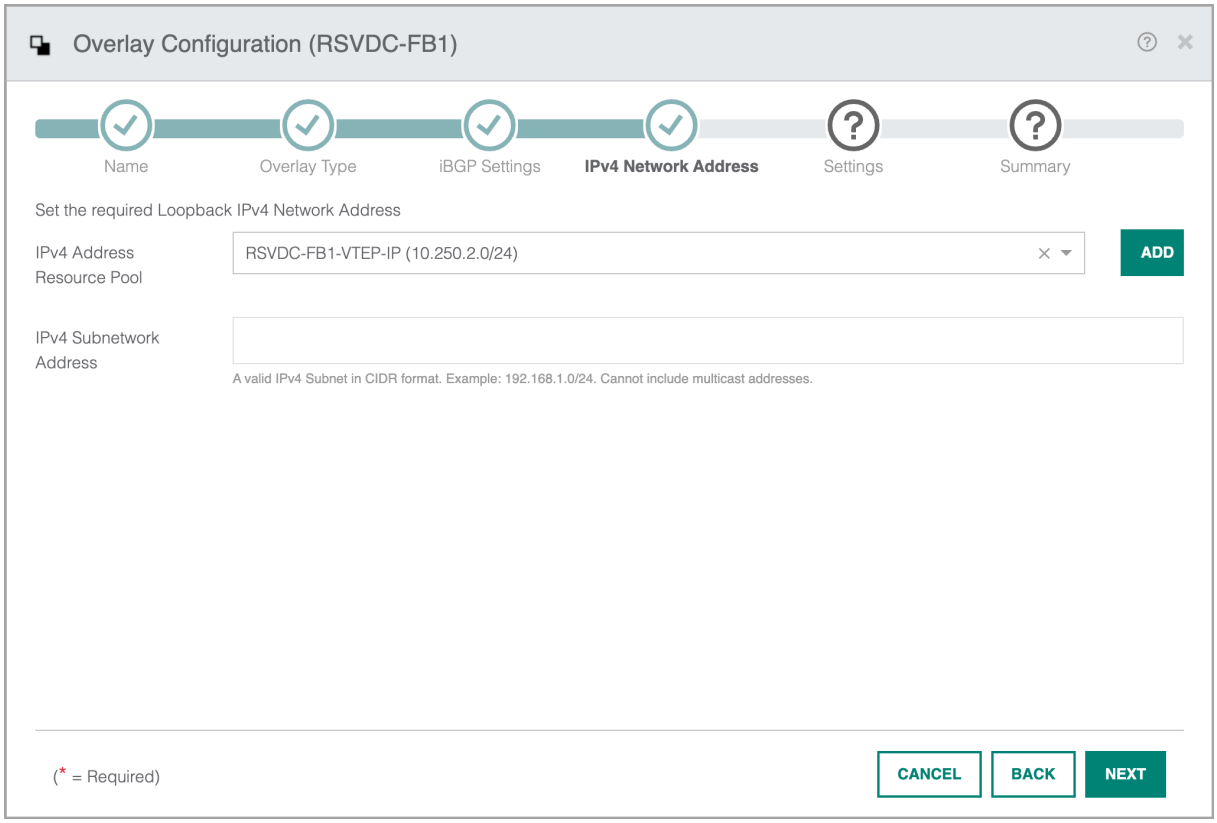


**Step 10** On the **Overlay Configuration Settings** page, leave the default values and click **NEXT**.

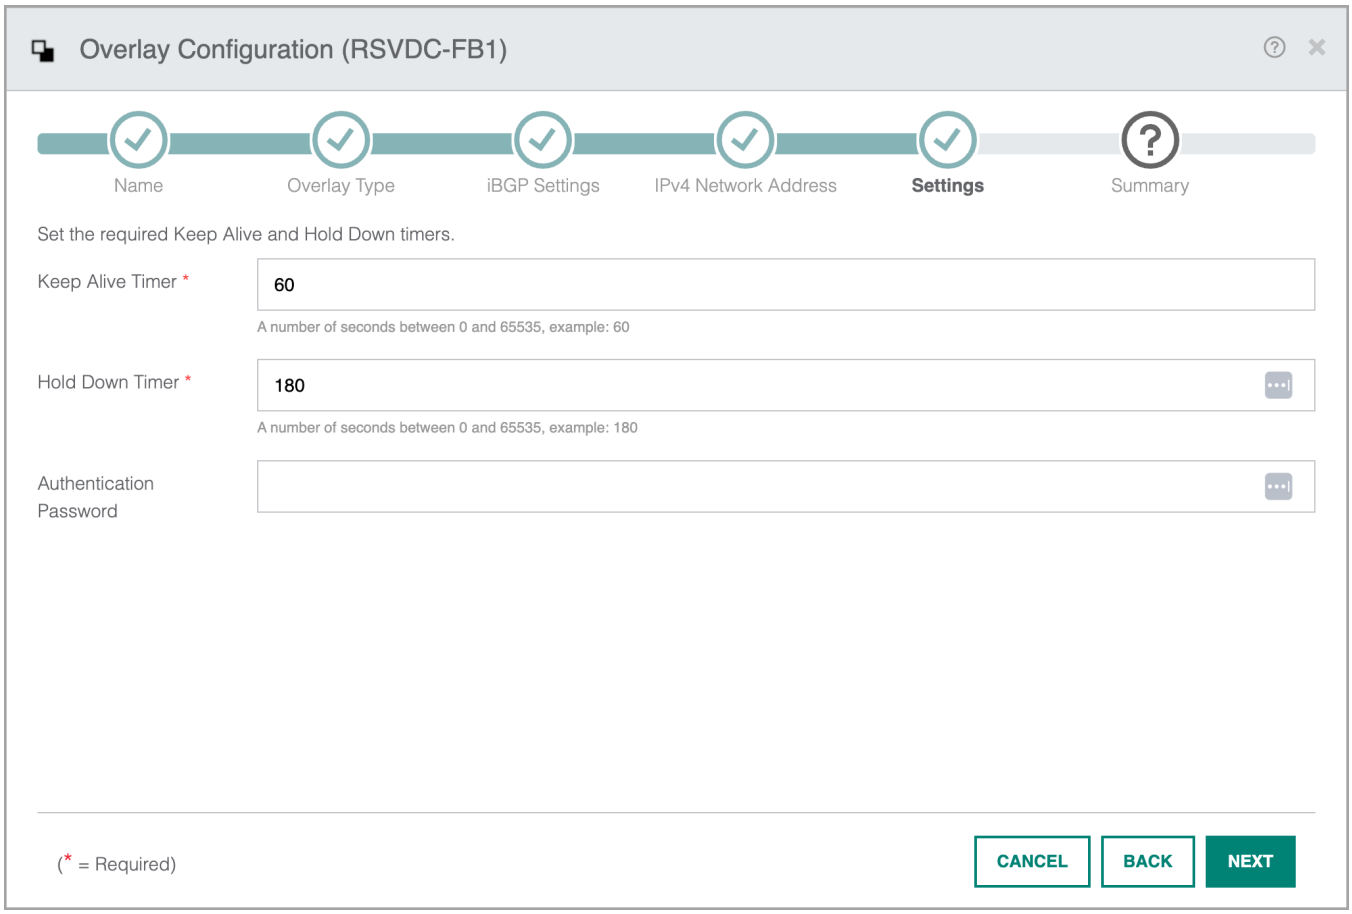


**Step 11** On the **Summary** page, verify that the iBGP information is correct, then click **APPLY**.

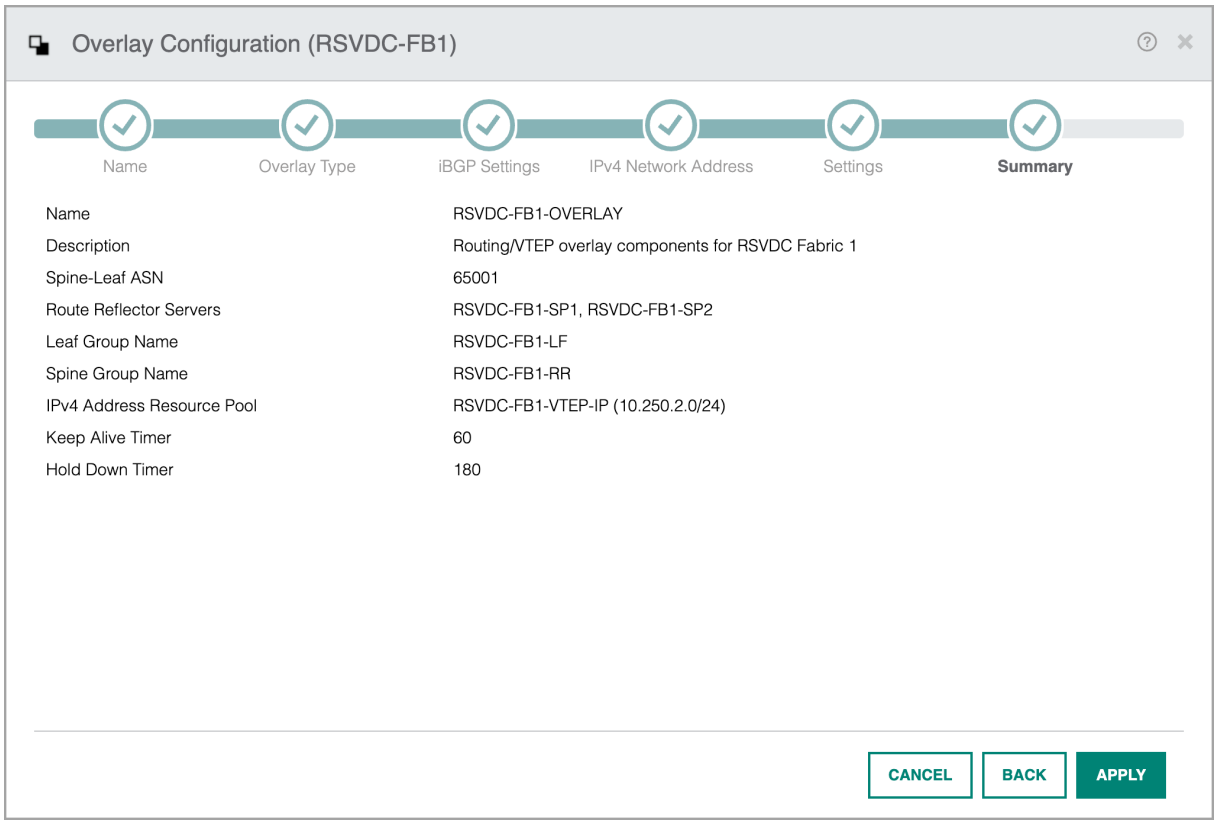

**Step 12** In the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 13** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select both route reflector spine switches* >
- **Commands:** *show bgp l2vpn evpn summary*

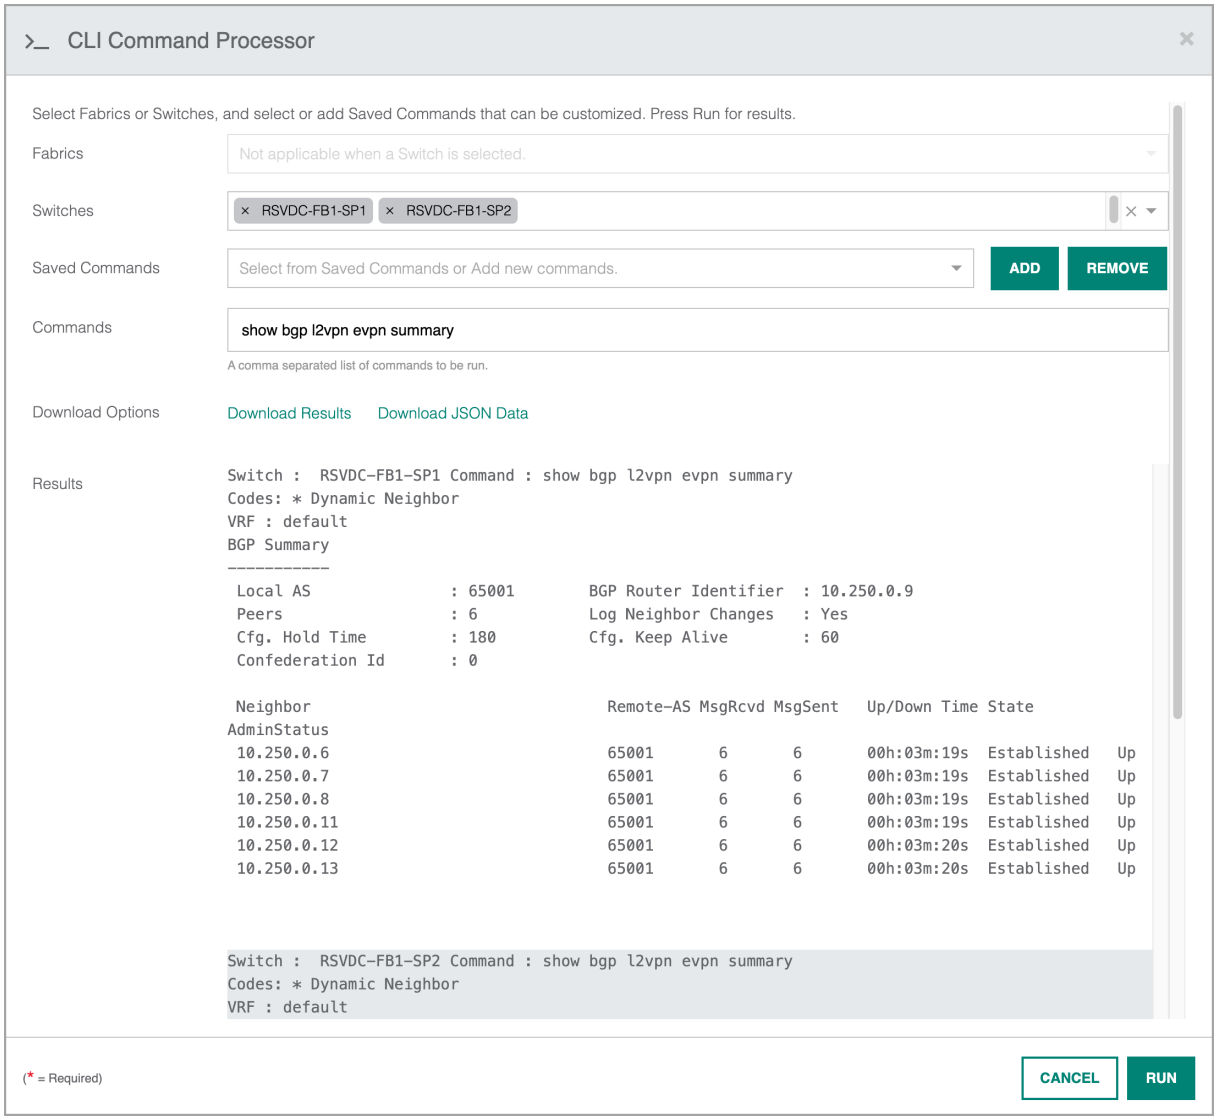


**Step 14** Verify that both route reflectors show an L2VPN EVPN neighbor relationship in the “Established” state for all leaf switches.

## Configure Overlay VRFs

An EVPN-VXLAN data center uses overlay VRFs to provide the Layer 3 virtualization and macro segmentation required for flexible and secure data centers. VRFs are distributed across all leaf switches. A VRF instance on one switch is associated to the same VRF on other leaf switches using a common L3 VNI and EVPN route-target, binding them together into one logical routing domain. VRFs are commonly used to segment networks by tenants and business intent.

Use the **Virtual Routing & Forwarding** workflow to create overlay network VRFs and associate a VRF with an L3 VNI and EVPN route-target. The VNI and route target for each set of overlay VRFs must be unique to preserve traffic separation.

This guide uses a production VRF and development VRF as an example of route table isolation. TCP/IP hosts in one VRF are expected to be isolated from hosts in the other VRF. The diagram below illustrates the logical VRF overlay across all leaf switches.

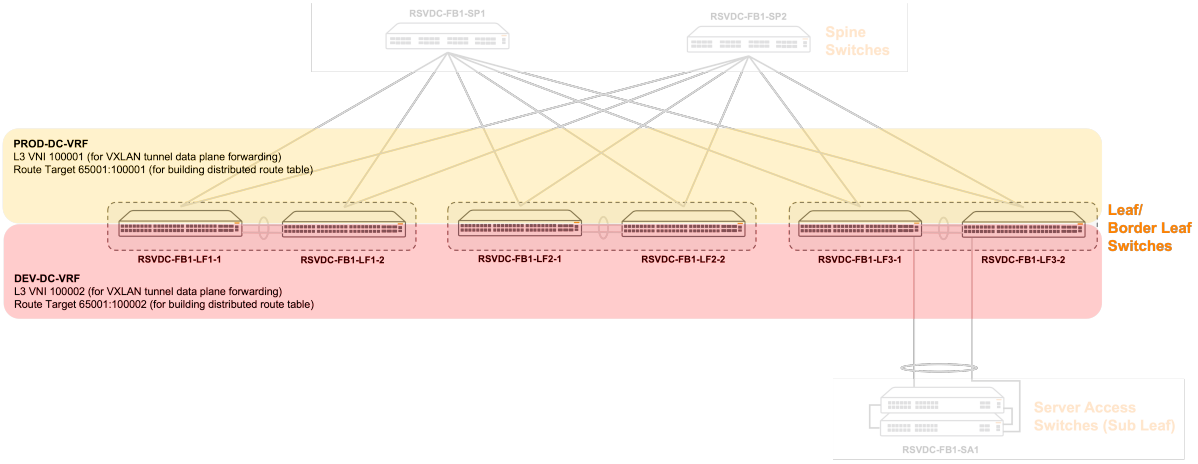


**NOTE:**

The diagram above depicts the border leaf switches at the same horizontal level as all other leaf switches. This placement of the border leaf pair is a cosmetic preference for easier depiction of virtualization across leaf switches. The deployed topology is consistent with previous diagrams, but without the pictorial emphasis of the special role of the border leaf handling data center north/south traffic.Hosts attached to server access switches can be connected to subnets in either VRF by VLAN extension from the leaf switch, but the server access switches do not contain their own VRF definition.

**Step 1** On the left menu, select **VRF**. If **VRF** does not appear in the left pane, select **Configuration > Routing > VRF** from the top menu.

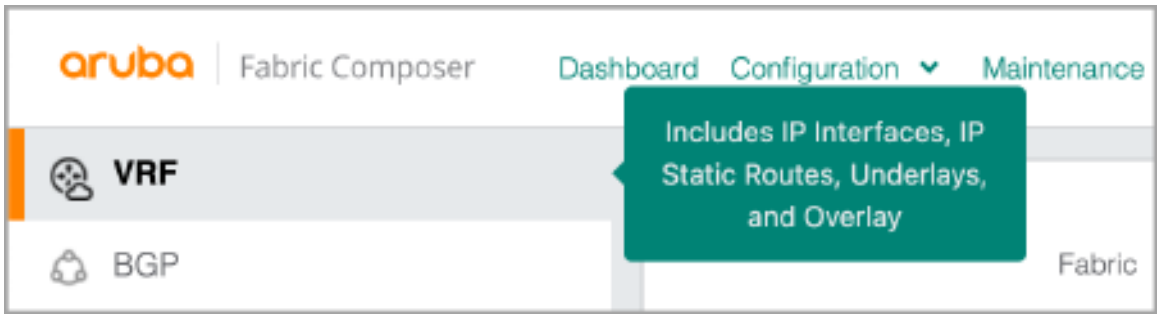


**Step 2** On the **ACTIONS** menu on the right, select **Add**.

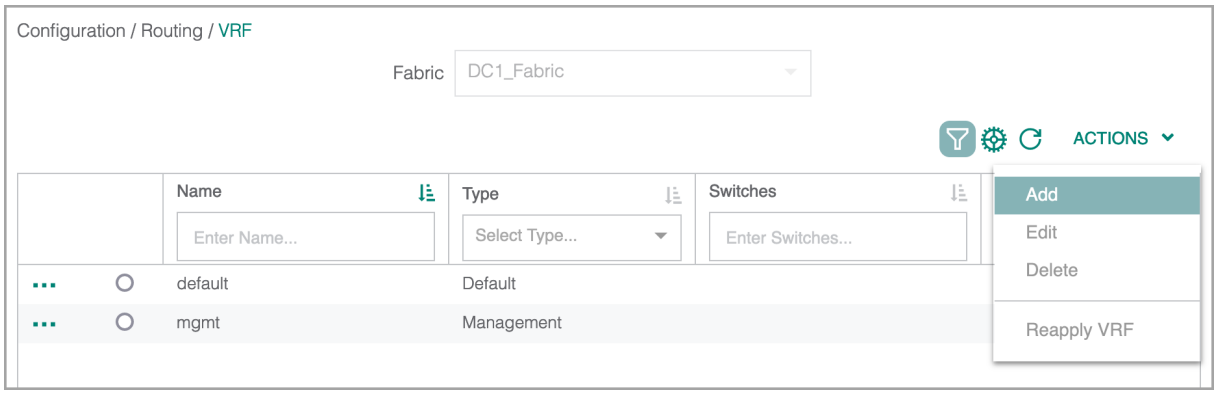


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

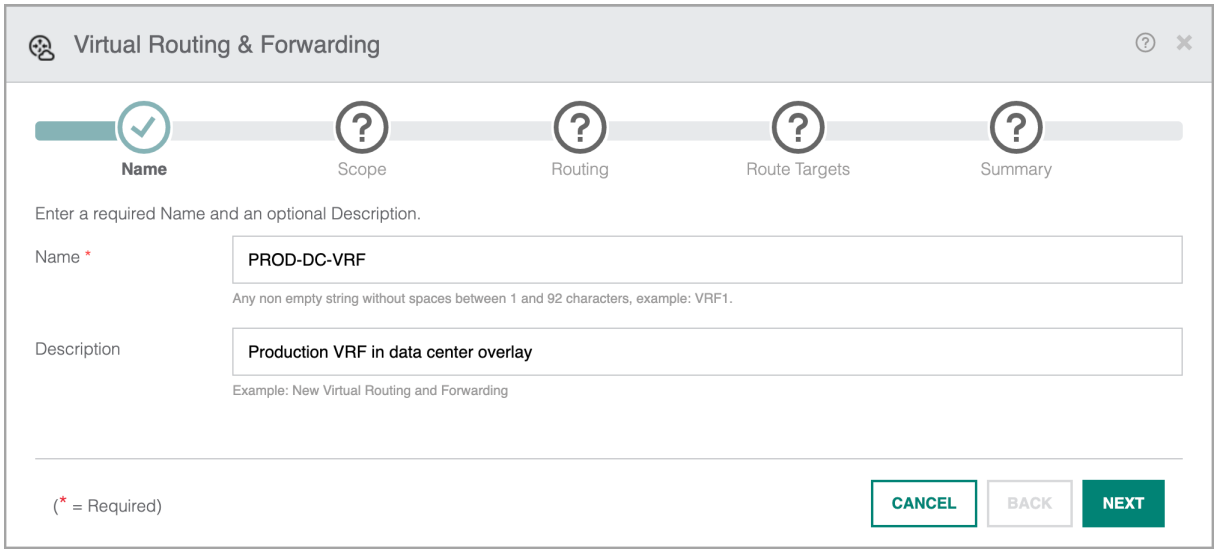


**Step 4** On the **Scope** page, uncheck **Apply the VRF to the entire Fabric and all Switches contained within it**. Select the VSX leaf pairs in the **Switches** field, then click **NEXT**.

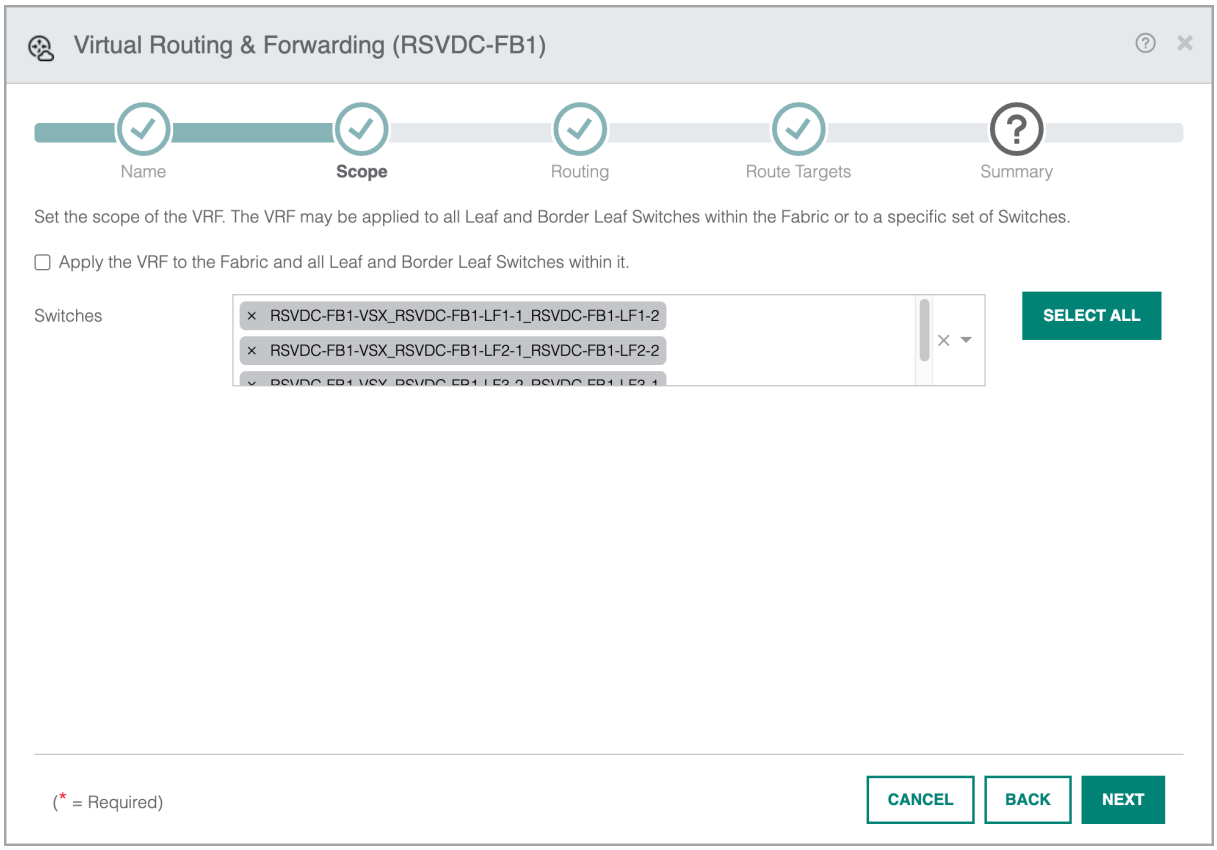


**NOTE:**

When a large number of leaf switches is present, click the **SELECT ALL** button to select all switches, then deselect spine and server access switches. Spine and server access switches do not participate in overlay virtualization and do not possess VTEPs, so overlay VRFs should not be configured on them.

**Step 5** On the **Routing** page, enter the following values to create a Layer 3 VNI and BGP route distinguisher * **L3 VNI**: 100001 * **Route Distinguisher**: *loopback1 : 1*

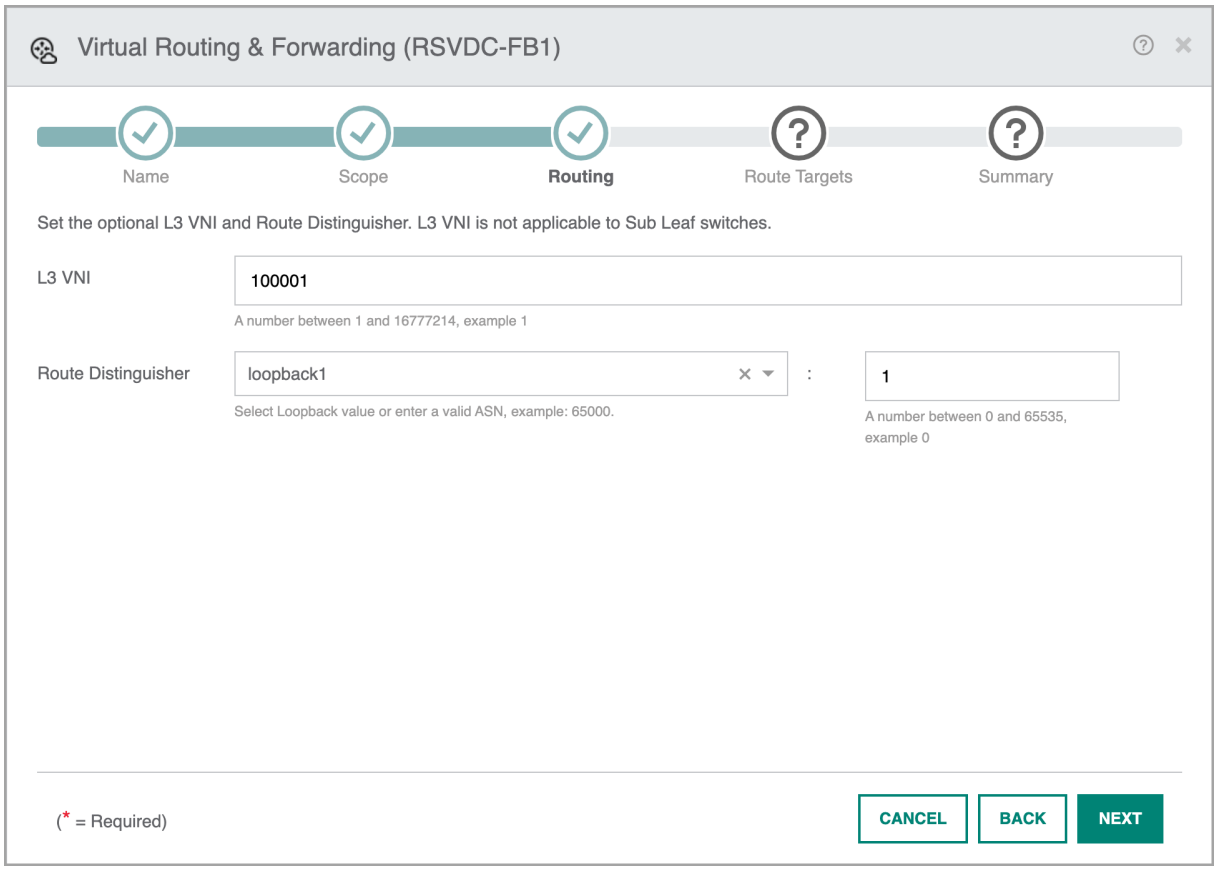


**NOTE:**

Refer to the "Overlay Connectivity and Addressing" section above for a VNI numbering reference. The Layer 3 VNI associates routes in an EVPN-VXLAN overlay with a VRF. The integer value in the **Route Distinguisher** should correlate to the VNI value without the addition of its 100,000 prefix for easier troubleshooting. The integer must be unique for each VRF.

**Step 6** On the **Virtual Routing & Forwarding Route Targets** page, assign the following settings to add an EVPN route-target to the VRF, then click **ADD**.

- **Route Target Mode:** *Both*
- **Route Target Ext-Community:** *65001:100001*
- **Address Family:** *EVPN*

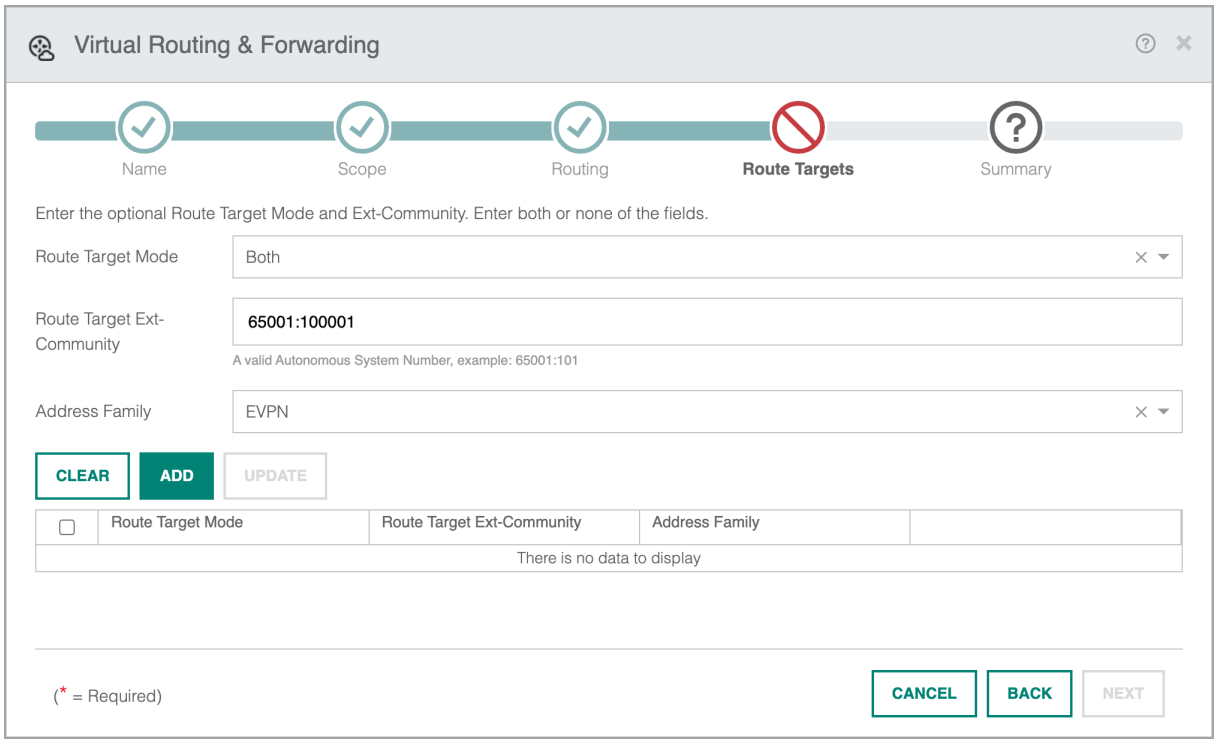


**NOTE:**

Setting **Route Target Mode** to *Both* exports local switch VRF routes to BGP with the **Route Target Ext-Community** value assigned as the route target and imports BGP routes into the local VRF route table advertised by other switches with the same value.For **Route Target Ext-Community**, enter the private autonomous system number used in the “Configure Overlay Network Routing” procedure and the L3 VNI, separated by a colon. The L3 VNI is used in the BGP route target for logical consistency with the VXLAN L3 VNI. The complete route target value uniquely identifies a set of VRFs.

**Step 7** Verify that the **Route Targets** information is correct and click **NEXT**.

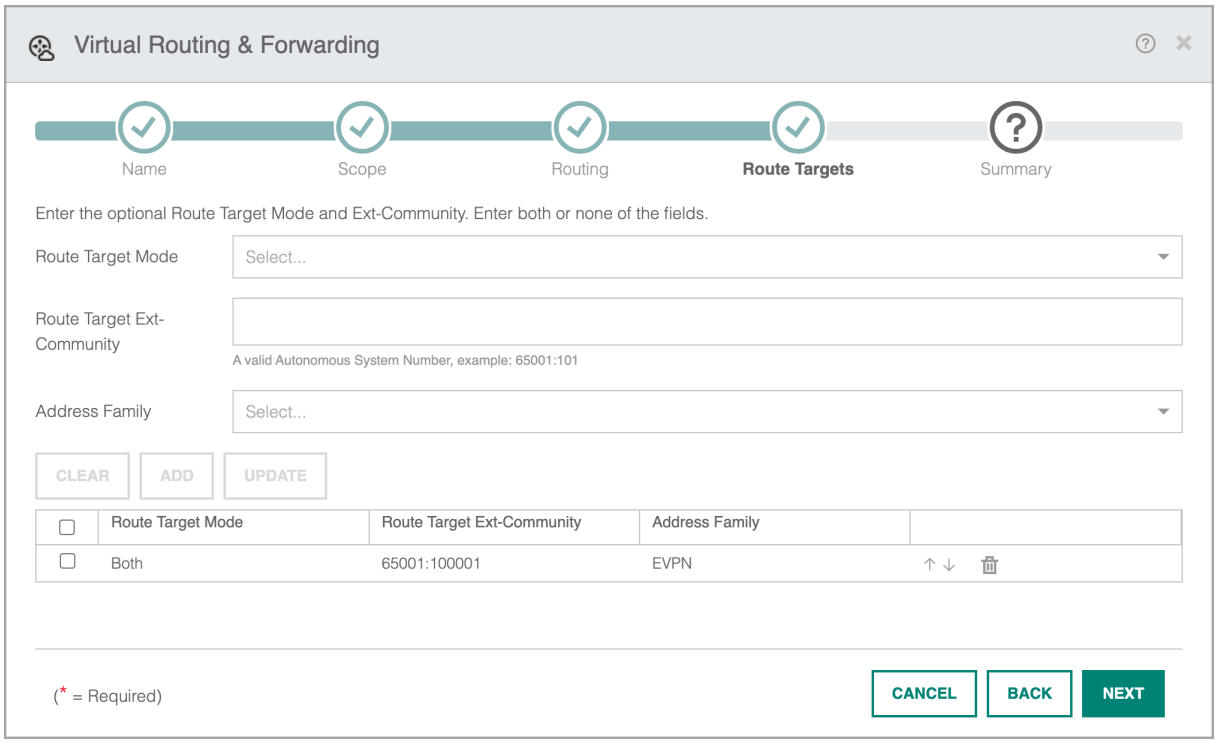


**Step 8** On the **Summary** page, verify that the complete set of VRF information is correct and click **APPLY**.

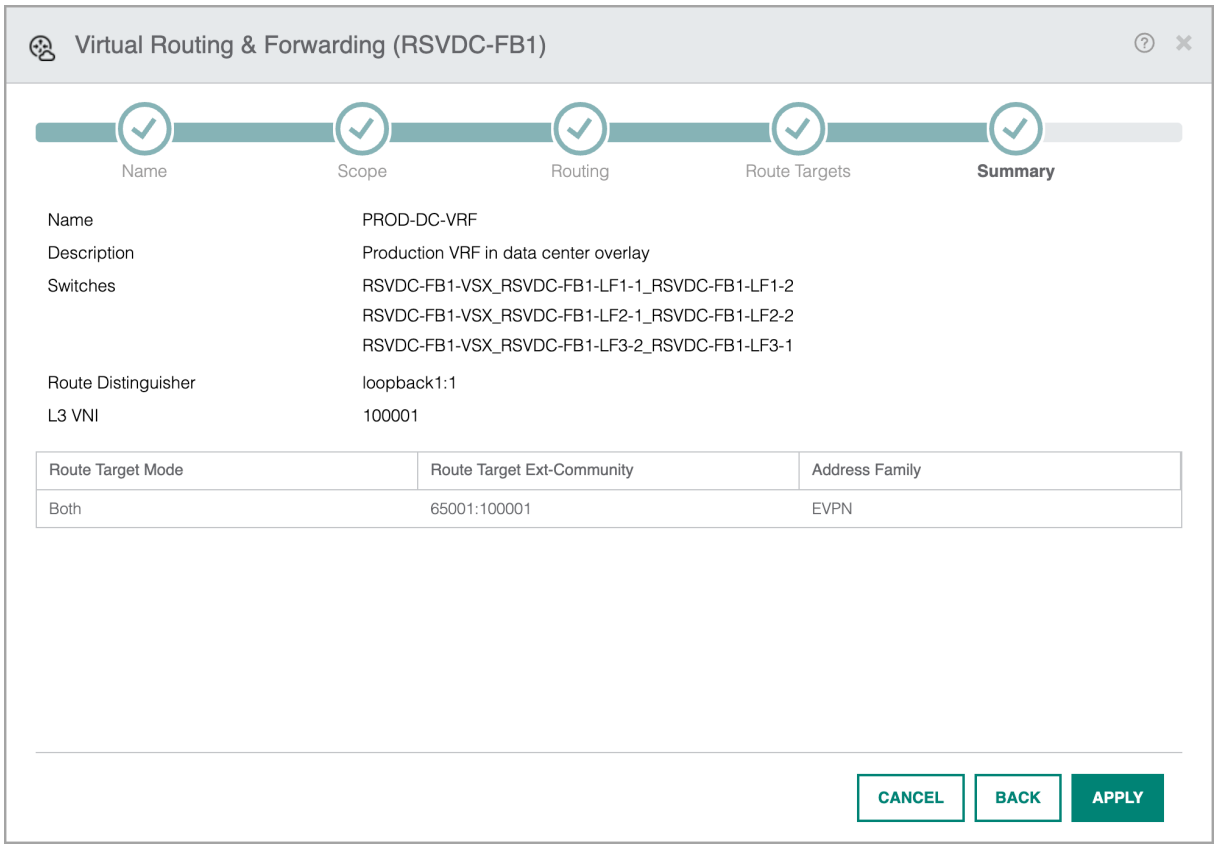


**Step 9** Repeat this procedure for each additional overlay VRF.

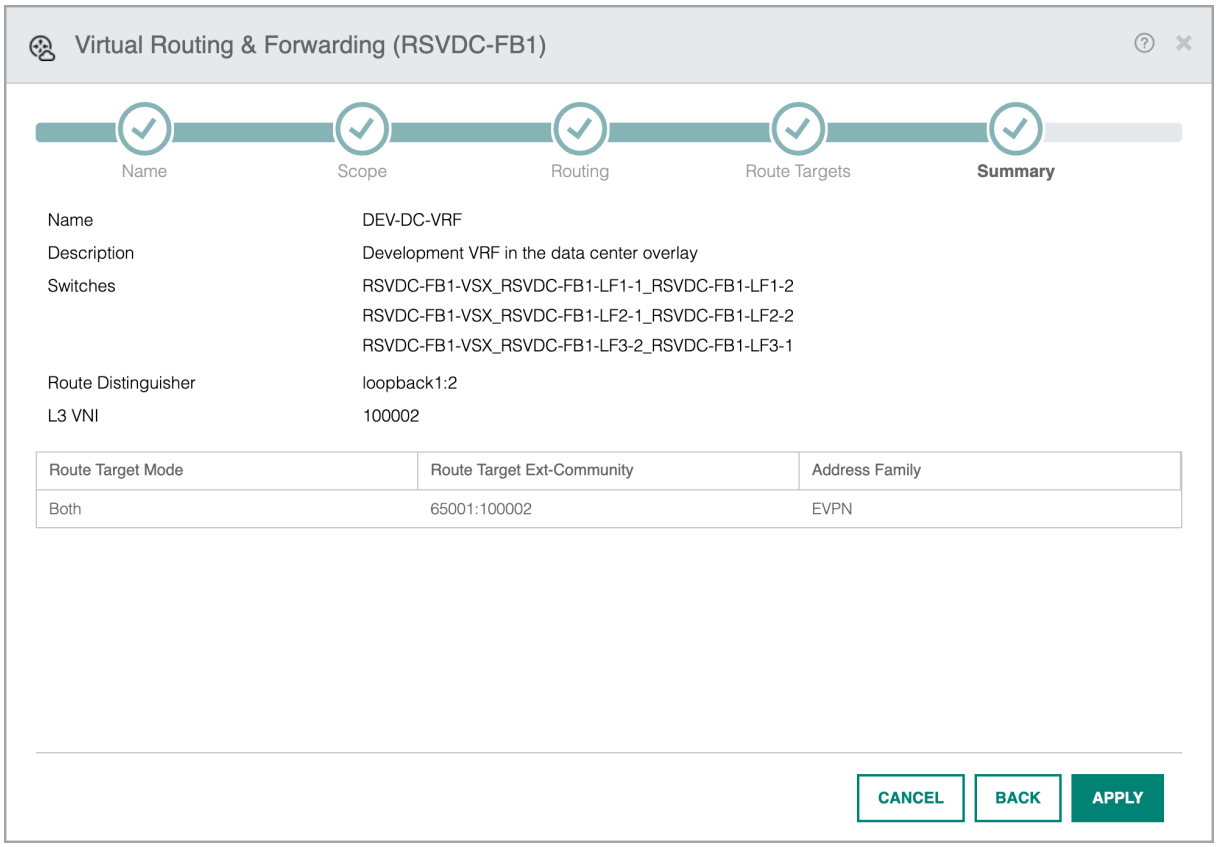


## Configure Overlay VLANs and SVIs

One or more VLANs within each VRF provide host connectivity. VLAN SVIs provide IP addressing within the fabric. The Fabric Composer **IP Interface** workflow creates consistent VLANs across all leaf switches within an overlay VRF. The workflow assigns an SVI IP address, a virtual gateway address, and a locally administered virtual MAC address to the VLAN interface on each leaf switch. Aruba Active Gateway permits the SVI IP and virtual gateway to be used on VSX leaf pairs.

The creation of VLANs and SVIs in this step is prerequisite to binding the VLANs across racks in logically contiguous Layer 2 domains in the next procedure. At the end of this procedure, each VLAN's broadcast domain is scoped to each VSX pair.

CX 10000 switches positioned in a border leaf role support east-west policy for attached hosts, when the switches are assigned a **leaf** switch profile. When IPsec or NAT features are enabled, the CX 10000 border leaf must be configured with a **spine** switch profile. After assigning the CX 10000 a **spine** switch profile, east-west policy enforcement is no longer supported, and directly attaching hosts to the border leaf is not recommended. This limitation does not apply to other switch models in the border leaf role. In this guide, CX 10000 switches positioned at the border leaf are not configured with host VLANs to support enabling IPsec in a separate procedure.

The diagram below illustrates the creation of VLANs on ToR VSX leaf pairs, except the CX 10000 border leaf.

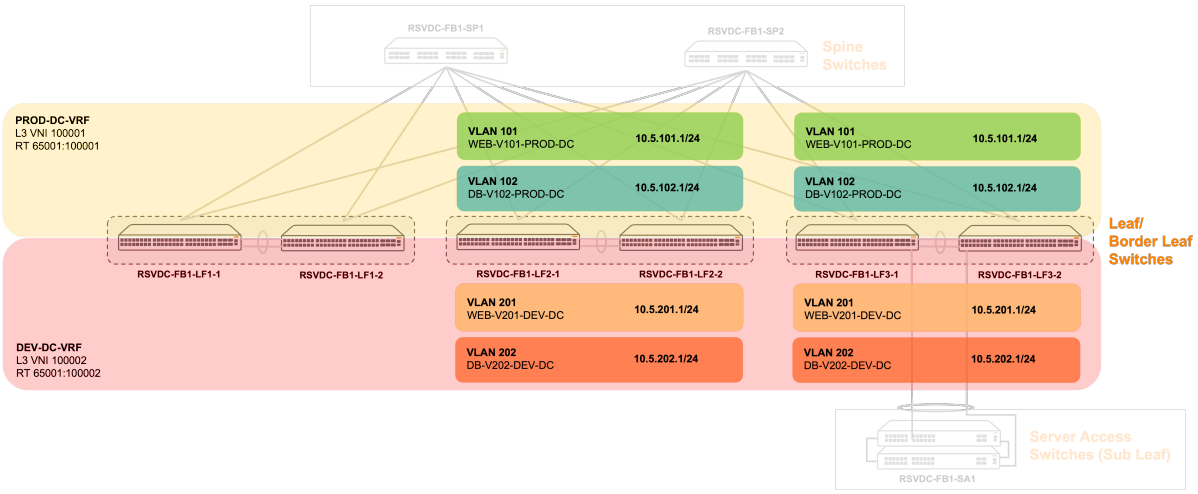


**Step 1** Confirm that the view is set to **Configuration/Routing/VRF**, then click the • • • symbol next to **PROD-DC-VRF** and select **IP Interfaces**.

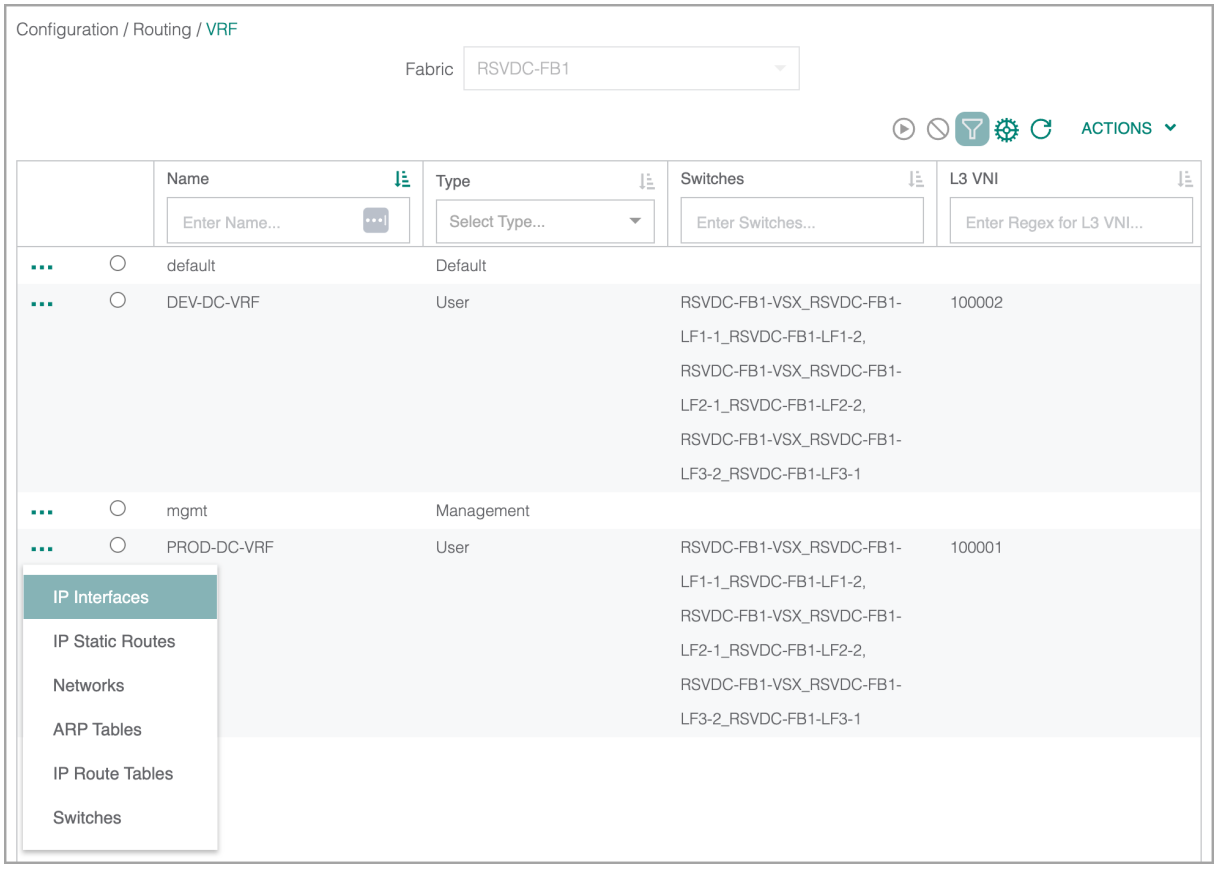


**NOTE:**

The • • • symbol is a shortcut to most options in the **ACTIONS** menu. This shortcut method is available in many Fabric Composer contexts. The IP Interfaces context also can be viewed by clicking the **PROD-DC-VRF** radio button and selecting **IP Interfaces** on the **ACTIONS** menu.

**Step 2** On the **Configuration/Routing/VRF/PROD-DC-VRF** page, select the right **ACTIONS** menu below **IP INTERFACES** and click **Add**.

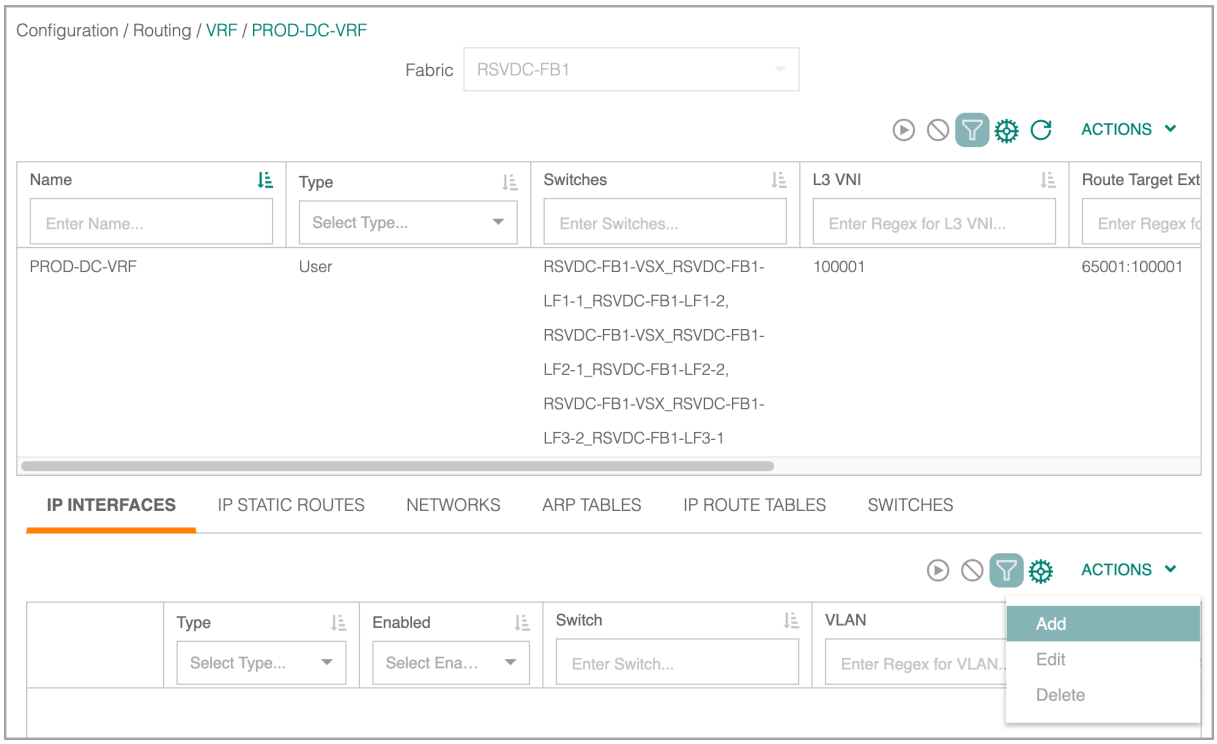


**Step 3** On the **IP Interfaces** page, assign the following values, then click **NEXT**.

- **Type:** *SVI*
- **VLAN:** *101*
- **Switches:** *< Select all leaf switches, except CX 10000 border leaf>*
- **IPv4 Subnetwork Address:** *10.5.101.0/24*
- **Switch Addresses:** *10.5.101.1*
- **Active Gateway IP Address:** *10.5.101.1*
- **Active Gateway MAC Address:** *02:00:0A:05:00:01*

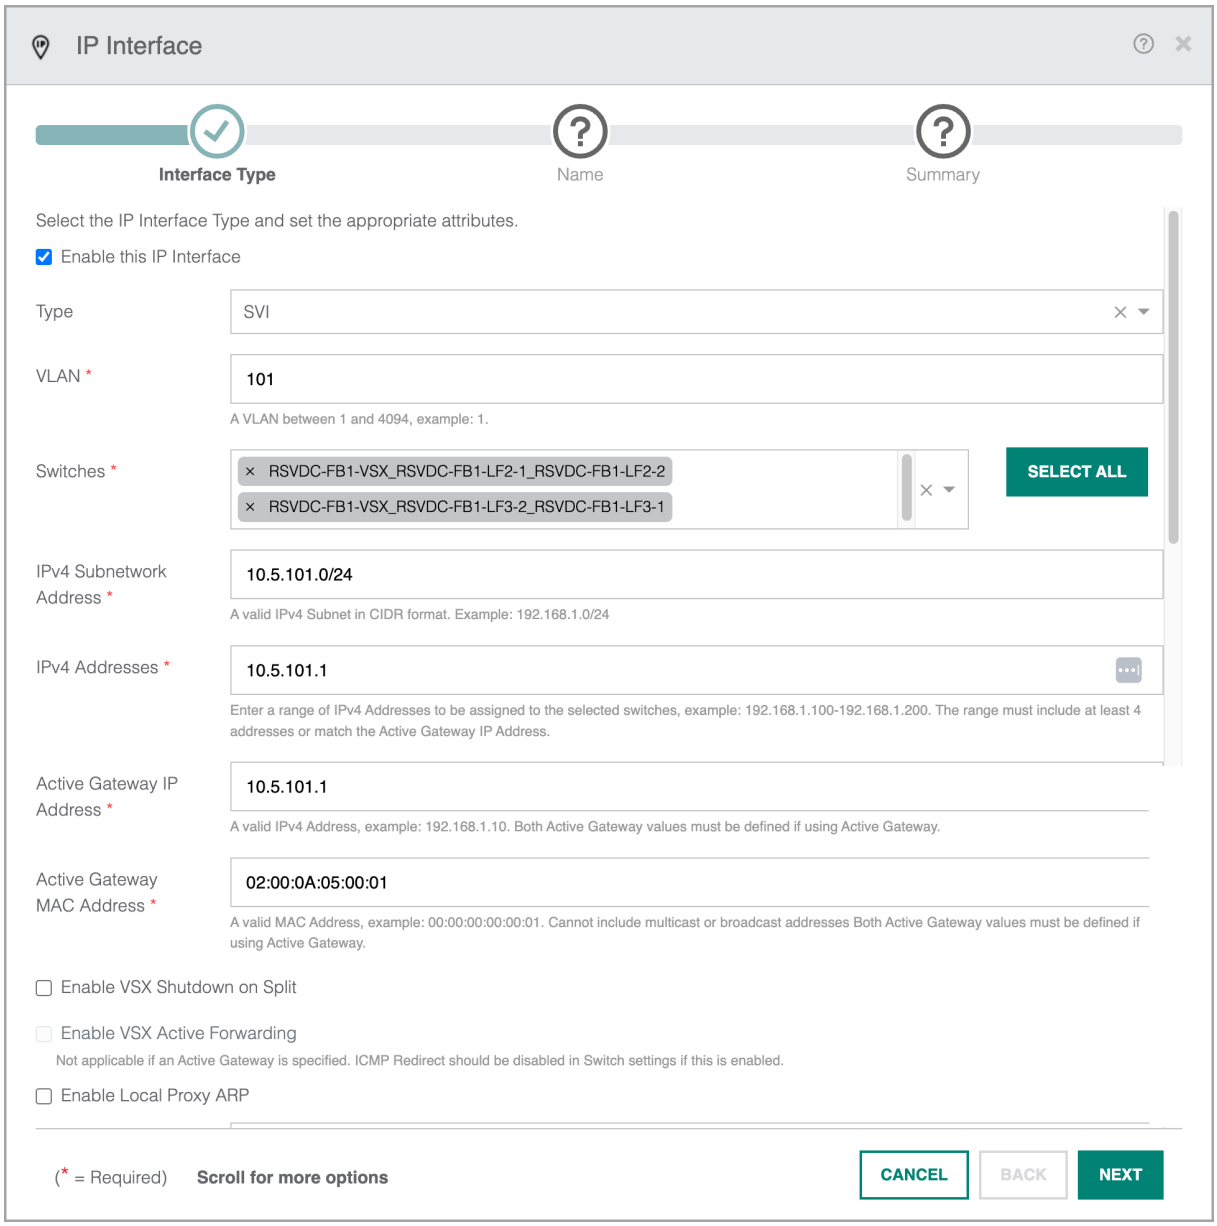


**NOTE:**

The **SELECT ALL** button selects all switches assigned to the VRF where the SVI interface will be created. The range provided for **IPv4 Addresses** and the **Active Gateway IP Address** must be from the same network range as the **IPv4 Subnetwork Address**. The **IPv4 Addresses** field value is used to assign an IP address to each SVI interface. AOS-CX 10.09 and above supports assigning the same IP address as both the SVI interface and the active gateway. This maximizes the number of IPs available to assign to attached network hosts.

**NOTE:**

The Active Gateway IP address is not supported as a source IP address when using the ping command. When assigning the same IP address to both the Active Gateway and VLAN SVI, the ping command must specify a unique source interface or IP address, such as a loopback assigned to the same VRF, to verify reachability.For example:`# ping 10.5.101.11 vrf PROD-DC-VRF source loopback11`

**Step 4** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

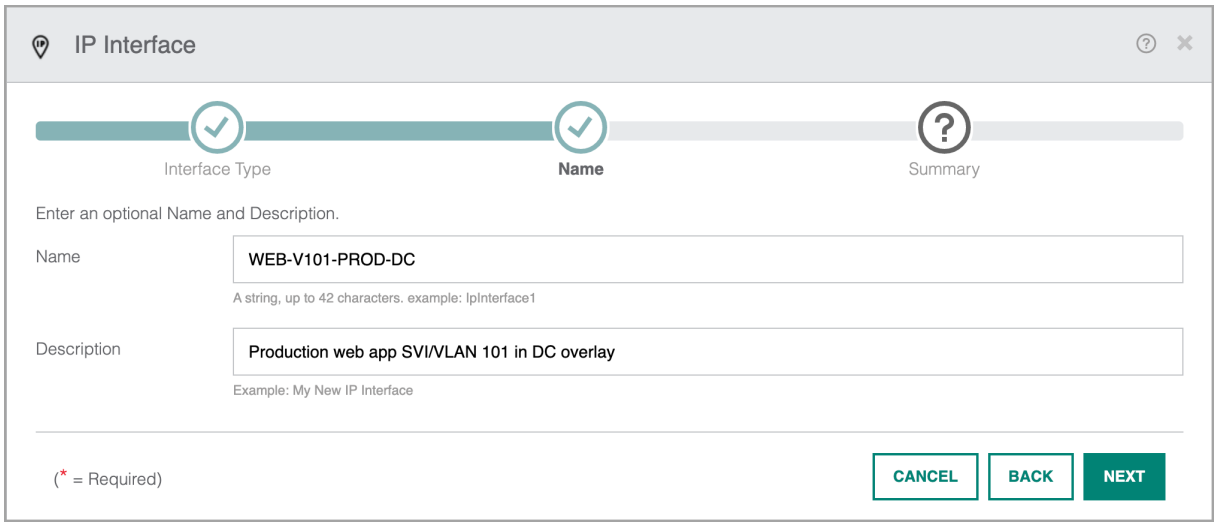


**NOTE:**

Including the associated VLAN ID and overlay VRF in the **Name** can be helpful during management operations.

**Step 5** On the **Summary** page, verify that the information is entered correctly and click **APPLY**.

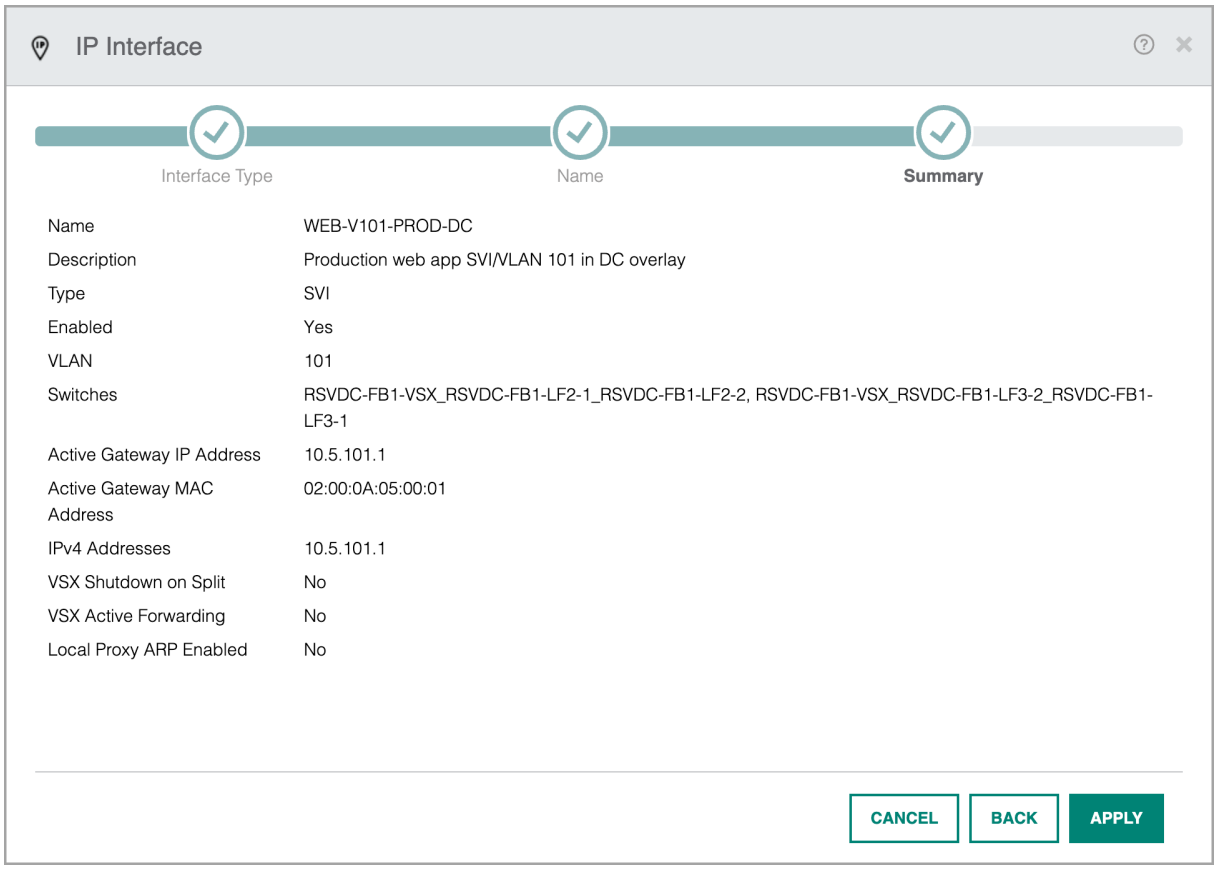


**Step 6** Repeat the procedure to create an additional overlay subnet in the production VRF using the following values:

| Name | Description | Type | VLAN | Switches | IPv4 Sub-network Address | IPv4 Addresses | Active Gateway IP Address | Active Gateway MAC Address |
|---|---|---|---|---|---|---|---|---|
| DB-V102-PROD-DC | Production database SVI/VLAN 102 DC overlay | SVI | 102 | < All non-border leaf switches > | 10.5.102.0/ | 10.5.102 | 10.5.102.1 | 02:00:0A:05:00:01 |

**Step 7** Repeat the procedure to create additional overlay subnets in the development VRF using the following values:

| Name | Description | Type | VLAN | Switches | IPv4 Sub-network Address | IPv4 Address | Active Gateway IP Address | Active Gateway MAC Address |
|---|---|---|---|---|---|---|---|---|
| WEB-V201-DEV-DC | Development web app SVI/VLAN 201 in DC overlay | SVI | 201 | < All non-border leaf switches > | 10.6.201.0/ | 10.6.2( | 10.6.201.1 | 02:00:0A:06:00:01 |
| DB-V202-DEV-DC | Development database SVI/VLAN 202 in DC overlay | SVI | 202 | < All non-border leaf switches > | 10.6.202.0/24 | 10.6.202.1 | 10.6.202.1 | 02:00:0A:06:00:01 |

**NOTE:**

Host connectivity can be extended to border leaf switches, when not using IPsec or NAT services on CX 10000 switches in the border leaf role.

## Configure EVPN Instances

An EVPN instance joins each previously created VLAN across leaf switches into a combined broadcast domain. This procedure defines two key attributes to logically bind each VLAN across the leaf switches. A VNI is assigned to each VLAN. MP-BGP associates host MACs to VNI values in its EVPN host advertisements to support VXLAN tunneling. An auto-assigned route target per VLAN also is defined. The VLAN route-target associates a MAC address with the appropriate VLAN at remote switches for the purpose of building bridge table MAC reachability. Route targets are included in MP-BGP EVPN host advertisements.

The Fabric Composer EVPN wizard maps VLAN IDs to L2 VNI values. A prefix value is provided for automatic generation of route targets. The EVPN wizard also creates an EVPN instance to associate route targets with VLANs. When using iBGP for the overlay control plane protocol, route targets can be assigned automatically. A resource pool is used to assign the EVPN system MAC addresses.

At the completion of this procedure, distributed L2 connectivity across leaf switches in the fabric is established, with the exception of the border leaf. Aruba active gateway permits the same IP address to be used on all leaf switches in the fabric for a VLAN. Overlay reachability between the border leaf and other leaf switches is routed. The diagram below illustrates the the logical binding of VLANs across leaf racks into logically contiguous broadcast domains.

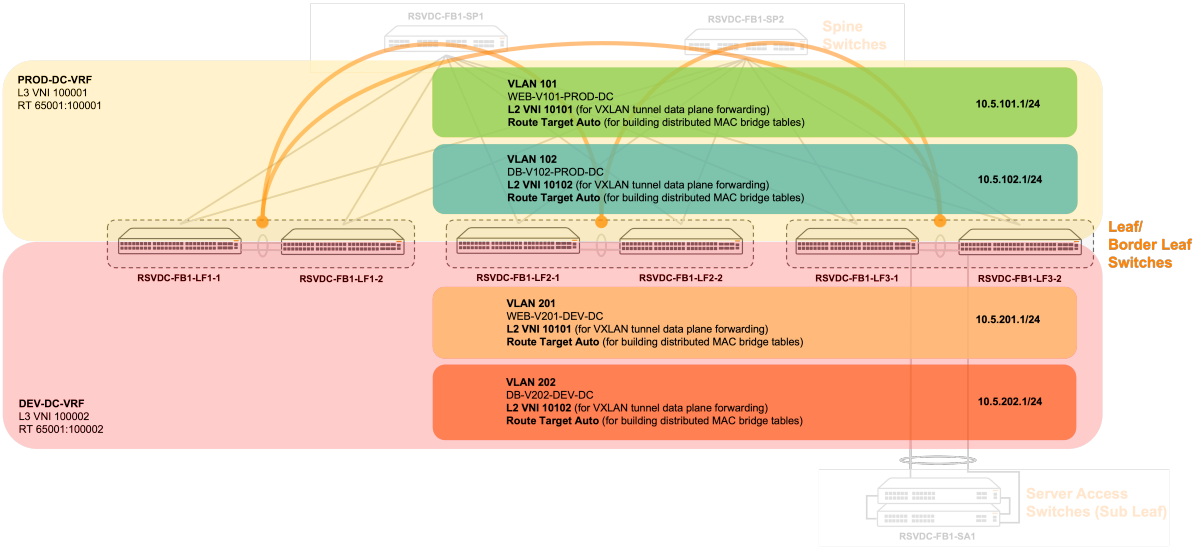


**Step 1** On the **Guided Setup** menu, select **EVPN CONFIGURATION** to start the **EVPN** workflow.

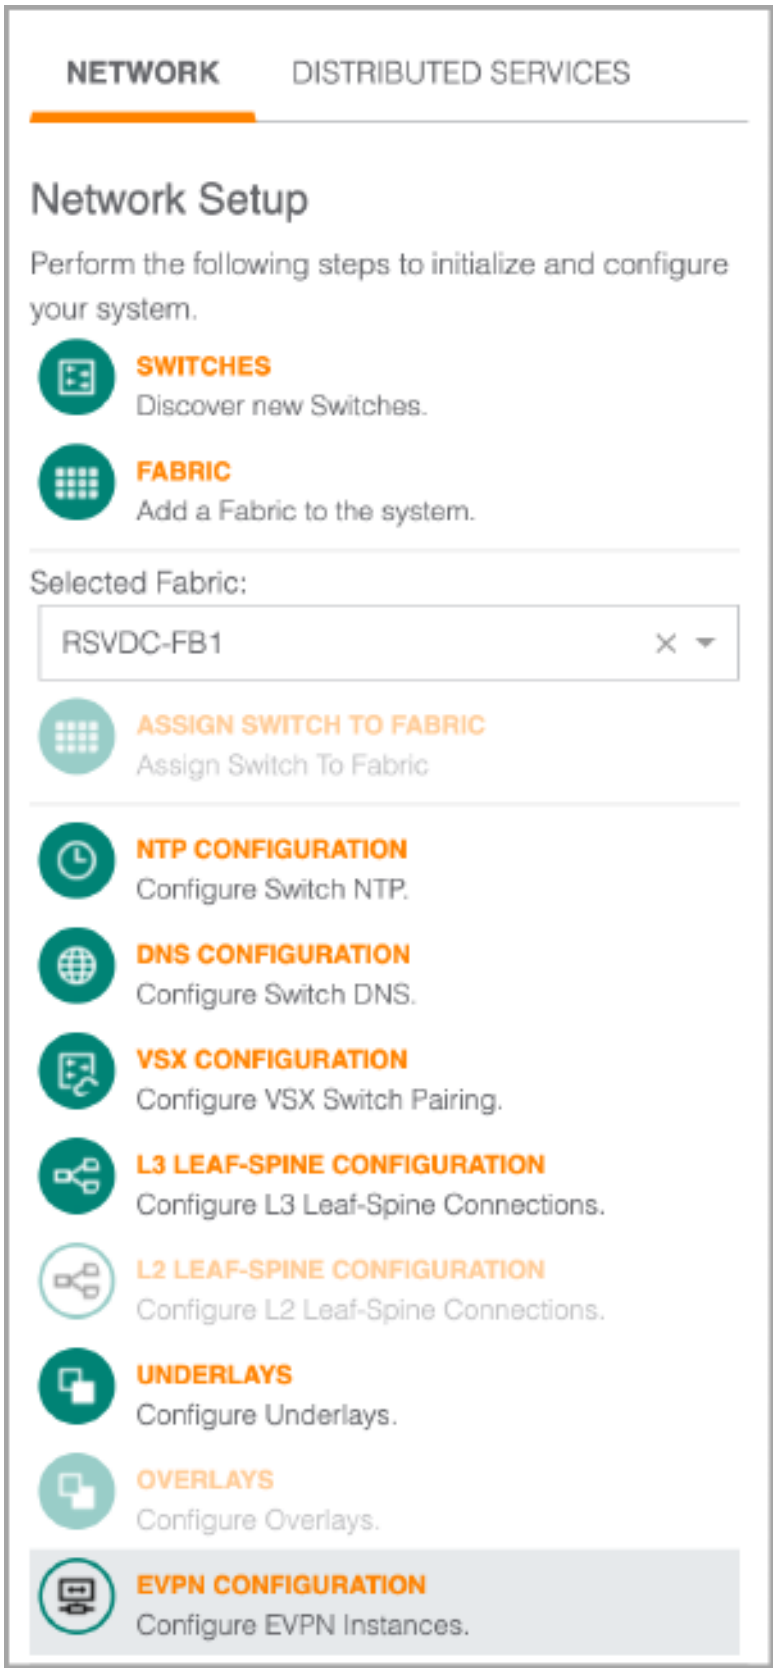


**Step 2** On the **Introduction** page, review the guidance and click **NEXT**.

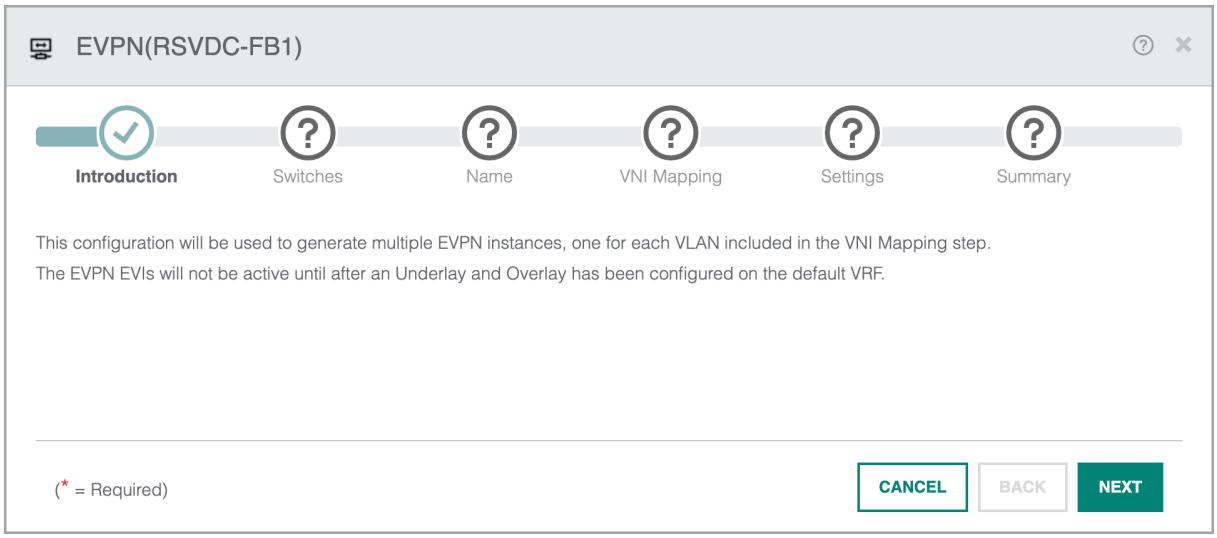


**NOTE:**

The prerequisites noted above were completed in previous steps.

**Step 3** On the **Switches** page, uncheck **Create EVPN instances across the entire Fabric and all Switches contained within it**, select all leaf switches except the border leaf, and click **NEXT**.

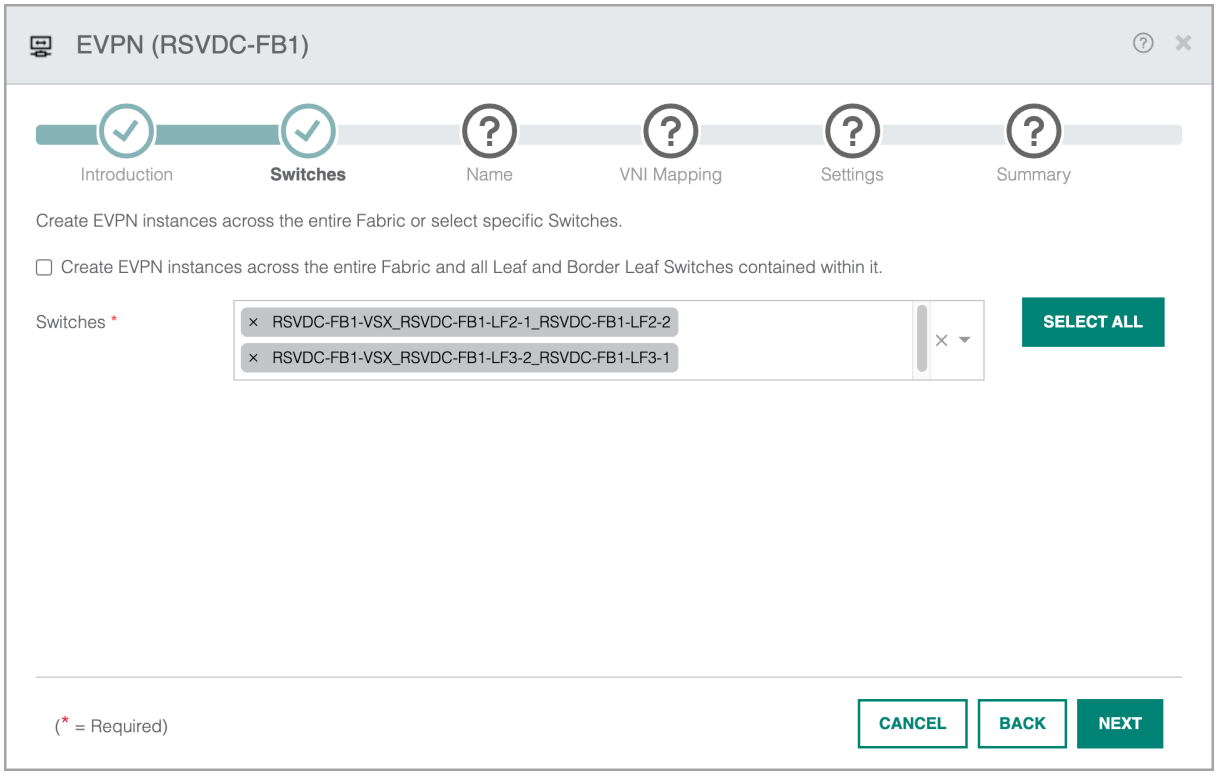


**Step 4** On the **Name** page, enter a **Name Prefix** and **Description**, then click **NEXT**.

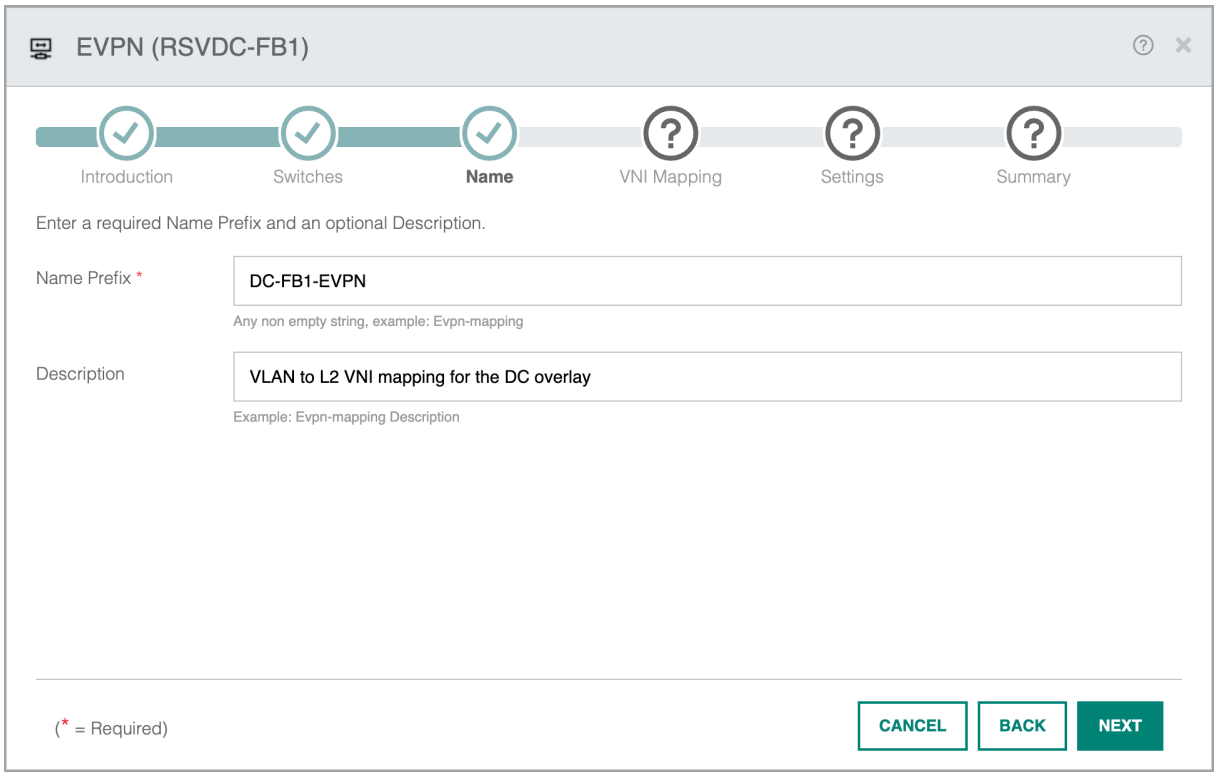


**Step 5** On the **VNI Mapping** page, enter one or more **VLANs** and a **Base L2VNI**, then click **NEXT**.

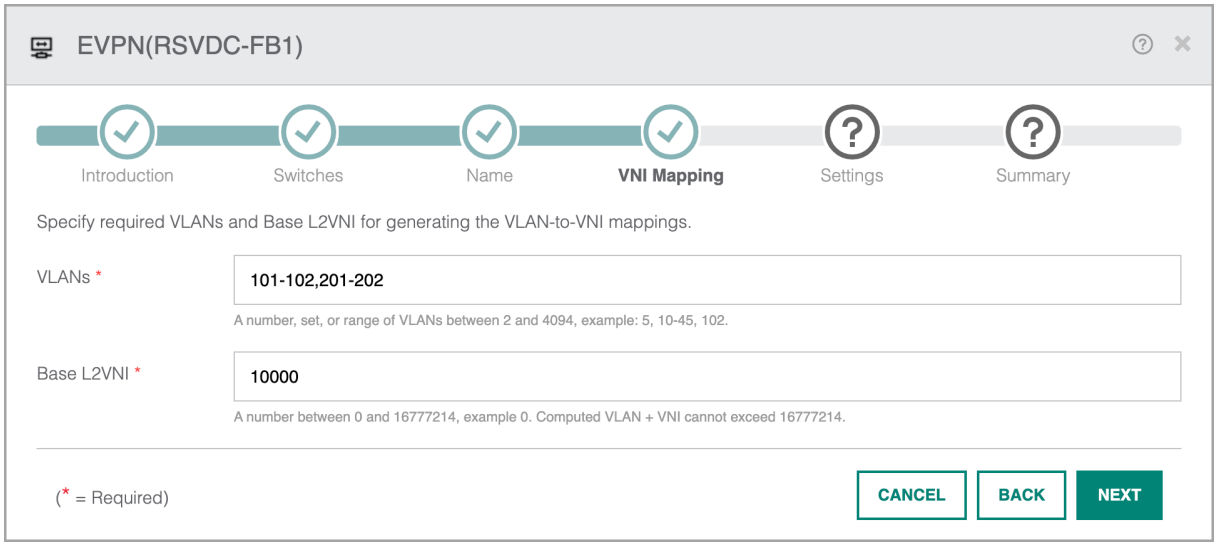


**NOTE:**

The **Base L2VNI** value is added to each VLAN ID to generate a unique **L2 VNI** associated to each VLAN automatically.

**Step 6** On the **Settings** page, click **ADD** to launch the **Resource Pool** wizard.

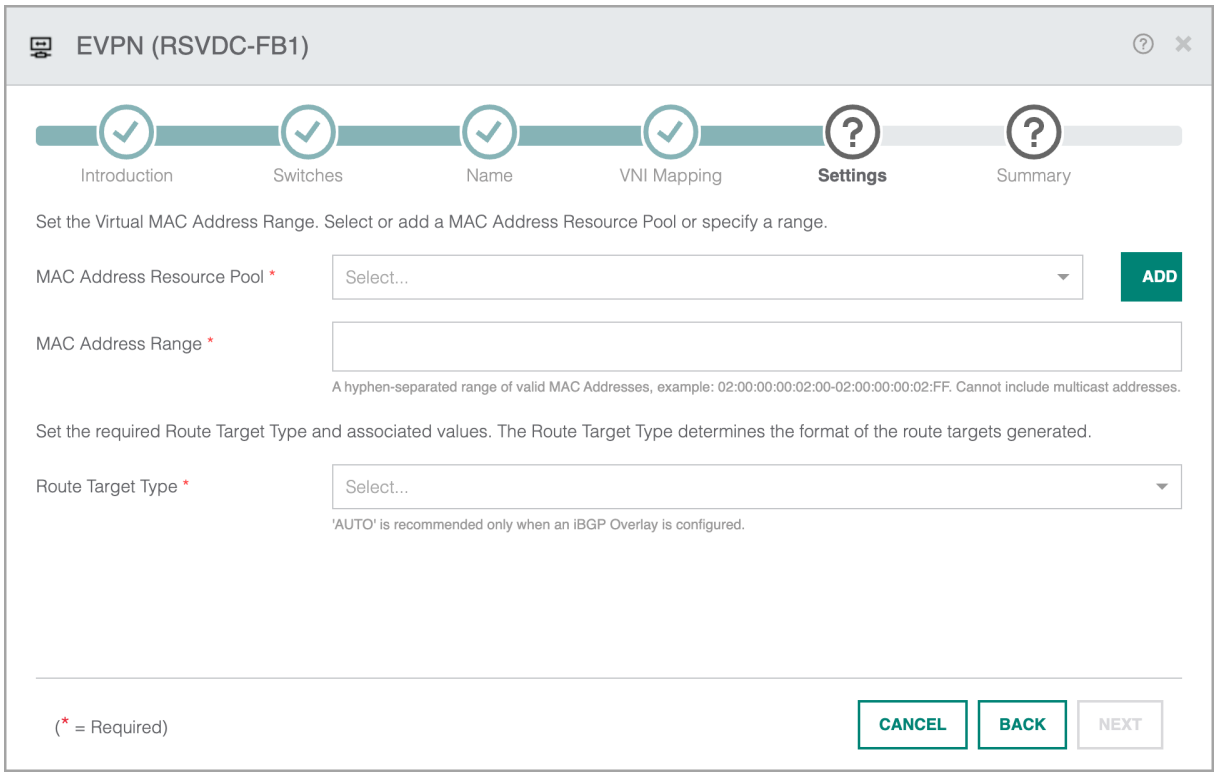


**Step 7 Resource Pool** wizard: On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

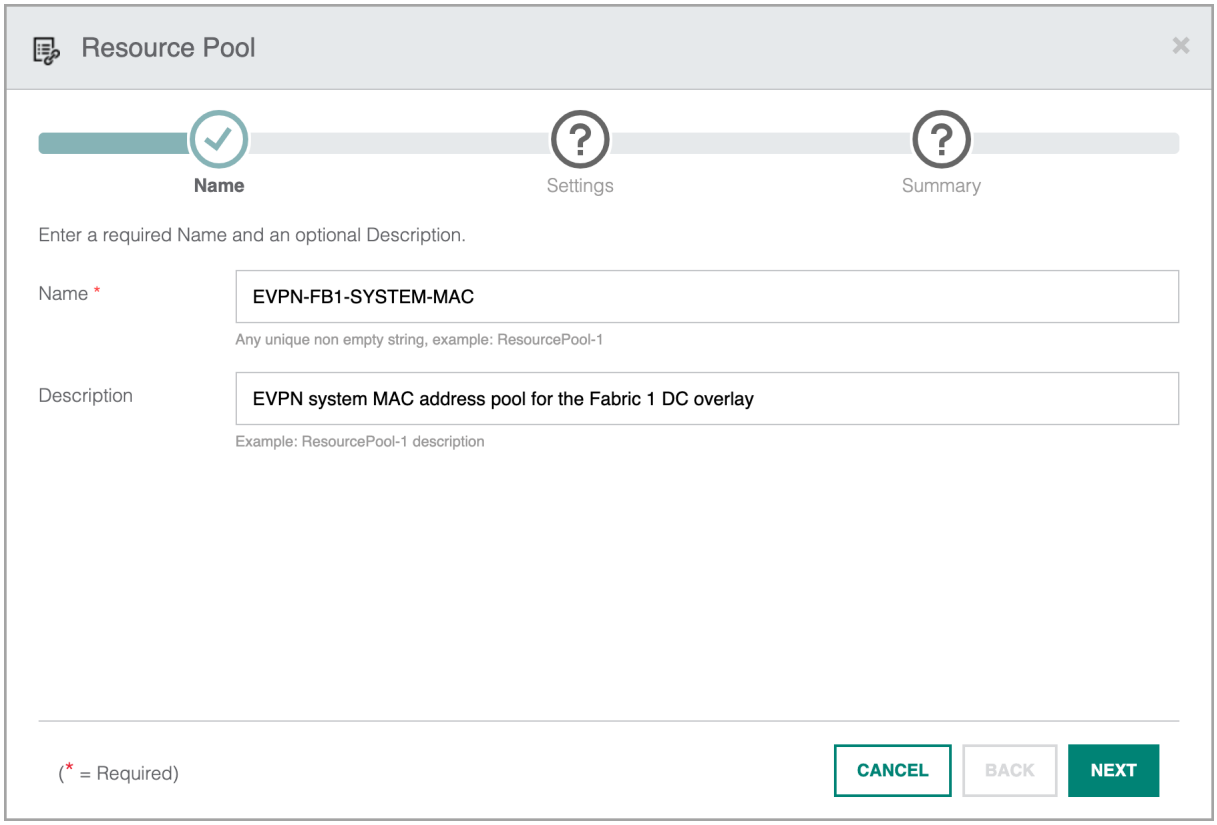


**Step 8 Resource Pool** wizard: On the **Settings** page, enter a MAC address range for System MAC Addresses in the **Resource Pool** field and click **NEXT**.

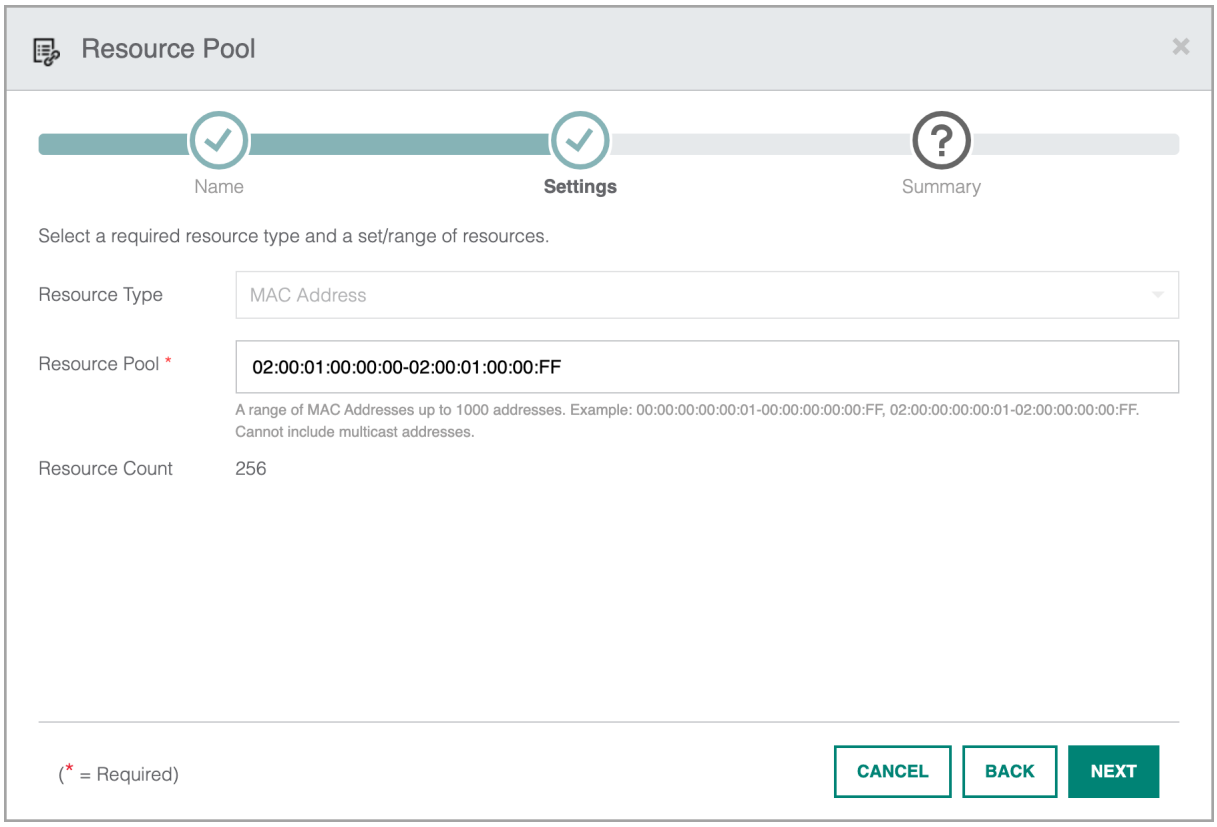


**Step 9 Resource Pool** wizard: On the **Summary** page, verify that the System MAC information is correct and click **APPLY**. The **Resource Pool** wizard closes and returns to the main **EVPN Configuration** workflow.

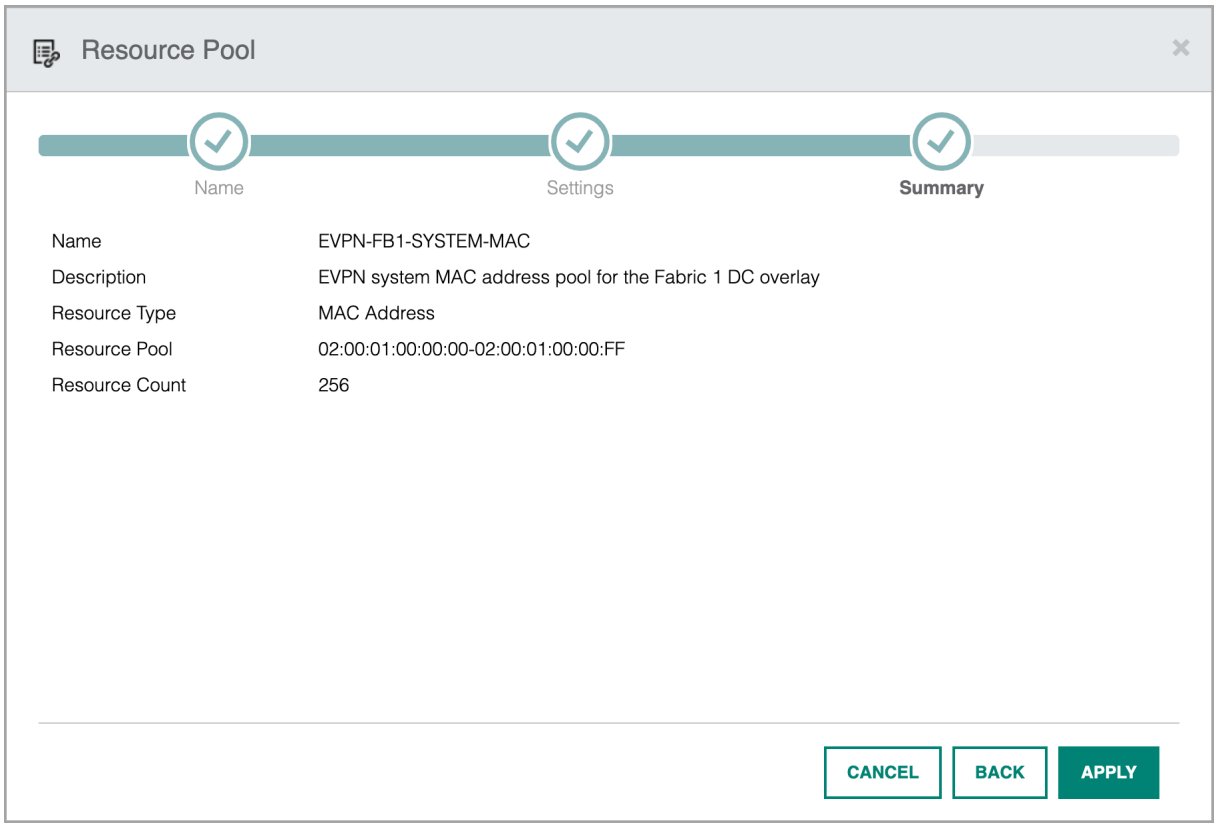

**Step 10** On the **Settings** page, verify that the **MAC Address Resource Pool** just created is selected, set the **Route Target Type** to *AUTO*, and click **NEXT**.

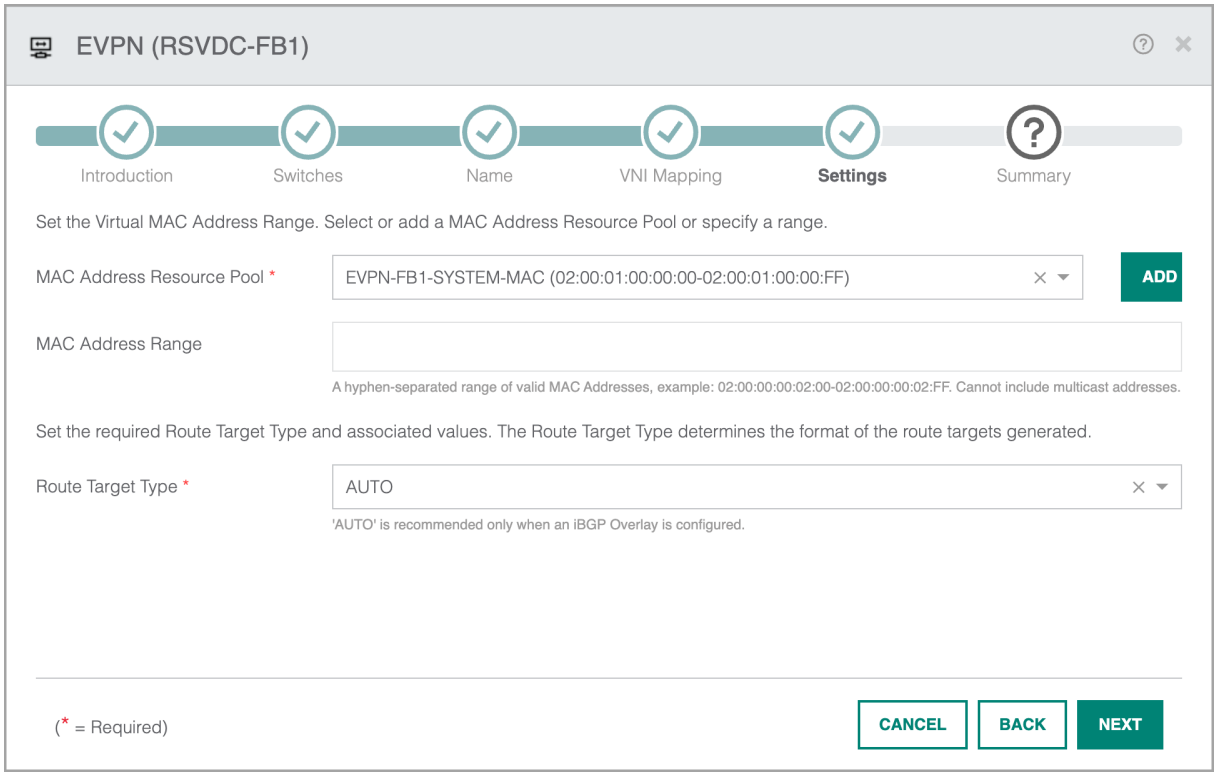


**NOTE:**

EVPN route targets can be set automatically by switches only when using an iBGP overlay.

**Step 11** On the **Summary** page, verify that the information is correct and click **APPLY**.

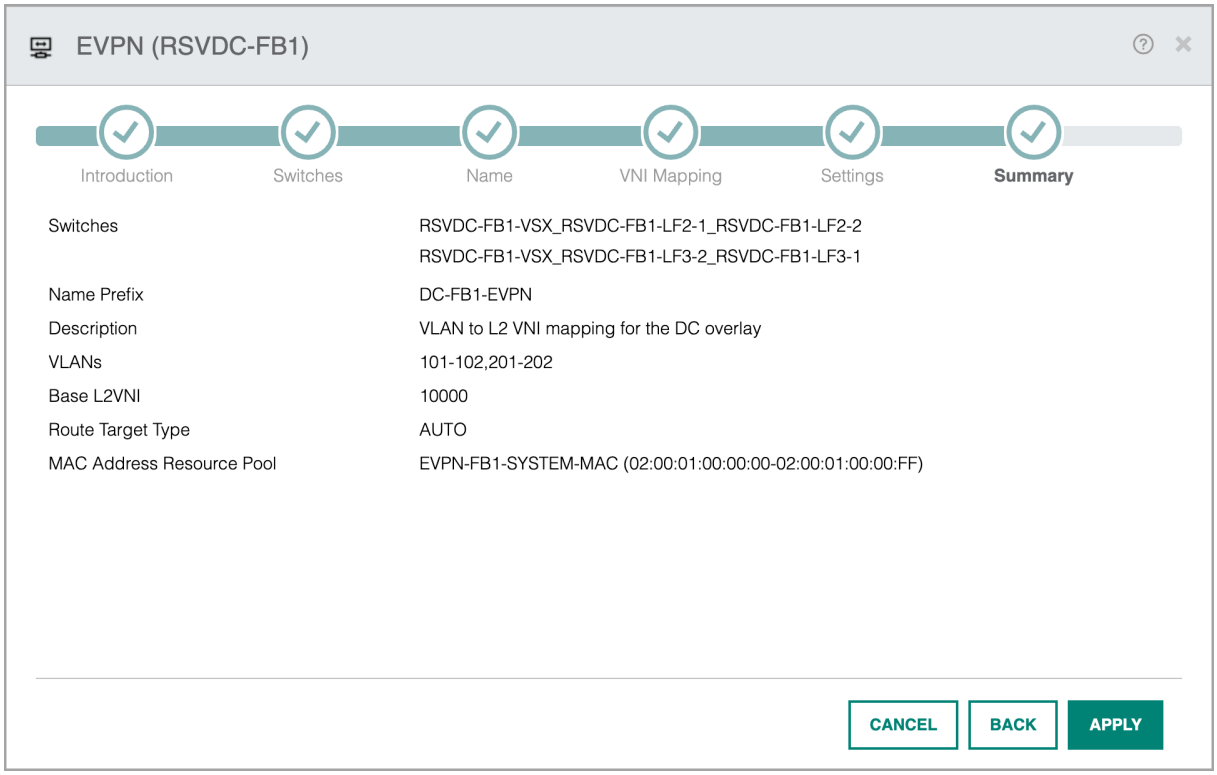

**Step 12** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 13** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select all non-border leaf switches* >
- **Commands:** *show interface vxlan vtep*

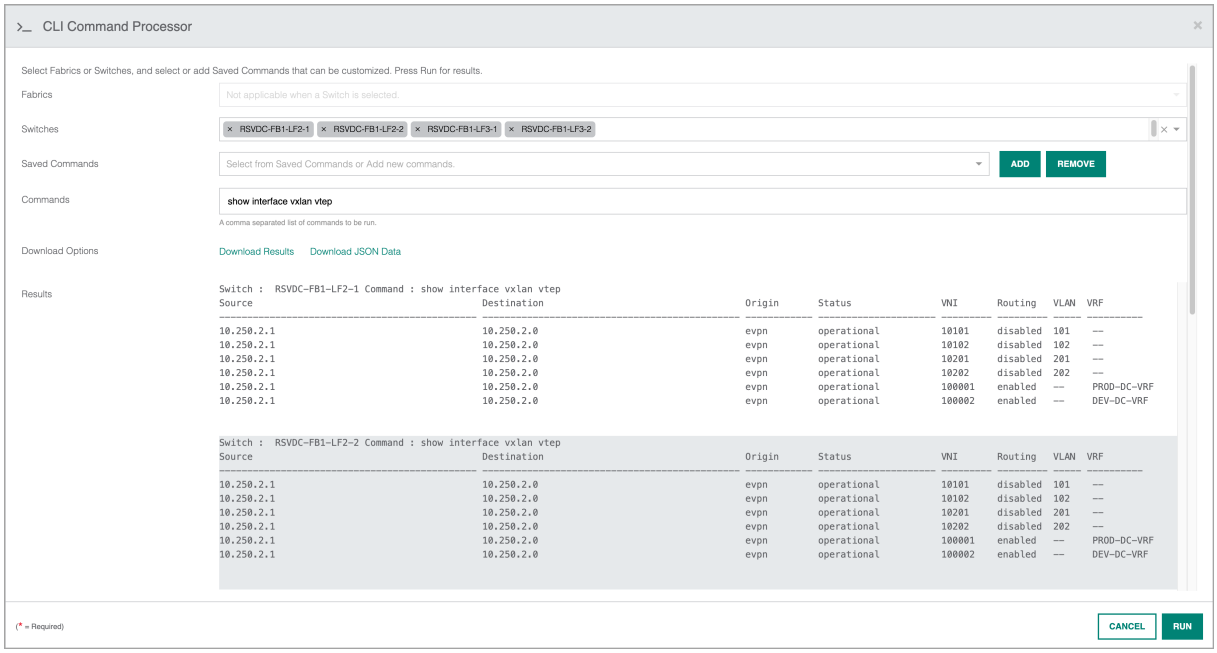


**Step 14** Verify that the output for each switch displays a remote VTEP to each non-border leaf switch pair for each VLAN.

The Guided Setup is now complete.

# Host Port Configuration

Use this section to configure the Port Groups and LACP Host LAG ports.

## Configure Port Groups

The SFP28 ports on the Aruba 10000-48Y6C switches (R8P13A and R8P14A) are organized into 12 groups of four ports each. The SFP28 default port speed is 25 Gb/s and must be set manually to 10 Gb/s, if required. Port groups can be configured on CX 8325, 8360, 10000, and 9300S series switches.

For additional details, find the *Installation and Getting Started Guide* for a specific switch model on the Aruba Support Portal. Go to the section: **Installing the switch > Install transceivers > Interface-Group operation**.

The following procedure configures a set of ports for 10 Gbps operation.

**Step 1** On the **Configuration** menu, select **Ports > Ports**.

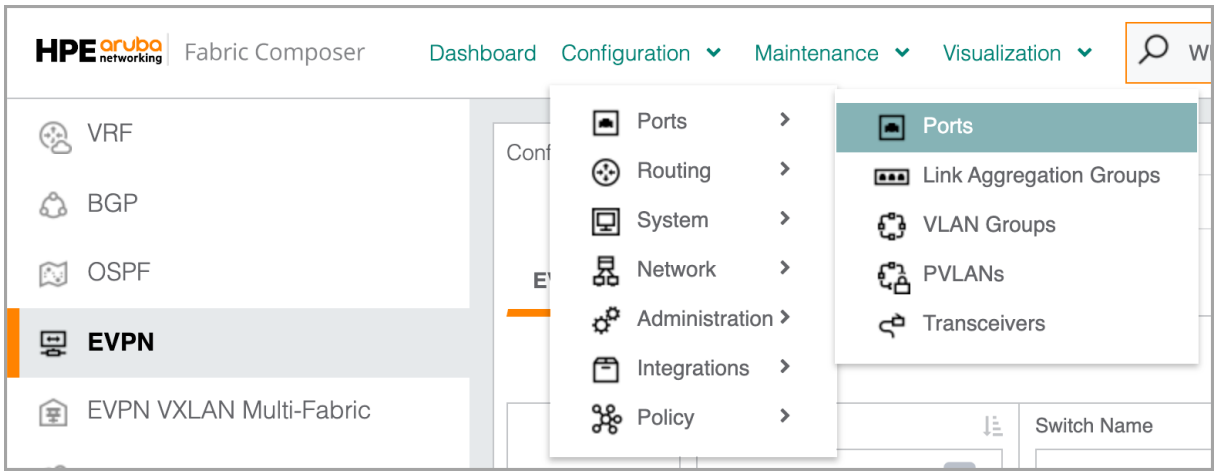


**Step 2** Select all switches that require port speed changes in the **Switch** field.

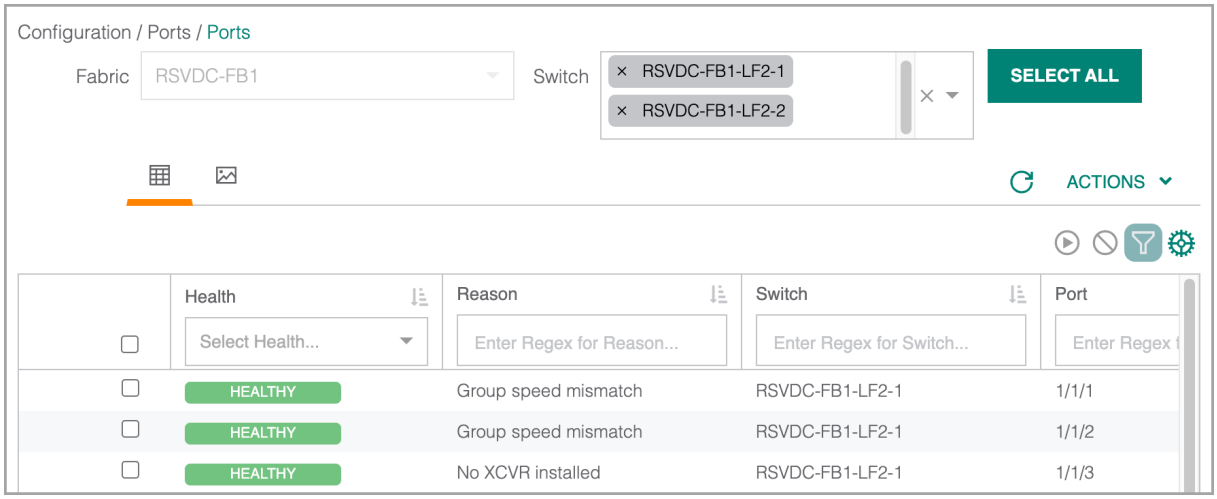


**Step 3** Enter **mismatch** in the **Reason** column's regex field and click the **Apply table filters** icon.

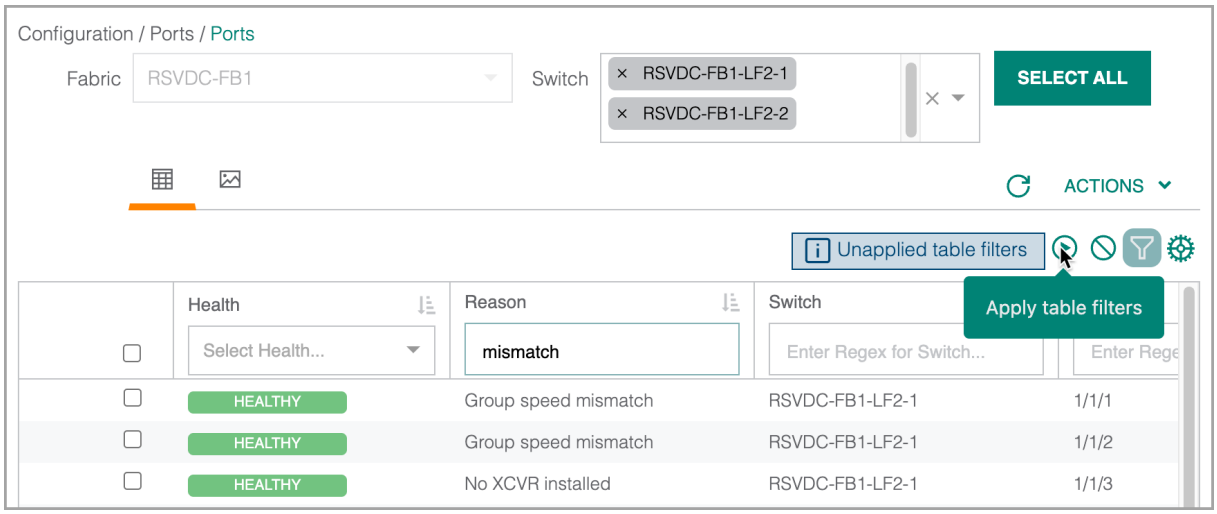


**NOTE:**

This step is optional. Cabling must be complete before this step so the switch can generate a speed mismatch status used for filtering.

**Step 4** Select an individual port in the port-group to be changed. On the right **ACTIONS** menu, select **Edit**.

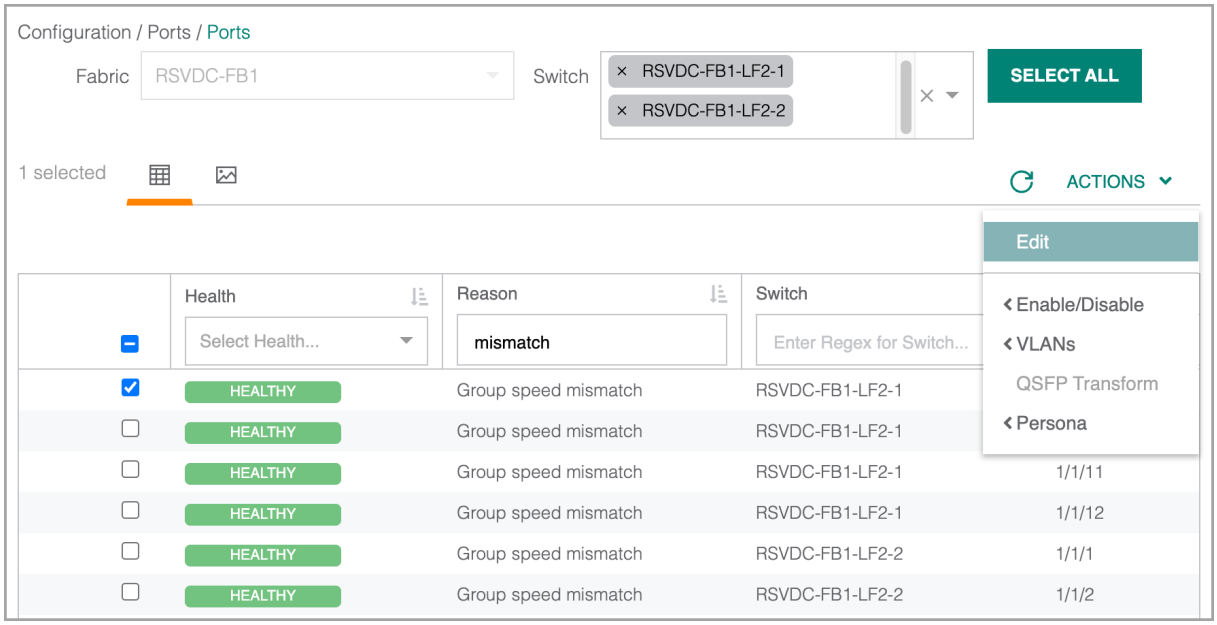


:::info The **Edit** option on the **ACTIONS** menu is available only when a single switch port is selected.

:::

**Step 5** On the **Ports** page, select the **Speed** tab. Select the appropriate value in the **Speed** dropdown, then click **APPLY**.

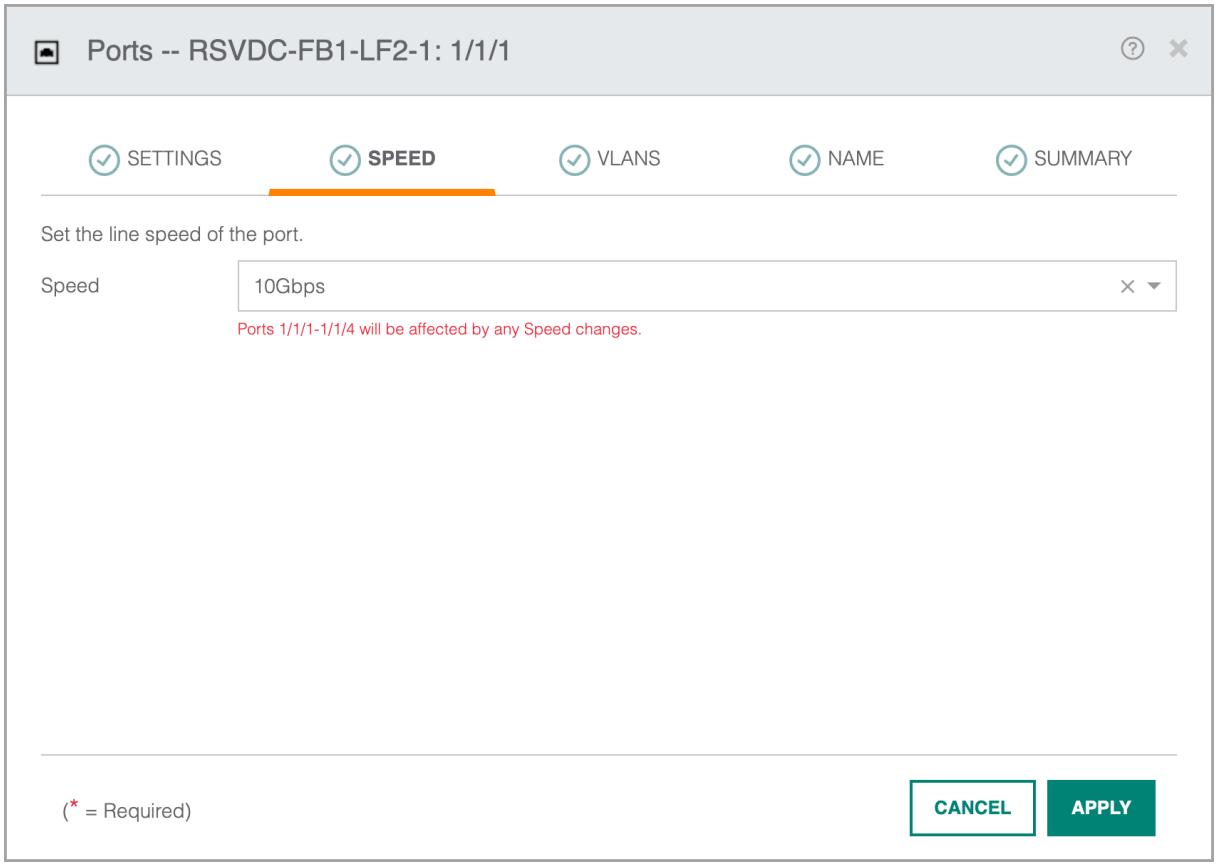

**NOTE:**

Observe the full list of ports affected by the speed change. Ensure that this is the correct speed setting for all listed ports.

**Step 6** Repeat the procedure for additional leaf switch ports requiring speed changes. Be sure to make changes to corresponding ports on VSX-paired switches supporting MC-LAGs.

**NOTE:**

The displayed port list of mismatched transceiver speeds is updated dynamically. It may be necessary to toggle the select all/deselect all checkbox in the upper left column to deselect the previously selected port after the update hides it from view.

## Multiple LACP MC-LAG Configuration

LACP link aggregation groups provide fault tolerance and efficient bandwidth utilization to physical hosts in the data center. The **Link Aggregation Group** wizard configures multi-chassis LAGs and LACP on fabric switches. Use the Fabric Composer **CLI Command Processor** to verify LAG interface state for LAG connected hosts.

Multiple MC-LAG creation can be applied to one or more switches using the Fabric Composer wizard. It enables quick setup of MC-LAGs across all leaf switches, when leaf switch models and cabling are consistent across the fabric. For example, a single pass of the Link Aggregation Groups wizard can configure all host leaf ports for MC-LAG given the following conditions:

- All leaf switches contain the same number of host facing ports.
- No more than one port per switch requires assignment to an individual MC-LAG.
- VLANs assigned to all MC-LAGs are consistent.

Configuration also is required on the connected hosts. Configuration varies by server platforms and operating systems and is not presented in this guide. Refer to the appropriate technical documentation for attached devices and operating systems.

**Step 1** On the **Configuration** menu, select **Ports > Link Aggregation Groups**.

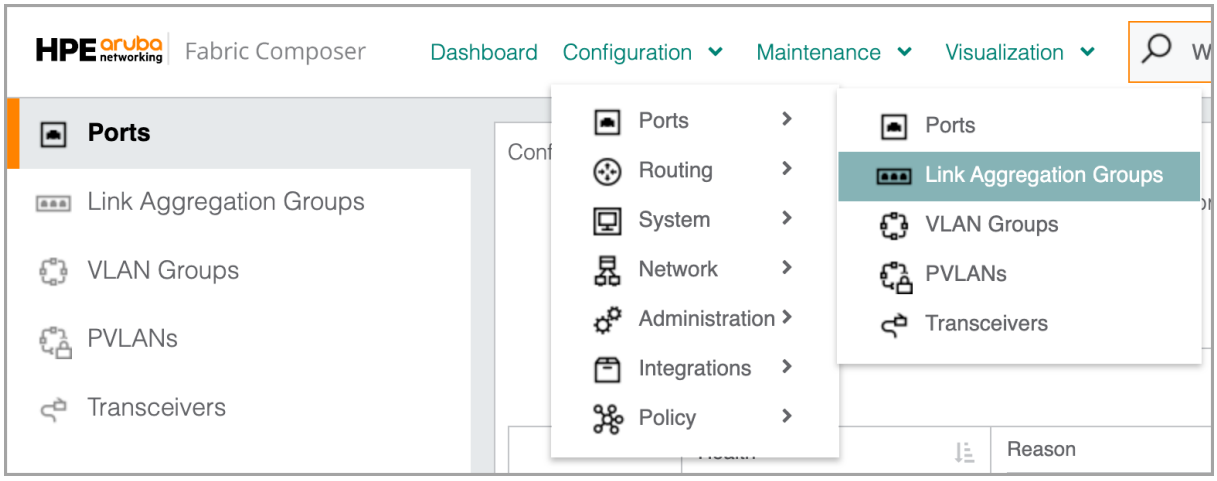

**Step 2** On the right **ACTIONS** menu, select **Add**.

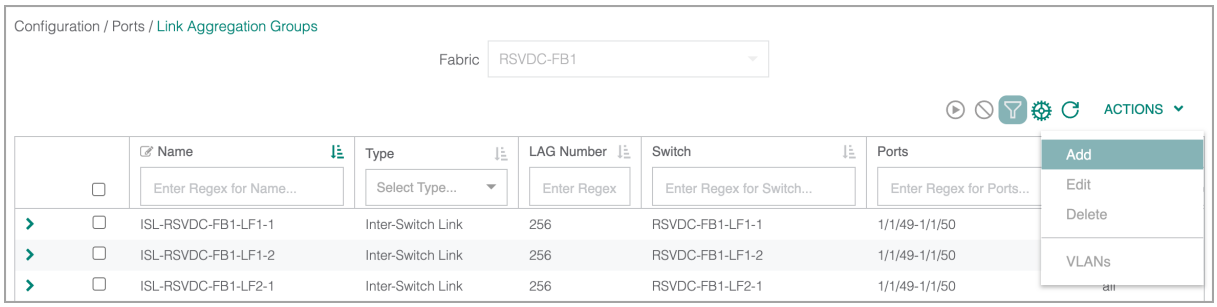


**Step 3** On the **Create Mode** page, select **Create multiple MLAGs for selected VSX Pairs** and click **NEXT**.

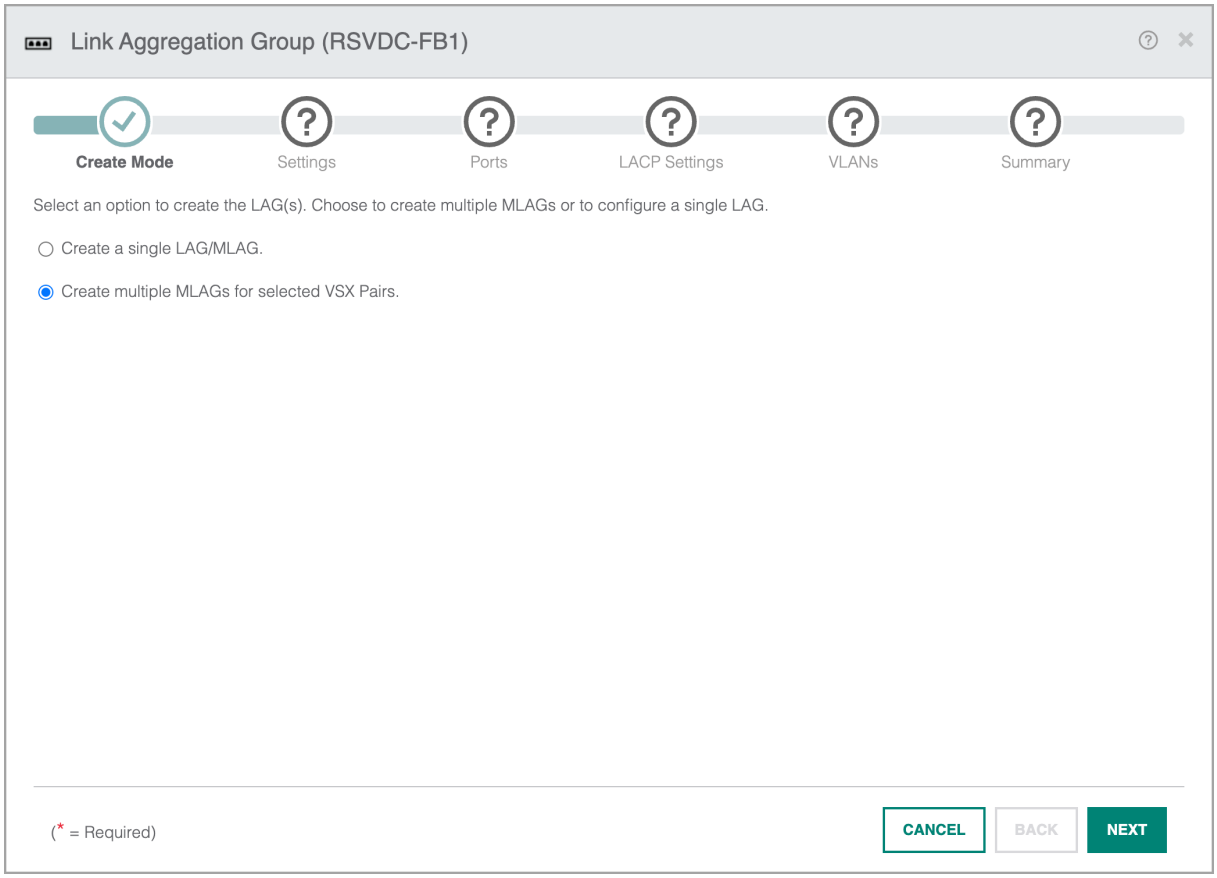


**Step 4** On the **Settings** page, enter a **Name Prefix**, **LAG Number Base**, then click **NEXT**.

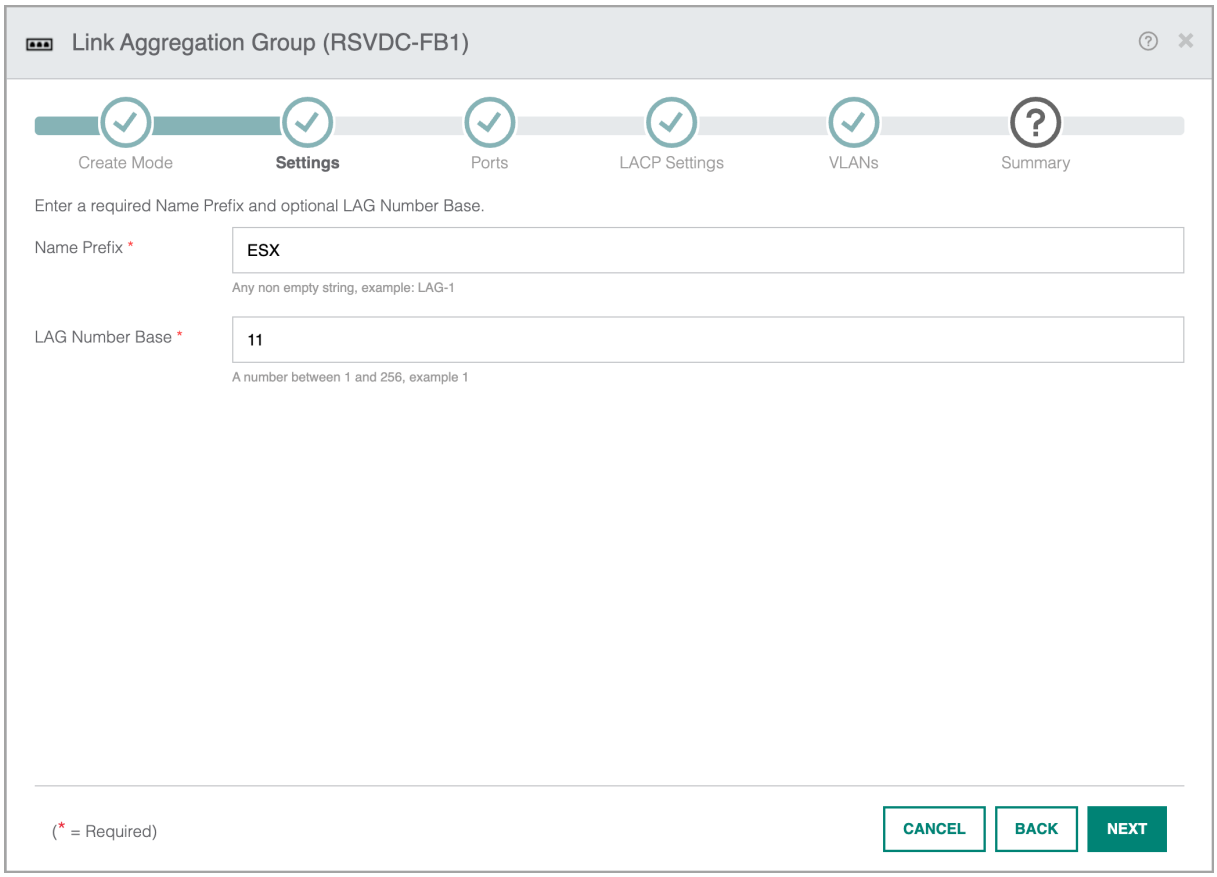


**NOTE:**

If individual hostname assignments are required per MC-LAG in place of a more general prefix-based naming convention, choose **Create a single LAG/MLAG** in the previous step to provide a unique a name per LAG.LAG index values are numbered sequentially beginning with the **LAG Number Base**. The wizard will not complete if a LAG value is already in use on any switch target. The server access/subleaf MC-LAG configured from the RSVDC-FB1-LF3-1 and RSVDC-FB1-LF3-2 VSX pair to the downstream RSVDC-FB1-LF3-SA1 switch uses LAG index 1 on all three switches. **LAG Number Base** 11 is chosen to avoid a conflict with the existing LAG.

**Step 5** On the **Ports** page, select one or more VSX-pairs of switches in the **VSX Pairs** field, enter the ports that are physically cabled for MC-LAG operation in **Ports**, then click **VALIDATE**.

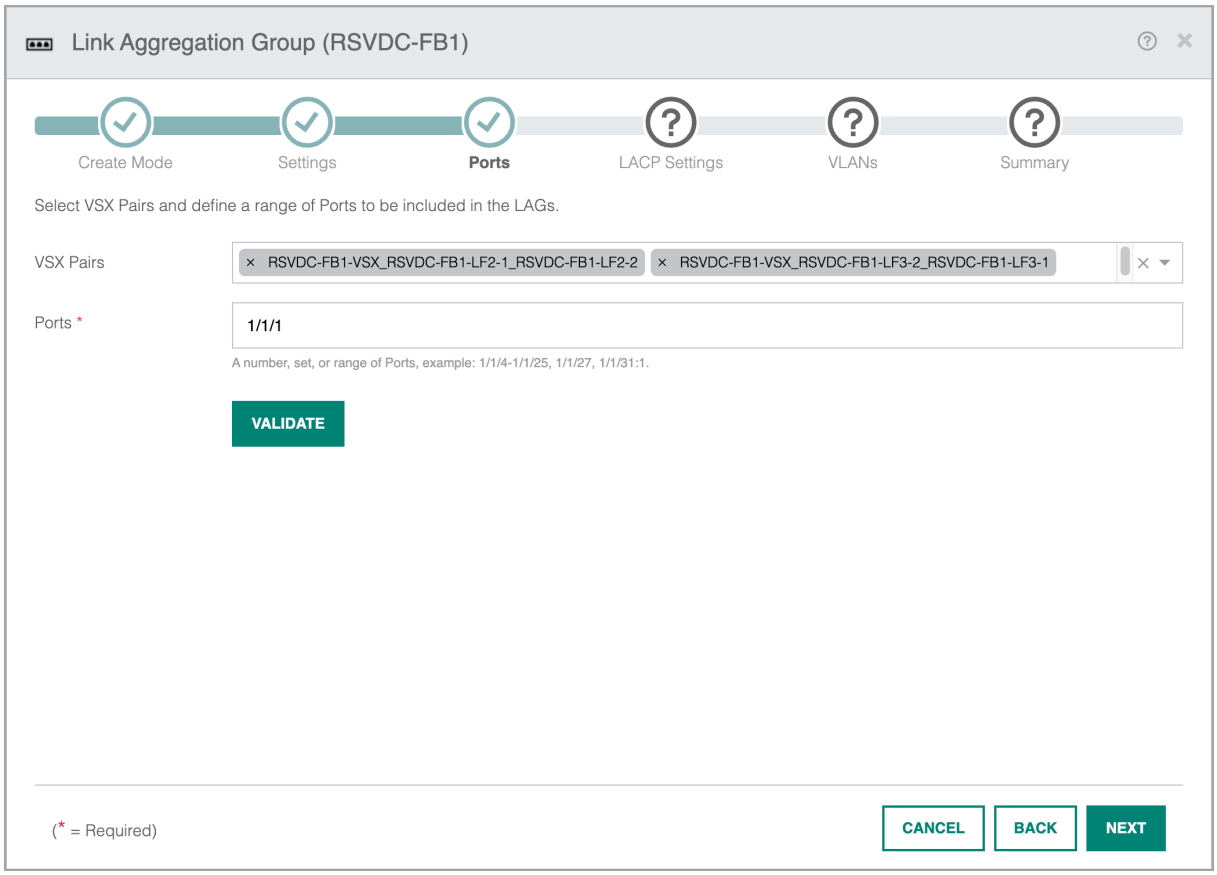


**Step 6** If Fabric Composer can validate that MC-LAG port configuration is consistent with LLDP neighbor data, a success message is presented.


**NOTE:**

Validation of hypervisor host connections requires previous assignment of physical host ports to LACP LAGs. Validation is intended to verify that the requested configuration is consistent with cabling to attached hosts. It is not required to continue the process of MC-LAG creation, if attached hosts are not configured or present. Fabric Composer configures an MC-LAG using port 1/1/1on RSVDC-FB1-LF2-1 and RSVDC-FB1-LF2-2 switches, and a second MC-LAG using port 1/1/1 on RSVDC-FB1-LF3-1 and RSVDC-FB1-LF3-2 switches, when using the values above. Specifying multiple ports will create additional MC-LAGs with corresponding port numbers between the VSX switch pairs.

**Step 7** On the **Ports** page, click **NEXT**.

**Step 8** On the **LACP Settings** page, check **Enable LACP Fallback**, leave other LACP settings at their default values, and click **NEXT**.

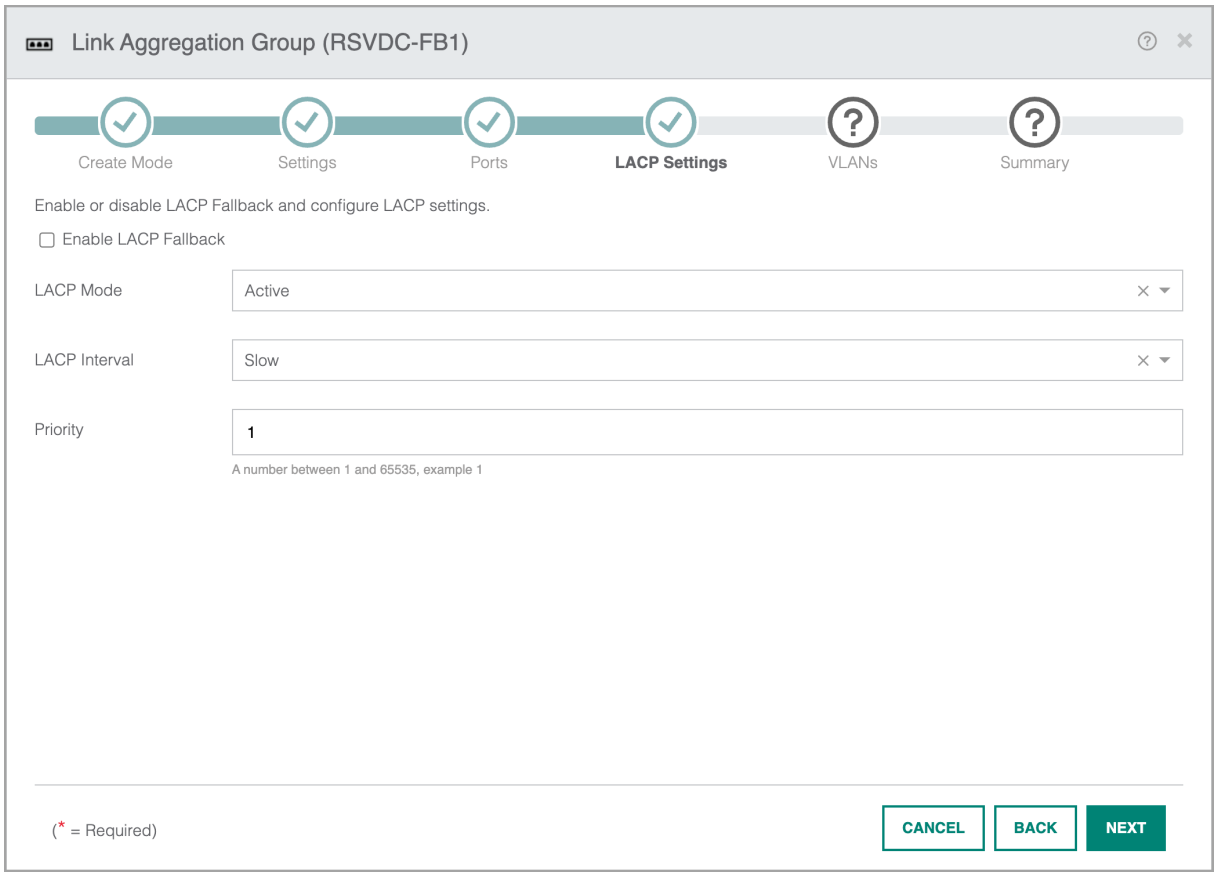


**Step 9** On the **VLANs** page, modify the untagged **Native VLAN** number if necessary, enter the tagged VLAN IDs in the **VLANs** field, then click **NEXT**.

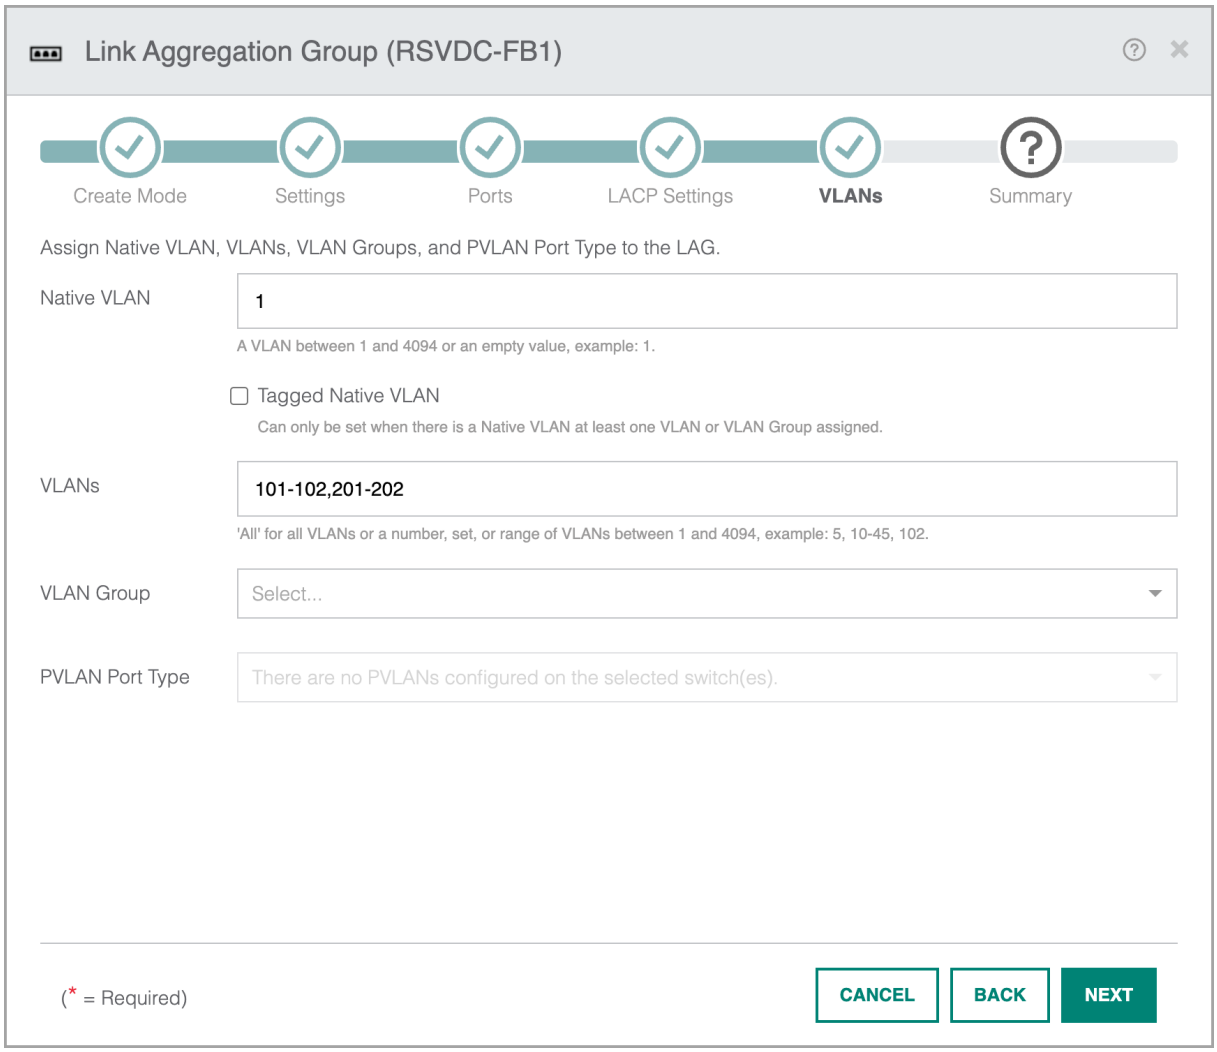


**Step 10** On the **Summary** page, confirm that the information is entered correctly and click **APPLY** to create the LAGs.

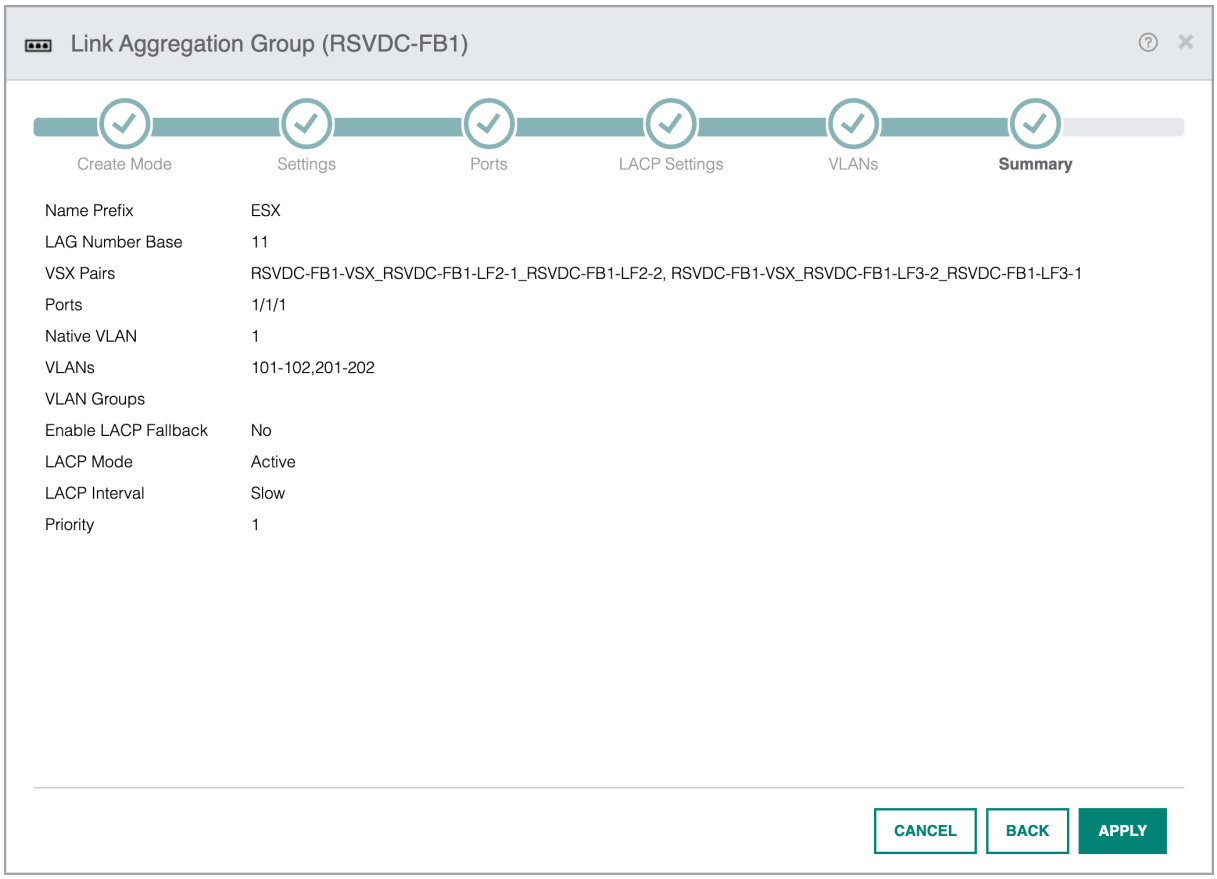


**Step 11** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 12** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** *< Select all switches with newly configured LAGs >*
- **Commands:** *show lacp interfaces*

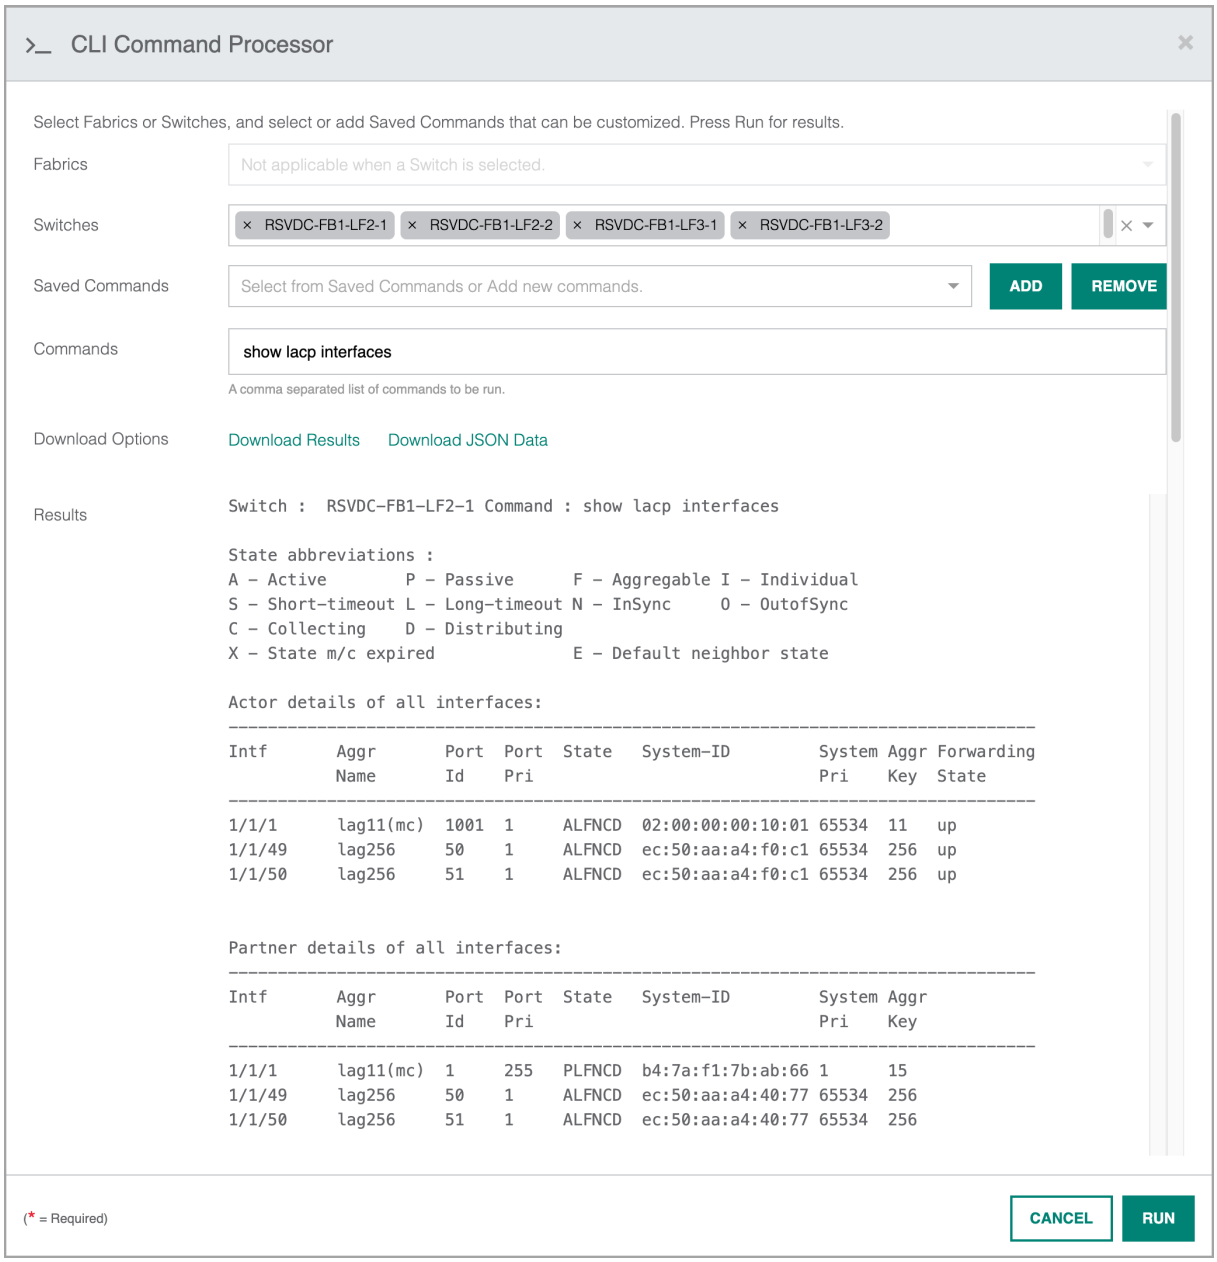


**Step 13** When a host is connected to the LAG, verify that each port assigned to one of the host LAGs created in this procedure has a **State** of “ALFNCD” for its local interfaces and “PLFNCD” for its partner interfaces. The **Forwarding State** should be “Up” for local interfaces.

**NOTE:**

A combination of VSX peer LAG interfaces and VSX multi-chassis LAG interfaces to hosts may be included in the command output. As shown above, the multi-chassis interfaces are denoted with **(mc)** after the LAG name. The **Actor** is the switch where the command was run. The **Partner** is the host at the other end of the LAG. The **State** column shows the expected values for a switch set to **Active** LACP mode and a host set to **Passive** LACP mode with a healthy LAG running.

## Single LACP MC-LAG Configuration

Individual LAGs are assigned for the following conditions:

- A unique LAG name is required in Fabric Composer.
- Assigned VLANs are unique to the LAG.
- More than one port per switch are assigned to an MC-LAG to increase capacity.

**Step 1** On the **Configuration** menu, select **Ports > Link Aggregation Groups**.

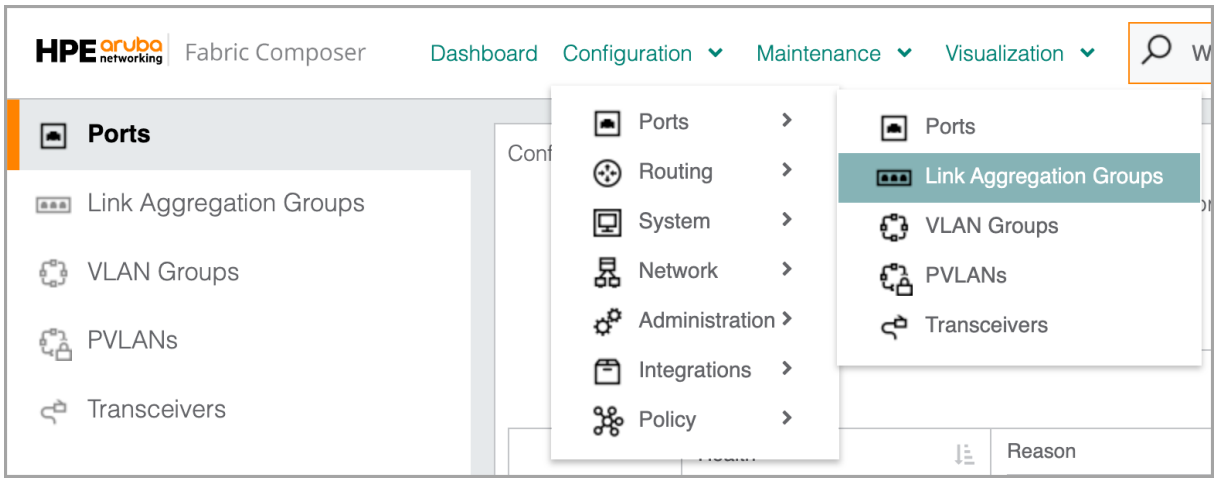


**Step 2** On the right **ACTIONS** menu, select **Add**.

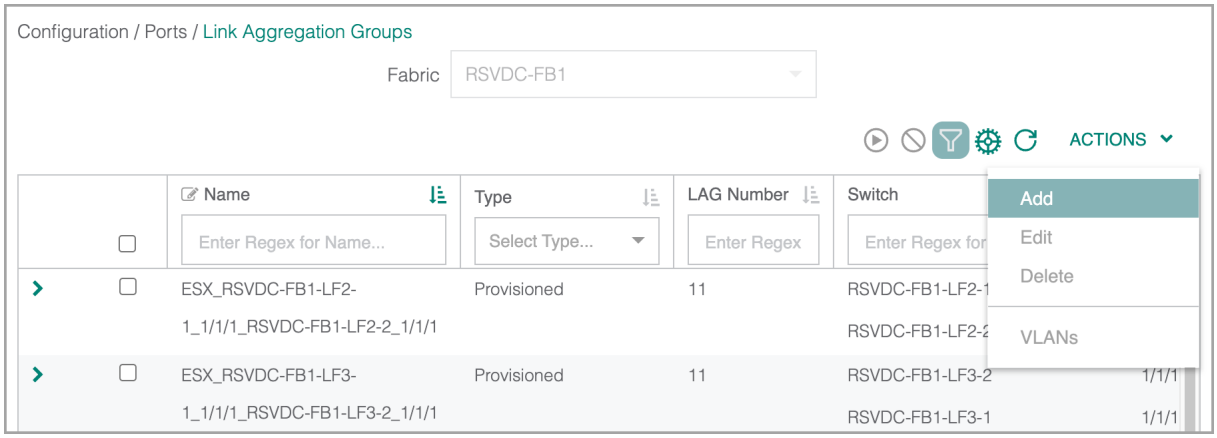


**Step 3** On the **Create Mode** page, leave **Create a single LAG/MLAG** selected and click **NEXT**.

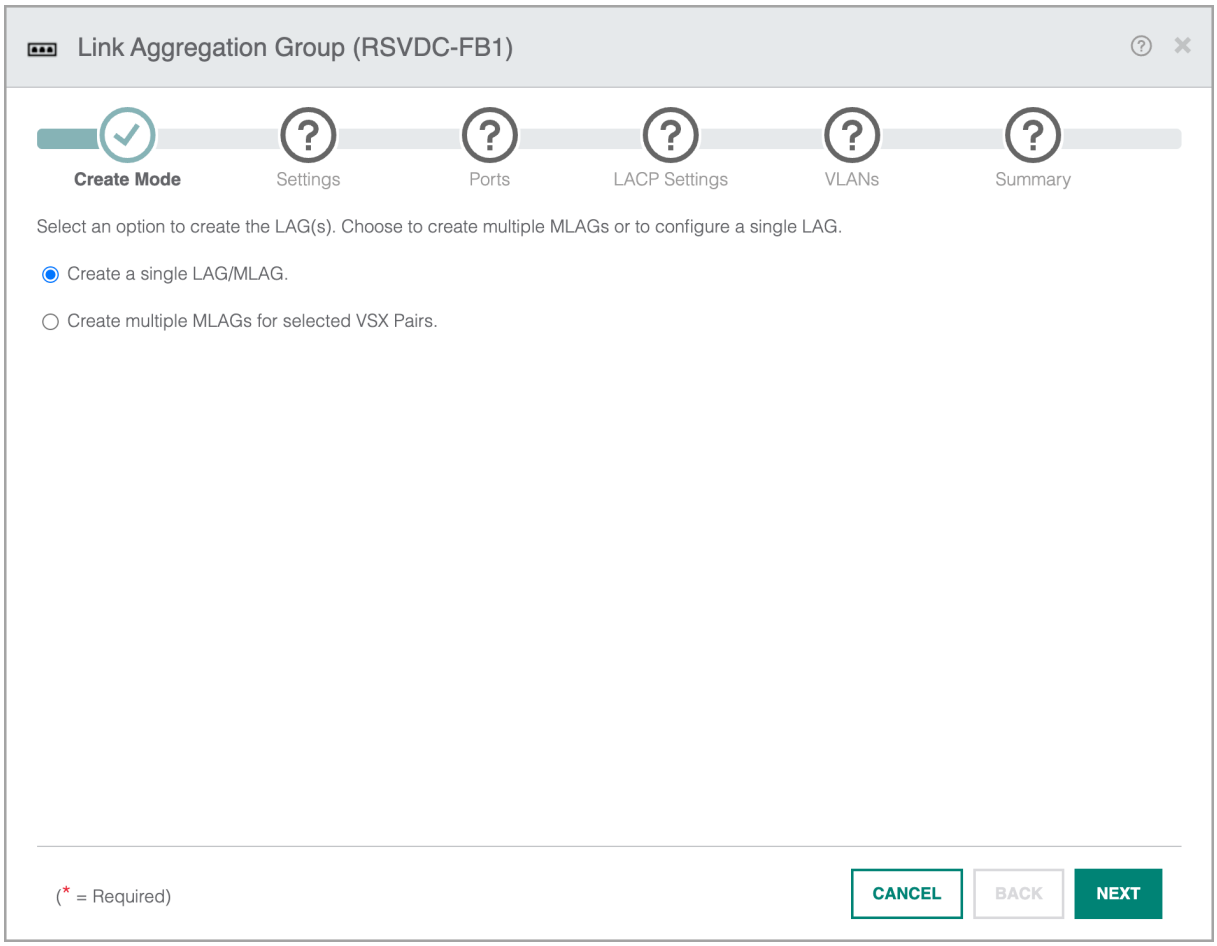


**Step 4** On the **Settings** page, enter a **Name**, **Description**, and **LAG Number**. Click **NEXT**.

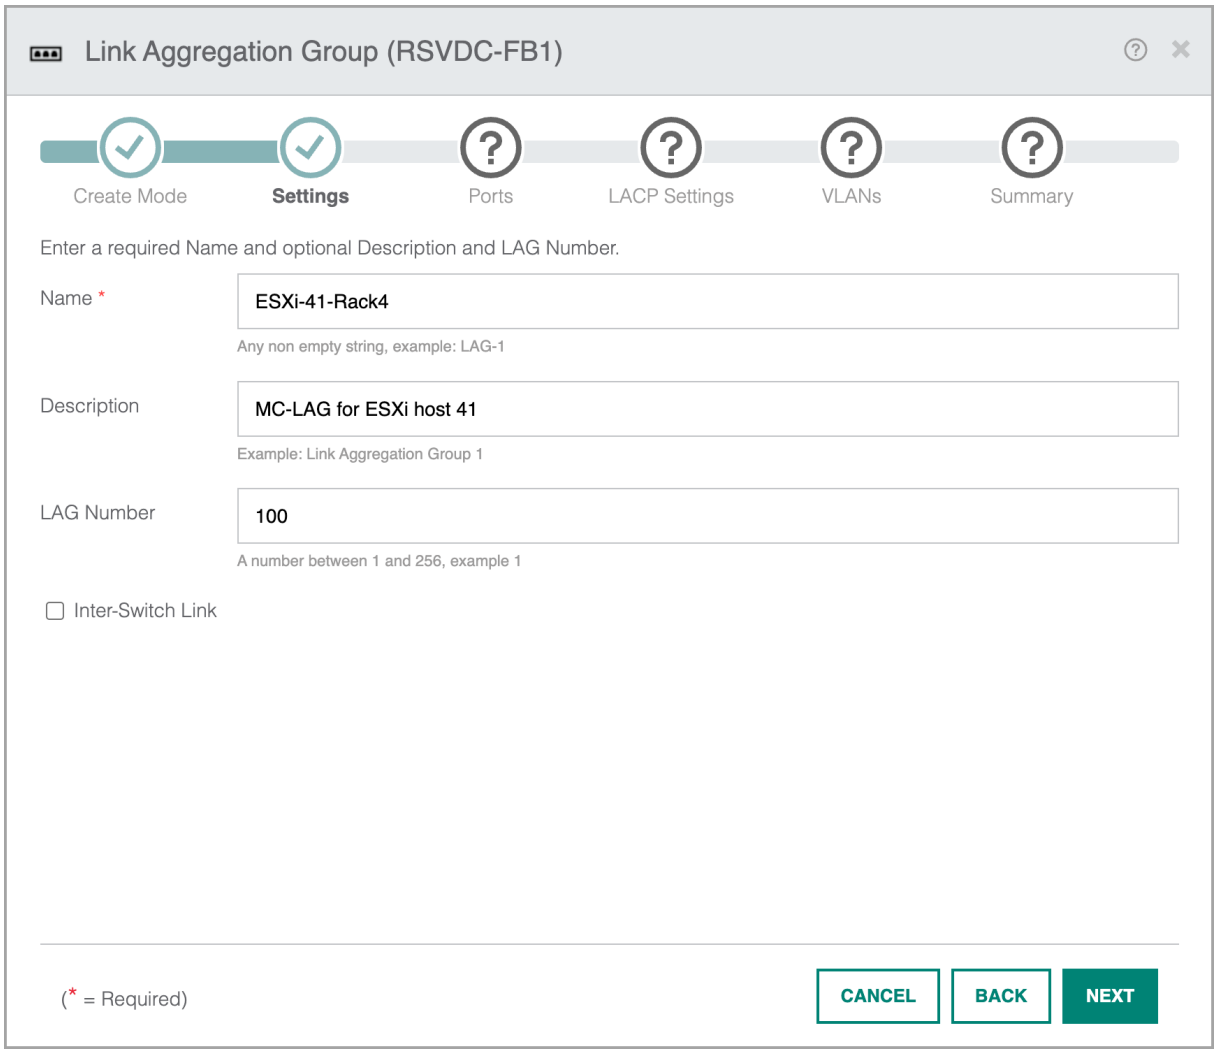


**NOTE:**

Consider using a **Name** that identifies the host and where it is connected.

**Step 5** On the **Ports** page, select a VSX-pair or a VSF stack of switches from the **LAG Switch Member** dropdown.

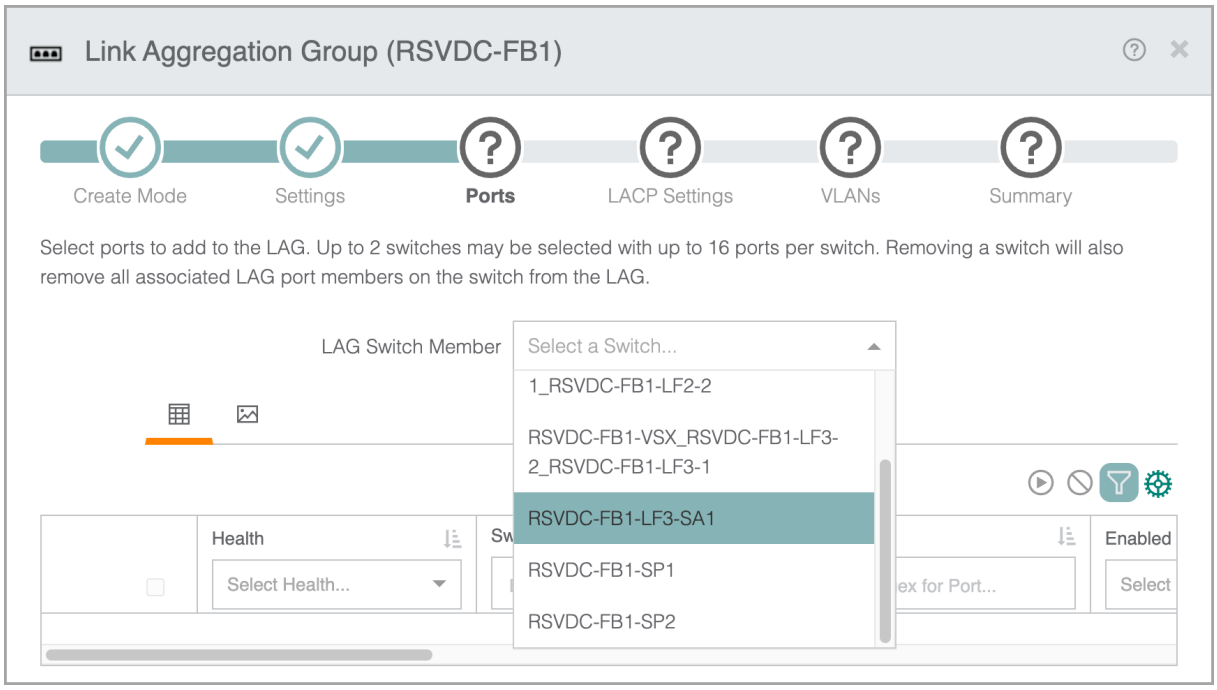


**Step 6** Click the **Switch View** mode icon to identify ports more easily.

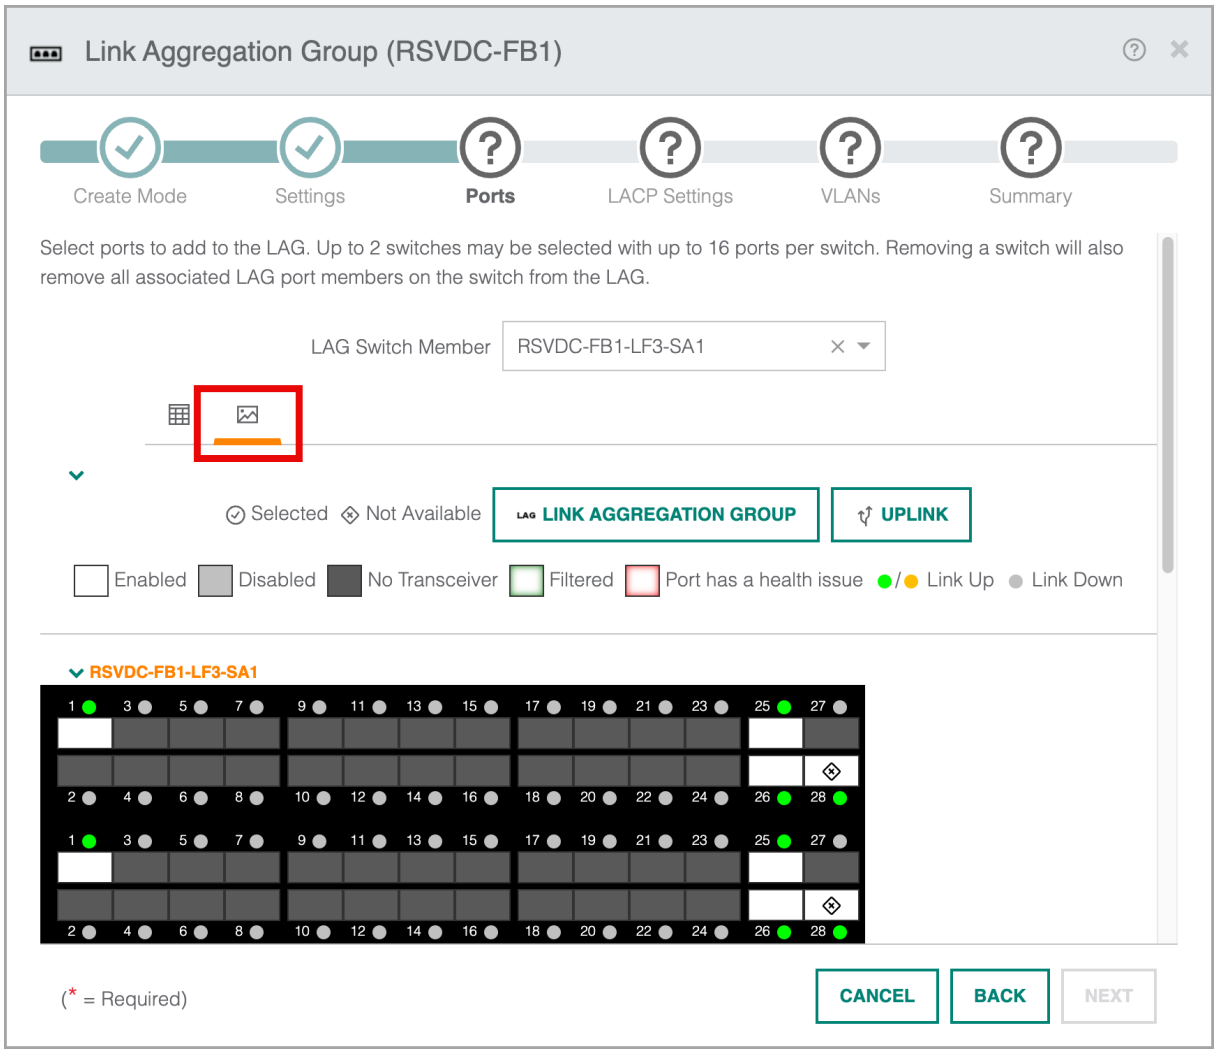


**Step 7** Click the port icons to add them as members of the link aggregation group and click **NEXT**.

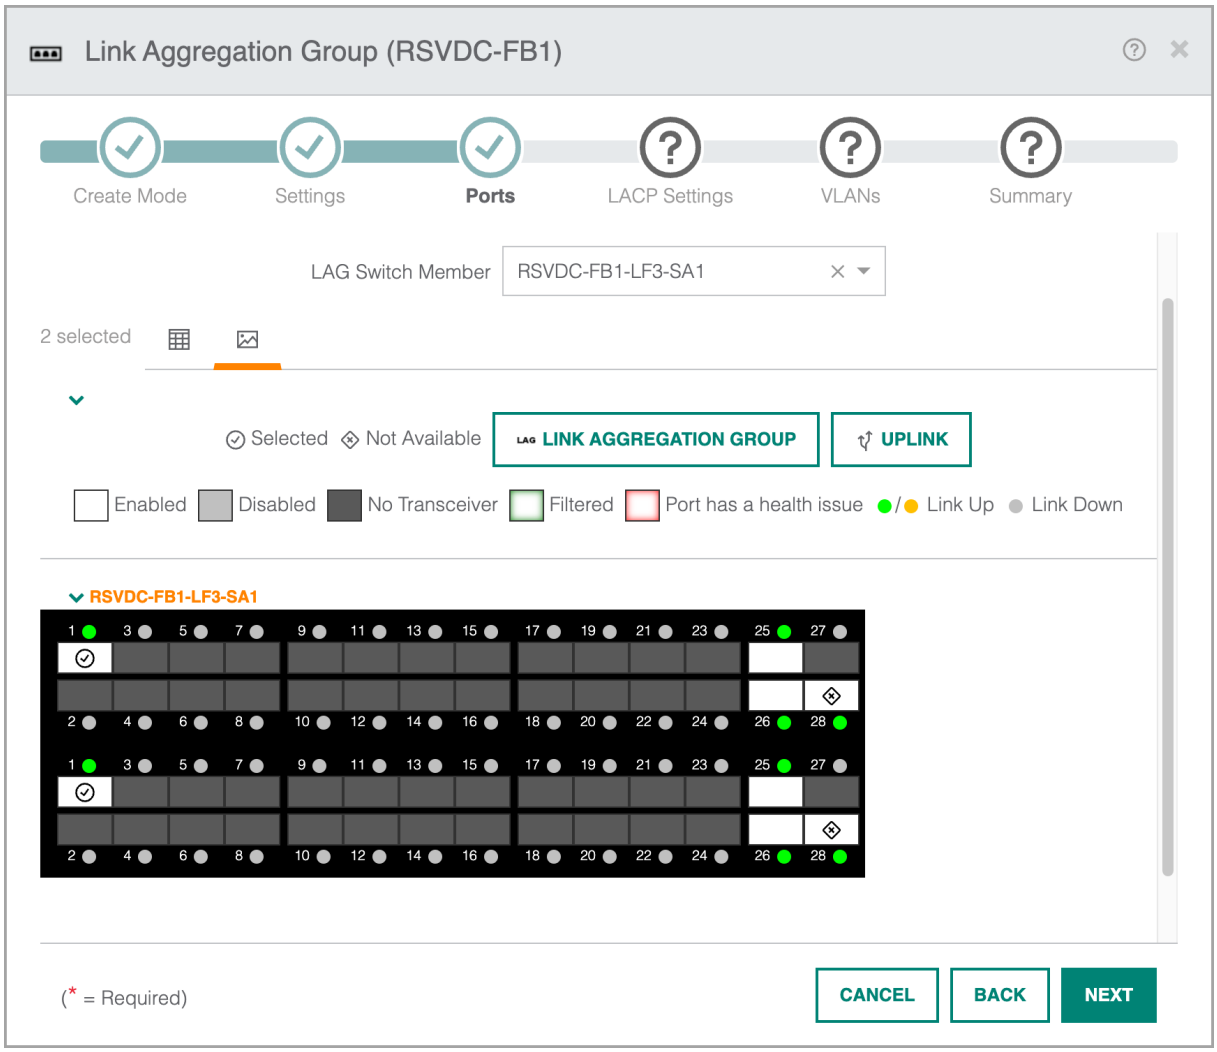


**NOTE:**

A checkmark appears on the newly selected ports. The diamond icon appears on ports not currently available for a new LAG group assignment. Select ports on both VSX switches or multiple VSF switches to create a functional multi-chassis LAG.

**Step 8** On the **LACP Settings** page, leave the settings at their defaults, and click **NEXT**.

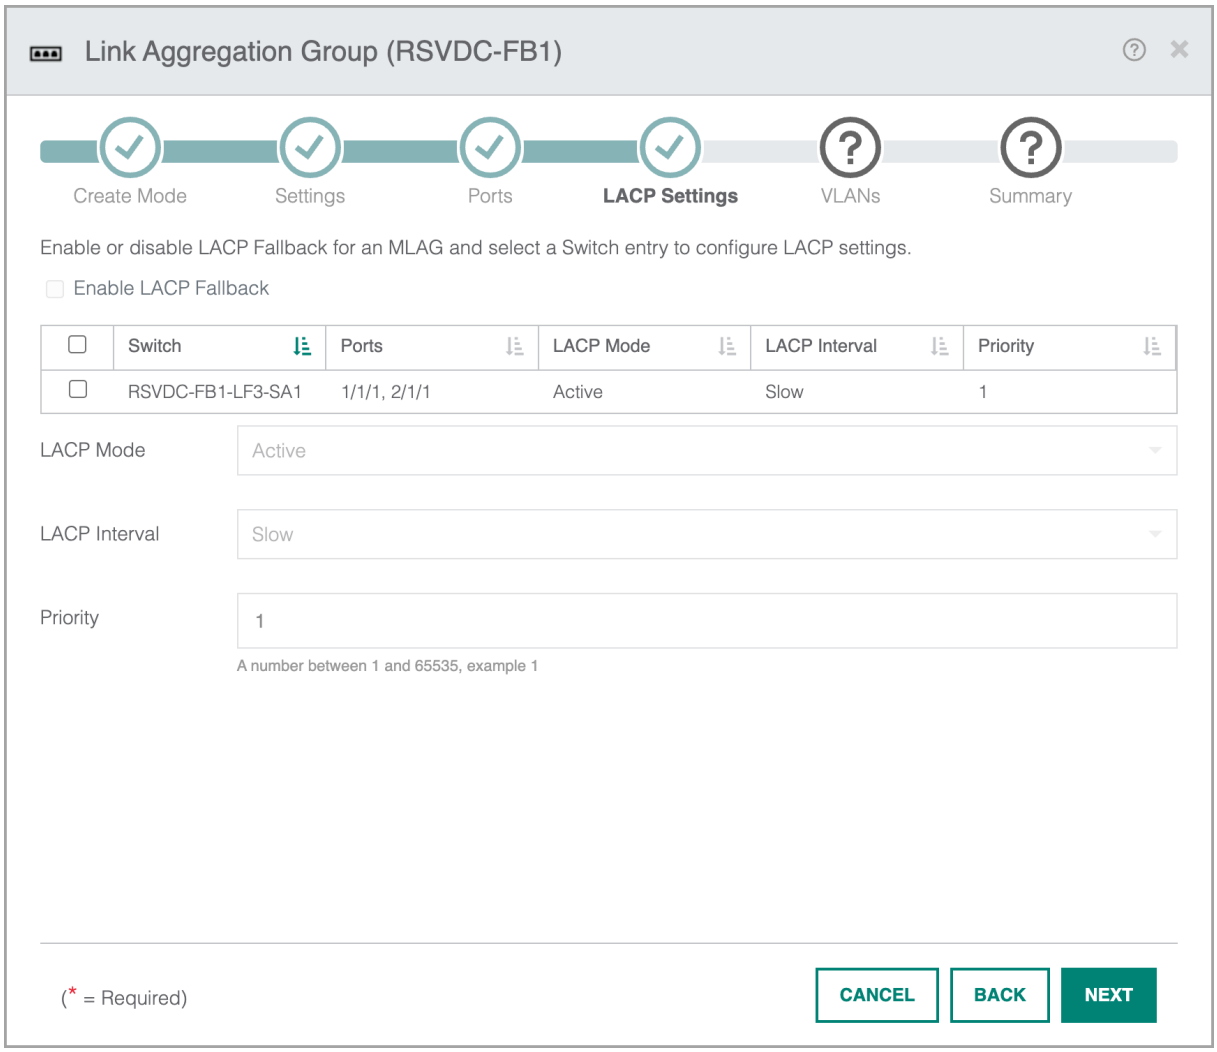


**NOTE:**

**Enable LACP Fallback** is auto-selected for VSX pairs. LACP fallback is not a valid option on VSF stacks. CX switches default to “Active” mode to ensure that LACP can be established regardless of the LACP configuration of the host platform. Using the default settings is recommended. Click the box next to one or both switch names to modify default values.

**Step 9** On the **VLANs** page, modify the untagged **Native VLAN** number if necessary, enter tagged VLAN IDs in the **VLANs** field, then click **NEXT**.

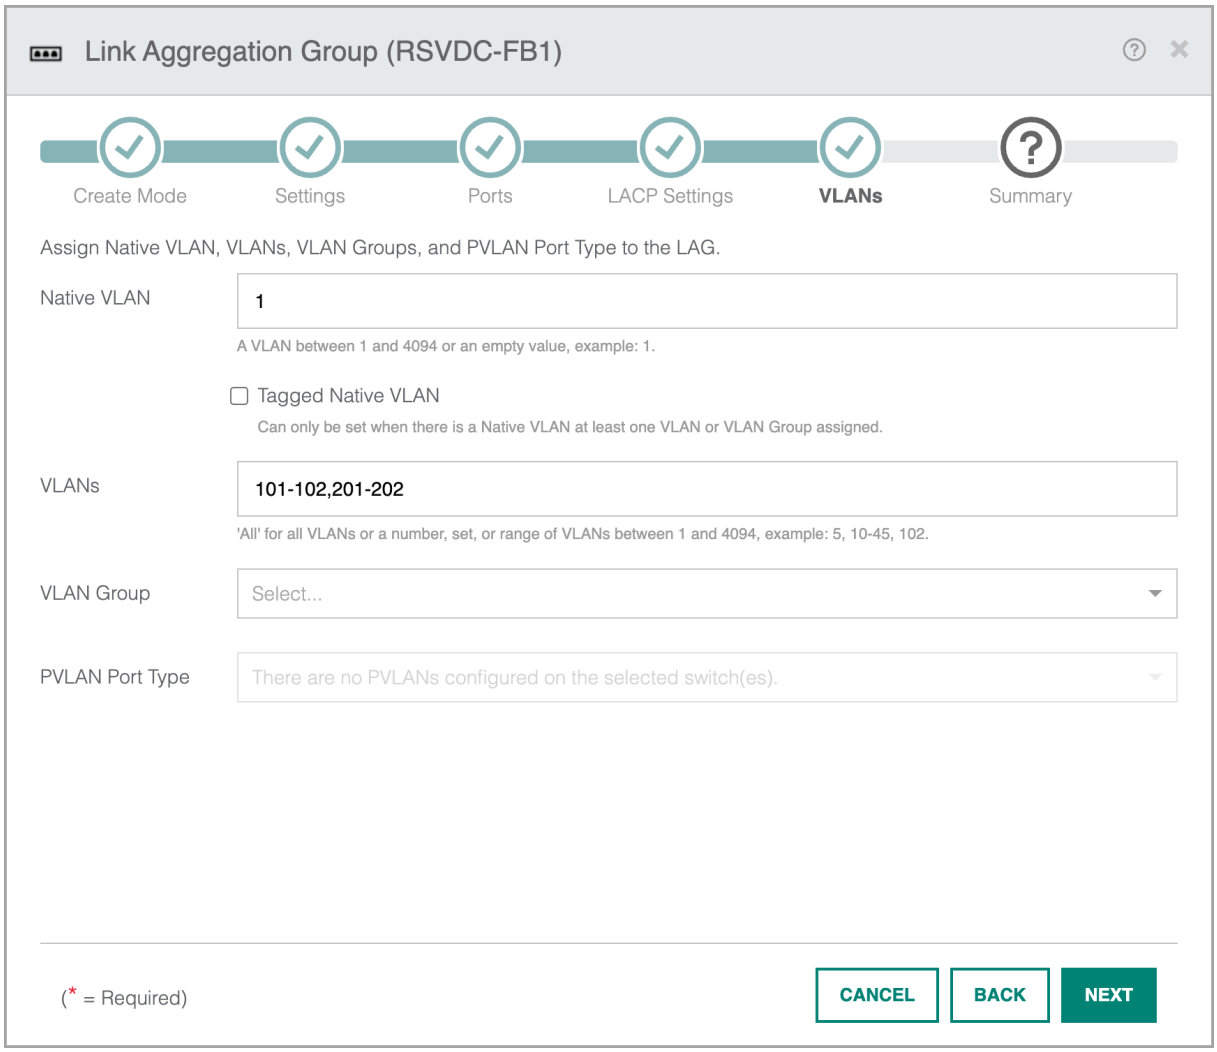


**Step 10** On the **Summary** page, confirm that the information is entered correctly and click **APPLY** to create the LAGs.

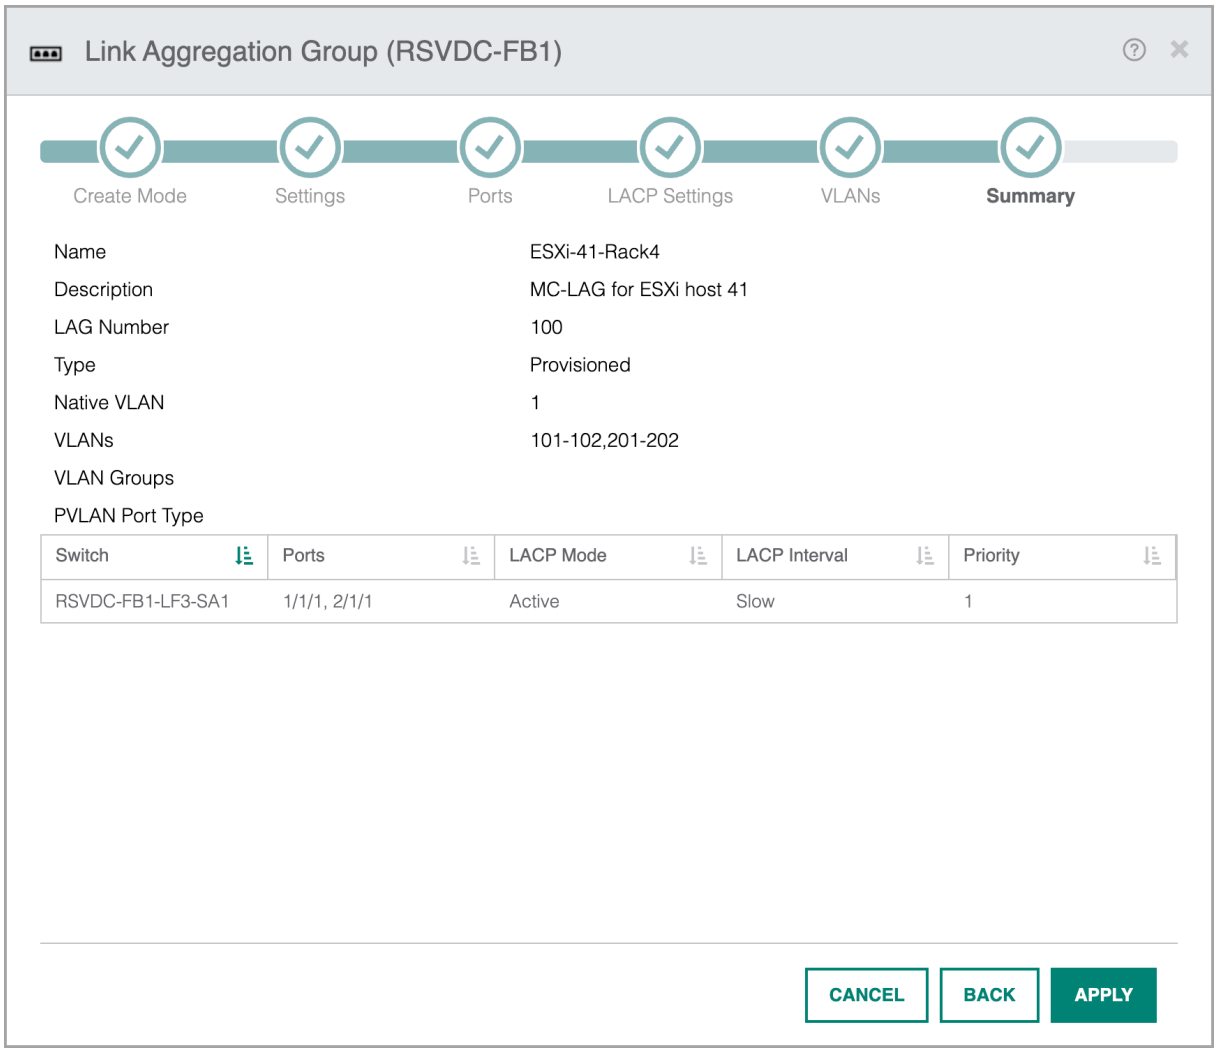


**Step 11** Repeat the procedure for each individual LAG connection in the fabric.

**Step 12** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 13** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select the switch LAG configured for the new LAG* >
- **Commands:** *show lacp interfaces*

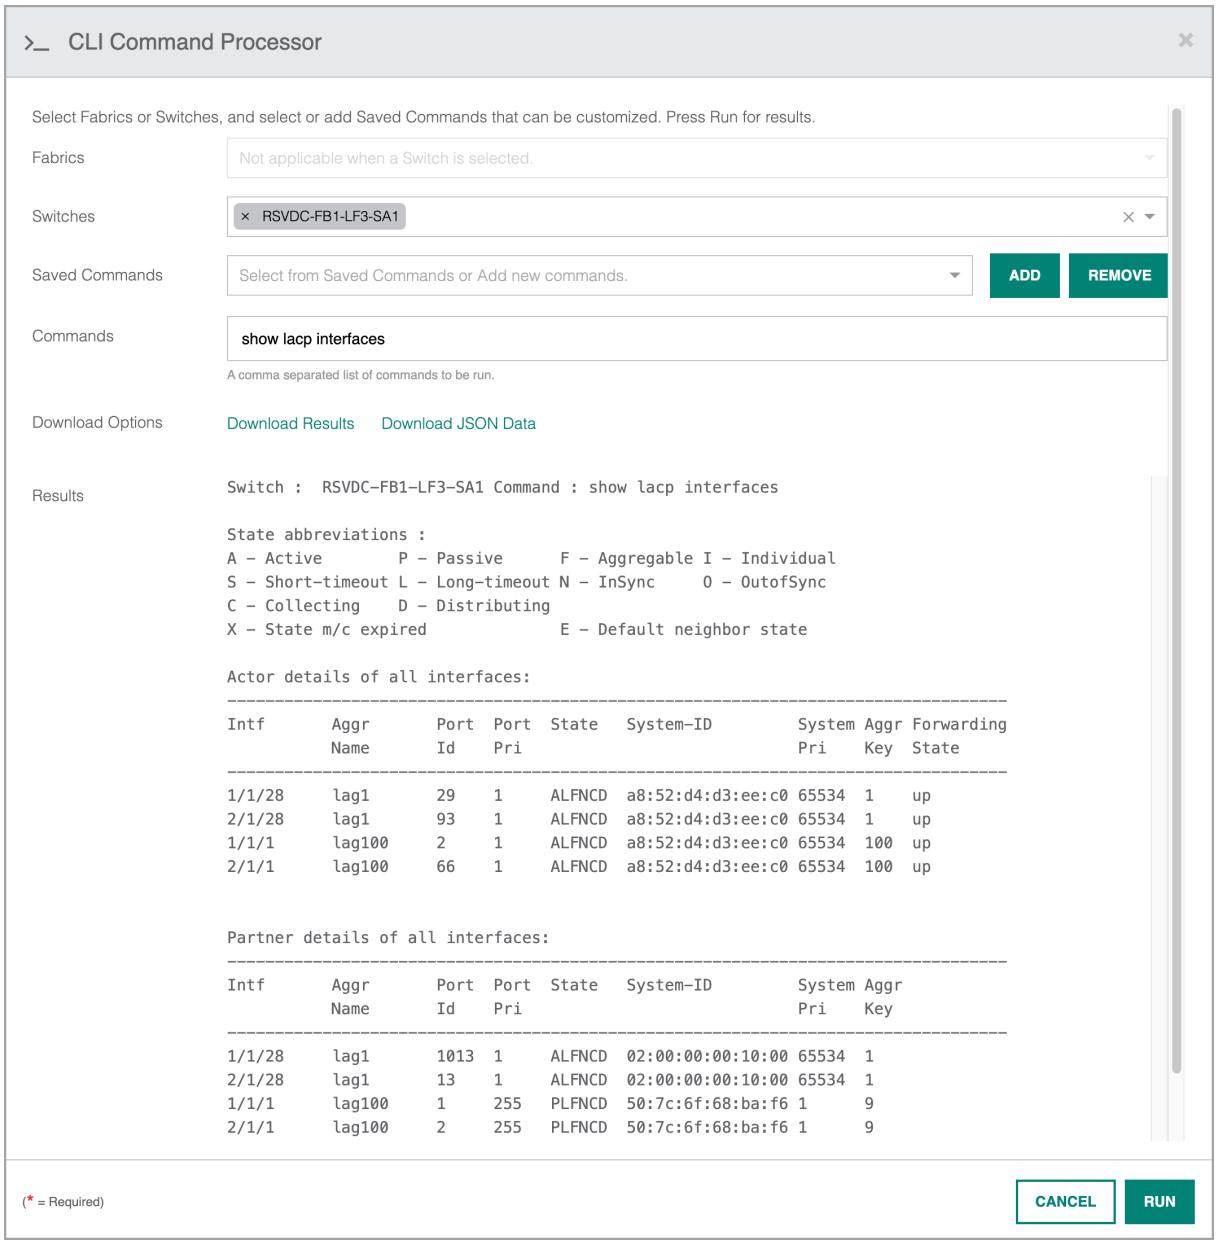


**Step 14** When a host is connected to the LAG, verify the port status for each member of the LAG. On a VSX pair, the local interface **State** should be **ALFNCD** and the partner interface should be **PLFNCD**. Verify that all interfaces in a LAG defined on a VSF stack have a **Sate** of **ALFNCD**, when connected to a host. The **Forwarding State** should be “Up” for local interfaces.

**NOTE:**

When assigning VLANs to a LAG on a server access (sub-leaf) switch, Fabric Composer automatically creates VLAN configuration for VLANs not present on the switch.

## Configure the Border Leaf

The border leaf is the ToR switch pair that connects the data center fabric to other networks such as a campus, WAN, or DMZ.

When connecting overlay networks to external networks, segmentation is preserved by establishing a distinct Layer 3 connection for each data center overlay VRF. A firewall often is used between the fabric hosts and an external network for policy enforcement, but this is not a requirement. A firewall also can be configured to permit traffic between VRFs based on policy. When connecting multiple overlay VRFs that require preserving route table separation upstream, the firewall must support VRFs or device virtualization.

The following diagram illustrates the topology and BGP peerings for connecting the production overlay VRF to an active/passive pair of upstream firewalls.

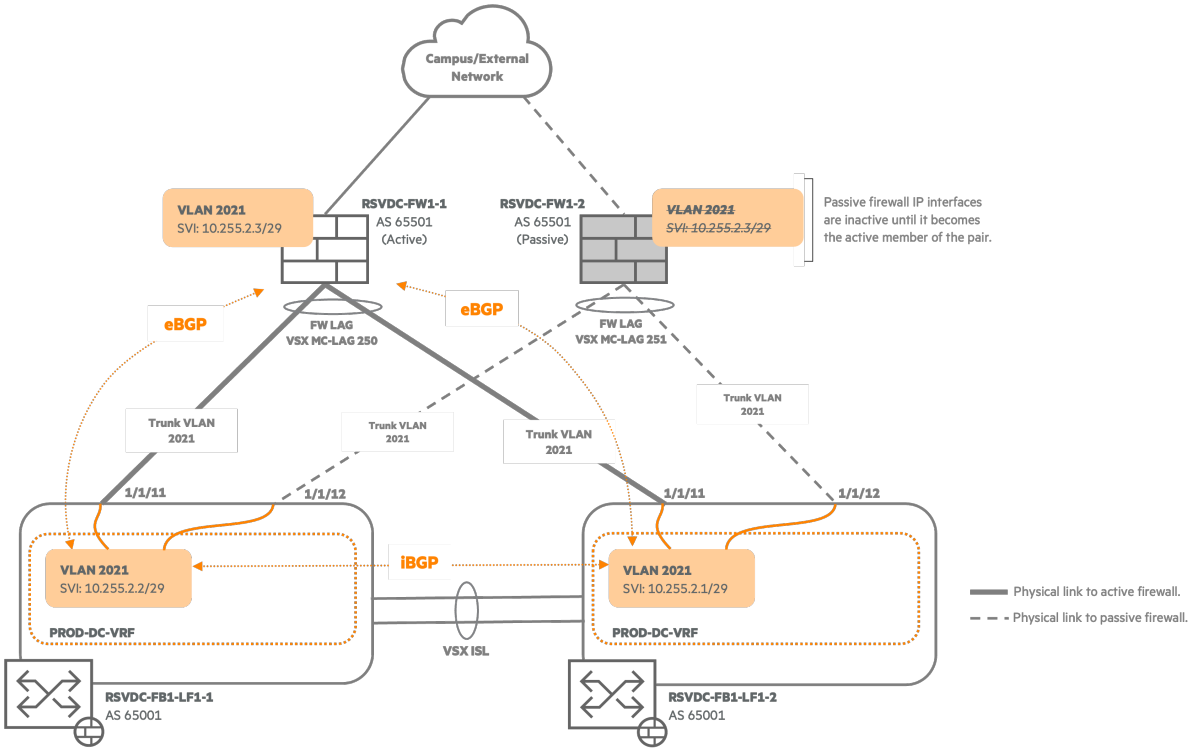


**Figure 4:** Border Leaf Topology Production VRF

An MC-LAG is used between the border leaf switch pair and each upstream firewall. This strategy provides network path redundancy to each firewall. When using an active/passive firewall, traffic is forwarded only to the active upstream firewall. Detailed firewall configuration is outside the scope of this document.

Each MC-LAG between the border leaf switches and the firewalls is an 802.1Q trunk, where one VLAN per VRF is tagged on the LAG. Tagging the same VLANs on both LAGs supports the active/passive operation of the firewall. Using VLAN tags when only one overlay VRF is present supports adding overlay VRFs in the future without additional cabling or changing port roles from access to trunk.

MP-BGP EVPN advertisements share host routes inside the data center (/32 IPv4 and /128 IPv6). EVPN host routes are commonly filtered to connections outside the data center. In the following example, only network prefixes containing overlay hosts are shared, which can be redistributed connected routes or learned within the fabric from type-5 EVPN route advertisements.

In this sample implementation, each overlay VRF on the border leaf switches learns a default route and a campus summary route from the firewalls. The border leaf shares learned external routes with other leaf switches by advertising a type-5 EVPN route.

The following diagram illustrates additional elements required when adding external connectivity to the development overlay VRF. The same set of physical links between the border leaf and the firewalls is used to connect both production and development overlay VRFs. A development VRF VLAN is tagged on the previously configured MC-LAG trunks between the border leaf switches and the firewalls to support an additional set of BGP peerings with the firewall.

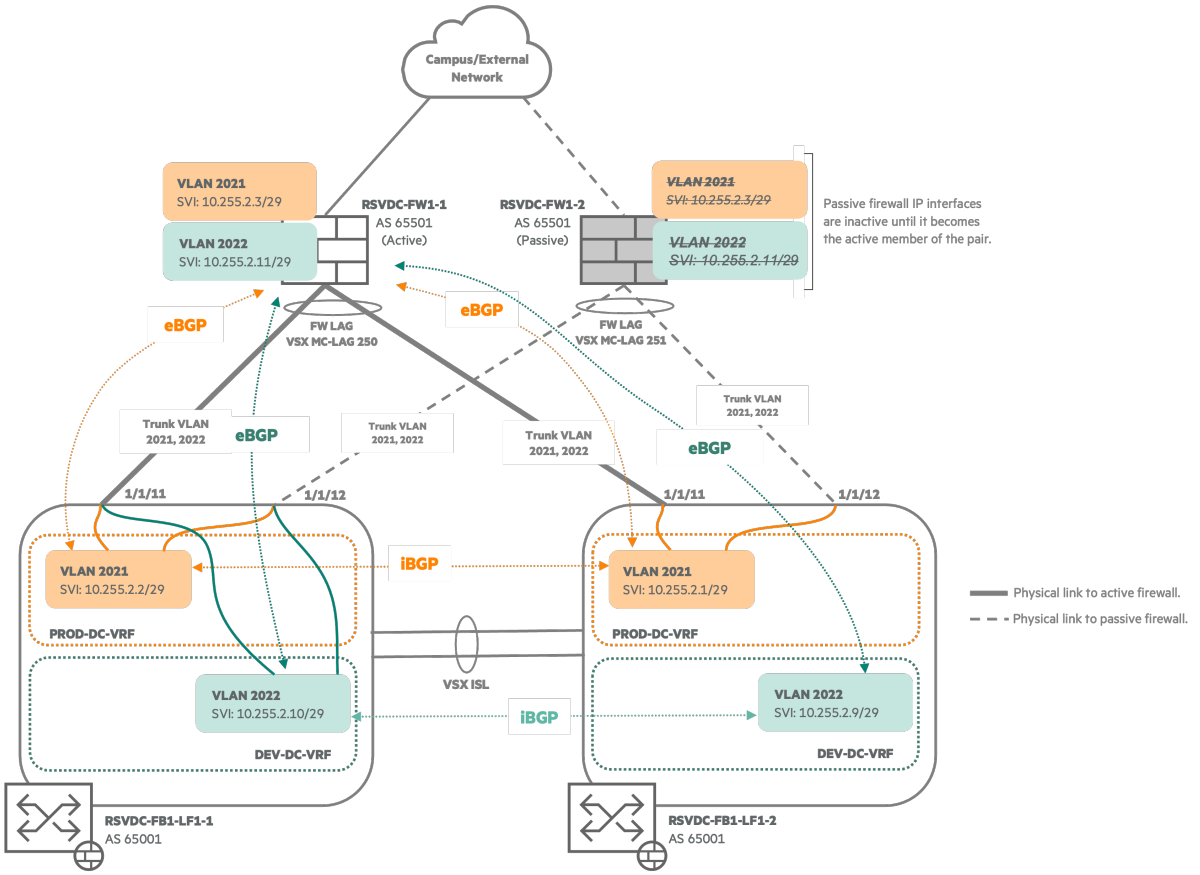


**Figure 5:** Border Leaf Topology Multiple VRFs

**NOTE:**

When using an Aruba CX 10000 in the border leaf role, physical ports connecting to external networks must be configured with *persona access*.

## Configure External Routing VLAN SVIs

In the configuration steps below, the VLAN SVIs are created to use in eBGP peerings between border leaf switches and the upstream active/passive firewall pair.

**Step 1** On the **Configuration** menu, select **Routing > VRF**.

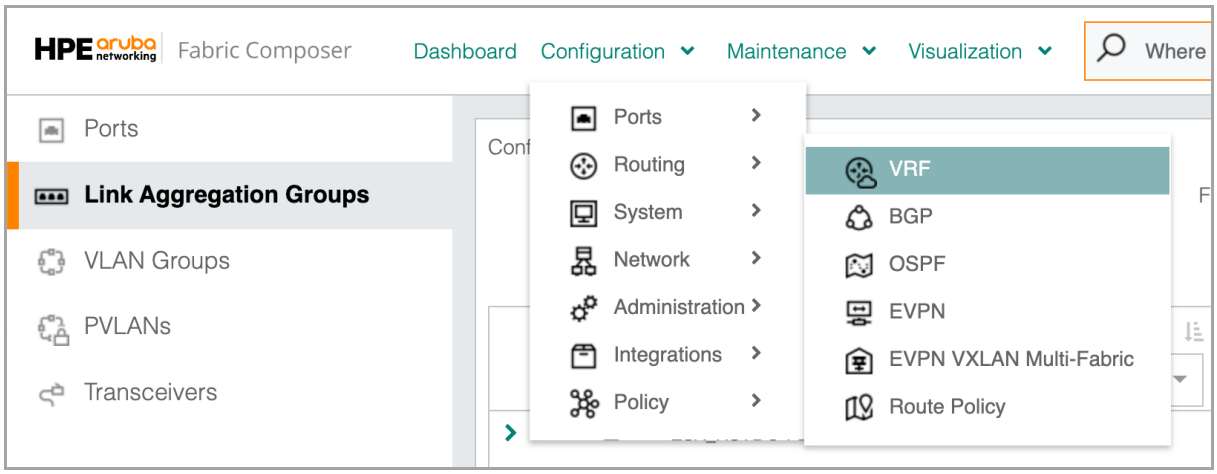


**Step 2** On the **Configuration > Routing > VRF** page, click the • • • symbol left of **PROD-DC-VRF** and select **IP Interfaces**.

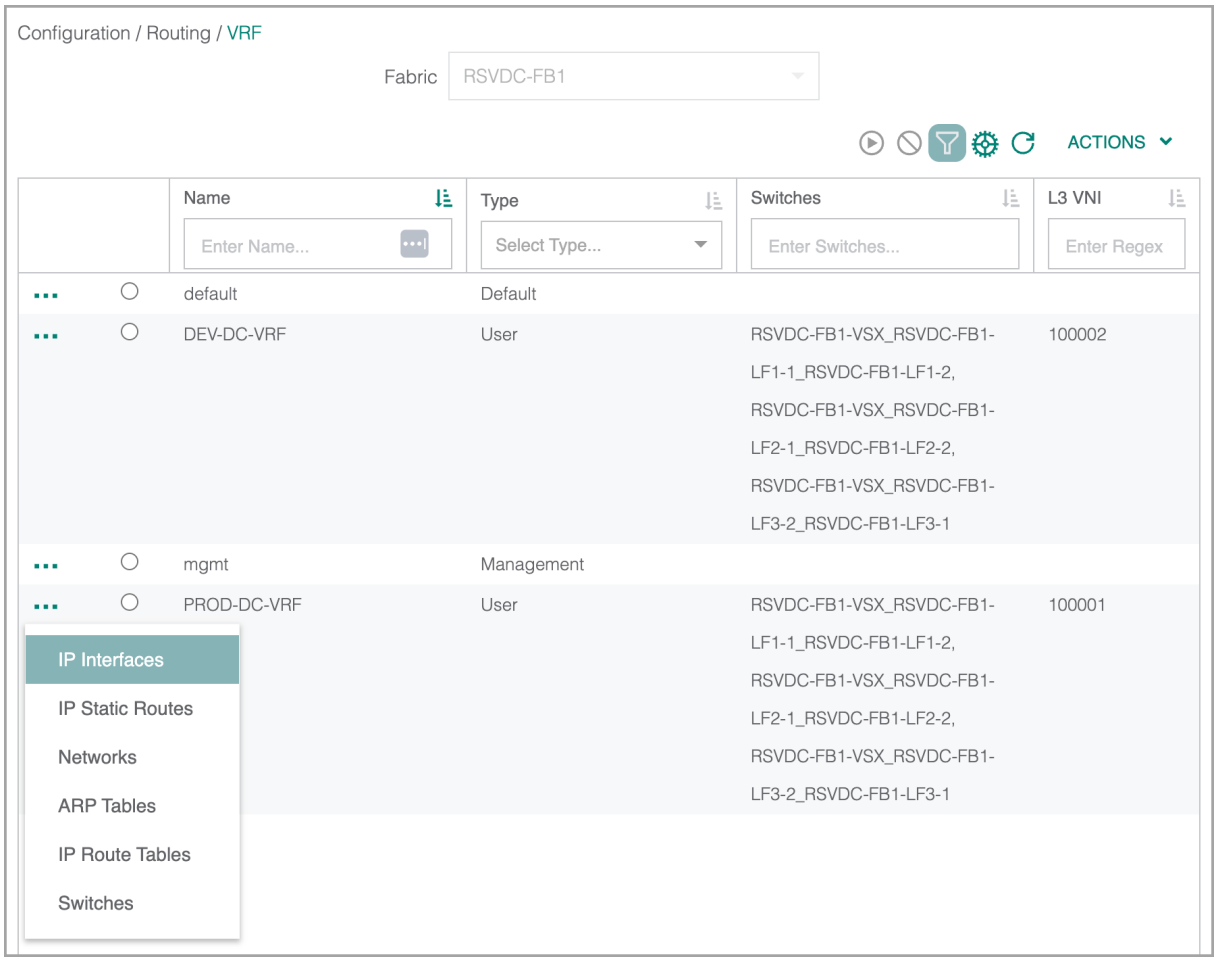


**Step 3** On the right **ACTIONS** menu of the **IP Interfaces** tab, select **Add** to launch the **IP Interfaces** wizard.

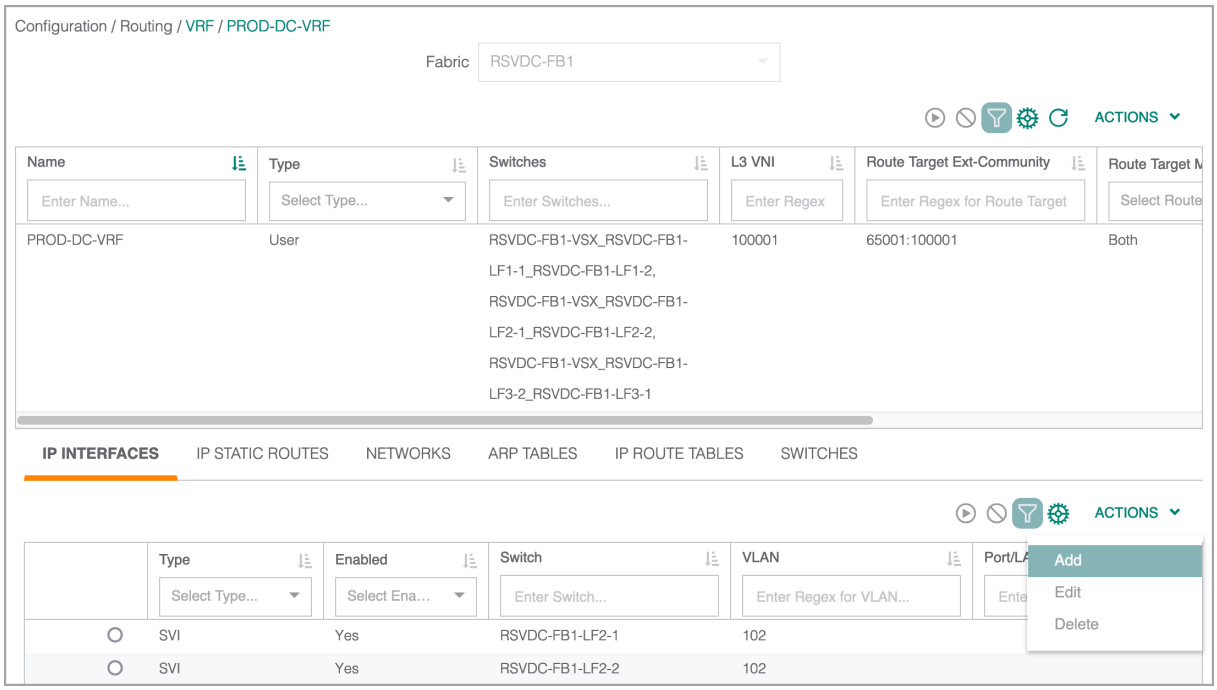


**Step 4** On the **IP Interfaces** page, enter the following values and click **NEXT**.

- **Type:** *SVI*
- **VLAN:** *2021*
- **Switches:** *< Select the border leaf VSX pair object >*
- **IPv4 Subnetwork Address:** *10.255.2.0/29*
- **IPv4 Addresses:** *10.255.2.1-10.255.2.2*
- **Active Gateway IP Address**: *< blank >*
- **Active Gateway MAC Address**: *< blank >*
- **Enable VSX Shutdown on Split**: *< unchecked >*
- **Enable VSX Active Forwarding**: *< unchecked >*
- **Enable Local Proxy ARP**: *< unchecked >*

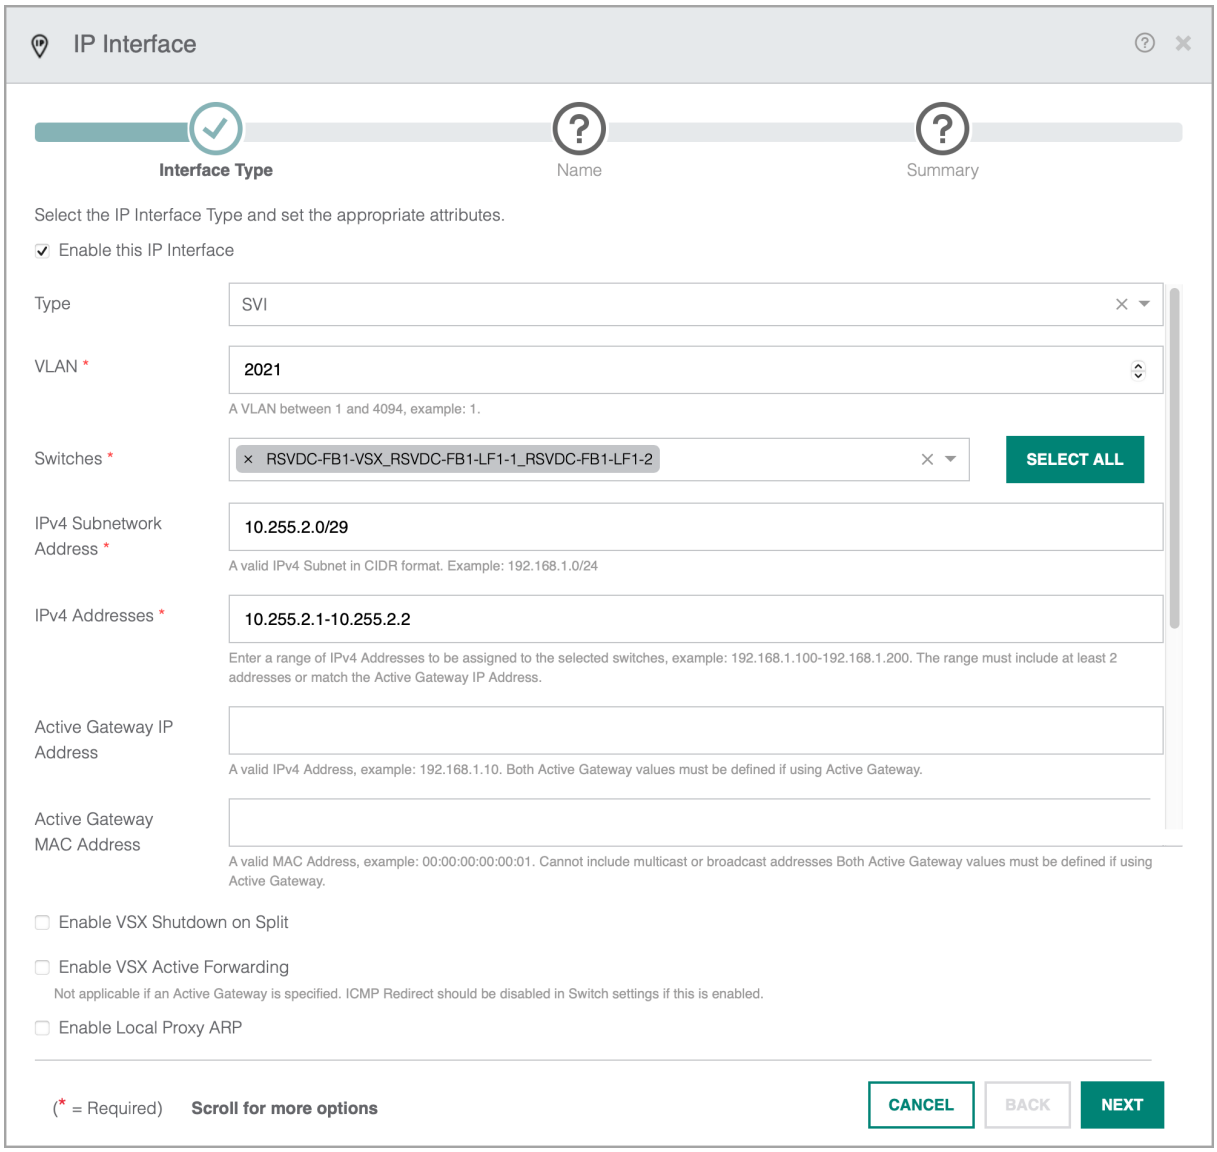


**Step 5** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

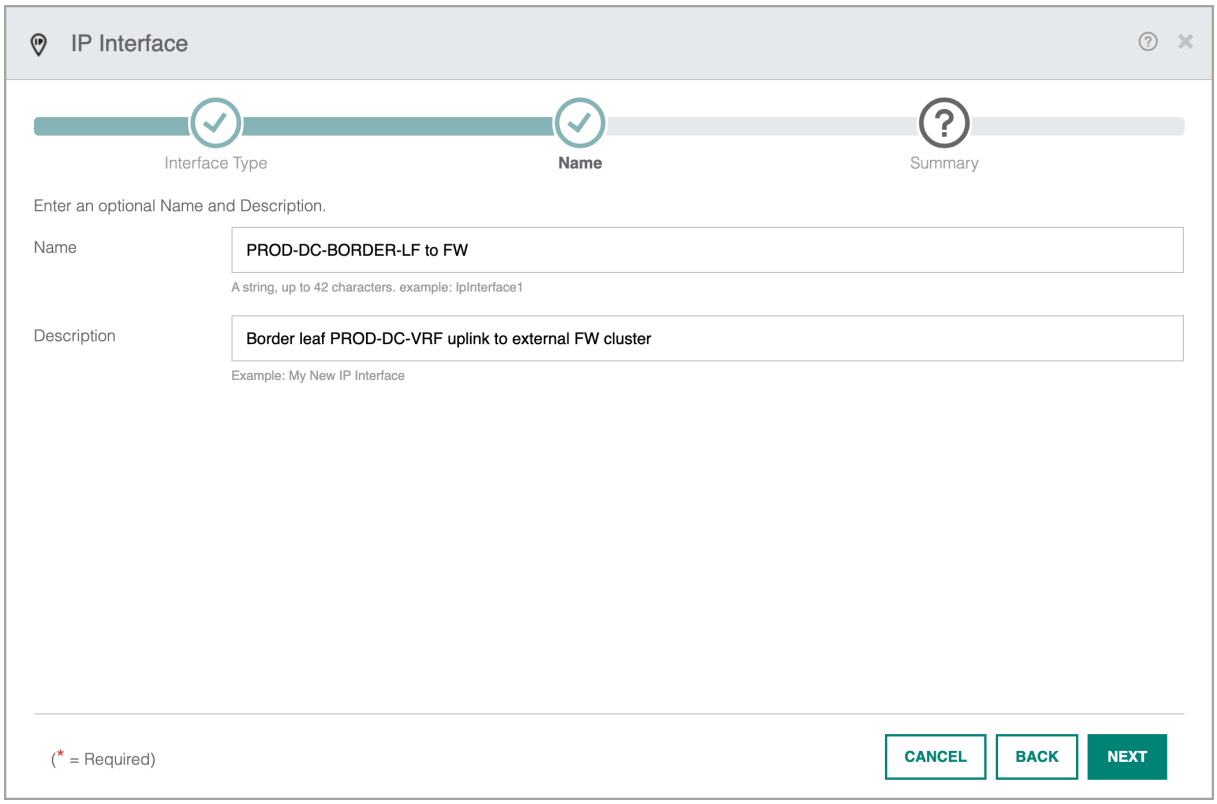


**Step 6** On the **Summary** page, review the interface settings and click **APPLY**.

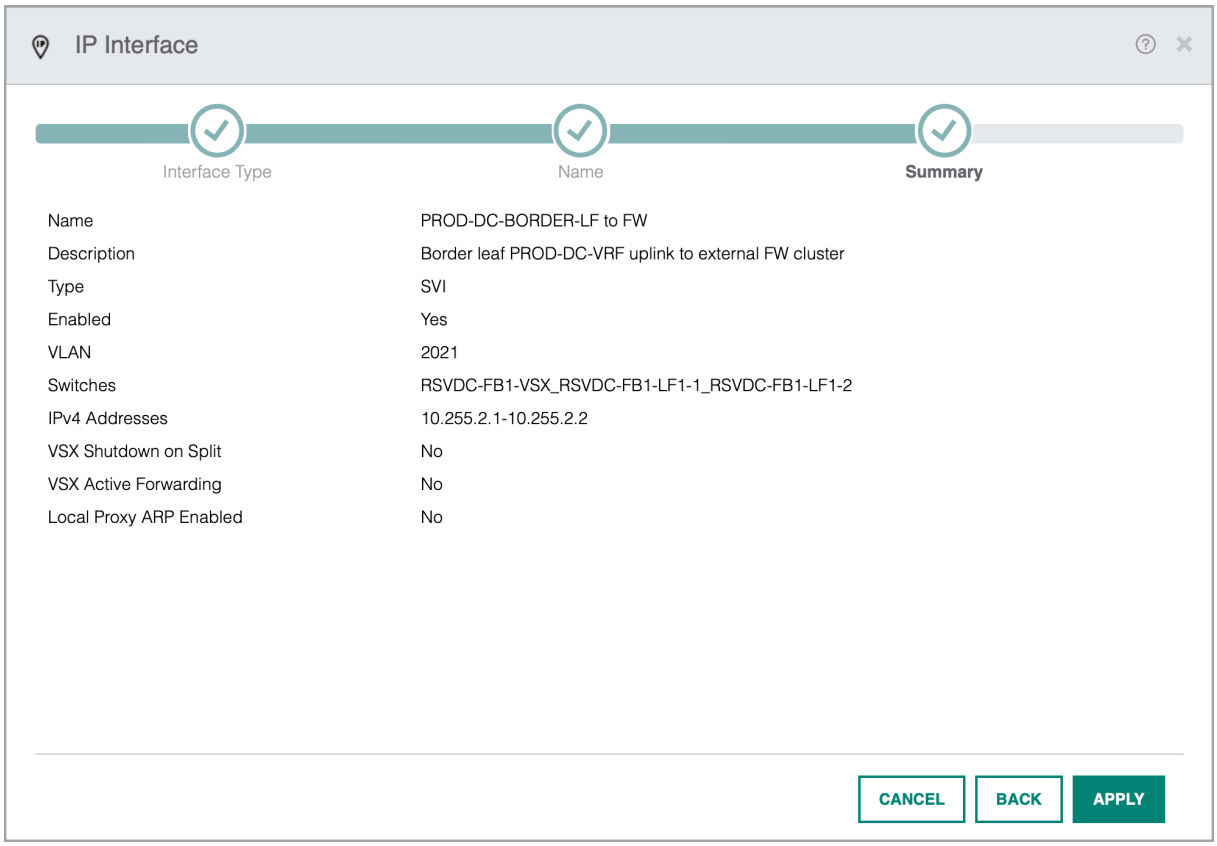


**Step 7** Repeat this procedure to create an additional VLAN and SVI interface for DEV-DC-VRF. In step 2, select **DEV-DC-VRF**, then create an SVI with the following values:

| Name | Description | Type | VLAN | Switches | IPv4 Sub-network Address | IPv4 Addresses |
|---|---|---|---|---|---|---|
| DEV-DC-BORDER-LF to FW | Border leaf DEV-DC-VRF uplink to external FW cluster | SVI | 202 | < Border leaf VSX pair object > | 10.255.2.8/2 | 10.255.2.9-10.255.2.10 |

## Create Border Leaf to Firewall MC-LAGs

A VSX-based MC-LAG is created to each individual firewall in the active/passive cluster from the border leaf switches.

**Step 1** On the **Configuration** menu, select **Ports > Link Aggregation Groups**.

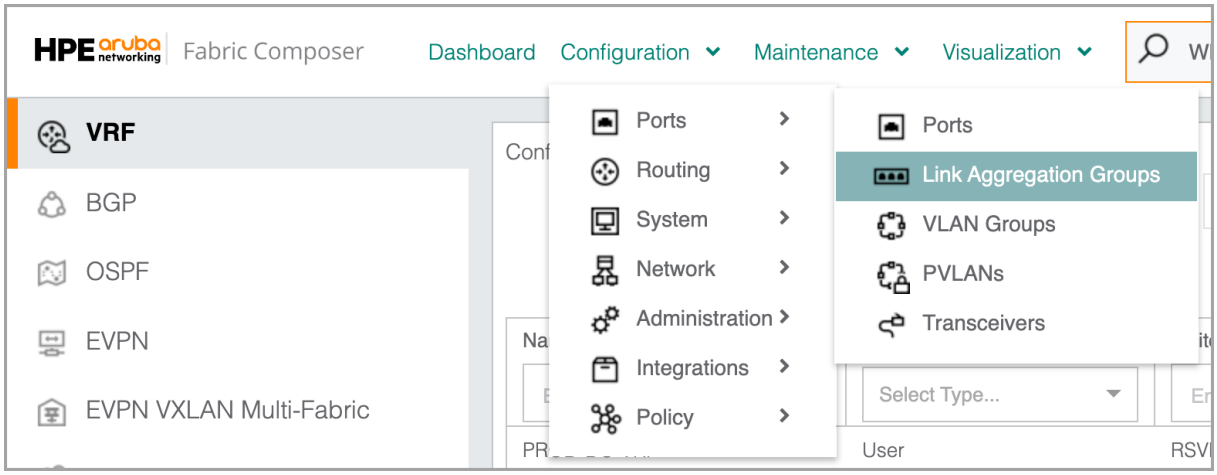


**Step 2** On the right **ACTIONS** menu, select **Add**.

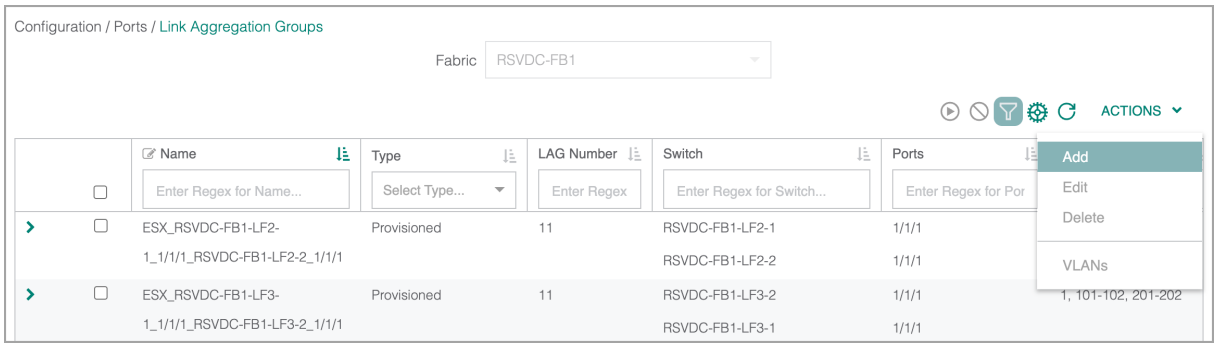


**Step 3** On the **Create Mode** page, leave **Create a single LAG/MLAG** selected and click **NEXT**.

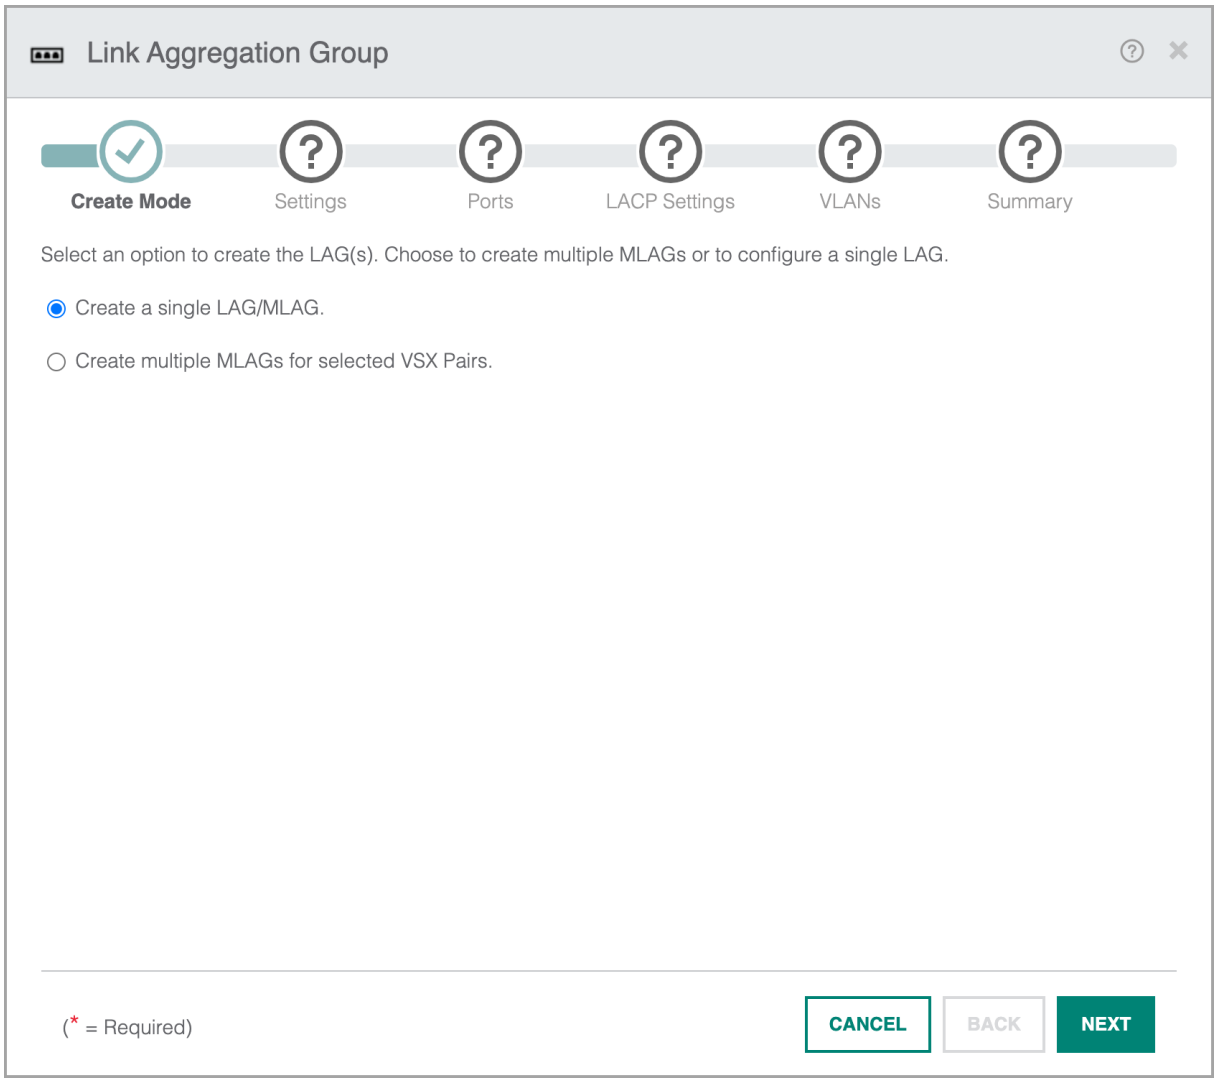


**Step 4** On the **Settings** page, enter the following values and click **NEXT**.

- **Name:** *RSVDC-BL to EXT-FW1*
- **Description:** *MC-LAG from border leaf switches to FW1 in firewall cluster*
- **LAG Number:** *251*

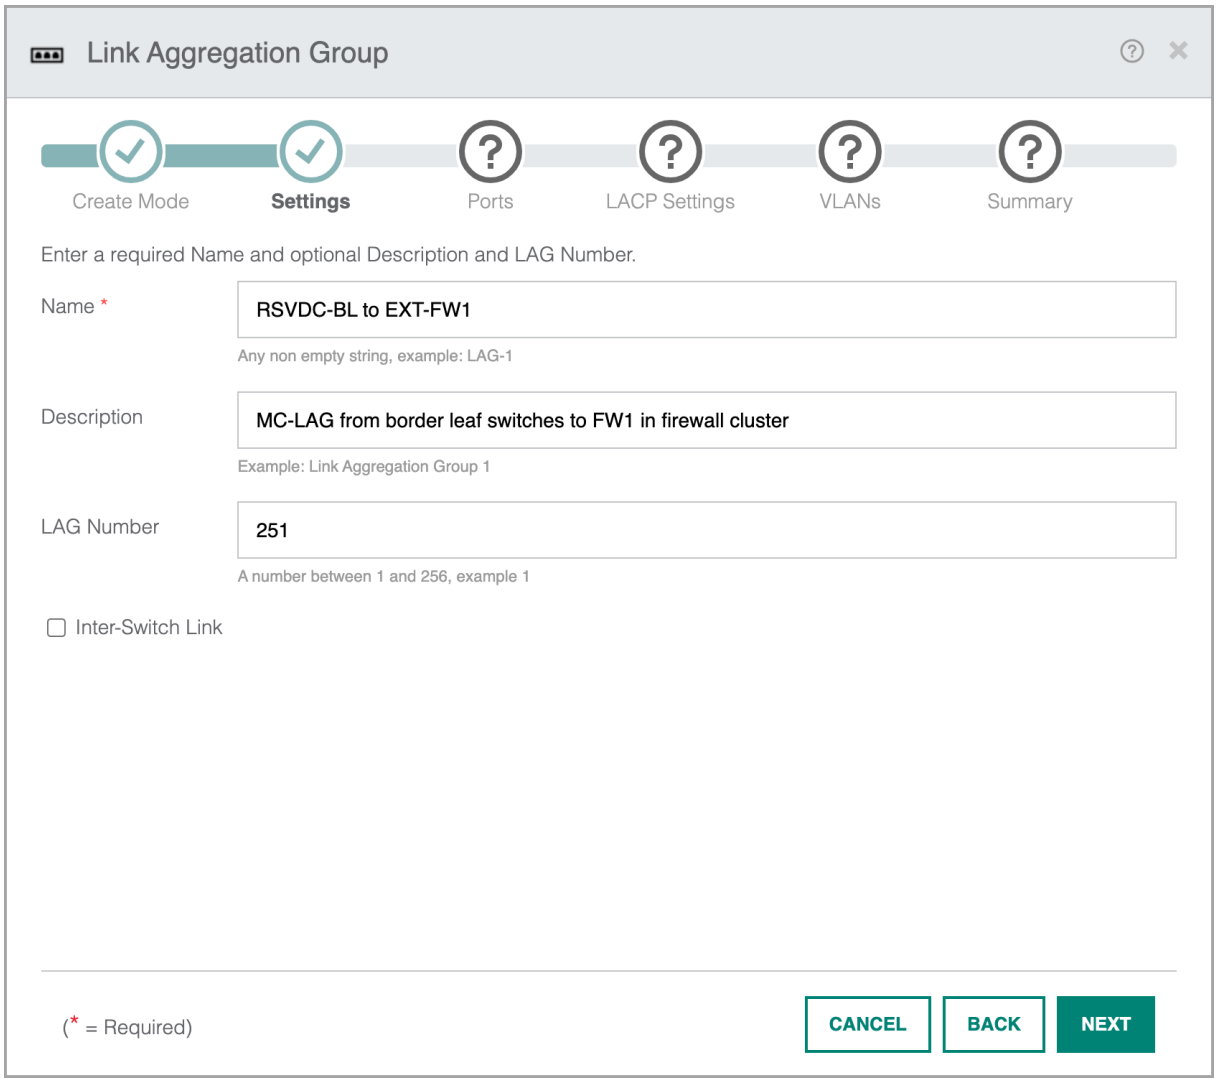


**Step 5** On the **Ports** page, select the border leaf VSX object from the **LAG Switch Member** dropdown.

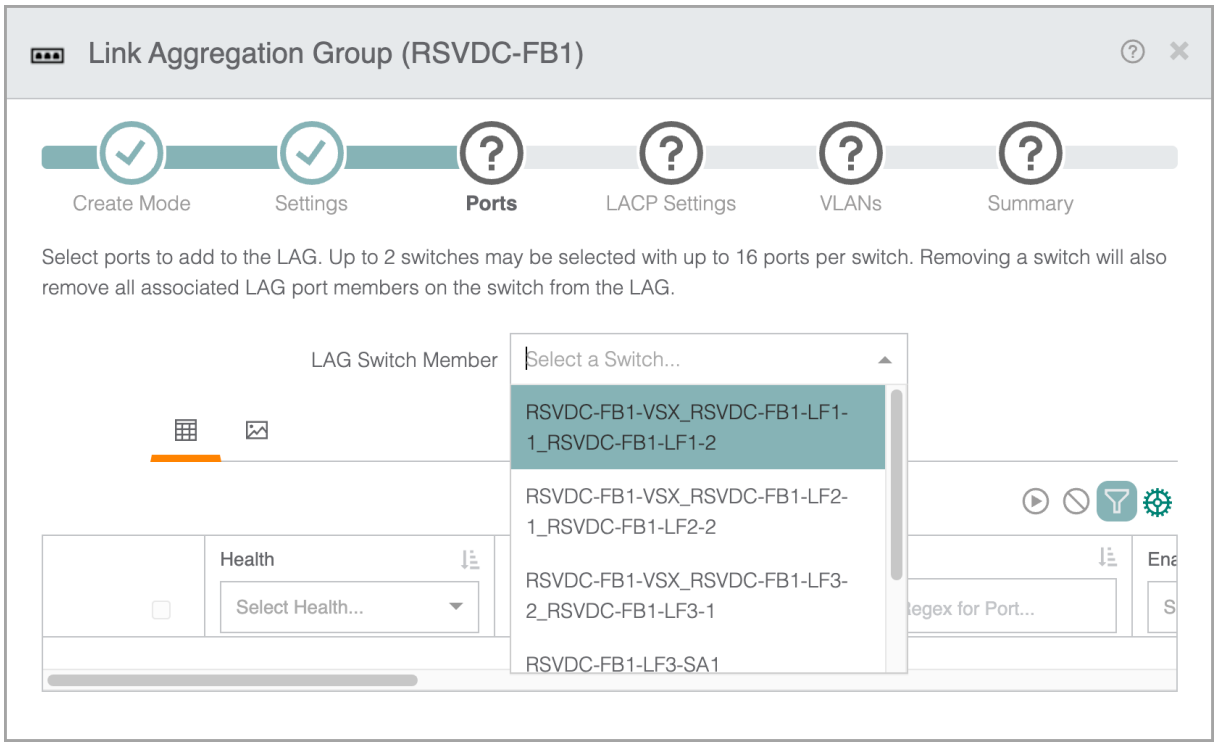


**Step 6** Click the **Switch View** mode icon to identify ports more easily. Click the port icons connected to the first firewall to add them as members of the link aggregation group and click **NEXT**.

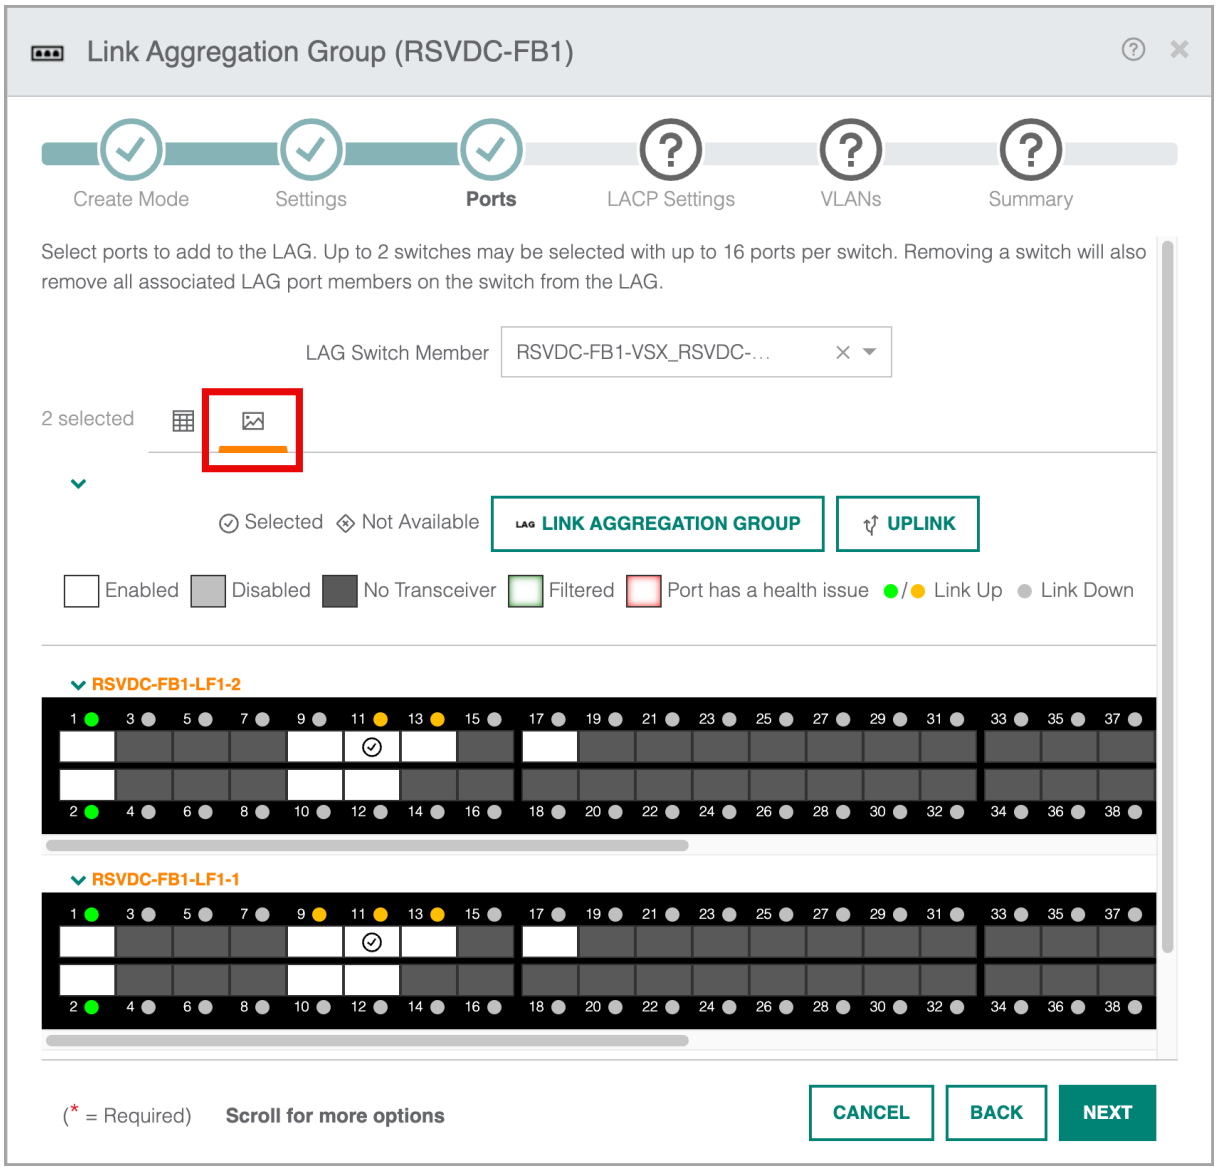


**NOTE:**

A checkmark appears on the newly selected ports. The diamond icon appears on ports that are not currently available for a new LAG group assignment.

**Step 7** On the **LACP Settings** page, leave all settings at their defaults and click **NEXT**.

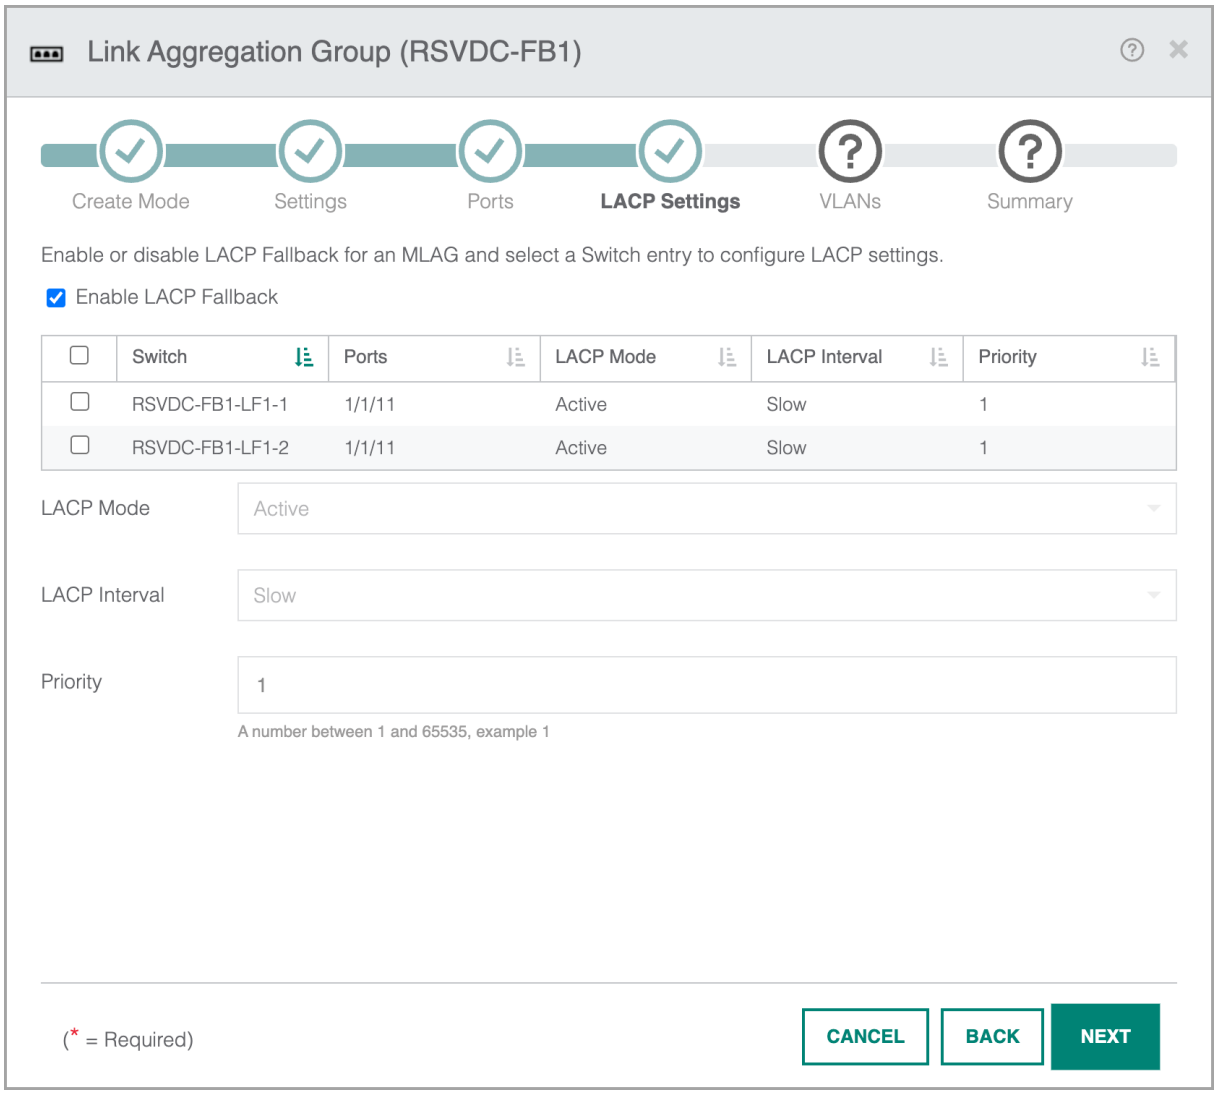


**Step 8** On the **VLANs** page, enter the VLAN ID for each VRF previously created to connect to external networks in the **VLANs** field, and click **NEXT**.

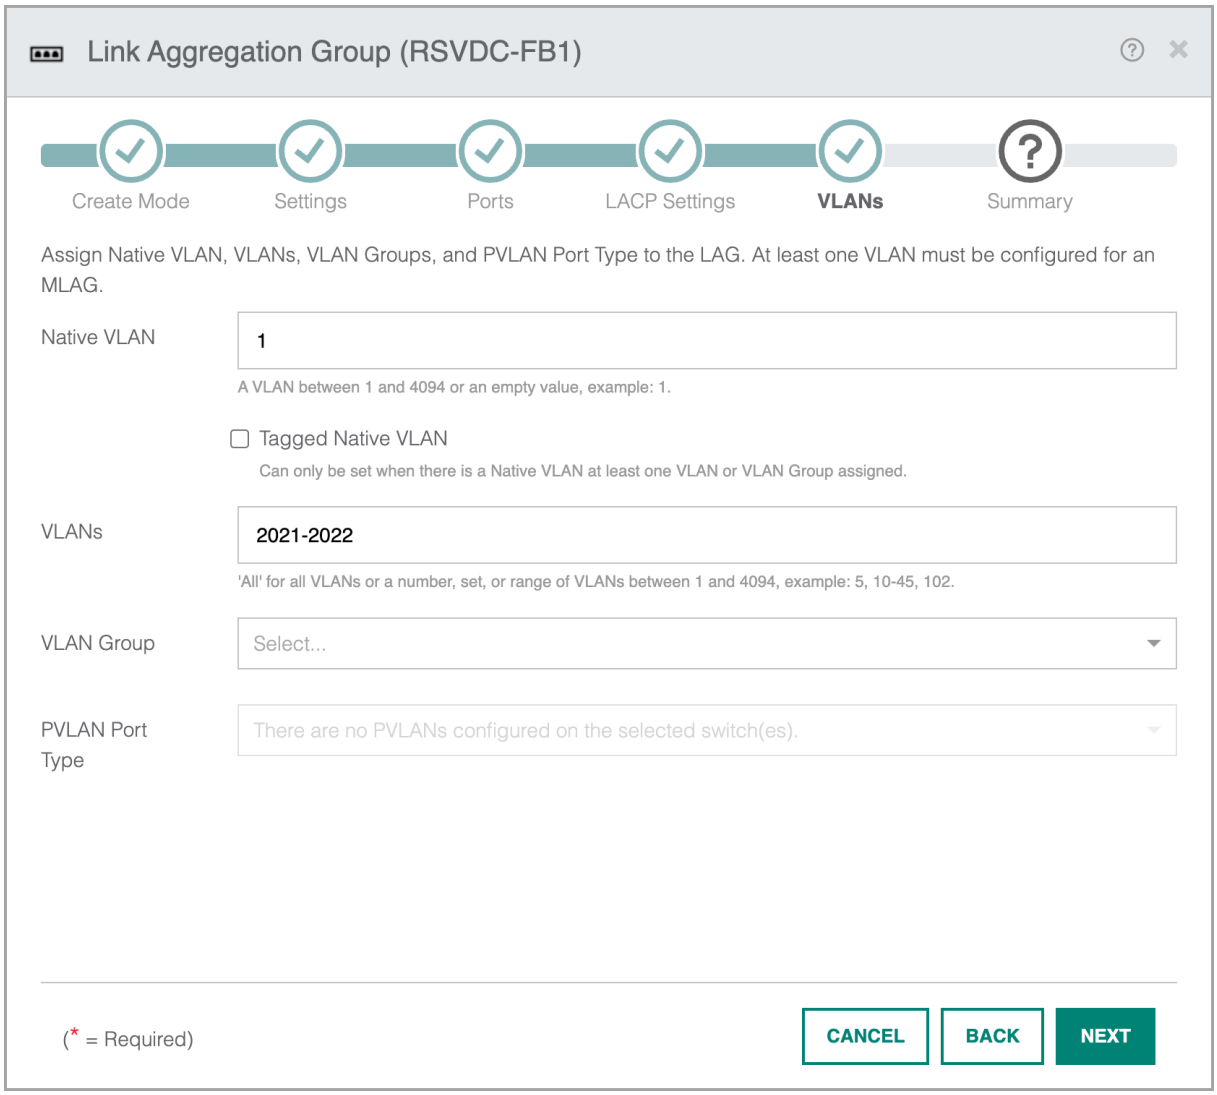


**Step 9** On the **Summary** page, review the link aggregation settings and click **APPLY**.

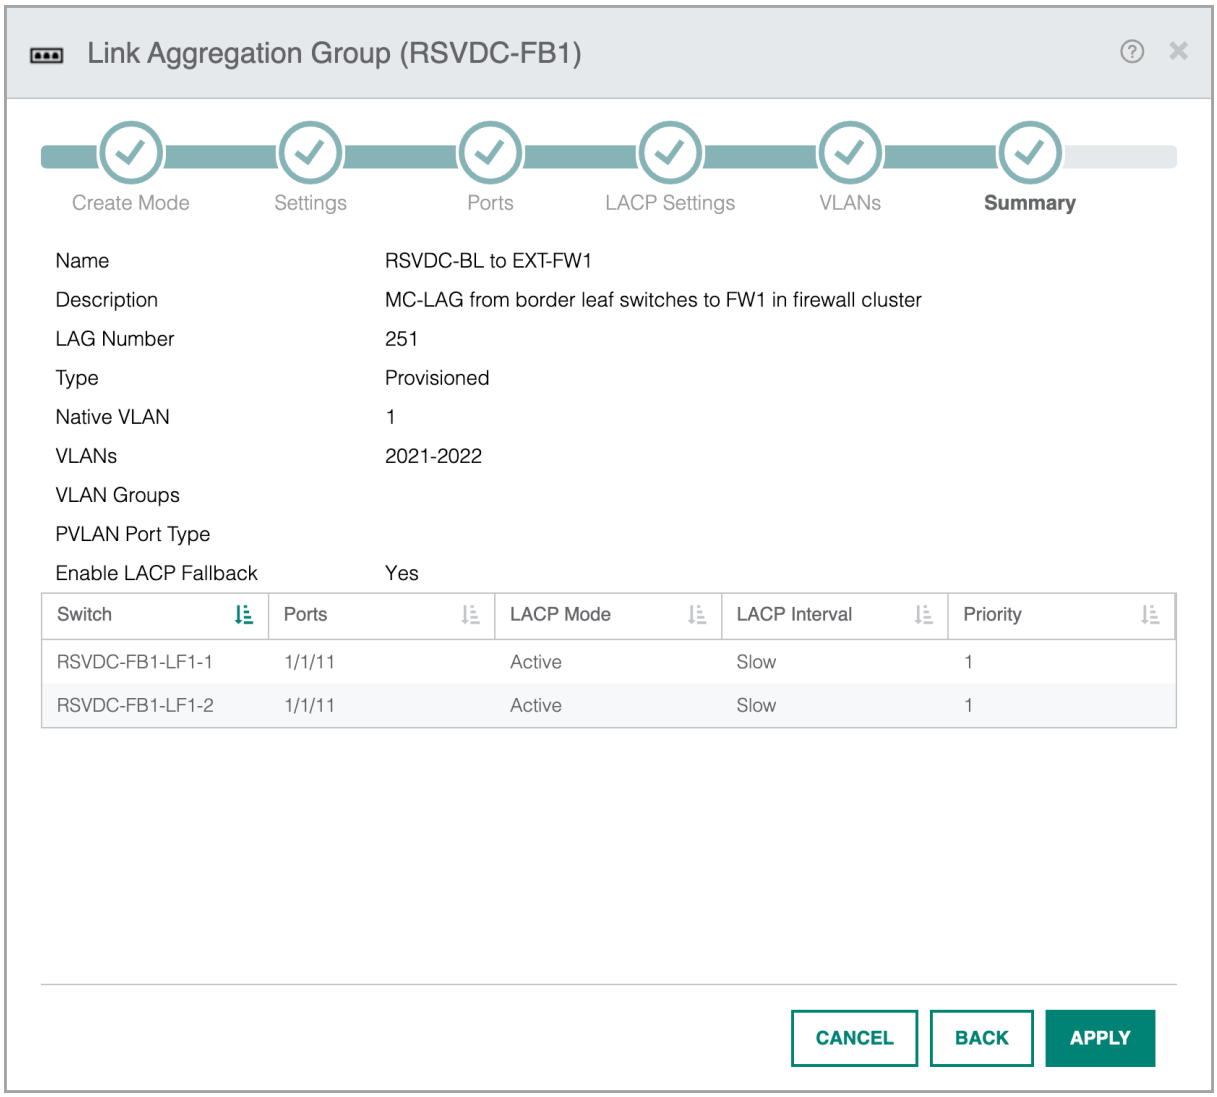


**Step 10** Repeat the procedure to create an additional MC-LAG on ports connecting the VSX border leaf switch pair to the second firewall using the following settings:

| Name | Description | LAG Number | Ports | LACP Settings | VLANs |
|---|---|---|---|---|---|
| RSVDC-BL to EXT-FW2 | MC-LAG from border leaf switches to FW2 in firewall cluster | 252 | 12 (on each switch member) | < Leave all defaults > | 2021-2022 |

## Configure Host Filter Prefix List

Host routes and point-to-point link prefixes should not be advertised to external networks. The following procedure creates a prefix list used in route policy to filter /31 and /32 IPv4 prefix advertisements.

**Step 1** On the **Configuration** menu, select **Routing > Route Policy**.

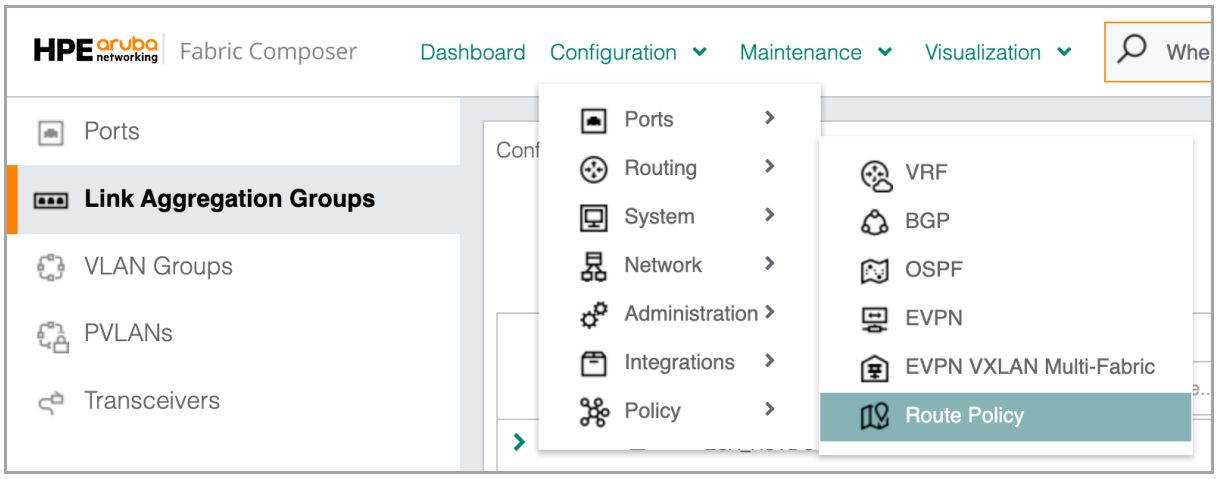


**Step 2** Click the **PREFIX LISTS** tab. On the right **ACTIONS** menu, select **Add**.

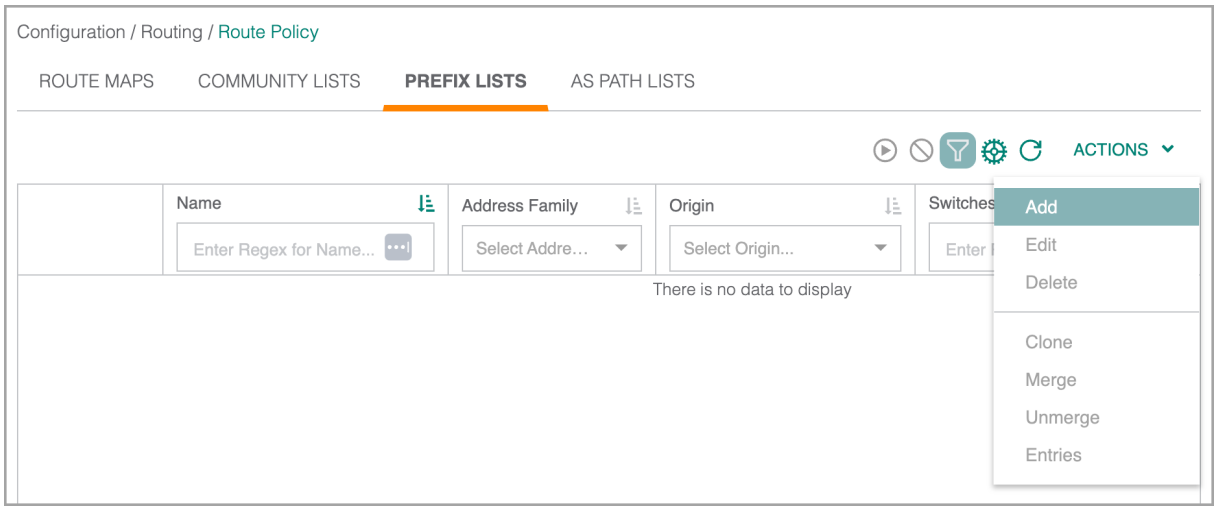


**Step 3** On the **Settings** page, enter a **Name** and **Description**, then click **NEXT**.

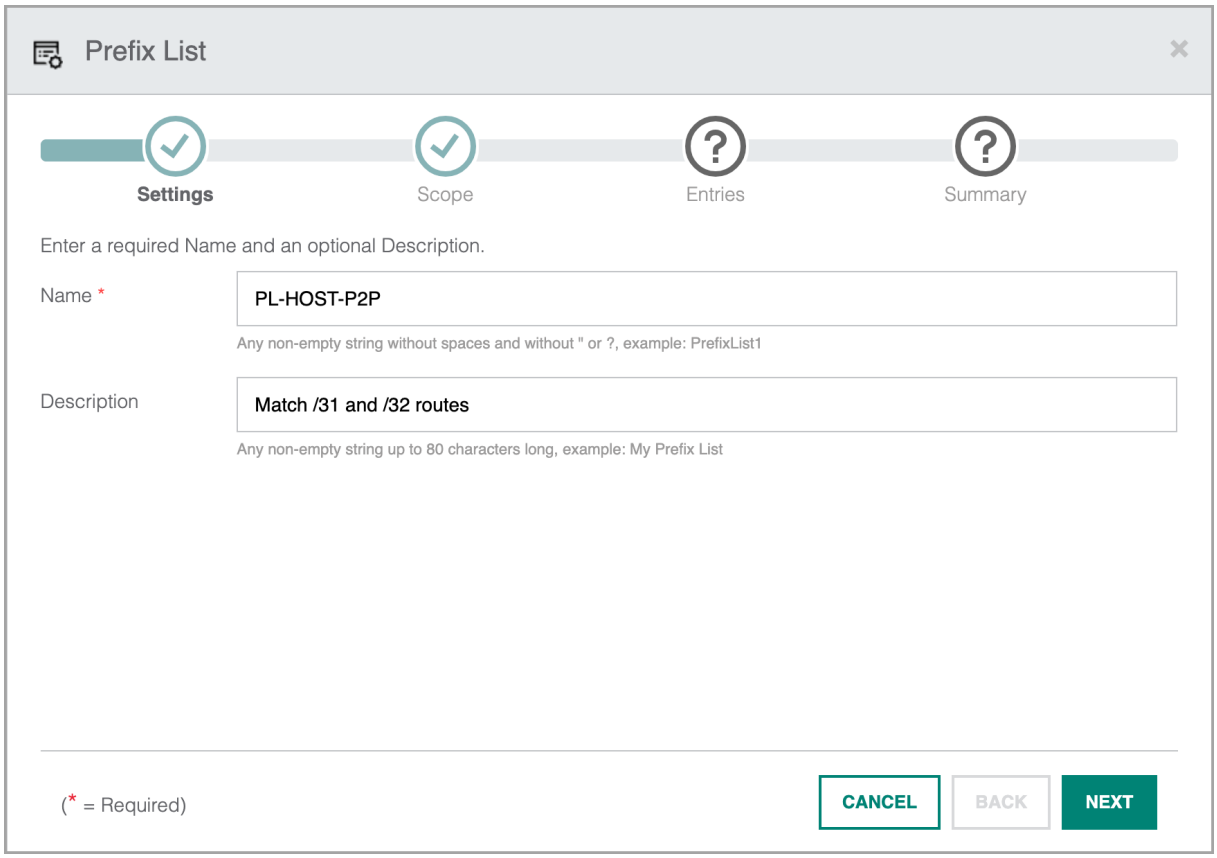


**NOTE:**

The **Name** value defines the name of the prefix list in Fabric Composer and on the switch.

**Step 4** On the **Scope** page, select the two border leaf switches in the **Switches** field, then click **NEXT**.

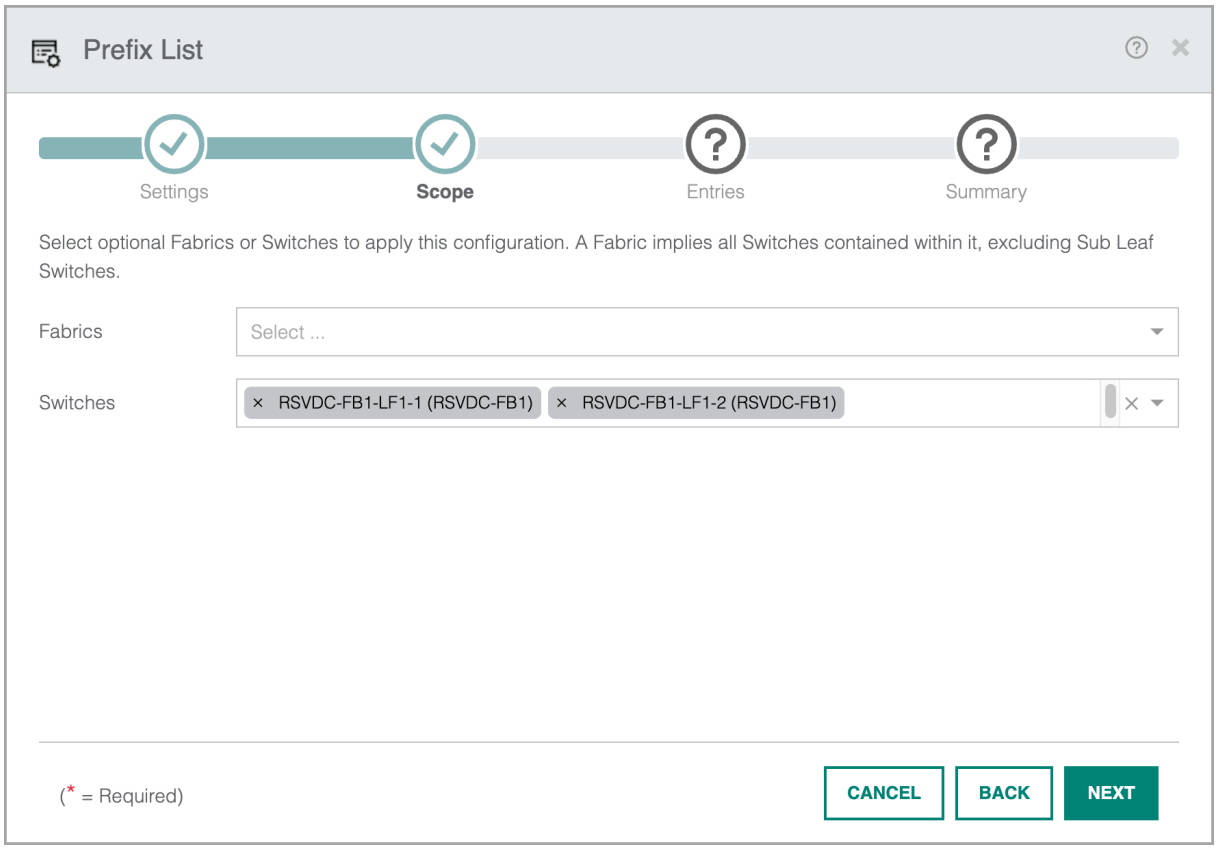


**Step 5** On the **Entries** page, enter the following non-default values and click **Add**

- **Action:** *Permit*
- **Prefix:** *0.0.0.0/0*
- **GE:** *31*

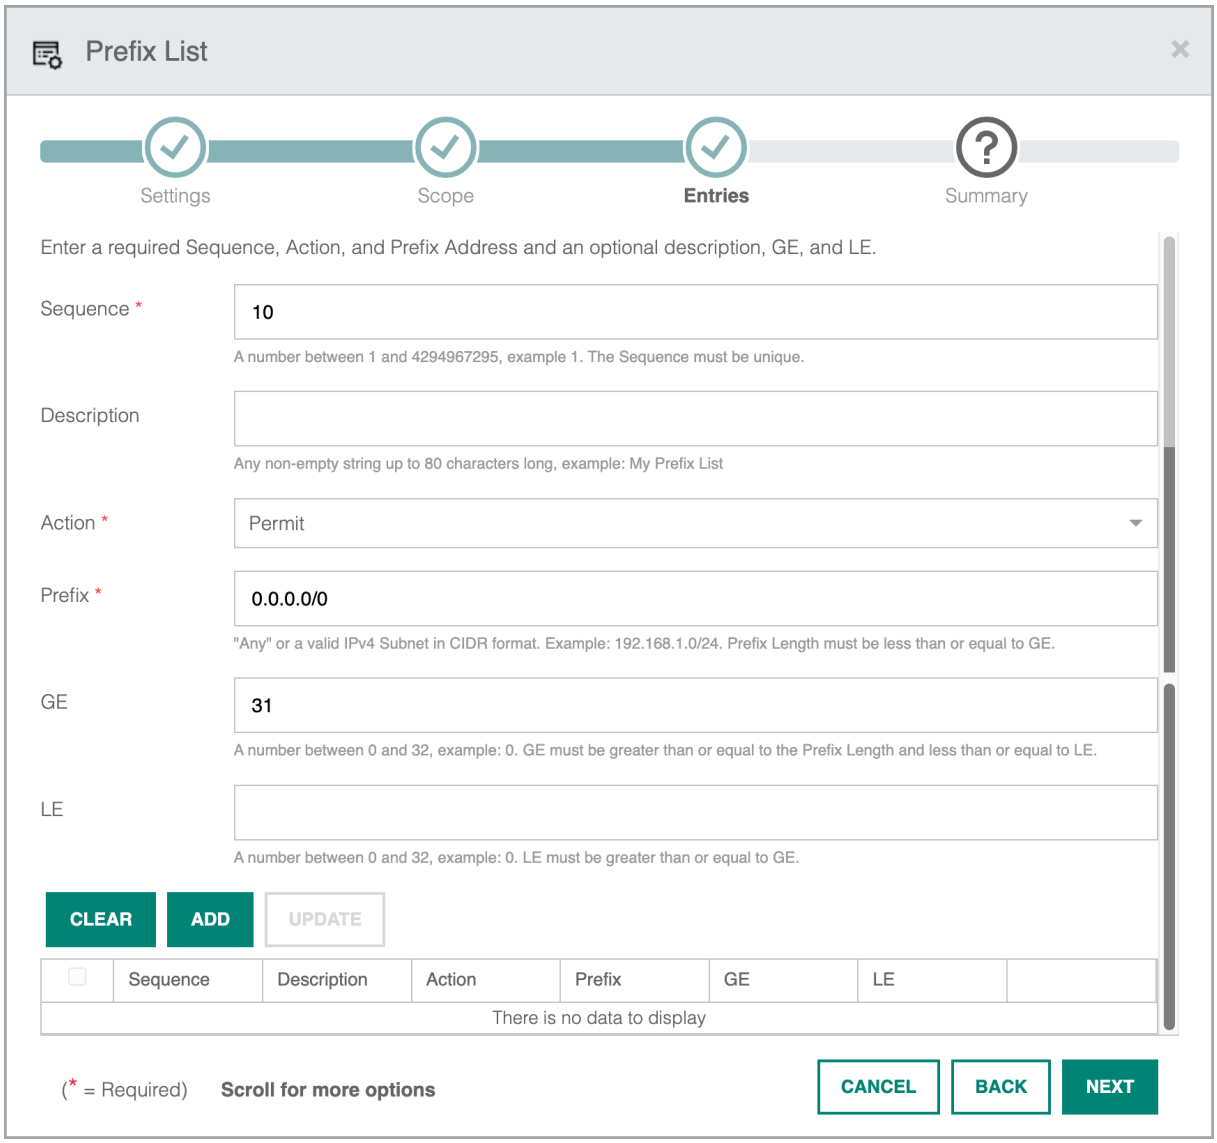


**Step 6** Click **NEXT**.

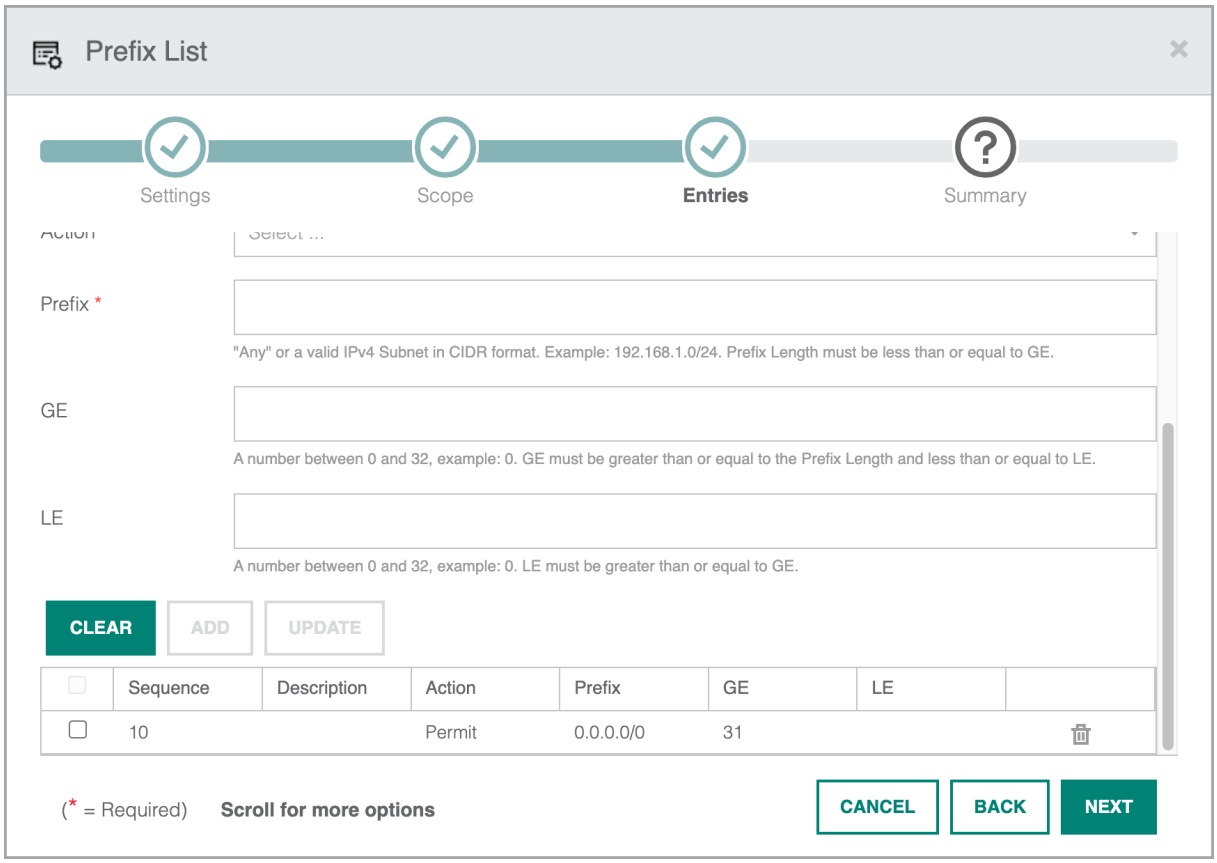


**Step 7** On the **Summary** page, review the prefix list settings and click **APPLY**.

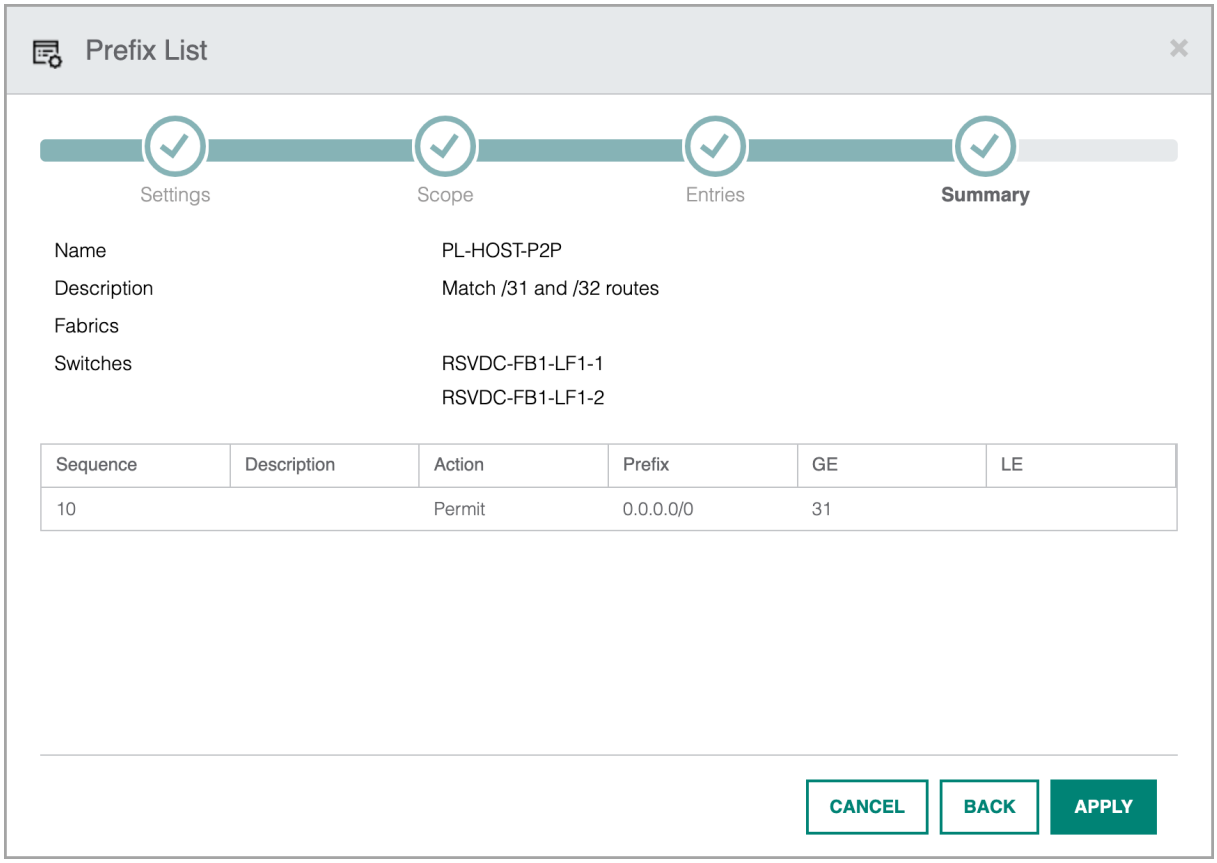

## Configure Campus AS Path List

An internal BGP peering is established between the border leaf pair to create a routed backup path to the upstream firewall. IP prefixes learned in the fabric should not be advertised in the overlay BGP peering between the border leaf pair to avoid a routing loop. The following procedure creates an AS path list that matches only prefix advertisements sourced from the upstream firewall and campus routers.

**Step 1** Click the **AS PATH LISTS** tab. On the right **ACTIONS** menu, select **Add**.

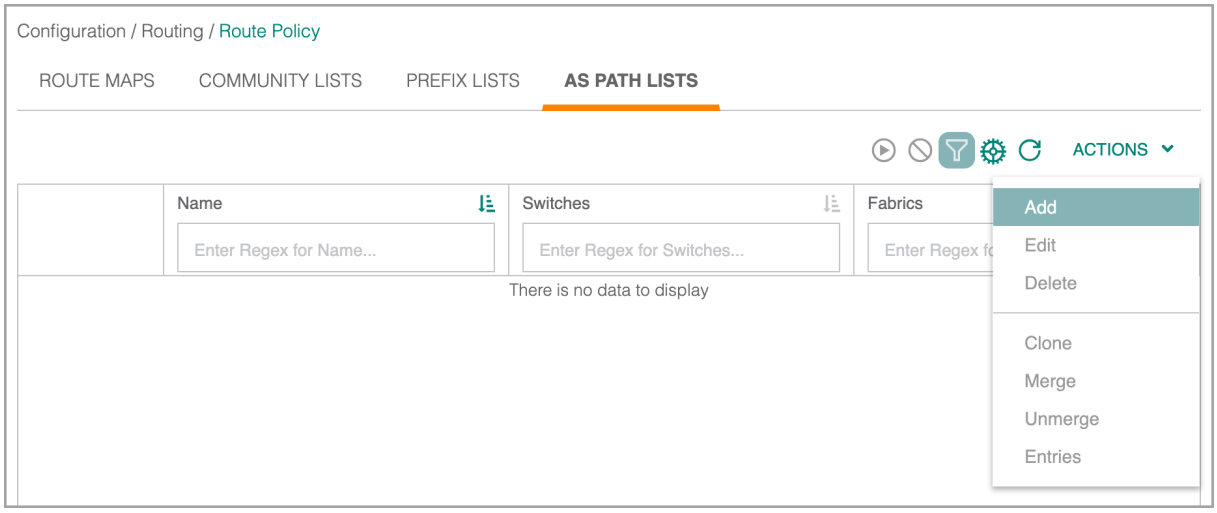


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

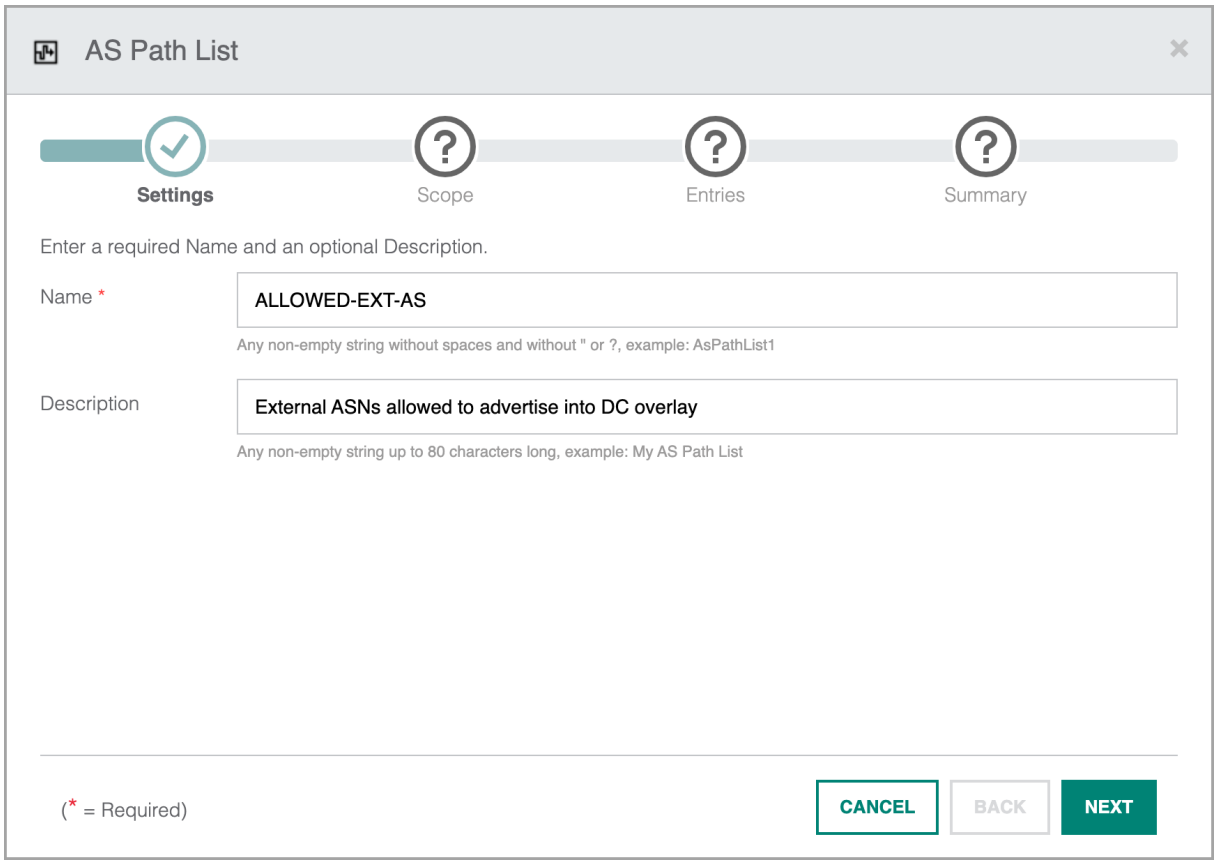


**Step 3** On the **Scope** page, select the two border leaf switches in the **Switches** field, then click **NEXT**.

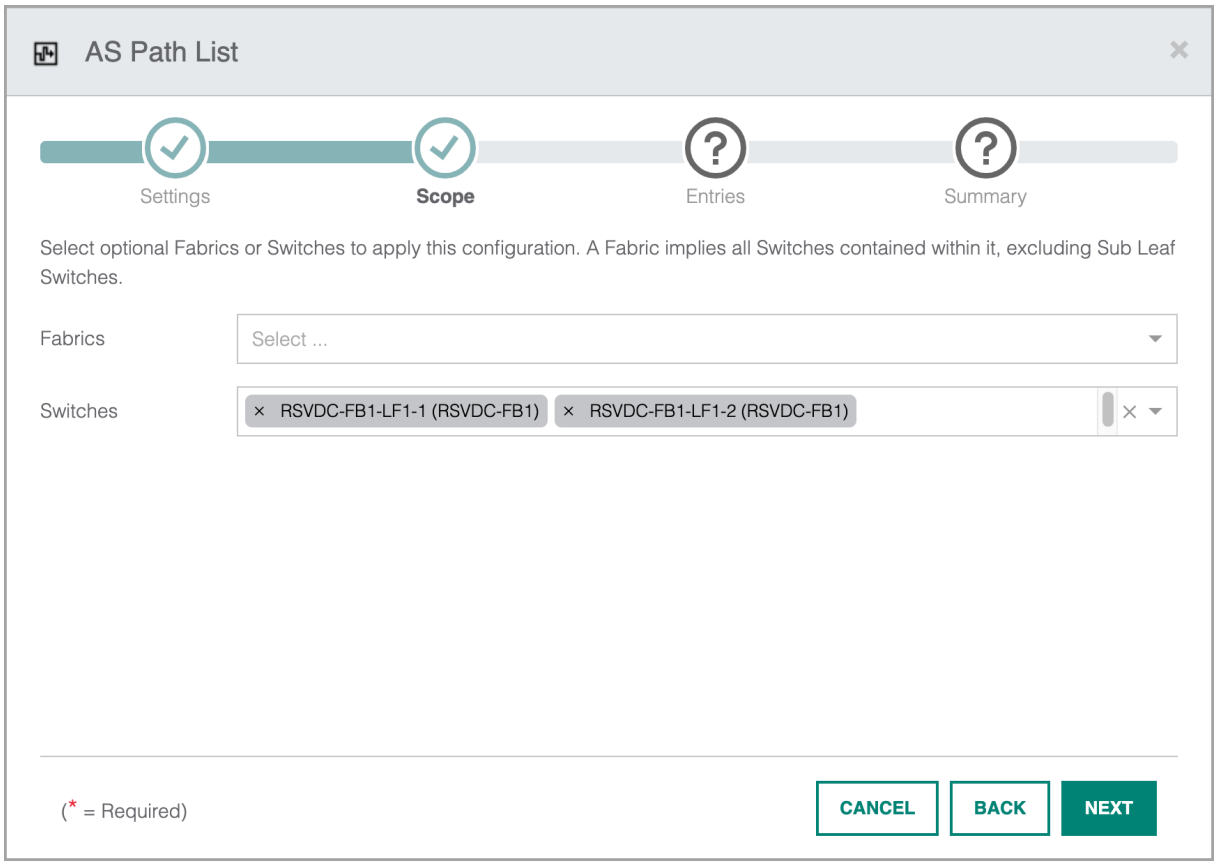


**Step 4** On the **Entries** page, enter the following values and click **ADD**.

- **Sequence**: *10*
- **Description**: *permit campus originated advertisements*
- **Action**: *Permit*
- **Regex**: *^65501 65000$*

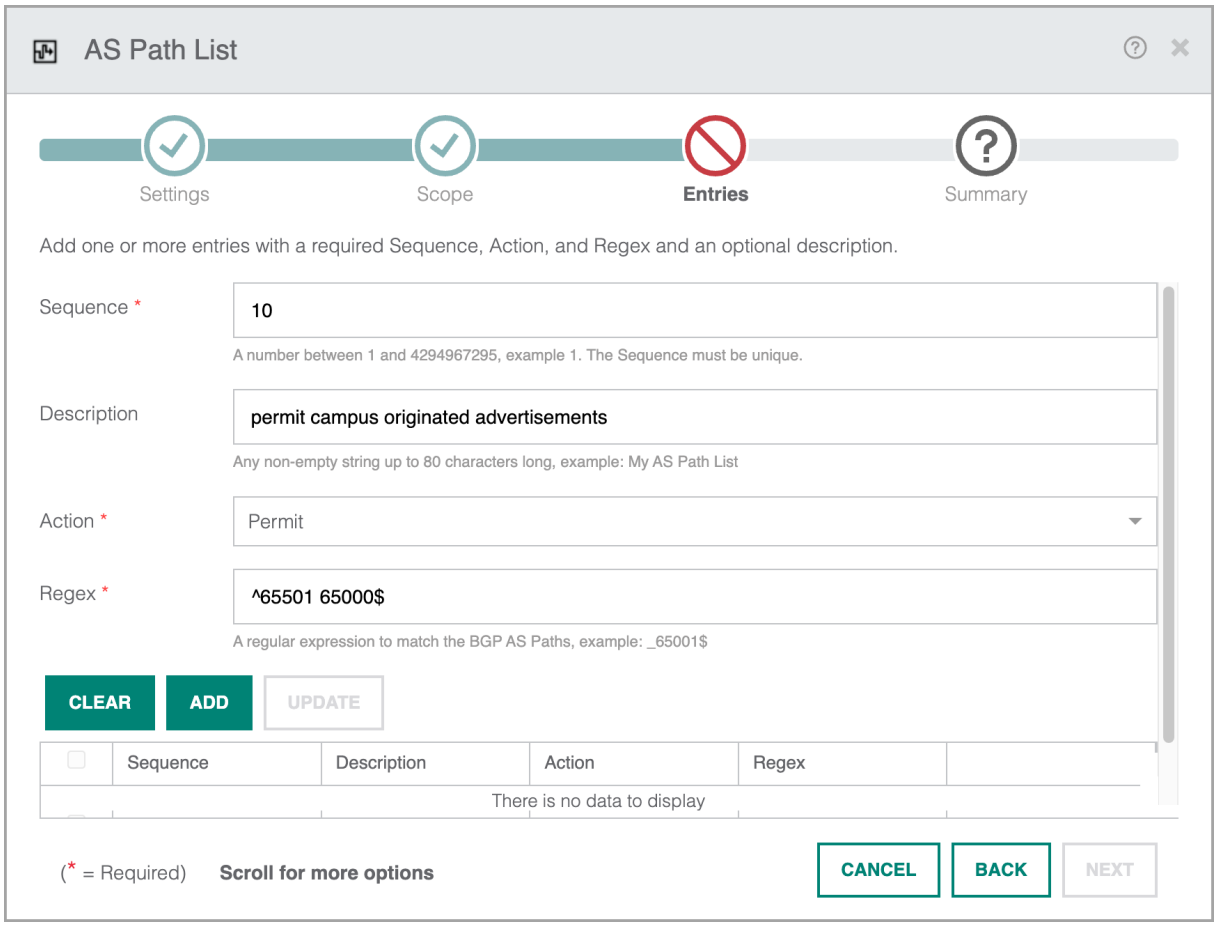


**NOTE:**

The **Regex** field value matches BGP advertisements originated by the campus AS (65000) that are received by the RSVDC fabric border leaf via the firewall AS (65501). Routes advertised by the campus that are received from other external AS numbers are not accepted.

**Step 5** On the **Entries** page, enter the following values and click **ADD**.

- **Sequence**: *20*
- **Description**: *permit firewall originated advertisements*
- **Action**: *Permit*
- **Regex**: *^65501$*

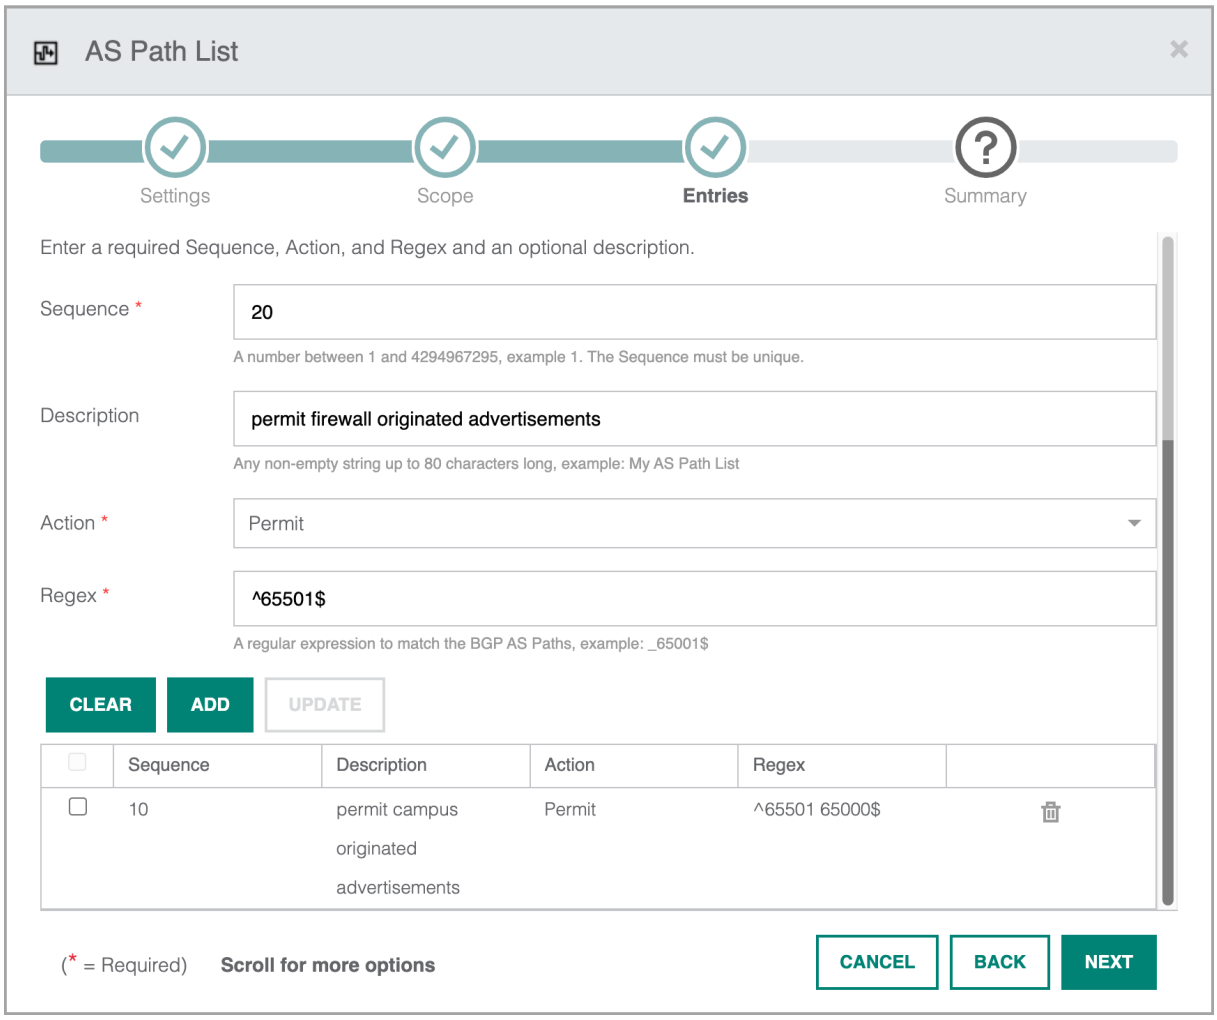


**NOTE:**

The **Regex** field value matches BGP advertisements originated by the firewall AS. In this example topology, the default route is originated by the firewall.

**Step 6** Click **NEXT**.

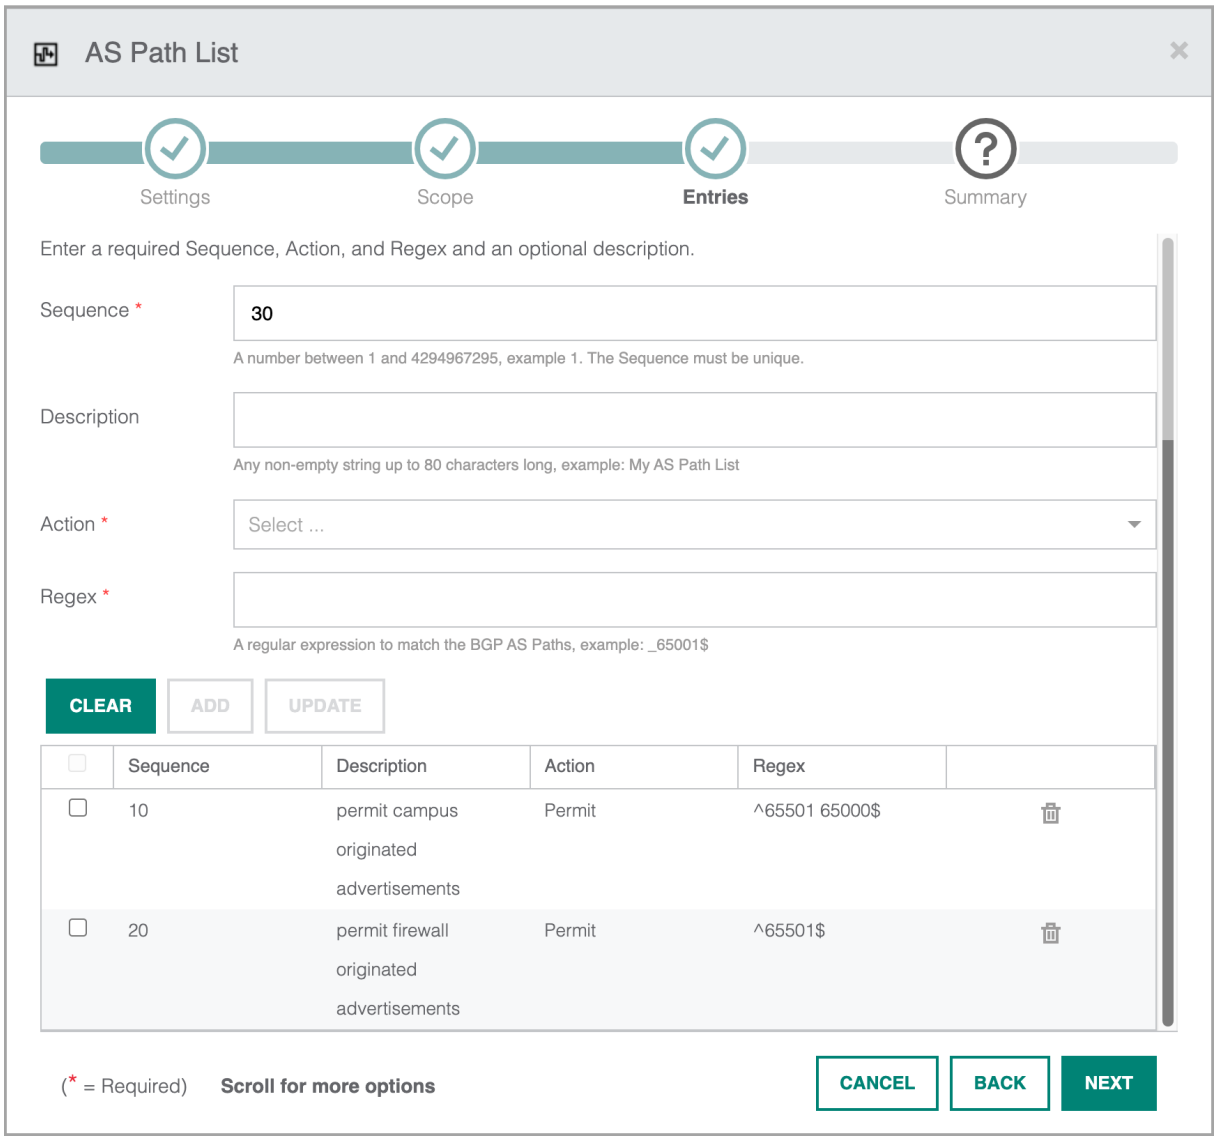


**Step 7** On the **Summary** page, verify the AS path list settings and click **APPLY**.

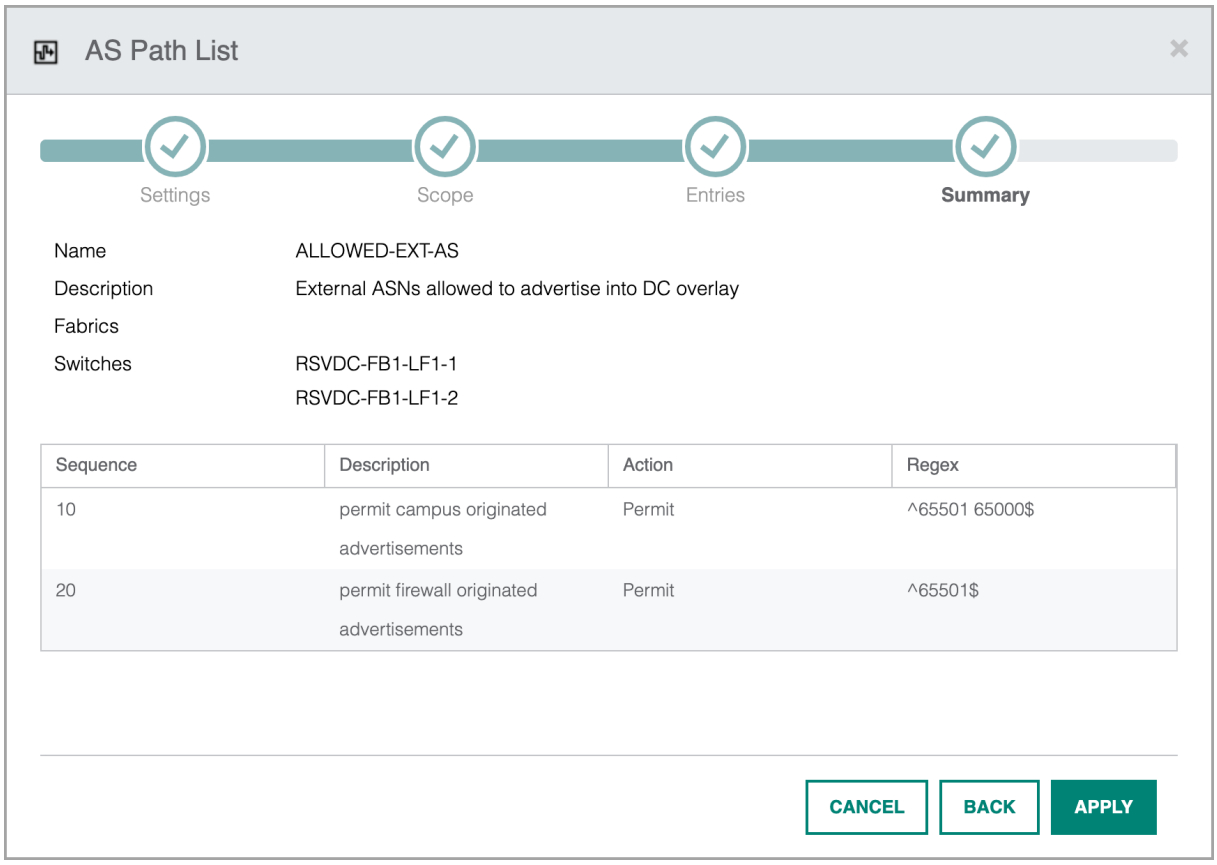


## Configure Firewall Route Map

The following procedure creates a route map that will be applied outbound to external BGP peers. The route map policy filters host and point-to-point prefixes using the previously created host filter prefix list.

**Step 1** On the **Configuration > Routing > Route Policy** page, click the **ROUTE MAPS** tab. On the right **ACTIONS** menu, select **Add**.

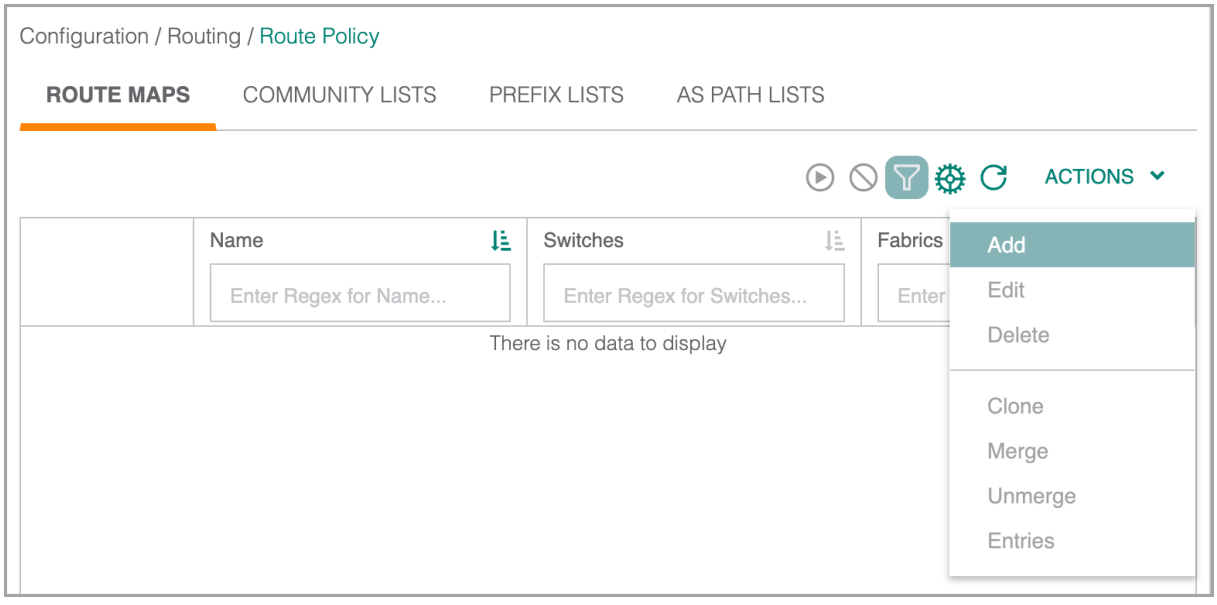

**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

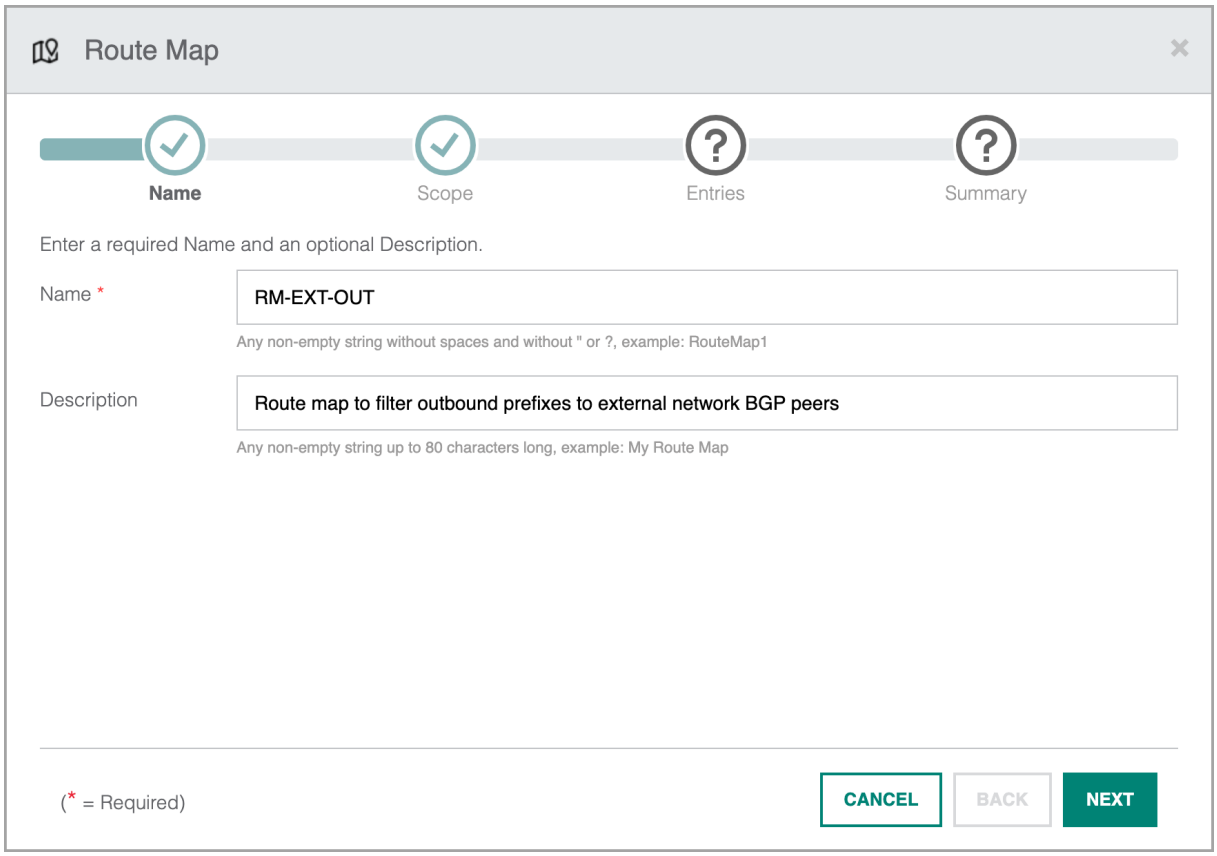


**Step 3** On the **Scope** page, select the two border leaf switches in the **Switches** field, then click **NEXT**.

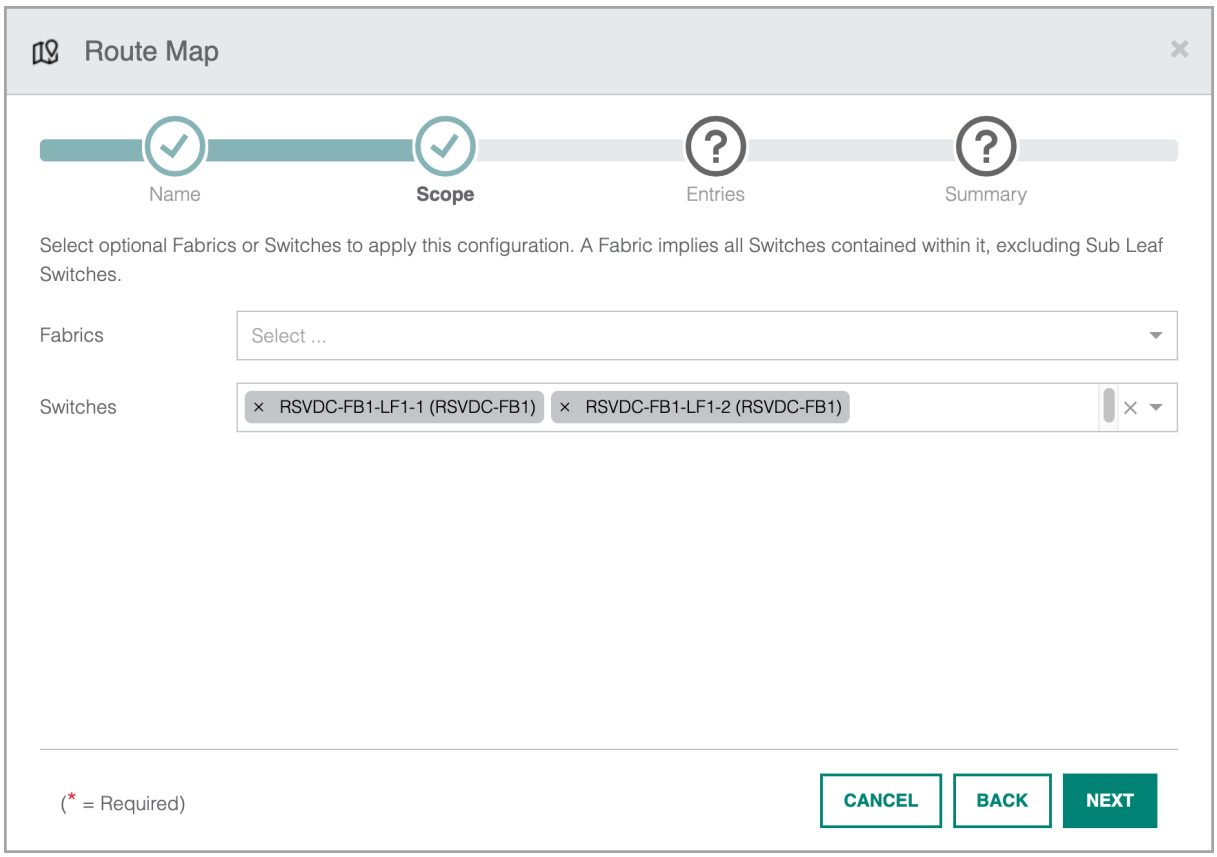

**Step 4** On the **Entries** page, click the right **ACTIONS** menu and select **Add** to launch the **Route Map Entries** wizard.

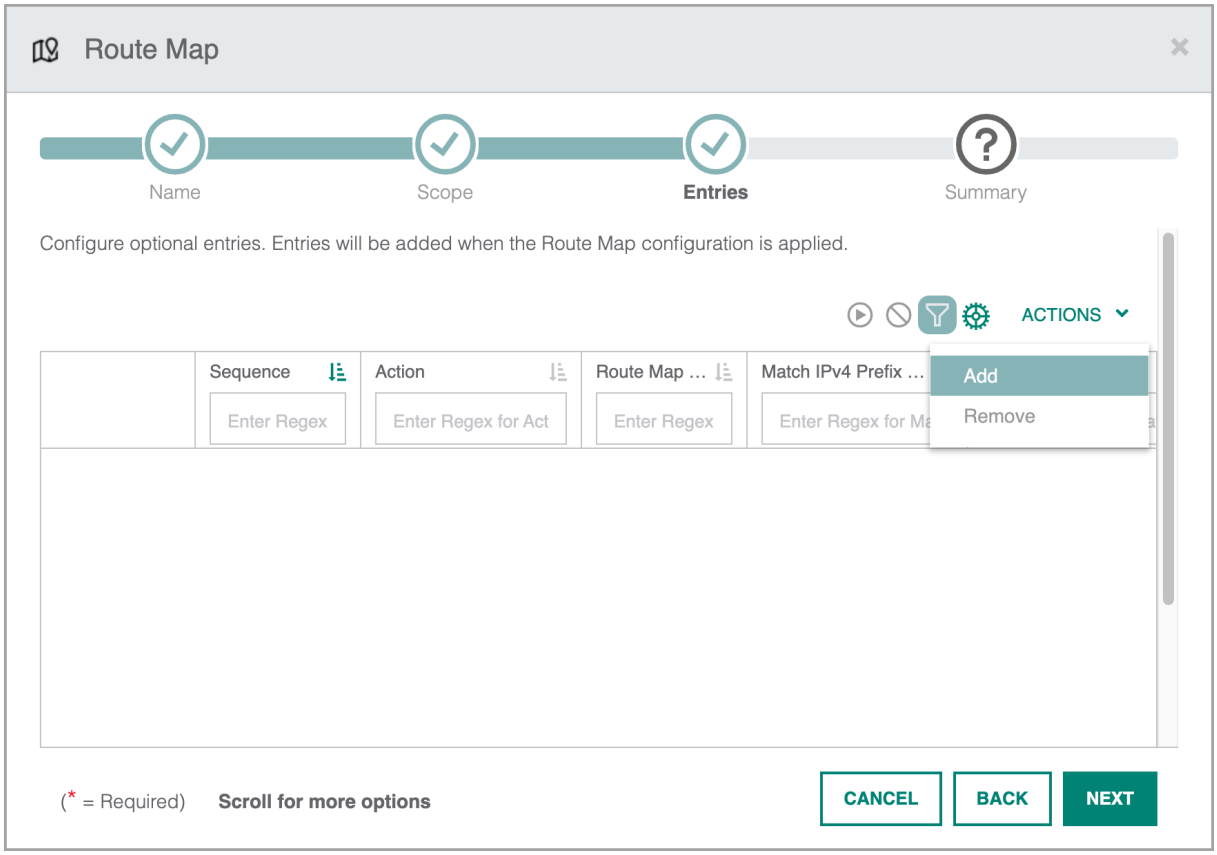


**Step 5 Route Map Entries** wizard: On the **Settings** page, enter the following non-default values and click **NEXT**.

- **Description:** *filter host and P2P prefixes*
- **Action:** *Deny*

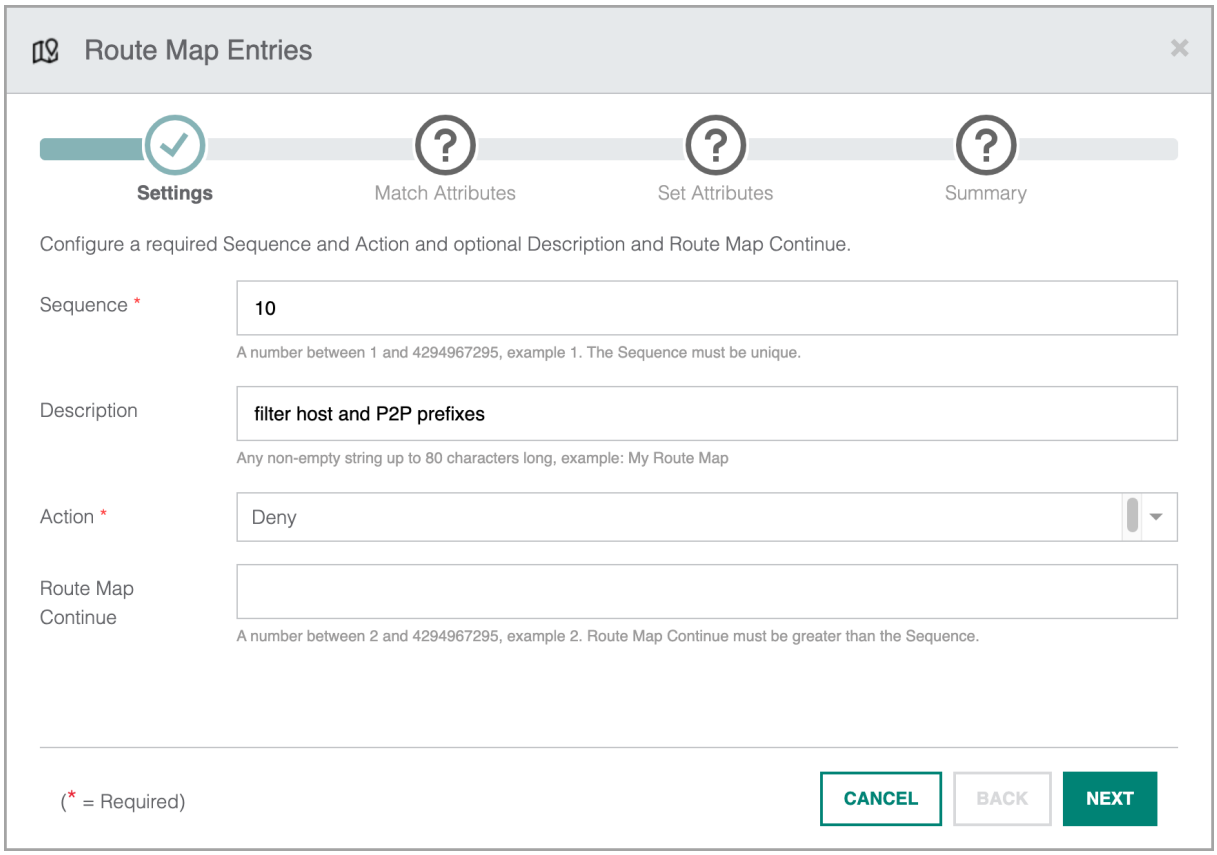


**Step 6 Route Map Entries** wizard: On the **Match Attributes** page, enter the following values and click **NEXT**.

- **Attributes:** *Match IPv4 Prefix List*
- **Match IPv4 Prefix List:** *PL-HOST-P2P*

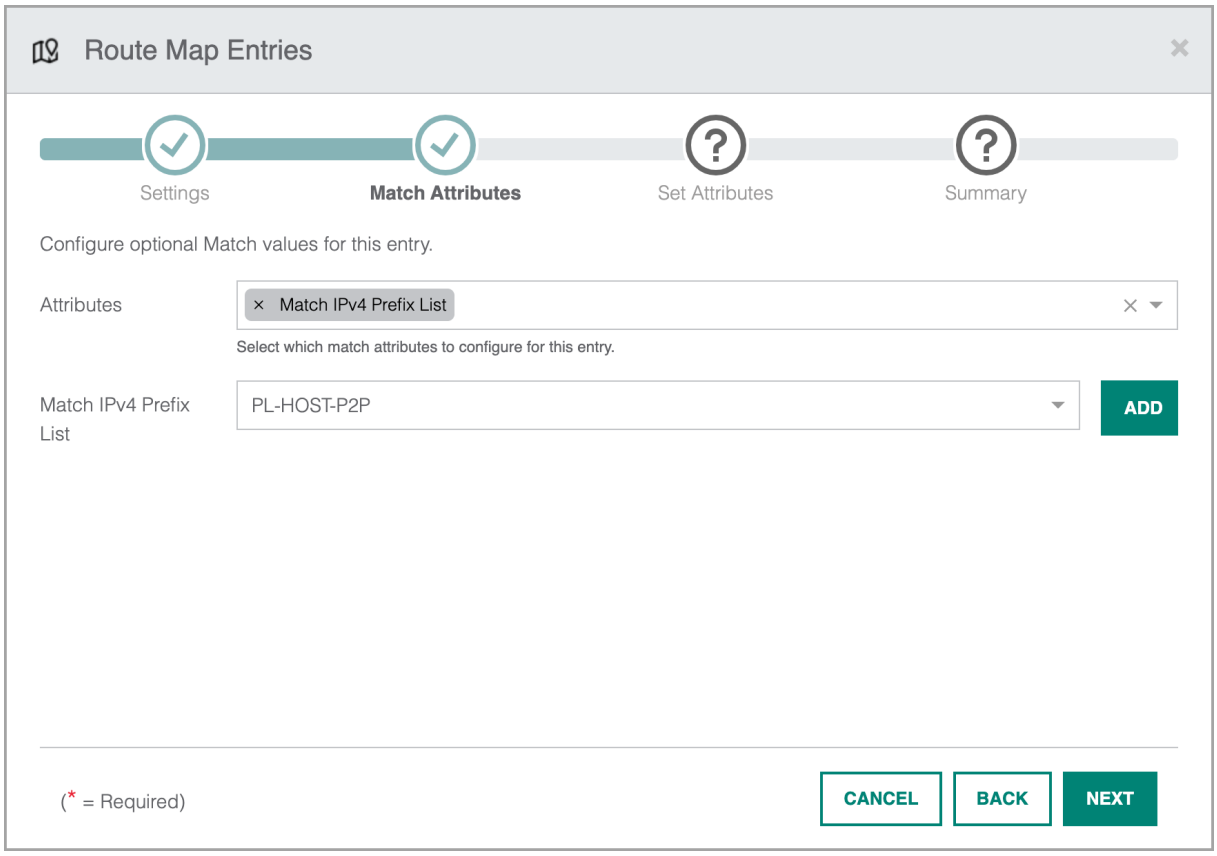


**Step 7 Route Map Entries** wizard: On the **Set Attributes** page, click **NEXT**.

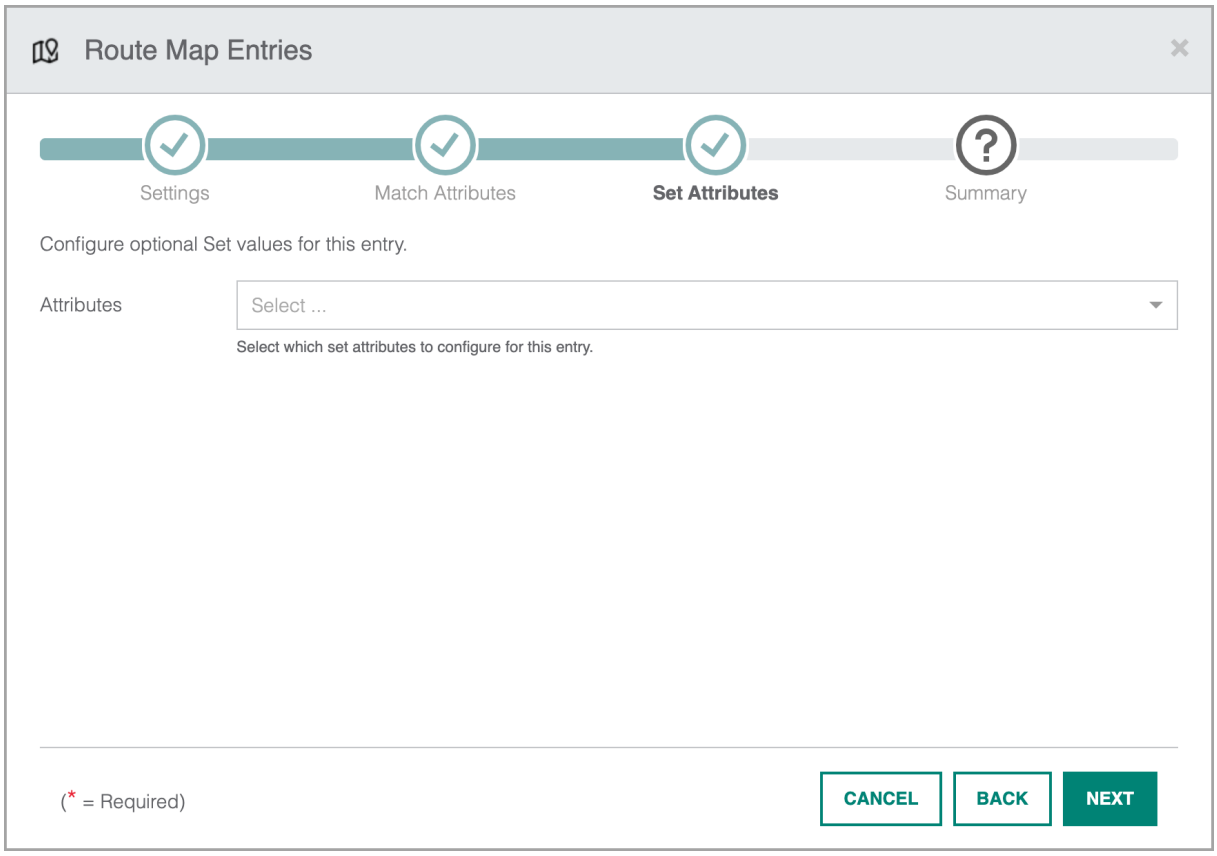

**Step 8 Route Map Entries** wizard: On the **Summary** page, review the route map entry settings and click **APPLY**.

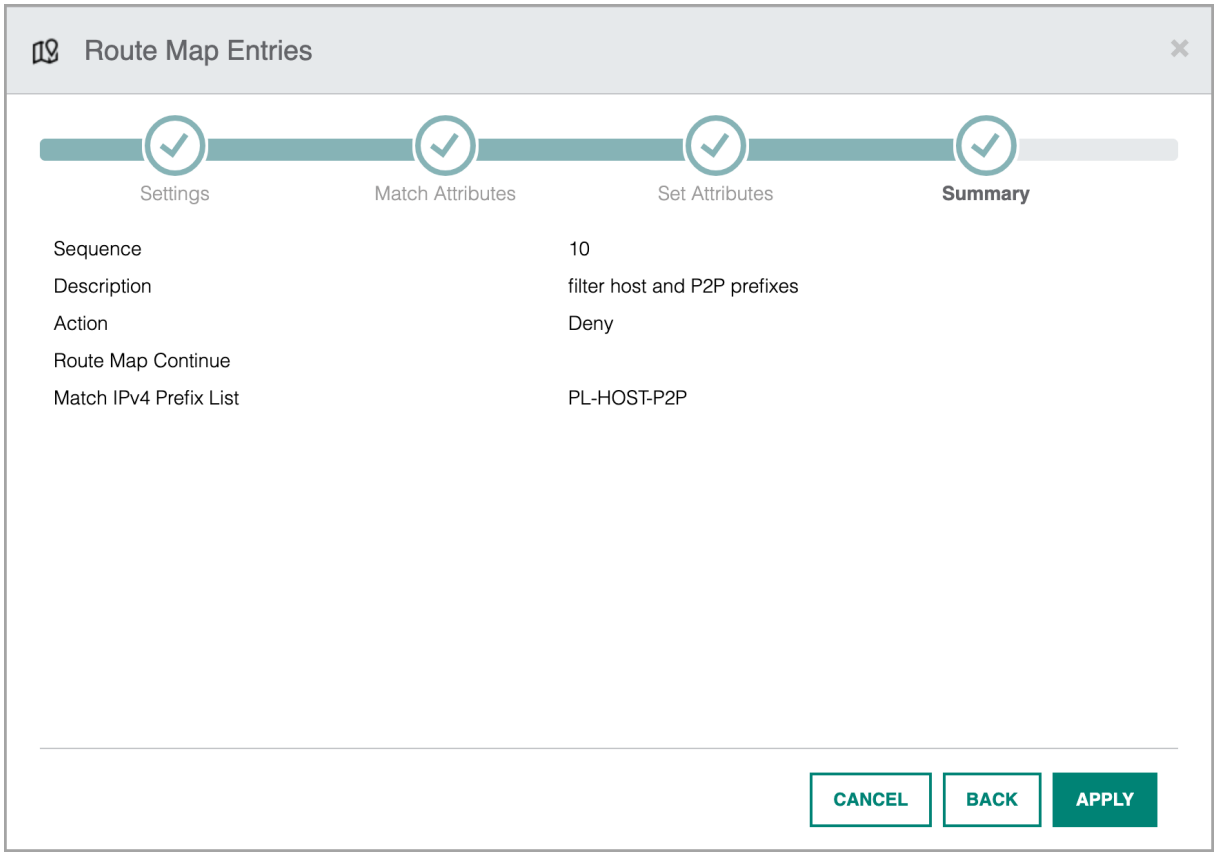


**Step 9** Create a second route map sequence. On the **Entries** page, click the right **ACTIONS** menu, and select **Add**.

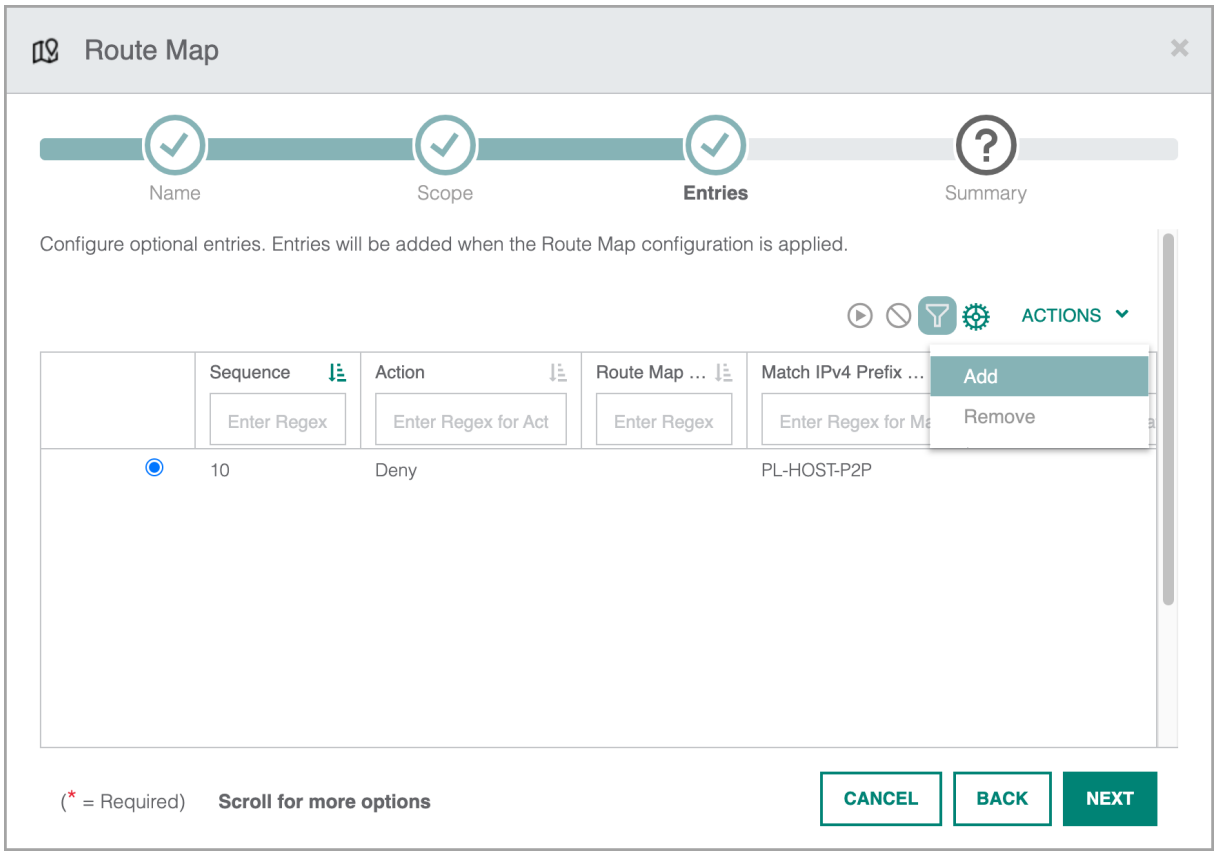


**Step 10 Route Map Entries** wizard: On the **Settings** page, set the **Action** field to *Permit* and click **NEXT**.

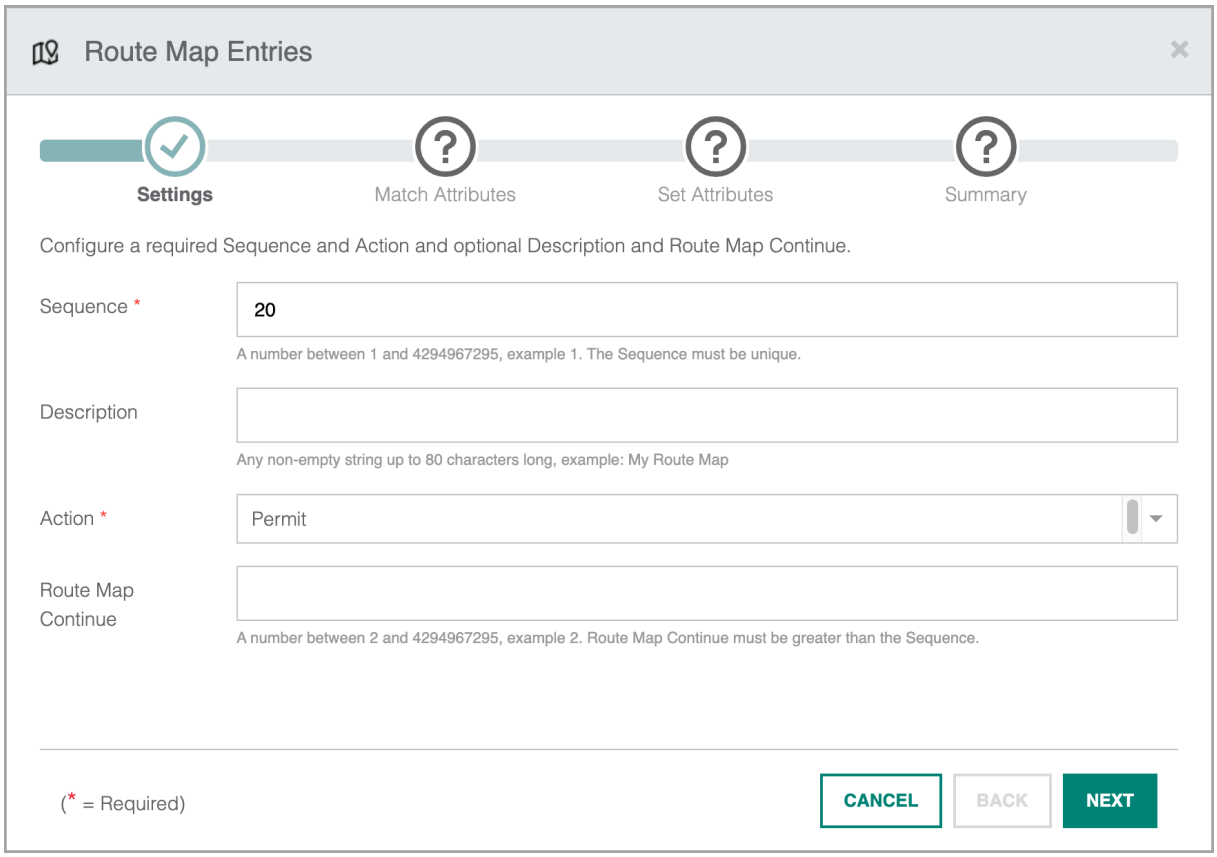


**Step 11 Route Map Entries** wizard: On the **Match Attributes** page, click **NEXT**.

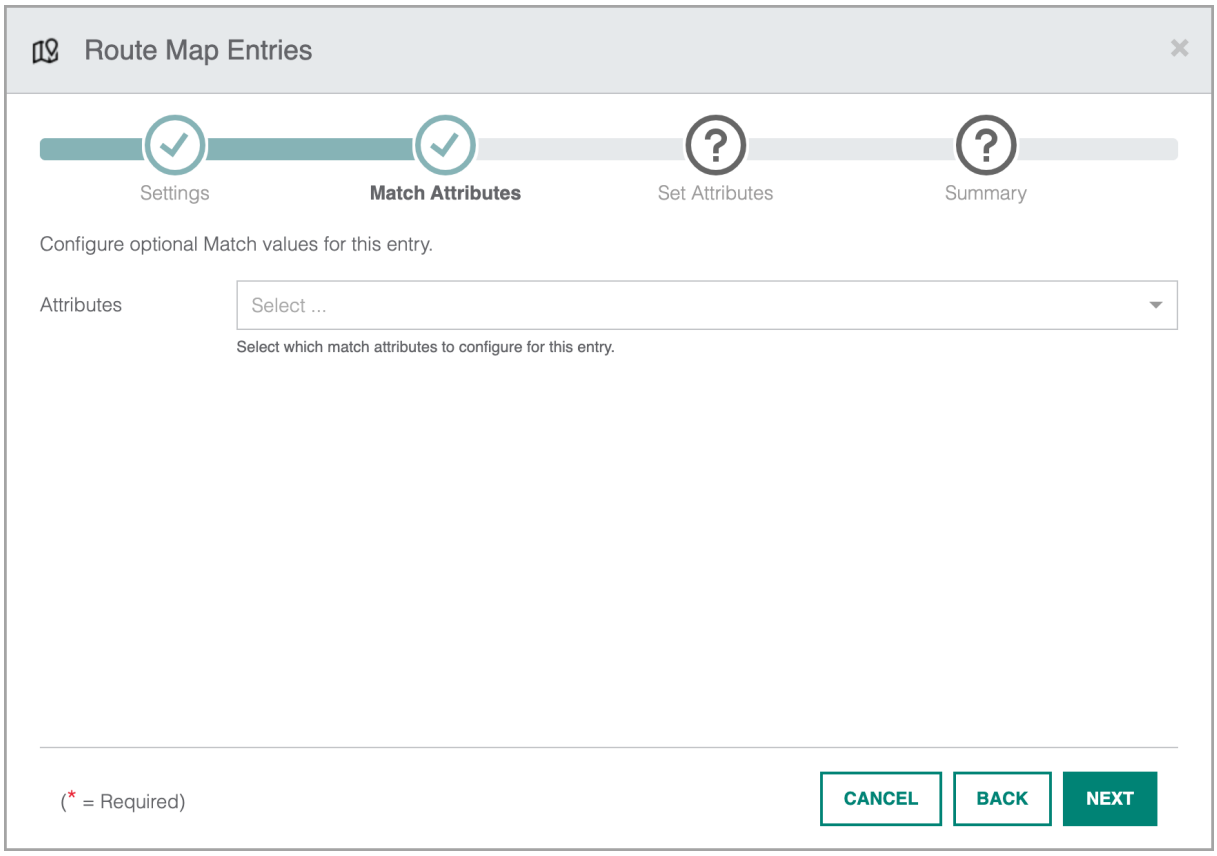


**Step 12 Route Map Entries** wizard: On the **Set Attributes** page, click **NEXT**.

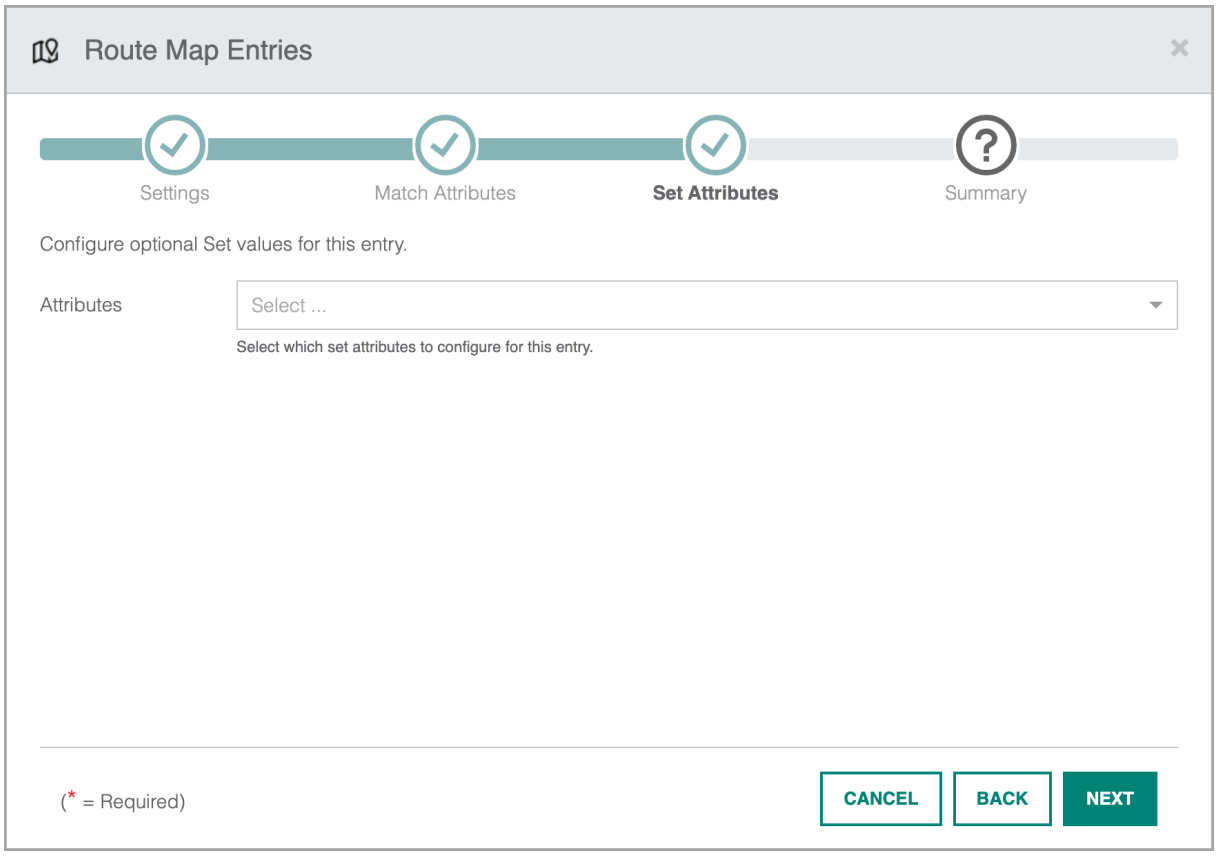

**Step 13 Route Map Entries** wizard: On the **Summary** page, review the route map entry settings and click **APPLY**.

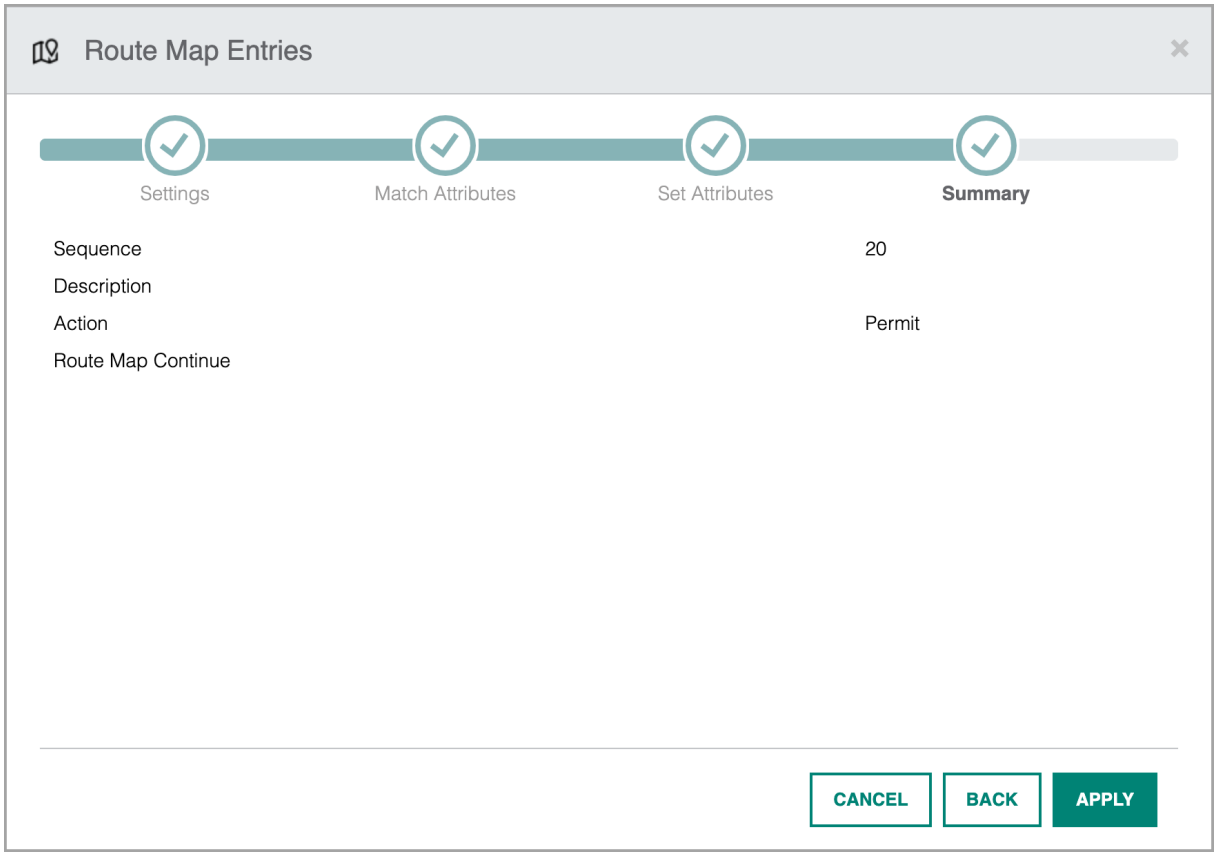


**Step 14** On the **Entries** page, click **NEXT**.

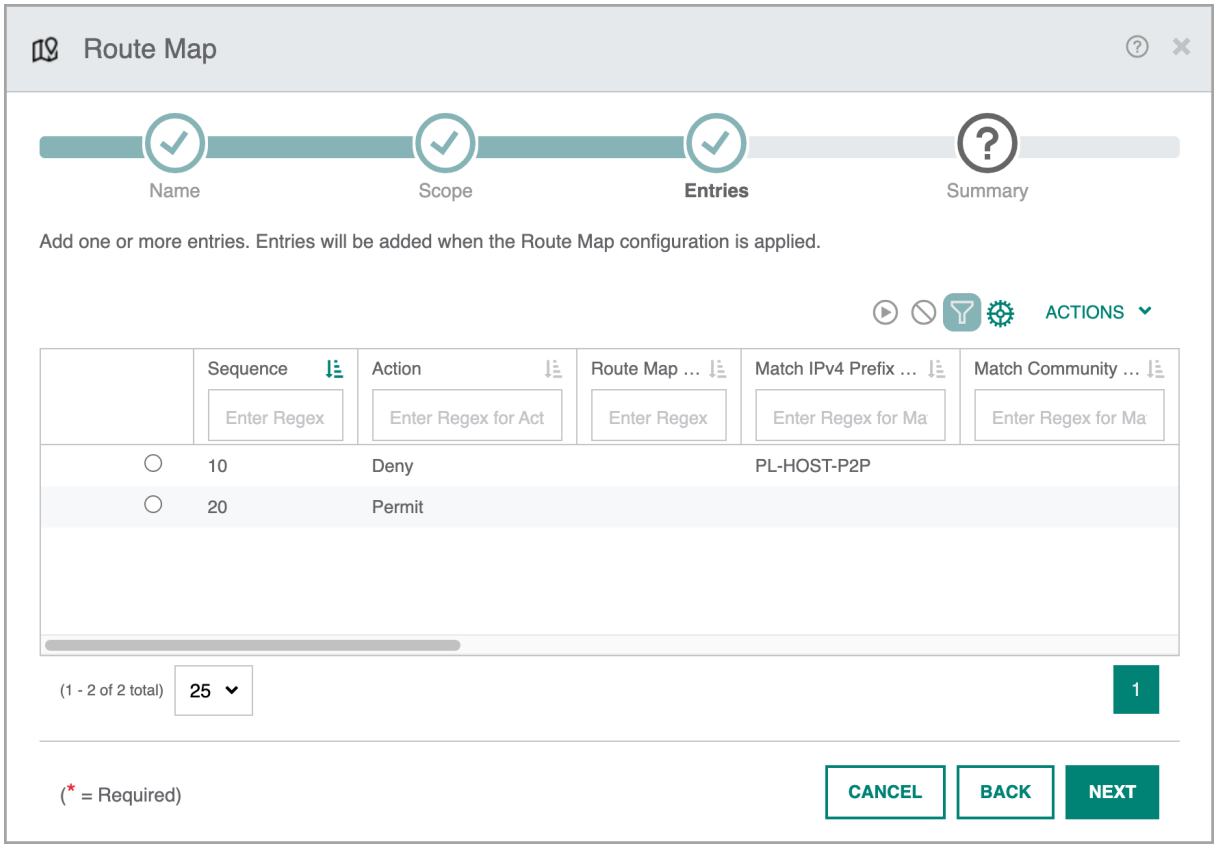

**Step 15** On the **Summary** page, review the route map settings and click **APPLY**.

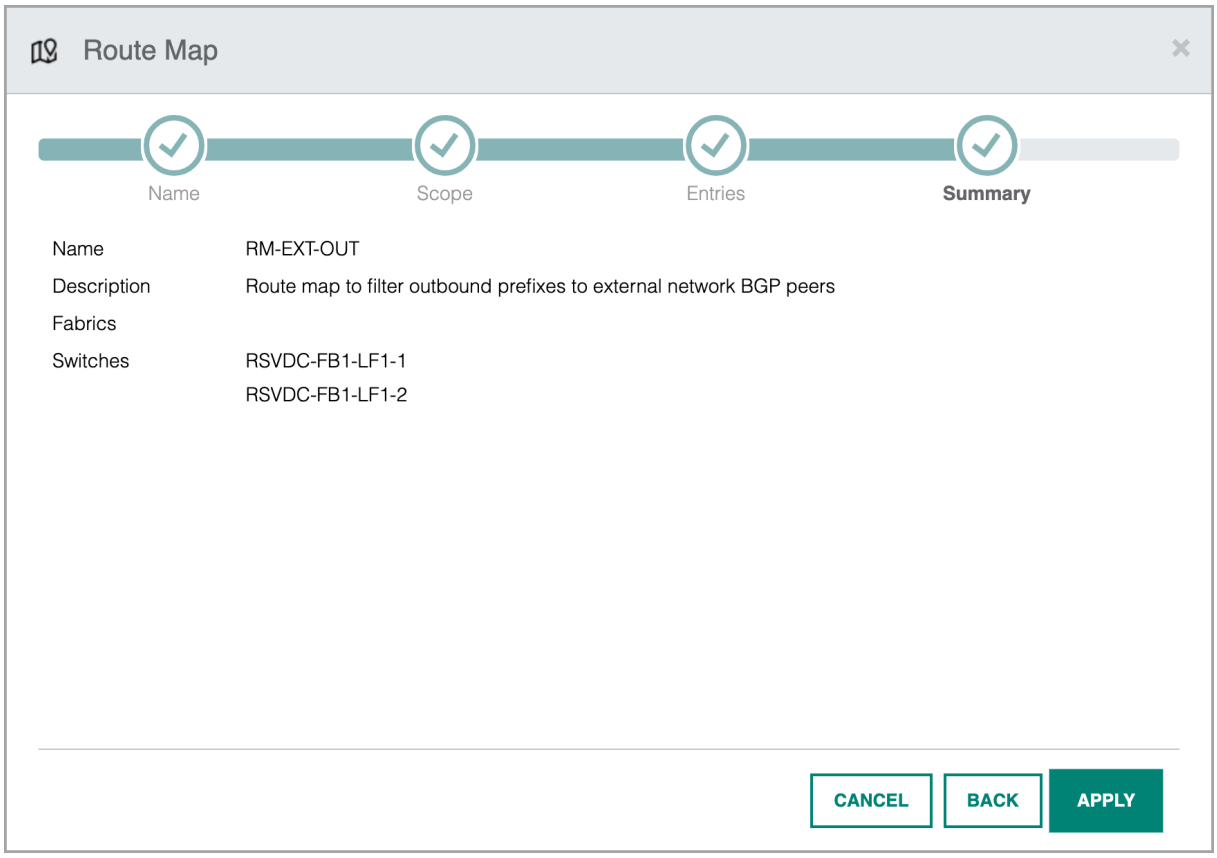


## Configure Internal Border Leaf Route Map

The following procedure creates a route map that is applied to the BGP peering between the border leaf switches. The route map permits advertising only prefixes originated by the campus AS number or the upstream firewall AS number.

**Step 1** On the right **ACTIONS** menu of the **ROUTE MAPS** tab, select **Add**.

**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

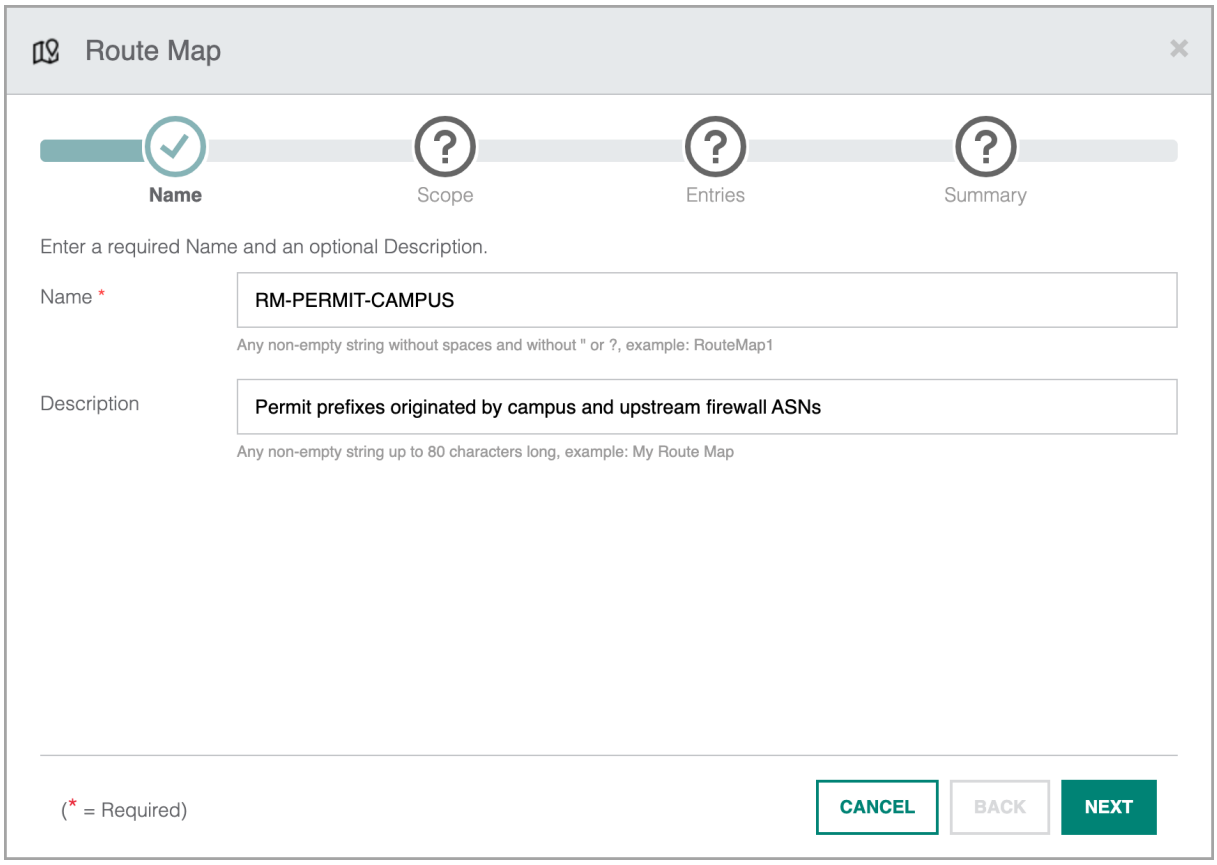


**Step 3** On the **Scope** page, select the two border leaf switches in the **Switches** field, then click **NEXT**.

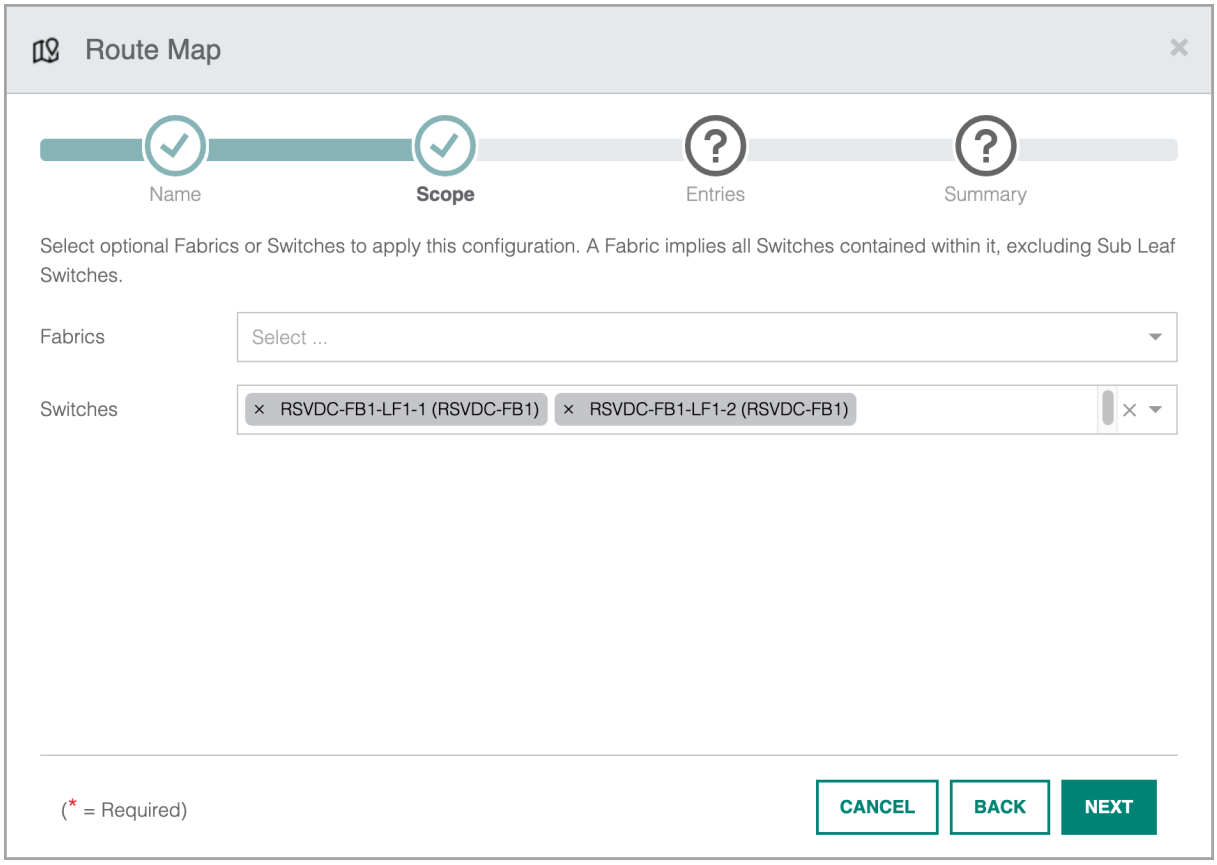

**Step 4** On the **Entries** page, click the right **ACTIONS** menu, and select **Add** to launch the **Route Map Entries** wizard.

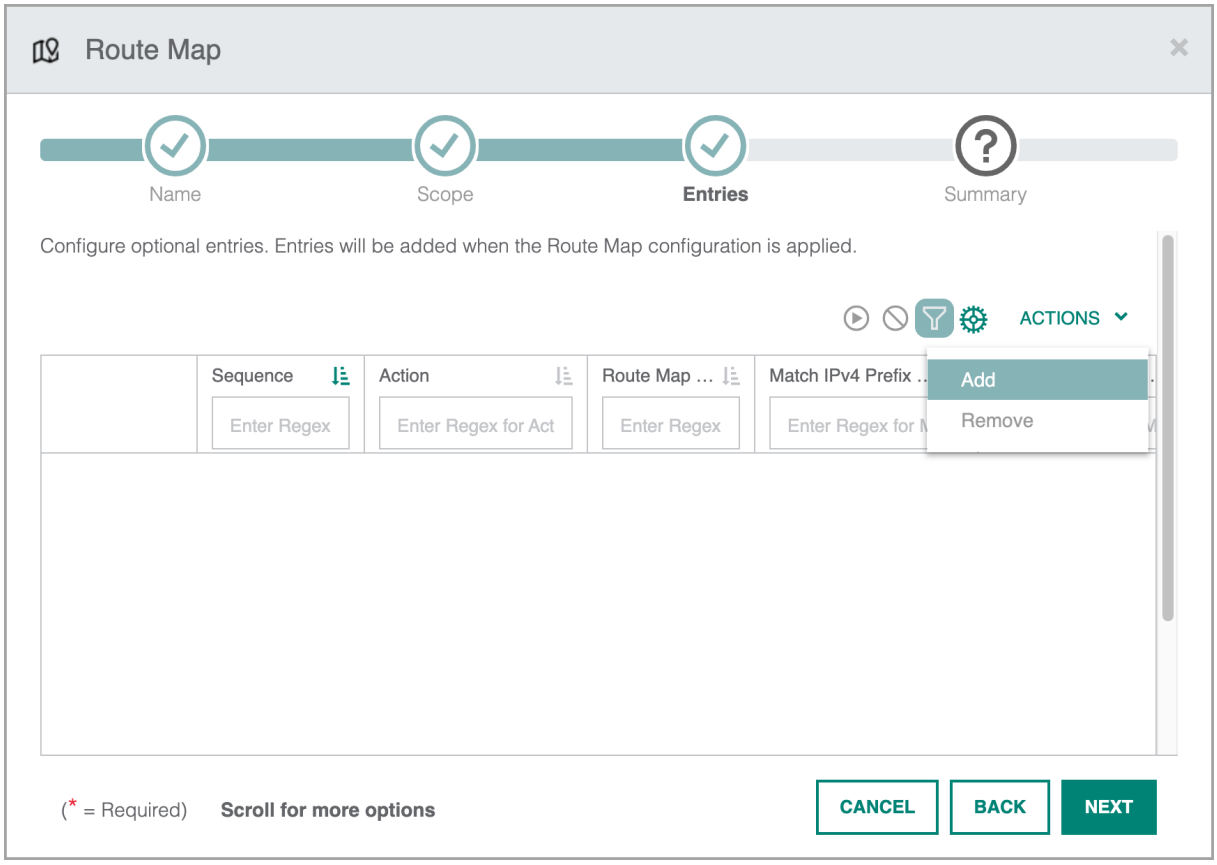


**Step 5 Route Map Entries** wizard: On the **Settings** page, enter the following non-default values and click **NEXT**.

- **Description:** *allow campus and firewall ASNs*
- **Action:** *Permit*

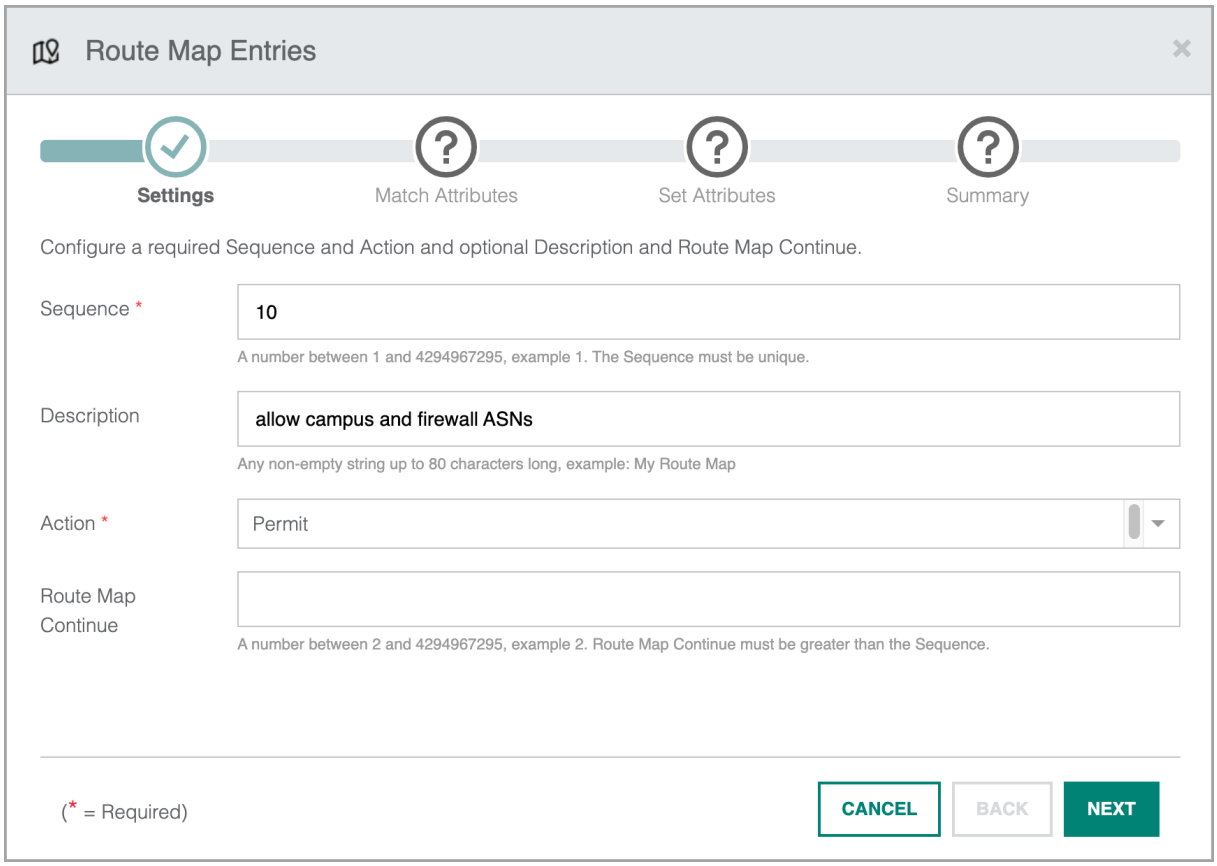


**Step 6 Route Map Entries** wizard: On the **Match Attributes** page, enter the following values and click **NEXT**.

- **Attributes:** *Match AS Path List*
- **Match IPv4 Prefix List:** *ALLOWED-EXT-AS*

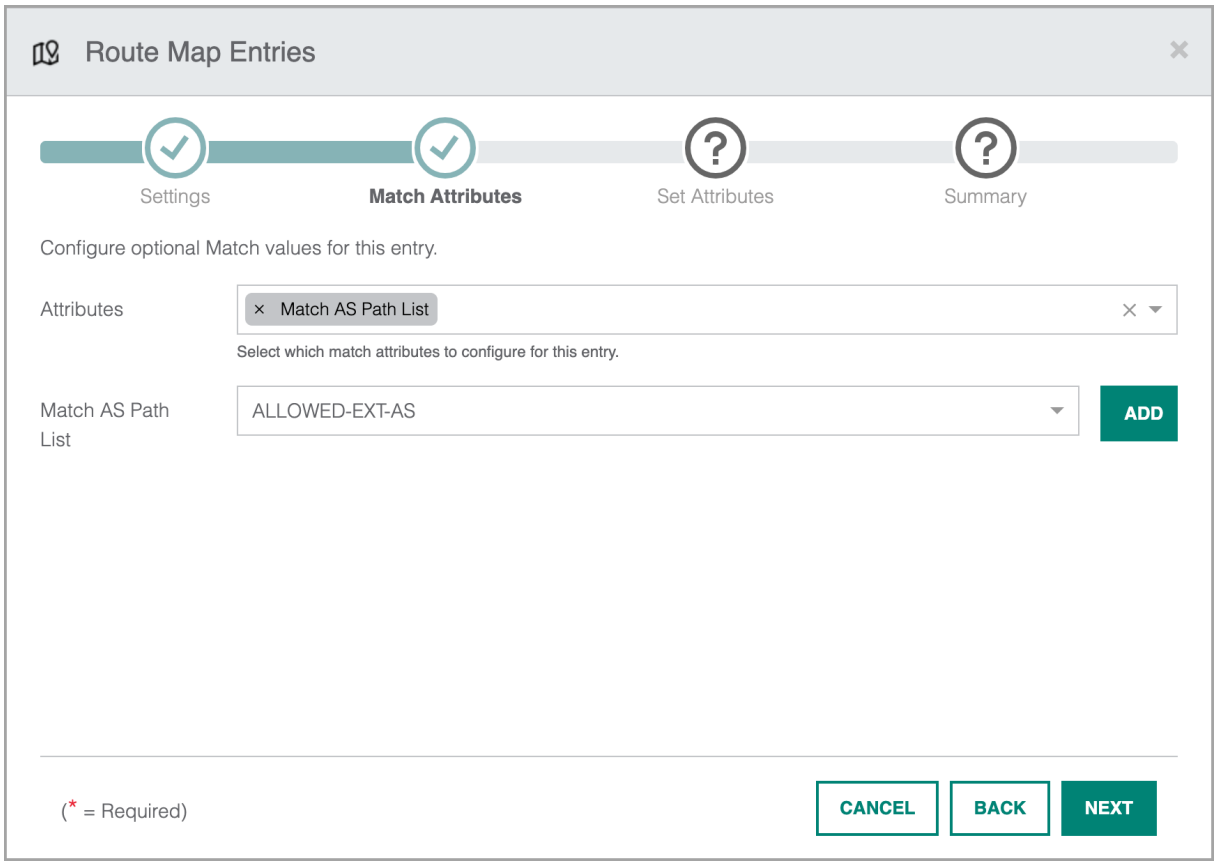


**Step 7 Route Map Entries** wizard: On the **Set Attributes** page, click **NEXT**.

**Step 8 Route Map Entries** wizard: On the **Summary** page, review the route map entry settings and click **APPLY**.

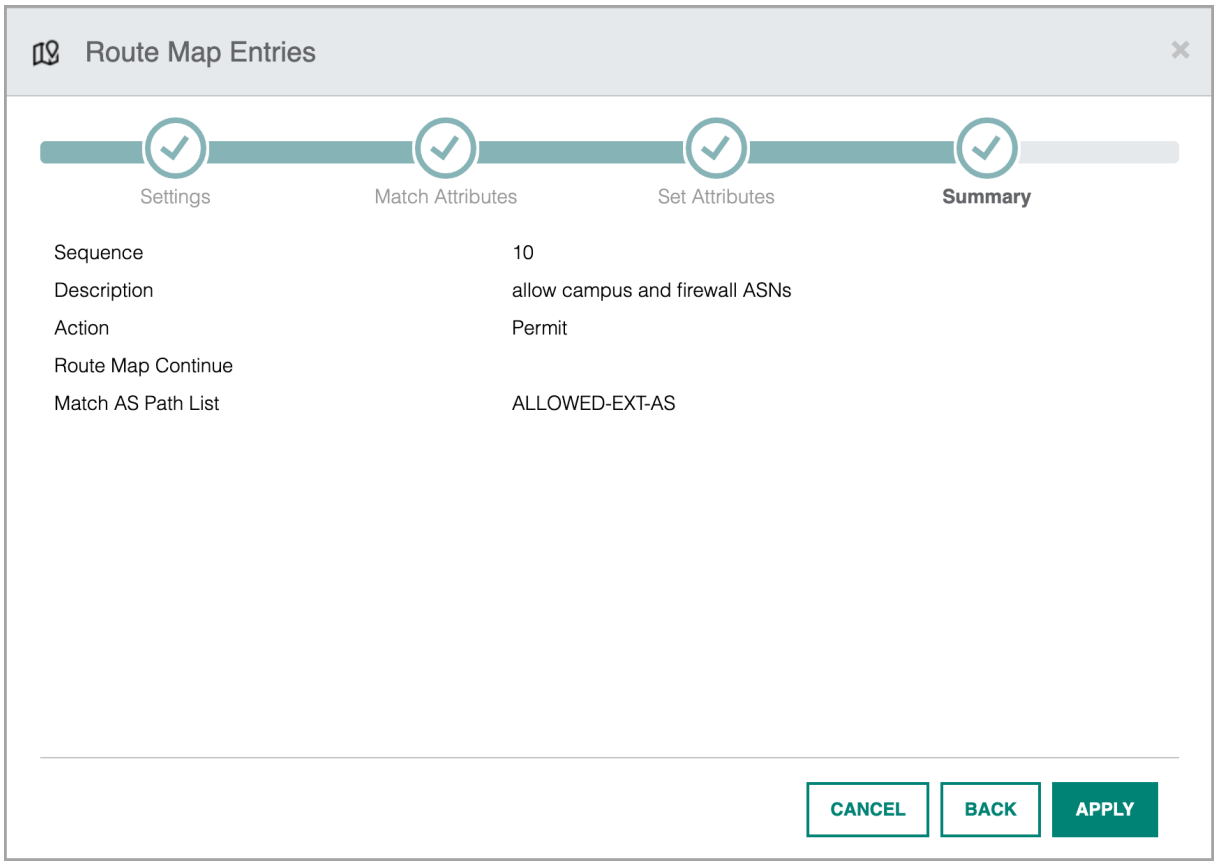


**Step 9** Create a second route map sequence. On the **Entries** page, click the right **ACTIONS** menu, and select **Add**.

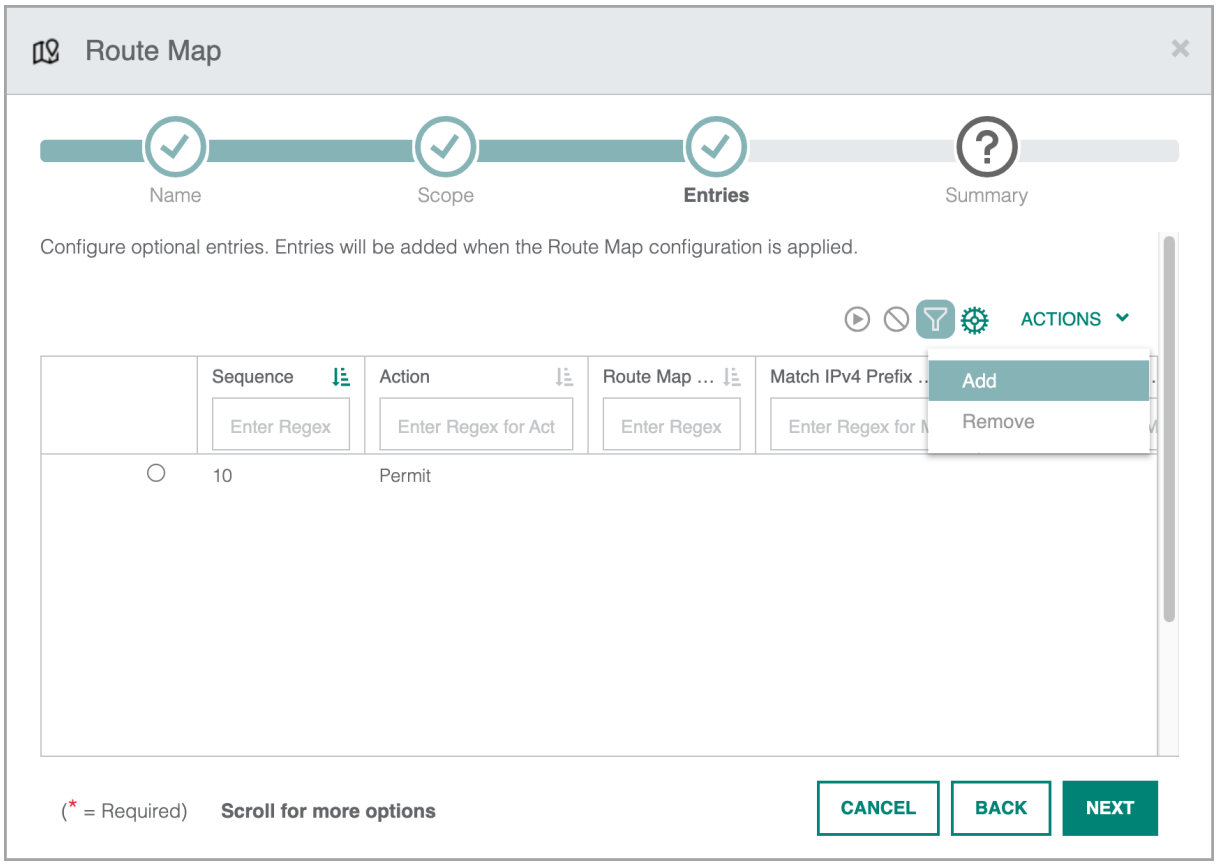

**Step 10 Route Map Entries** wizard: On the **Settings** page, set the **Action** field to **Deny** and click **NEXT**.

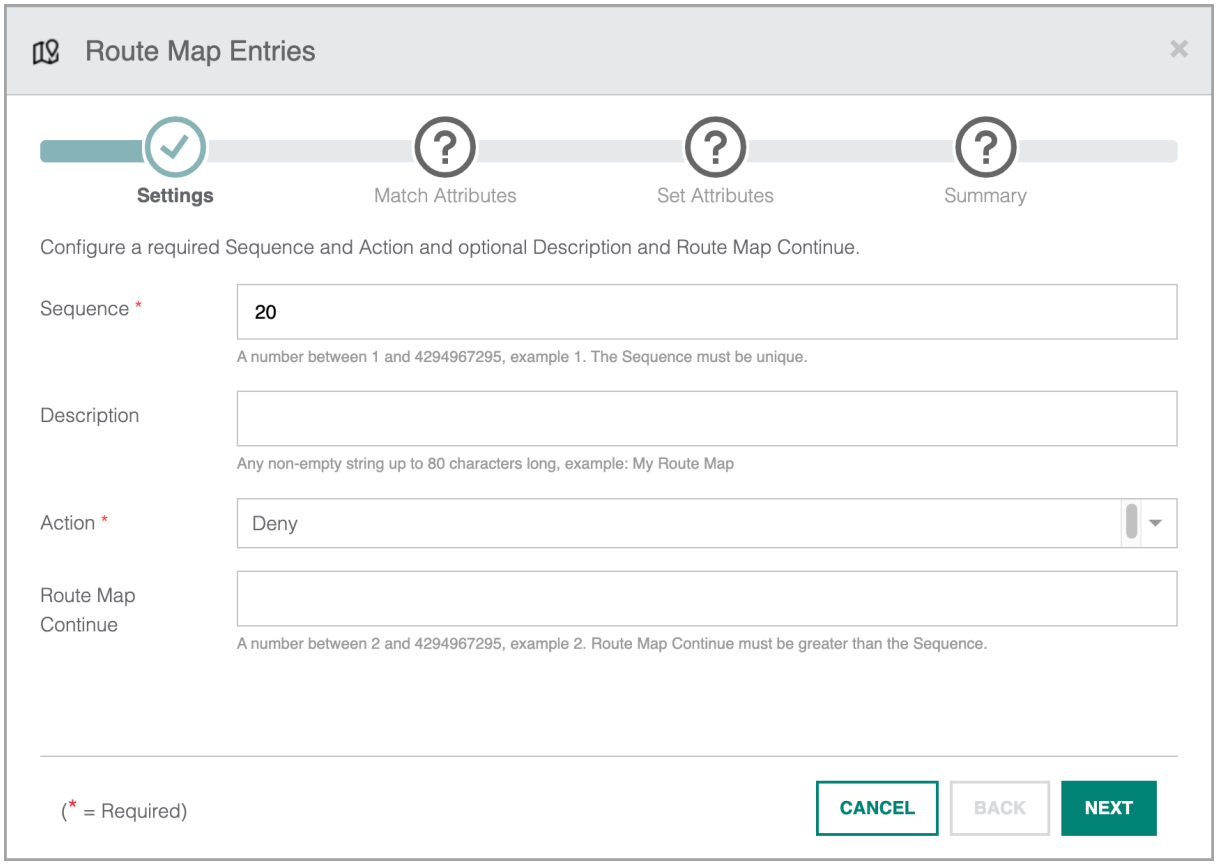


**Step 11 Route Map Entries** wizard: On the **Match Attributes** page, click **NEXT**.

**Step 12 Route Map Entries** wizard: On the **Set Attributes** page, click **NEXT**.

**Step 13 Route Map Entries** wizard: On the **Summary** page, review the route map entry settings and click **APPLY**.

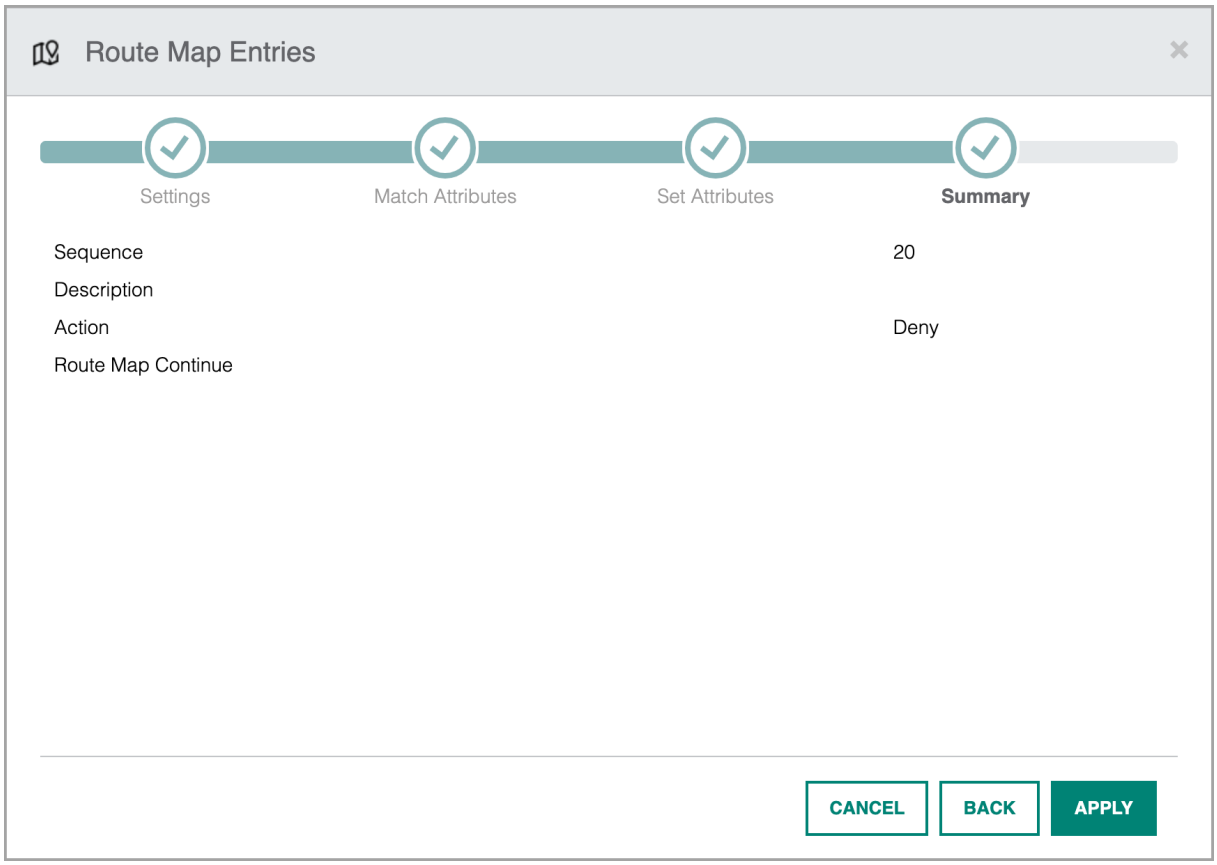


**Step 14** On the **Entries** page, click **NEXT**.

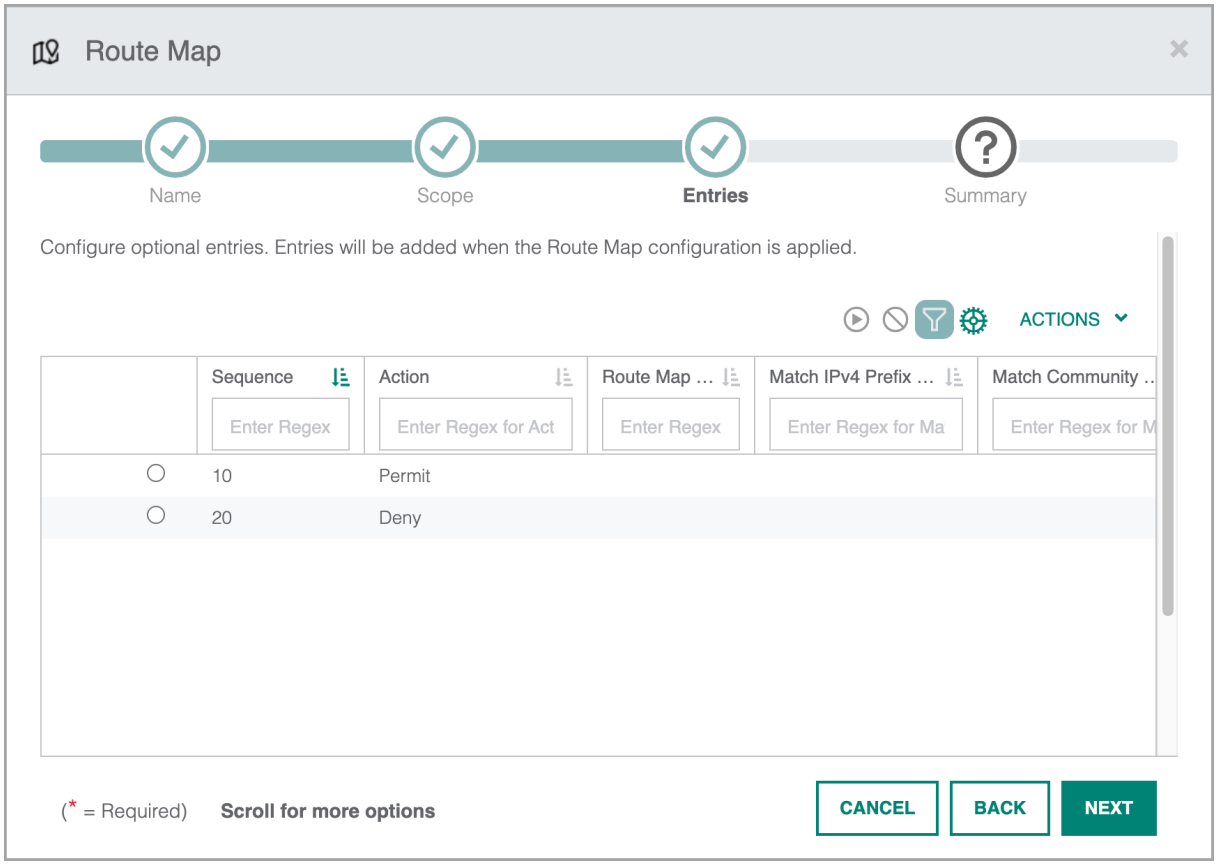


**Step 15** On the **Summary** page, review the route map settings and click **APPLY**.

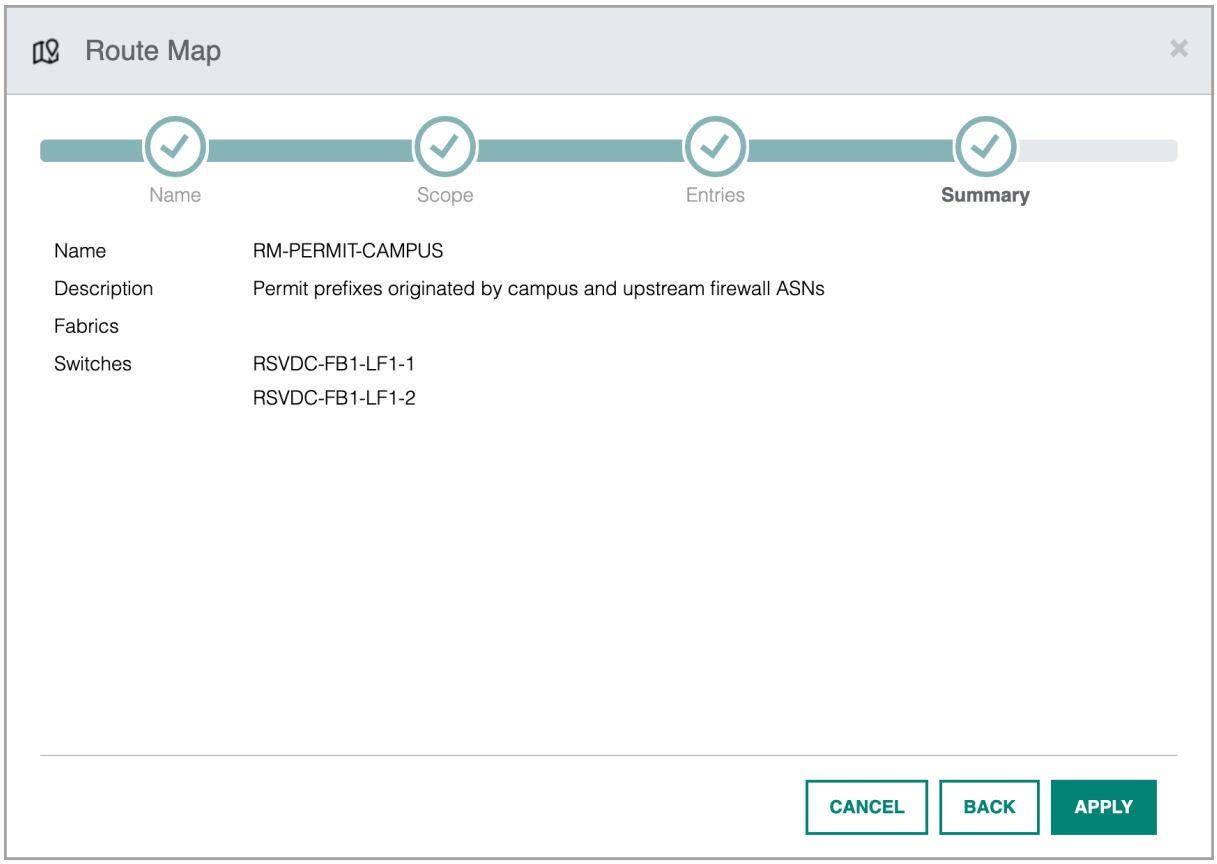


## Configure Border Leaf BGP Peerings

The following procedure configures the eBGP peerings between the border leaf switches and the upstream firewalls with a route map applied to filter host routes and point-to-point link prefixes. A single BGP peering is defined to the upstream firewalls, which is established only with the active firewall in the active/passive pair.

**Step 1** On the left navigation menu, click **BGP**. Click the **PROD-DC-VRF** radio button. On the right **ACTIONS** menu, select **Edit**.

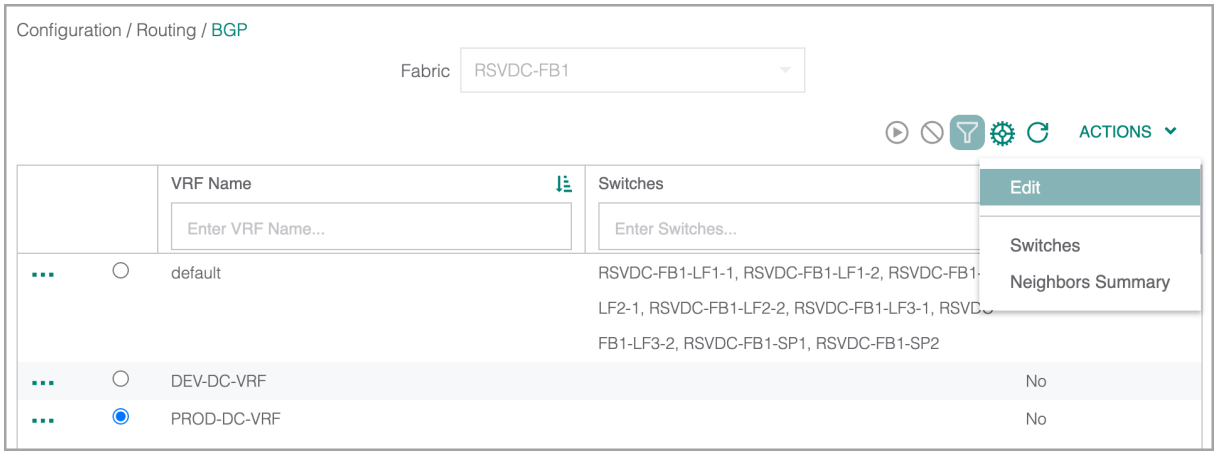


**Step 2** On the **Settings** page, check **Enable BGP on PROD-DC-VRF**, check **Redistribute Loopback**, and click **APPLY**.

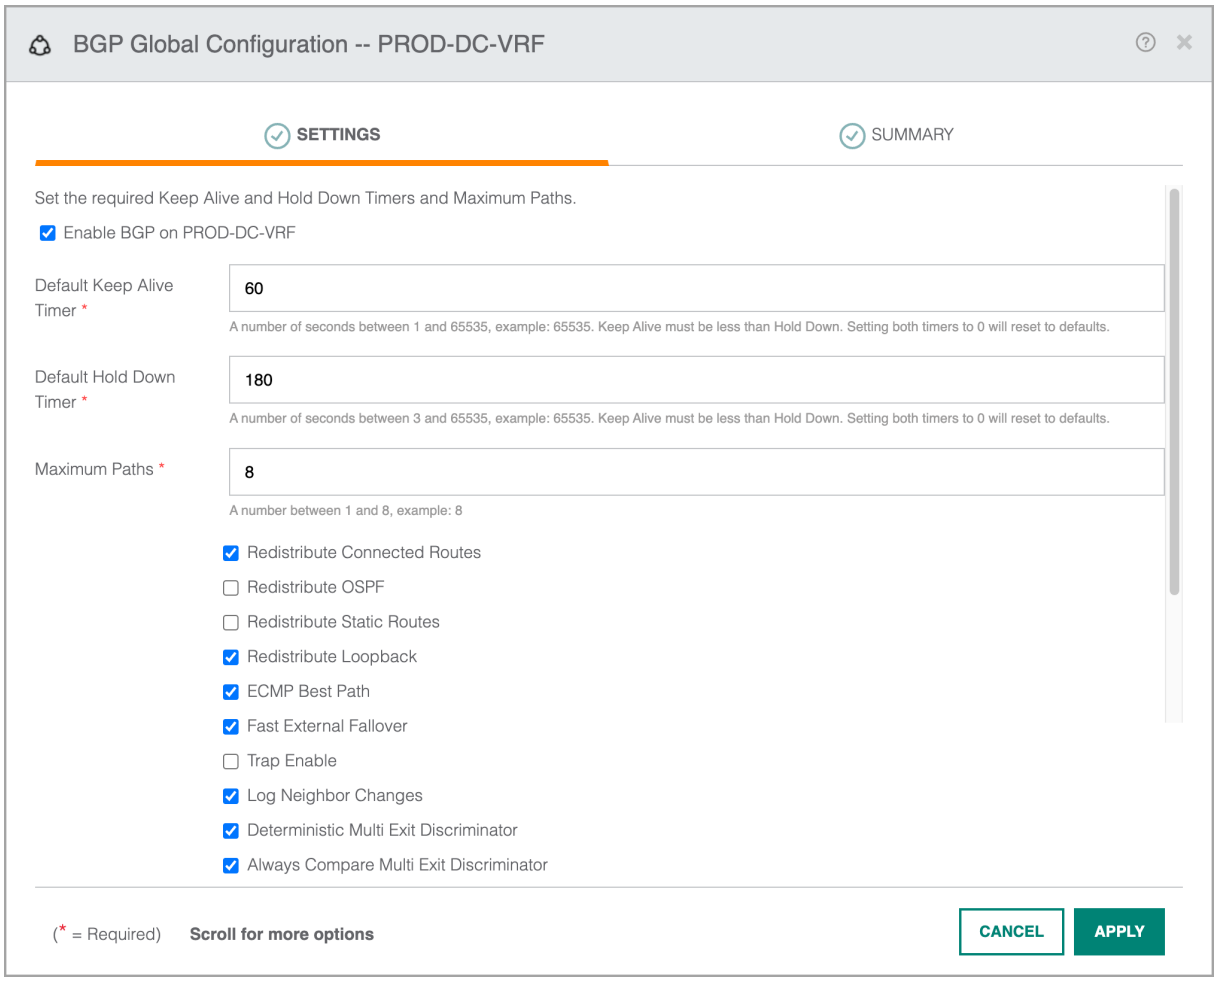


**NOTE:**

The loopback interfaces that are redistributed are created in the **Assign Unique Overlay Loopbacks** procedure later in this guide.

**Step 3** On the **Configuration > Routing > BGP** page, click the • • • symbol left of **PROD-DC-VRF** and select **Switches**.

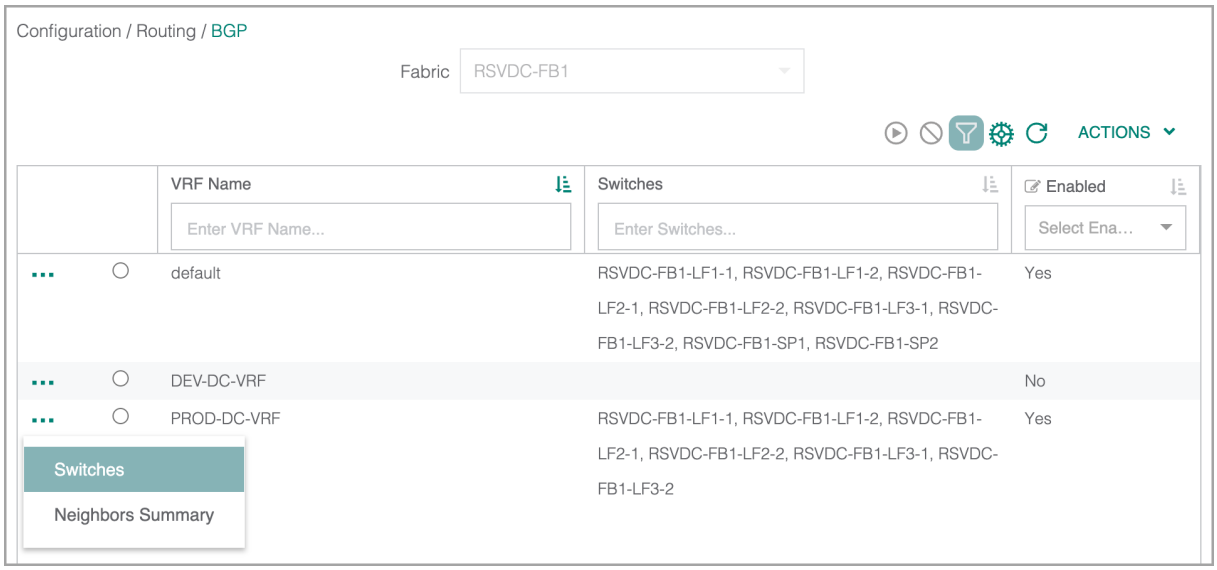

**Step 4** On the **SWITCHES** tab, click • • • next to **RSVDC-FB1-LF1-1** and select **Neighbors**.

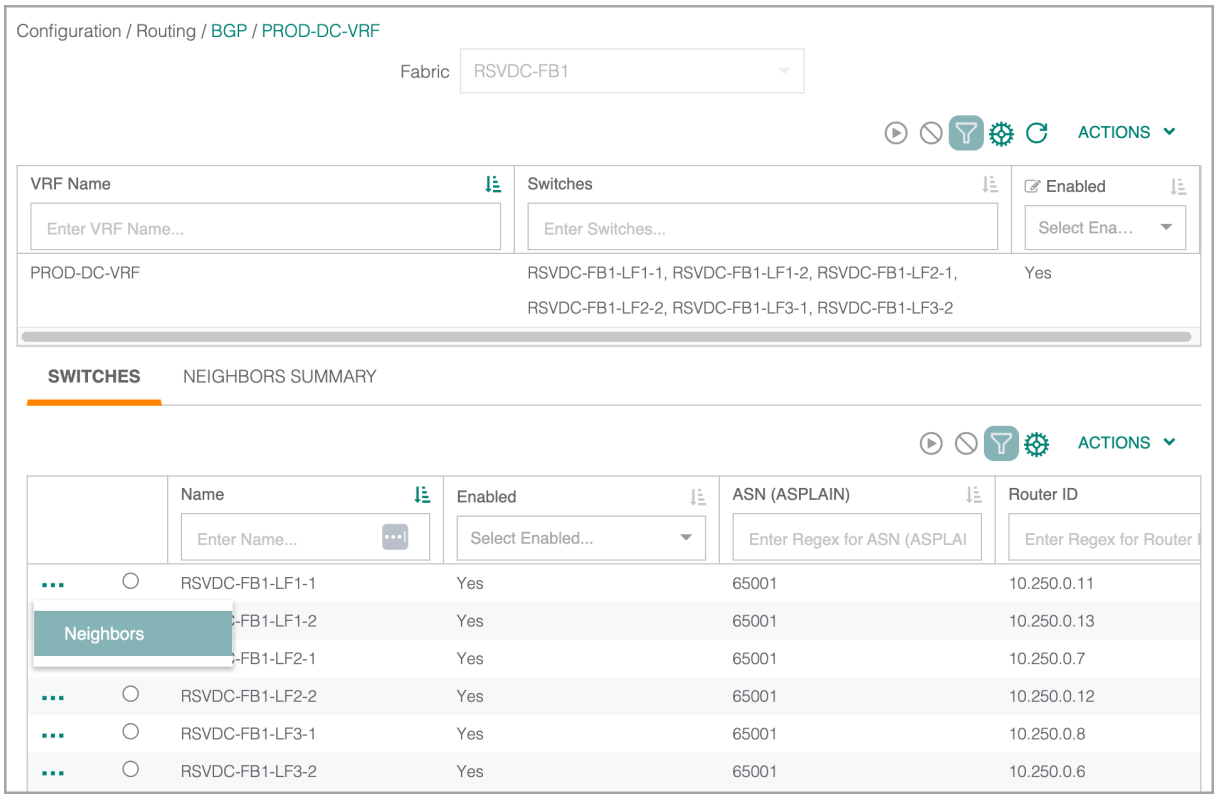


**Step 5** On the right **ACTIONS** menu of the **NEIGHBORS** tab, select **Add**.

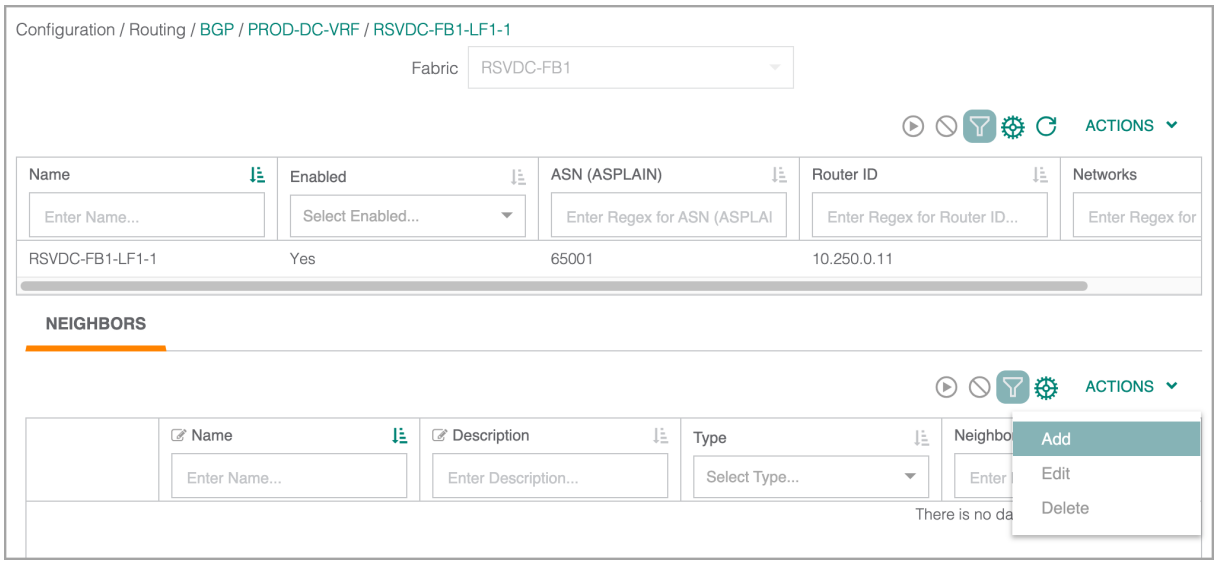


**Step 6** On the **Settings** page, enter the following non-default values and click **NEXT**.

- **Neighbor AS Number:** *65501*
- **IP Address:** *10.255.2.3*
- **IPv4 Route Map In:** *RM-PERMIT-CAMPUS*
- **IPv4 Route Map Out:** *RM-EXT-OUT*
- **Enable Bidirectional Forwarding Detection (BFD) Fall Over**: *< checked >*

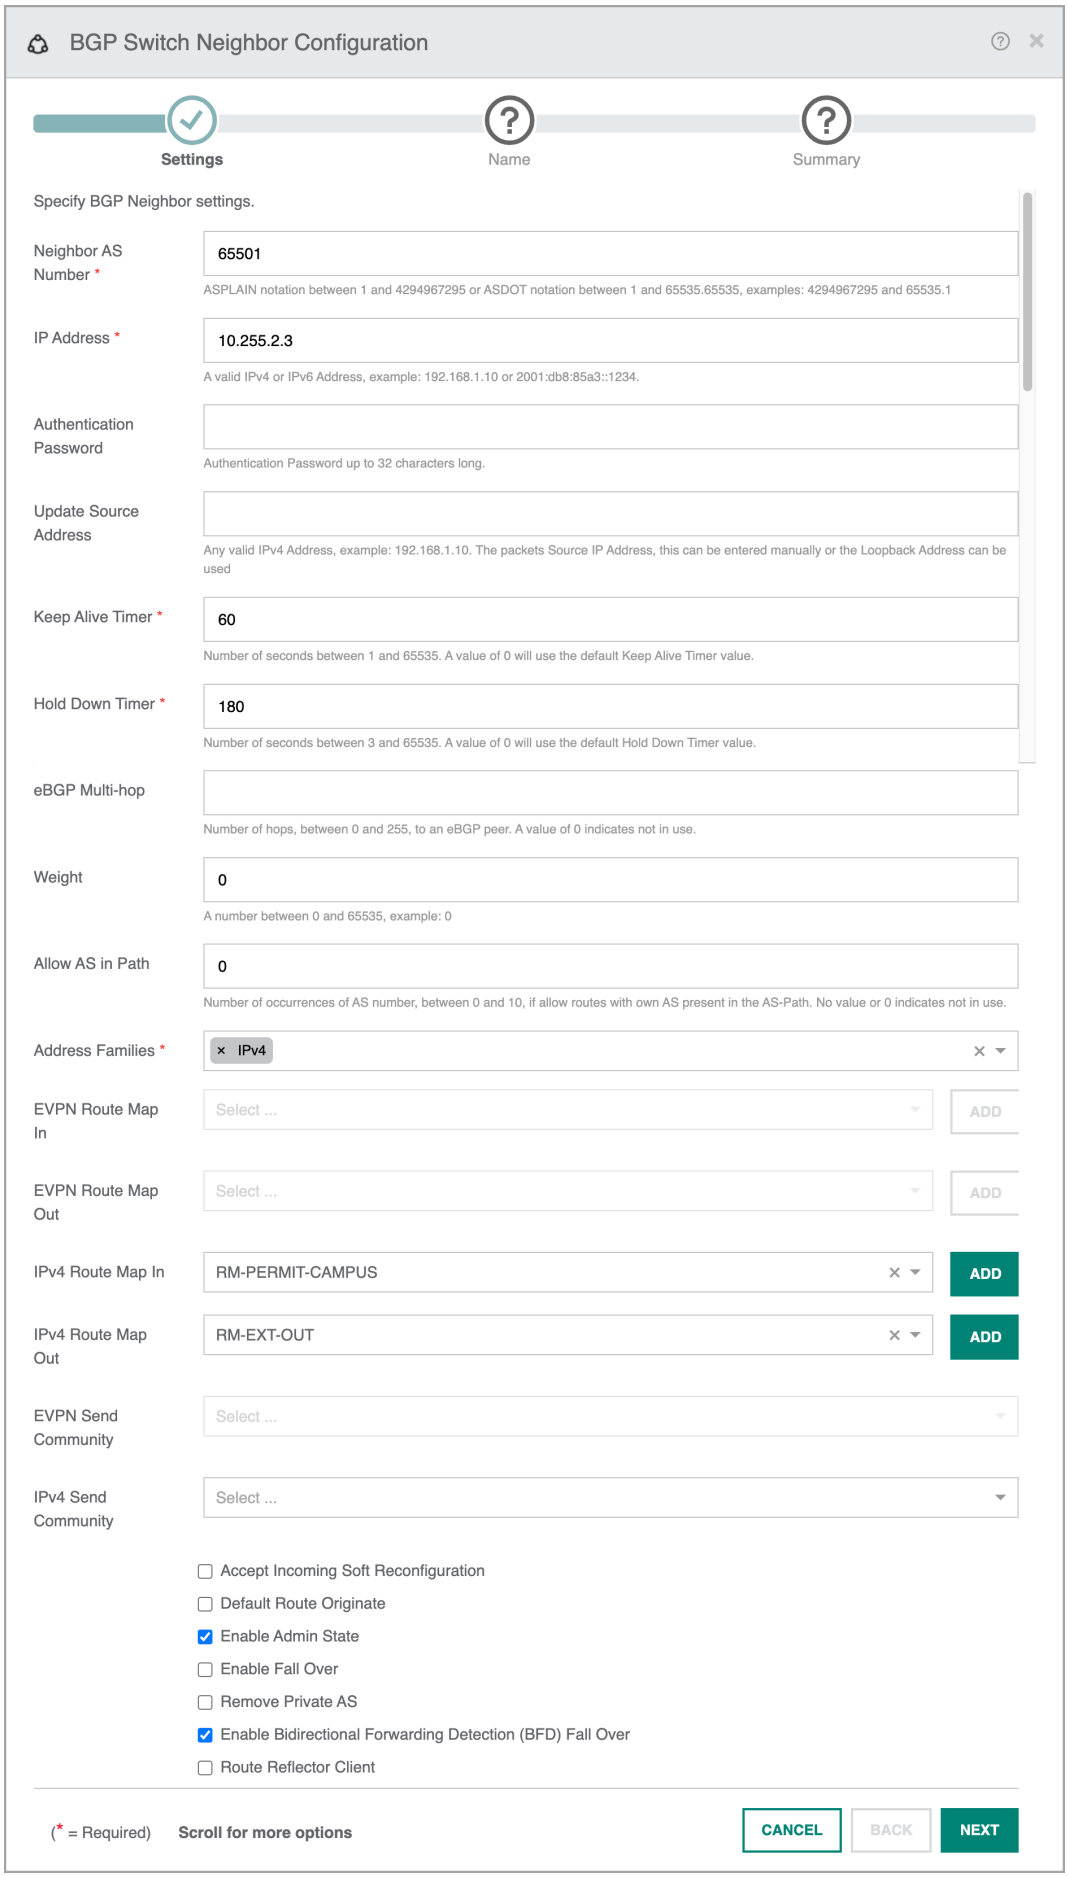
BGP Switch Neighbor Configuration
Settings
Name
Summary
Specify BGP Neighbor settings.
Neighbor AS Number *
65501
ASPLAIN notation between 1 and 4294967295 or ASDOT notation between 1 and 65535.65535, examples: 4294967295 and 65535.1
IP Address *
10.255.2.3
A valid IPv4 or IPv6 Address, example: 192.168.1.10 or 2001:db8:85a3::1234.
Authentication Password
Authentication Password up to 32 characters long.
Update Source Address
Any valid IPv4 Address, example: 192.168.1.10. The packets Source IP Address, this can be entered manually or the Loopback Address can be used
Keep Alive Timer *
60
Number of seconds between 1 and 65535. A value of 0 will use the default Keep Alive Timer value.
Hold Down Timer *
180
Number of seconds between 3 and 65535. A value of 0 will use the default Hold Down Timer value.
eBGP Multi-hop
Number of hops, between 0 and 255, to an eBGP peer. A value of 0 indicates not in use.
Weight
0
A number between 0 and 65535, example: 0
Allow AS in Path
0
Number of occurrences of AS number, between 0 and 10, if allow routes with own AS present in the AS-Path. No value or 0 indicates not in use.
Address Families *
IPv4
EVPN Route Map In
Select ...
ADD
EVPN Route Map Out
Select ...
ADD
IPv4 Route Map In
RM-PERMIT-CAMPUS
ADD
IPv4 Route Map Out
RM-EXT-OUT
ADD
EVPN Send Community
Select ...
IPv4 Send Community
Select ...
Accept Incoming Soft Reconfiguration
Default Route Originate
Enable Admin State
Enable Fall Over
Remove Private AS
Enable Bidirectional Forwarding Detection (BFD) Fall Over
Route Reflector Client
(* = Required)
Scroll for more options
CANCEL
BACK
NEXT

**Step 7** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

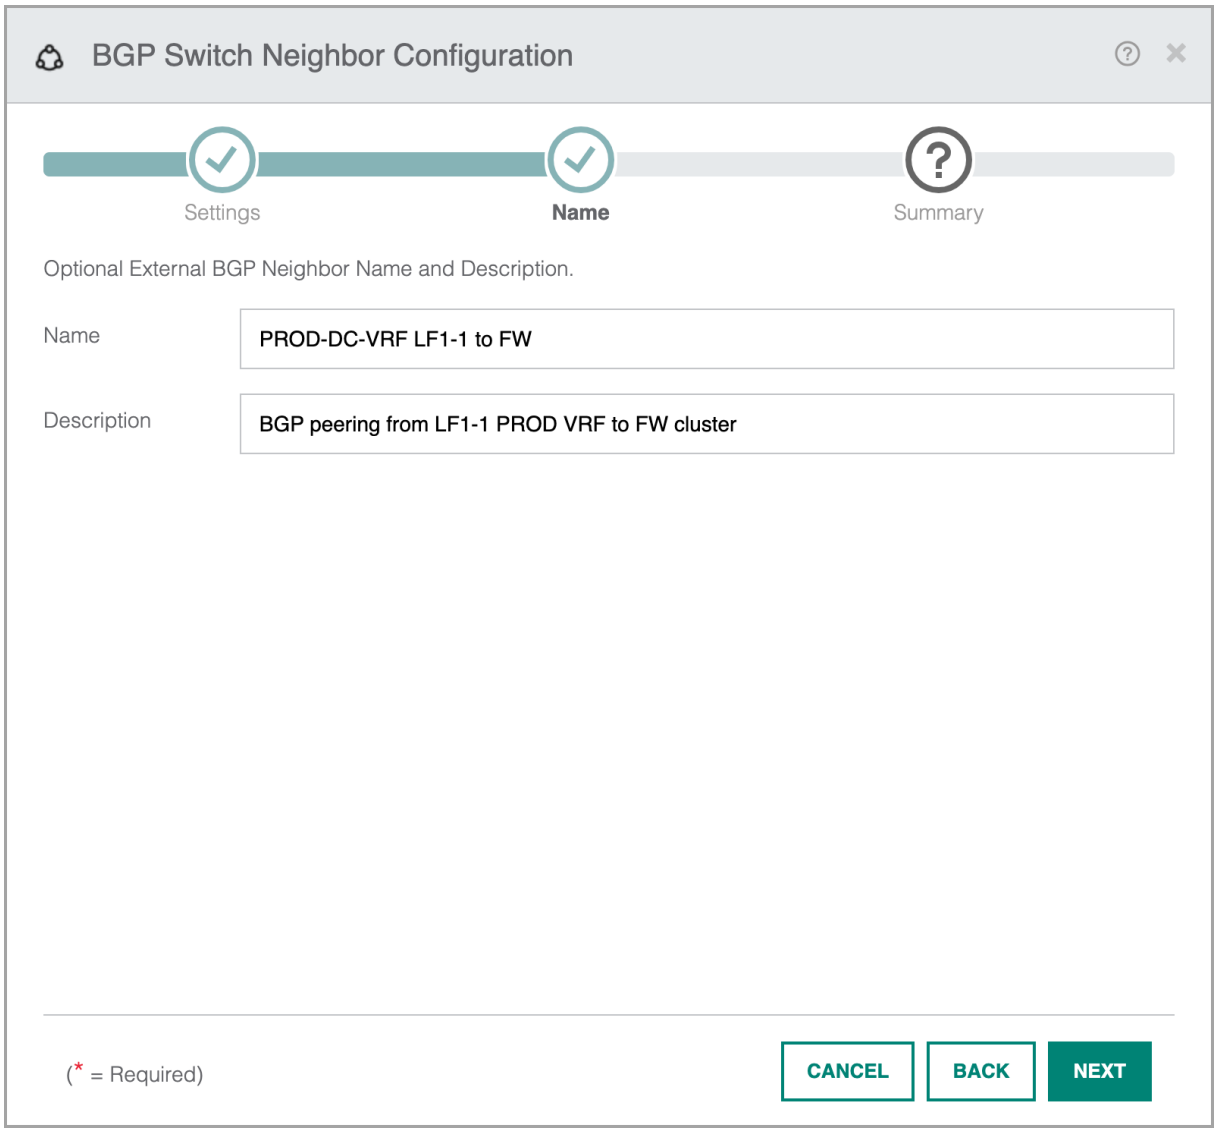


**Step 8** On the **Summary** page, review the BGP neighbor settings and click **APPLY**.

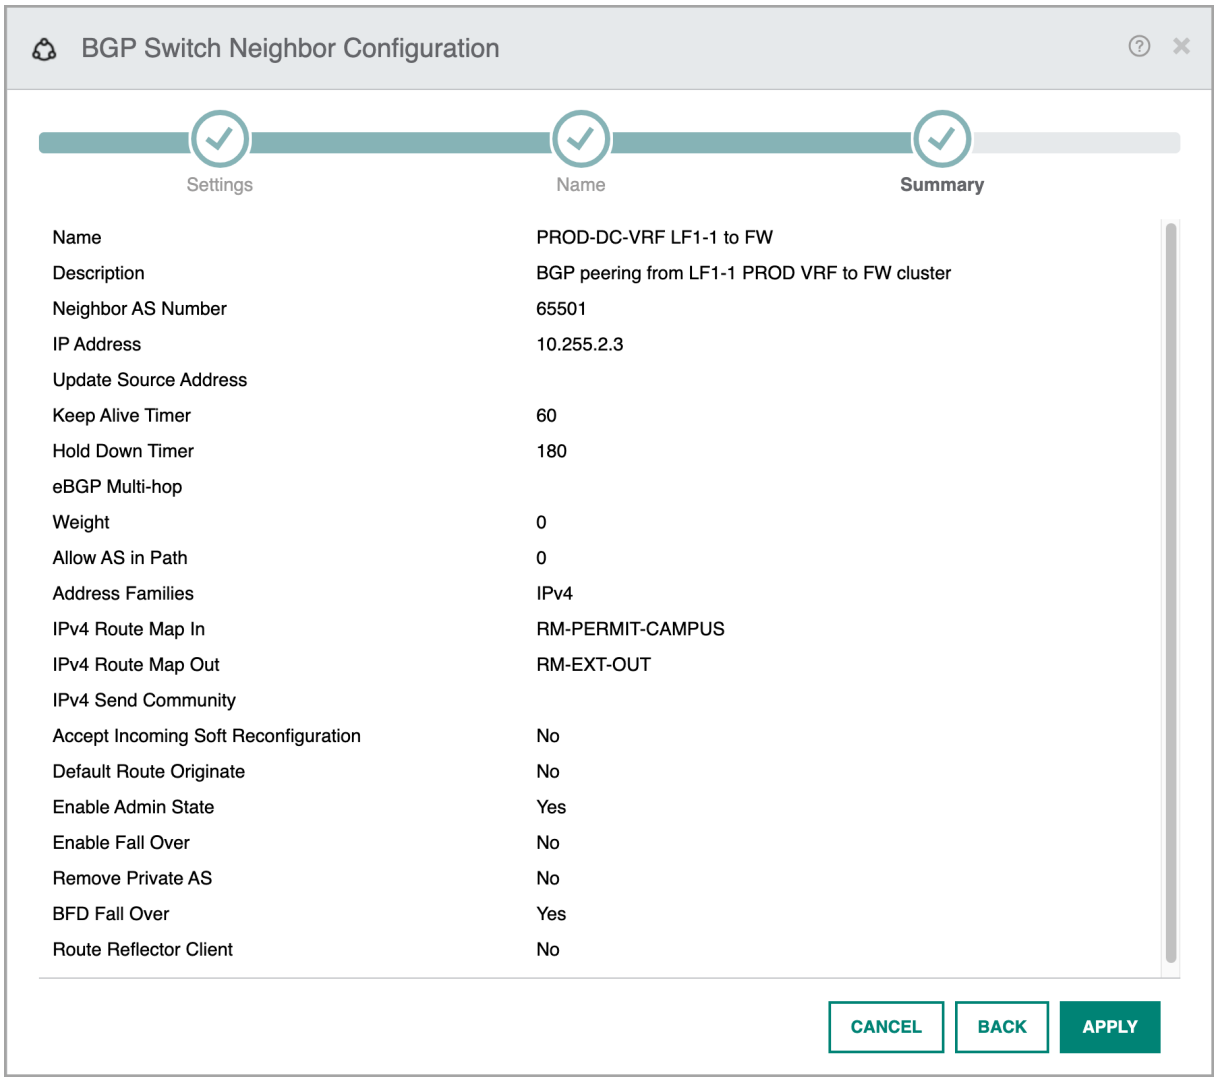


**Step 9** Repeat steps 5 to 8 to add an iBGP peering between the border leaf switches in the production VRF with the following non-default settings:

| Name | Description | Neighbor ASN | IP Addresses | IPv4 Route Map Out |
|---|---|---|---|---|
| PROD-DC-VRF LF1-1 to LF1-2 | PROD VRF peering between border leaf switches | 65001 | 10.255.2.1 | RM-PERMIT-CAMPUS |

**Step 10** In the top left current context path, click **PROD-DC-VRF**.

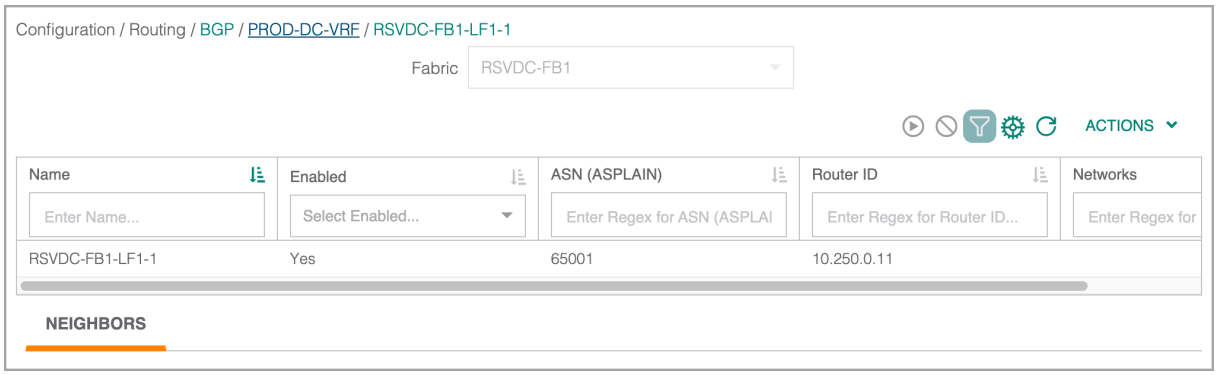


**Step 12** On the **SWITCHES** tab, click • • • next to **RSVDC-FB1-LF1-2** and select **Neighbors**.

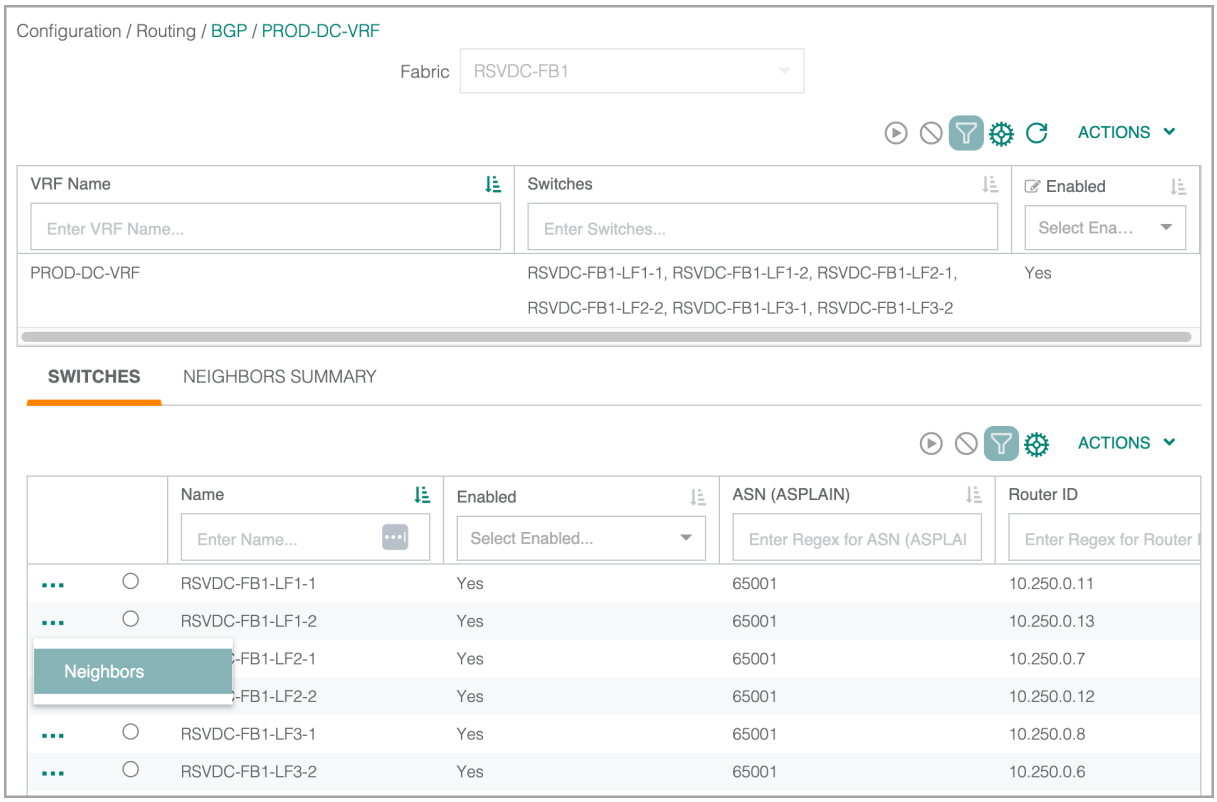


**Step 13** Repeat steps 6 to 9 to create additional BGP peerings on RSVDC-FB1-LF1-2 with the following settings:

| Name | Description | Neighbor ASN | IP Addresses | IPv4 Route Map In | IPv4 Route Map Out | BFD |
|---|---|---|---|---|---|---|
| PROD-DC-VRF LF1-2 to FW | BGP peering from LF1-2 PROD VRF to FW cluster | 65501 | 10.255.2. | RM-PERMIT-CAMPUS | RM-EXT-OUT | < checked > |
| PROD-DC-VRF LF1-2 to LF1-1 | PROD VRF peering between border leaf switches | 65001 | 10.255.2.2 | | RM-PERMIT-CAMPUS | < unchecked > |

**Step 13** Repeat this procedure for each overlay VRF network that requires external connectivity. Reachability between overlay VRFs is governed by policy at the upstream firewall. Strict overlay route table separation can be maintained by connecting to discrete VRFs or virtual firewall contexts on the upstream firewall.

| Name | Description | Neighbor ASN | IP Addresses | IPv4 Route Map In | IPv4 Route Map Out | BFD |
|---|---|---|---|---|---|---|
| DEV-DC-VRF LF1-1 to FW | BGP peering from LF1-1 DEV VRF to FW cluster | 65501 | 10.255.2. | RM-PERMIT-CAMPUS | RM-EXT-OUT | < checked > |
| DEV-DC-VRF LF1-1 to LF1-2 | DEV VRF peering between border leaf switches | 65001 | 10.255.2.9 | | RM-PERMIT-CAMPUS | < unchecked > |
| DEV-DC-VRF LF1-2 to FW | BGP peering from LF1-2 DEV VRF to FW cluster | 65501 | 10.255.2. | RM-PERMIT-CAMPUS | RM-EXT-OUT | < checked > |
| DEV-DC-VRF LF1-2 to LF1-1 | DEV VRF peering between border leaf switches | 65001 | 10.255.2.10 | | RM-PERMIT-CAMPUS | < unchecked > |

## Verify Border Leaf Routing

**Step 1** In the top-left current context path, click **BGP**.

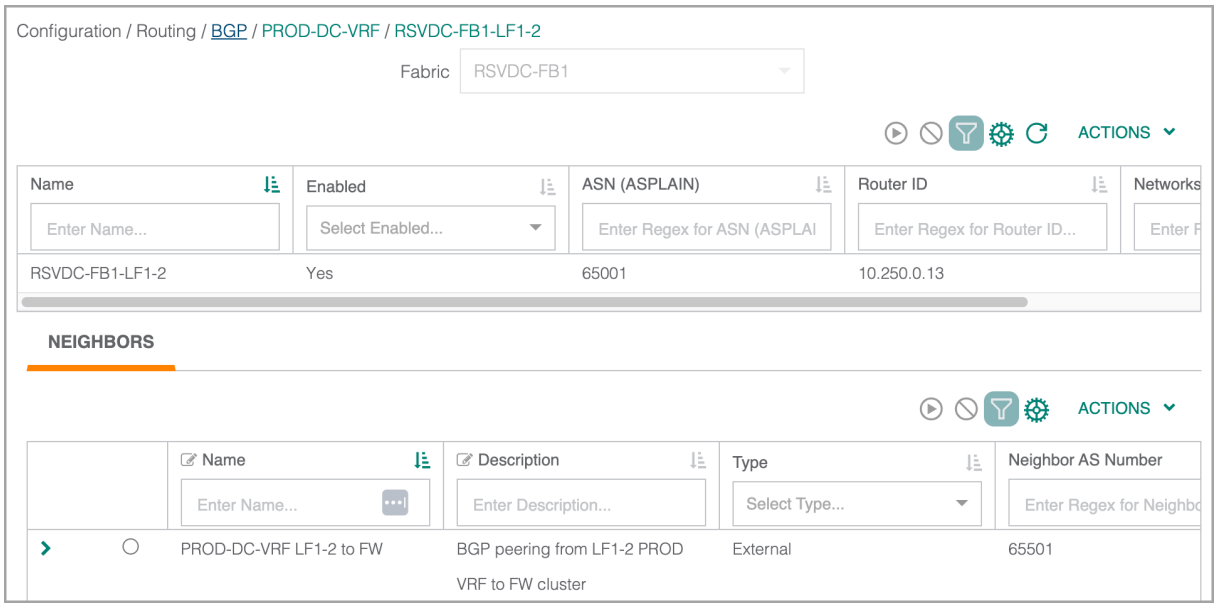

**NOTE:**

To display information and the current state of an individual BGP peering, click the expansion icon (**>**) at the beginning of the row for each BGP peer definition. After a BGP peering is defined, the Fabric Composer web page may require a refresh to display the expansion icon.

**Step 2** Click • • • next to **PROD-DC-VRF** and select **Neighbors Summary**.

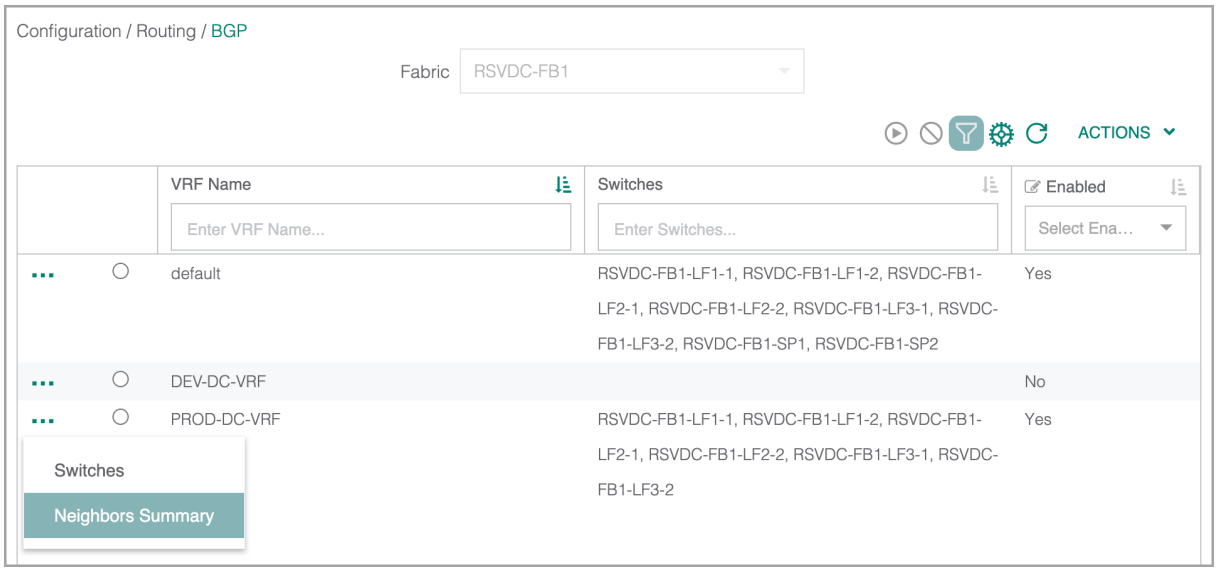


**Step 3** In the **NEIGHBORS SUMMARY** window, verify that each peering displays **Established** in the **State** column.

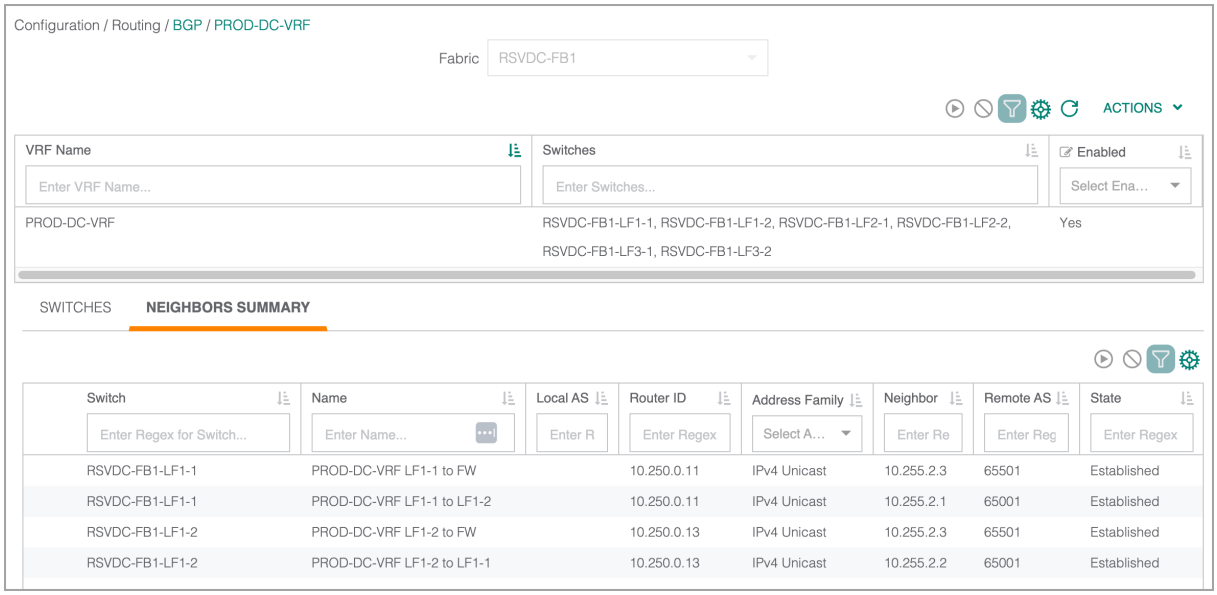


**Step 4** Repeat steps 1 to 3 for each overlay VRF.

**Step 5** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

**Step 6** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select all leaf switches* >

- **Commands:** *show ip route bgp vrf PROD-DC-VRF*

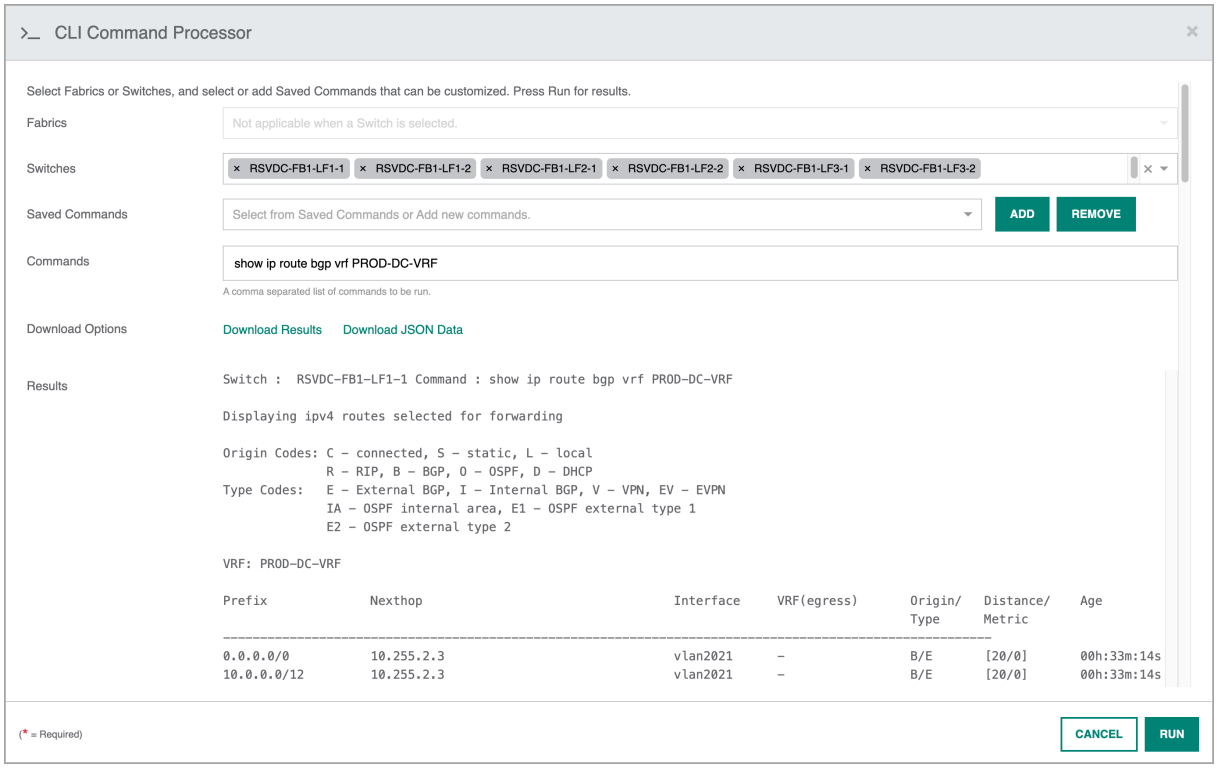


**Step 7** Verify that there is a default route and campus summary route learned on all leaf switches in the production VRF. The border leaf switch routes use the upstream firewall IP as a next hop. The remaining leaf switches use a next hop of the border leaf Anycast VTEP, learned via BGP EVPN type-5 advertisements.

**NOTE:**

The prefixes advertised into an overlay fabric vary based on the environment. A default route is often the only learned prefix required. The campus summary route is used in the Validation Solution Guide's multifabric configuration.

**Step 8** Repeat steps 6 to 7 for each overlay VRF.

## Configure Overlay Test Loopbacks

A unique loopback IP per switch in each overlay VRF is required to verify connectivity to directly attached hosts and reachability through the overlay.

Sourcing a ping from a switch to one of its directly attached hosts is a common method to verify reachability at the point of attachment. When a VSX leaf pair provides redundant links to attached hosts, the return data path from the host may not be the same link as the originating traffic. By default, a switch will use the SVI interface IP for the VLAN connecting the downstream host. The same VLAN SVI IP address is configured on both VSX member switches to conserve IP address space. If the response to a ping originated by one member of the VSX pair is received by the other member, the response will be dropped, because the switch receiving the response has no state for the ping conversation.

Sourcing a ping from a unique IP address that is present only on one of the VSX switch pair members resolves the issue, when combined with a static route to share reachability of the unique IP between the VSX pair members. On any leaf switch, one loopback IP address in a VRF can be used to test reachability to all locally attached hosts in all subnets associated locally with that VRF. If a response is received by the non-originating member of the VSX pair, the destination IP address is not local to the VSX member, so the route table is consulted and the static route is used to forward the ping response to the originating member of the pair. A unique IP loopback per VRF per switch is required for full testing capability.

A similar problem exists when sourcing a ping from a switch to verify overlay reachability. Up to this point, the only IP interfaces configured in the overlay are VLAN SVIs. For each VLAN, the same SVI IP address is assigned to all leaf switches. Sourcing a ping with a VLAN SVI to an IP host connected to another leaf in the fabric will result in the ping response being dropped at the remote host's point of attachment, which also owns the destination IP address in the ping response. In the case where a ping response can be received by the originating VTEP, it is not guaranteed that the response will be received by the originating switch. A VSX pair of switches represent a single logical VTEP. If the response to a ping originated by one member of the VSX pair is received by the other member, the response will be dropped for the same reason noted above.

Sourcing a ping from an IP address that is unique to an individual switch resolves the reachability problem in the overlay as well.

The following procedure configures a transit VLAN in each overlay VRF between leaf switch VSX pairs, unique loopback on each switch in both the PROD and DEV VRFs, and a static route using the transit VLAN for loopback reachability between the pair. The loopback address can be used to test overlay reachability to directly attached hosts and remote IP destinations.

The following procedure configures the required elements for one VSX leaf pair.

## Configure Overlay Transit VLAN

Each redundant pair of ToR switches requires a transit VLAN in the overlay VRF to enable routed reachability to its VSX partner's IP loopback address in the same VRF.

**Step 1** On the **Configuration** menu, select **Routing > VRF**.

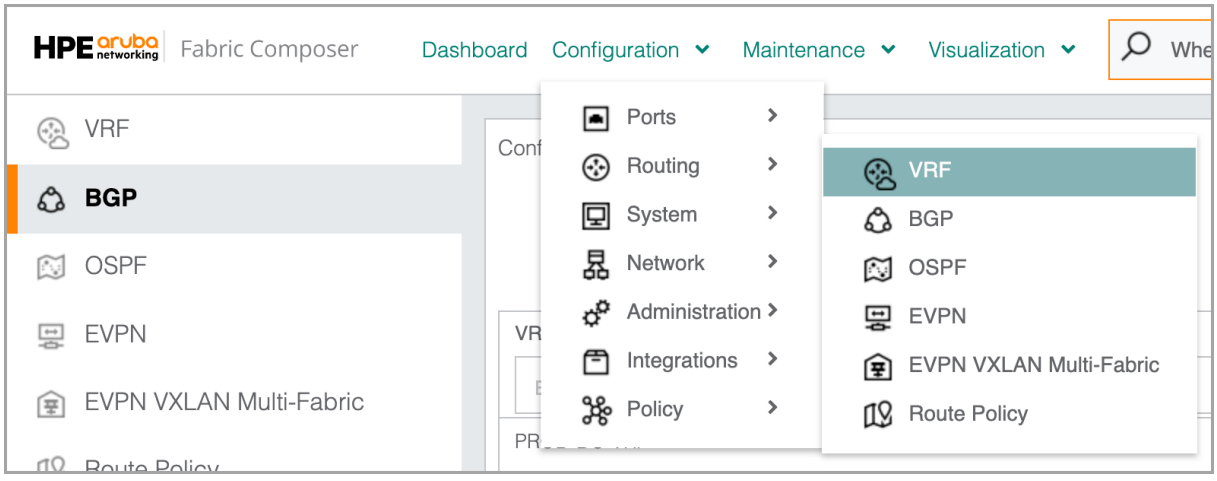

**Step 2** Click • • • next to **PROD-DC-VRF** and select **IP Interfaces**.

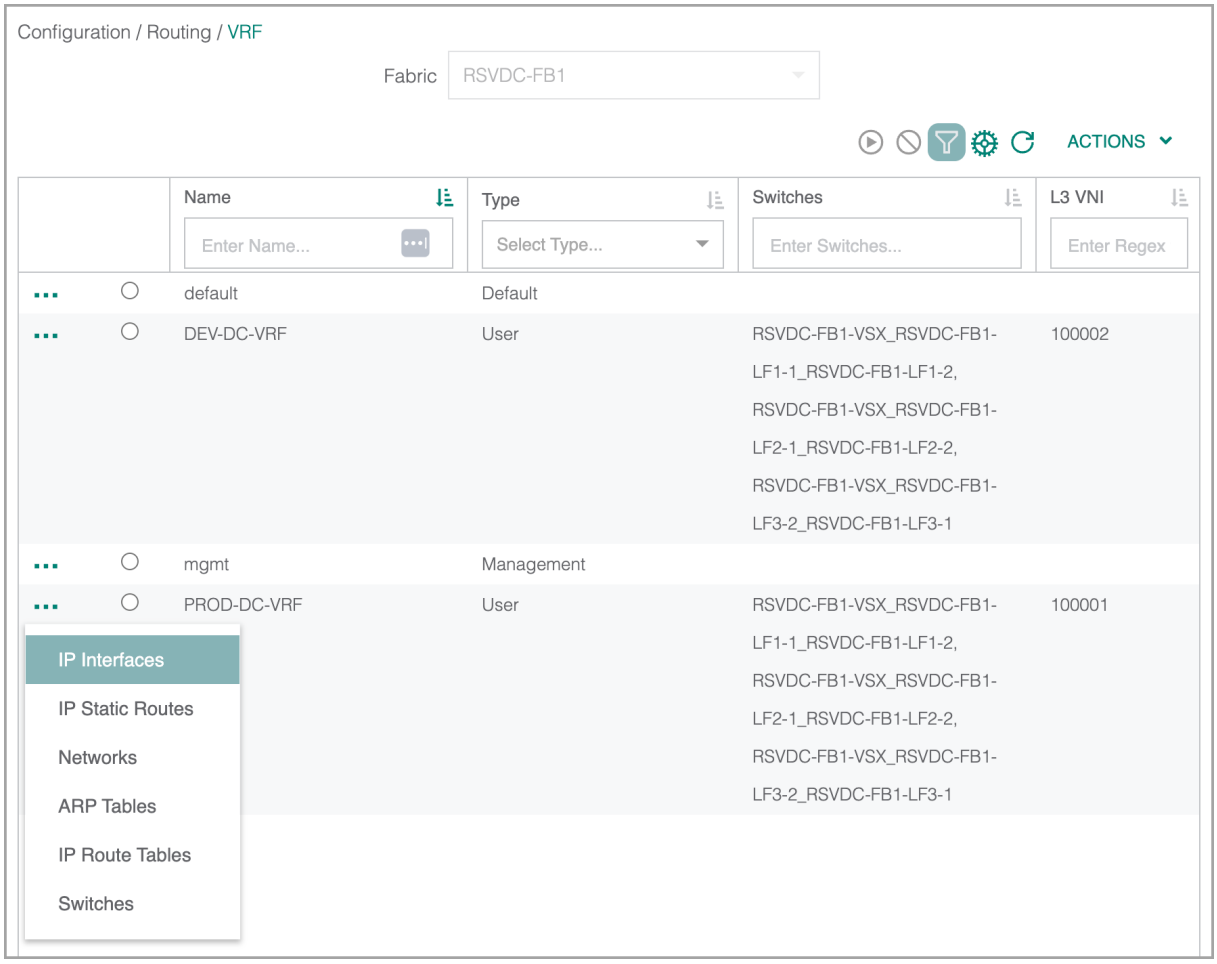


**Step 3** On the right **ACTIONS** menu, select **Add**.

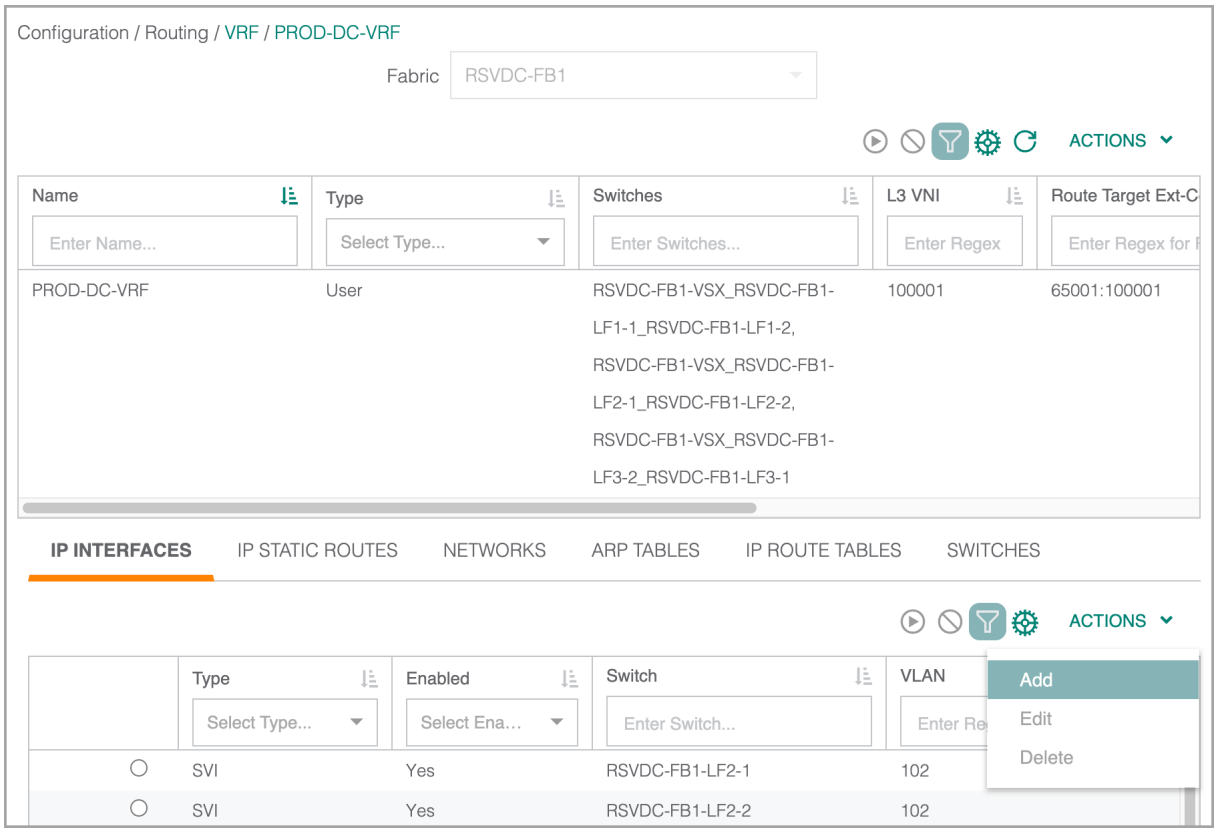


**Step 4** On the **Interface Type** page, assign the following non-default values and click **NEXT**.

- **VLAN**: *3001*
- **Switches**: *< Select a VSX leaf switch pair >*
- **IPv4 Subnetwork Address**: *< Assign a /31 block of addresses >*
- **IPv4 Address**: *< Assign the range of 2 IP address that comprise the /31 subnet >*

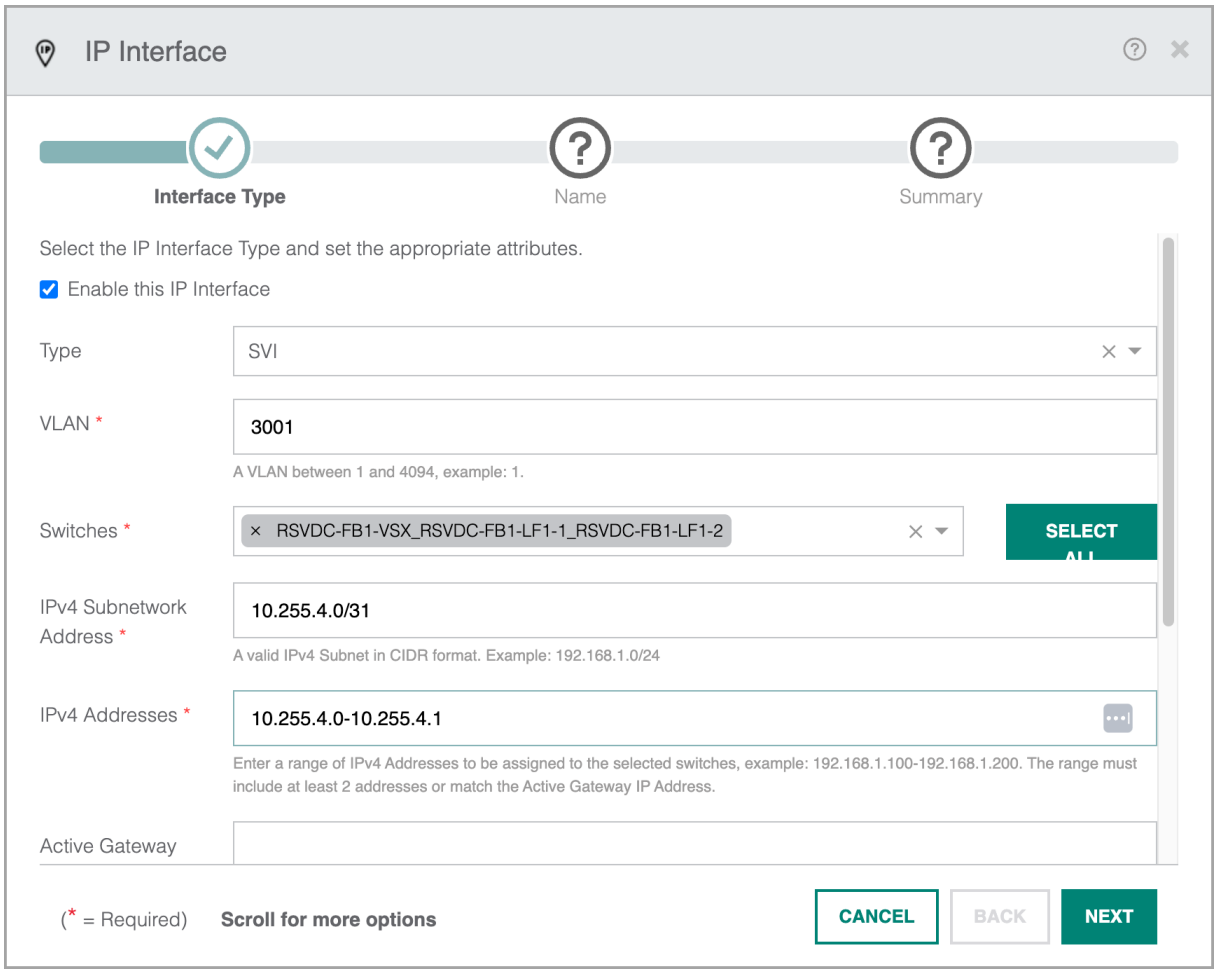


**Step 5** On the **Name** page, assign a **Name** and **Description**, then click **NEXT**.

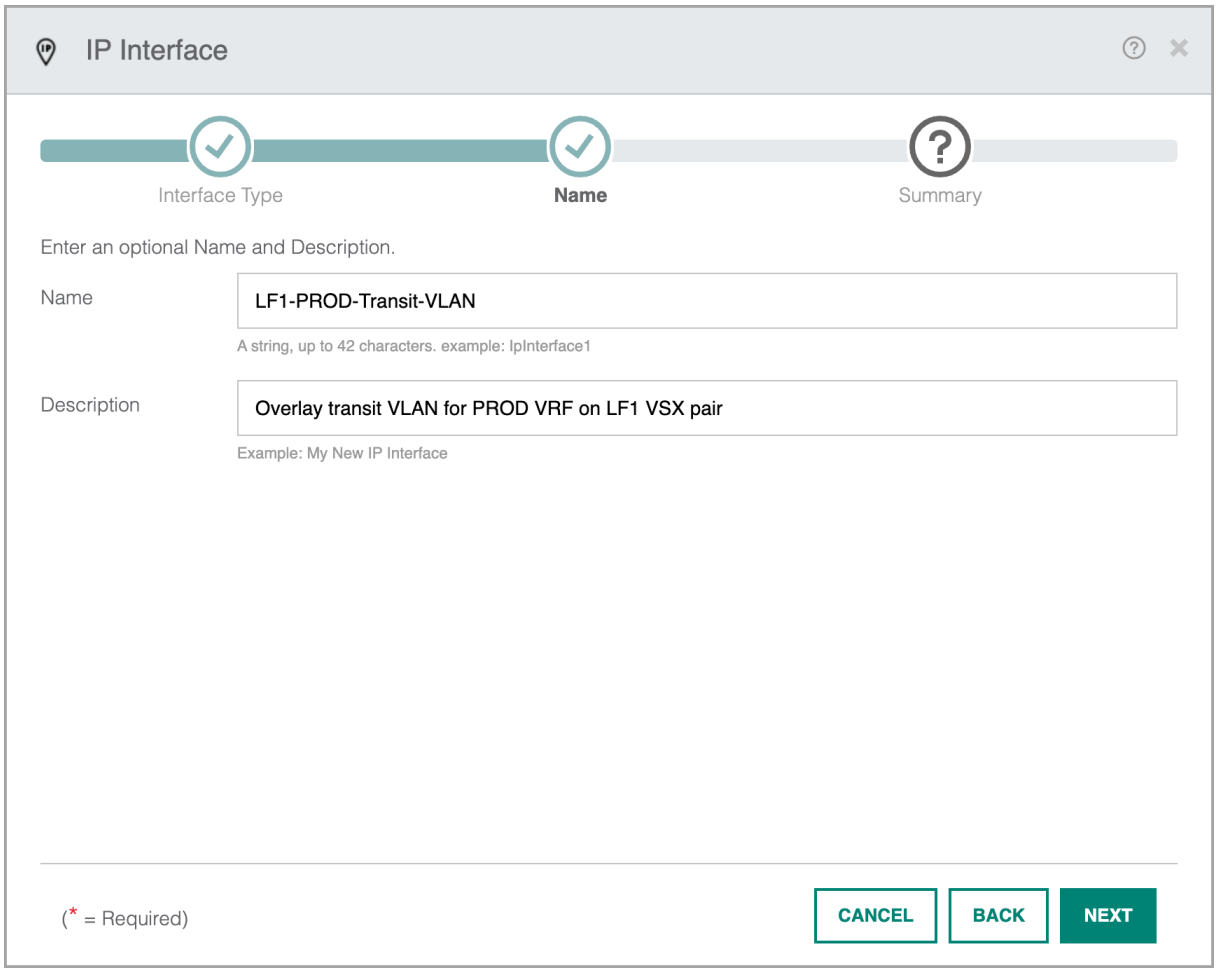


**Step 6** On the **Summary** page, review the transit VLAN settings and click **APPLY**.

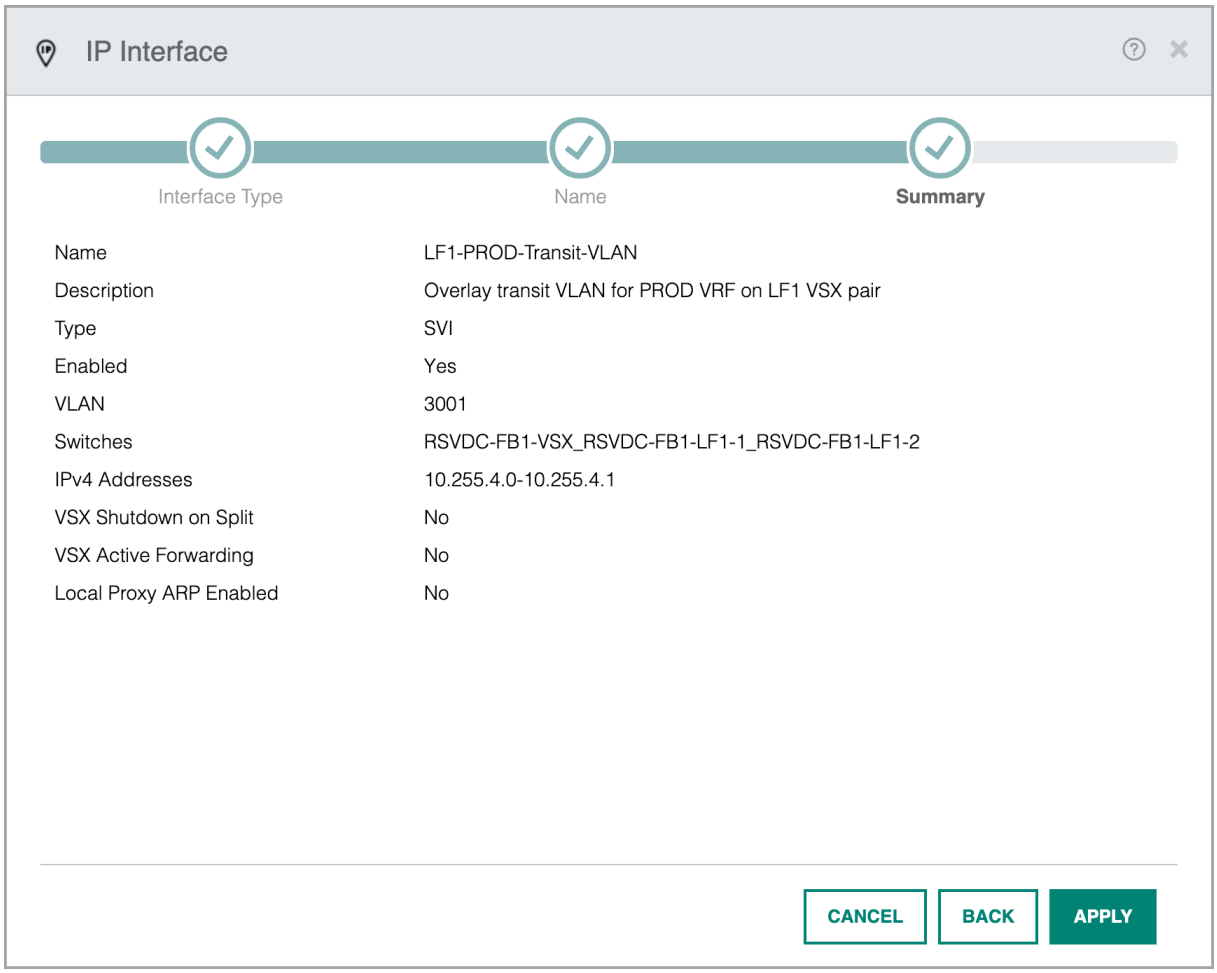


## Assign Unique Overlay Loopbacks

**Step 1** On the right **ACTIONS** menu of the **Configuration > Routing > VRF > PROD-DC-VRF** page, select **Add**.

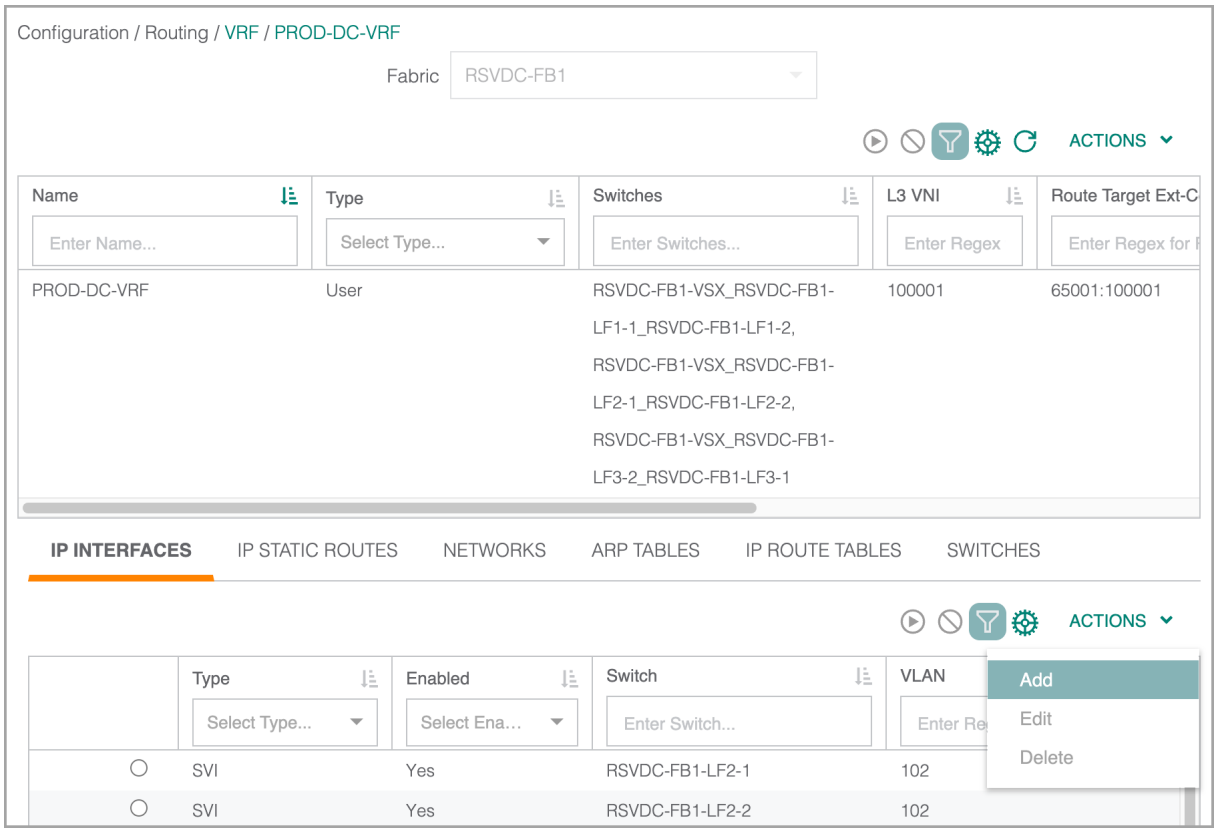


**Step 2** On the **Interface Type** page, enter the following non-default values:

- **Type**: *Loopback*
- **Loopback Name**: *< An unused loopback interface value >*
- **Switch**: *< Select an individual switch in the VSX pair where a transit VLAN was created >*

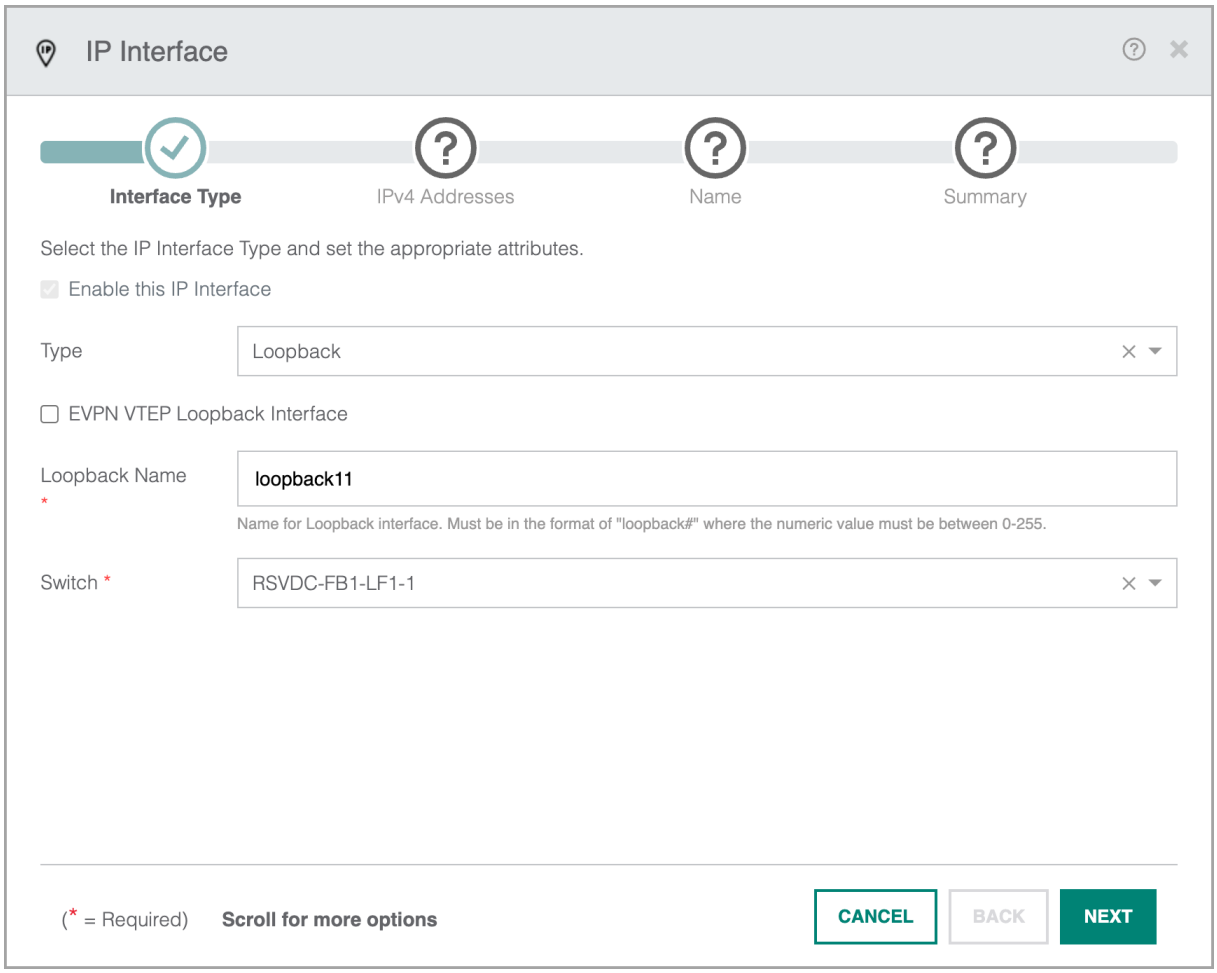


**Step 3** On the **IPv4 Addresses** page, enter a /32 host address and click **NEXT**.

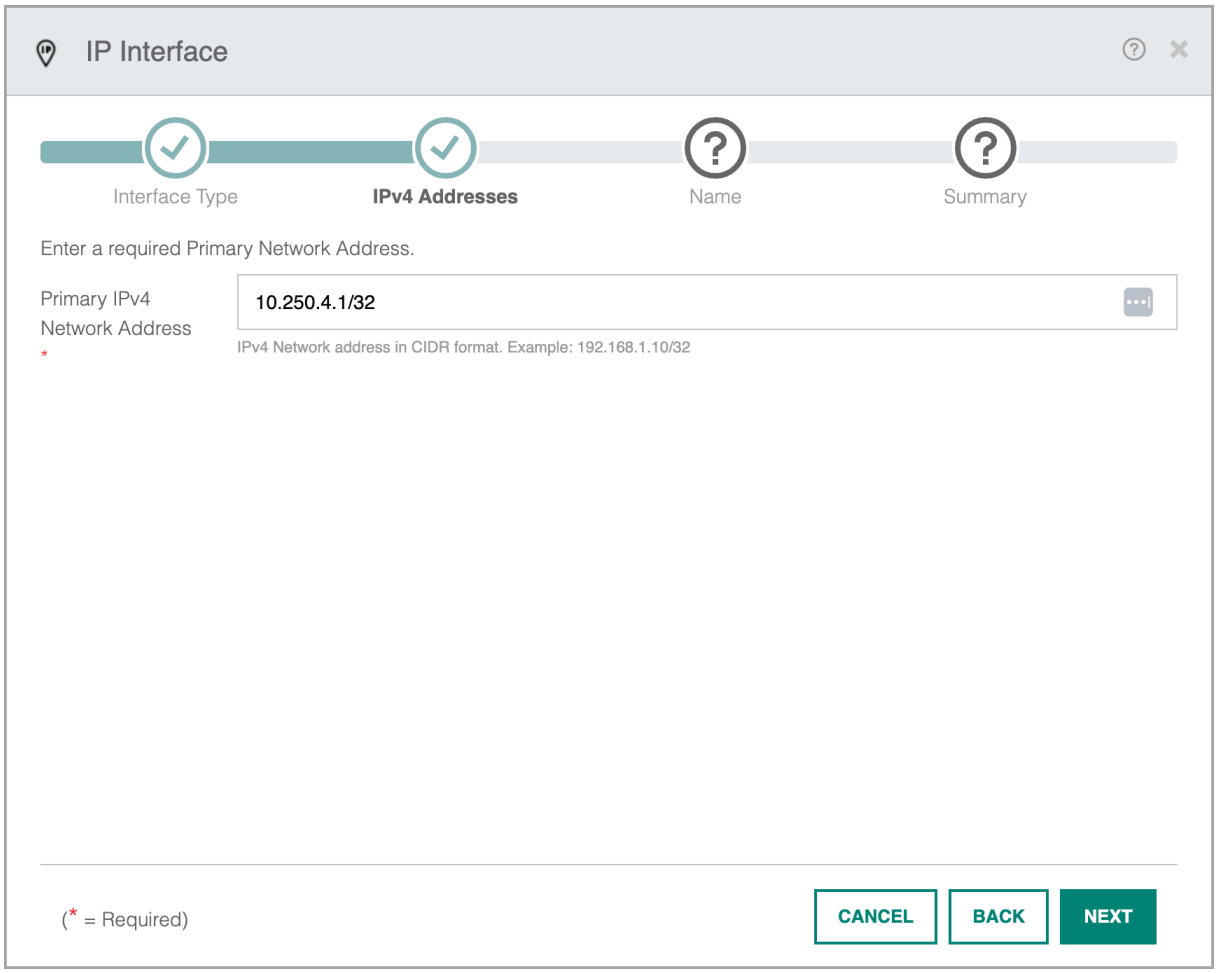


**Step 4** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

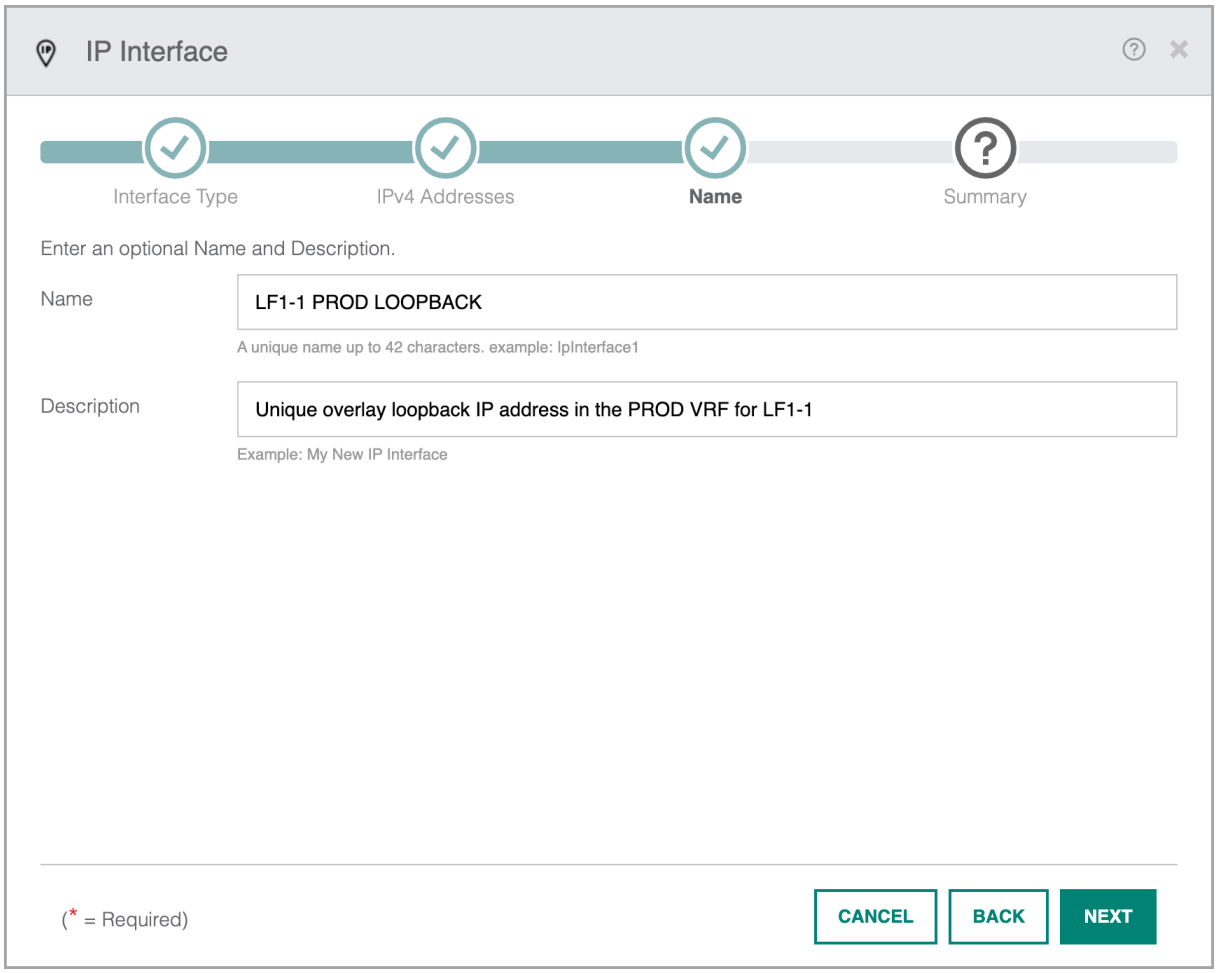


**Step 5** On the **Summary** page, review the loopback interface settings and click **APPLY**

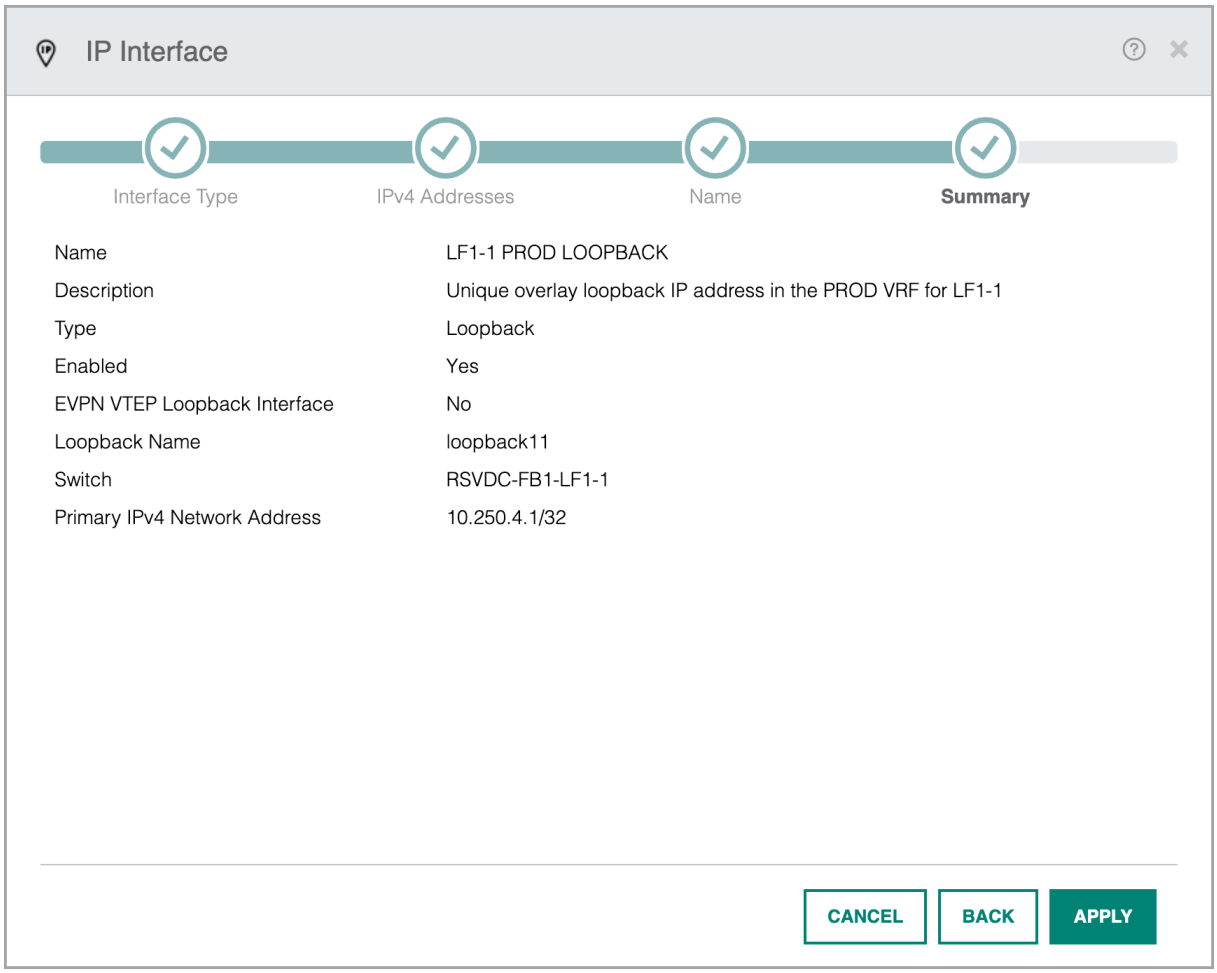


**Step 6** Repeat steps 1 to 5 to assign a loopback interface to the VSX partner switch with the following non-default values:

| Name | Description | Type | Loopback Name | Switch | Primary IPv4 Network Address |
|---|---|---|---|---|---|
| LF1-2 PROD LOOPBACK | Unique overlay loopback IP address in the PROD VRF for LF1-2 | Loopb | Loopback | RSVDC-FB1-LF1-2 | 10.250.4.0/32 |

## Configure Static Route for VSX Routed Loopback Reachability

**Step 1** On the **Configuration > Routing > VRF > PROD-DC-VRF** page, click **IP STATIC ROUTES**. On the right **ACTIONS** menu, select **Add**.

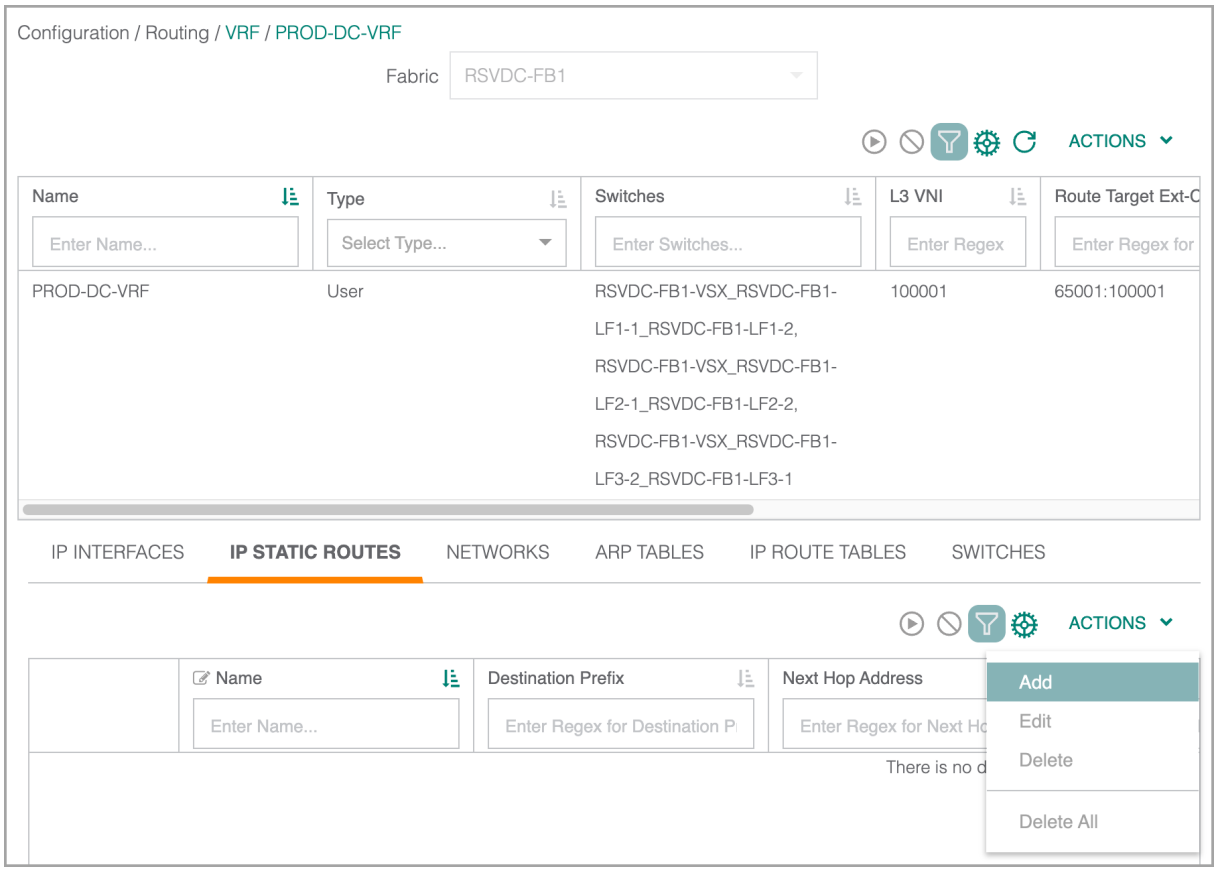


**Step 2** On the **Route** page, enter the following values and click **NEXT**.

- **Destination Prefix**: *< Host IP prefix of VSX peer's overlay loopback interface >*
- **Next Hop Address**: *< IP address of VSX peer's transit VLAN interface >*
- **Switch**: *< Individual VSX member target for static route >*

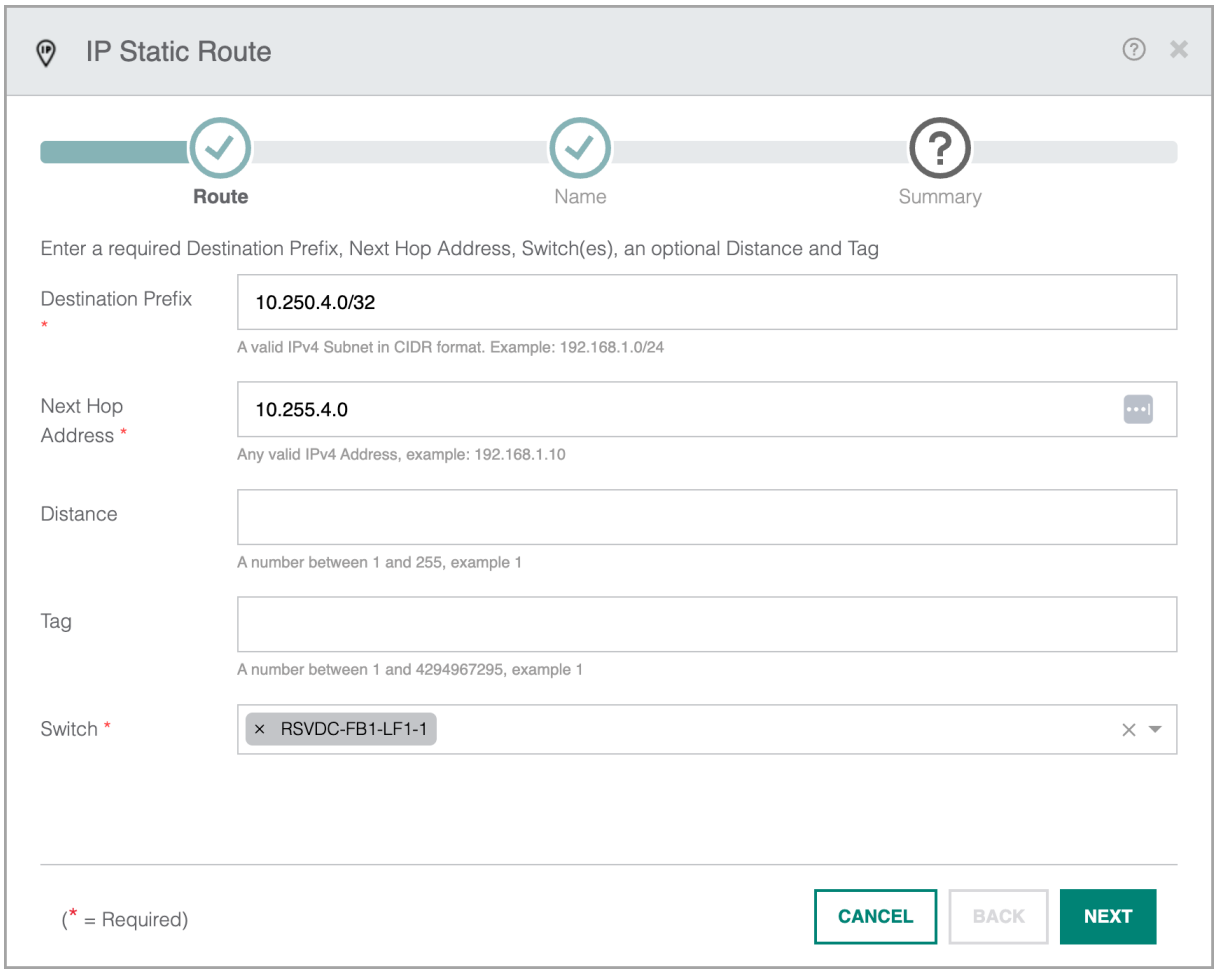


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

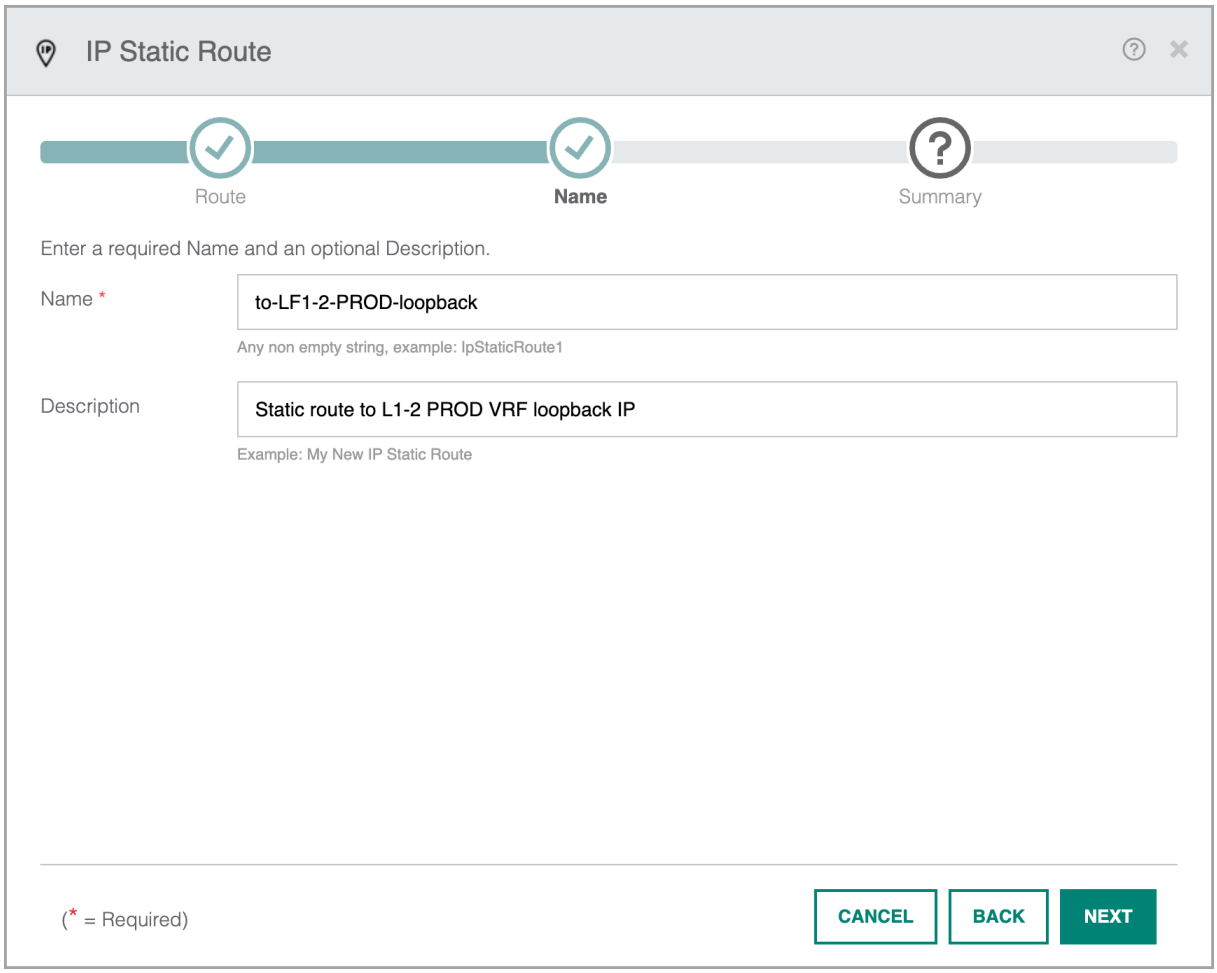


**Step 4** On the **Summary** page, review the static route settings and click **APPLY**.

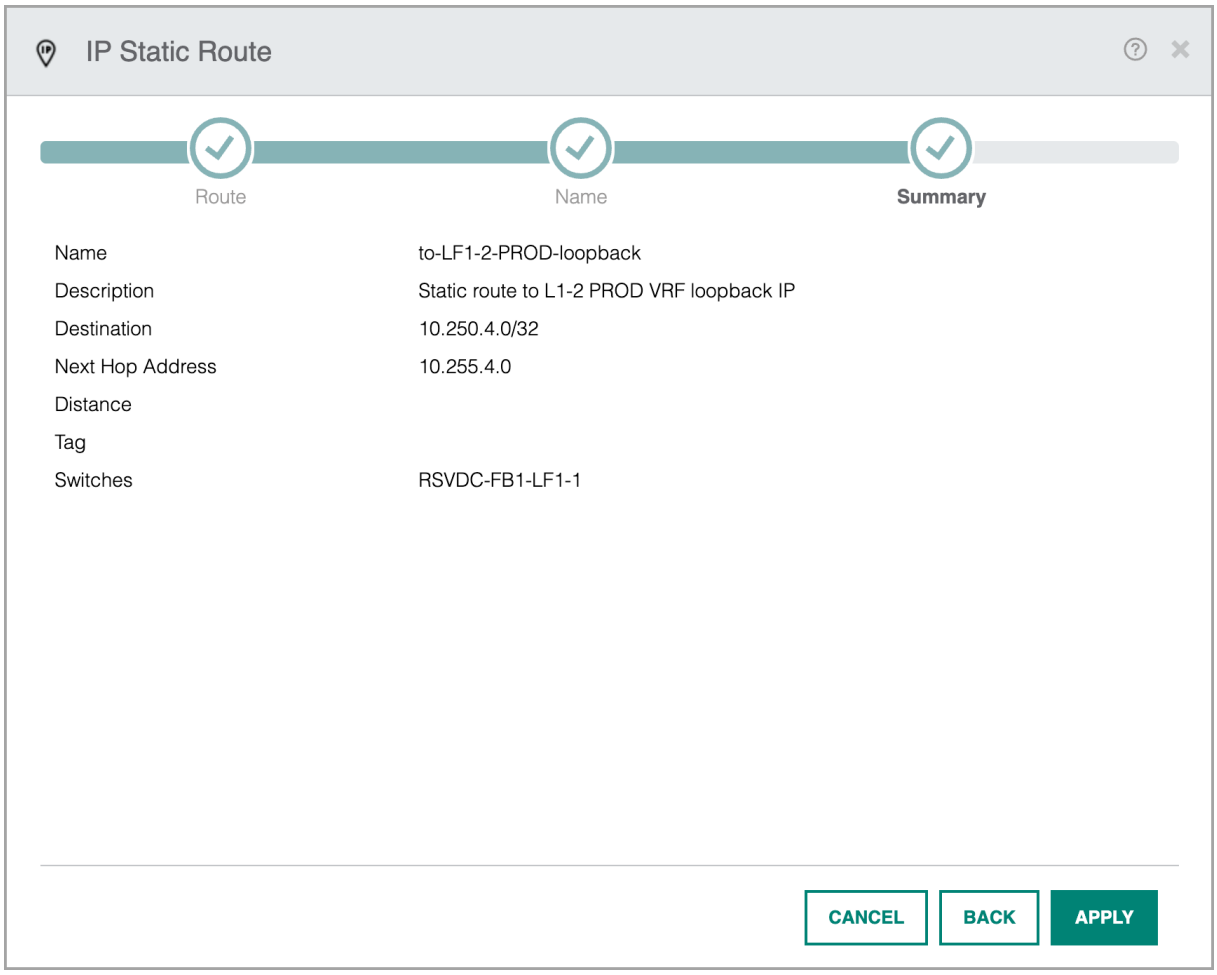


**Step 5** Repeat steps 1 to 4 to create a static route to the VSX peer's loopback interface using the following non-default values:

| Name | Description | Destination | Next Hop Address | Switches |
|---|---|---|---|---|
| to-LF1-1-PROD-loopback | Static route to L1-1 PROD VRF loopback IP | 10.250.4.1/3 | 10.255.4.1 | RSVDC-FB1-LF1-2 |

Repeat the **Configure Overlay Test Loopbacks** procedure for each VRF in the overlay on each VSX leaf pair in the network. For standalone leaf switches, perform only the **Assign Loopbacks to Individual Switches** steps in this section for each VRF in the overlay.

## Configure Loopback Summary Route

A summary static route for the collective set of loopback interfaces is configured on the border leaf to advertise loopback reachability to external networks. This static route points to null. A summary route is created for each VRF in the overlay.

**Step 1** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Configuration Editor**.

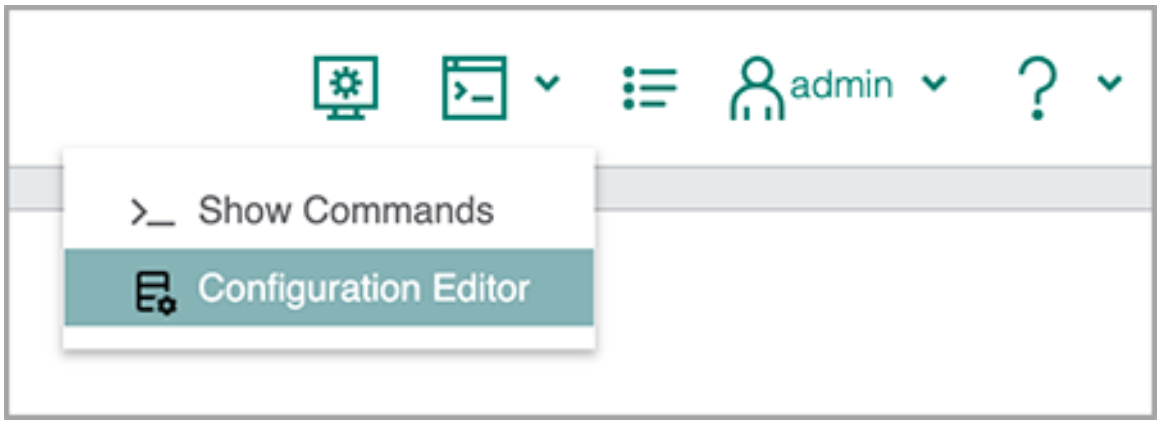


**Step 2** On the **Configuration Editor** page, select the two border leaf switches in the **Switch** field.

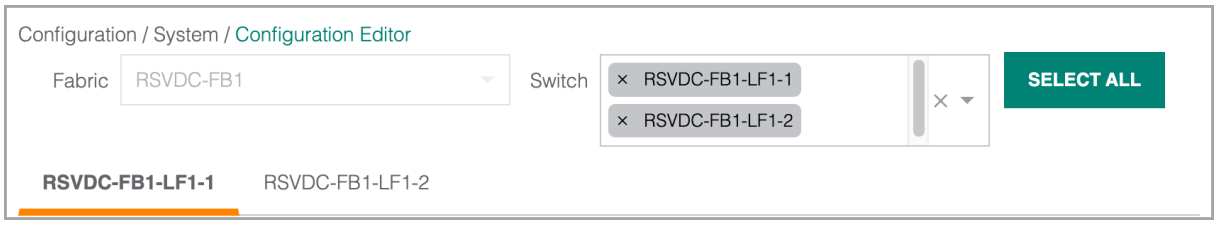


**Step 3** Enter the summary static route configuration for the loopback addresses in the same section where the previously created static routes appear on both switch tabs, and click **VALIDATE ALL**.

```
ip route 10.250.4.0/24 nullroute vrf PROD-DC-VRF
ip route 10.250.5.0/24 nullroute vrf DEV-DC-VRF
```

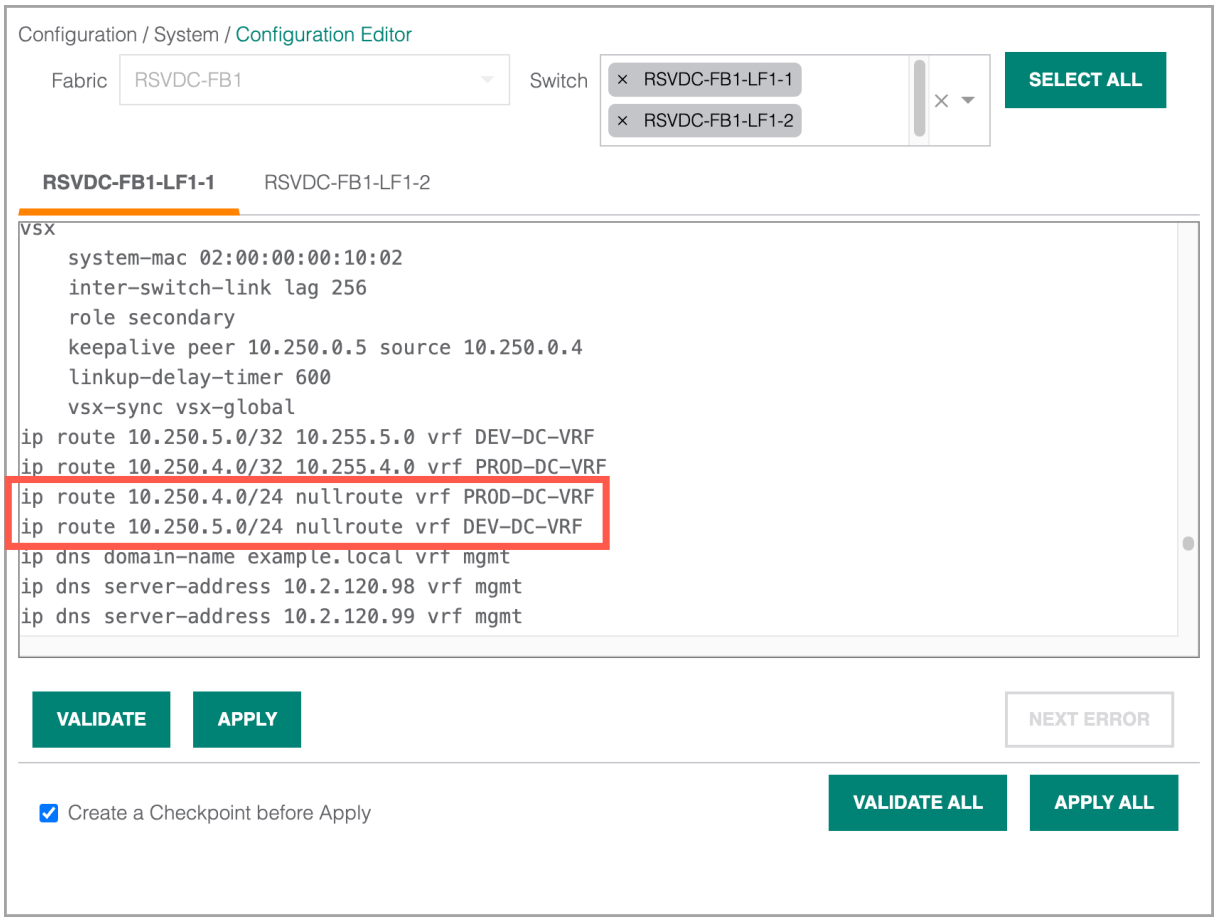

**NOTE:**

Fabric Composer 7.0.X and previous versions require creating static null routes using Fabric Composer's **Configuration Editor** or the switch CLI.

If the configuration is valid, a **Success** message is presented.


**Step 4** Click **APPLY ALL**.

**NOTE:**

**Create Checkpoint before Apply** is selected by default. Fabric Composer creates a checkpoint on the switch to restore the switch configuration back to its state prior to the change.

If the checkpoint was successfully created, the **Success** message indicates the time of its creation.

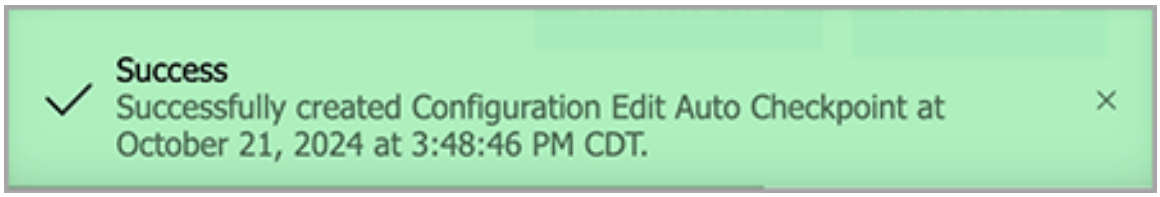


If the configuration is applied successfully, a **Success** message is presented.

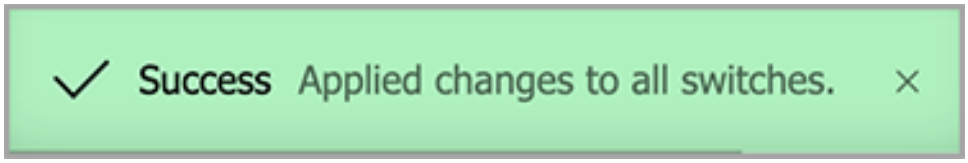


## Redistribute Static Routes on the Border Leaf

The loopback summary route created in the previous procedure is redistributed into BGP for advertisement to campus. Redistribution is applied to each VRF.

**Step 1** On the **Configuration** menu, select **Routing > VRF**.

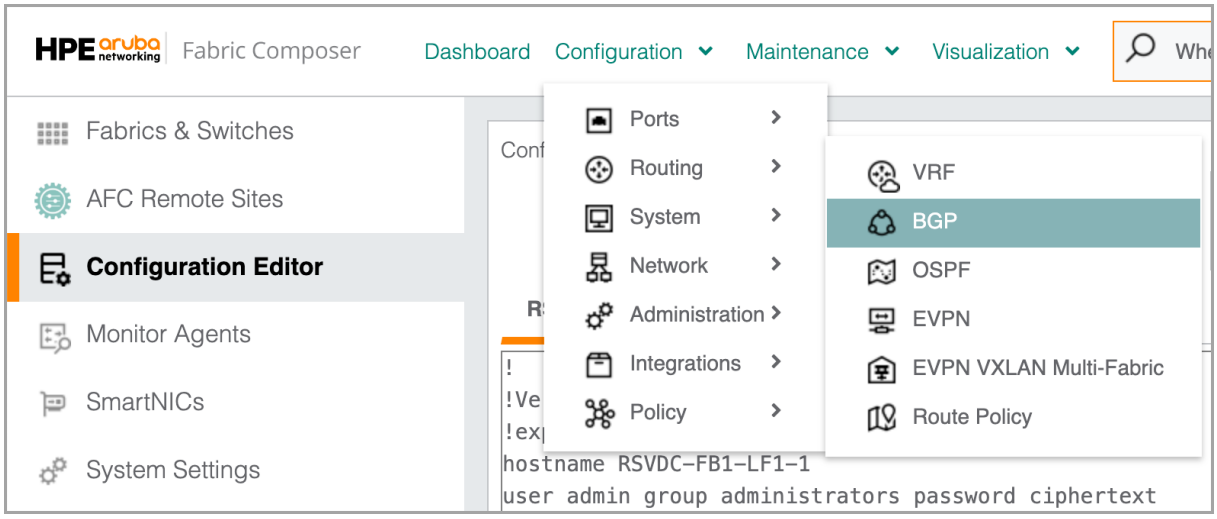

**Step 2** On the **Configuration > Routing > BGP** page, right-click • • • next to **PROD-DC-VRF** and select **Switches**.

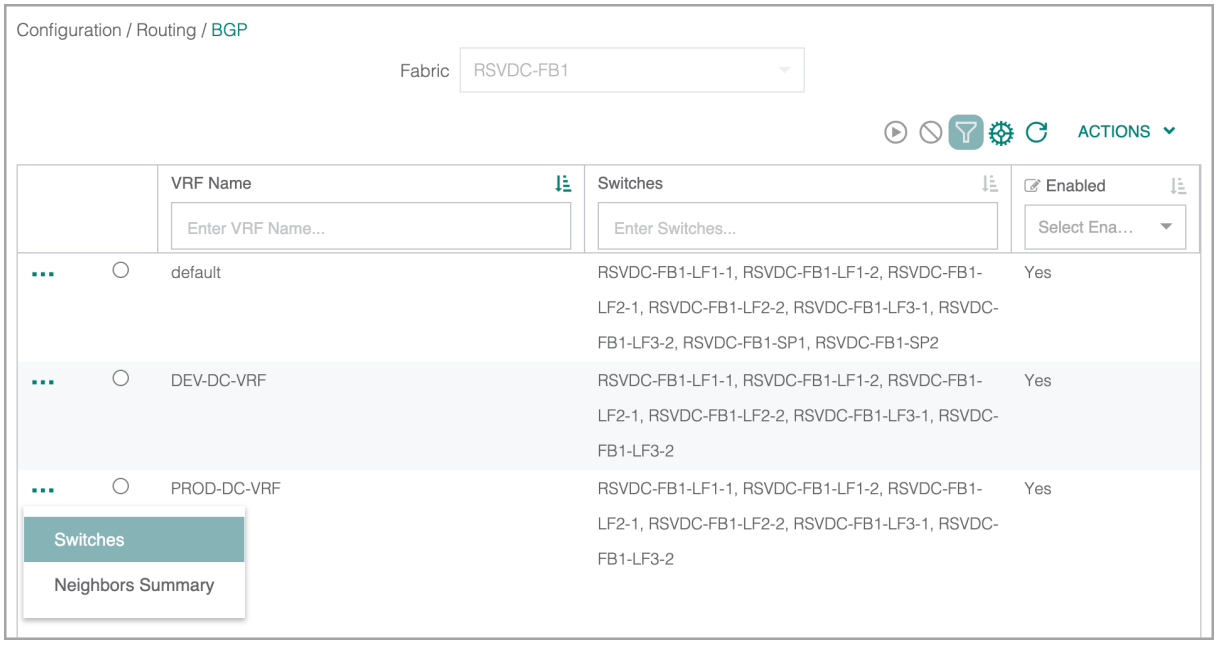


**Step 3** Click the radio button for the **RSVDC-FB1-LF1-1** border leaf switch. On the right **ACTIONS** menu, select **Edit**.

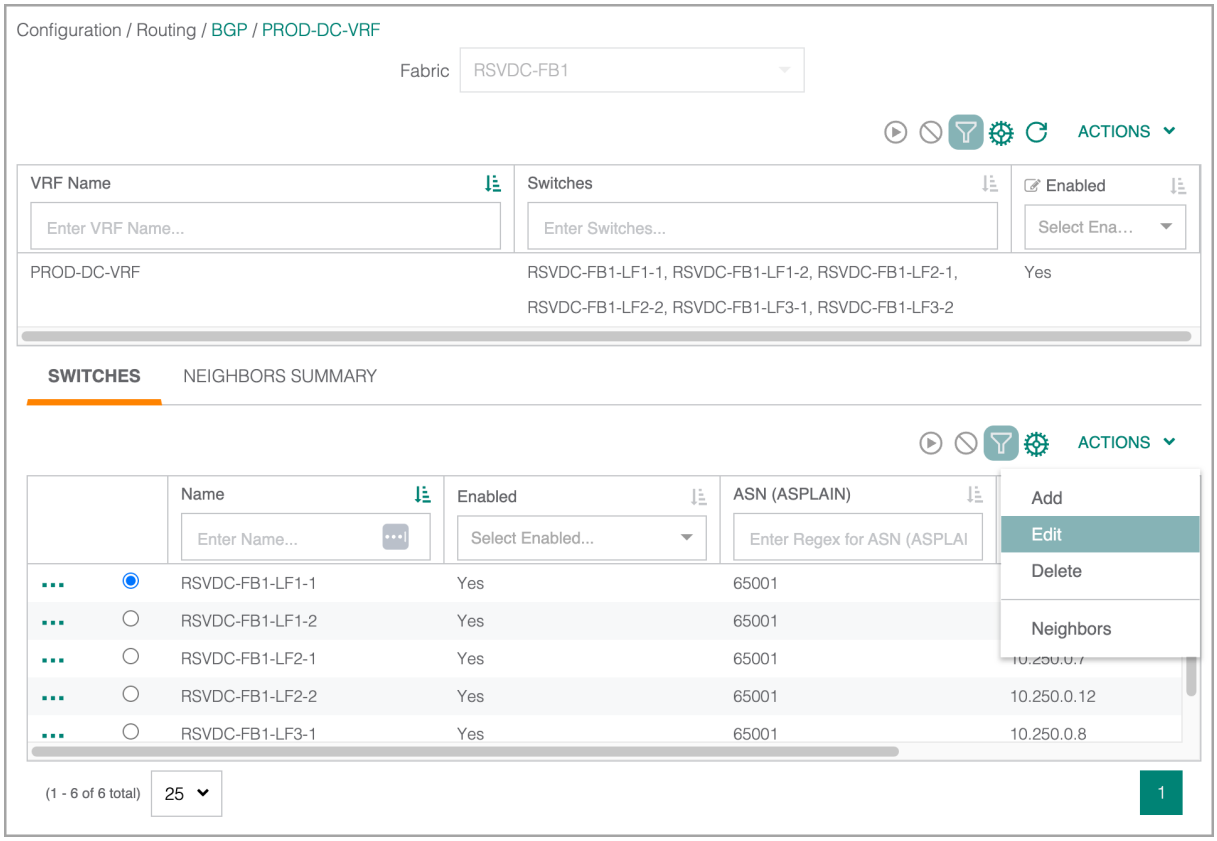


**Step 4** Click the **REDISTRIBUTE ROUTES** tab, click the **Redistribute Static Routes** option, and click **APPLY**.

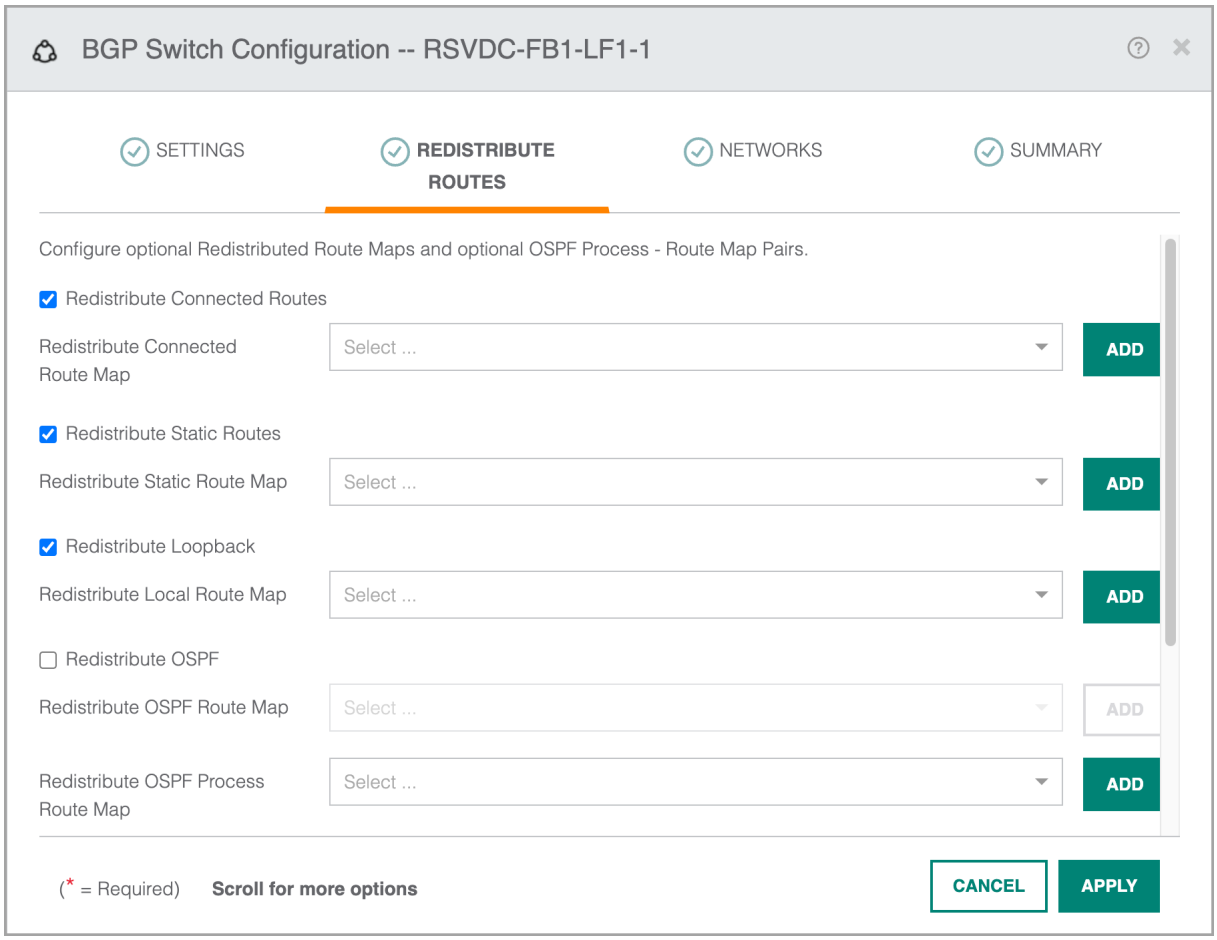


If the BGP configuration is updated, a **Success** message is presented.


**Step 5** Repeat steps 3 to 4 for the second border leaf to redistribute the static route on the second border leaf.

**Step 6** Click **BGP** in the path at the top of the main configuration window and repeat steps 2 to 5 for each overlay VRF.


## Verify Overlay Test Loopback

On each leaf switch, ping a directly connected host in the overlay using the loopback interface in the host's VRF as a ping source to verify the overlay test loopback is working.

```
ping 10.5.101.121 vrf PROD-DC-VRF source loopback11
```

```
RSVDC-FB1-LF2-1# ping 10.5.101.121 vrf PROD-DC-VRF source loopback11
PING 10.5.101.121 (10.5.101.121) from 10.250.4.3 : 100(128) bytes of data.
108 bytes from 10.5.101.121: icmp_seq=1 ttl=63 time=0.497 ms
108 bytes from 10.5.101.121: icmp_seq=2 ttl=63 time=0.450 ms
108 bytes from 10.5.101.121: icmp_seq=3 ttl=63 time=0.434 ms
108 bytes from 10.5.101.121: icmp_seq=4 ttl=63 time=0.406 ms
108 bytes from 10.5.101.121: icmp_seq=5 ttl=63 time=0.456 ms

--- 10.5.101.121 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4115ms
rtt min/avg/max/mdev = 0.406/0.448/0.497/0.029 ms
RSVDC-FB1-LF2-1#
```

# Configure Overlay IP Multicast

Protocol Independent Multicast–Sparse Mode (PIM-SM) is configured to build multicast route state between VTEPs within the data center and to external networks. PIM-SM is required for both sources and listeners in the data center. Internet Group Management Protocol (IGMP) manages known multicast listener state on data center leaf switches. IGMP snooping is configured to optimize Layer 2 forwarding of multicast traffic to only ports with interested listeners on leaf and server access switches.

In this guide, the PIM-SM rendezvous point (RP) is located outside the data center fabric in the campus network. The RP is learned with PIM-SM's Bootstrap Router (BSR) mechanism.

Configuration of multicast can be done at the command line of the switch or using Fabric Composer's **Configuration Editor**. Enter the configuration in code blocks in the procedures below to enable multicast in the EVPN-VXLAN overlay. The code blocks may include existing configuration to set context and existing descriptions to assist the reader.

## Configure Overlay PIM Multicast

The configuration examples in each step should be applied to all leaf switches, except where noted that configuration is only applied to border leaf switches.

**Step 1** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Configuration Editor**.

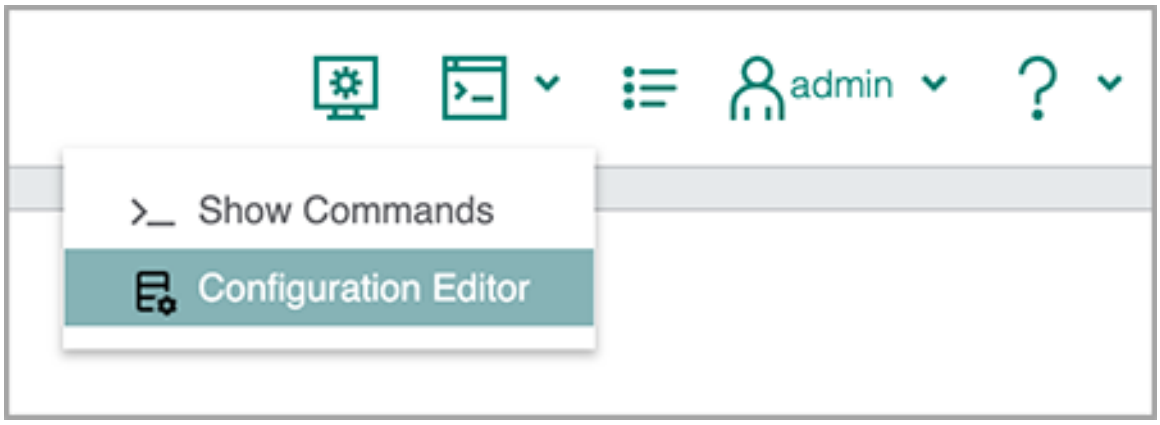


**Step 2** In the **Switch** field, select the border leaf and all other leaf switches.

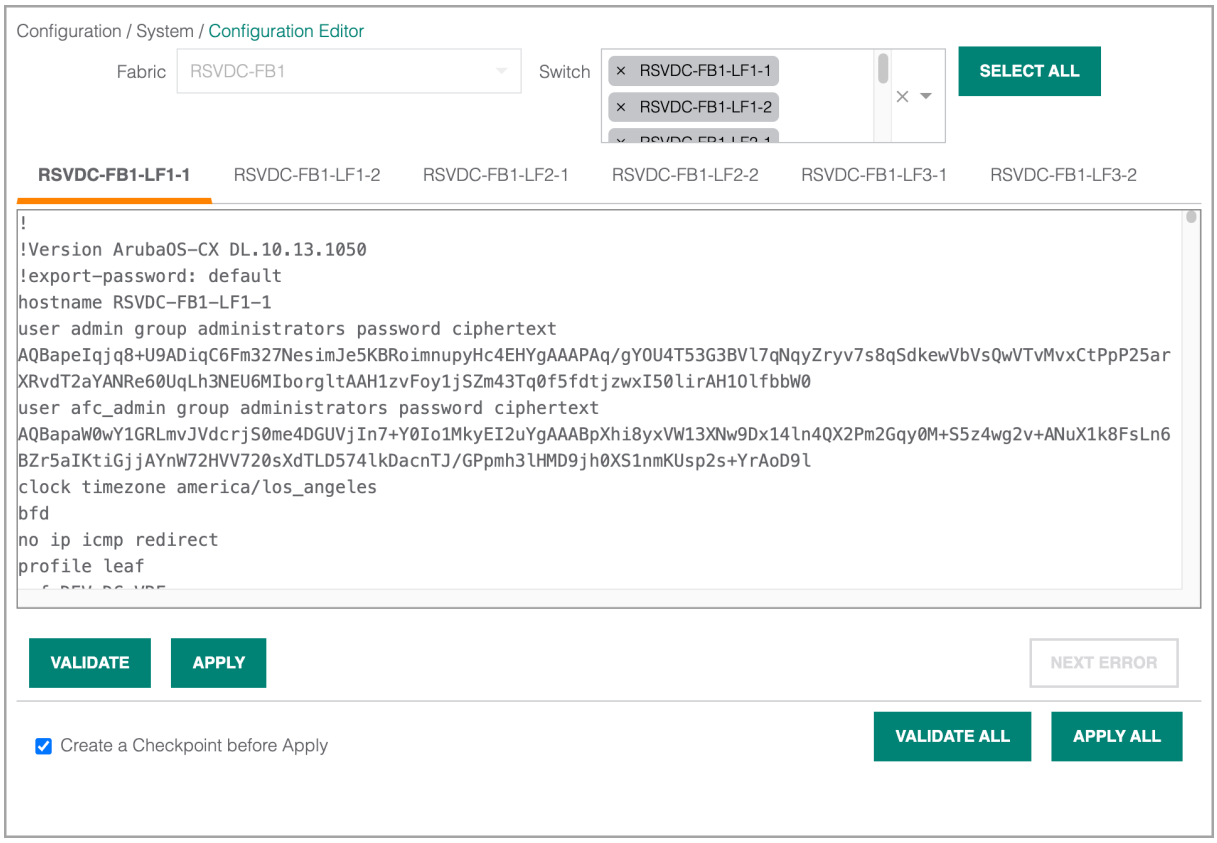


**NOTE:**

The spine switches do not contain multicast configuration for the overlay. Layer 2 server access switches are configured in a separate procedure in this guide. If there are few spine and server access switches, click **SELECT ALL** to select all switches in the fabric and deselect spine and server access switches.

**Step 3** Enable PIM routing in each overlay VRF.

```
router pim vrf PROD-DC-VRF
    enable
    register-source loopback11
router pim vrf DEV-DC-VRF
    enable
    register-source loopback12
```

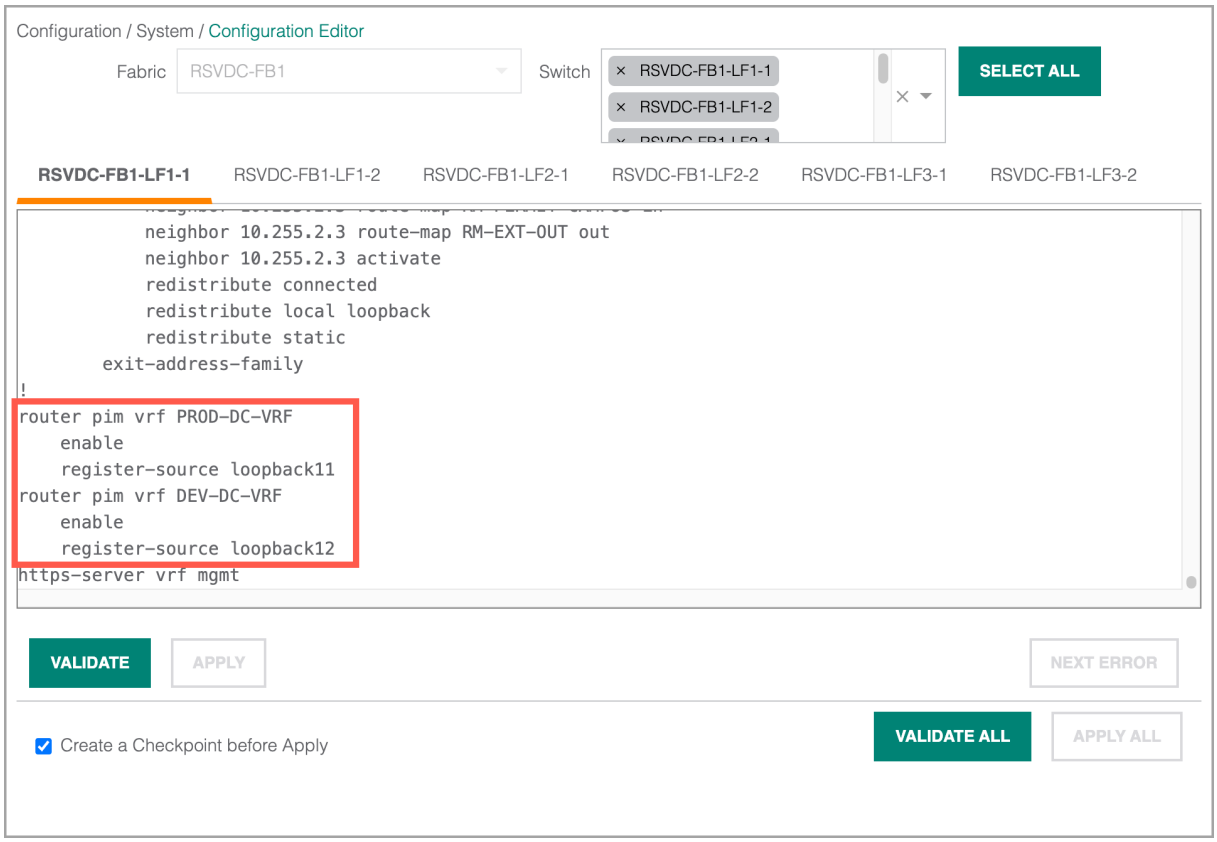


**NOTE:**

In the configuration above, the **register-source** command instructs PIM to send register messages to the RP using the unique overlay loopback IP configured in the previous **Assign Unique Overlay Loopbacks** procedure. This is required to send register-stop messages originated by the RP to reach the individual leaf switch originating PIM register messages.When entering configuration in Fabric Composer's **Configuration Editor**, new configuration must be entered below any references to other configuration elements. In this example, VRF names and loopback interfaces are referenced in the PIM router configuration, which requires placing the new PIM router config after those elements are defined in the existing configuration.

**Step 4** Enable PIM for each unique overlay loopback interface by adding **ip pim-sparse enable** in each configuration stanza. Do not modify existing configuration.

```
interface loopback 11
    ip pim-sparse enable
interface loopback 12
    ip pim-sparse enable
```

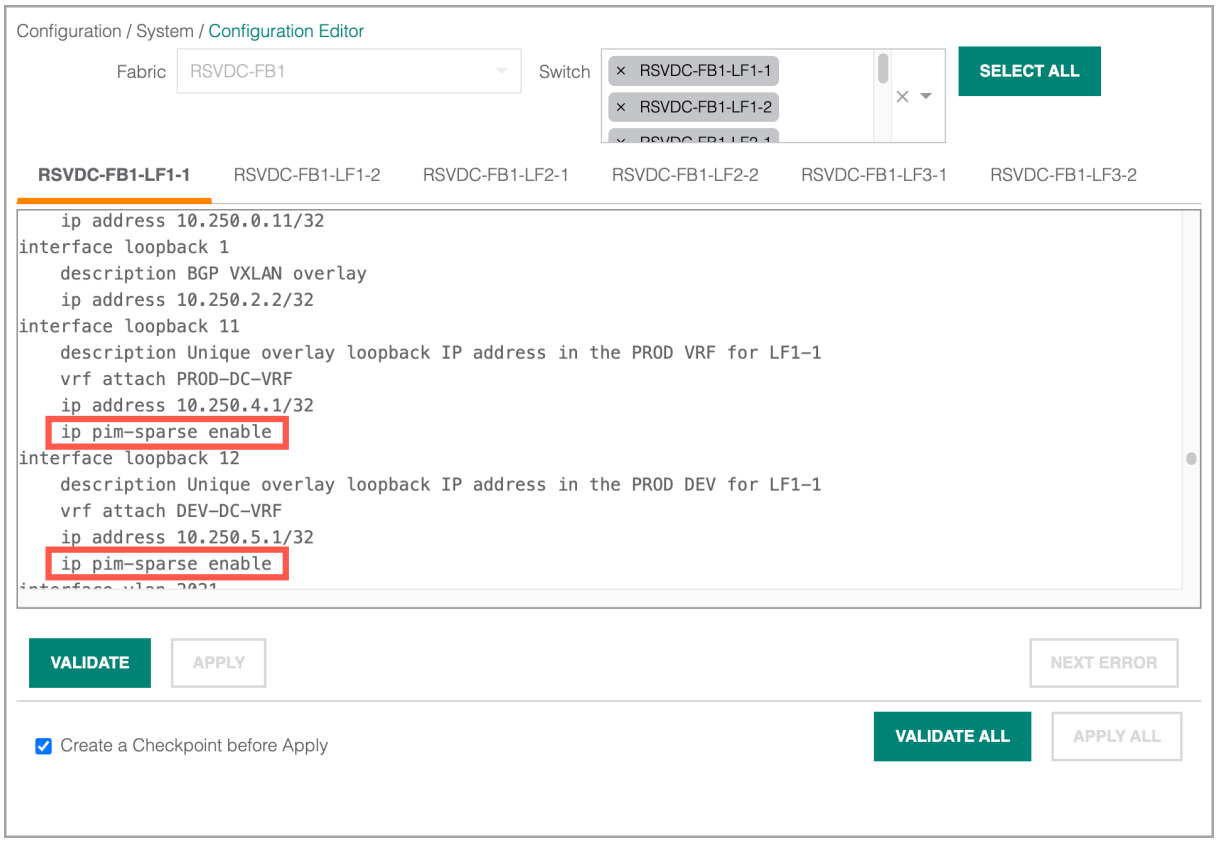


**Step 5** Enable PIM on overlay VLAN SVIs. This includes adding PIM to all data center host VLANs and overlay transit VLANs. On the border leaf, this includes campus routed interfaces.

On the border leaf, add the following configuration lines:

```
interface vlan2021
    description Border leaf PROD-DC-VRF uplink to external FW cluster
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan2022
    description Border leaf DEV-DC-VRF uplink to external FW cluster
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan3001
    description Overlay transit VLAN for PROD VRF
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan3002
    description Overlay transit VLAN for DEV VRF
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
```

On all other leaf switches, add the following configuration lines:

```
interface vlan101
    description Production web app SVI/VLAN 101 in DC overlay
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan102
    description Production database SVI/VLAN 102 DC overlay
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan201
    description Development web app SVI/VLAN 201 in DC overlay
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan202
    description Development database  SVI/VLAN 202 in DC overlay
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan3001
    description Overlay transit VLAN for PROD VRF
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
interface vlan3002
    description Overlay transit VLAN for DEV VRF
    ip pim-sparse enable
    ip pim-sparse vsx-virtual-neighbor
```

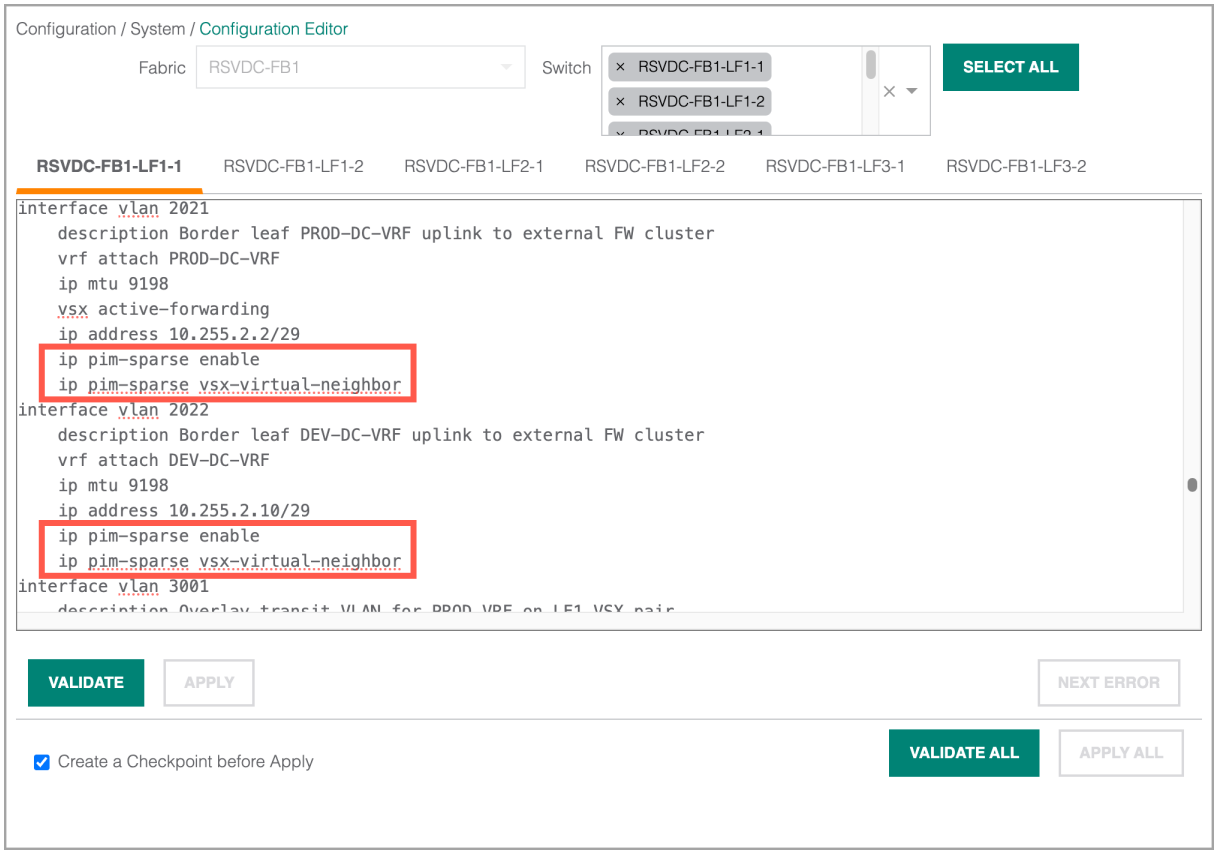


**NOTE:**

The border leaf switches in our example are dedicated to the border leaf function and do not include overlay host VLANs. When overlay VLANs are present on border leaf switches, configure PIM on those VLAN interfaces.

**Step 6** Click **VALIDATE ALL**.

**Step 7** Click **APPLY ALL**.

## Verify Overlay PIM

**Step 1** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

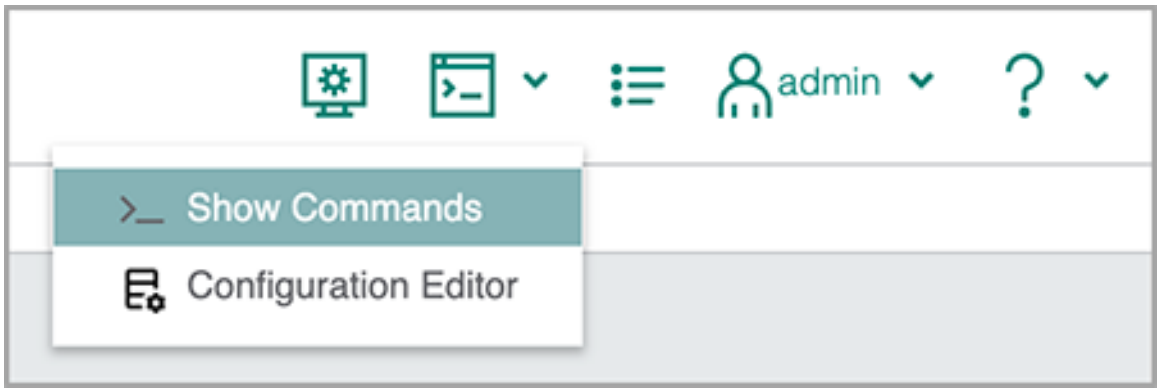


**Step 2** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select all leaf switches* >
- **Commands:** *show ip pim neighbor brief all-vrfs*

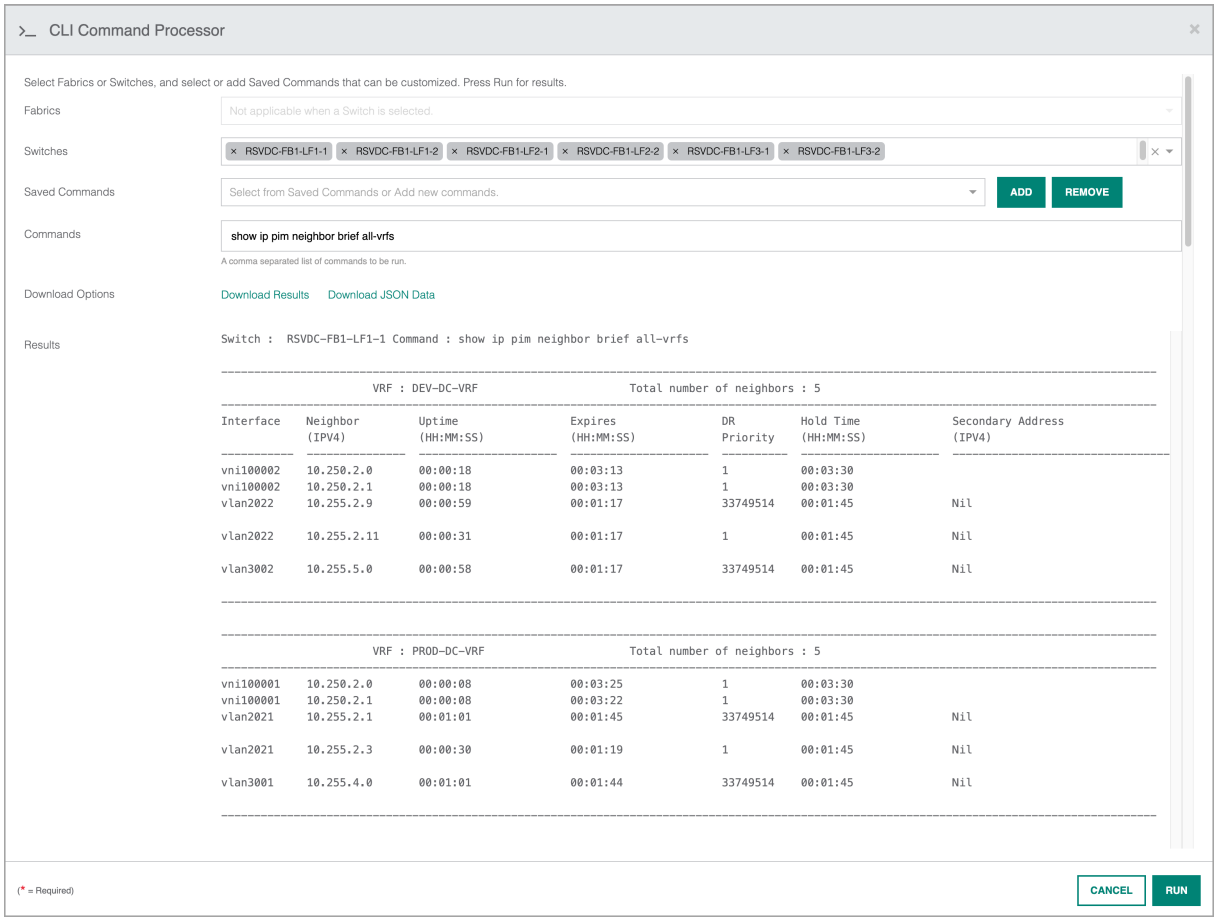

**Step 3** Review the output to verify that the following PIM neighbor adjacencies are established:

- Each VRF logical L3 VNI interface has a PIM neighbor relationship with each other VTEP in the EVPN-VXLAN fabric.
- Each VRF overlay transit VLAN has a PIM neighbor adjacency.
- Each host facing VLAN has a PIM neighbor adjacency on all VSX redundant leaf switches.
- Two PIM adjacencies are formed on the border leaf VLAN that supports external routed connectivity. One adjacency is with the peer VSX switch and the second is with the external firewall.

**NOTE:**

Overlay PIM adjacencies formed between logical L3 VNI interfaces take longer to establish than PIM adjacencies between switches. It may take a minute for the logical adjacencies to form.The RP in the fabric is learned from PIM BSR messages received by the border leaf switches from the external network.

**Step 4** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select all leaf switches* >
- **Commands:** *show ip pim rp-set all-vrfs*

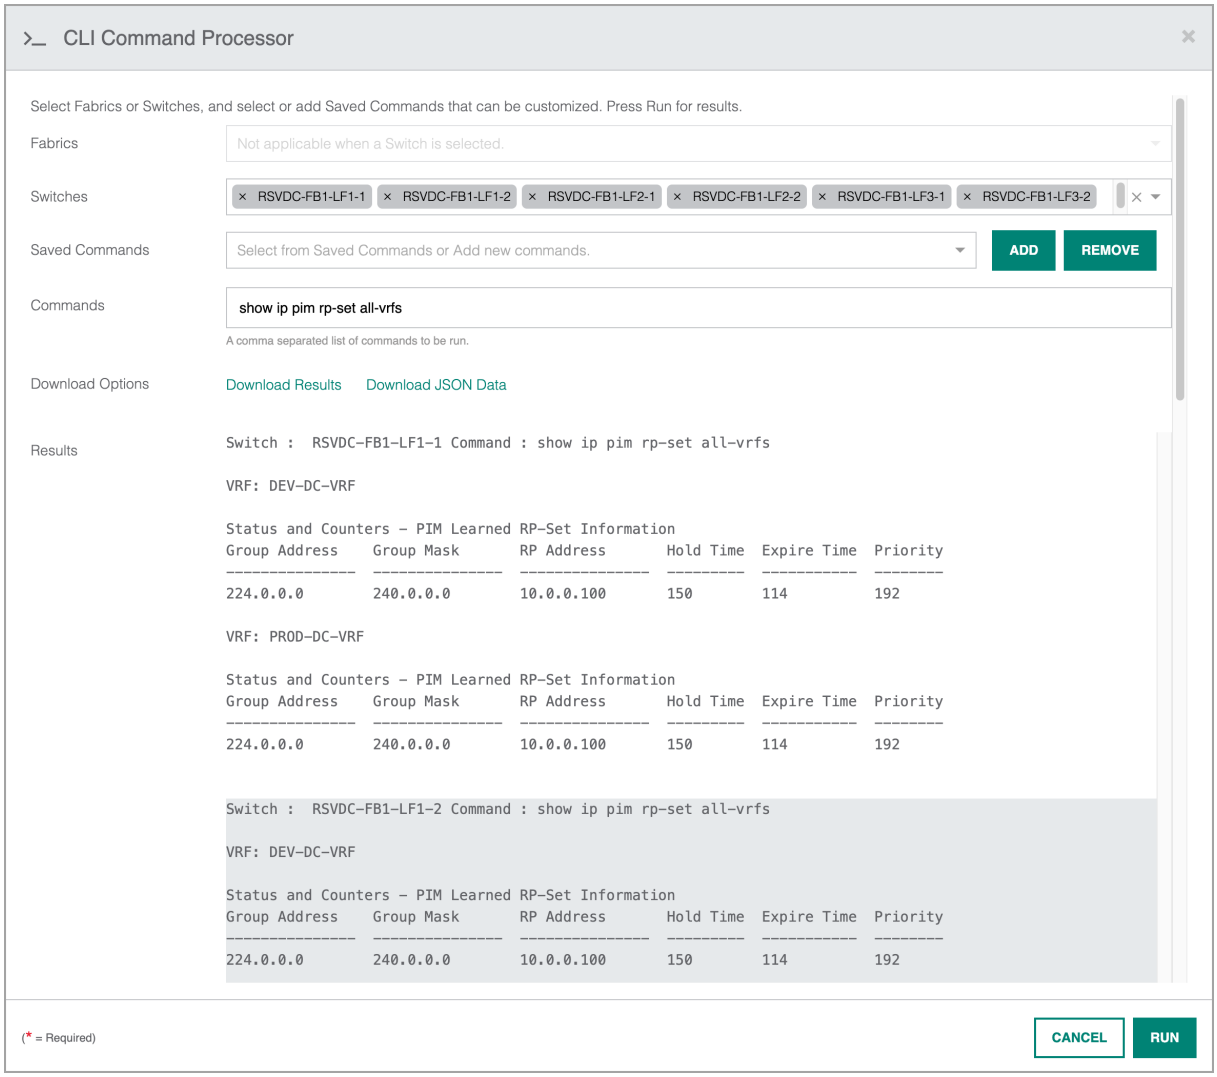

**Step 5** Review the output to verify that the campus RP is learned in all overlay VRFs on all leaf switches.

## Redistribute Local SVI into EVPN

Redistribute local SVI interfaces into EVPN instances on leaf switches to distribute system-MAC values. This ensures proper distribution of IGMP querier information throughout the fabric.

**Step 1** On the **Configuration** menu, select **Routing > EVPN**.

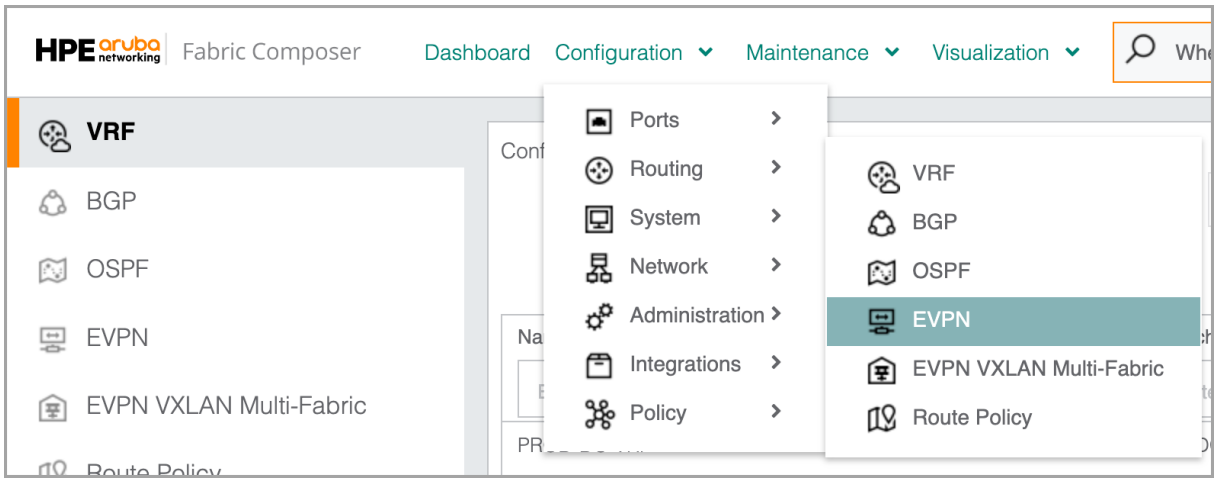


**Step 2** On the right **ACTIONS** menu, select **Settings**.

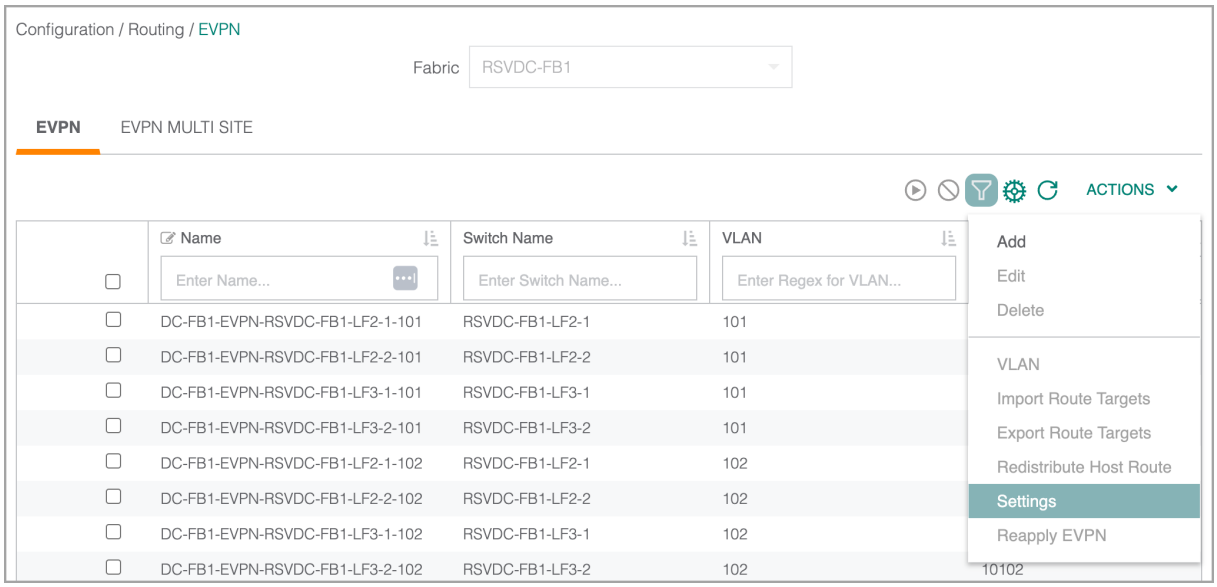


**Step 3** On the **EVPN Settings Page**, assign the following values and click **APPLY**.

- **Enable ARP Suppression**: < *checked* >
- **Redistribute Local MAC Address**: < *unchecked* >
- **Redistribute Local SVI**: < *checked* >
- **Apply the EVPN Settings across the entire Fabric and all Switches contained within it**: < *unchecked* >
- **Switches**: < *Select leaf switches containing EVPN mapped overlay VLANs* >

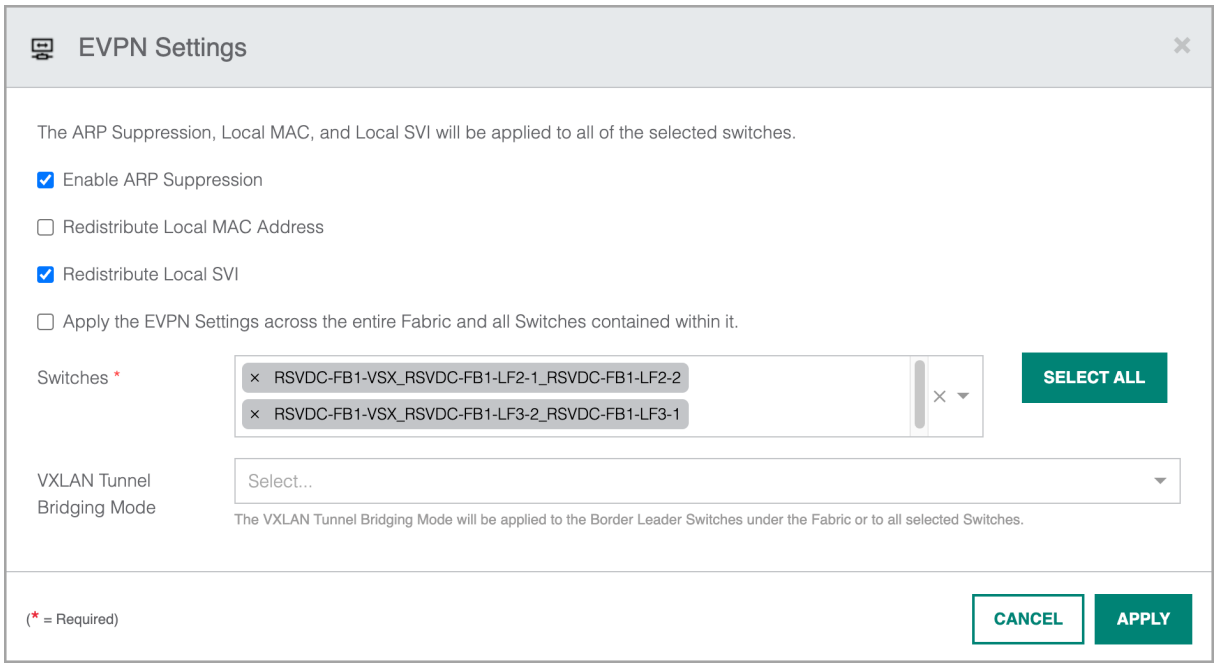


**NOTE:**

ARP suppression was enabled when originally creating the EVPN instance. Checking **Enable ARP Suppression** when changing EVPN settings is required, because existing settings will be overwritten using the values specified on the **EVPN Settings** page after clicking **APPLY**. When routed multicast is not performed in the overlay, select **Redistribute Local MAC Address** to enable system-MAC propagation.

## Configure Overlay IGMP and IGMP Snooping

IGMP is configured on all leaf switches, and IGMP snooping is configured on both leaf switches and server access switches.

**Step 1** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Configuration Editor**.

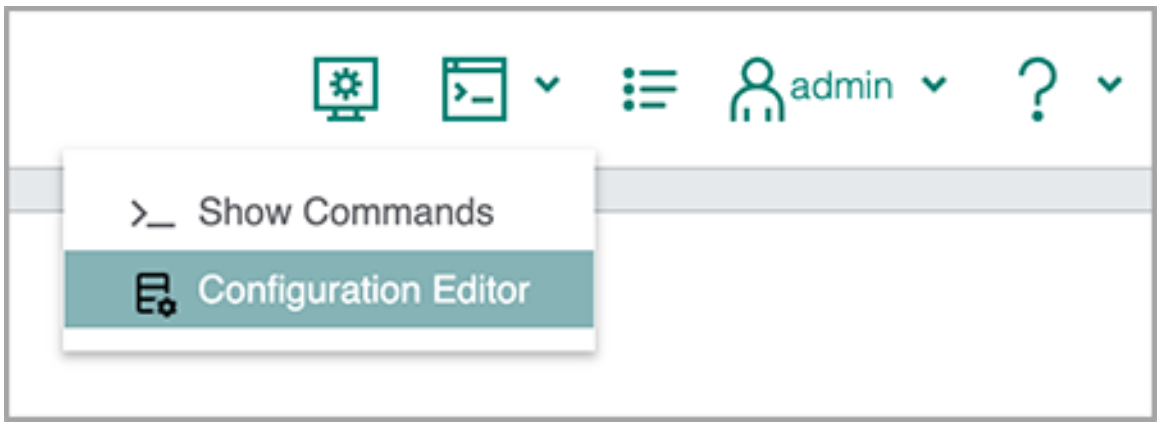


**Step 2** In the **Switch** field, select all other leaf switches with data center overlay VLANs and server access switches

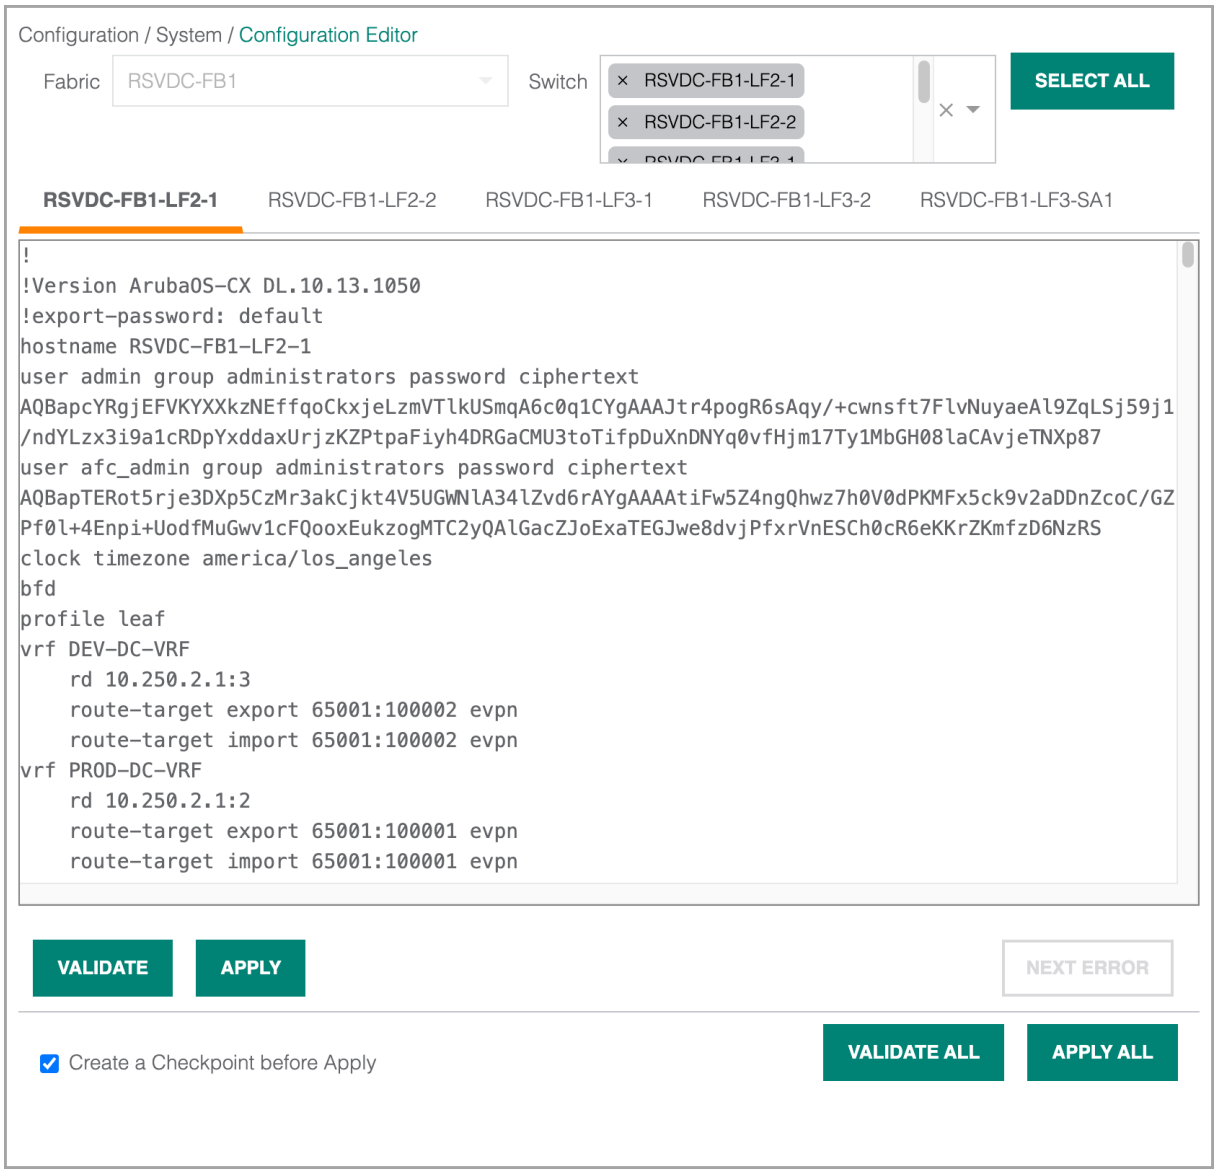


**NOTE:**

The border leaf switches in our example are dedicated to the border leaf function and do not include overlay host VLANs. Configure IGMP on border leaf switches, when overlay host VLANs are present.

**Step 3** Enable IGMP on all overlay VLAN interfaces on each leaf switch.

```
interface vlan 101
    ip igmp enable
interface vlan 102
    ip igmp enable
interface vlan 201
    ip igmp enable
interface vlan 202
    ip igmp enable
```

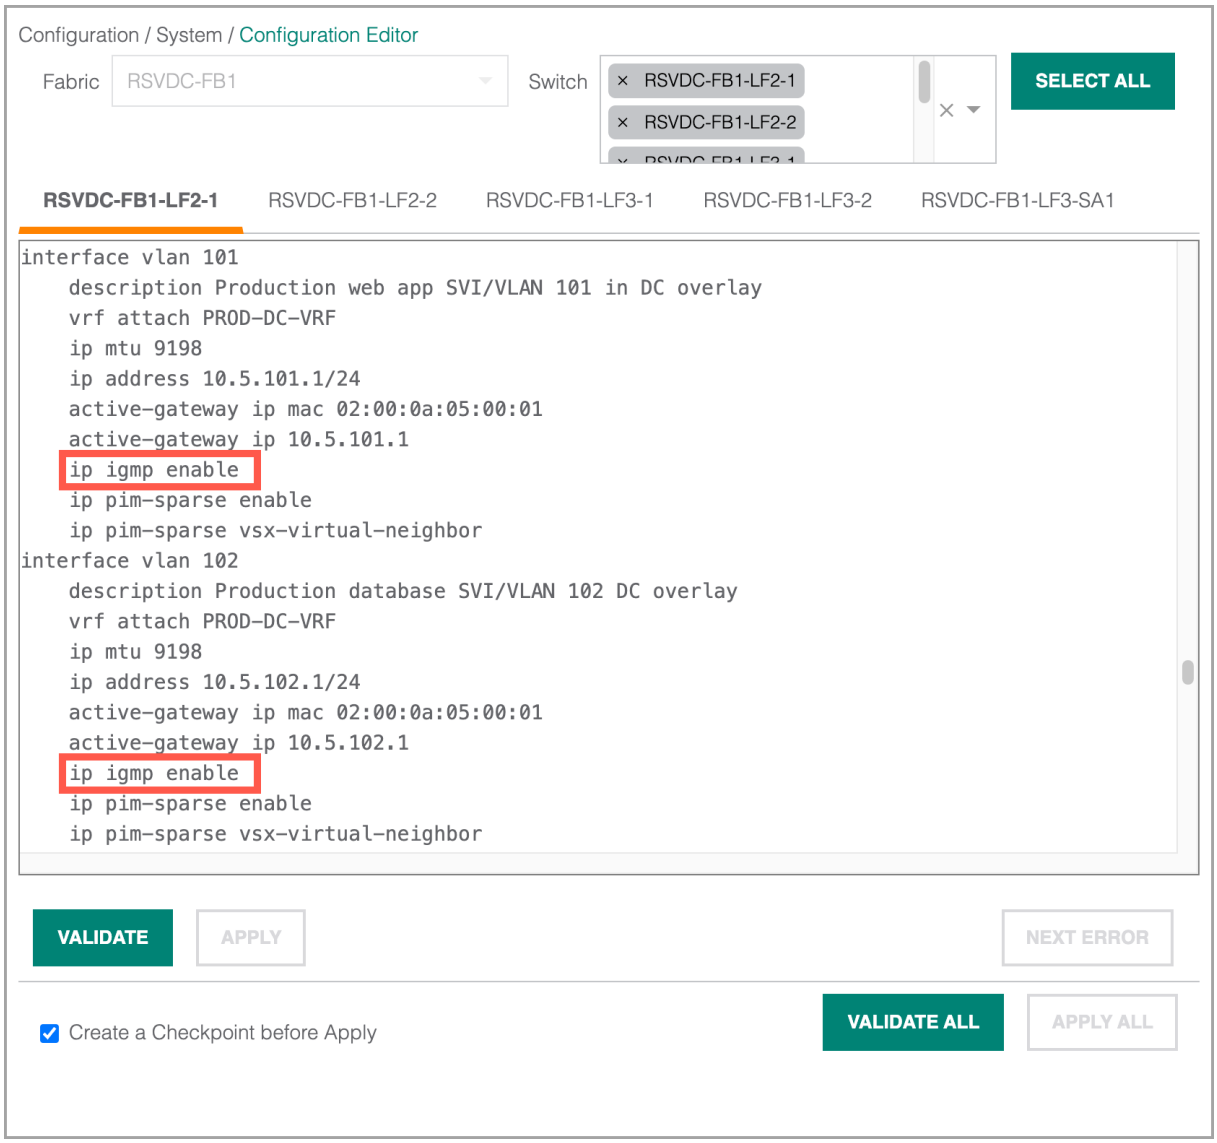


**Step 4** Enable IGMP snooping on all overlay VLANs for both leaf switches and server access switches.

```
vlan 101
    ip igmp snooping enable
vlan 102
    ip igmp snooping enable
vlan 201
    ip igmp snooping enable
vlan 202
    ip igmp snooping enable
```

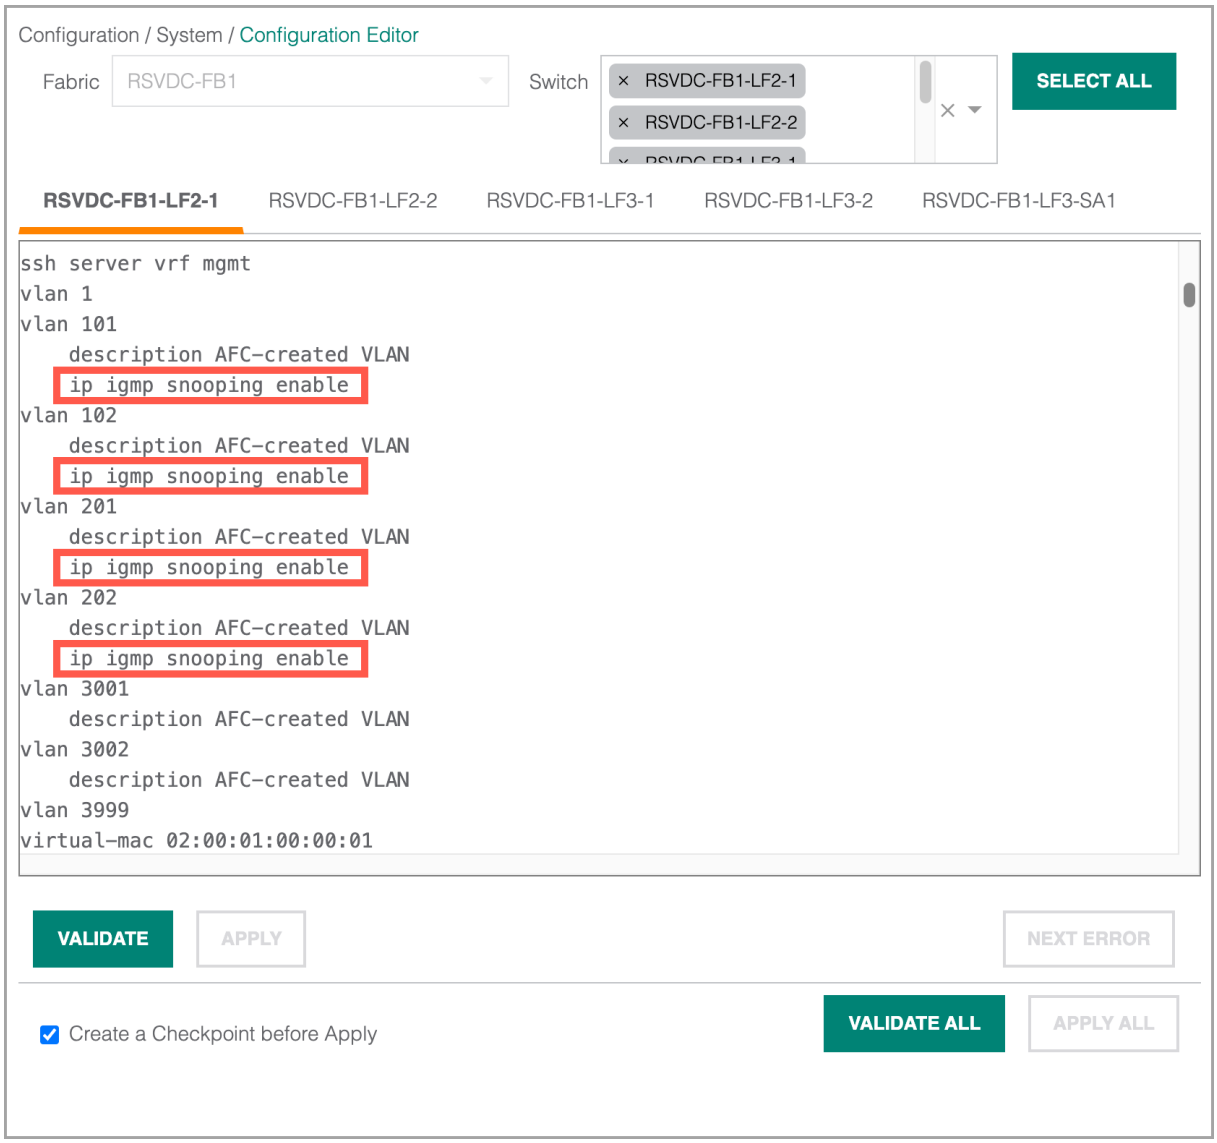


**Step 5** Click **VALIDATE ALL**.

**Step 6** Click **APPLY ALL**.

## Verify Overlay IGMP and IGMP Snooping

On all leaf and server access switches, start a multicast listener for a multicast group with an active source, then use the follow procedure to verify IGMP and IGMP snooping optimizations.

**Step 1** On the menu bar at the top right of the Fabric Composer window, click the **CLI Commands** icon and select **Show Commands**.

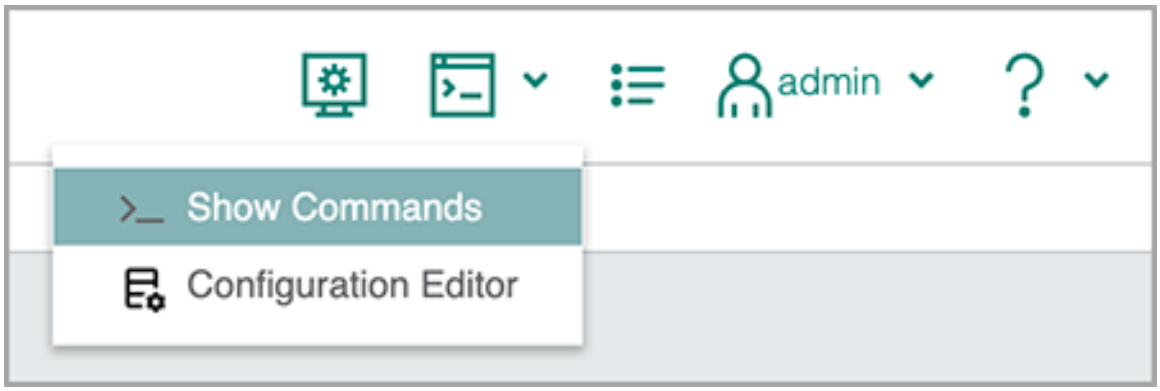


**Step 2** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** *< Select all leaf switches with overlay VLANs >*
- **Commands:** *show ip igmp groups all-vrfs*

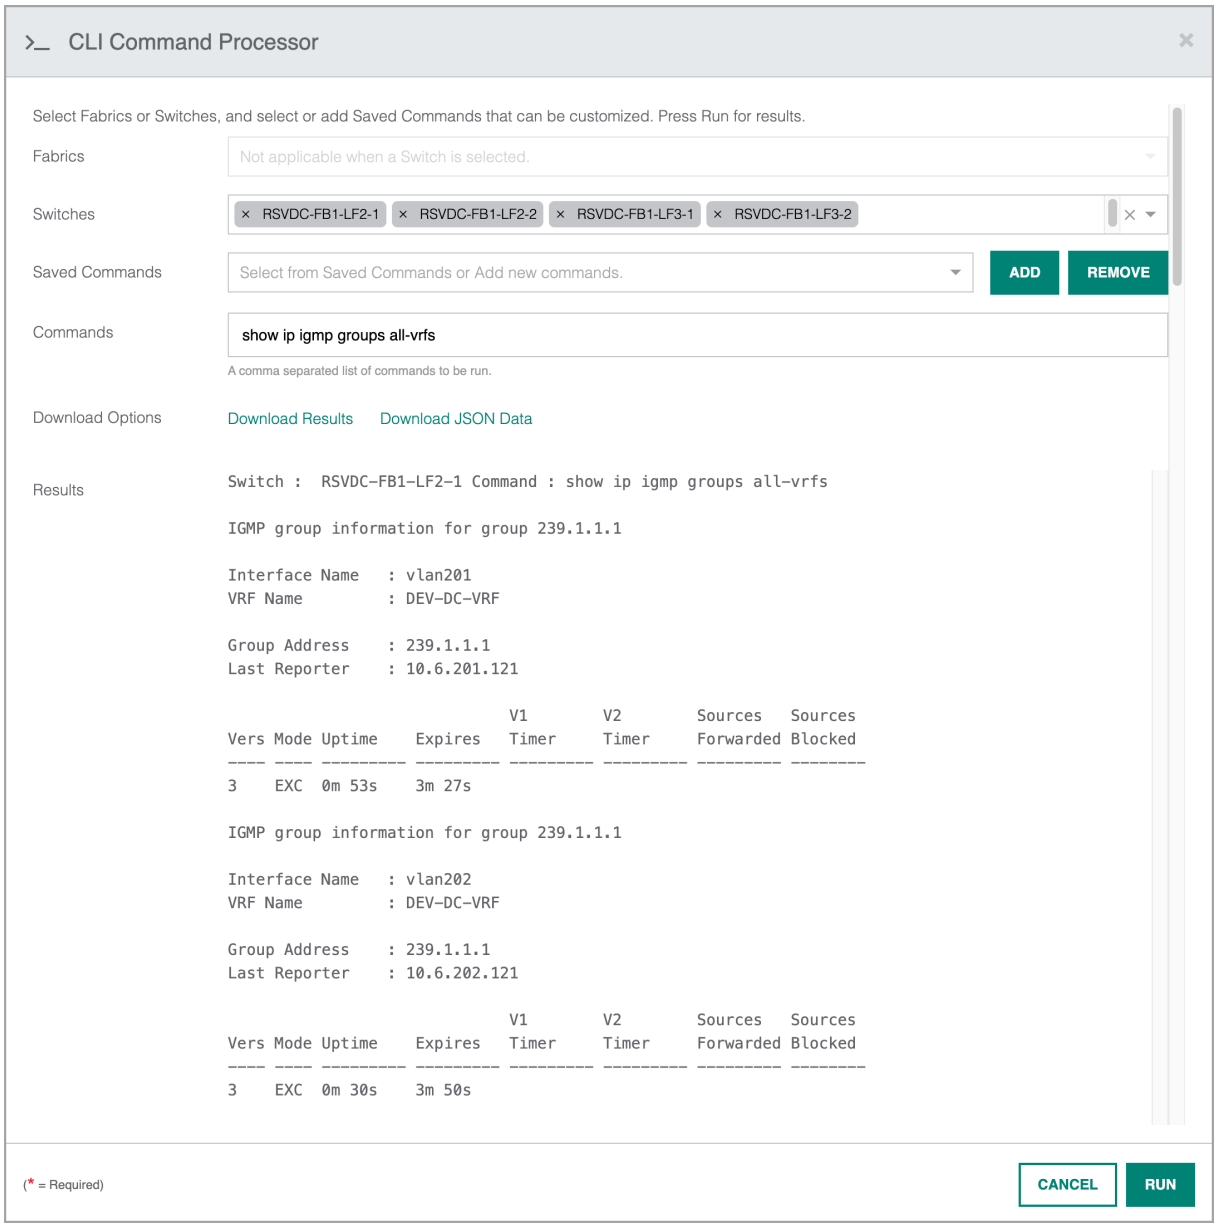

**Step 3** Verify that the multicast group is learned on each switch on the VLAN corresponding with the attached listening hosts.

**Step 4** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** *< Select all leaf switches with overlay VLANs and server access switches >*
- **Commands:** *show ip igmp snooping groups*

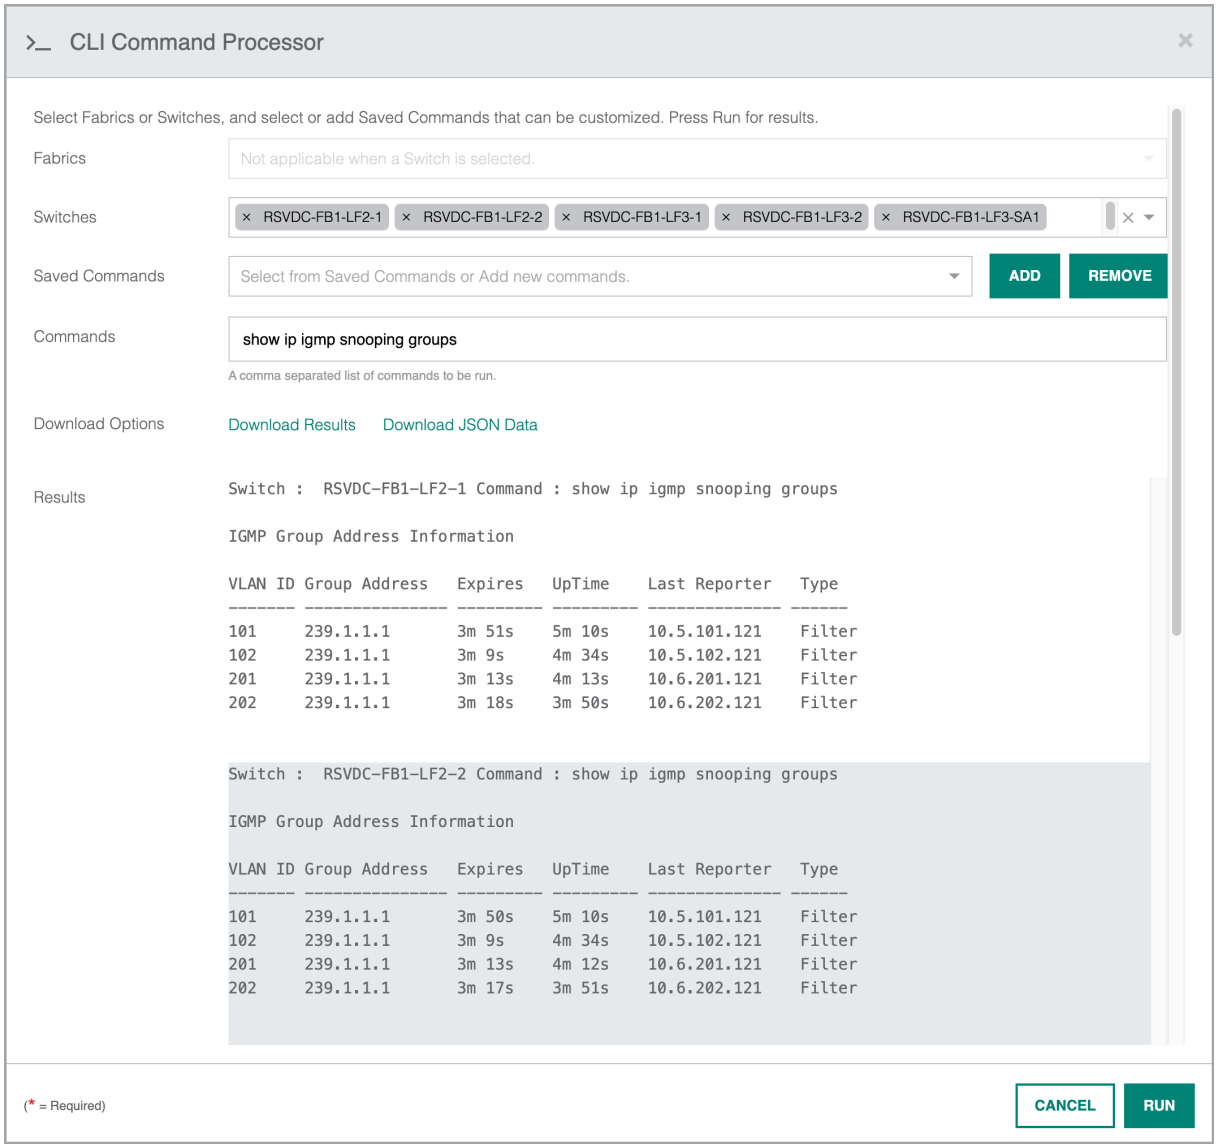


**Step 5** Verify that that IGMP snooping has state for all VLANs with a listener.

## VMWare vSphere Integration

VMware vSphere integration enables VMware host and virtual machine visualization within Fabric Composer. This procedure also enables automated switch port provisioning of VLANs based on how the vSwitch and VMs are setup.

**Step 1** On the **Configuration** menu, select **Integrations > VMware vSphere**.

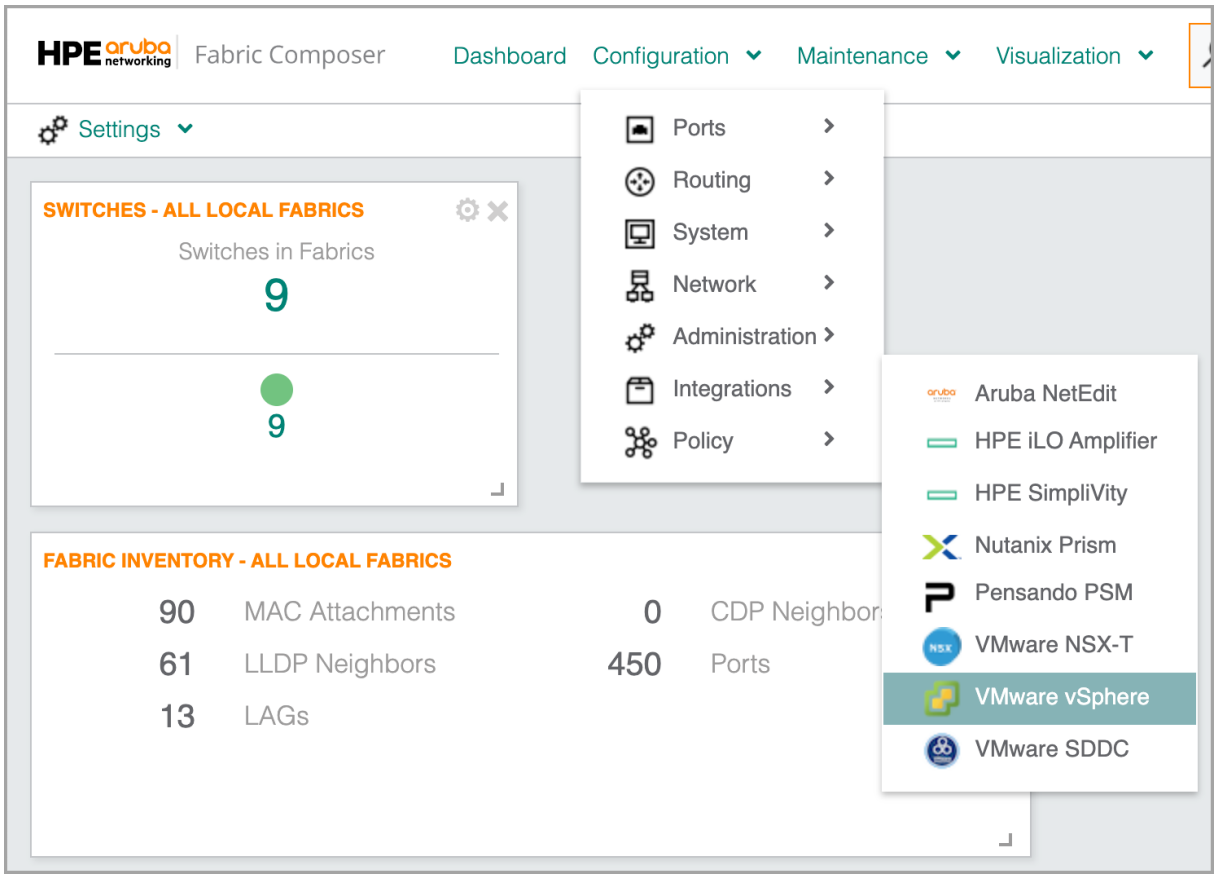


**Step 2** On the right **ACTIONS** menu, click **Add** to start the **VMware vSphere** wizard.

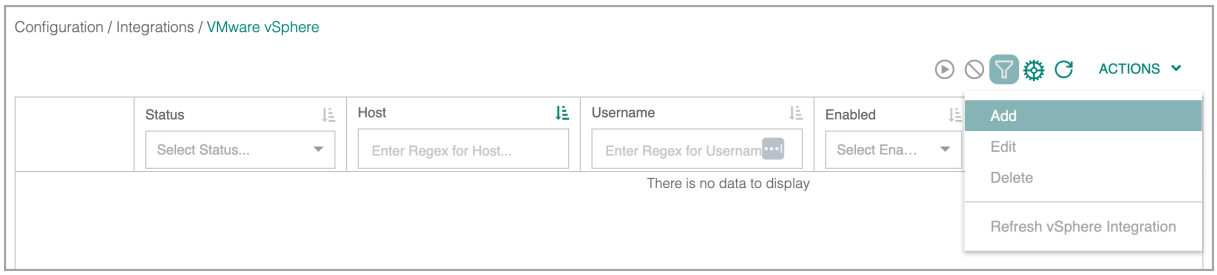


**Step 3** On the **Host** page, assign the following settings:

- **Name:** *Example-vSphere1*
- **Description:** *Example vSphere Integration*
- **Host:** *rsvdc-vcenter.example.local*
- **Username:** *administrator@example.local*
- **Password:** *< password >*
- **Validate SSL/TLS certificates for Aruba Fabric Composer:** *unchecked*
- **Enable this configuration:** *checkmark*

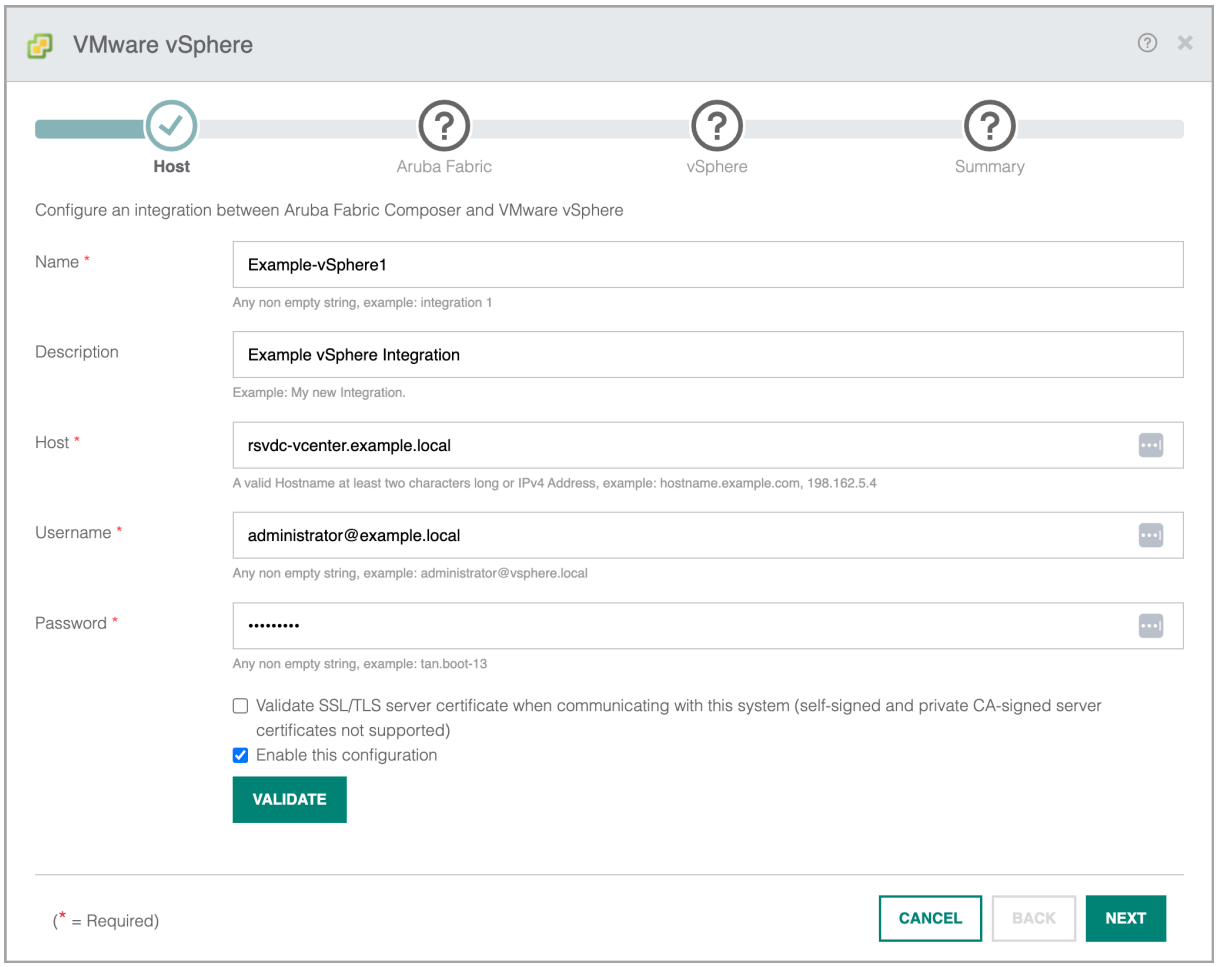


**NOTE:**

**Host** is the resolvable hostname or IP address of the vCenter server. **Username** is the name of an administrator account on the vCenter server. **Password** is the password for the administrator account on the vCenter server.

**Step 4** Click **VALIDATE** to verify that the provided credentials are correct. A green success message appears at the bottom right. Click **NEXT**.


**Step 5** On the **Aruba Fabric** page, choose from the two options below and enter a **VLAN Range**, check **Automated PVLAN provisioning for ESX hosts direction connected to the fabric** and enter a VLAN range. Check **Automated Endpoint Group Provisioning**, then click **NEXT**.

- If the hosts are directly connected from the NIC to the switch, select **Automated VLAN provisioning for ESX hosts directly connected to the fabric**.
- If host infrastructure is HPE Synergy or another chassis with an integrated switch solution, select **Automated VLAN provisioning for ESX hosts connected through intermediate switches**.

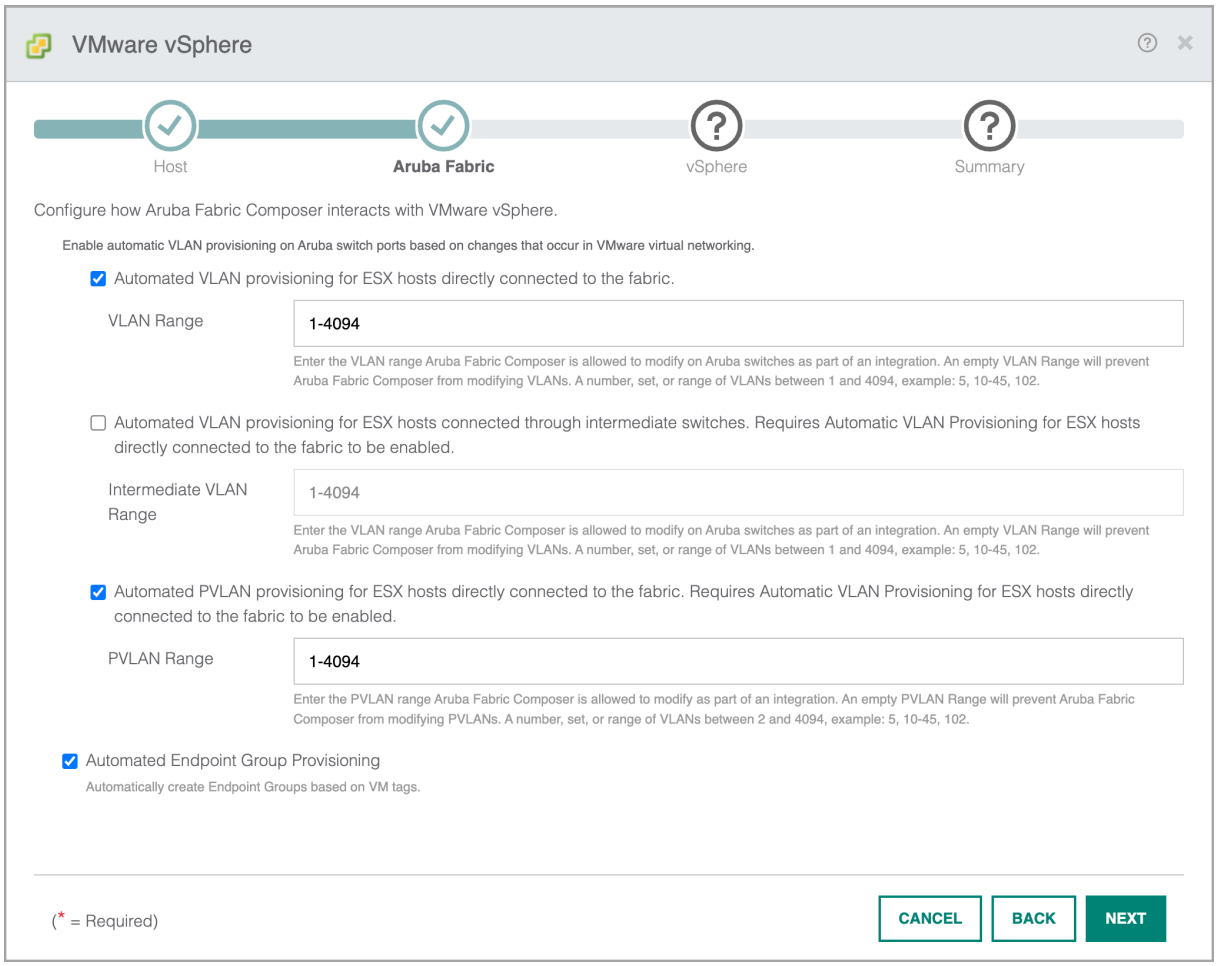


**NOTE:**

**Automated PVLAN provisioning for ESX hosts directly connected to the fabric** is a prerequisite for microsegmentation automations built into Fabric Composer. **Automated Endpoint Group Provisioning** enables assigning VMs dynamically to firewall policy using VM tags. The IP addresses used in the policy are modified dynamically in the future, if a VM IP changes or the VMs associated with the tag change. For additional details on all options, refer to the *HP Aruba Networking Fabric Composer User Guide*.

**Step 6** On the **vSphere** page, click the checkbox for **Discovery protocols** and click **NEXT**.

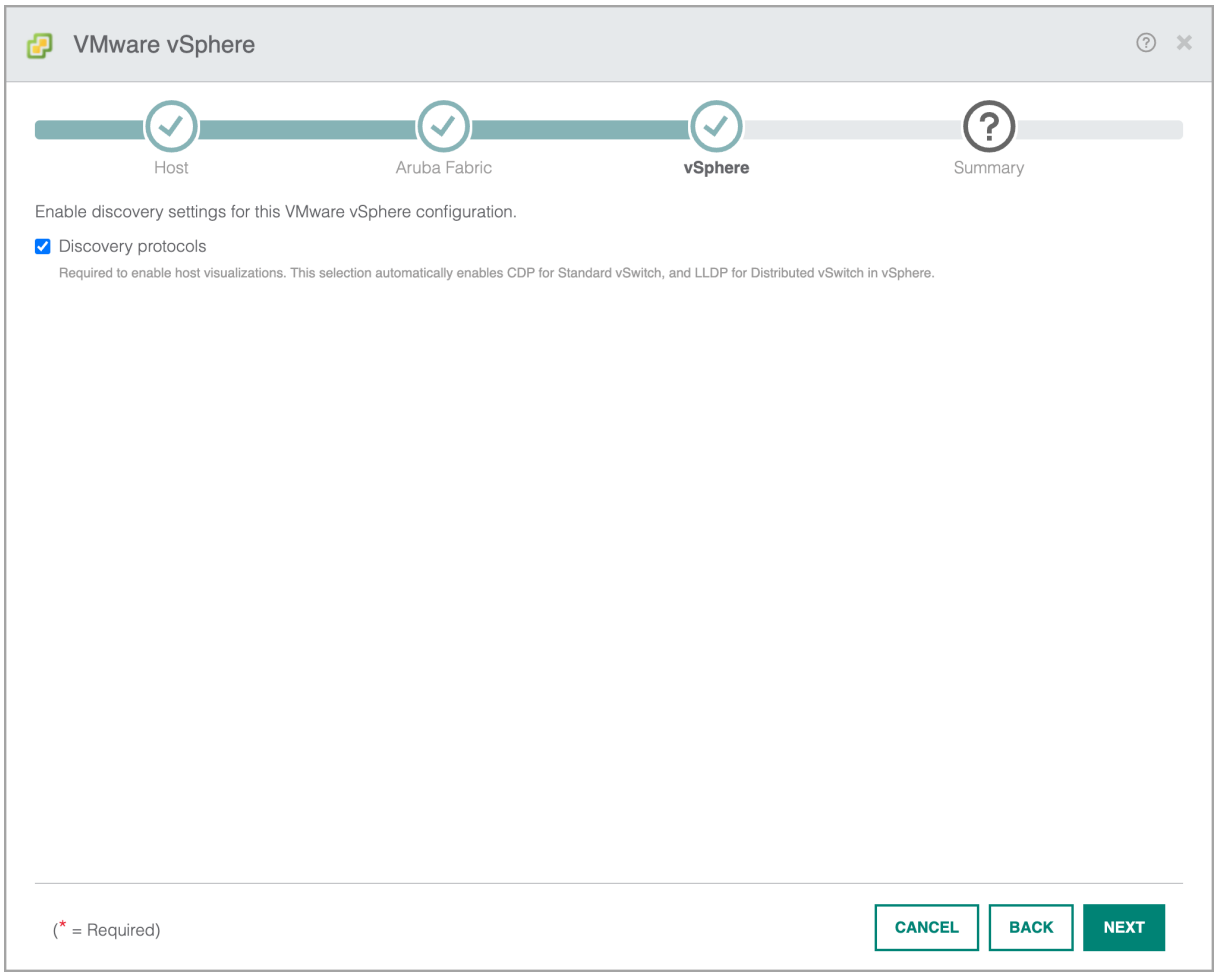


**CAUTION:**

If **Discovery protocols** is not enabled, the VMware integration cannot display virtual switches correctly.

**Step 7** On the **Summary** page, confirm that the information is entered correctly and click **APPLY**.

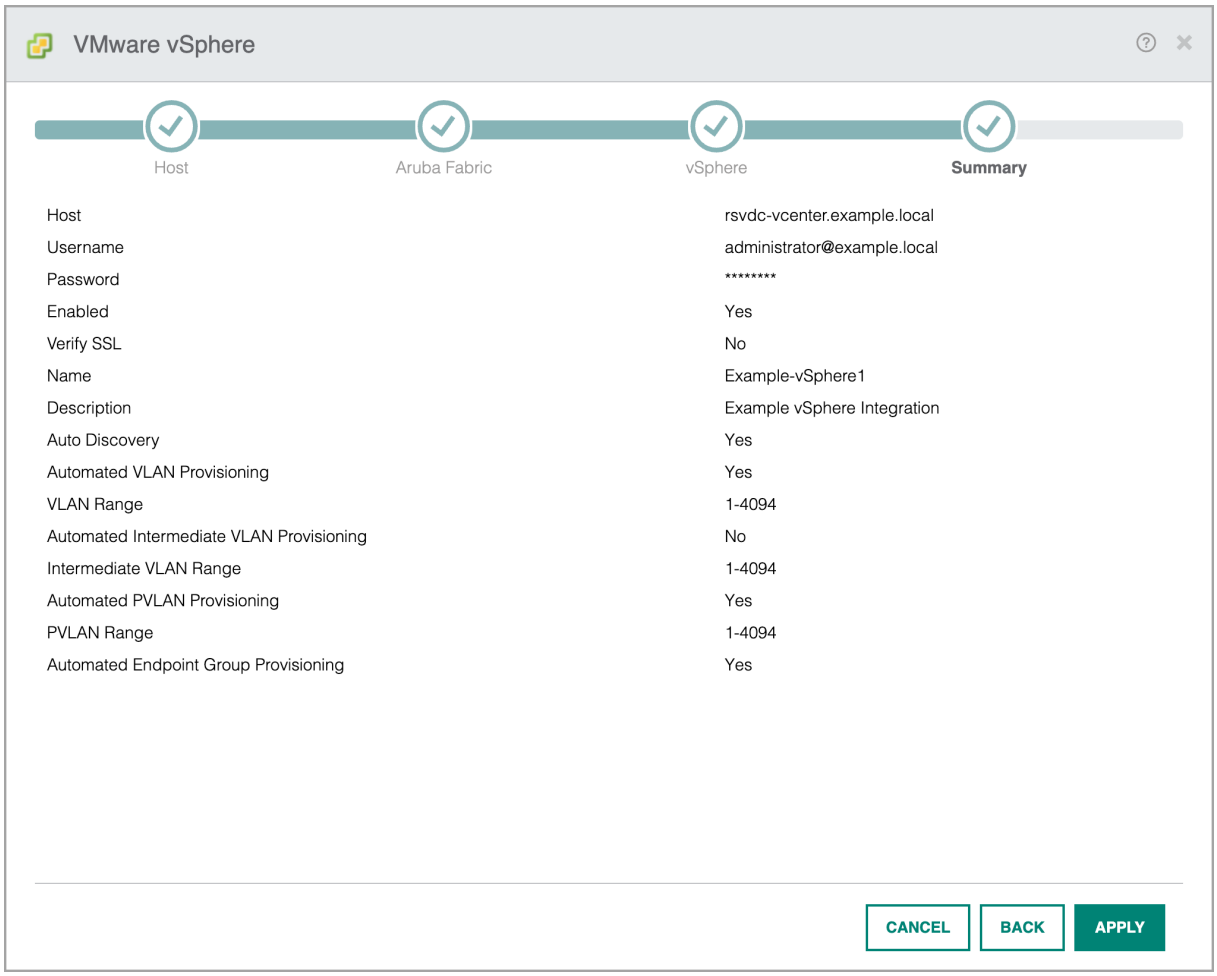


**Step 8** Go to **Visualization > Hosts**.

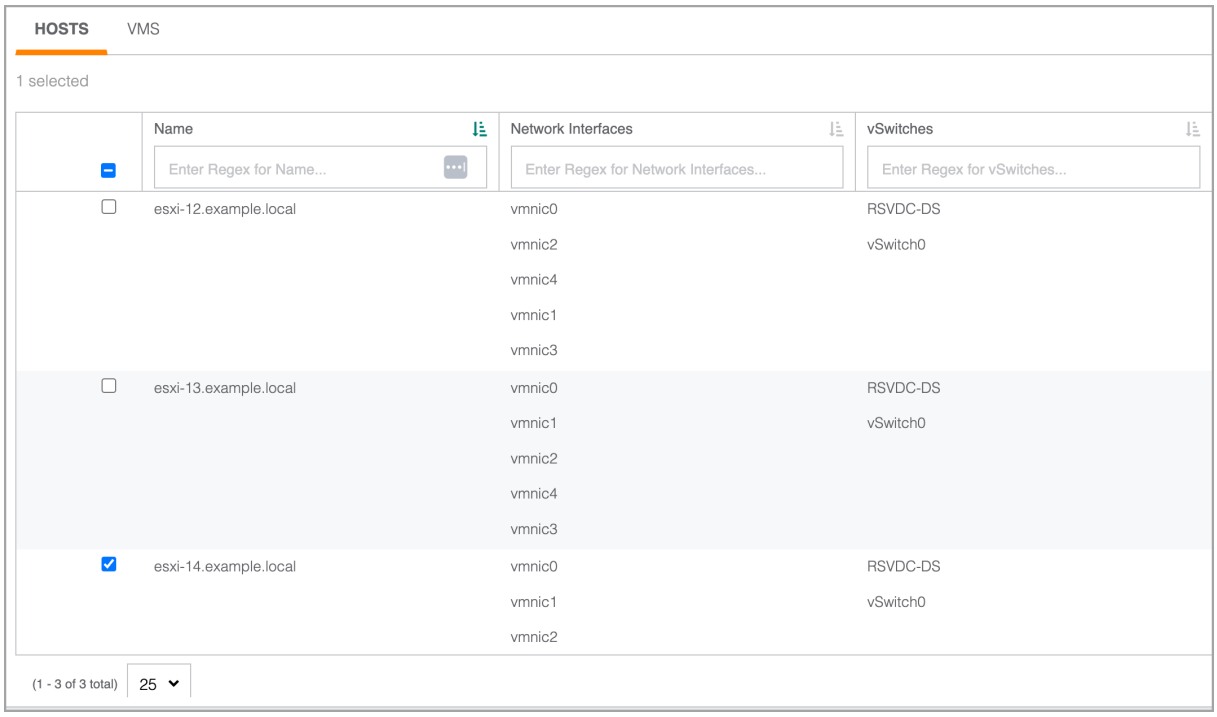


**Step 9** Select the checkbox next to the **Name** of an ESXi VM host to add it to the visualization window.

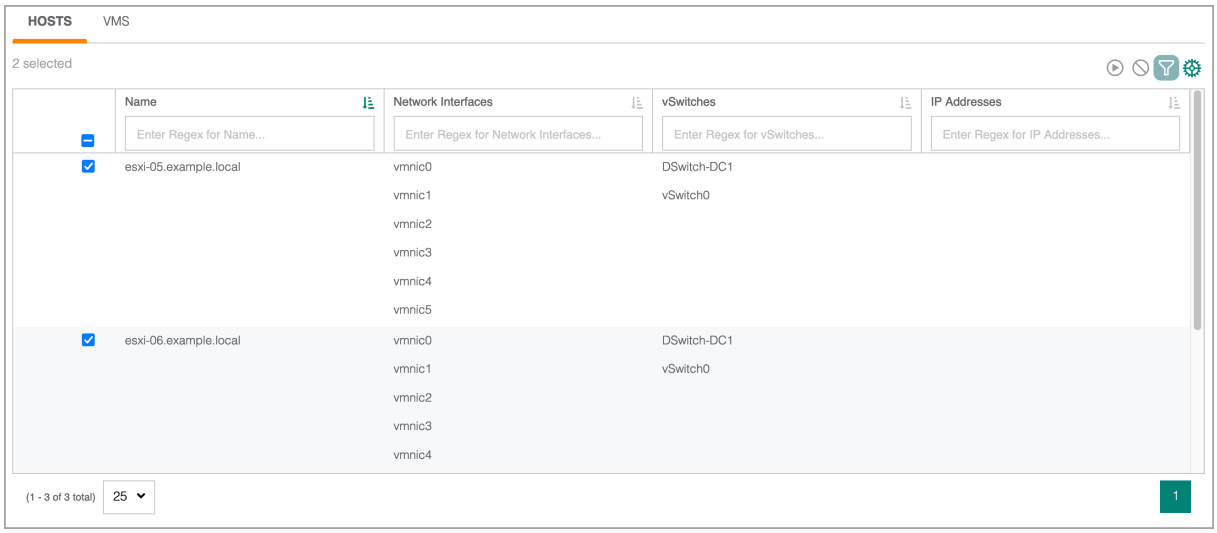


**Step 10** Verify the connectivity displayed from the hypervisor layer to the leaf switches.

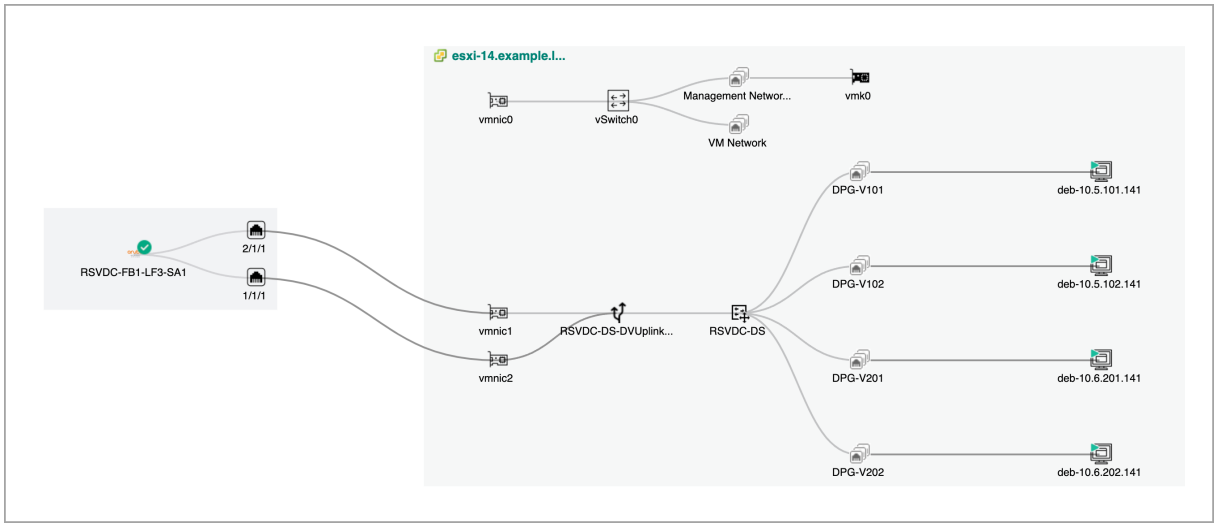

# Secure an Aruba EVPN Fabric

AFC orchestrated security policy is applied to east-west traffic using CX 10000 switches and AMD Pensando's Policy and Services Manager in this example fabric. ACL policy also can be applied to east-west and north-south traffic.

## Overview

This chapter describes enforcing policy in the Aruba ESP data center network using the Aruba CX 10000 switch platform's stateful firewall. Policy is implemented using the AMD Pensando Policy and Services Manager (PSM) orchestrated by Aruba Fabric Composer (AFC). Stateful firewall policy filters traffic between VLANs, between hosts in the same VLAN, and between VM guests assigned to the same hypervisor (microsegmentation).

The Aruba CX 10000 series switch is a data-center-class Distributed Services Switch (DSS). It includes hardware dedicated to performing stateful firewall functions on data center host traffic. Configuring and managing these capabilities require deployment of PSM. AFC is recommended to manage PSM. Refer to the Aruba ESP Data Center Design VSG and the *Pensando Policy and Services Manager for Aruba CX 10000: User Guide* for additional details.

**NOTE:**

Aruba CX 10000 switches running AOS-CX 10.13 and above require a feature pack license to enable firewall features. Information on the installation and purchase of feature packs can be found in the Feature Pack Ordering Guide and Deployment Guide.

## Configure PSM Integration with Aruba Fabric Composer

Use this procedure to associate a PSM cluster with a fabric in AFC to enable centralized management of firewall policy.

**Step 1** On the menu bar at the top right of the AFC user interface, select **Guided Setup**.

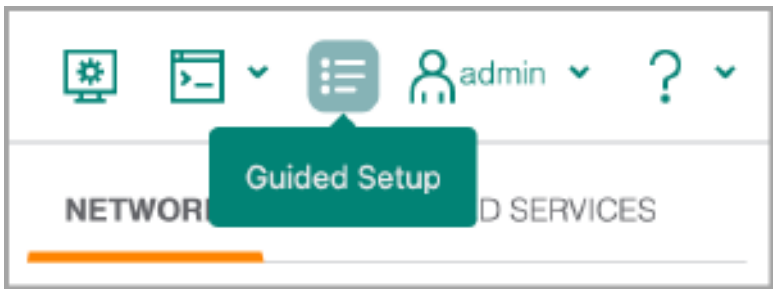


**Step 2** At the top of the **Guided Setup** window, click the **Distributed Services** tab, then click **PENSANDO PSM**.

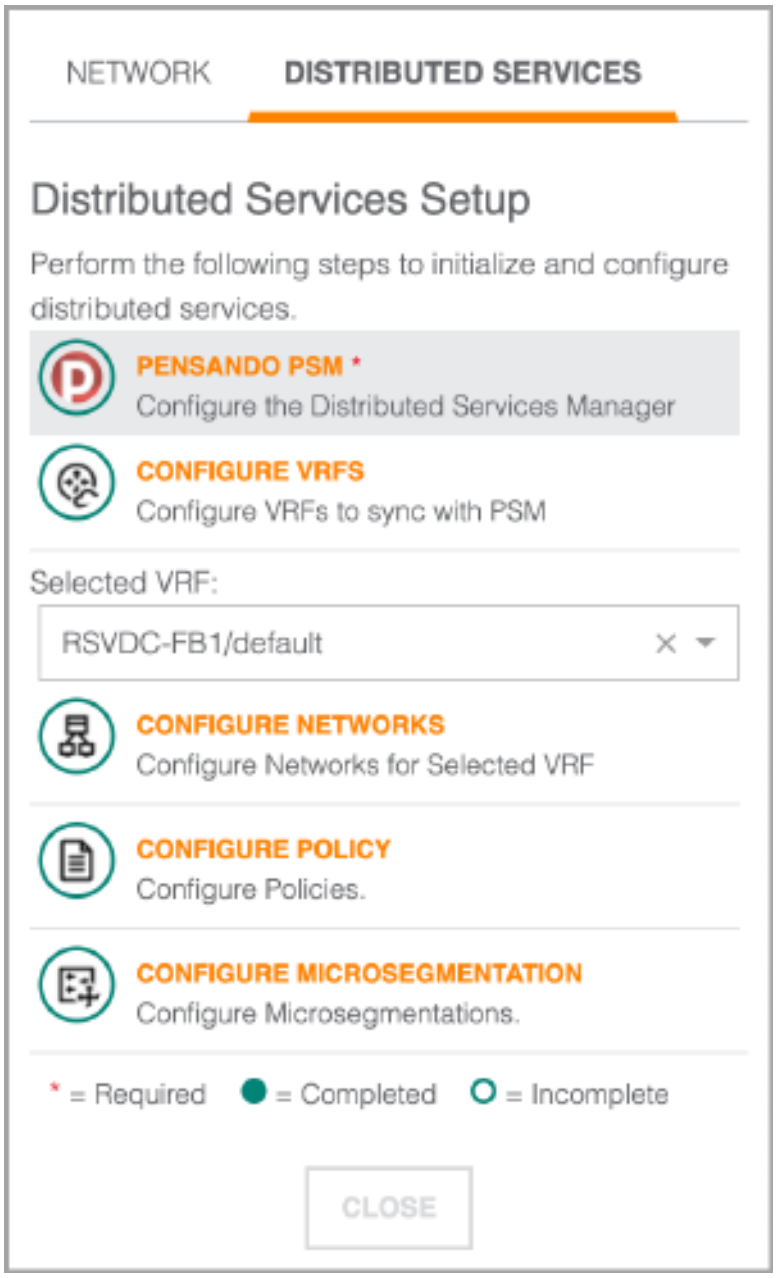


**Step 3** On the **Host** page, enter the following values and click **VALIDATE**.

- **Name:** *RSVDC-FB1-PSM*
- **Description:** *PSM cluster for RSVDC fabric 1*
- **Host:** *172.16.104.51*
- **Username:** *admin*
- **Password:** *< Password for the PSM admin user >*

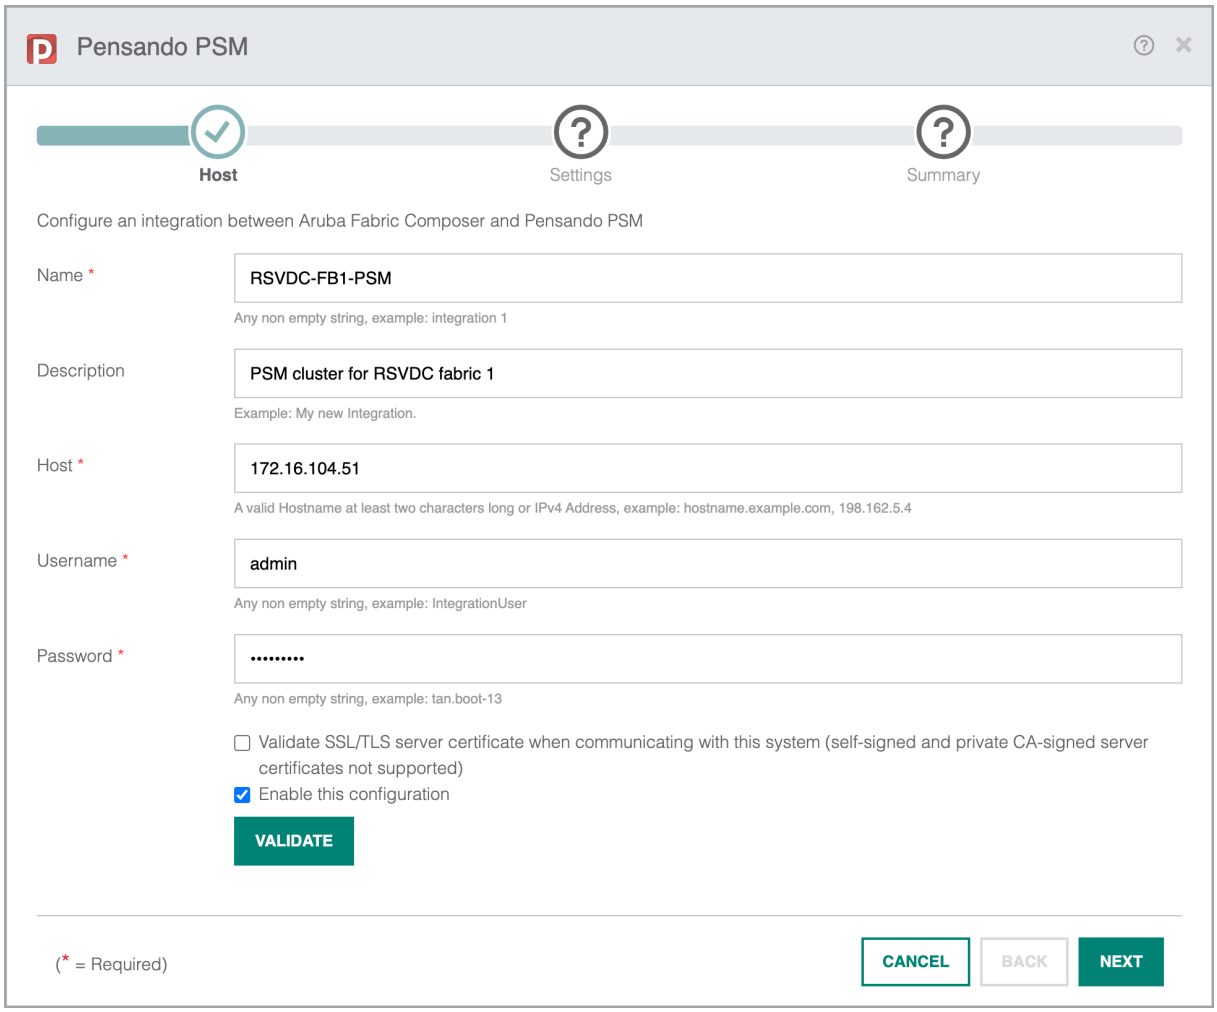


**NOTE:**

The **Host** field requires only a single DNS name or IP address of one PSM VM. The IP addresses of the remaining cluster members are discovered automatically.

**Step 4** After the validation success message appears, click **NEXT**.


**Step 5** On the **Settings** page, select the fabric. Leave other settings at their default and click **NEXT**.

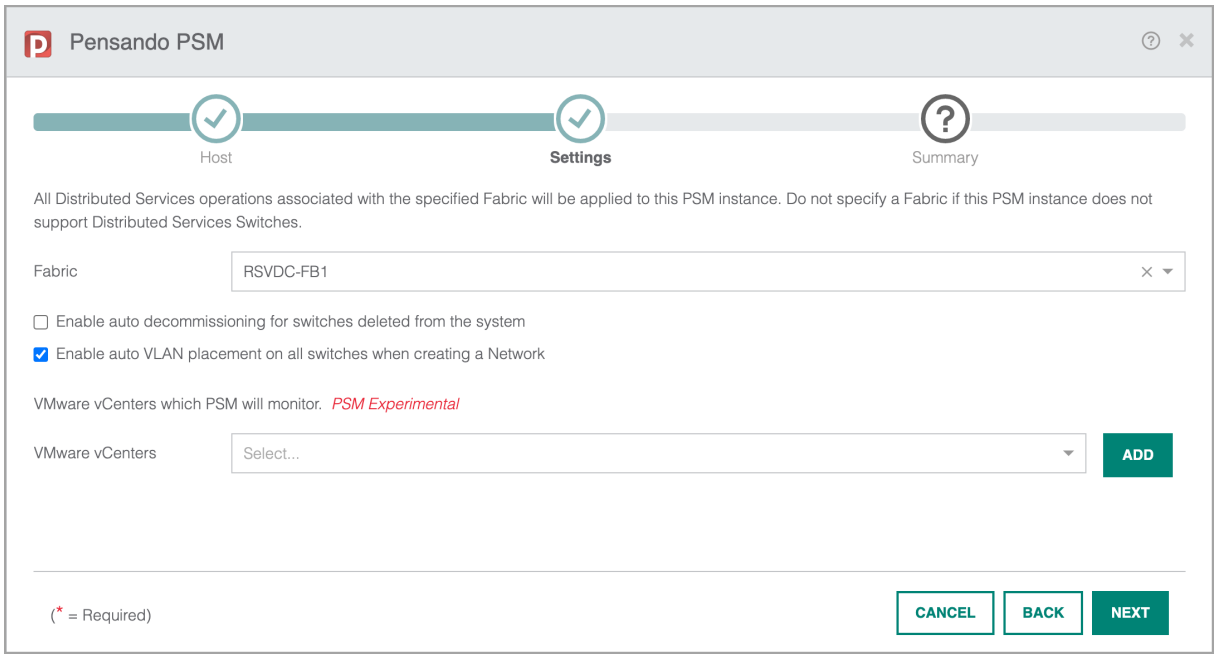


**Step 6** On the **Summary** page, verify that the PSM cluster settings are correct and click **APPLY**.

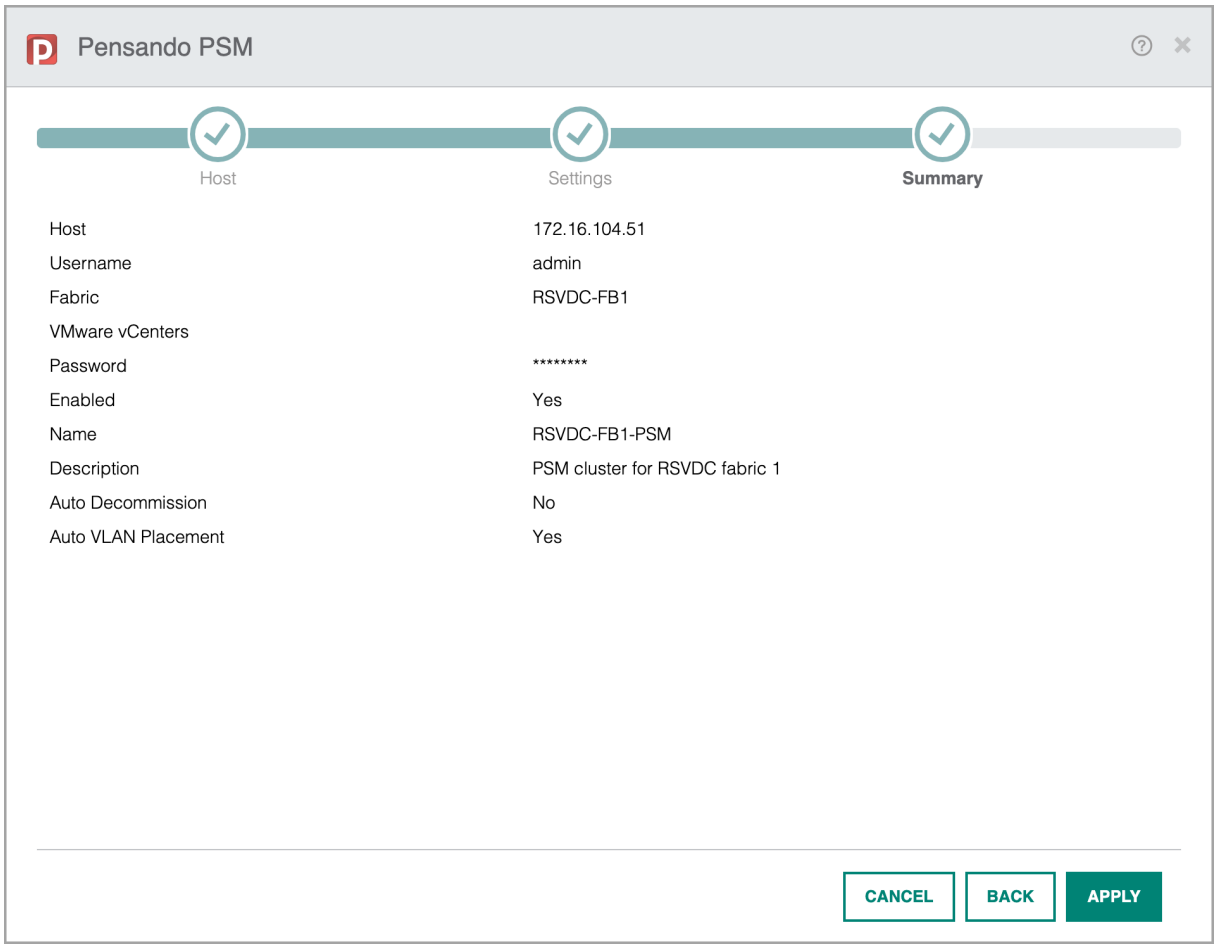

**NOTE:**

Aruba CX 10000 switches that are members of the same fabric make a join request to the PSM cluster at the completion of this step. AFC instructs PSM to admit CX 10000 switches, when auto-admission to the PSM cluster has been set to false.

**Step 7** On the AFC top right menu bar, click the **CLI Commands** icon.

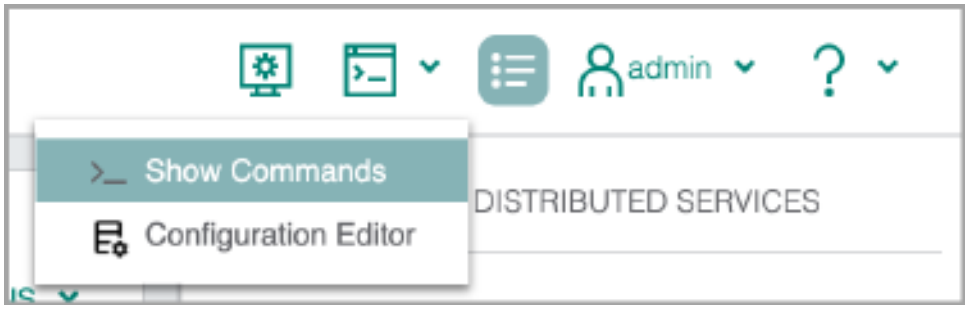


**Step 8** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** < *Select Aruba CX 10000 fabric leaf switches* >
- **Commands:** *show psm*

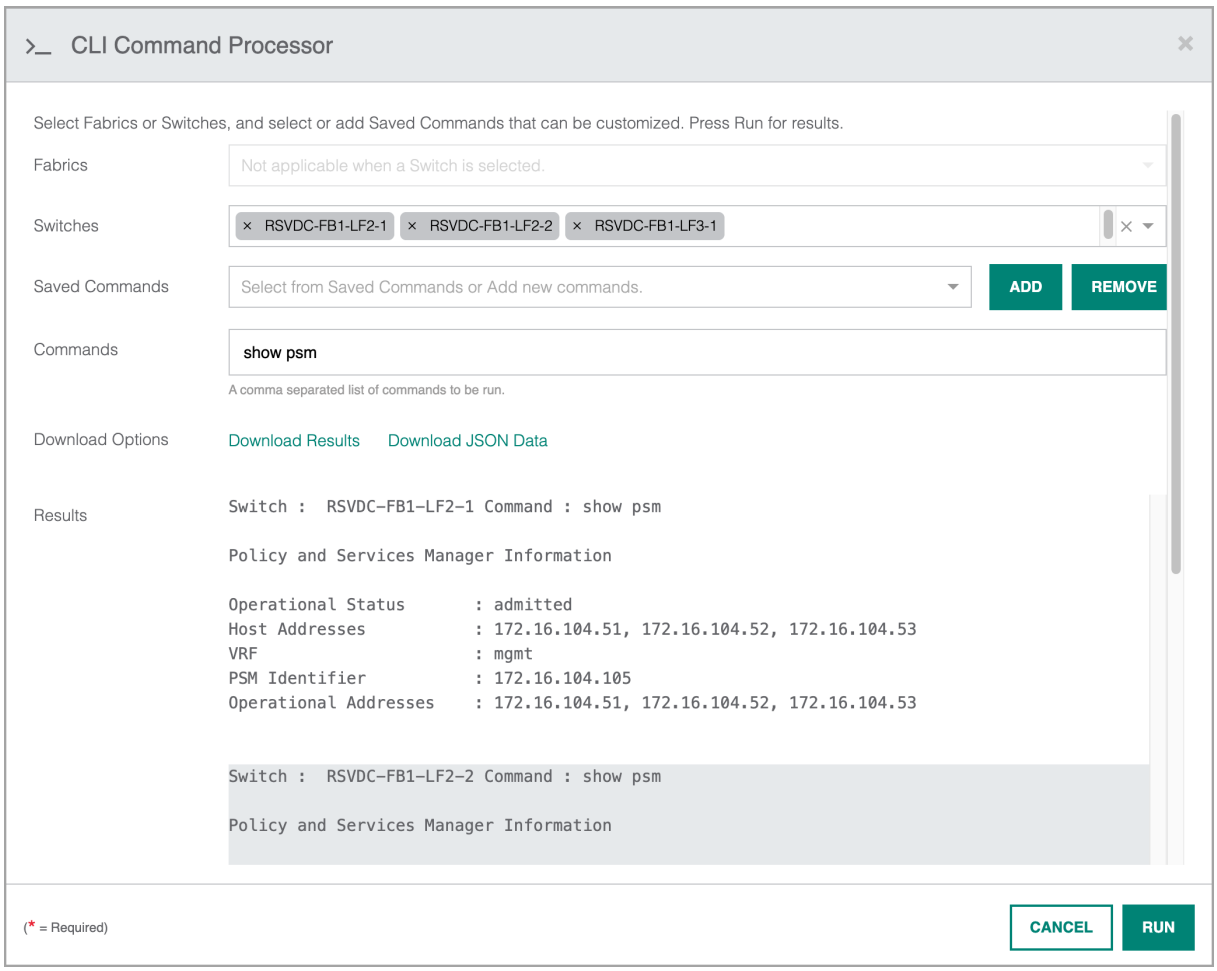


**Step 9** Verify that each CX 10000 has the following values in the command output.

- **Operational Status:** *admitted*

- **Host Addresses:** < *IP address list of PSM cluster members* >
- **VRF:** *mgmt*
- **Operational Addresses:** < *IP address list of PSM cluster members* >

# Configure Macro Firewall Policy

A firewall policy is a collection of rules applied to a PSM *Network* object or *Virtual Routing and Forwarding (VRF)* objects. When PSM policy is applied to a DSS switch, a *Network* object corresponds to a VLAN on the switch, and a *VRF* object corresponds to a VRF on the switch.

A policy contains a set of rules that specify the traffic allowed or denied between endpoint groups. Each rule contains service qualifiers or applications (sets of service qualifiers) that identify traffic type by port number. An implicit “deny all” rule is applied at the end of every policy.

The following firewall policy example restricts traffic allowed to the database server VLAN. It is applied to a PSM *Network*. For information on choosing between applying a policy to a *Network* or a *VRF*, refer to the Network Design section of the Data Center Design VSG.

The primary purpose of the sample policy is to protect database servers. At the completion of this process, web application servers are allowed to send MSSQL, ICMP, and traceroute traffic to the database servers. Other traffic is denied.

The process defines the following: - PSM *Network* objects to inform CX 10000 switches which VLANs to redirect to the firewall engine - Endpoint group objects to specify IP ranges applied to firewall rules - Firewall rules to specify allowed or denied traffic - Firewall policy that stitches the above components together

The web servers are distributed across both CX 10000 switches and non-DSS switch models in this example. The database servers are connected only to a CX 10000 leaf pairs.

## Configure PSM Networks

In this procedure, PSM *Network* objects are created to specify which VLANs redirect traffic for firewall policy enforcement. Creating a PSM *Network* in AFC will not create a VLAN on the switch. A corresponding VLAN and SVI must be defined, if they do not already exist. For complete configuration steps for prerequisite VRFs, VLANs, and SVIs, refer to Deploying the Fabric earlier in this guide.

**Step 1** On the **Distributed Services** tab of **Guided Setup**, choose **RSVDC-FB1/PROD-DC-VRF** in the **Selected VRF** dropdown. Click **CONFIGURE NETWORKS**.

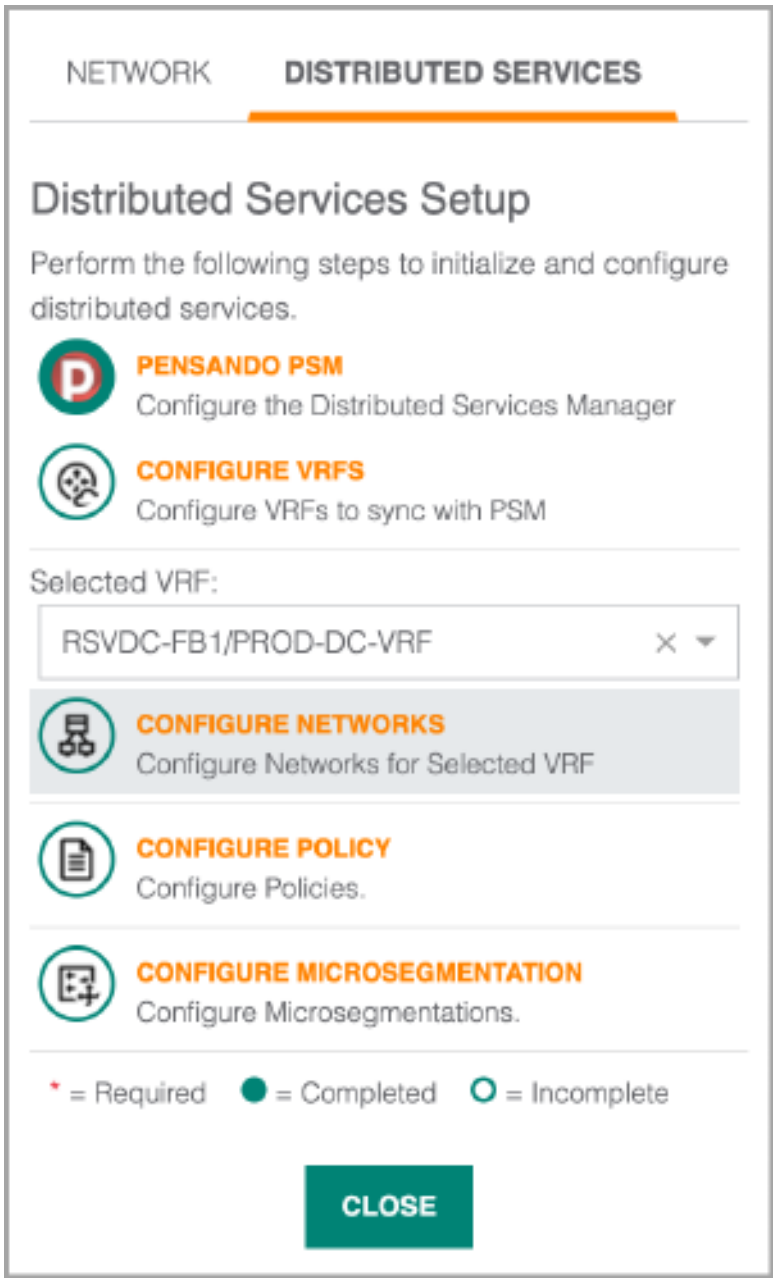


**Step 2** To add a PSM *Network* for the web server VLAN, click the right **ACTIONS** menu and select **Add** to start the **Network** workflow.

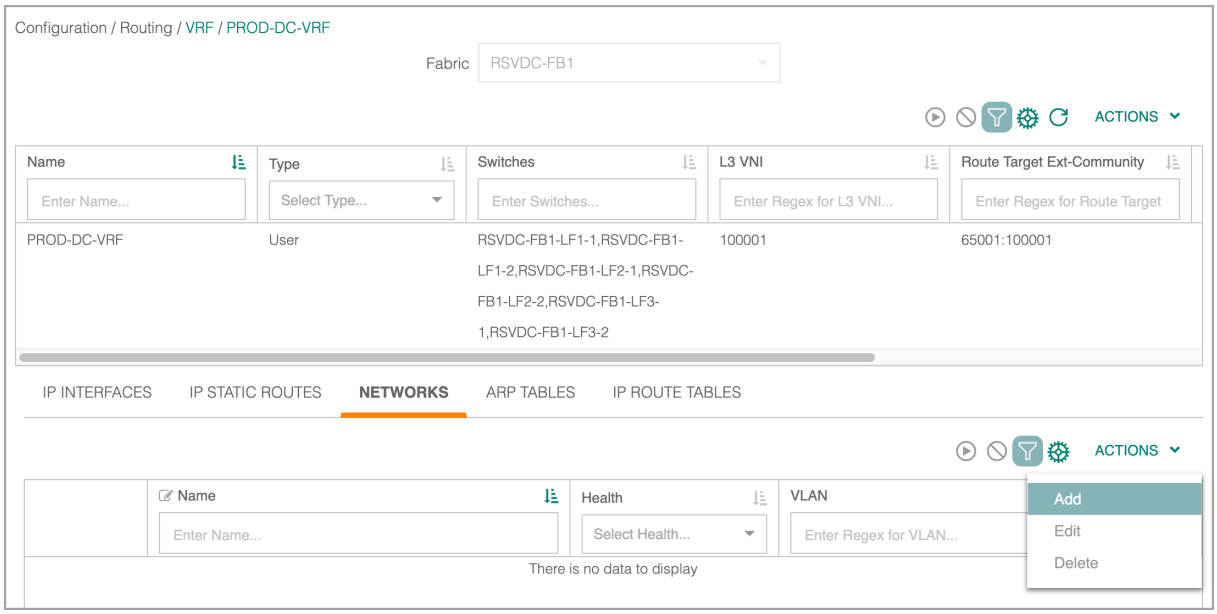


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

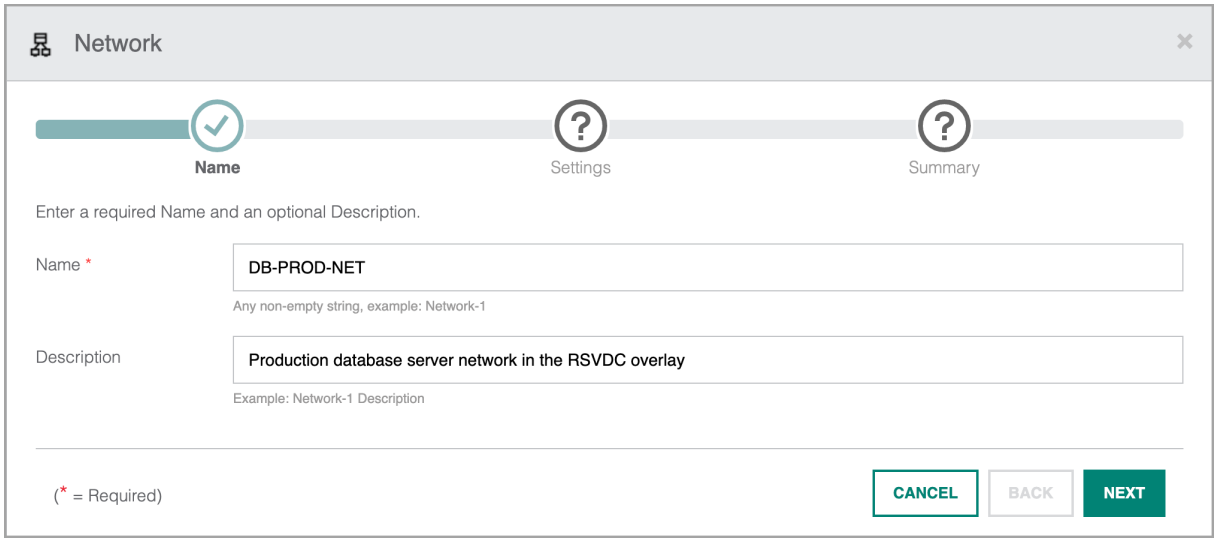


**Step 4** On the **Settings** page, enter the database server VLAN ID in the **VLAN** field to associate it with the *Network* object, then click **NEXT**.

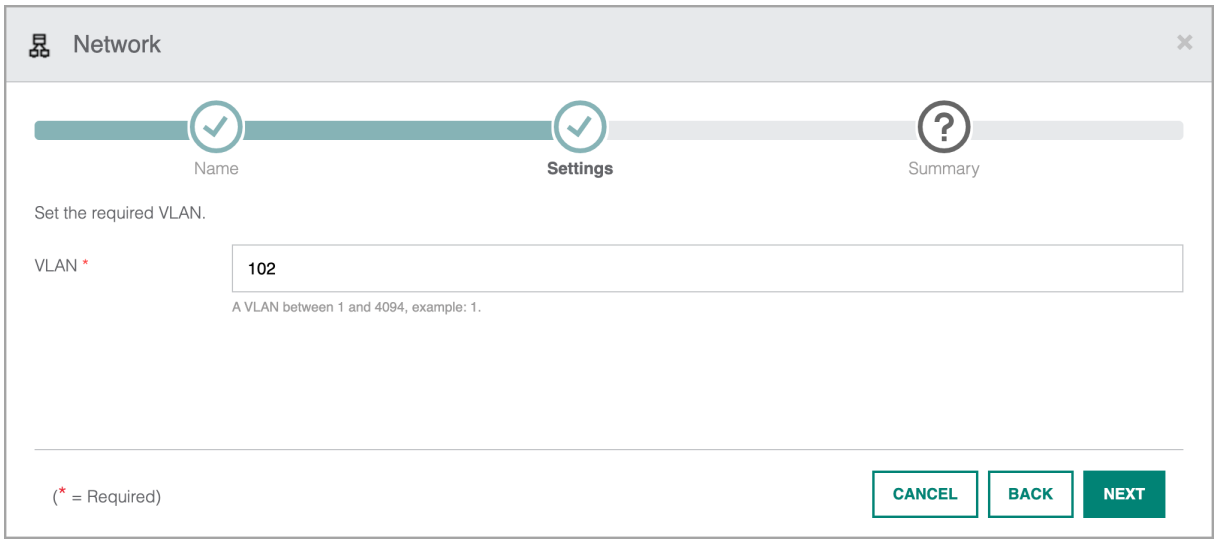


**Step 5** On the **Summary** page, verify that the information is correct and click **APPLY**.

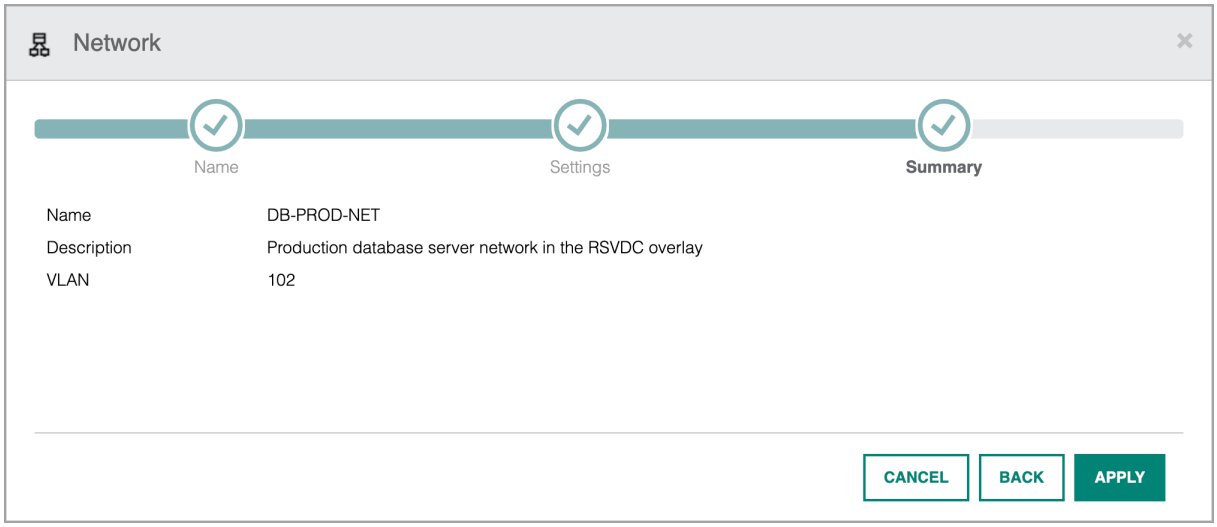


**Step 6** Repeat the procedure to define a *Network* object for every VLAN in PROD-DC-VRF.

**CAUTION:**

Communication failures may result between two hosts in different VLANs, if a *Network* object is not associated to both VLANs. It is best practice to define a *Network* for every VLAN in a VRF.

## Review Dynamic Endpoint Groups

Endpoint group objects represent source and destination IP addresses referenced in firewall rules. Dynamic endpoint groups are autopopulated based on the assignment of vSphere tags to VMs identified through AFC's vCenter integration.

All VMs assigned the same tag are autopopulated as a member of the dynamic group. When member IP addresses change or tag assignments are changed in vCenter, dynamic group objects are updated automatically.

**Step 1** On the **Configuration** menu, select **Policy > Endpoint Groups**.

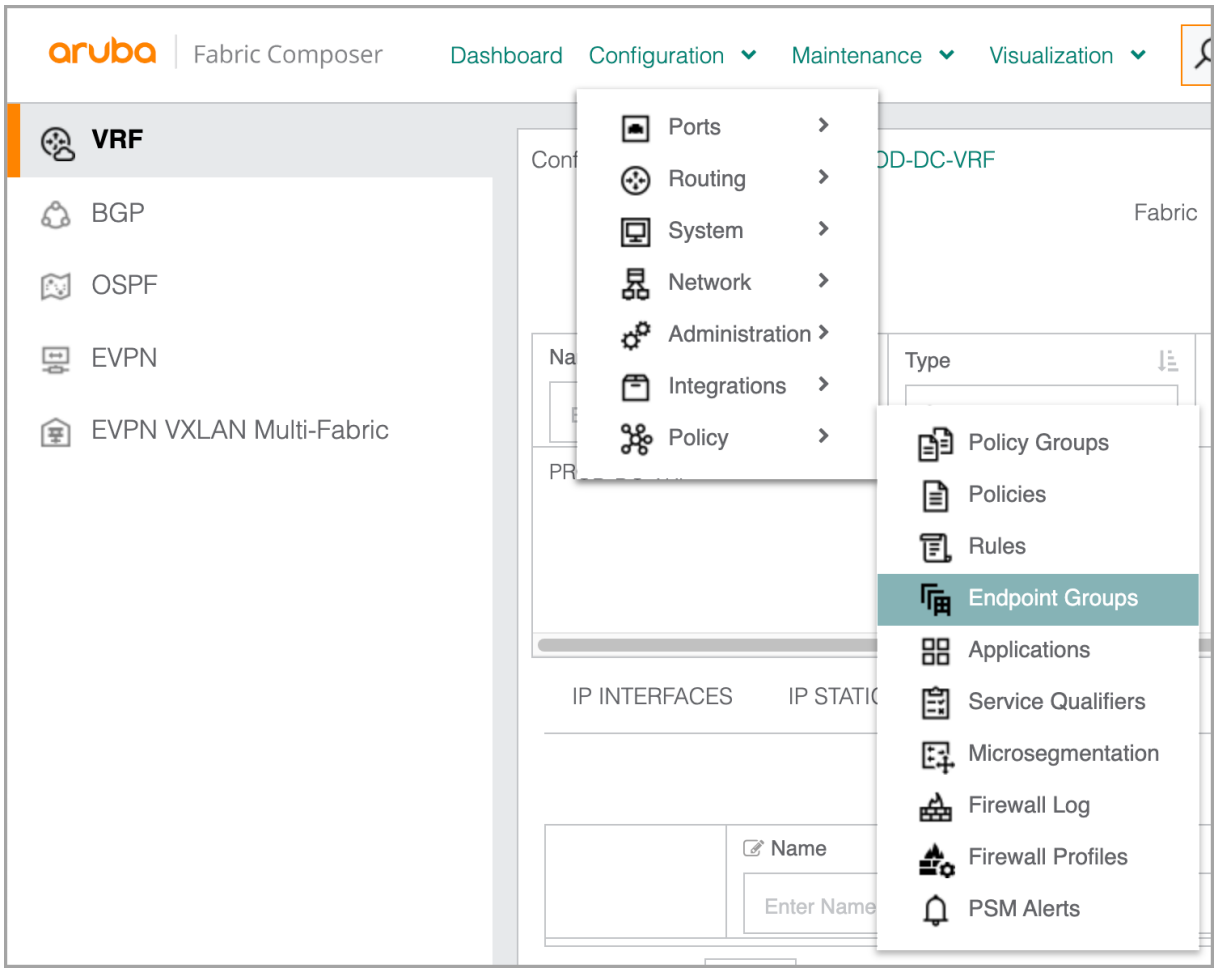


**Step 2** Verify that the VMs assigned tags in vSphere are populated in the dynamic endpoint groups.

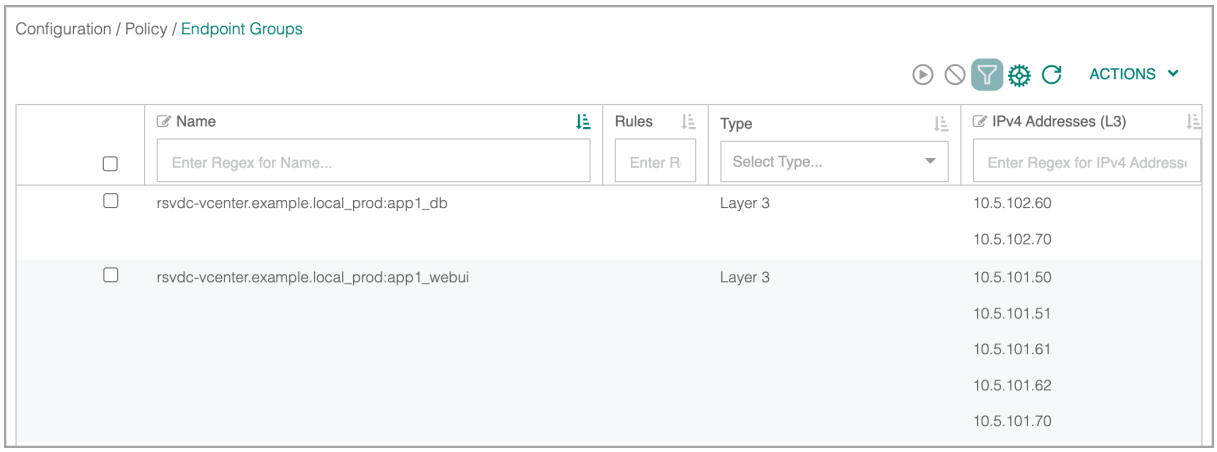


**NOTE:**

If no dynamic endpoint groups are present, verify that **Automated Endpoint Group Provisioning** is enabled in **Configuration > Integration > VMware vSphere**. Verify that VMs in vSphere have been assigned tags and are currently running. Refer to the Deploying the Fabric section for more information on integration.

## Configure Static Endpoint Groups

Additional endpoint groups are created to represent IP addresses not automatically populated in a dynamic group.

**Step 1** To create an endpoint group for the network services subnet, in the **Configuration > Policy > Endpoint Groups** context, click the right **ACTIONS** menu and select **Add**.

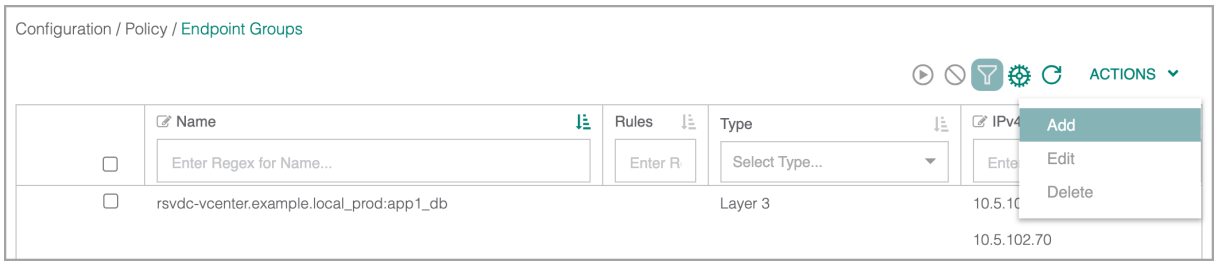


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

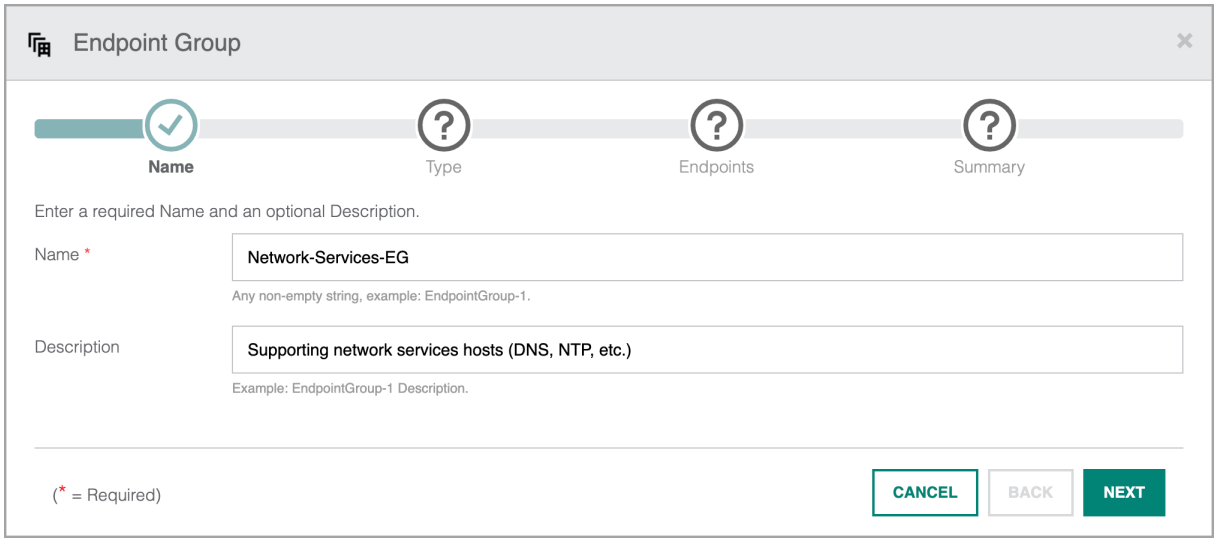


**Step 3** On the **Type** page, leave **Type** set to the default **Layer 3** value and click **NEXT**.

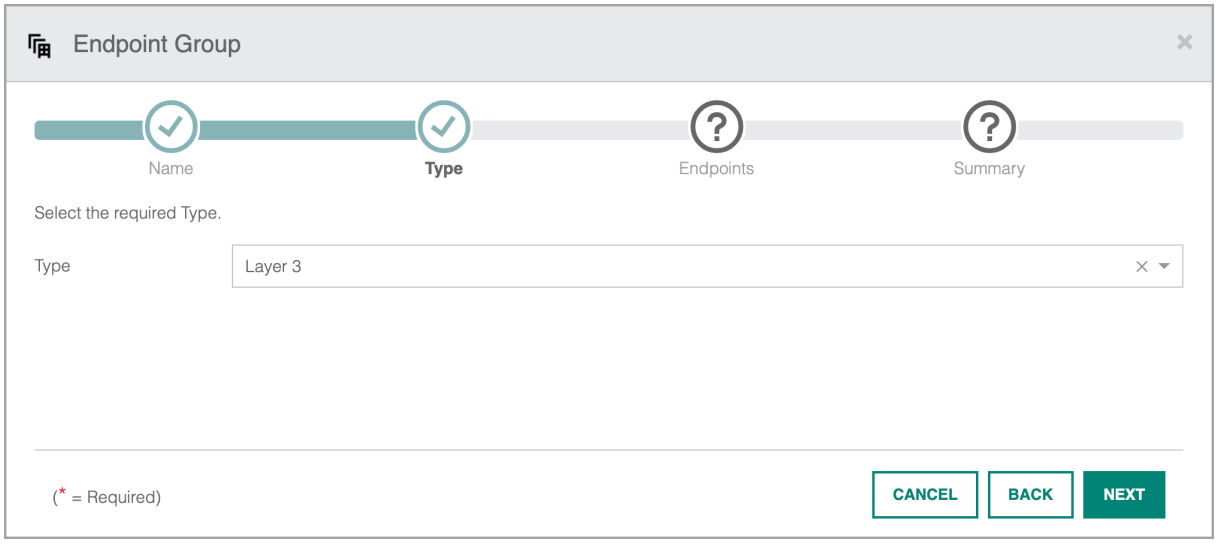


**Step 4** On the **Endpoints** page, click the **Manual Endpoint entry** radio button, enter the web server subnet value in the **IPv4 Network Address** field, and click **ADD**.

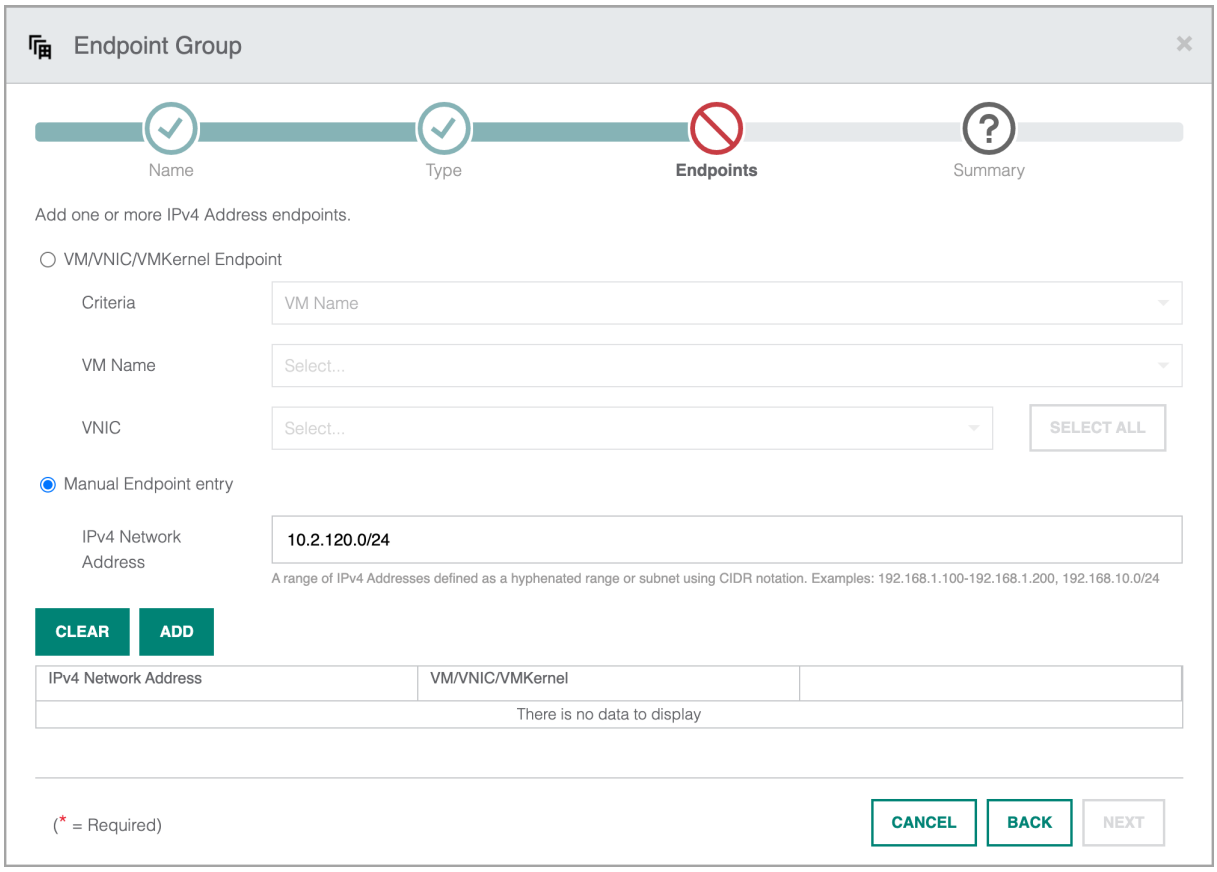


**NOTE:**

A single endpoint group can be a collection of IP ranges and individual addresses. Static endpoint groups can be made using AFC's vSphere integration. When making static assignments in the **Endpoint Group** wizard based on a VM Name or VM Tag, the IP addresses added to the endpoint group are established at the time of the endpoint group's creation. Future changes to VM IP addresses or tag assignments are not updated automatically in the address object. Use dynamic endpoint groups to automatically update VM IP assignments.

**Step 5** Verify that the subnet was added to the list at the bottom of the **Endpoint Group** window and click **NEXT**.

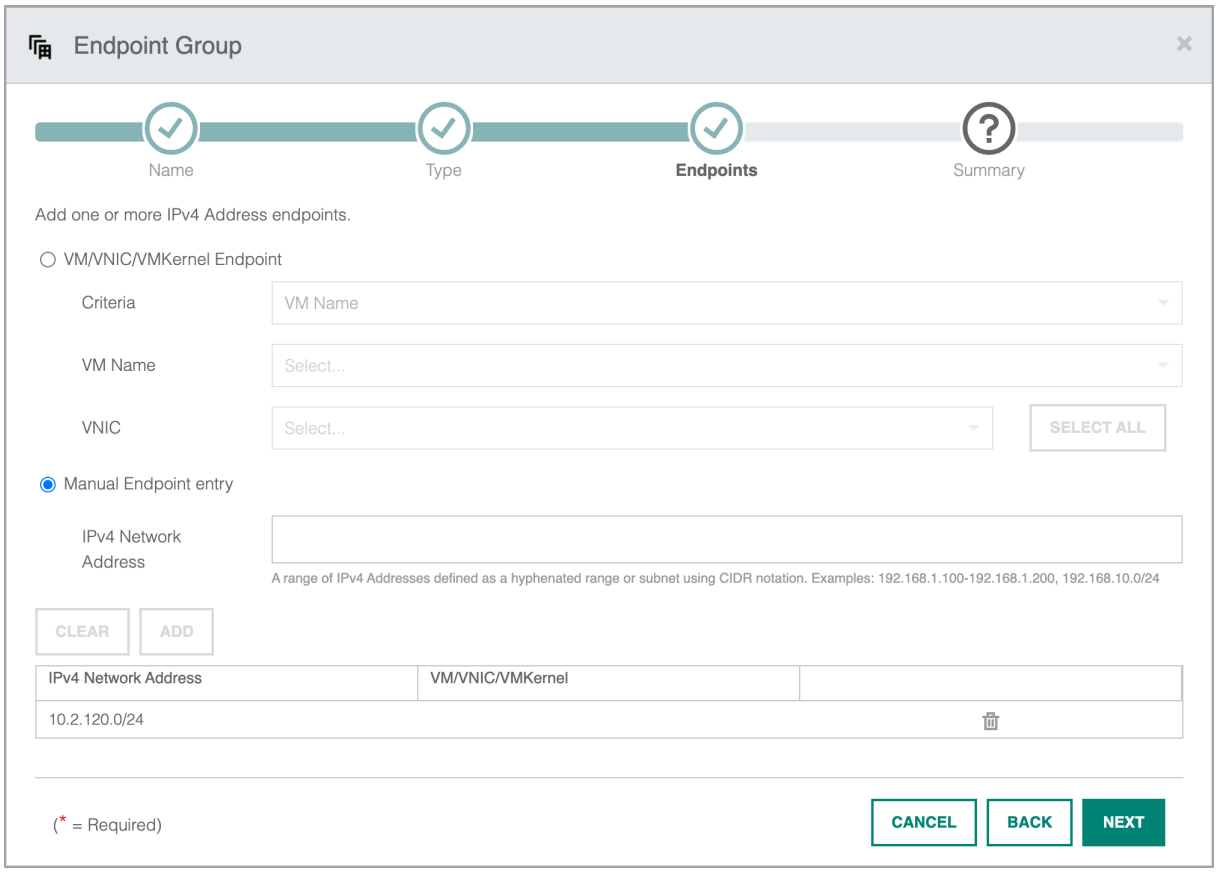


**Step 6** On the **Summary** page, verify that all information is correct and click **APPLY**.

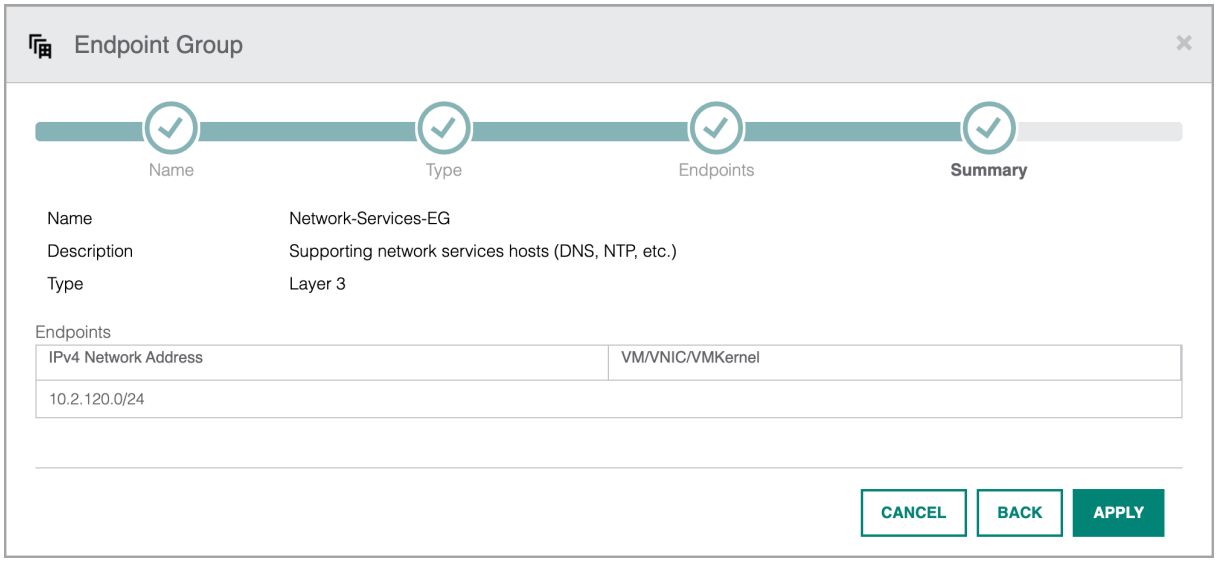


**Step 7** Repeat the procedure to create additional endpoint groups to be used in firewall rules using the following values:

| Name | Description | Type | Address |
|---|---|---|---|
| PROD-DC-VRF-Summary-EG | Summary IP prefix for PROD-DC-VRF in the DC overlay | Layer 3 | 10.5.0.0/16 |

| Name | Description | Type | Address |
|---|---|---|---|
| Campus-Admins-EG | Campus IP subnet used by application administrators | Layer 3 | 10.254.1.0/24 |
| All-Hosts-EG | All hosts endpoint group | Layer 3 | 0.0.0.0/0 |

## Configure Firewall Rules

In the following procedure, firewall rules are created to define the inbound traffic allowed to the backend database server VLAN.

In addition to allowing communication from frontend web servers, rules that enable network service functions and troubleshooting protocols are defined. This is required because any traffic not explicitly allowed in firewall policy is blocked by an implicit deny rule.

It is important to consider the policy direction, when defining rules. In this example, an ingress policy is used to filter both DSS and non-DSS sourced traffic to the database servers. Policy direction and design vary based on requirements and administrator preferences.

An individual rule identifies source and destination endpoint groups, a service qualifier or application to specify the traffic type, and an action of “allow” or “deny”. A rule does not allow or deny traffic until it is added to a policy.

**Step 1** In the AFC left navigation pane, click **Rules**.

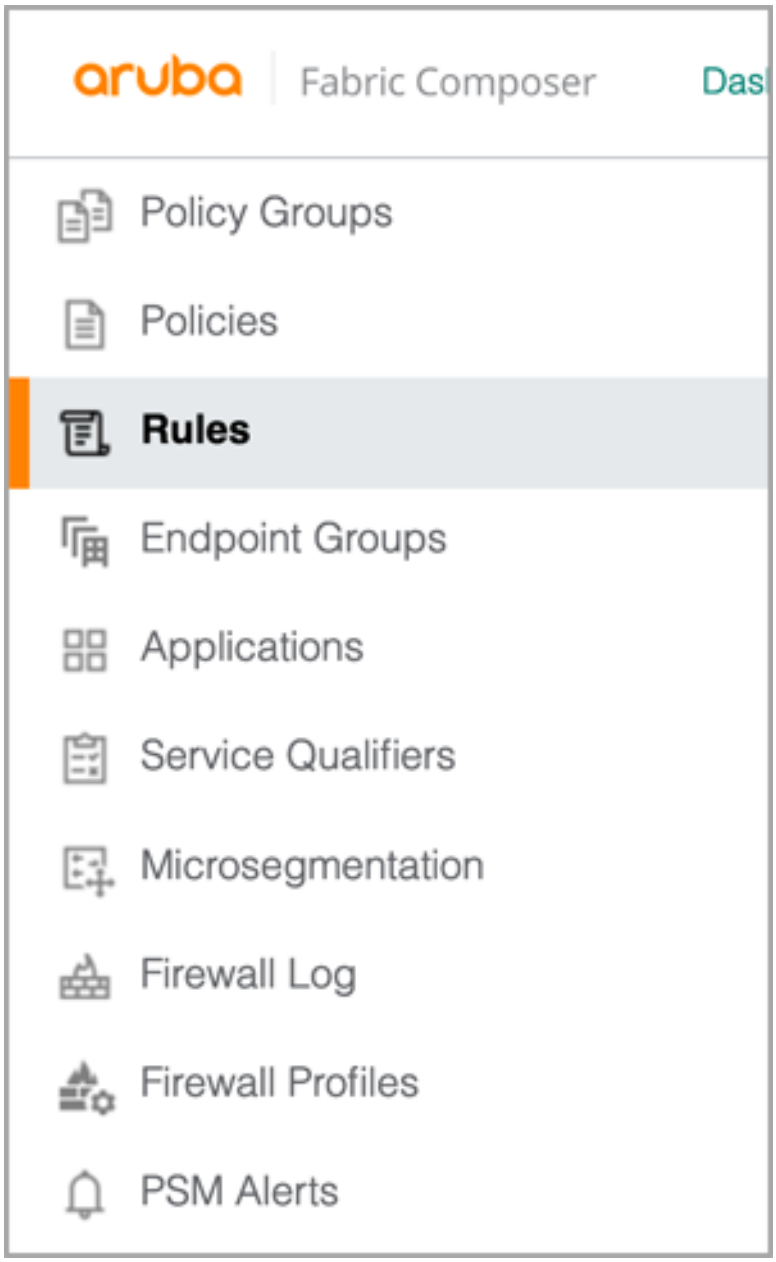


**Step 2** To create a rule that allows web servers to reach database servers using the MSSQL protocol, click the right **ACTIONS** menu and select **Add**.

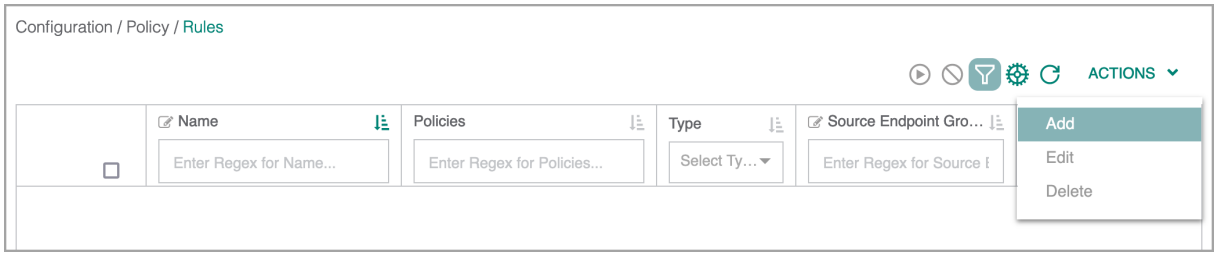


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

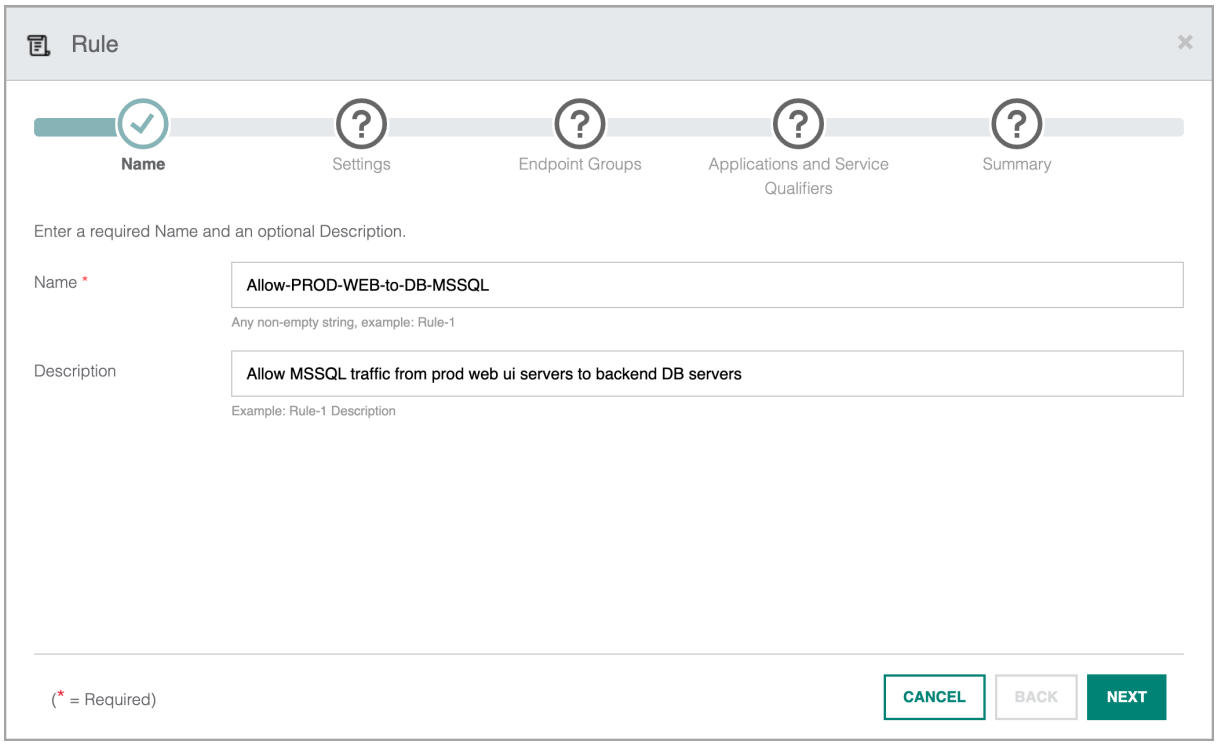


**NOTE:**

A helpful rule naming convention is to include the action of allow or deny, endpoint group, and service in the rule name. This simplifies policy construction later.

**Step 4** On the **Settings** page, leave the **Type** and **Action** fields at their default values and click **NEXT**.

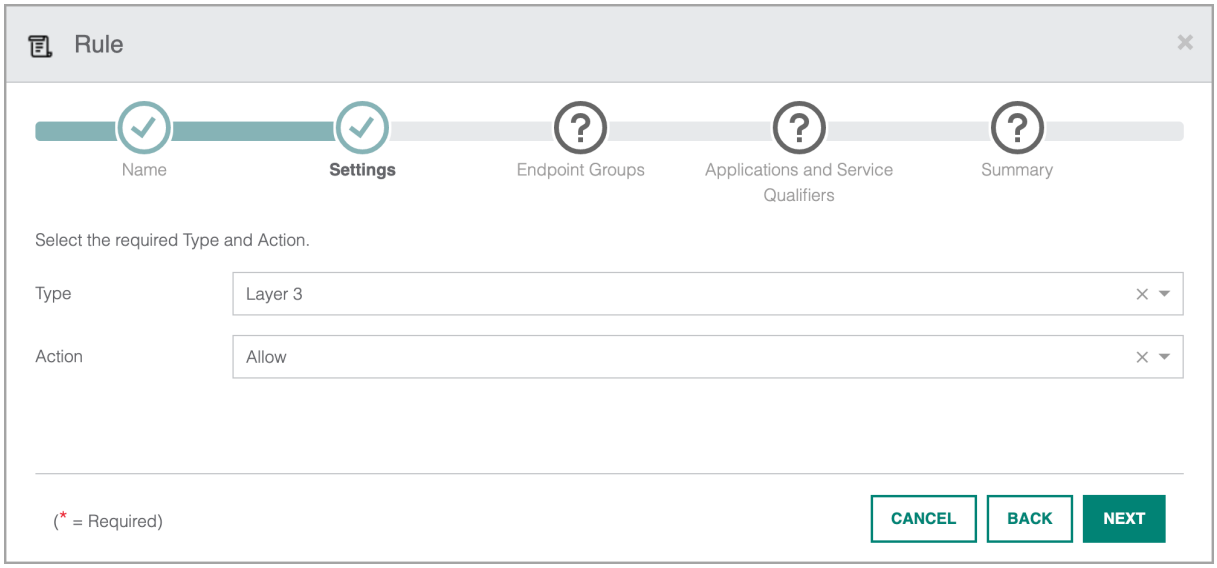


**Step 5** On the **Endpoint Groups** page, select the endpoint groups dynamically created using VM tags as follows, then click **NEXT**.

- **Source Endpoint Groups:** *rsvdc-vcenter.example.local_prod:app1_webui*
- **Destination Endpoint Groups:** *rsvdc-vcenter.example.local_prod:app1_db*

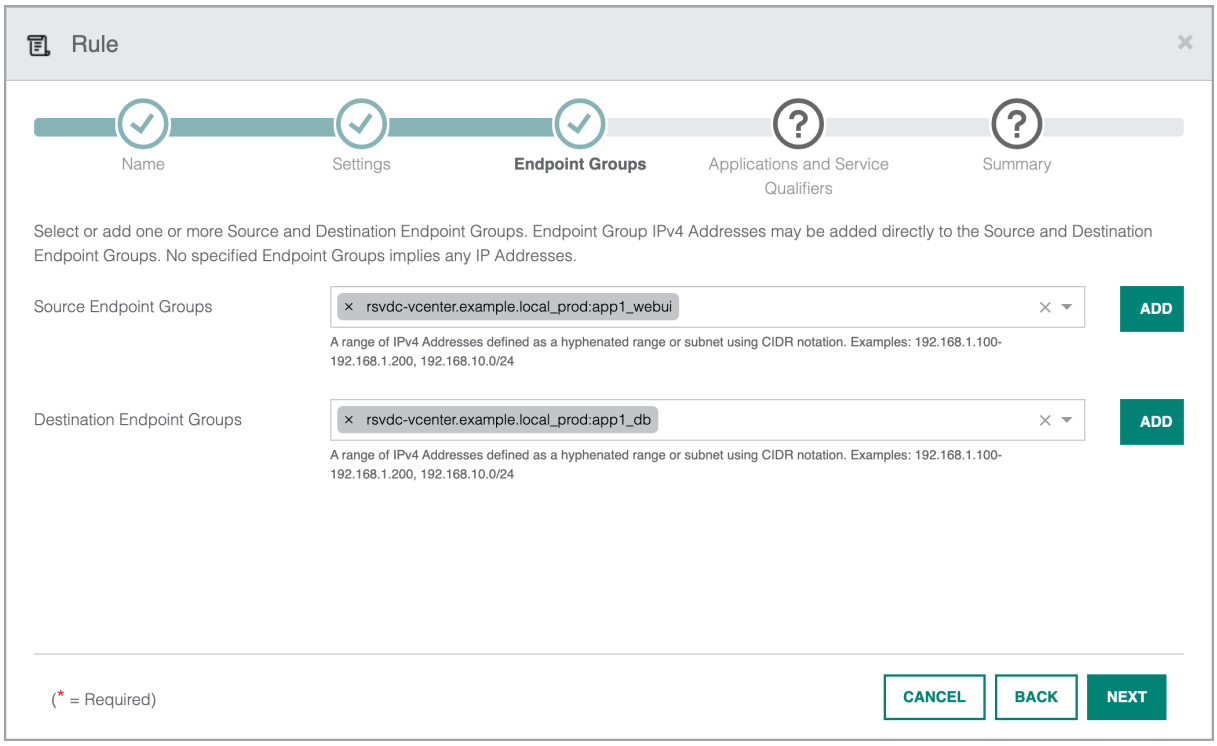


**NOTE:**

AFC automatically updates IP addresses in a firewall rule that uses dynamically created endpoint groups when a VM administrator changes IP assignments or VM tags in vCenter.

**Step 6** On the **Applications and Service Qualifiers** page, begin typing the predefined *mssql_server* name in the **Service Qualifier** field, click **mssql_server** when it appears in the list, then click **NEXT**.

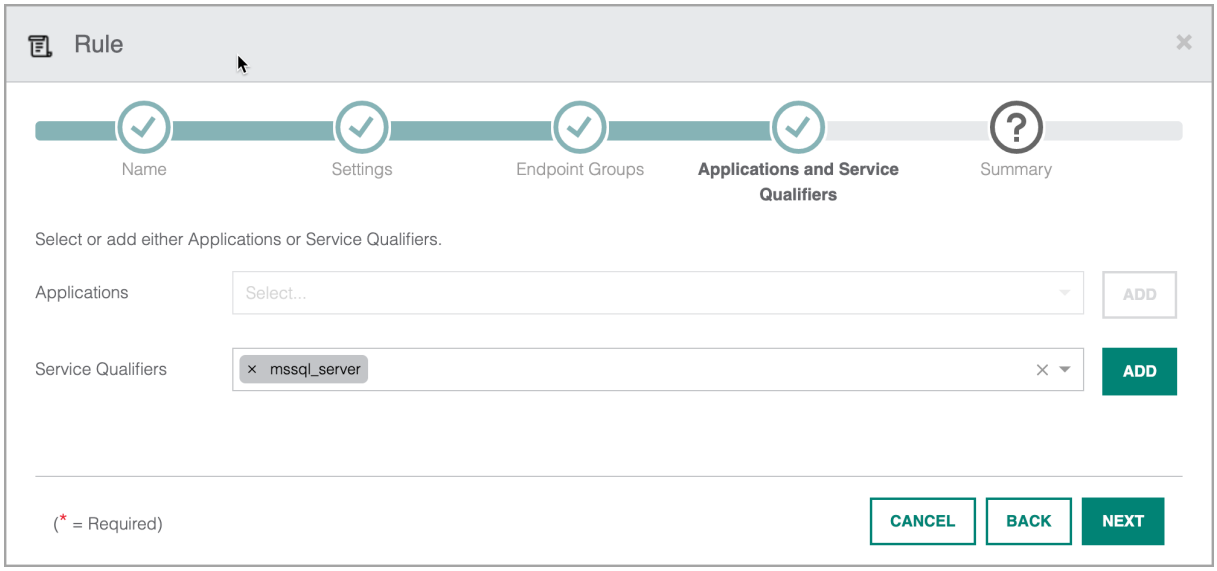


**Step 7** On the **Summary** page, verify that all information is correct and click **APPLY**.

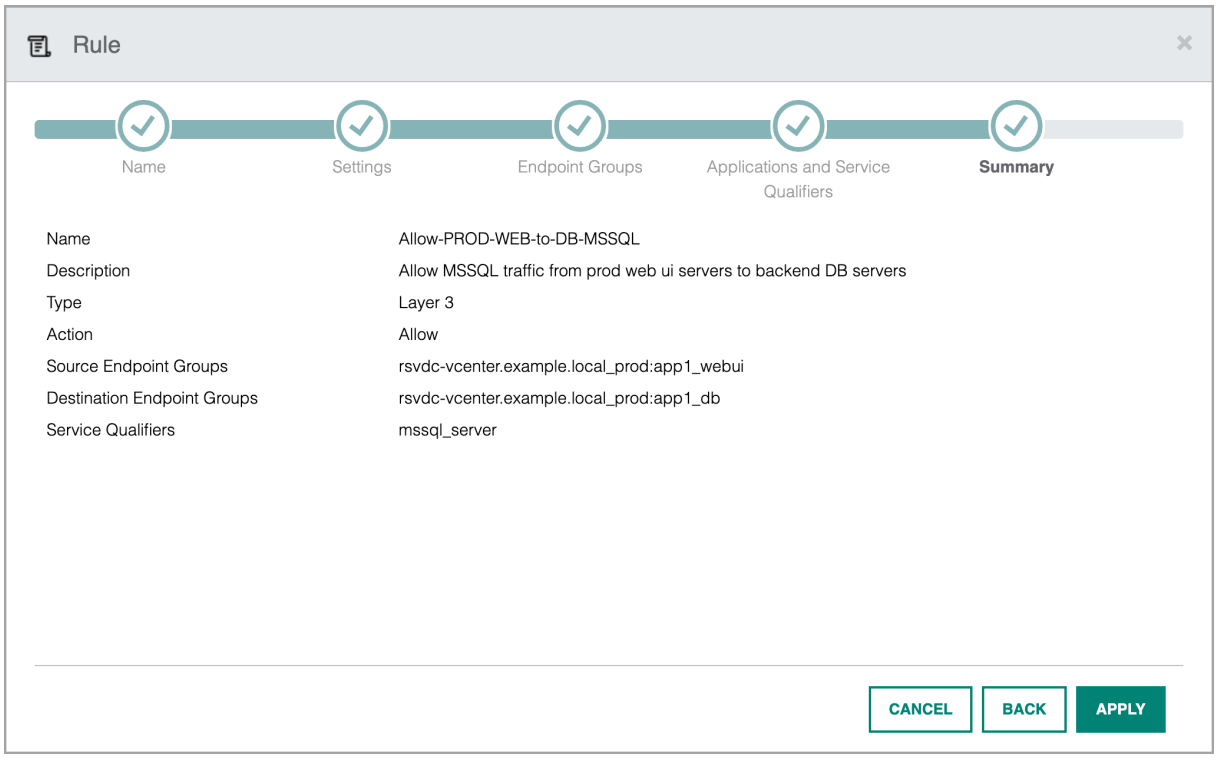


**Step 8** Repeat the procedure to create rules with the following values for use in the the database server *Network* ingress policy.

| Name | Description | Type | Action | Endpoint Groups | Applications and Service Qualifiers |
|---|---|---|---|---|---|
| Allow-All-ICMP | Allow ICMP from all hosts to all hosts | Layer 3 | Allow | **Source**: All-Hosts-EG **Destination**: All-Hosts-EG | **Service Qualifiers**: icmp |
| Allow-All-UDP-Traceroute | Allow UDP-based traceroute from all hosts to all hosts | Layer 3 | Allow | **Source**: All-Hosts-EG **Destination**: All-Hosts-EG | **Service Qualifiers**: traceroute |
| Allow-Admins-to-PROD-DC-VRF | Allow campus admin net hosts to production VRF in DC overlay | Layer 3 | Allow | **Source**: Campus-Admins-EG **Destination**: PROD-DC-VRF-Summary-EG | **Service Qualifiers**: all |

**NOTE:**

Reduce the total number of rules created by defining rules that can be applied to more than one policy.A new rule can be cloned from an existing rule by selecting **ACTIONS > Clone**, when minor modifications are required for similar rules.

## Configure Ingress Policy

When using an EVPN-VXLAN overlay, a PSM ingress policy filters all traffic arriving from other switches inside the fabric (DSS or non-DSS). This includes both Layer 2 and Layer 3 VXLAN forwarded traffic, and traffic originating from outside the data center fabric. An ingress policy also applies to routed traffic between hosts attached to the same CX 10000 switch, when using AOS-CX 10.10.1000 or later.

By default, an ingress policy does not apply to traffic between hosts in the same subnet attached to the same CX 10000 switch. This type of policy application can be achieved when using a PVLAN-based microsegmentation strategy, as described in ESXi VM Microsegmentation. Egress policy also can be applied to filter traffic between hosts attached to the same CX 10000 that are in the same VLAN and subnet, as long as the traffic traverses the switch.

The diagram below illustrates the components that protect the database servers from non-MSSQL traffic sourced by web servers:

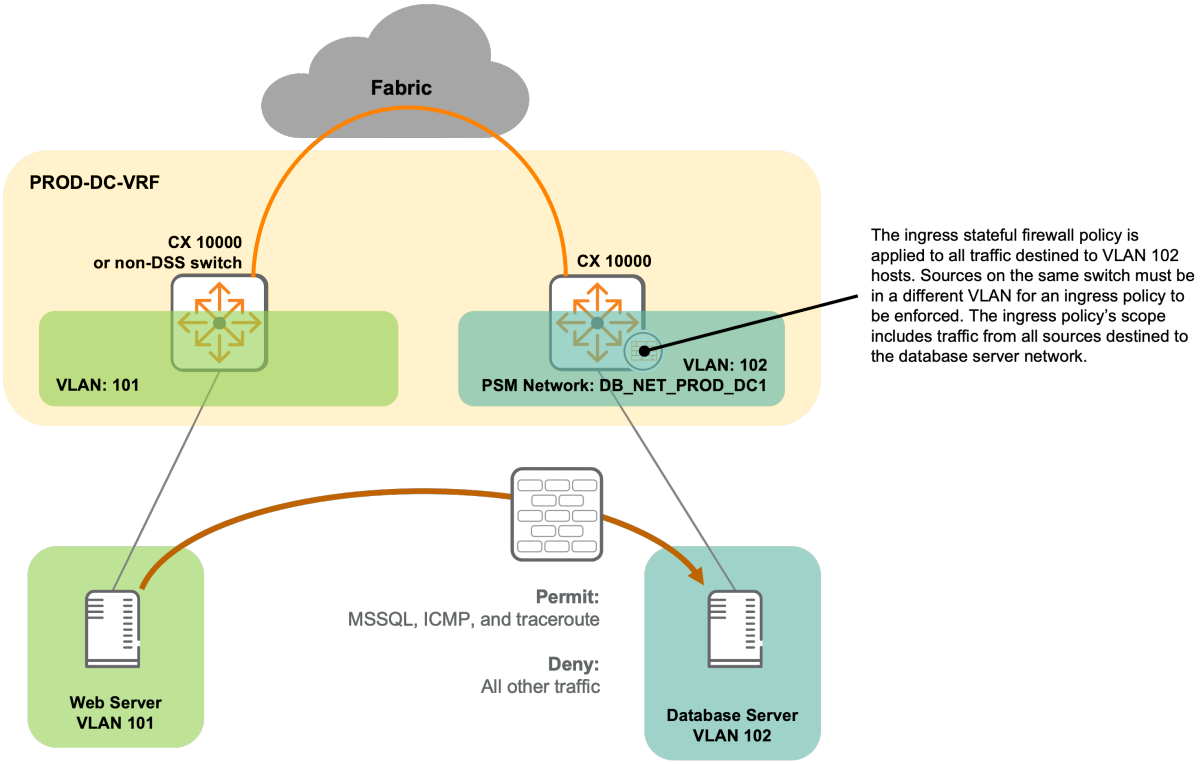


**Figure 6:** Ingress Policy Diagram

In the following procedure, the previously created rules are assigned to a policy to filter traffic destined for the database server VLAN.

**Step 1** In **Configuration** / **Policy** / **Policies**, click the right **ACTIONS** menu and select **Add**.

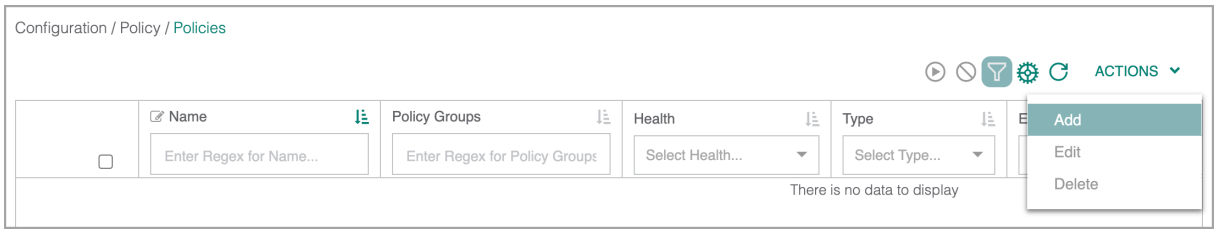


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

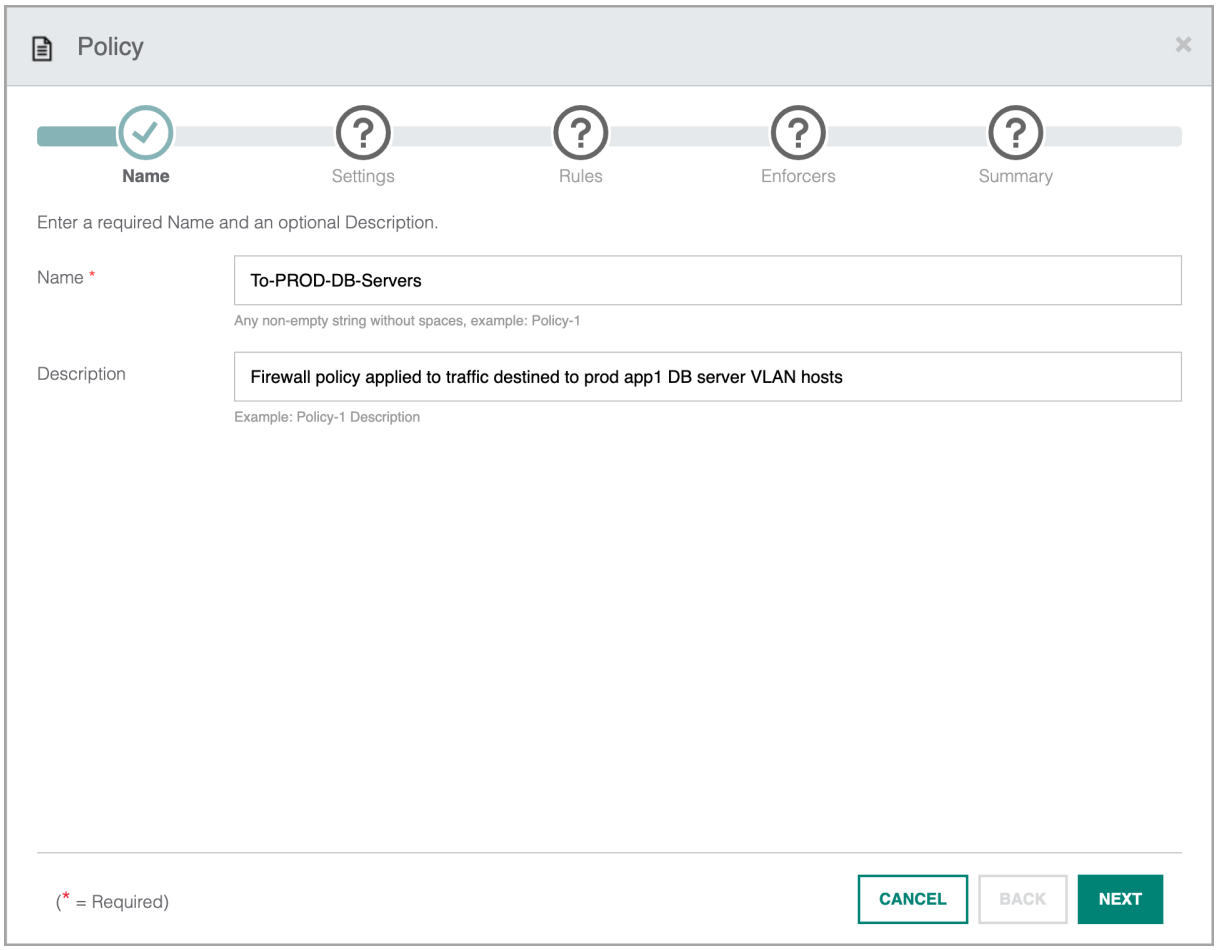


**Step 3** On the **Settings** page, leave the default value of *Distributed Firewall* for **Type** selected and click **NEXT**.

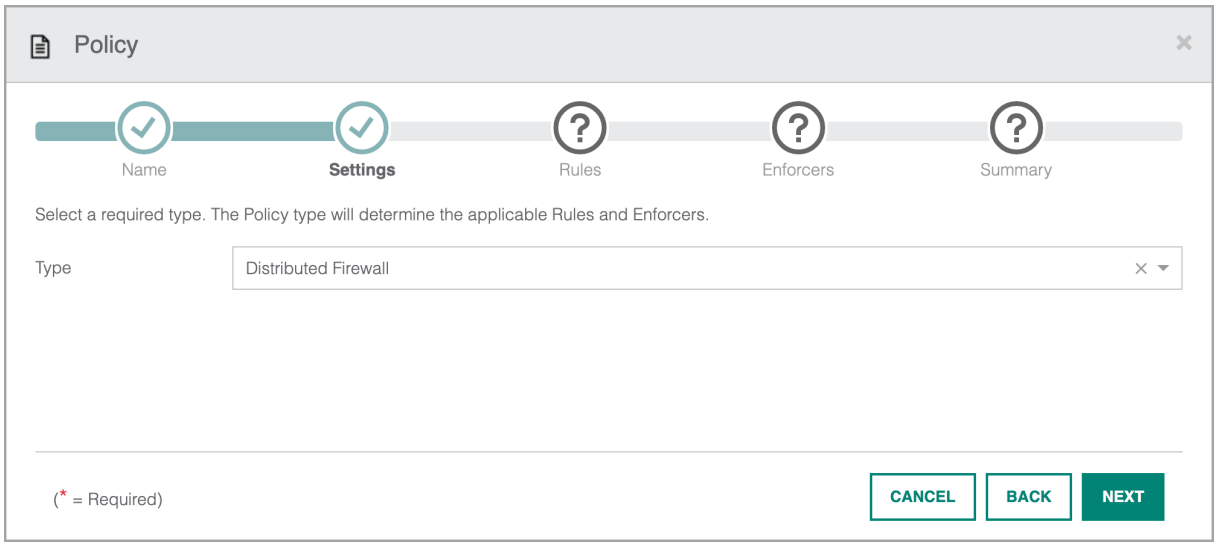


**Step 4** On the **Policy Rules** page, click the right **ACTIONS** menu and select **Add > Existing**.

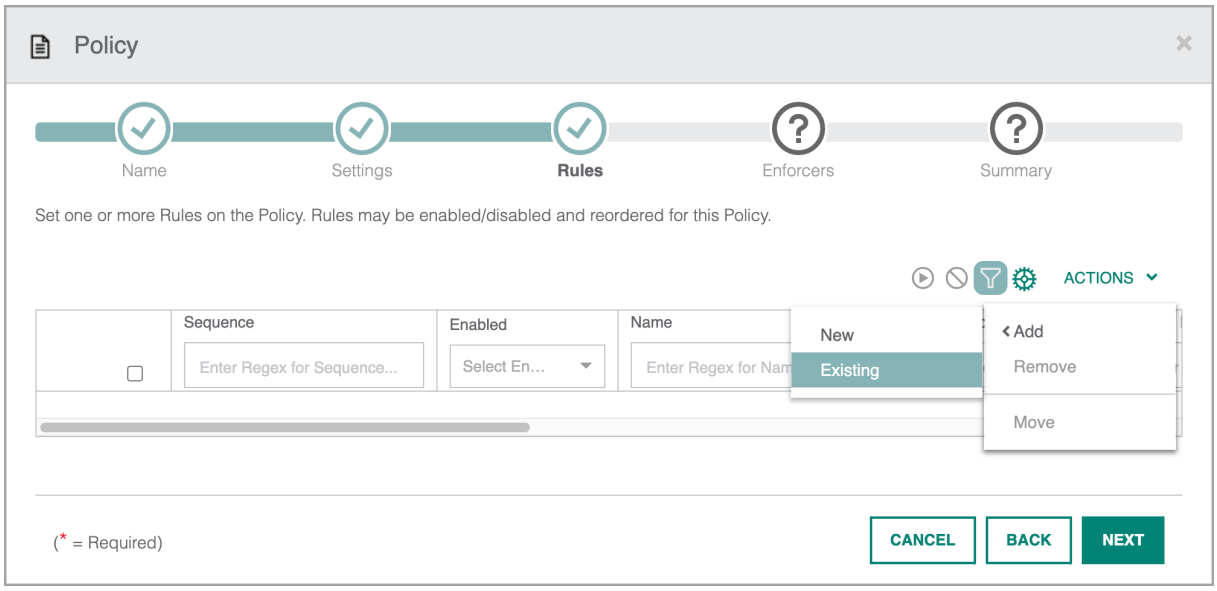


**Step 5** In the **Select Rules** window, click the checkboxes for rules allowing traffic destined to the database servers in the following order, then click **APPLY**.

- Allow-PROD-WEB-to-DB-MSSQL
- Allow-Admins-to-PROD-DC-VRF
- Allow-All-ICMP
- Allow-All-UDP-Traceroute

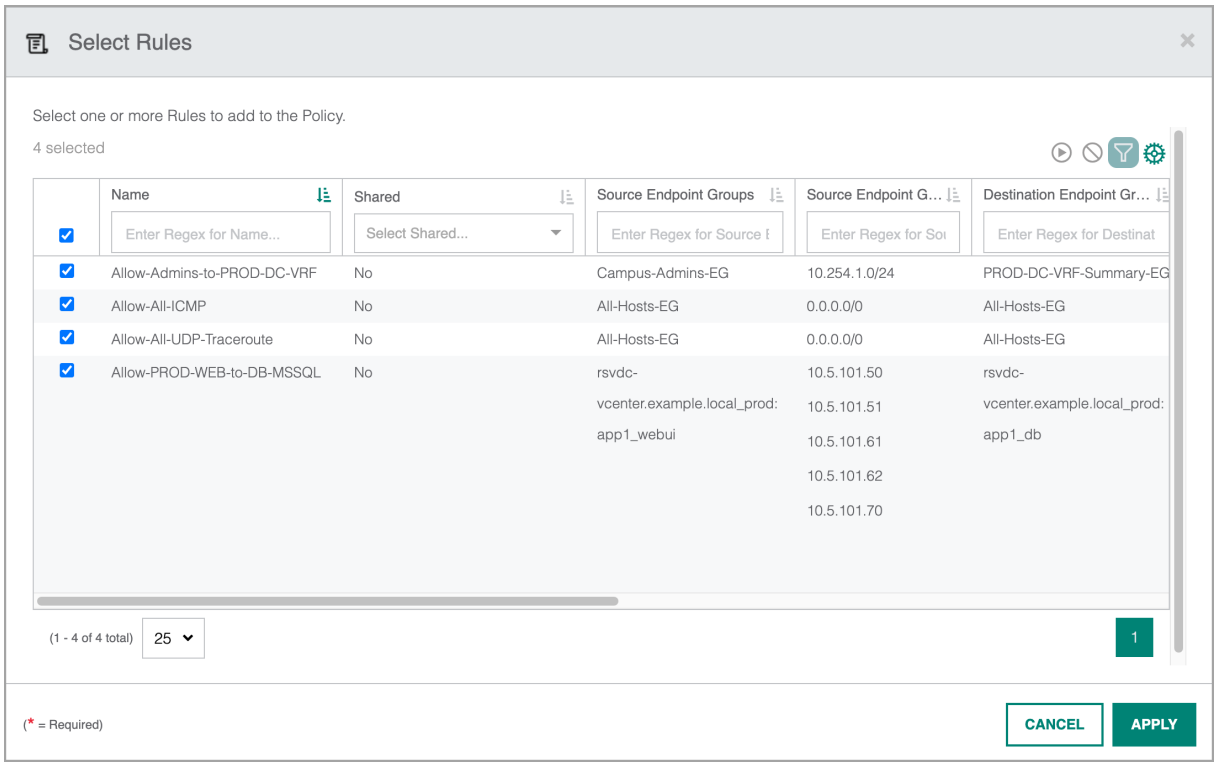


**NOTE:**

After applying an ingress policy for the database server VLAN, all traffic destined for a host in the VLAN must be explicitly allowed by a firewall rule in the policy or it is blocked. The only exception is traffic originated in the database server VLAN where the source and destination are attached to the same CX 10000 switch.

**Step 6** Verify that the rule set is in the desired order and click **NEXT**.

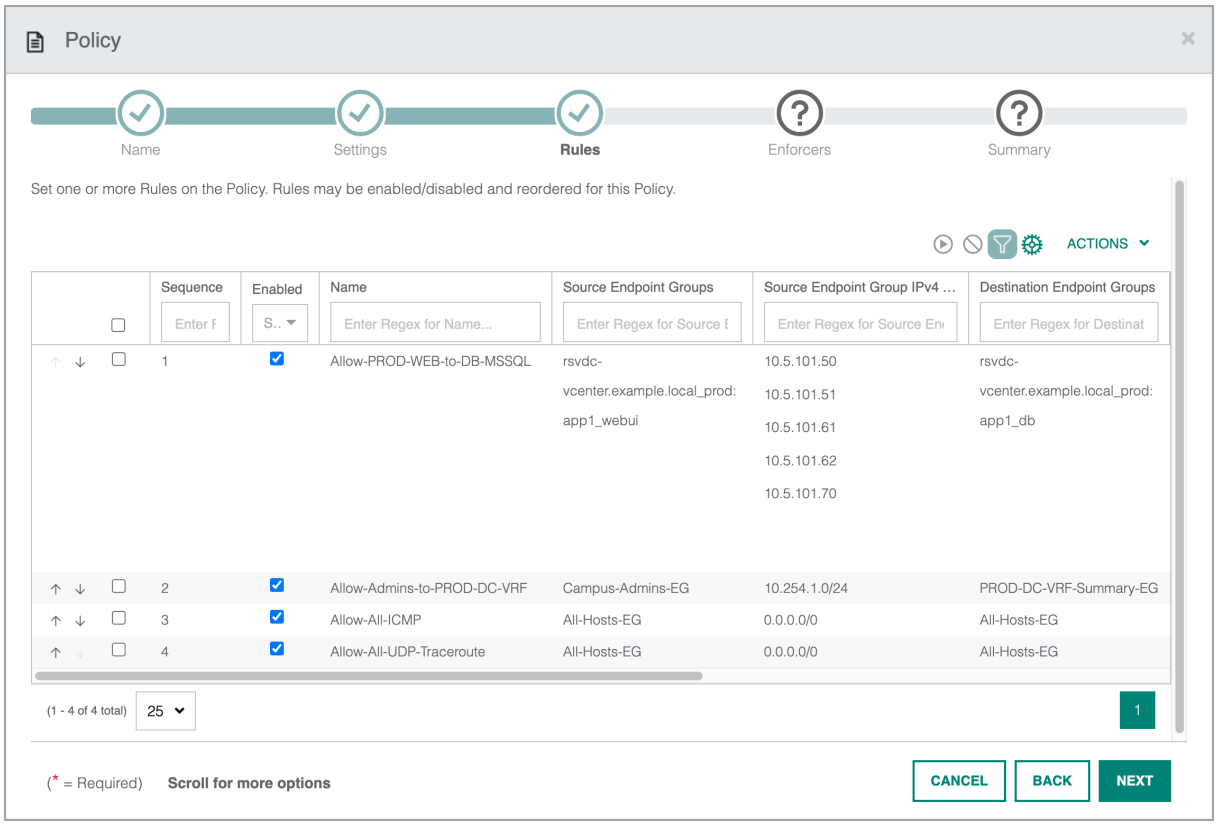


**Step 7** On the **Enforcers** page, select the following values. Click the **ADD** button near the bottom left of the window (below the dropdowns).

- **Fabric:** *RSVDC-FB1*
- **Direction:** *Ingress*
- **VRF:** *PROD-DC-VRF*
- **Networks:** *DB-NET-PROD - VLAN: 102*

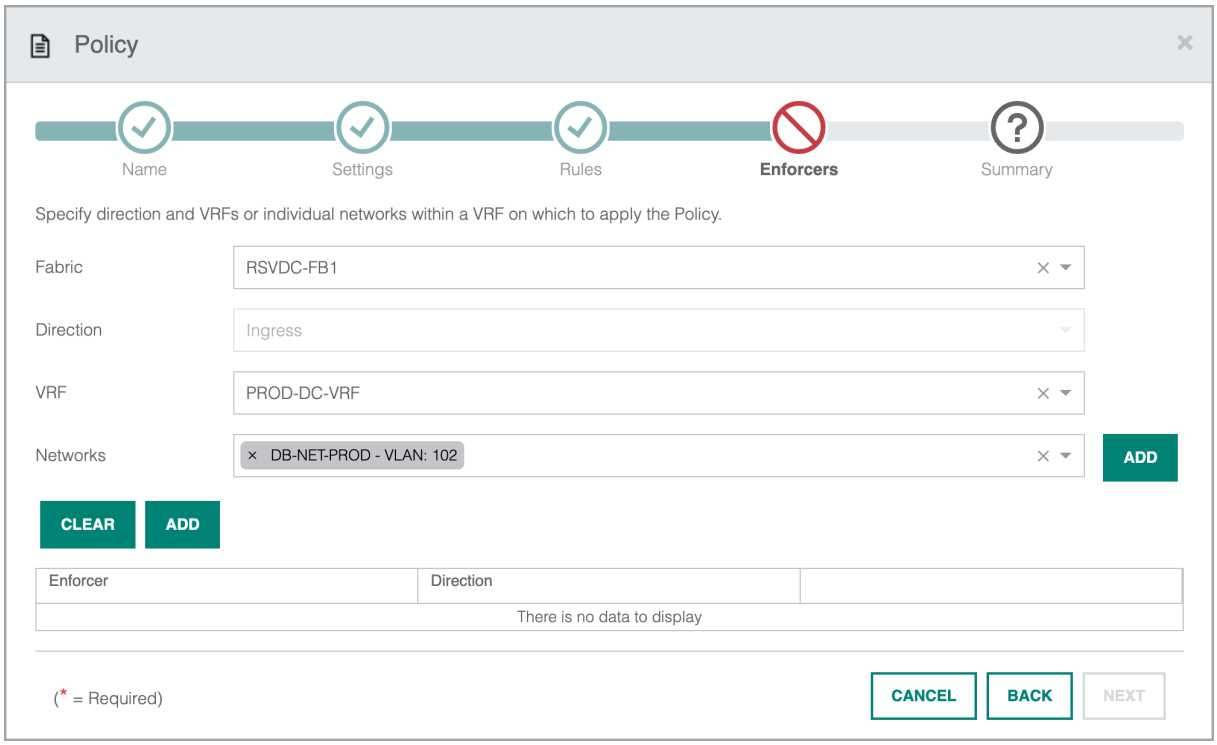

**Step 8** Verify that the enforcer information was added correctly and click **NEXT**.

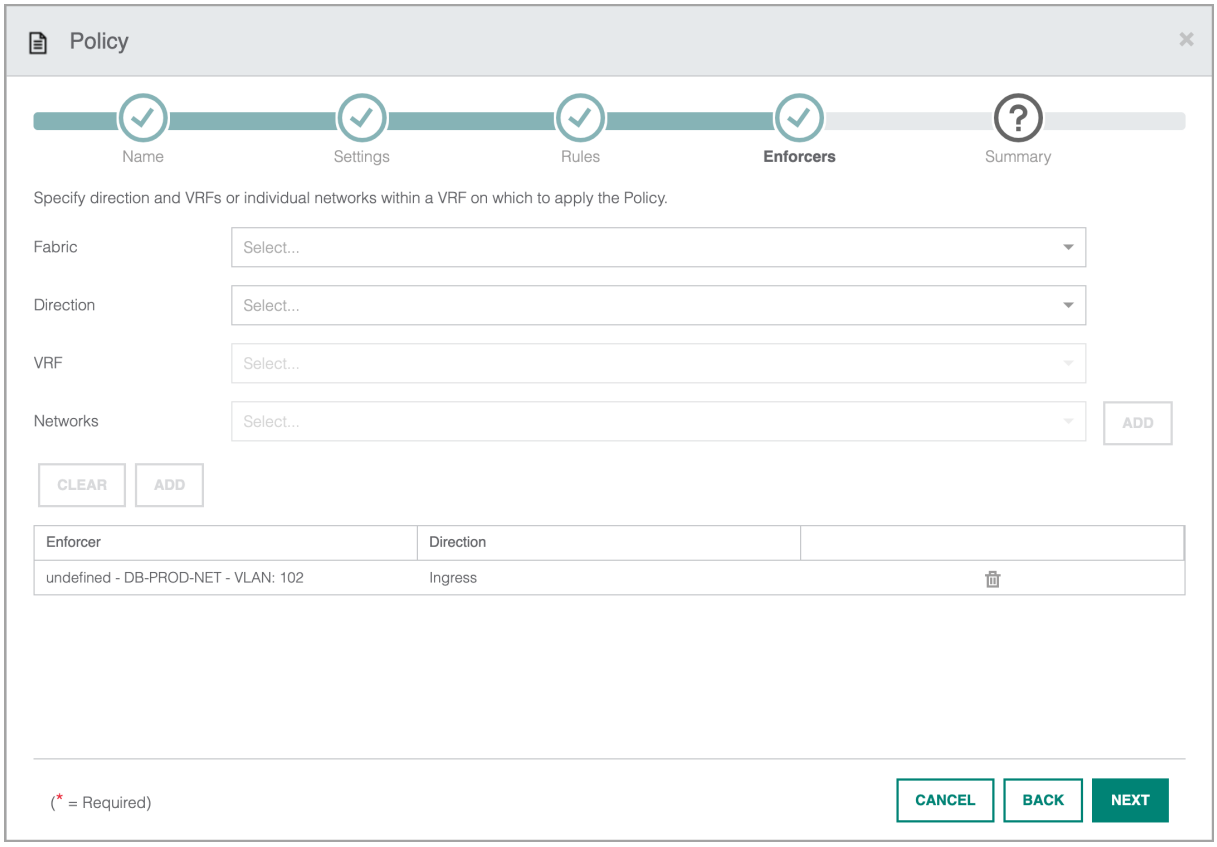


**Step 9** On the **Summary** page, verify that the information is correct and click **APPLY**.

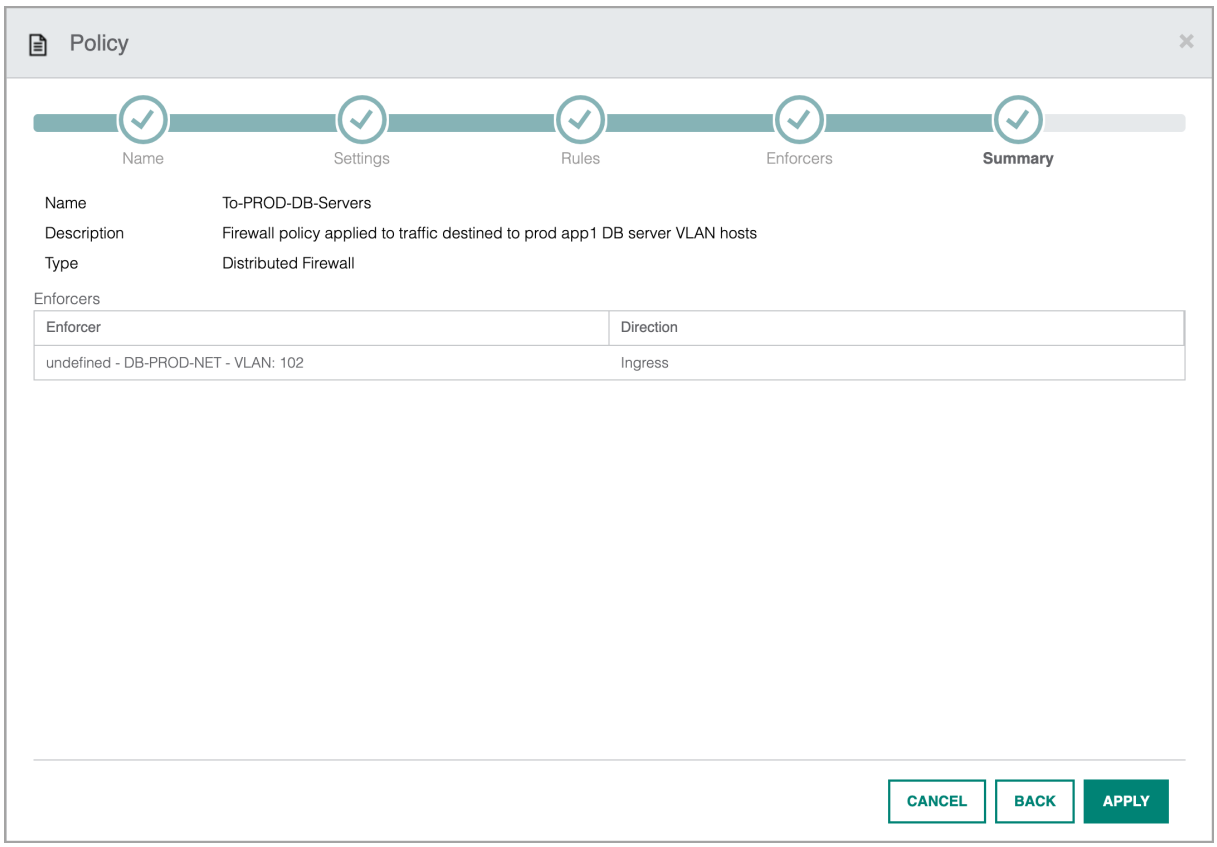

The frontend web application server to backend database server policy configuration is now complete.

## VM Microsegmentation

Microsegmentation enables the application of firewall policy between VM guests on the same hypervisor.

A PVLAN-based configuration forces all traffic from microsegmented VMs to the upstream CX 10000 switch for firewall policy inspection. Without the PVLAN setup, a hypervisor forwards traffic directly between hosts in the same VLAN. Egress and ingress policy can be applied.

The following diagram illustrates the networking components used to support a microsegmentation policy:

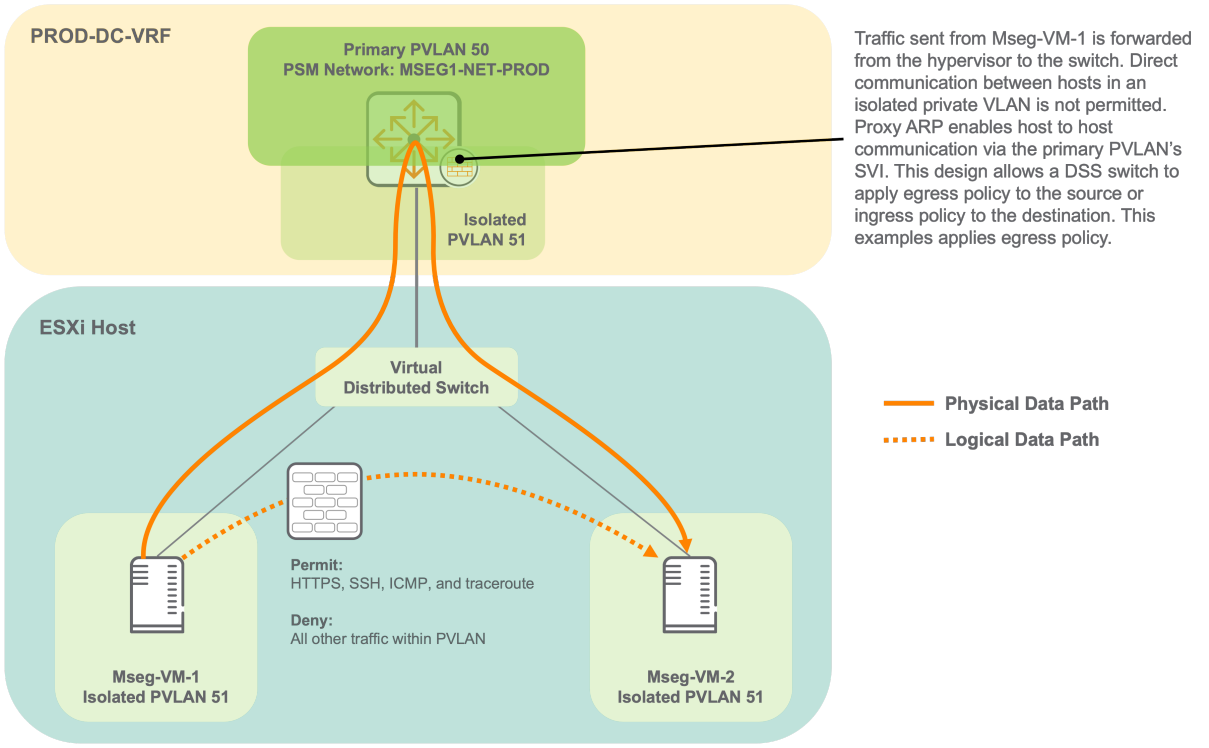


**Figure 7:** Microsegmentation Policy Diagram

In this procedure, egress firewall policy is applied between two VMs on the same ESXi host. Egress policy brings all traffic initiated by hosts in the VLAN into scope for policy rule creation.

AFC is used to create a virtual distributed switch (vDS) and microsegmentation PVLAN in vSphere. A corresponding PVLAN structure is created on CX 10000 switches attached to the ESXi host. An SVI is created with proxy ARP on the primary PVLAN to allow proxied communication between VMs placed in an isolated PVLAN, and a firewall policy is applied to the primary PVLAN *Network* to specify allowed traffic between hosts in the isolated PVLAN. In this example implementation, the microsegmentation is created across multiple CX 10000 switches.

The AFC microsegmentation workflow uses a series of wizards to complete microsegmentation set-up in a single fluid process. In addition to creating some of the vCenter configuration, the following related config components can be modified within the same workflow:

- PSM *Network* definition for the primary PVLAN
- VLAN SVI for the primary PVLAN
- CX 10000 MC-LAG configuration
- Policy endpoint groups
- Policy applications and service qualifiers
- Policy rules
- Firewall policy
- vCenter vDS and virtual port groups
- vCenter LAG/LACP configuration
- Assignment of distributed port groups to vCenter VMs

## Configure Policy Applications

Applications in AFC provide an administrator the flexibility to bundle multiple service qualifiers together into a set of protocols. The definition can represent a broad set of requirements for an application or logically group any collective set of protocols together for firewall or ACL policy purposes.

The following procedure defines an application that represents a set of network services that must be reachable from the microsegmented data center hosts.

**Step 1** In the AFC left navigation pane, click **Applications**.

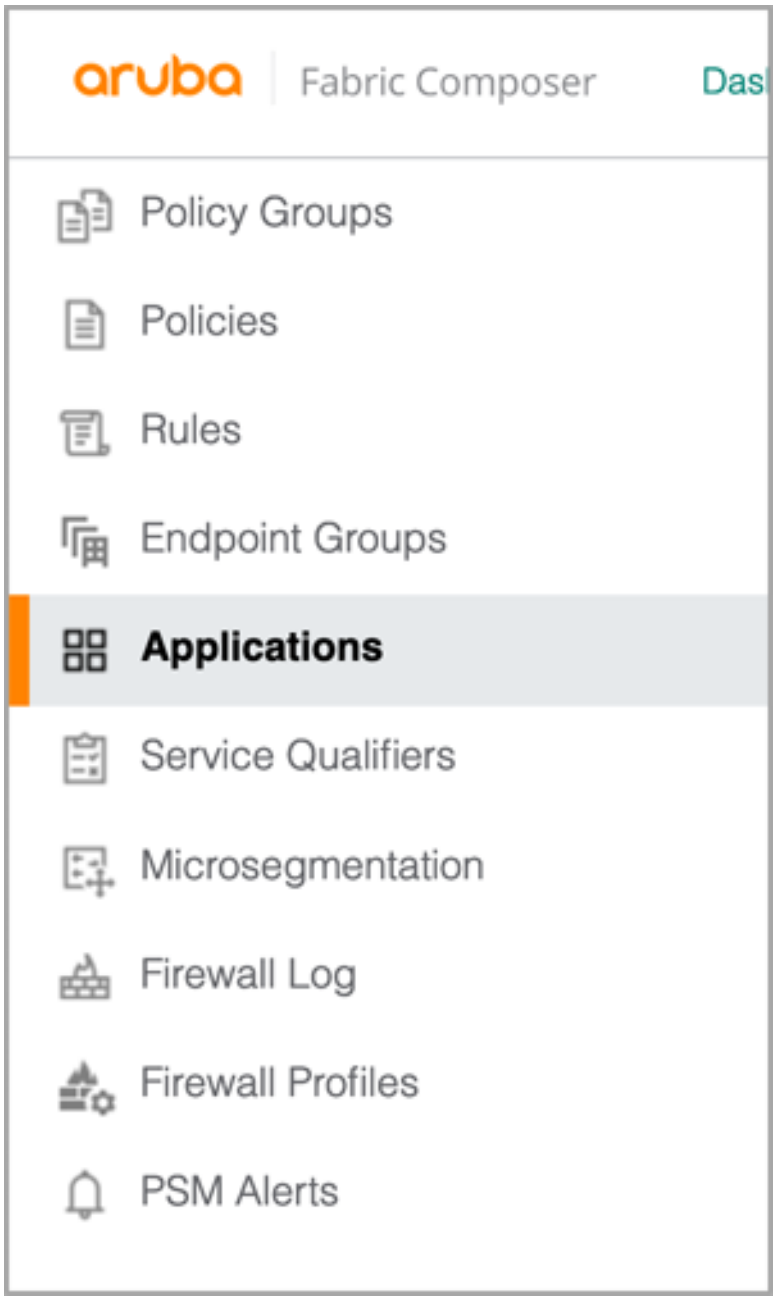


**Step 2** Click the right **ACTIONS** menu and select **Add**.

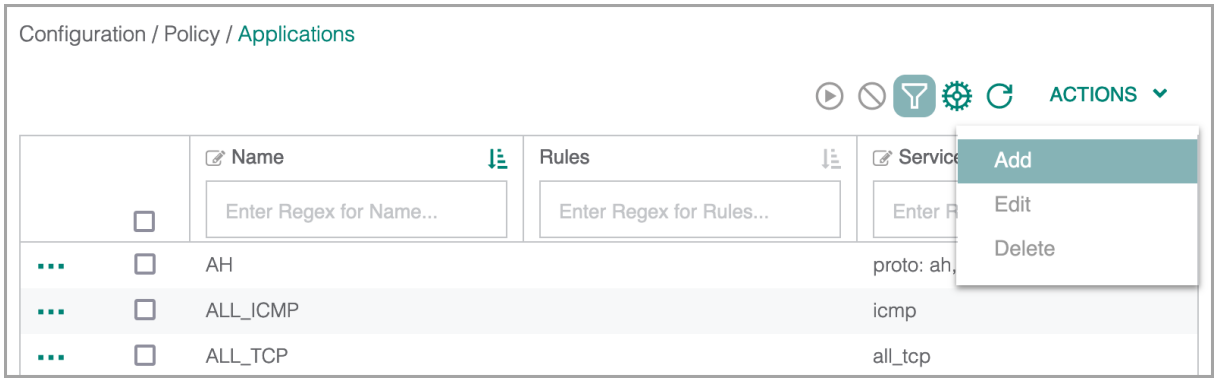


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

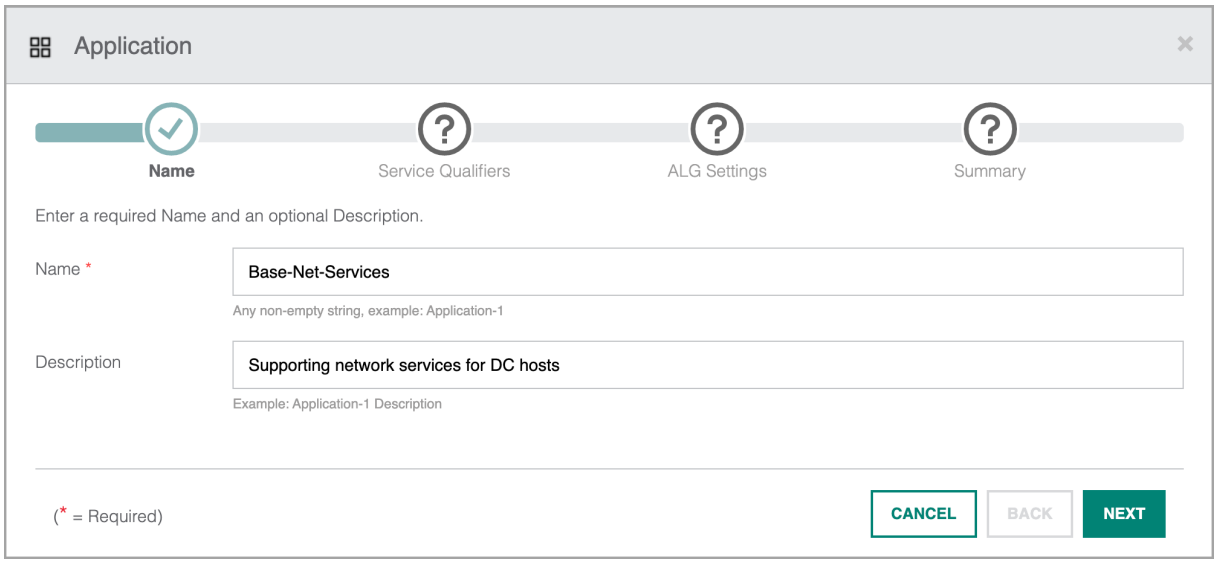


**Step 4** On the **Service Qualifiers** page, begin to enter the name of a service qualifier. A list of matching qualifiers is displayed. Press TAB or ENTER when a selection is highlighted to autocomplete the entry, or click the qualifier name in the list. Continue selecting service qualifiers until the following selections appear in the **Qualifiers** field, then click **NEXT**.

- syslog
- dns
- ntp
- radius_auth_udp
- radius_acct_udp
- ldaps

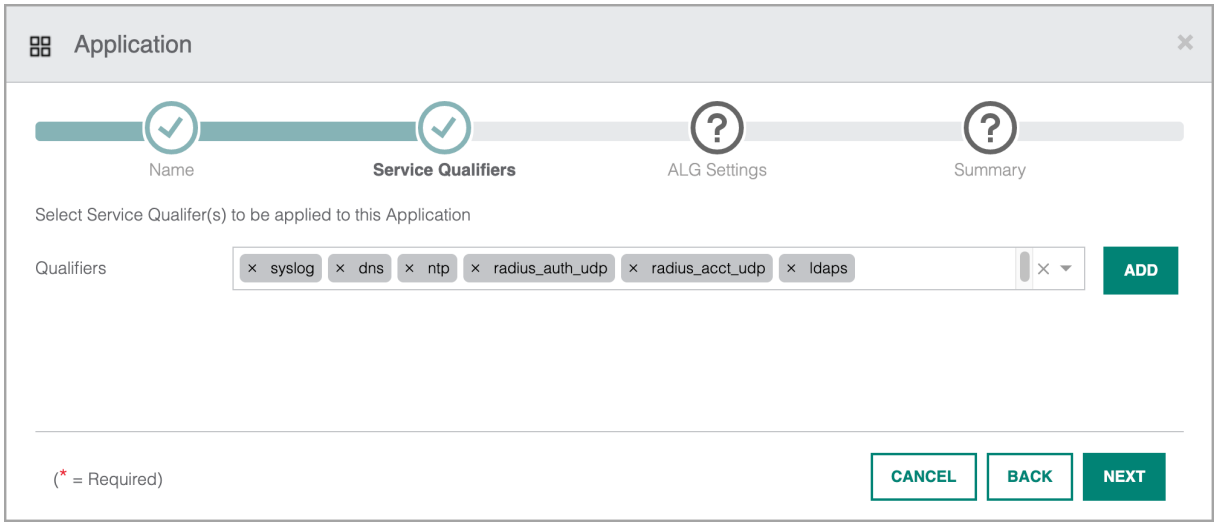


**NOTE:**

If a predefined service qualifier does not exist for the desired protocol and port combination, click **ADD** to create a new service qualifier.

**Step 5** On the **ALG Settings** page, leave **Type** set to the default value of *None* and click **NEXT**.

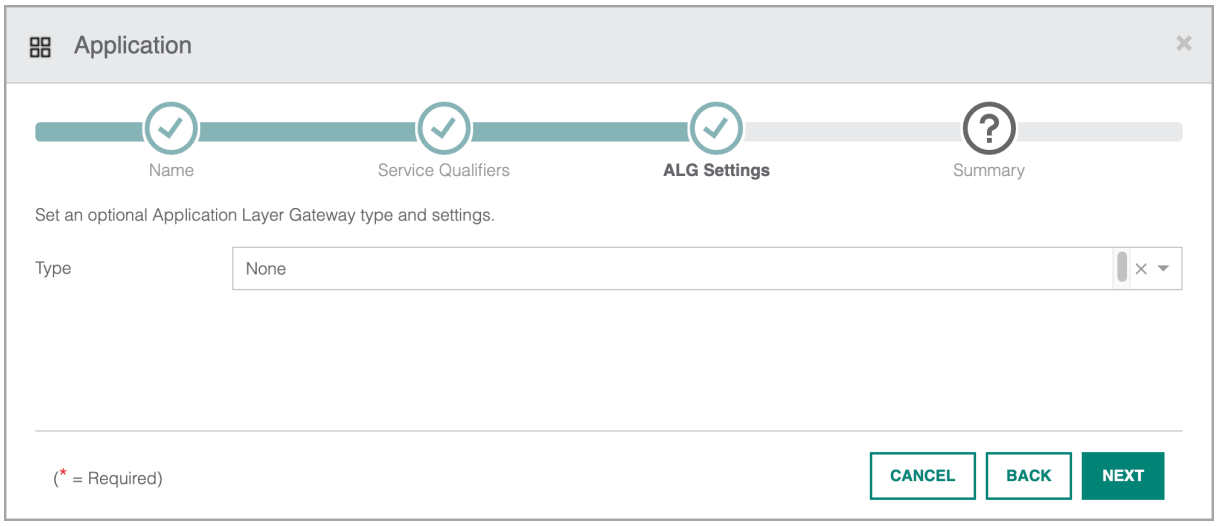


**Step 6** On the **Summary** page, verify that all information is correct and click **APPLY**.

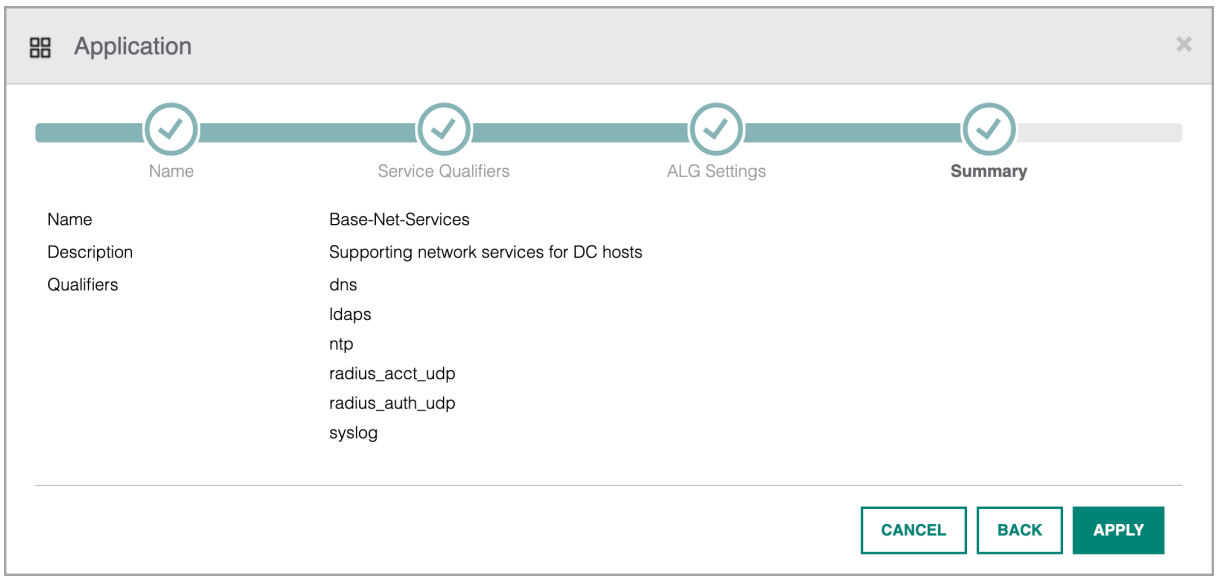


## Configure Microsegmentation Rules

Using the procedure in the Configure Firewall Rules section, create the firewall rules below. The AFC Microsegmentation wizard supports rule creation as a part of its workflow. Predefining rules simplifies the microsegmentation process.

| Name | Description | Type | Action | Endpoint Groups | Applications and Service Qualifiers |
|---|---|---|---|---|---|
| Allow-Base-Net-Services | Allow DC hosts to reach supporting network services | Layer 3 | Allow | **Source**: PROD-DC-VRF-Summary-EG **Destination**: Network-Services-EG | **Applications:** Base-Net-Services |

| Name | Description | Type | Action | Endpoint Groups | Applications and Service Qualifiers |
|---|---|---|---|---|---|
| Allow-MSEG1-HTTPS | Allow HTTPS between MSEG1 hosts | Layer 3 | Allow | **Source**: rsvdc-vcenter.example.local_rsvdc:mseg1 **Destination**: rsvdc-vcenter.example.local_rsvdc:mseg1 | **Service Qualifiers**: https |
| Allow-MSEG1-SSH | Allow SSH/SCP between MSEG1 hosts | Layer 3 | Allow | **Source**: rsvdc-vcenter.example.local_rsvdc: **Destination**: rsvdc-vcenter.example.local_rsvdc: | **Service Qualifiers**: ssh : |

**NOTE:**

Rules between microsegmented endpoints can be granular to an individual level. In the Allow-MSEG1 examples above, the same rule is applied to allow communication between the collective set of VMs assigned to the microsegmentation. It is best practice to apply policy to sets of hosts when possible to minimize the number of rules.In this example implementation, dynamic endpoint groups are used in microsegmentation firewall rules.

## ESXi Microsegmentation Using AFC

Microsegmentation uses PSM firewall policy to filter traffic between VMs installed on the same hypervisor, when attached to a CX 10000 switch.

The AFC **Create Microsegmentation** workflow automates microsegmentation in a vSphere environment.

The workflow below can be used for a new ESXi host deployment or on an existing host. The ESXi host must have at least one available VNIC that is not currently assigned to a virtual switch. Only unassigned VNICs are displayed and available in the wizard. If unused VNICs are not available on the target ESXi host, microsegmentation can be created by running appropriate AFC wizards individually after configuring PVLANs in vCenter.

**Step 1** On the **DISTRIBUTED SERVICES** tab of **Guided Setup**. Select **RSVDC-FB1/PROD-DC-VRF** in the **Selected VRF** field and click **CONFIGURE MICROSEGMENTATION**.

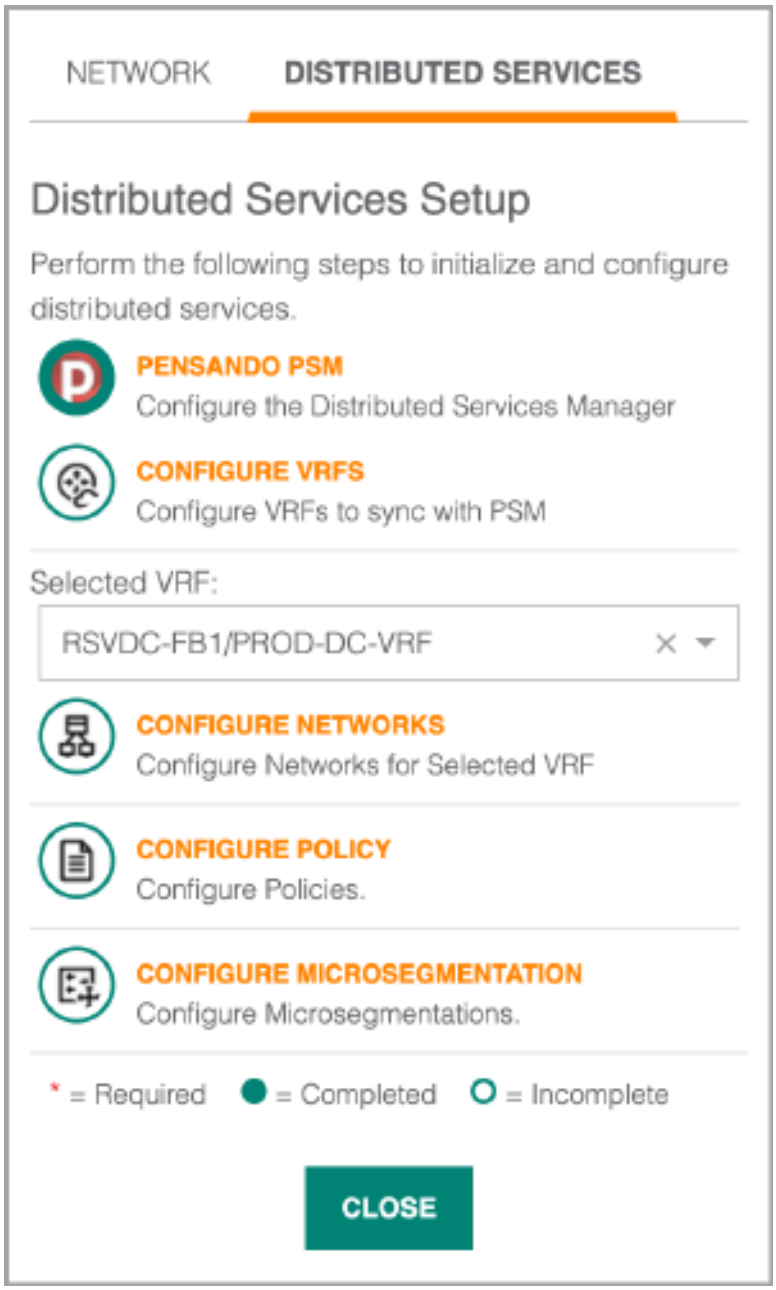


**Step 2** On the **Settings** page, enter the following values to select an ESXi host, then click **ADD**.

- **Name:** *MSEG1*
- **Host Addresses:** < *ESXi host target for microsegmentation* >
- **NICs:** *vmnic3 vmnic4*

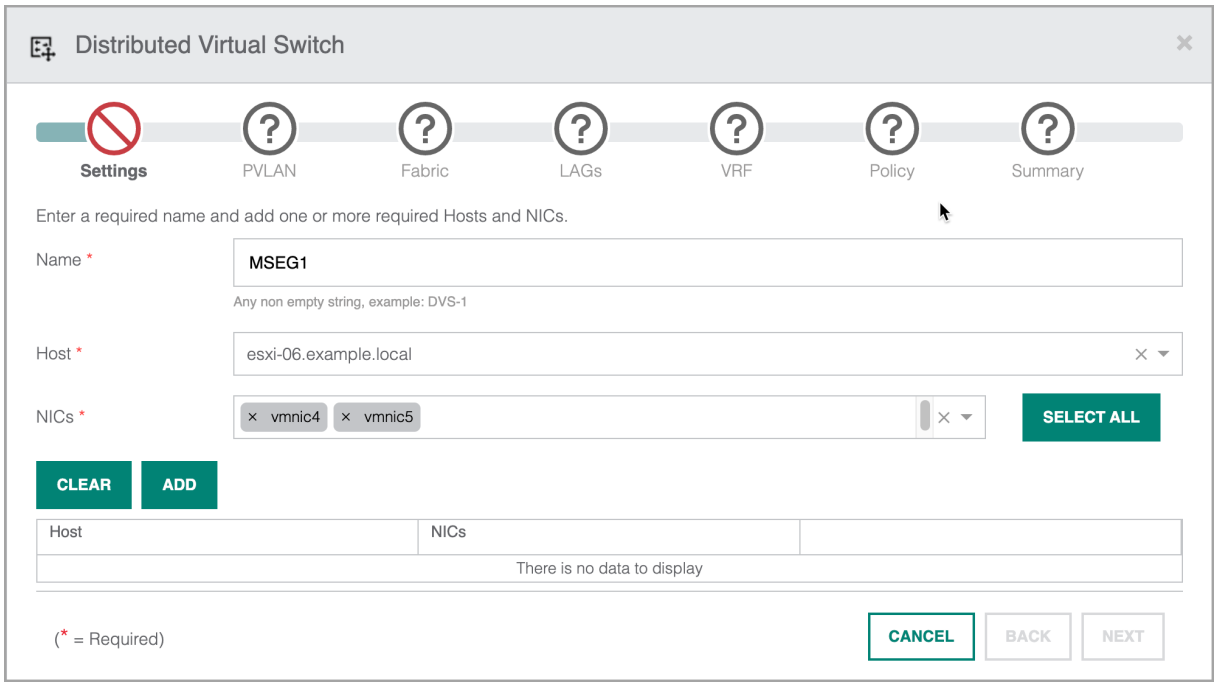


**Step 3** Repeat the step above to assign additional ESXi hosts to the microsegmentation instance. All assigned ESXi hosts must be attached to a CX 10000 switch.

**Step 4** Verify that all ESXi hosts were added to the list at the bottom of the dialogue box and click **NEXT**.

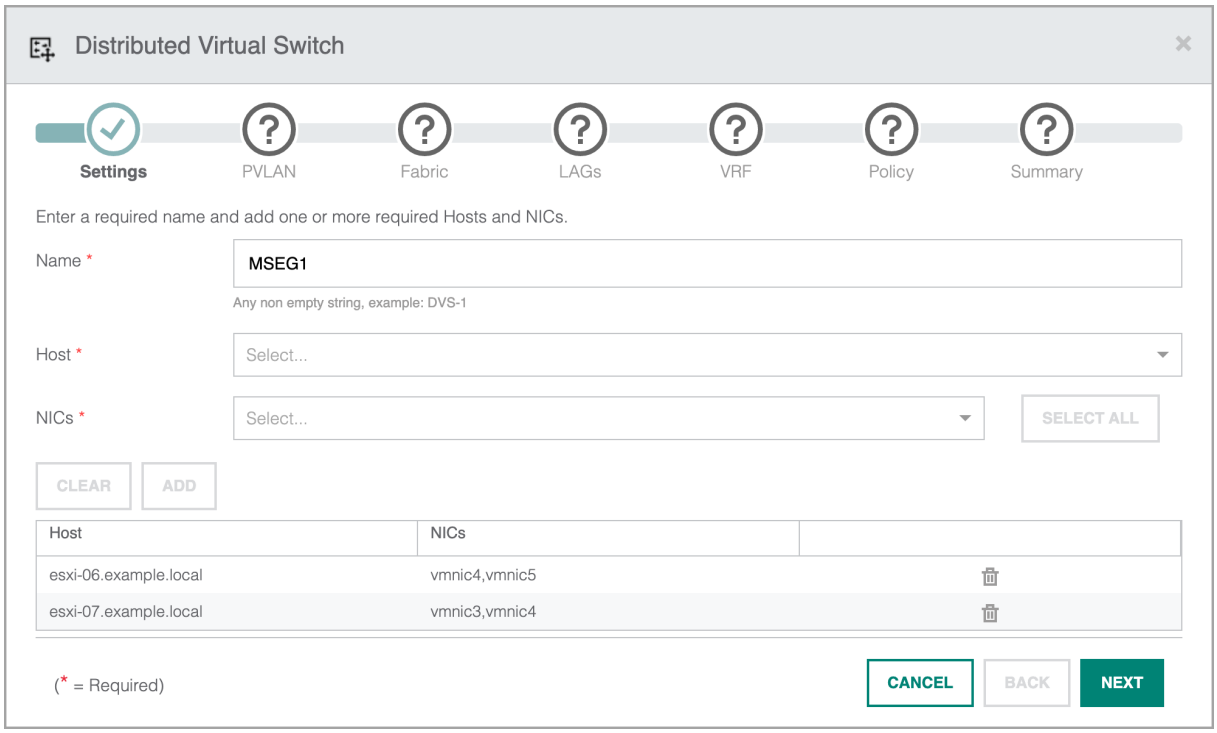


**Step 5** On the **PVLAN** page, enter the following values. Next to **Isolated VLAN VNICs**, click **SELECT VNICS**.

- **Portgroup Name Prefix:** *DPG*
- **Primary VLAN:** *50*
- **Isolated VLAN:** *51*

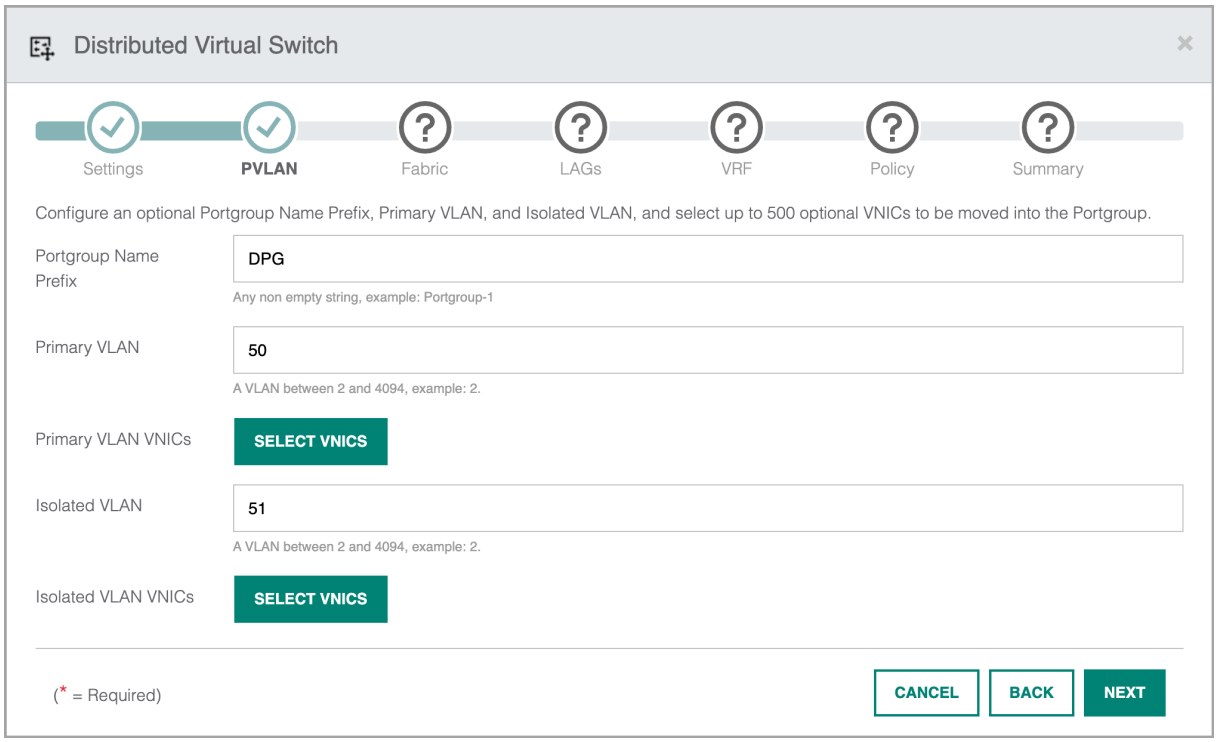


**Step 6** Check the VM virtual network adapters to be assigned to the PVLAN distributed port group, then click **APPLY**.

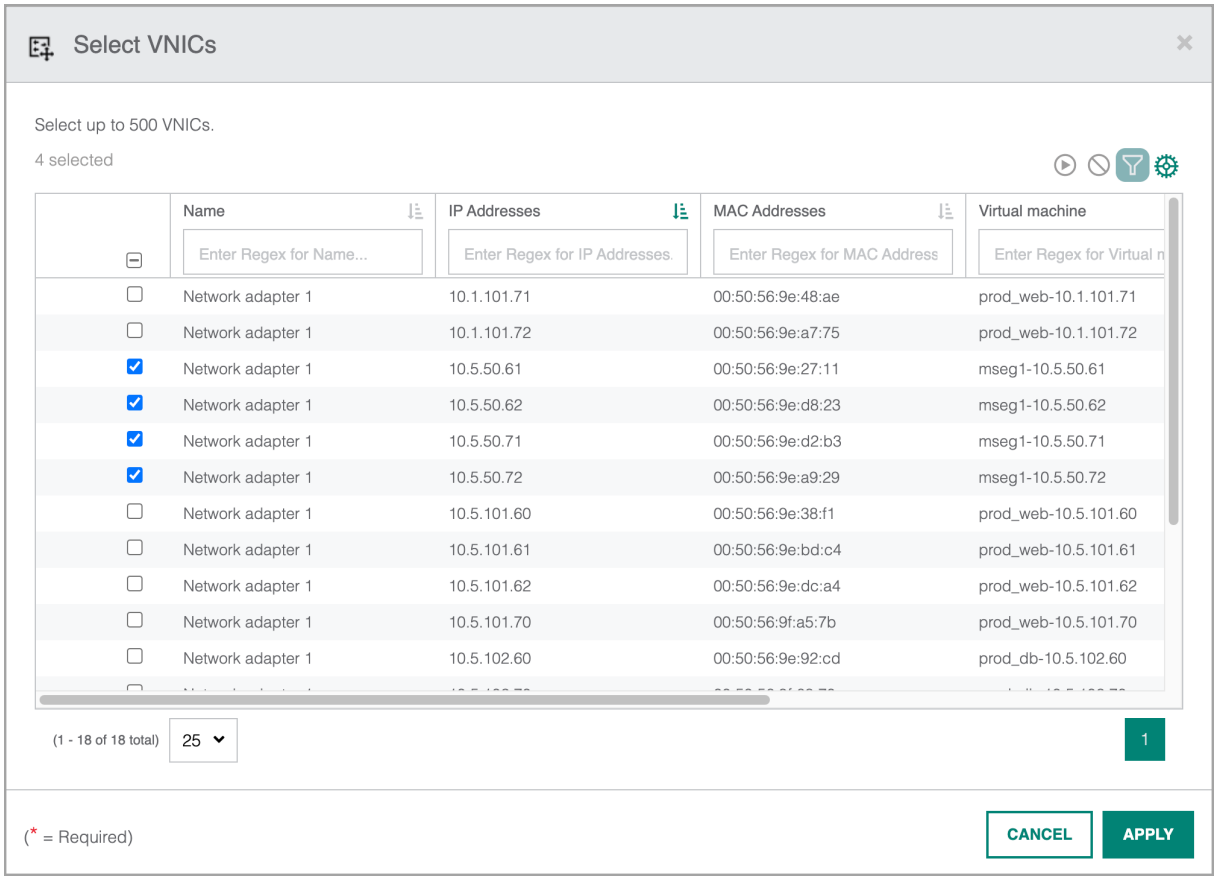

**NOTE:**

Each column supports filtering and sorting to manage the displayed VNICs in large VM deployments.

**Step 7** Verify that the number of **Isolated VLAN NICs** is correct, then click **NEXT**.

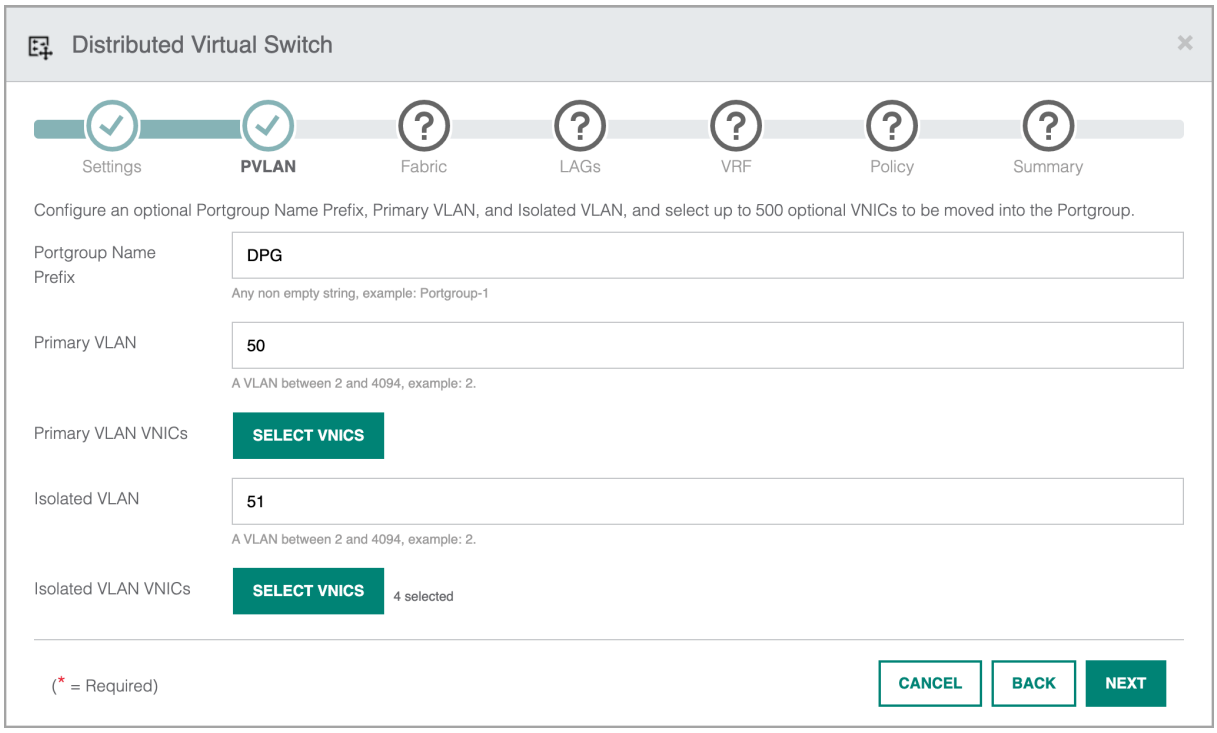


**Step 8** On the **Fabric** page, select the data center fabric and click **NEXT**.

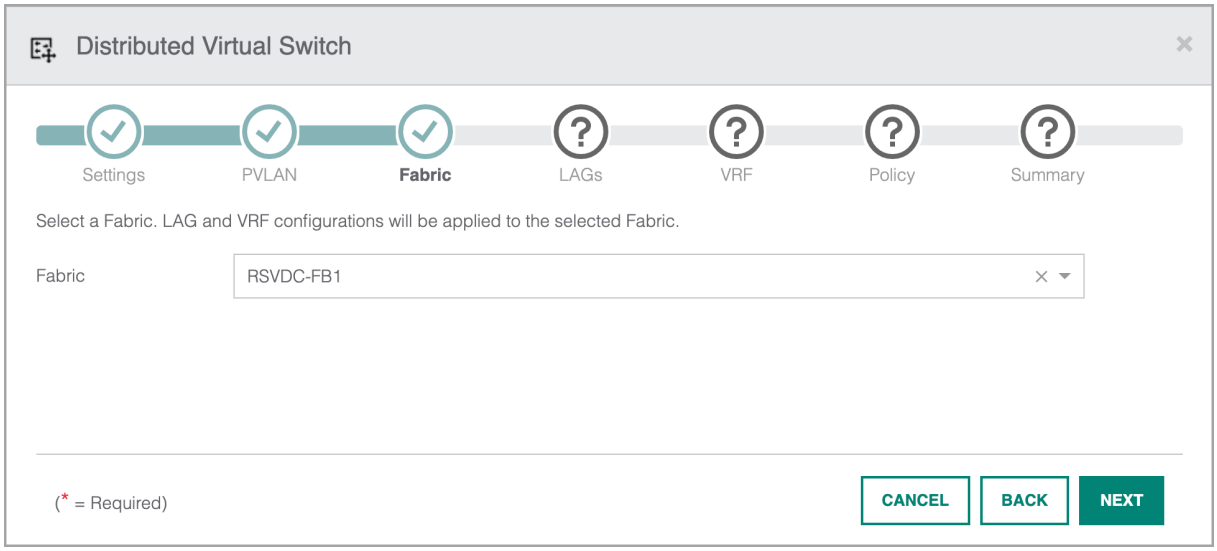


**Step 9** On the **LAGs** page, enter the following VM values. Next to **Switch LAG**, click **ADD**.

- **vSphere LAG Name**: *MSEG1*
- **Host:** *< ESXi host >*
- **Host NICs:** *< VM NICs connected to switch MC-LAG >*

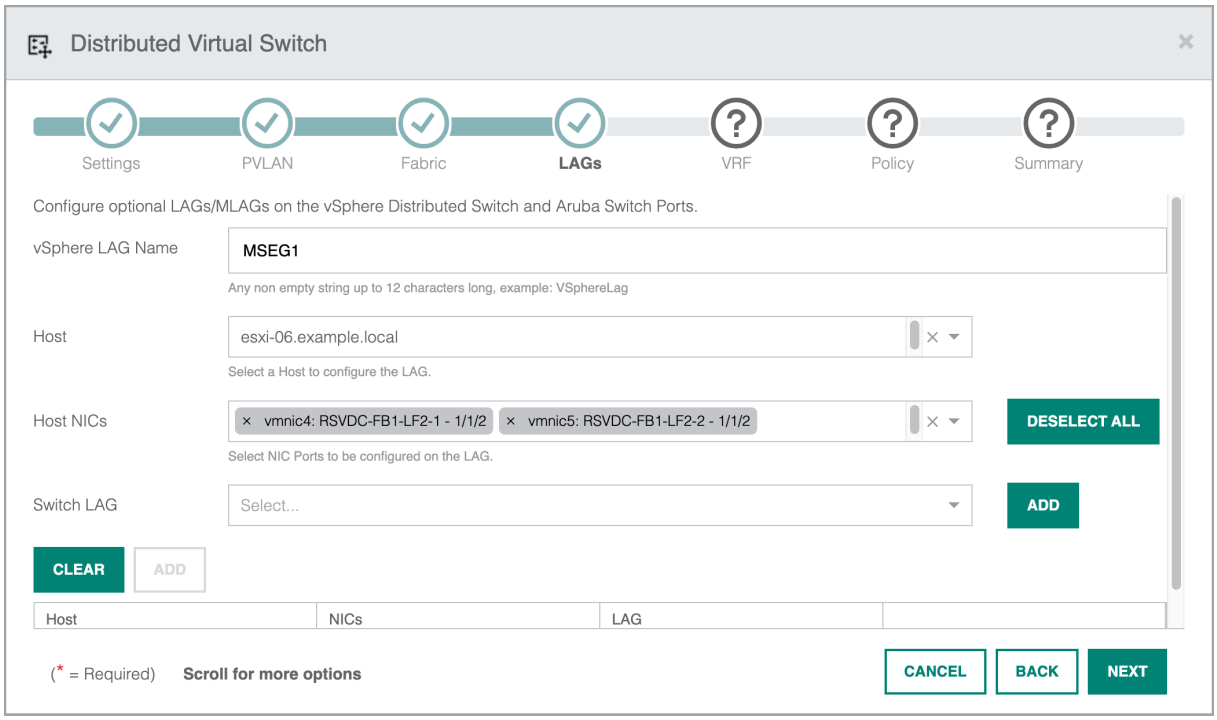


**Step 10 LAG** wizard: On the **Settings** page, enter a **Name**, **Description**, and **LAG Number**. Click **NEXT**.

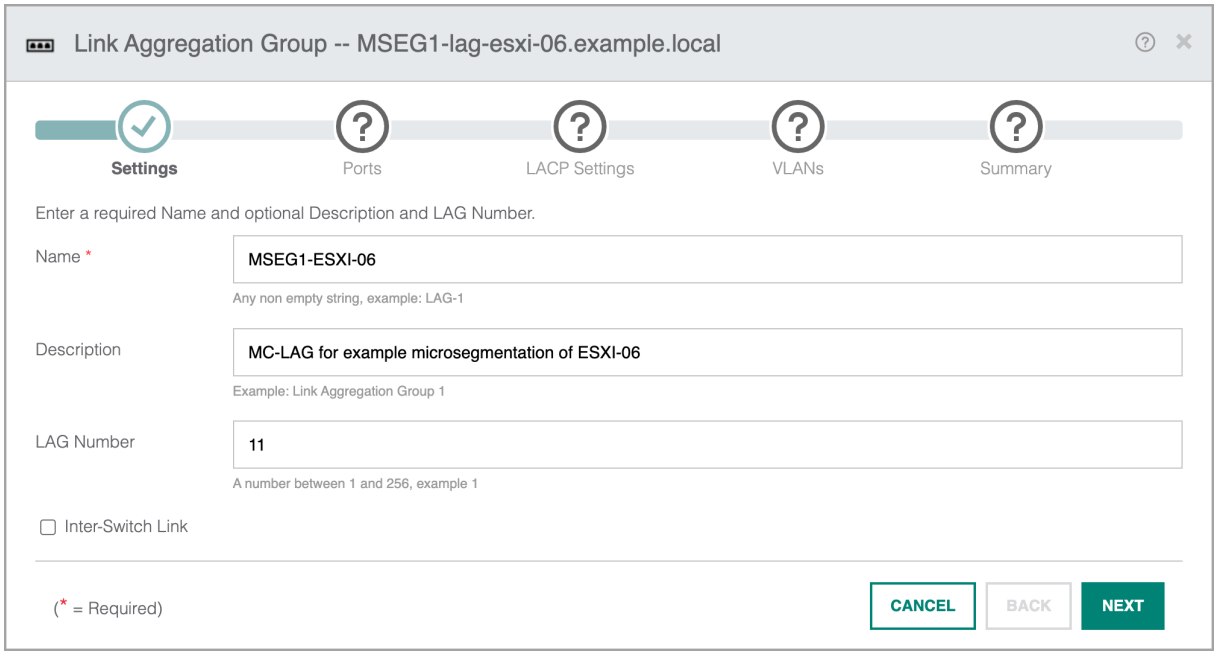


**Step 11 LAG** wizard: On the **Ports** page, select the switch corresponding to the ESXi host, verify that ports are preselected correctly, and click **NEXT**.

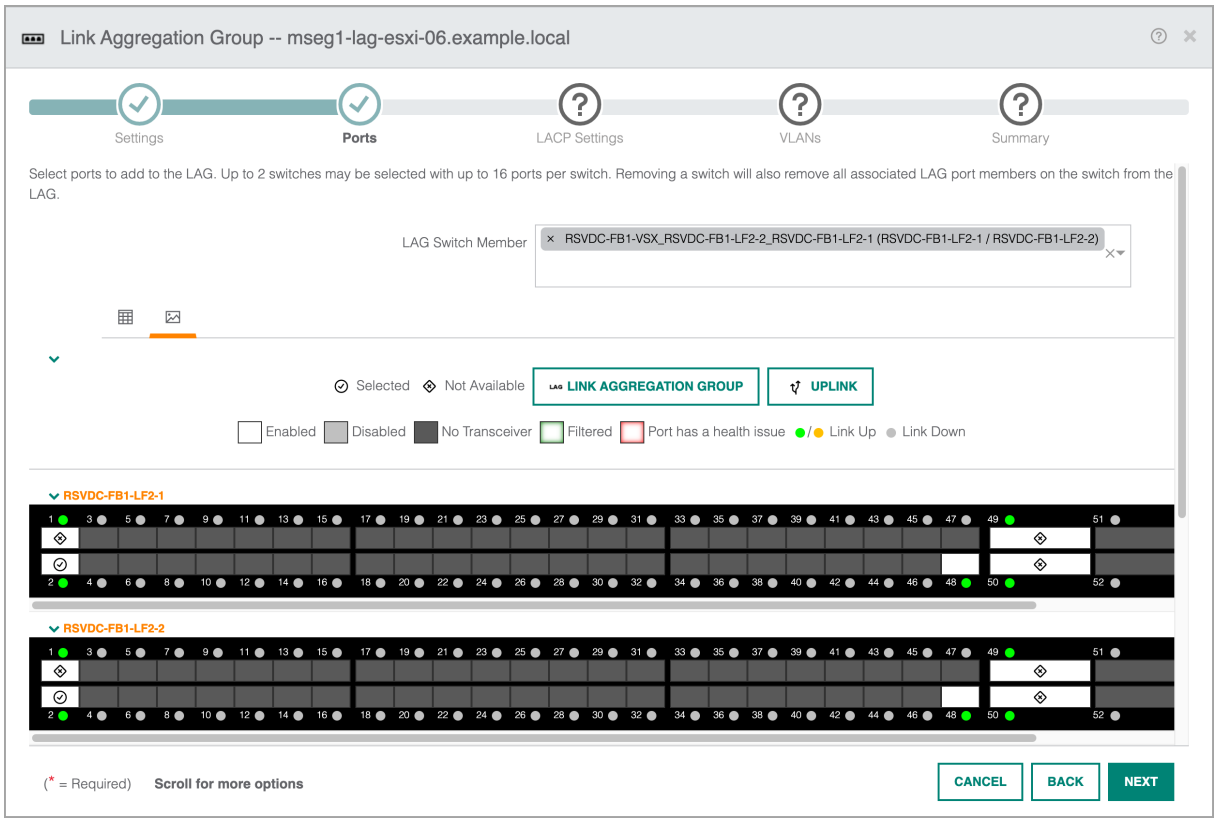


**Step 12 LAG** wizard: On the **LACP Settings** page, leave settings at their defaults and click **NEXT**.

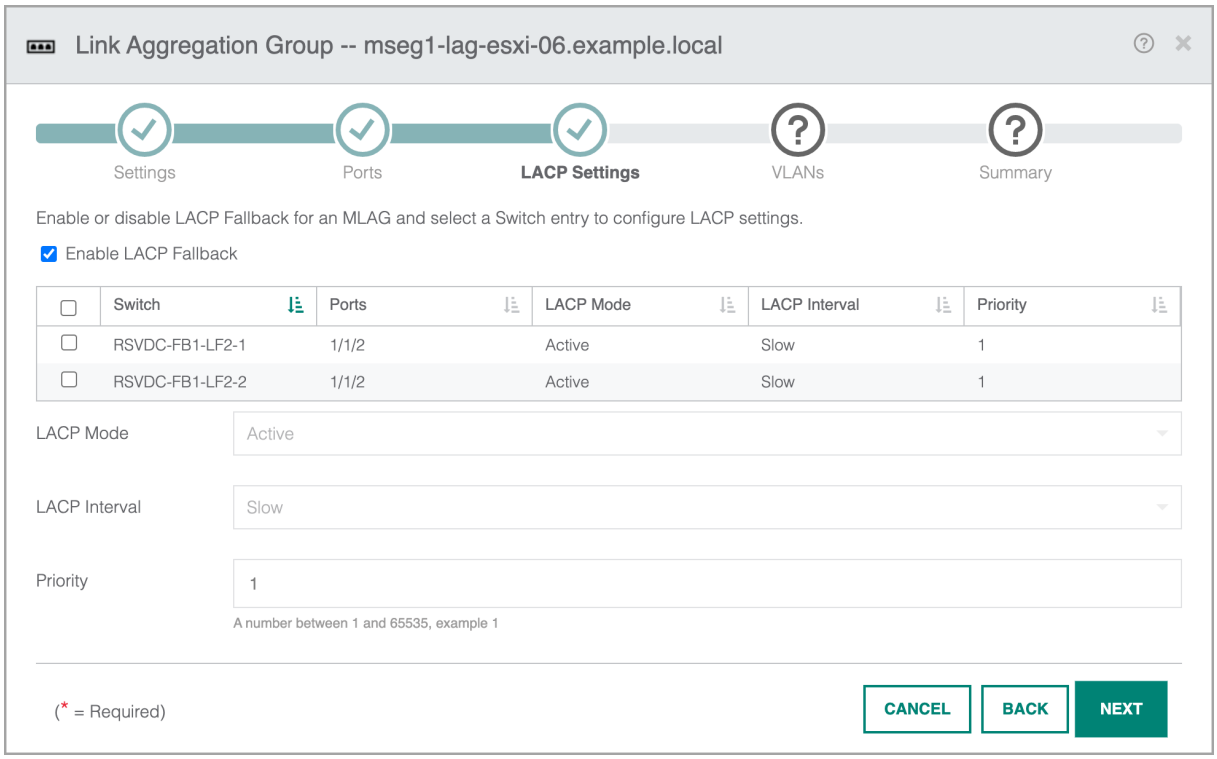


**Step 13 LAG** wizard: On the **VLANs** page, enter both primary and isolated VLAN IDs in the **VLANs** field, then click **NEXT**.

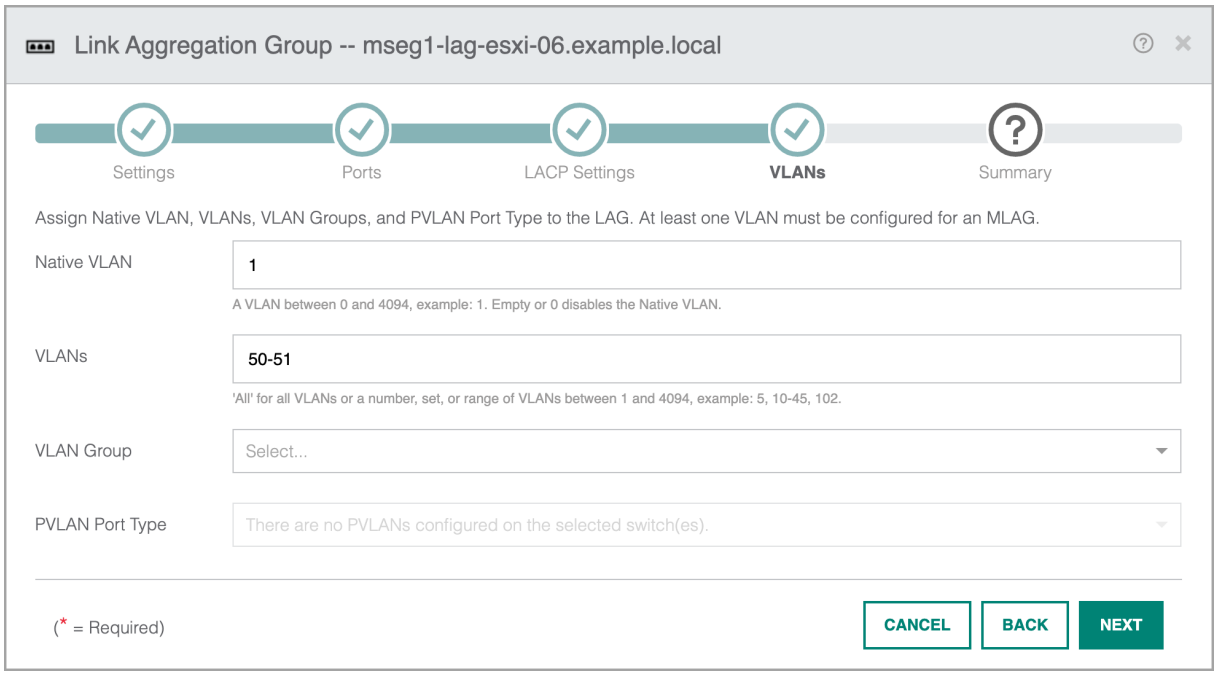


**Step 14 LAG** wizard: On the **Summary** page, verify that all information is correct and click **APPLY**. The wizard closes and returns to the main **Distributed Virtual Switch** workflow.

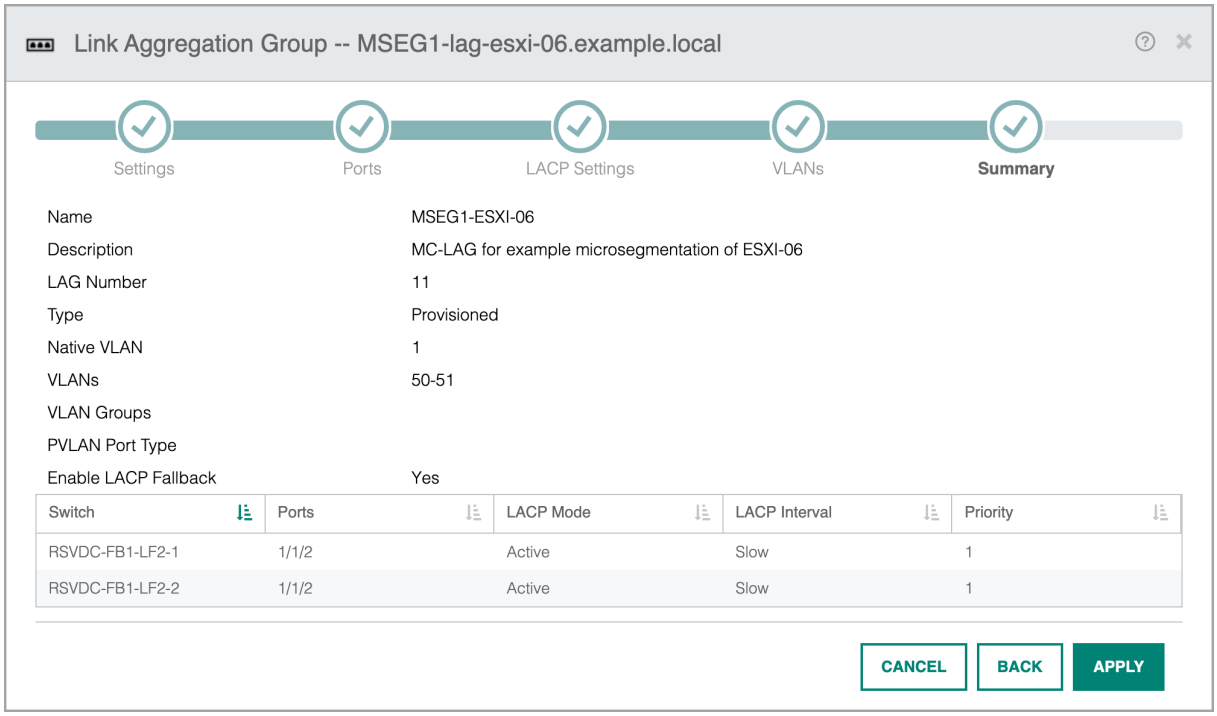


**Step 15** On the **LAGs** page, in the lower left, click **ADD**.

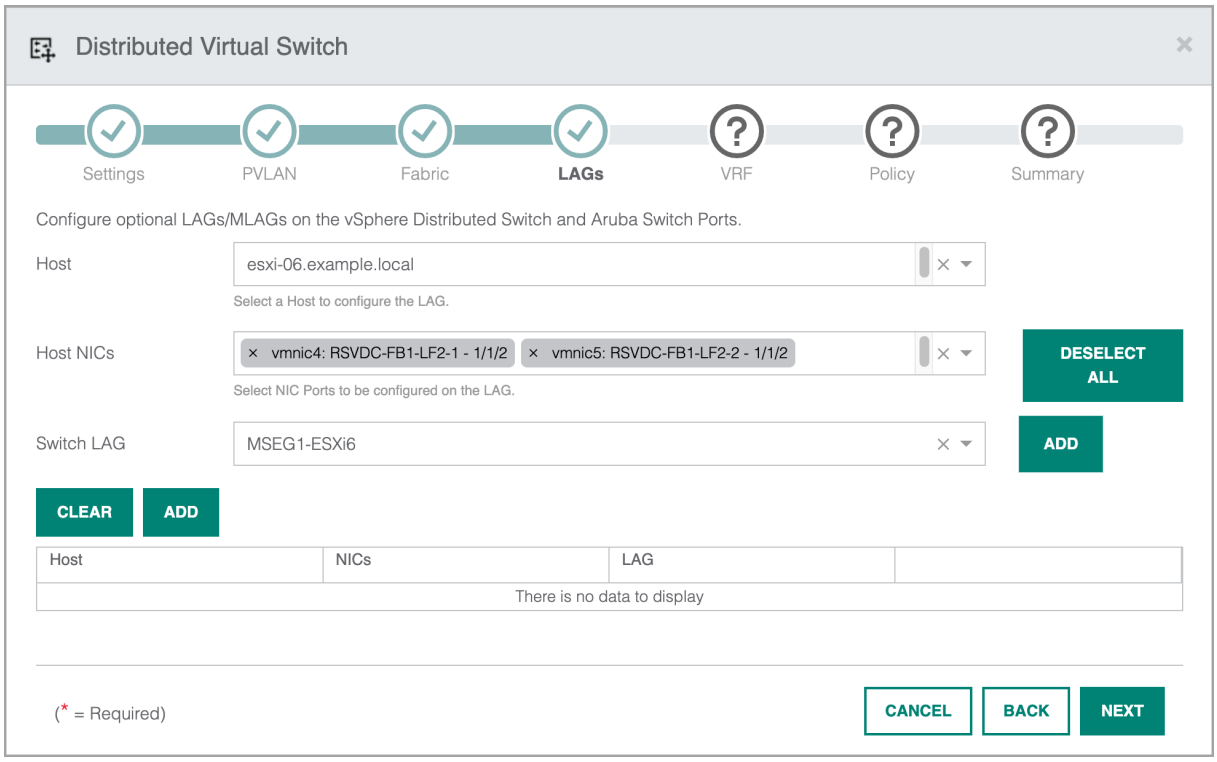


**Step 16** Repeat steps 9–15 to create additional MC-LAGs for each VM included in the microsegmentation, then click **NEXT**.

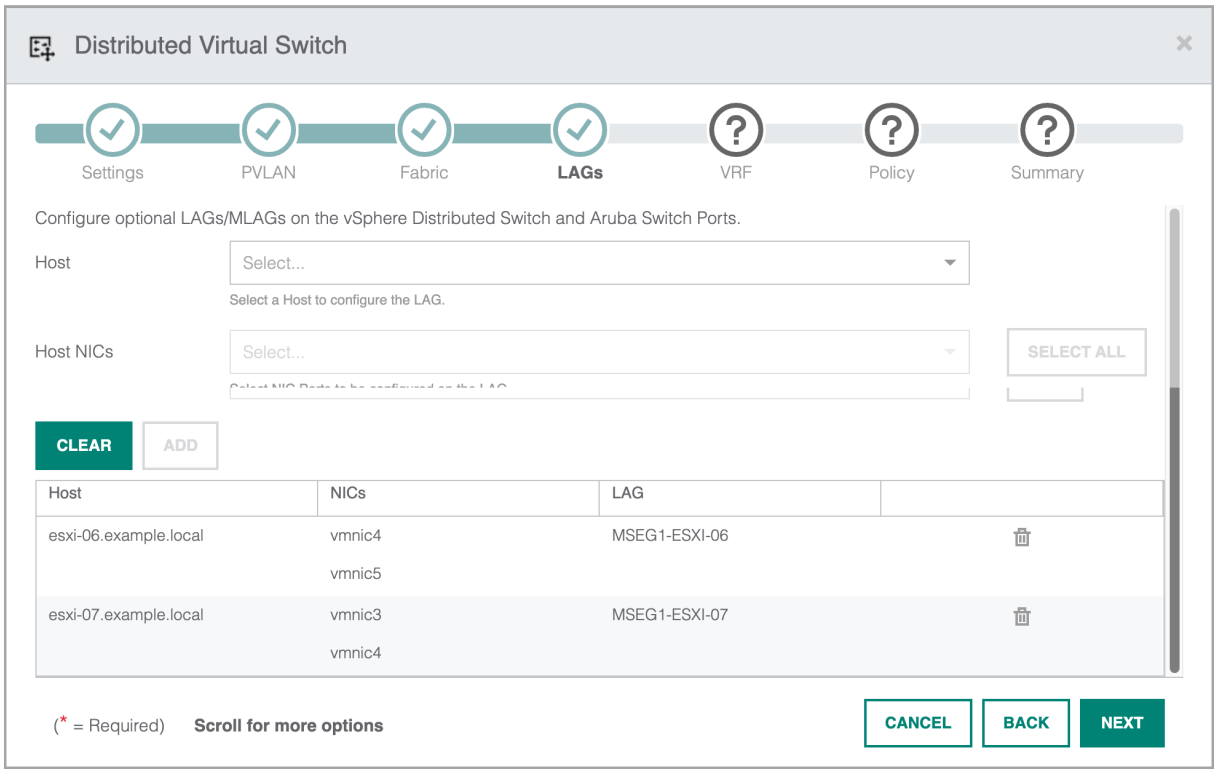


**Step 17** On the **VRF** page, select the VRF, then click **ADD** on the **Network** field to launch the **Network** wizard.

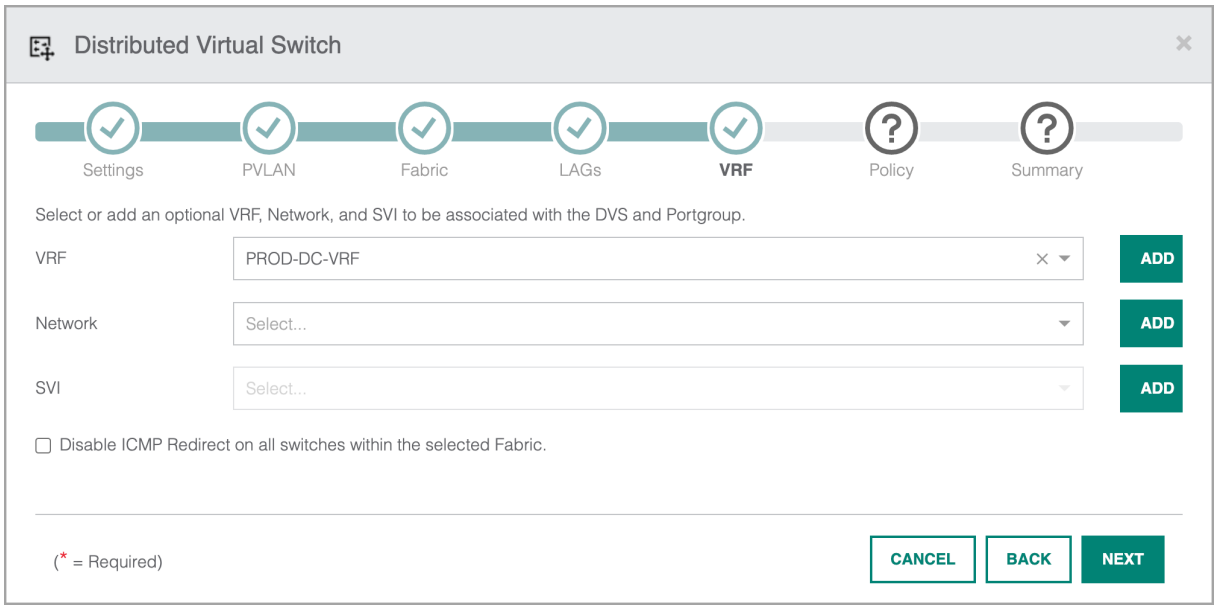


**Step 18 Network** wizard: On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

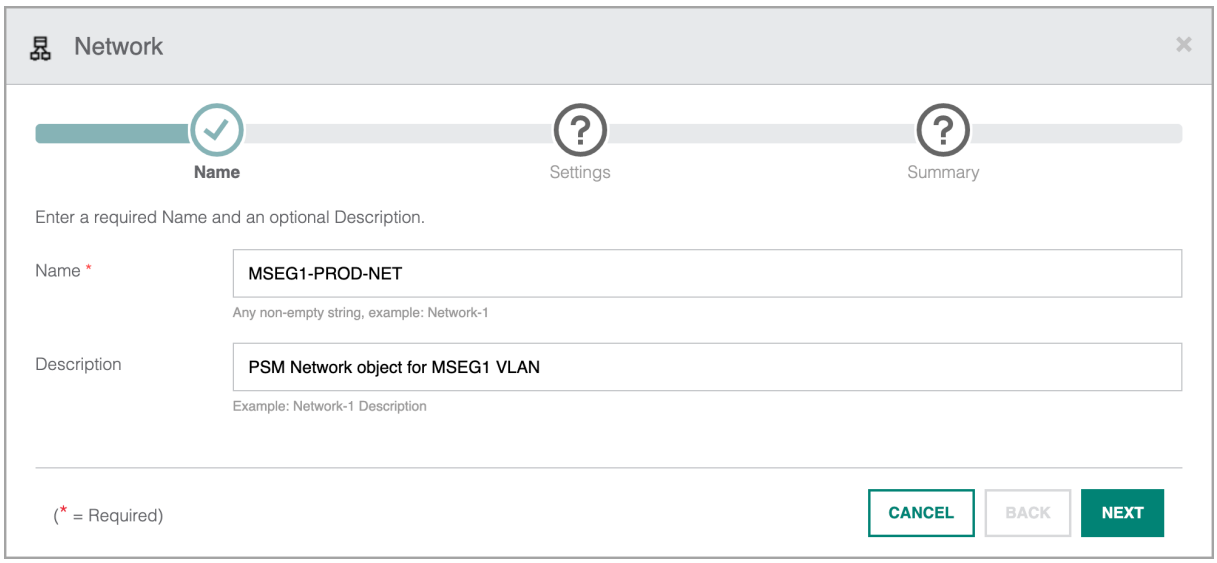


**Step 19 Network** wizard: On the **Settings** page, review the auto-selected primary PVLAN value and click **NEXT**.

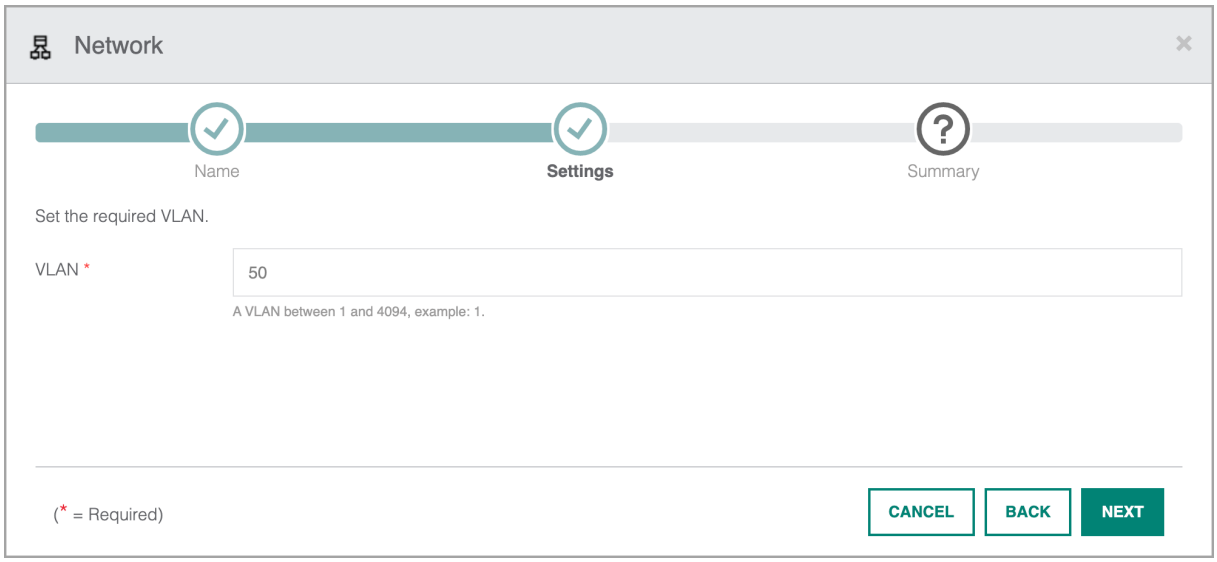


**NOTE:**

When creating a PSM Network within the microsegmentation workflow, the VLAN is automatically set to the primary PVLAN ID. This value cannot be modified. If a different VLAN ID is present, click **CANCEL** to exit the **Network** wizard, and click **BACK** to the **PVLANs** page for review.

**Step 20 Network** wizard: On the **Summary** page, verify that the settings are correct and click **APPLY**. The wizard closes and returns to the main **Distributed Virtual Switch** workflow.

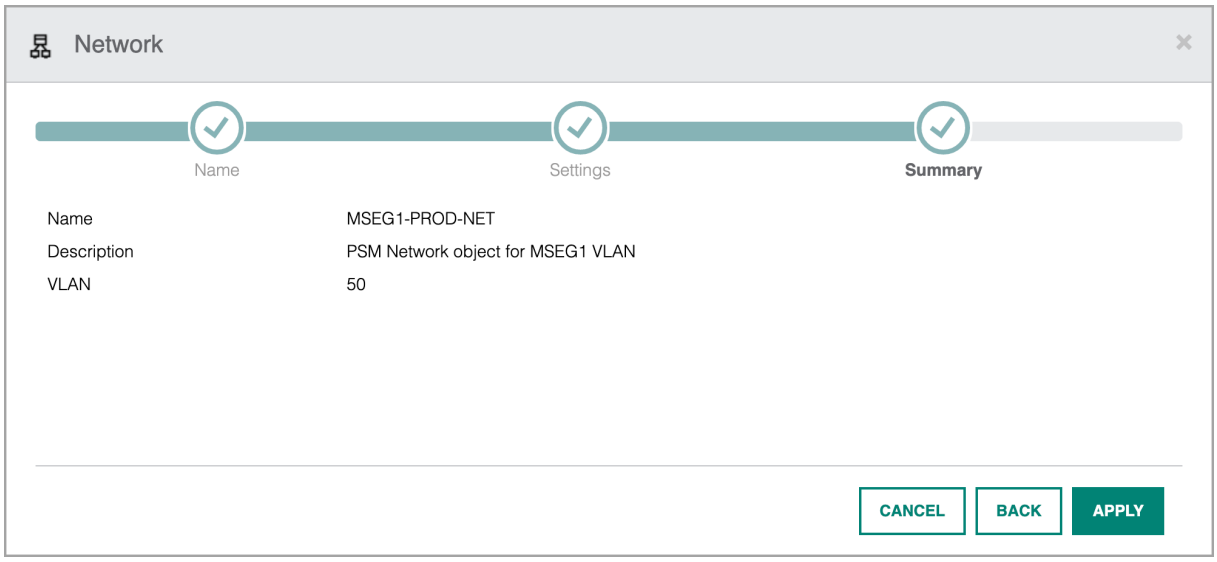


**Step 21** Verify that the new PSM *Network* is populated in the **Network Field**, then click **ADD** on the **SVI** field to launch the **IP Interface** wizard.

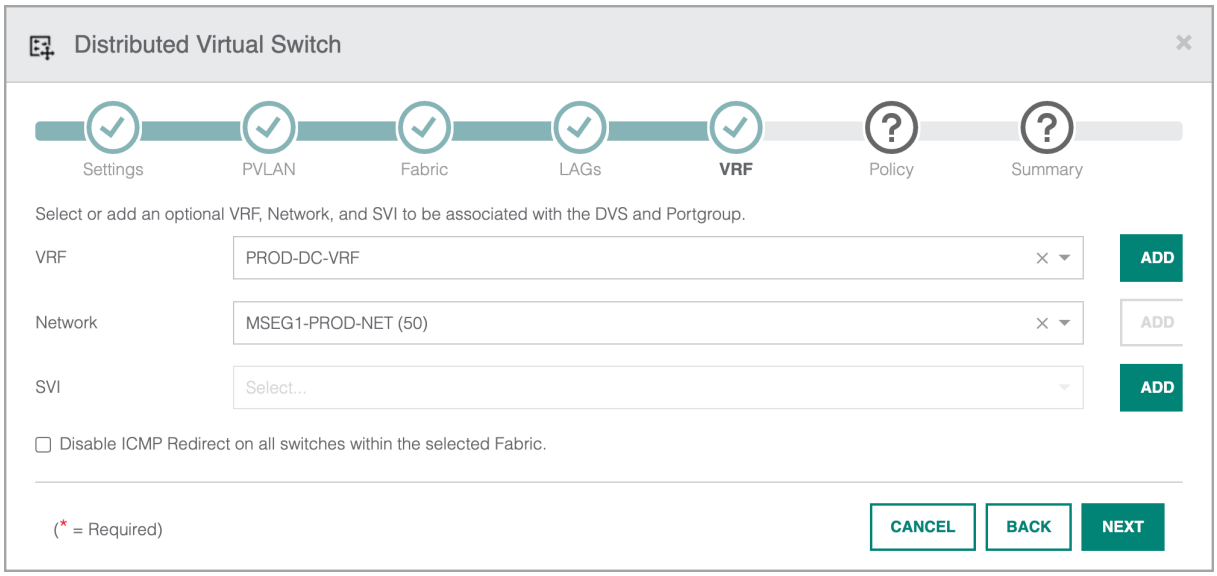


**Step 22 IP Interface** wizard: On the **Interface Type** page, review the auto-populated values for **Type**, **VLAN**, and **Enable Local Proxy ARP** fields. Enter the following values, then click **NEXT**.

- **Switches:** *< CX 10000 switches attached to the ESXi microsegmentation >*
- **IPv4 Subnetwork Address:** *10.5.50.0/24*
- **IPv4 Addresses:** *10.5.50.1*
- **Active Gateway IP Address:** *10.5.50.1*
- **Active Gateway IP MAC Address:** *02:00:0A:05:00:01*

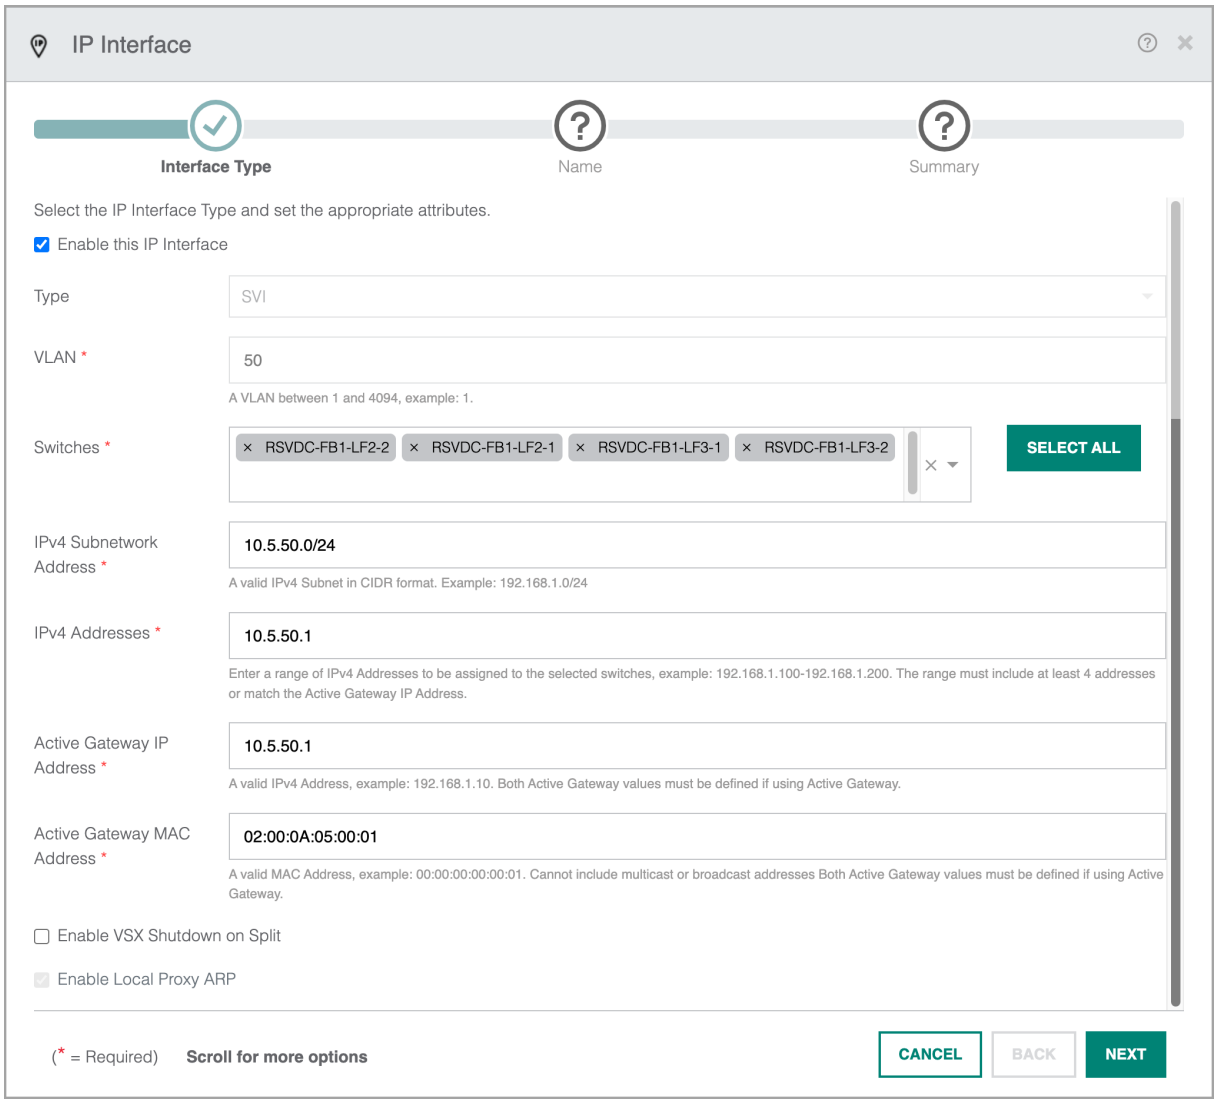


**NOTE:**

**Enable Local Proxy ARP** is checked by the wizard and cannot be modified. Proxy ARP is required to enable communication between VMs assigned to the same isolated PVLAN.

**Step 23 IP Interface** wizard: On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

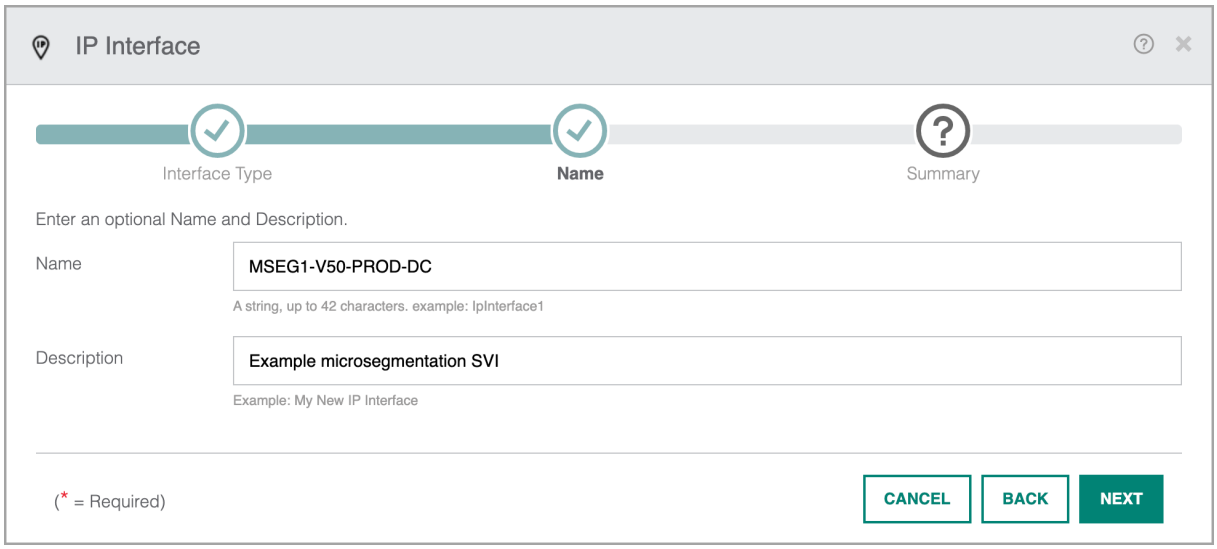


**Step 24 IP Interface** wizard: On the **Summary** page, verify that all information is correct, then click **APPLY**. The wizard closes and returns to the main **Distributed Virtual Switch** workflow.

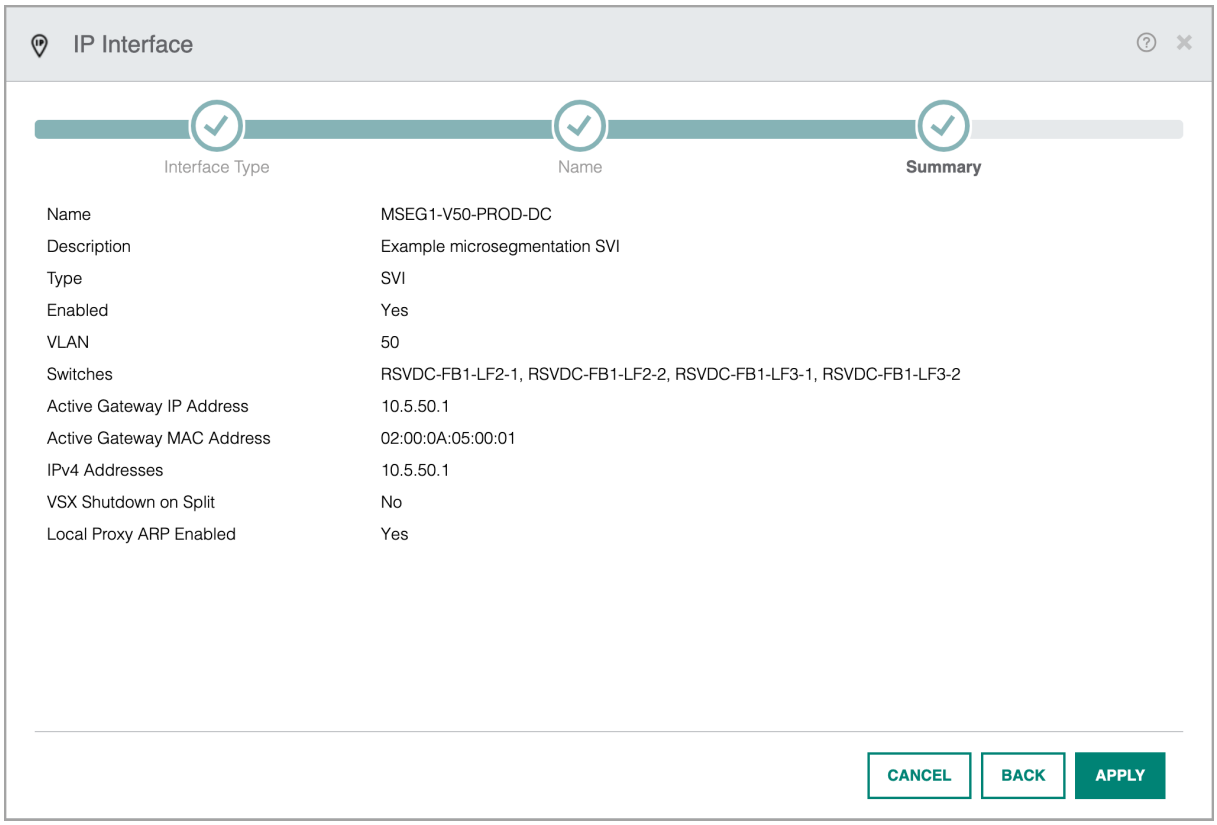


**Step 25** Verify that the new SVI was populated in the **SVI** field and click **NEXT** to proceed.

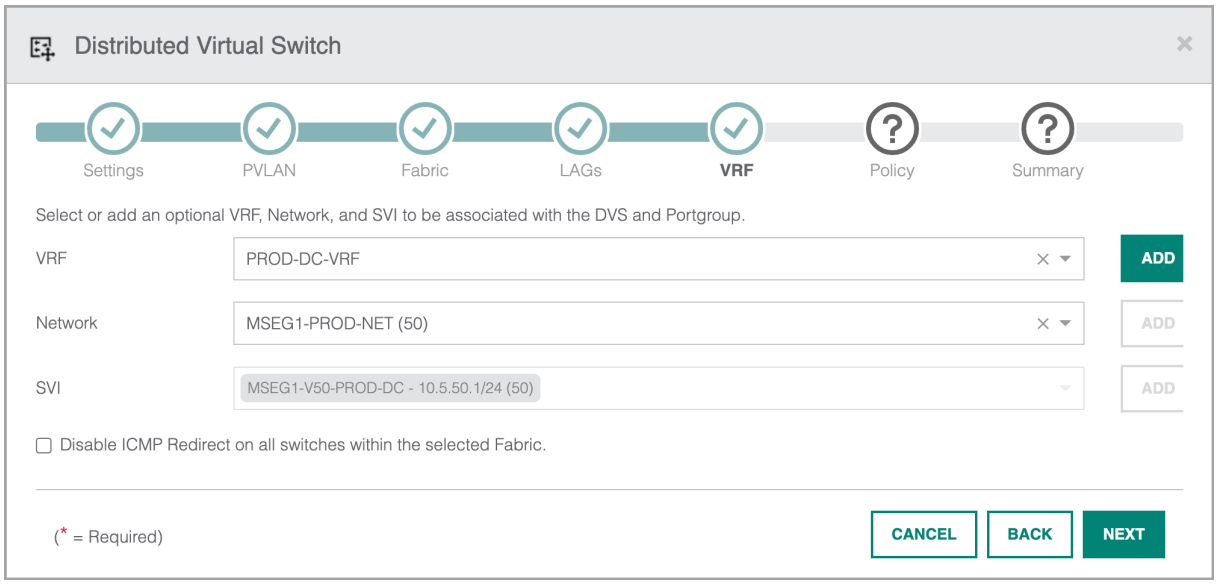


**Step 26** On the **Policy** page, click **ADD** in the **Policies** field to start the **Policy** wizard.

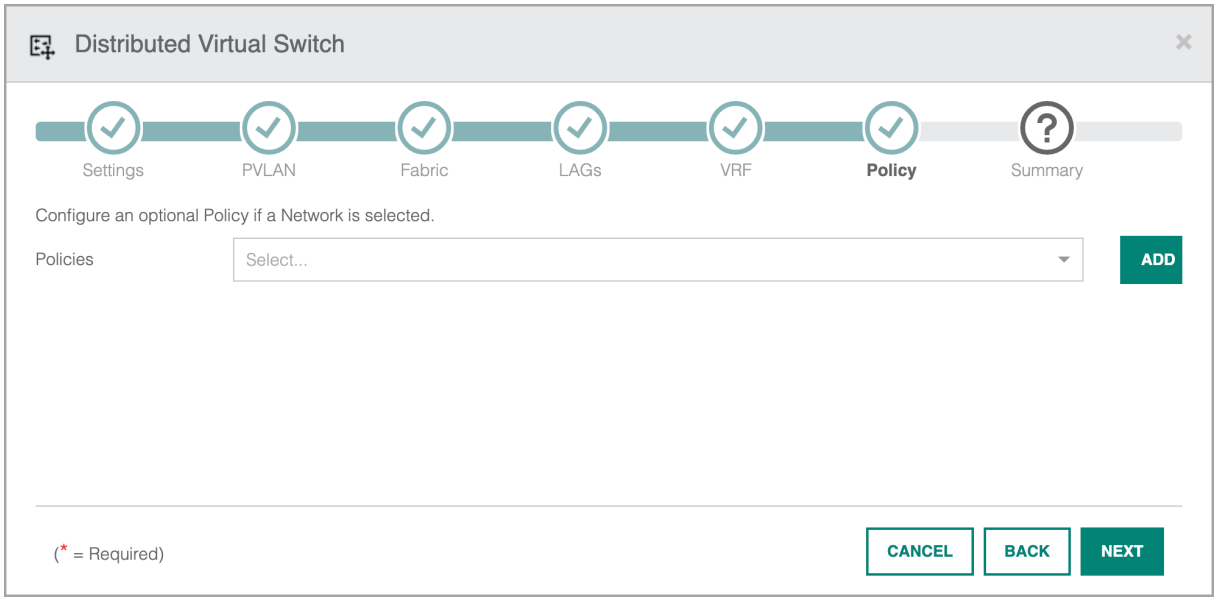


**Step 27 Policy** wizard: On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

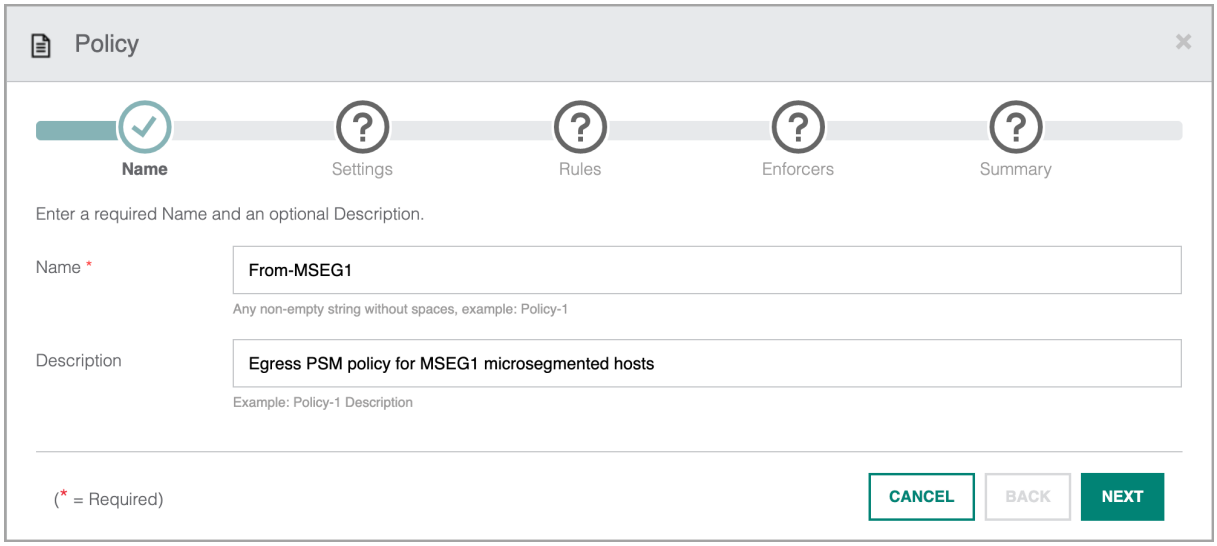

**Step 28 Policy** wizard: On the **Settings** page, verify that the prepopulated **Type** field value is set to *Distributed Firewall* and click **NEXT**.

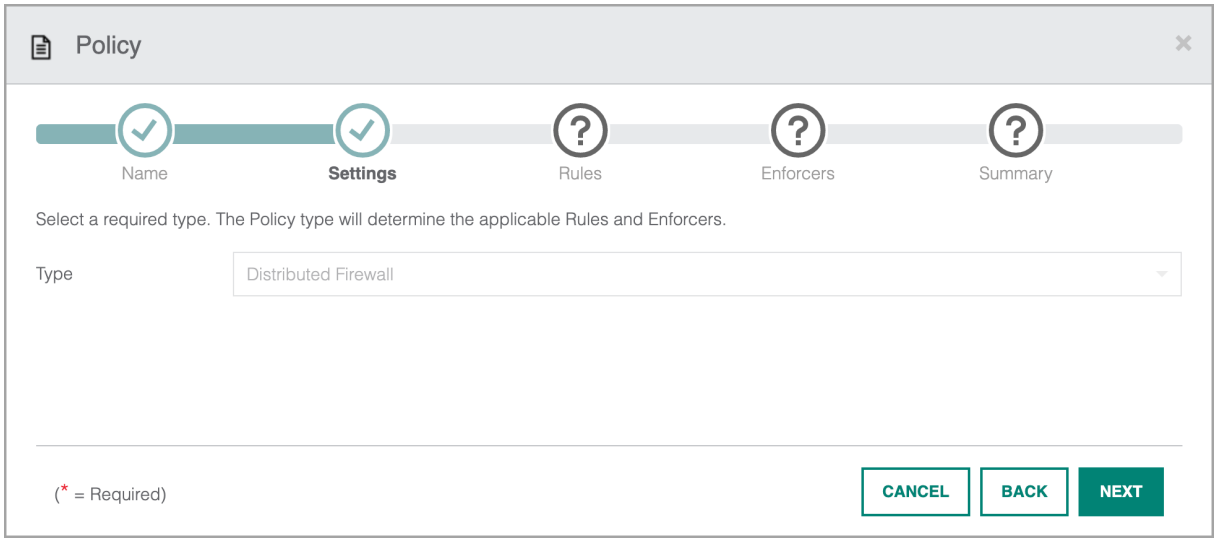


**Step 29 Policy** wizard: Click **ACTIONS**, then select **Add > Existing**.

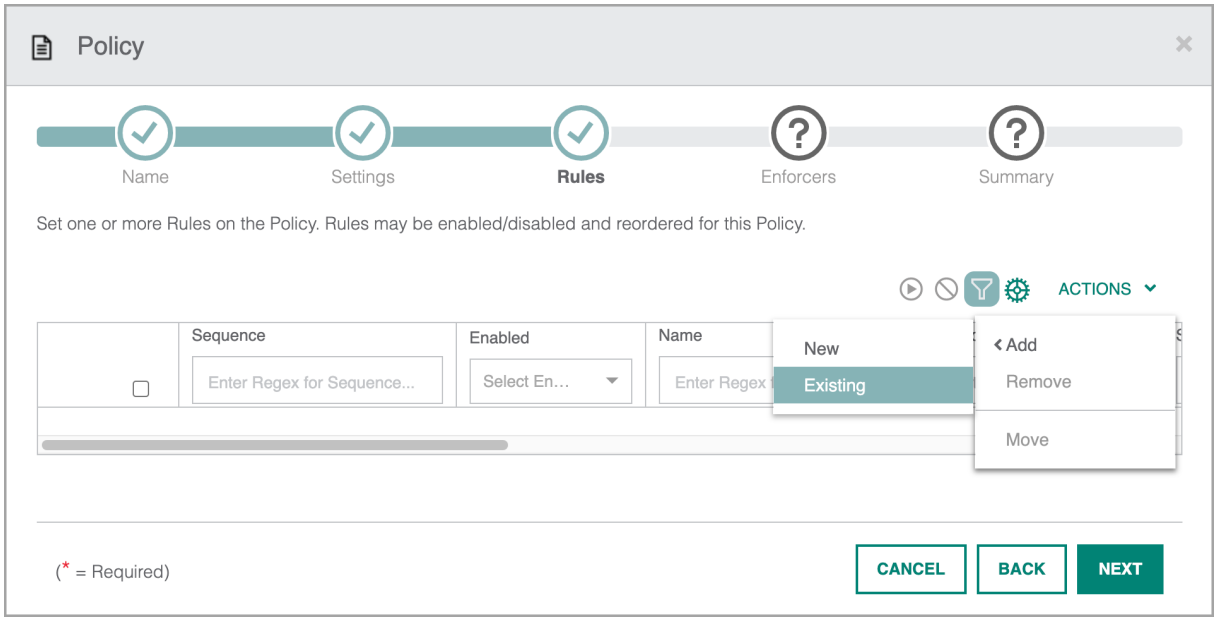


**Step 30 Policy** wizard. In the **Select Rules** window, select the checkbox for existing rules in the following order, and click **APPLY**.

- *Allow-MSEG1-HTTPS*
- *Allow-MSEG1-SSH*
- *Allow-Base-Net-Services*
- *Allow-All-ICMP*
- *Allow-All-UDP-Traceroute*

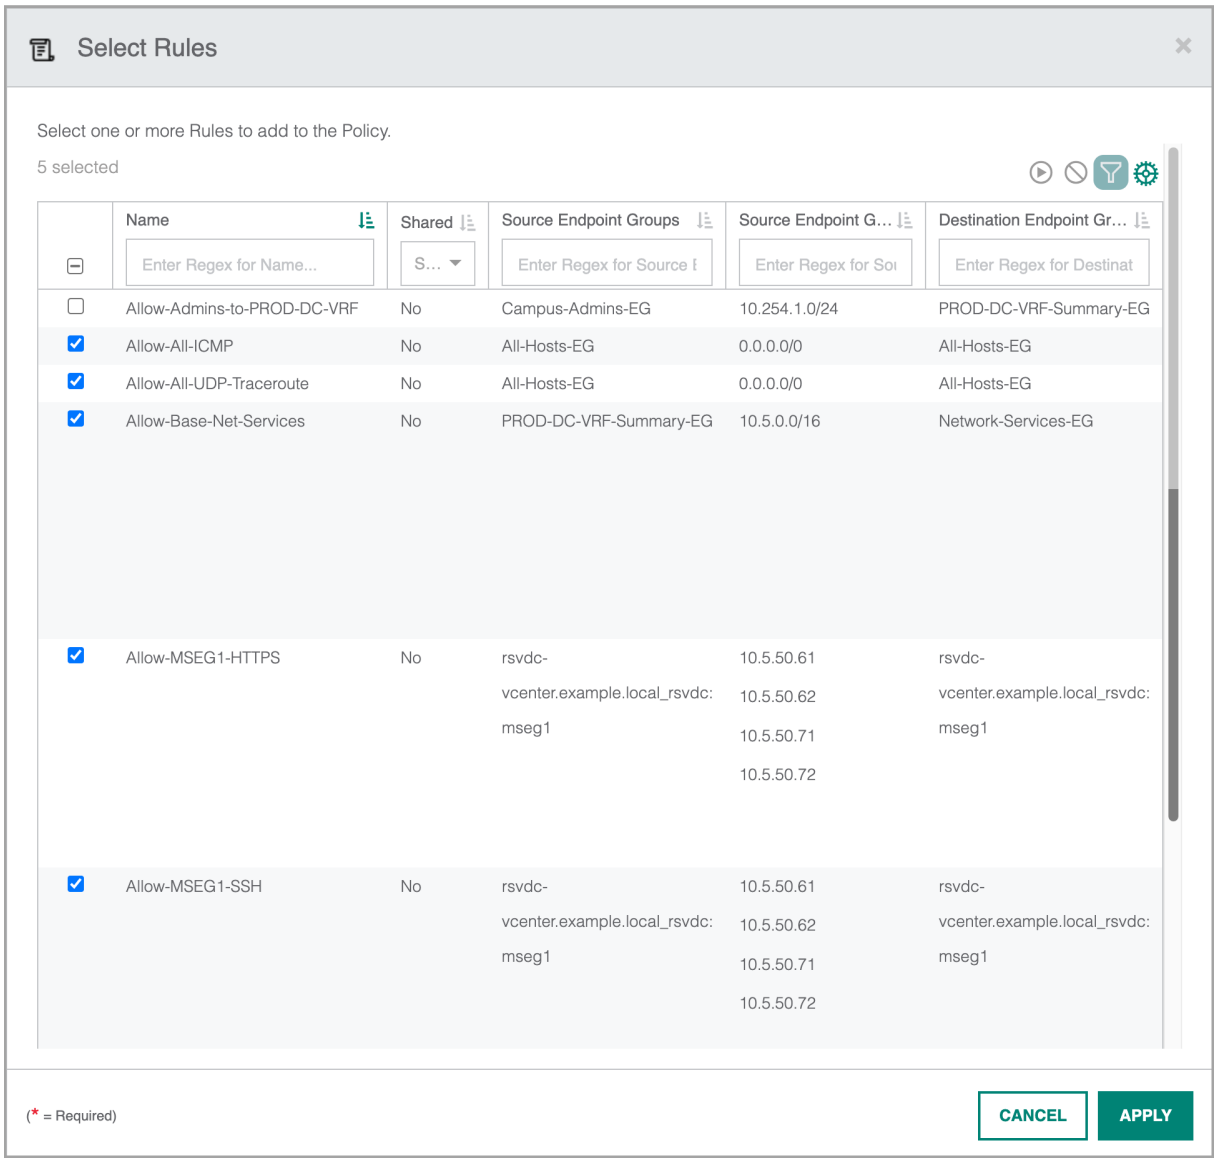


**Step 31 Policy** wizard: Verify the rule order matches requirements, then click **NEXT**.

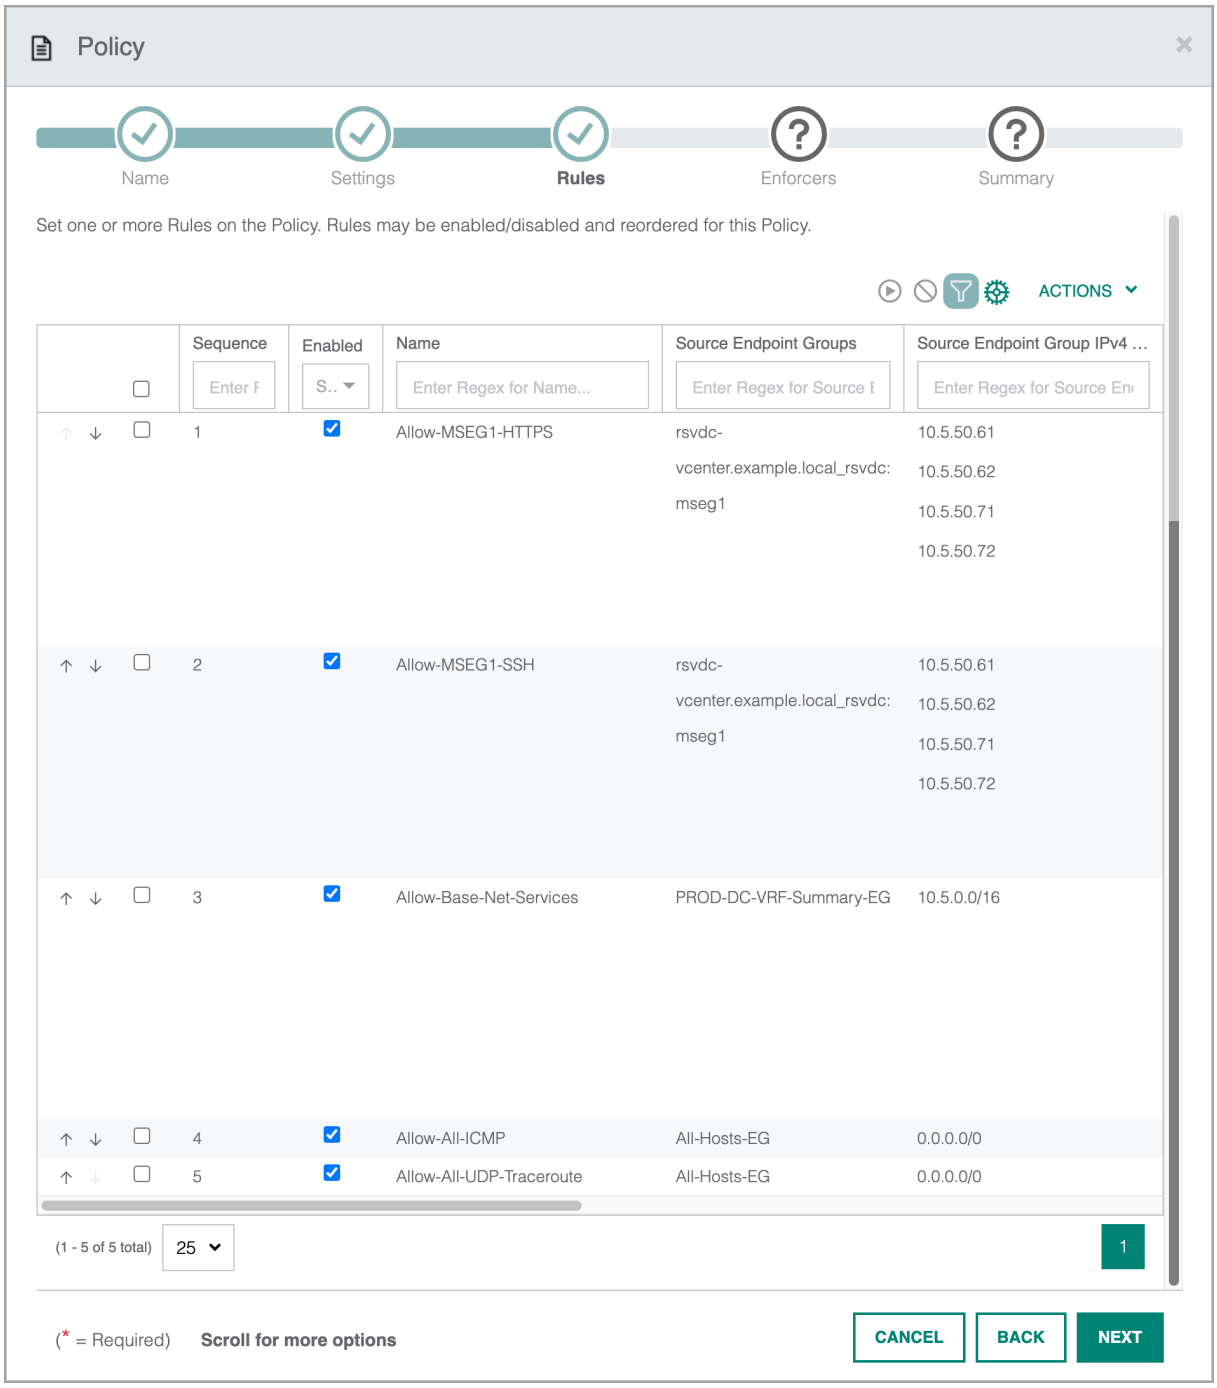


**Step 32 Policy** wizard: On the **Enforcers** page, verify the pre-populated egress enforcer is listed, then click **NEXT**.

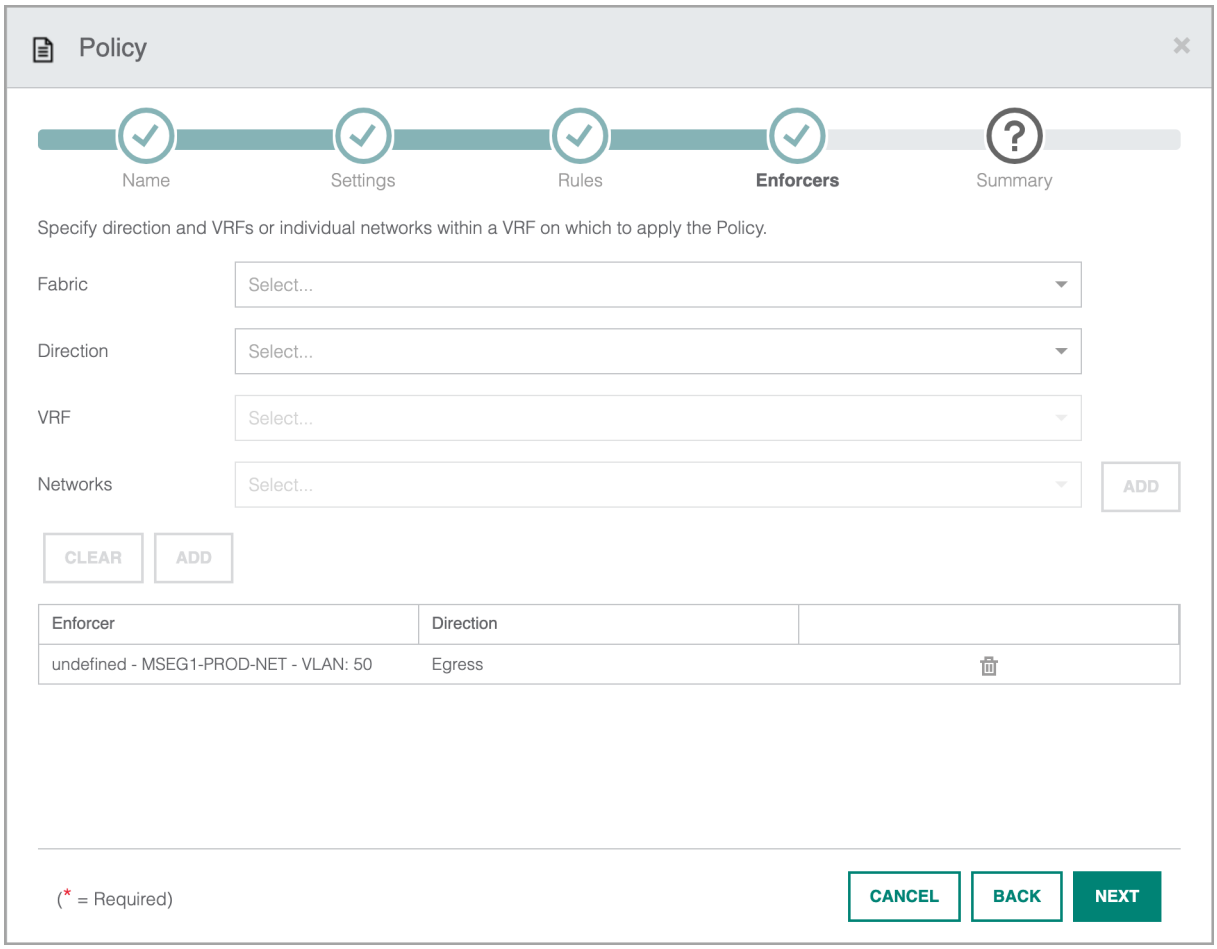


**NOTE:**

Policy direction and design vary based on requirements and administrator preferences. When all hosts are attached to CX 10000 switches, egress policy is recommended to filter east-west traffic. If an ingress policy is preferred, the pre-populated egress enforcer can be removed by clicking the **trash can** icon.

**Step 33 Policy** wizard: On the **Summary** page, verify that all information is correct, then click **APPLY**. The **Policy** wizard closes, returning to the main **Distributed Virtual Switch** workflow.

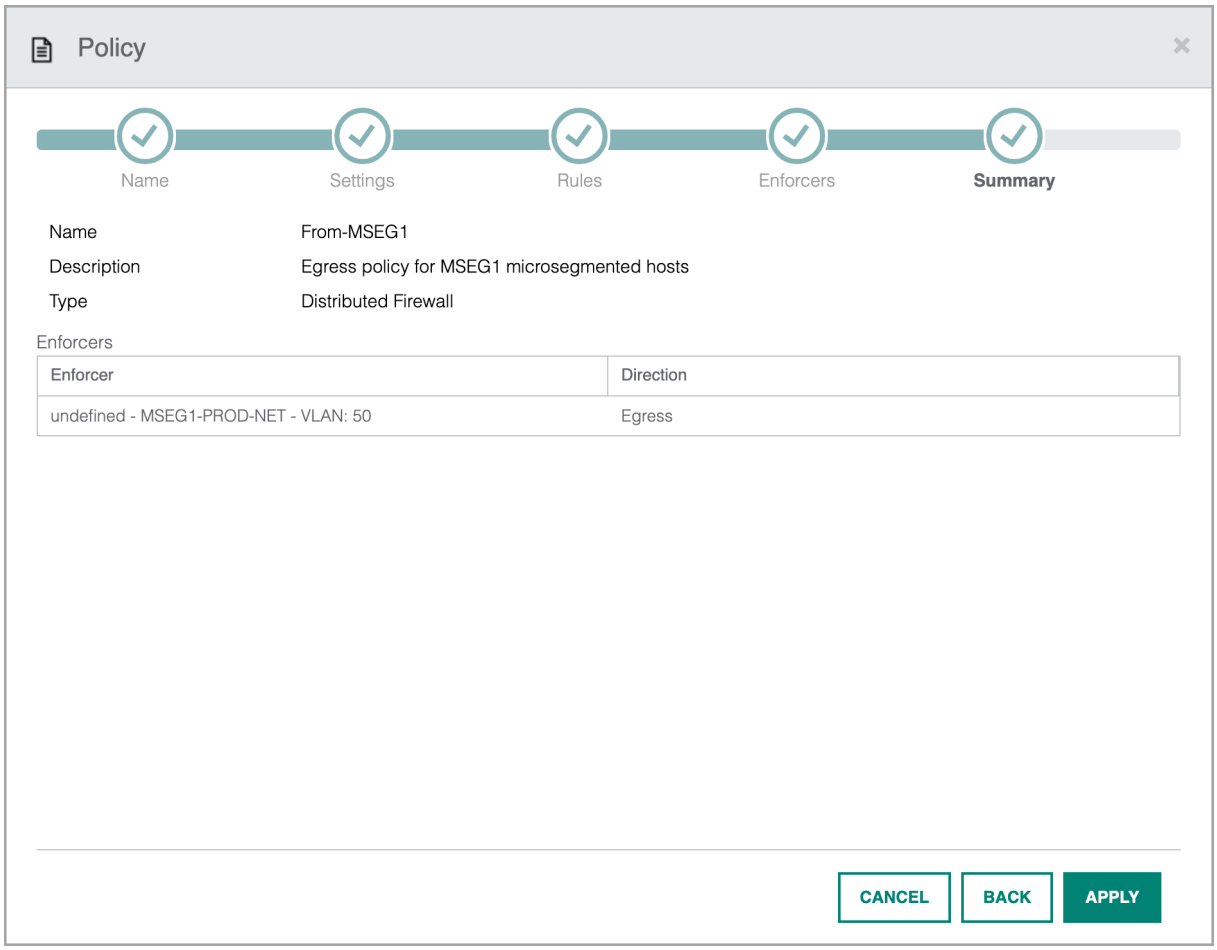


**Step 34** On the **Policy** page, verify that the newly created policy is populated in the **Policies** field, then click **NEXT**.

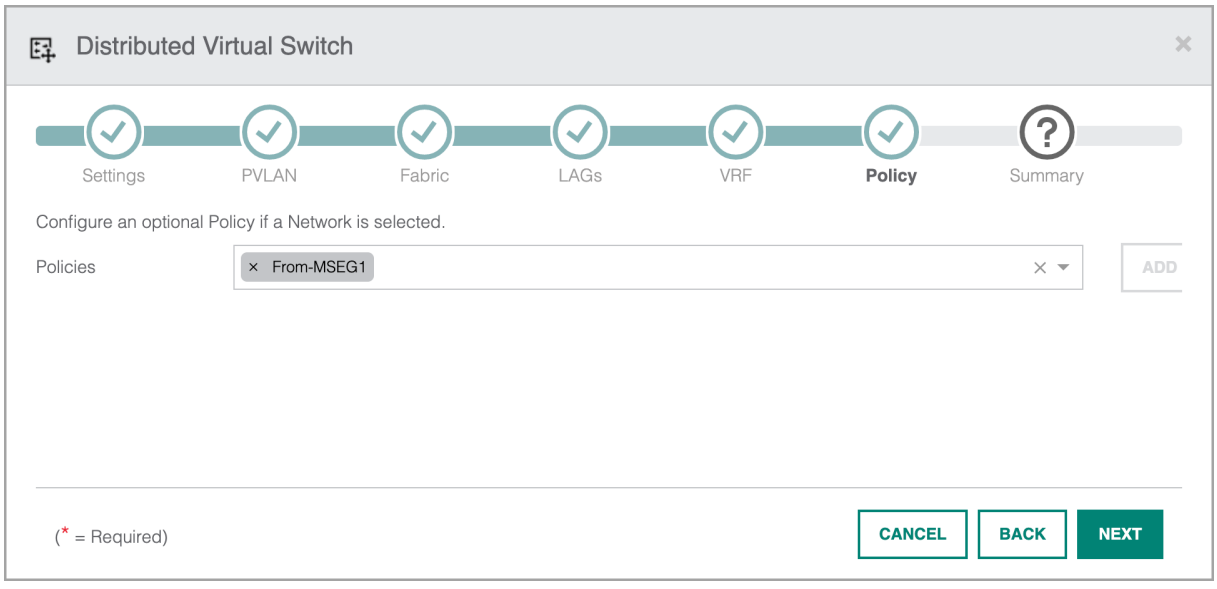


**Step 35** On the **Summary** page, verify that all information is correct, then click **APPLY**.

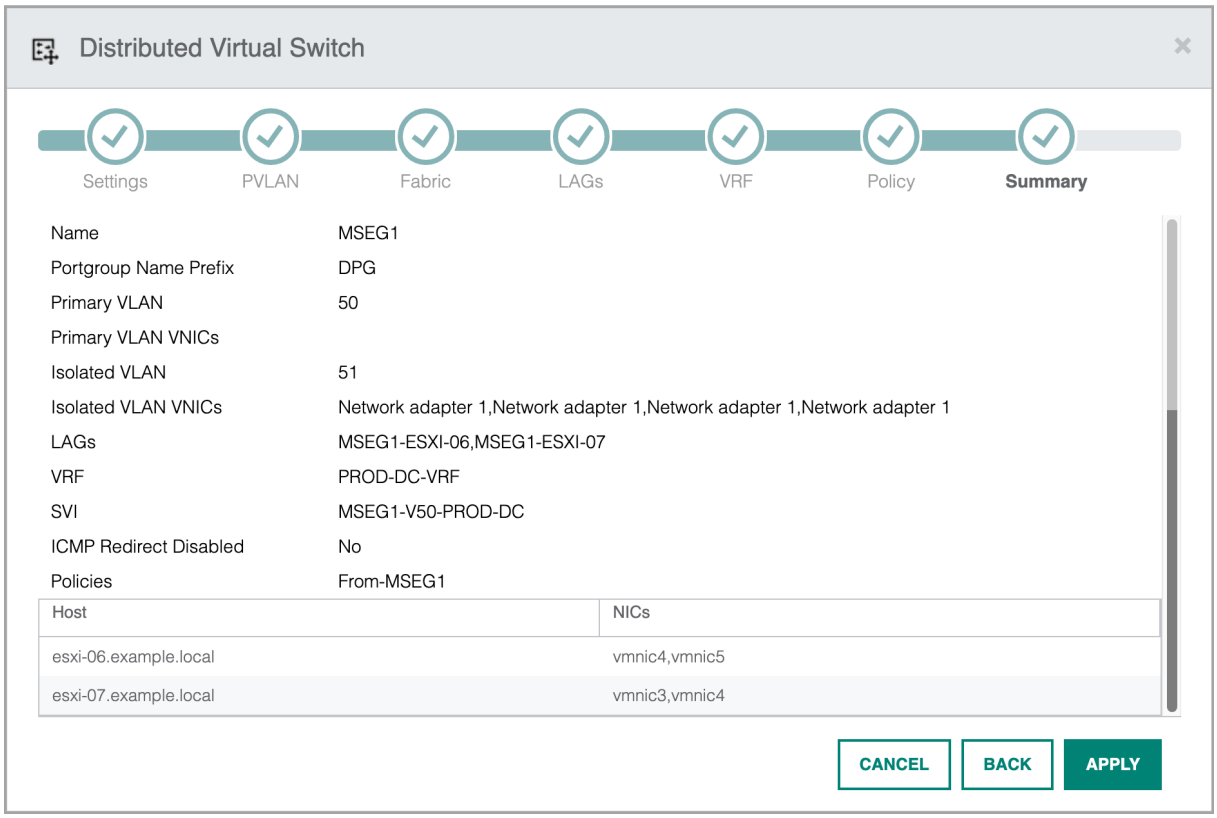


At the completion of this step, AFC creates the PVLAN and LAG/LACP configuration on all included ESXi hosts and the MC-LAGs on uplink ToR switches.

## Add the Primary PVLAN to the EVPN

The primary PVLAN is added to the EVPN config to advertise reachability of the microsegmented hosts in the fabric. A microsegmentation can exist in a single rack or multiple racks in the fabric.

**Step 1** On the **Configuration** menu, select **Routing > EVPN**.

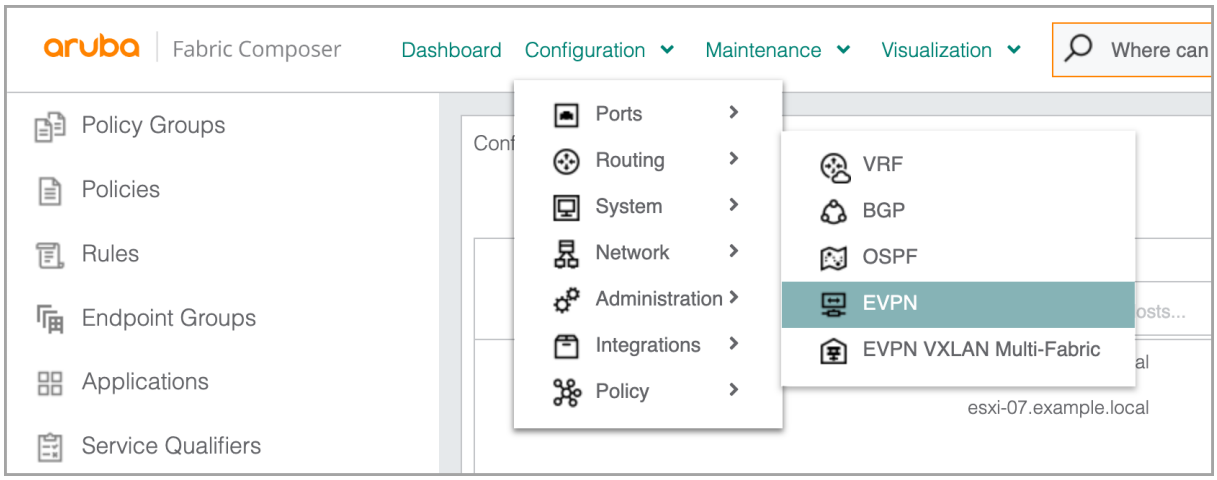


**Step 2** To modify the EVPN configuration, click the right **ACTIONS** menu and select **Add**.

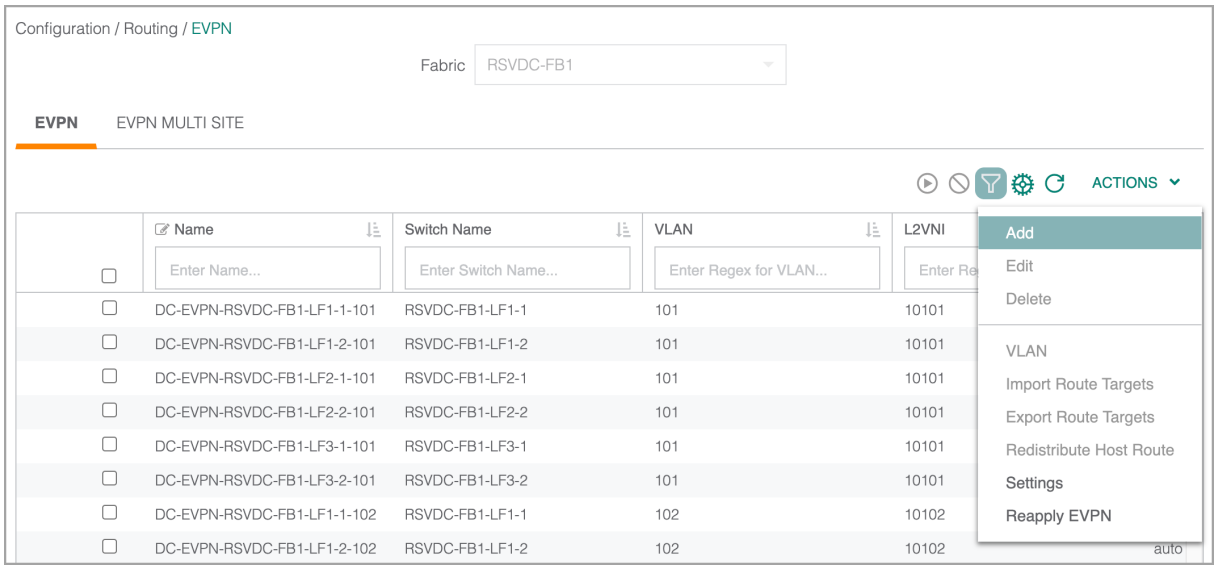


**Step 3** On the **Introduction** page, read the configuration note and click **NEXT**.

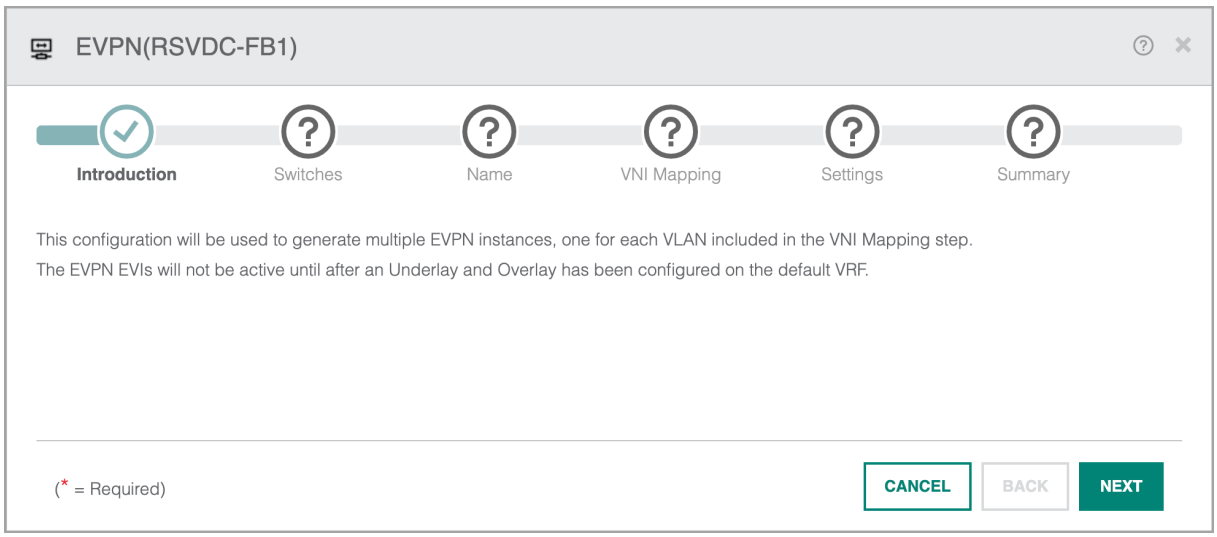


**Step 4** On the **Switches** page, unselect the **Create EVPN instance across the entire Fabric and all Switches contained within it** checkbox. Specify all DSS switches connected to the PVLAN microsegmentation, then click **NEXT**.

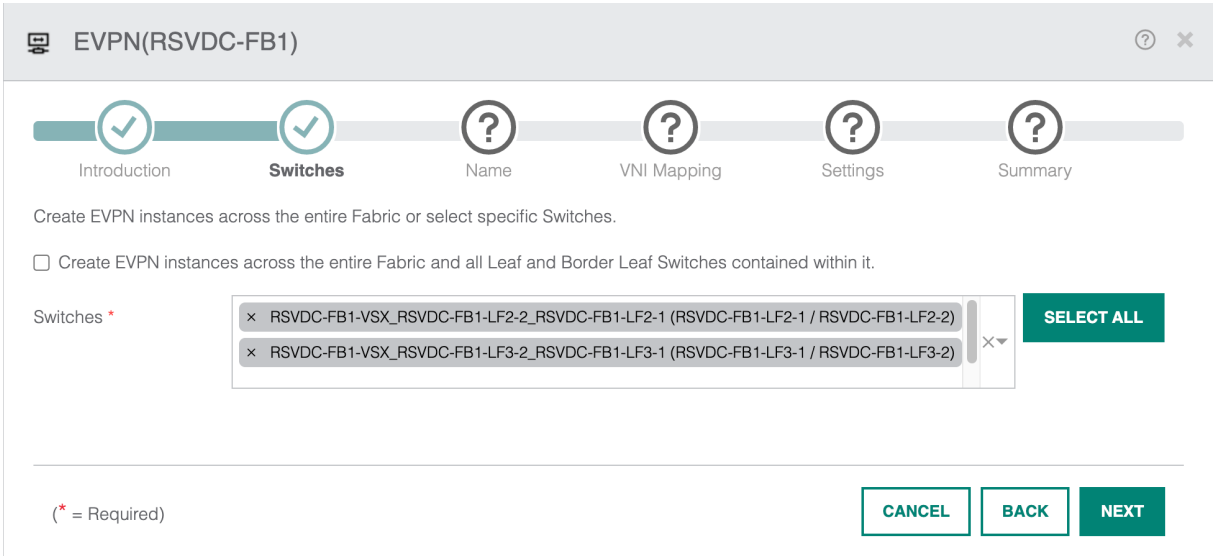


**Step 5** On the **Name** page, enter the **Name Prefix** previously used to create an EVPN map and click **NEXT**.

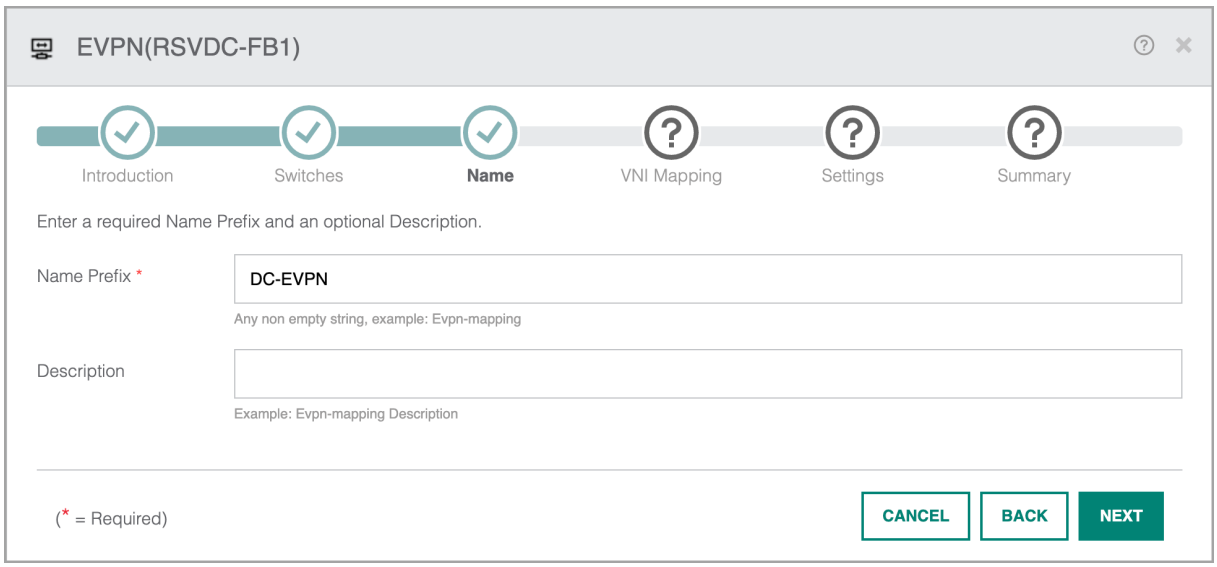


**NOTE:**

It is best practice to use the same prefix value for all EVPN configuration within a single fabric.

**Step 6** On the **VNI Mapping** page, the previously used value for **Base L2VNI** is autopopulated. Enter the microsegmentation primary PVLAN in the **VLANs** field, then click **NEXT**.

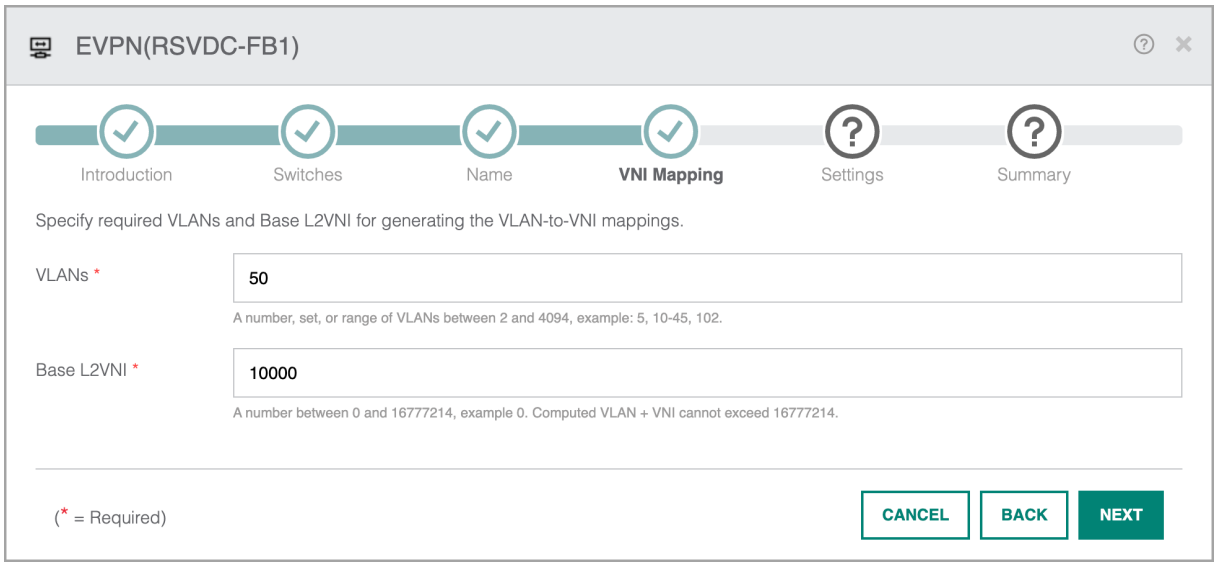


**Step 7** On the **Settings** page, the previously used MAC resource pool is autopopulated in the **MAC Address Resource Pool** field. Specify *AUTO* as the **Route Target Type** and click **NEXT**.

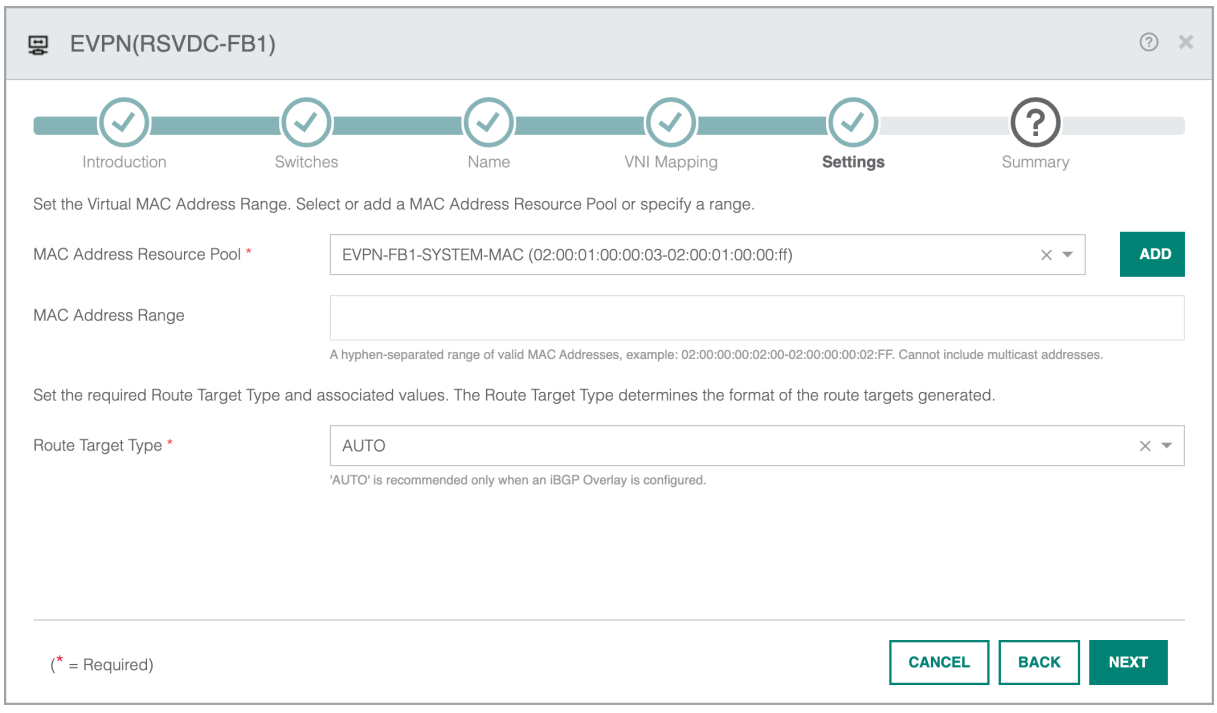


**Step 8** On the **Summary** page, verify that all values are correct and click **APPLY**.

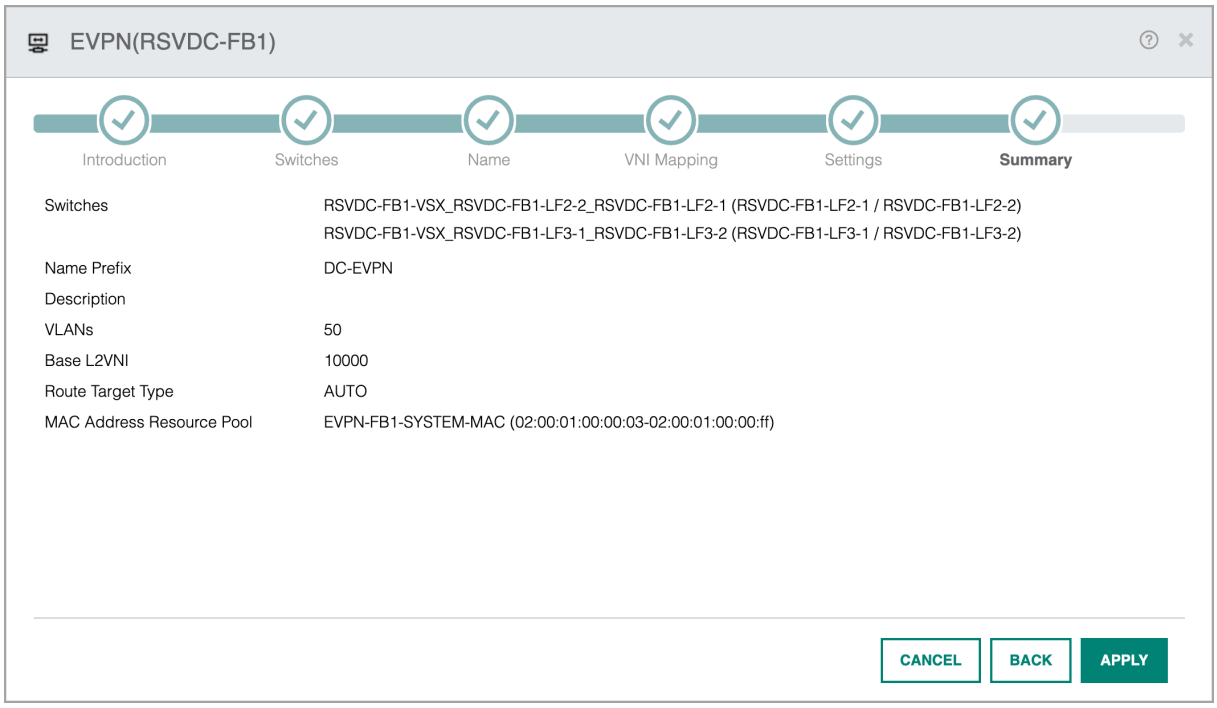


## Externally Advertise Microsegmented Networks

During the AFC EVPN-VXLAN creation, BGP was enabled for the PROD-DC VRF. AFC configures redistribution of connected interfaces for all leaf switches where the VRF is configured, when BGP is enabled. This configuration installs the primary PVLAN SVI prefix into BGP, which facilities the advertisement of the primary PVLAN network to campus and other fabrics without the need for additional steps.

# Verify Policy in PSM

AFC policy is pushed to AMD Pensando's PSM, which installs the policy on CX 10000 DPUs.

## Verify Networks

**Step 1** On **PSM**, expand **Tenants** in the left pane and click **Networks**.

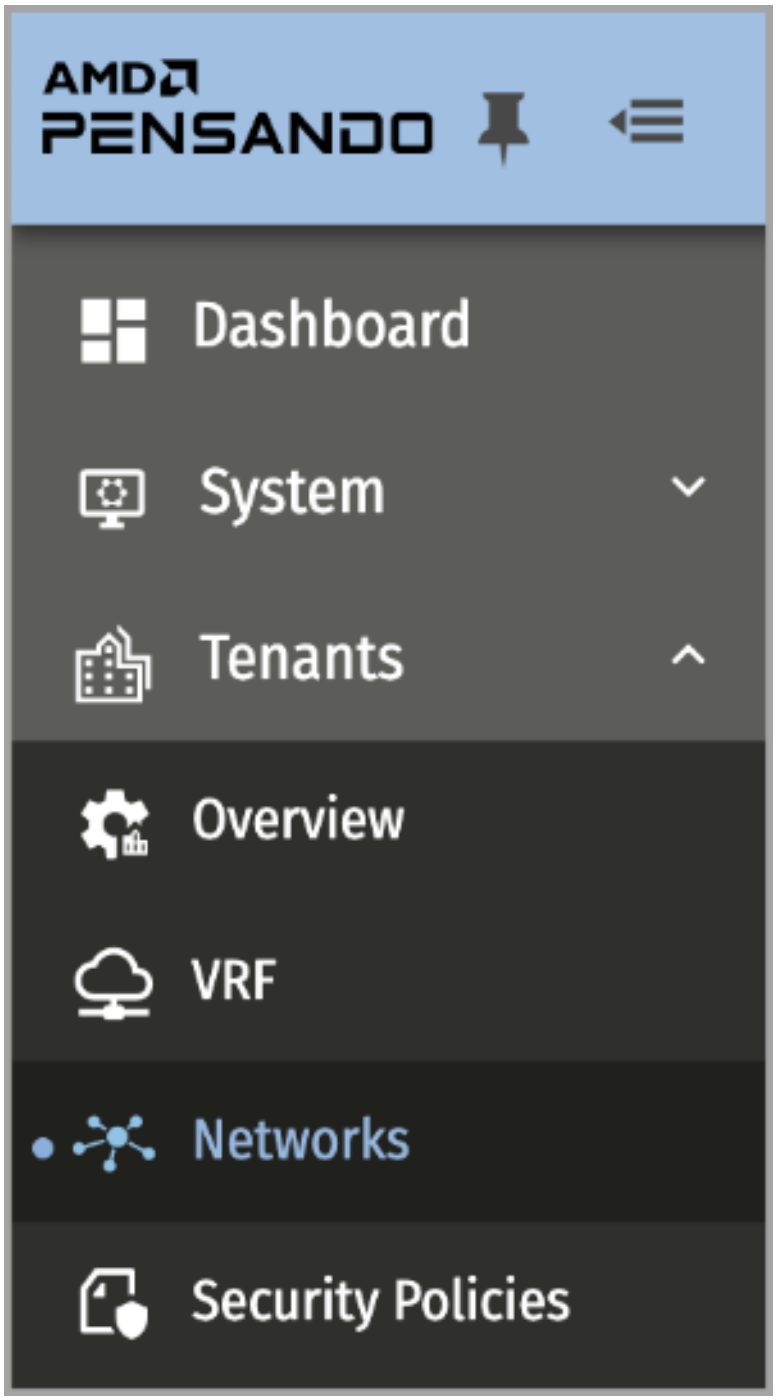


**Step 2** Verify that the PSM *Networks* displayed have the correct ingress and egress policies applied.

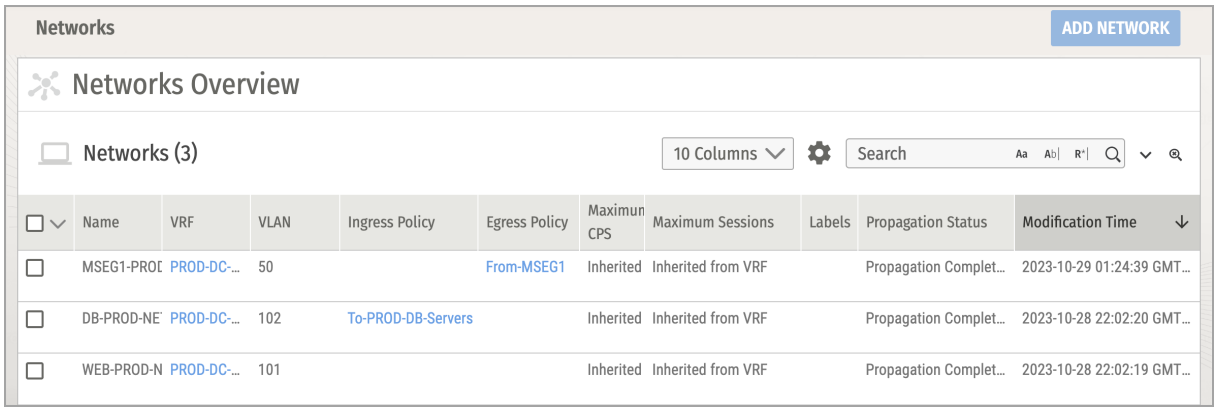


**NOTE:**

Each defined *Network* should display “Propagation Complete” in the **Propagation Status** column, indicating that all CX 10000 switches have installed the displayed policy.

**Step 3** To view a complete list of policies, on the left navigation menu, click **Security Policies**.

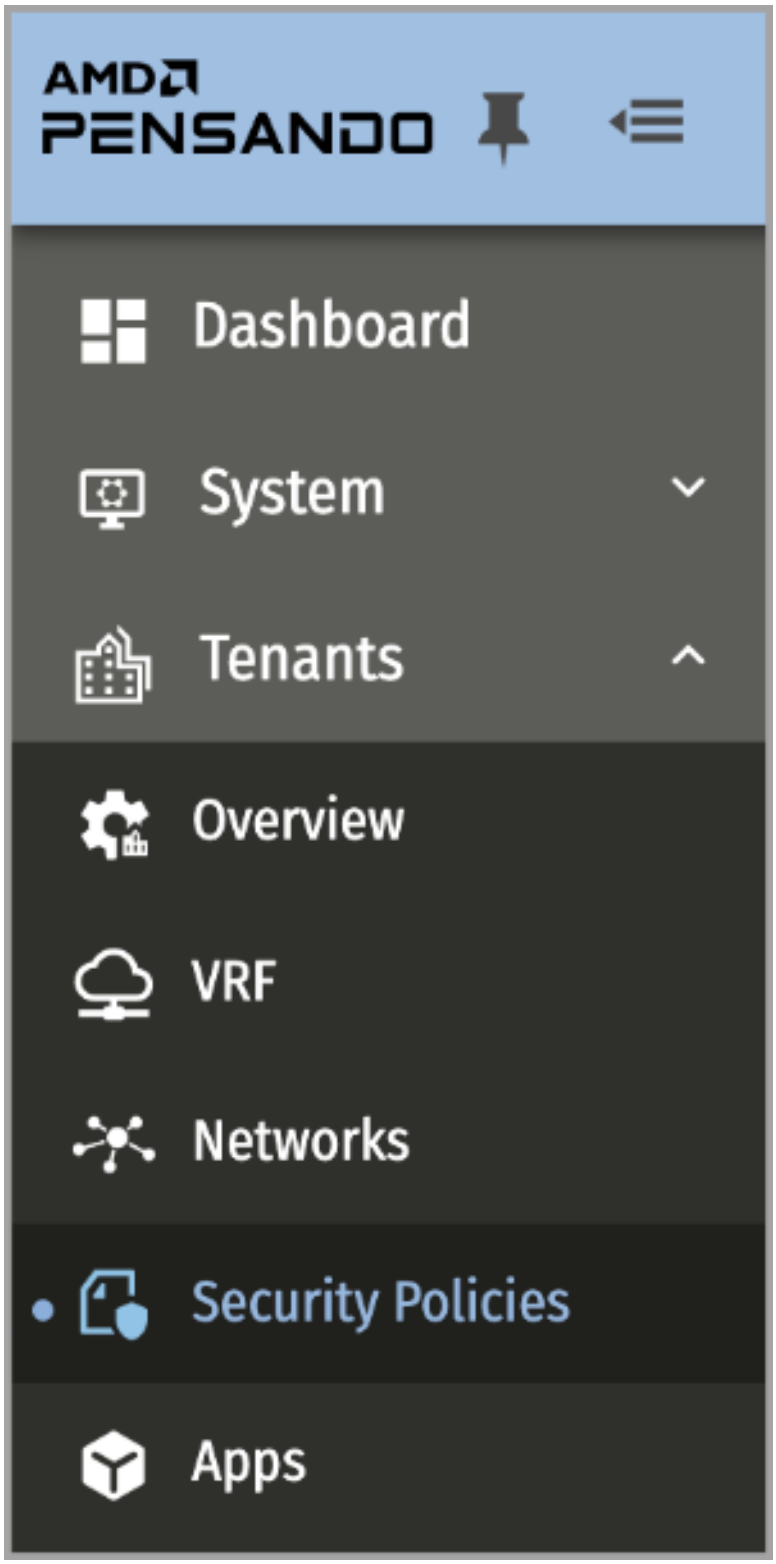


**Step 4** Review the list of policies. To view the rules included in a policy, click a policy name.

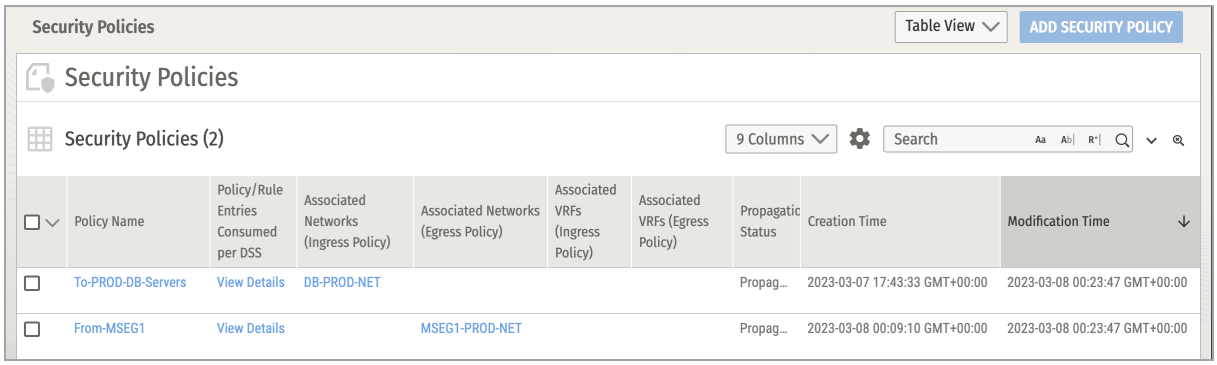


**Step 5** To view the full text of a rule in the displayed rule set, mouse-over a rule in the list.

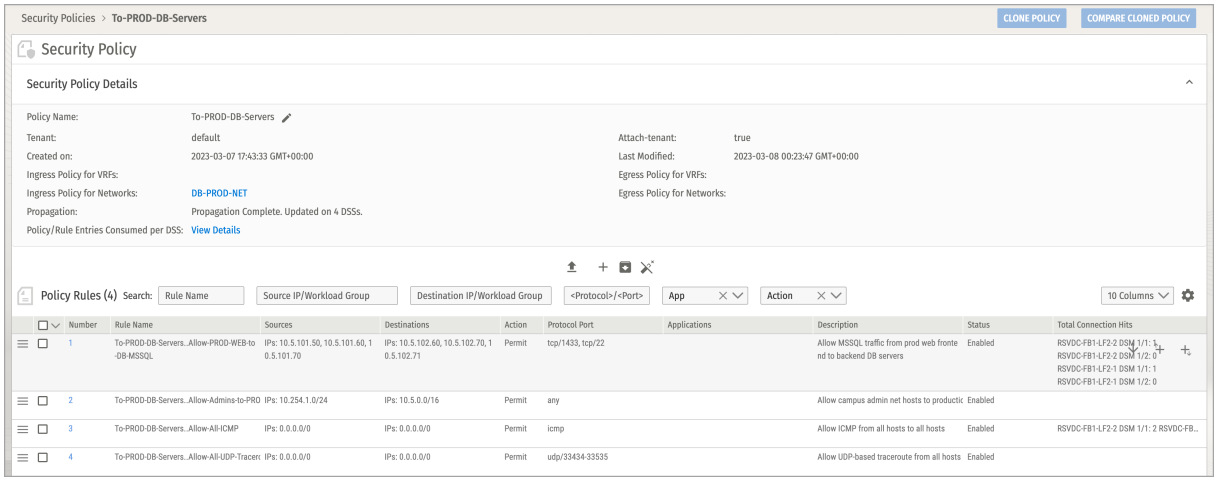

# AFC Multifabric Configuration

Combining multiple EVPN fabrics into a single overlay allows sharing Layer 2 and Layer 3 reachability between data center pods at the same site as well as more distant data center locations.

Developing an effective multifabric strategy supports growth, location diversity, and disaster recovery. Aruba Fabric Composer (AFC) can simplify the configuration process and automate building a multifabric EVPN-VXLAN overlay.

Multiple data center fabrics and locations can be combined into a single overlay topology. This guide focuses on the common usage of combining two data center locations.

Additional considerations are required when adding three or more fabrics or when multiple fabrics are present at a single site which co-exist with fabrics at remote locations. Details on supporting these configurations can be found in the *AOS-CX EVPN VXLAN Guide* in the Aruba Support Portal.

The primary data center fabric in this guide is named RSVDC-FB1, which was previously created. The new fabric established in a colocation facility is named RSVCO-FB2.

## Second Fabric Guidance

Use the process outlined in the AFC EVPN-VXLAN Configuration guide to create a second fabric. The same AFC instance can be used to build the new fabric, if IP connectivity can be established between AFC and the new fabric location with a latency of 50 ms or less. A second AFC instance can be deployed and linked to the first instance, if latency exceeds this value.

The diagram below illustrates the physical topology of the second fabric used in this guide. The new fabric is located at a colocation facility and identified as RSVCO-FB2.

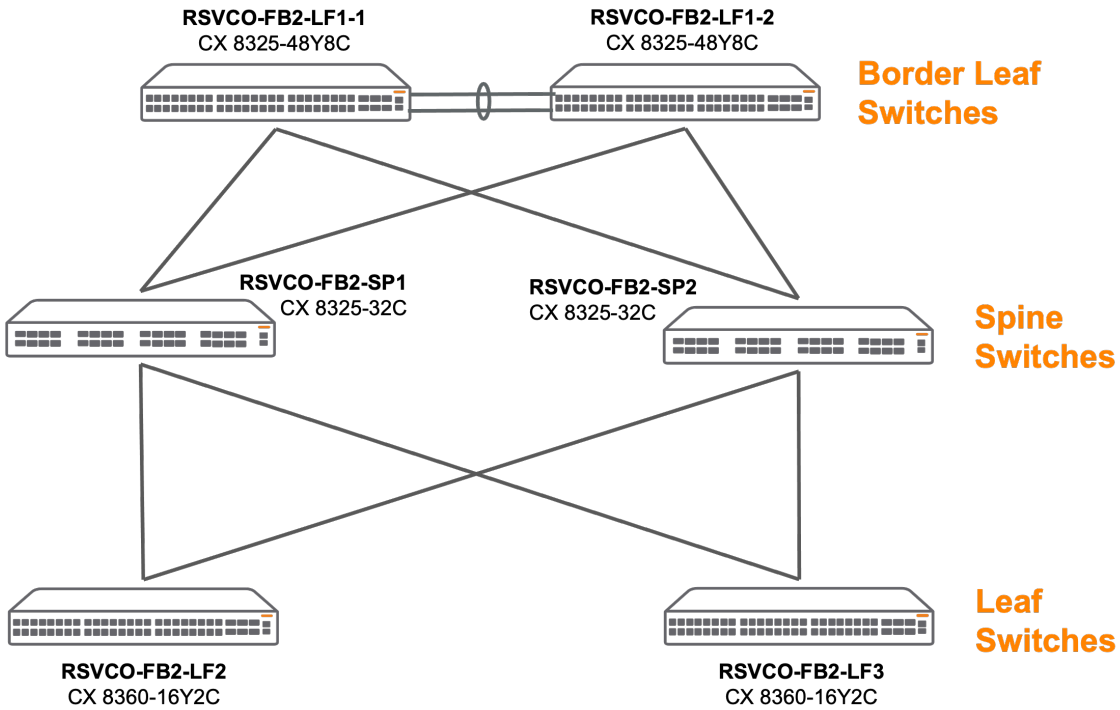

The RSVCO-FB2 fabric uses unique IP numbering for loopbacks and VTEPs. The following IP ranges are used to build the second fabric.

| Purpose | Description | Example |
|---|---|---|
| Leaf-Spine IP address block | An IPv4 address block used to create /31, point-to-point layer 3 links between leaf and spine switches. | 10.255.4.0/24 |
| Routed loopback and VSX transit VLAN IP address block | An IPv4 address block used to allocate unique loopback addresses (/32) for each switch and transit-routed VLAN between redundant ToRs (/31) | 10.250.3.0/24 |
| VTEP loopback IP address block | An IPv4 address block used to assign VTEP loopback addresses (/32) on VSX redundant ToRs | 10.250.4.0/24 |

The following network elements are logically contiguous between fabrics.

| Network Element | Description |
|---|---|
| PROD-DC-VRF | A virtual routing and forwarding table defined for the data center production network. |
| VLAN 101 | Production web server VLAN and SVI. |
| VLAN 102 | Production database server VLAN and SVI. |

## VXLAN Data Plane Configuration

The PROD-DC-VRF is Layer 3, extended across the RSVDC and RSVCO data center fabrics. VLANs 101 and 102 are Layer 2, extended between the fabrics. To allow a contiguous overlay for network segments extended between fabrics, VXLAN VNI values in each data center must agree. The same L3 VXLAN VNI value for PROD-DC-VRF must be defined in both fabrics.

To assign the same L2 VXLAN VNIs to both fabrics, assign the same **Base L2VNI** value in the **EVPN CONFIGURATION** wizard.

## Underlay WAN IP

Border leaf switches typically establish underlay route peerings between fabrics, but other switch members of a fabric can perform this function.

IP addresses must be assigned to WAN interfaces in both fabrics. When Layer 2 WAN services or dark fiber are used, IP addresses are assigned for peering directly between switches at each site. The following IP addresses are reserved for underlay interfaces connecting to the sample Layer2 metro Ethernet circuit: 10.255.6.0/29. When using Layer 3 WAN services, IP addresses are assigned to allow direct peering with the service provider.

# Configure Multifabric Underlay Connectivity

A multifabric underlay serves the same purpose as the underlay within a single fabric. It shares IP loopbacks to enable MP-BGP EVPN peerings and VTEP reachability.

## WAN Underlay Connectivity

Several methods can be used to establish connectivity between fabrics at different sites. Dark fiber and most metro Ethernet services support jumbo frame capabilities. Fragmentation of VXLAN encapsulated traffic is not supported. The WAN path between fabrics must accommodate an increase of 50 bytes in MTU over the encapsulated traffic. A WAN path MTU of 1600 is recommended.

This guide uses a Layer 2 metro Ethernet service with multi-port customer premise equipment (CPE). The border leaf switches for each fabric use a single physical connection to the CPE at their respective locations.

The diagram below illustrates the sample connectivity between data center locations. RSVDC Fabric 1 is the fabric created in the AFC EVPN-VXLAN Configuration guide. RSVCO Fabric 2 is the second fabric located in the colocation facility.

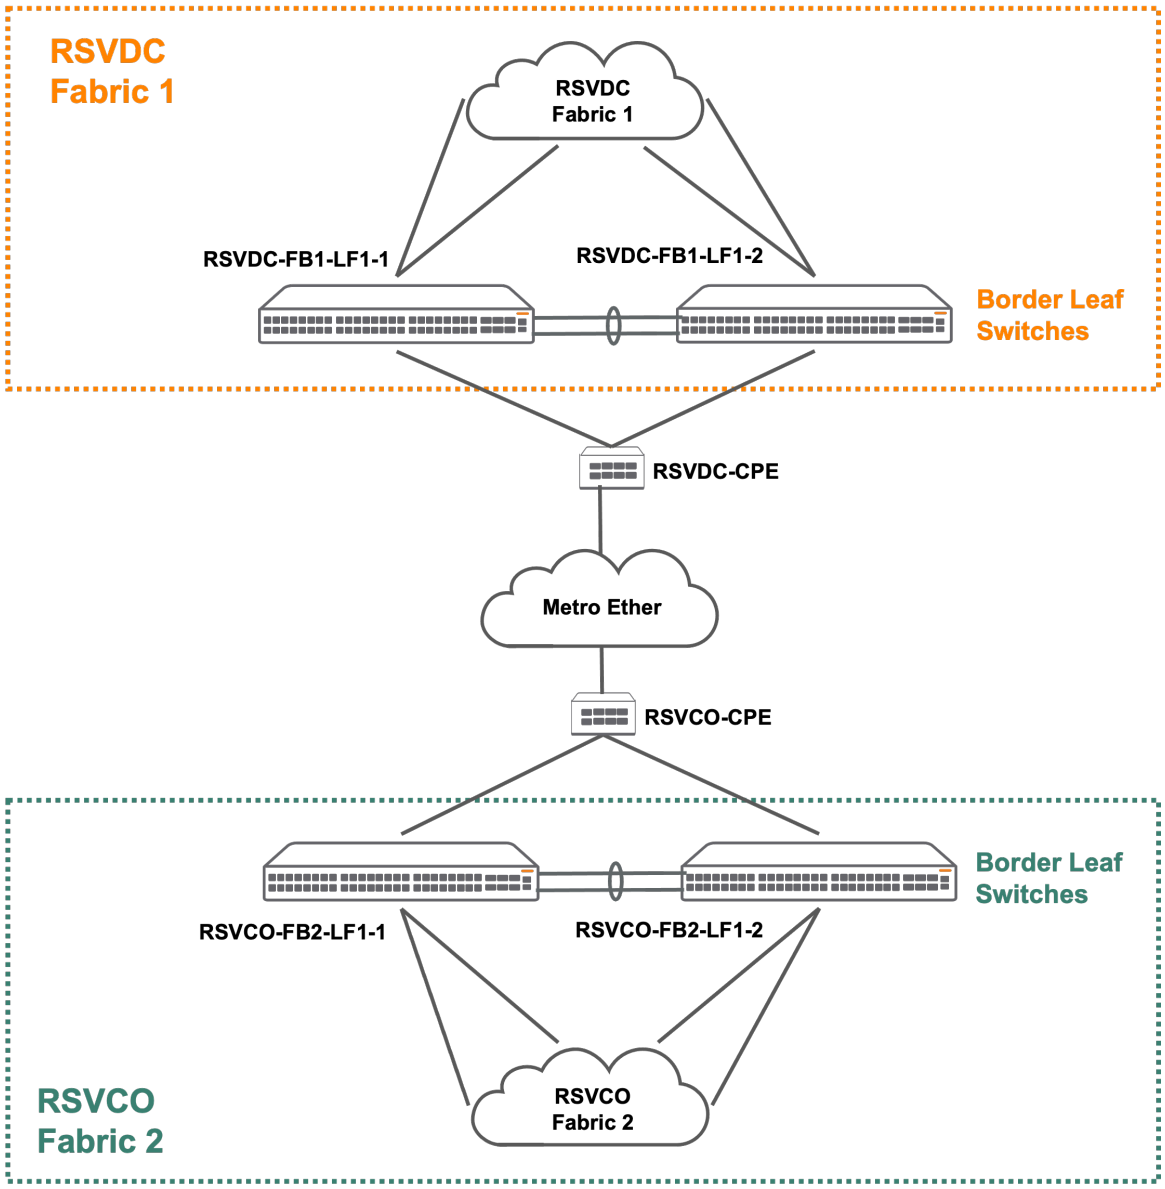


Improved WAN resiliency can be realized by using multiple CPE devices at each location, by provisioning a second metro Ethernet circuit, or by using multiple dark fiber links.

## Assign WAN IP Addresses

A single IP subnet is used over the Layer 2 metro Ethernet service. IP addresses are assigned to physical switch interfaces. Routed interfaces are used to avoid Layer 2 loops.

**Step 1** Select **Configuration > Routing > VRF** on the top menu.

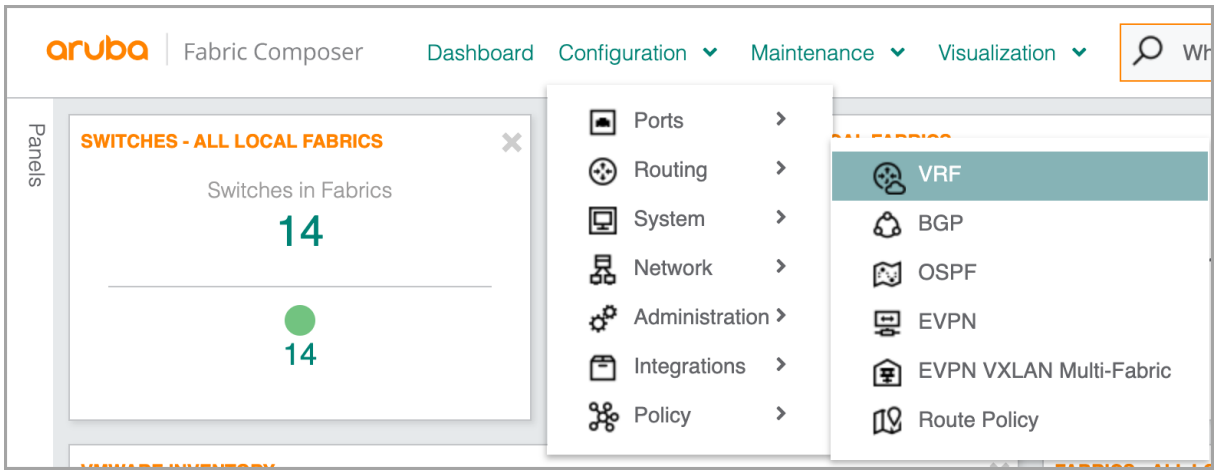


**Step 2** Select **RSVDC-FB1** in the **Fabric** menu, click the • • • symbol next to **default**, and select **IP Interfaces**.

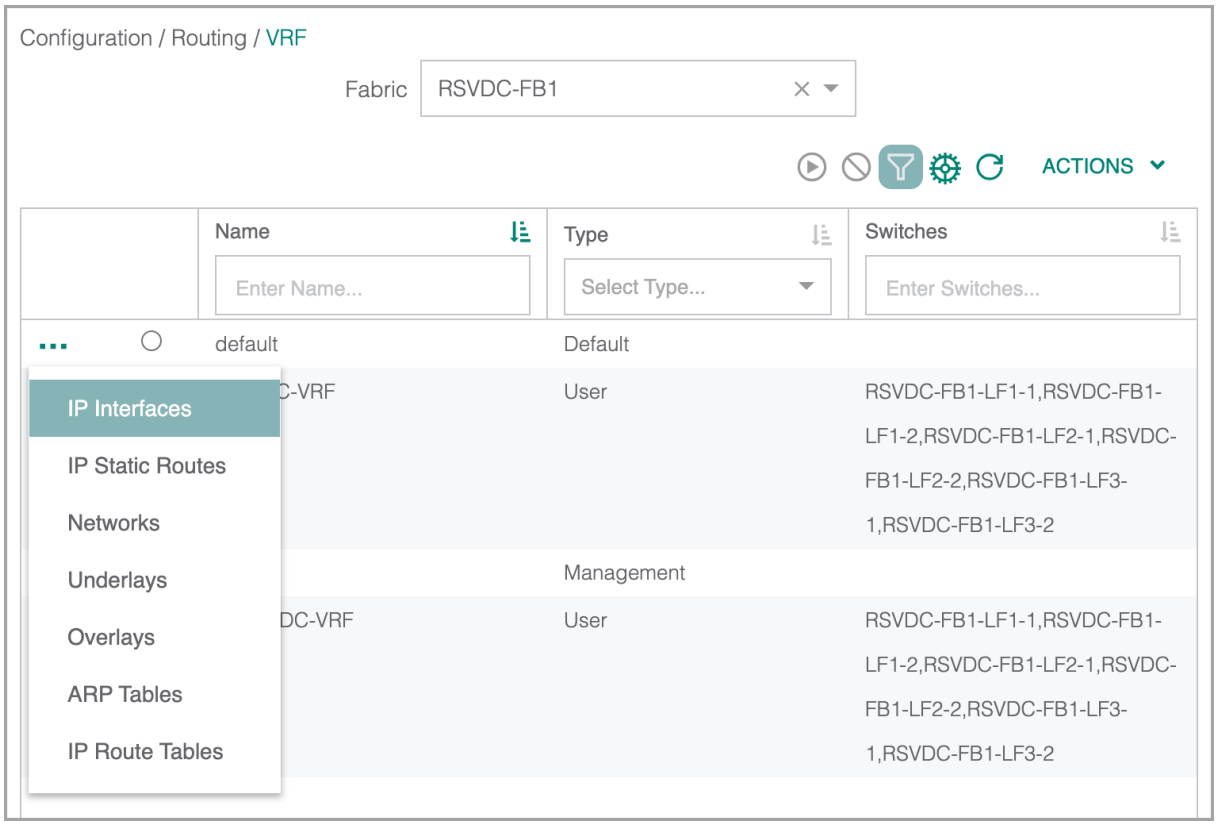


**Step 3** On the lower **ACTIONS** menu of the **IP INTERFACES** tab, select **Add**.

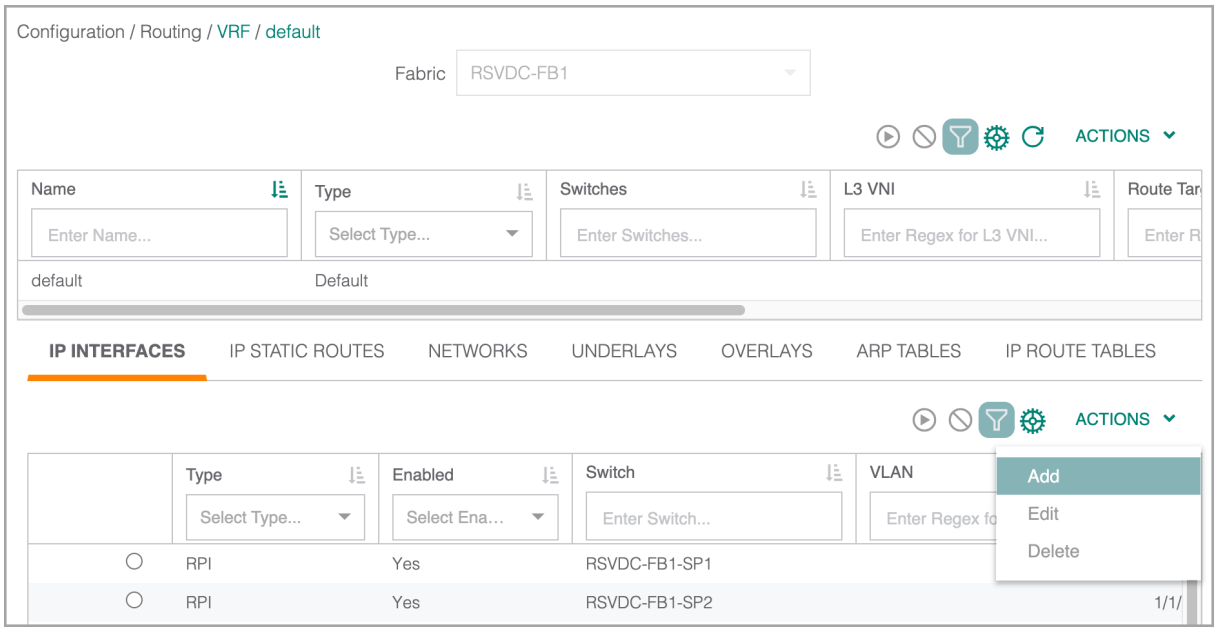


**Step 4** On the **Interface Type** page, assign the following values:

- **Type:** *RPI*
- **Switch:** *RSVDC-FB1-LF1-1*
- **Port/LAG:** *1/1/13*

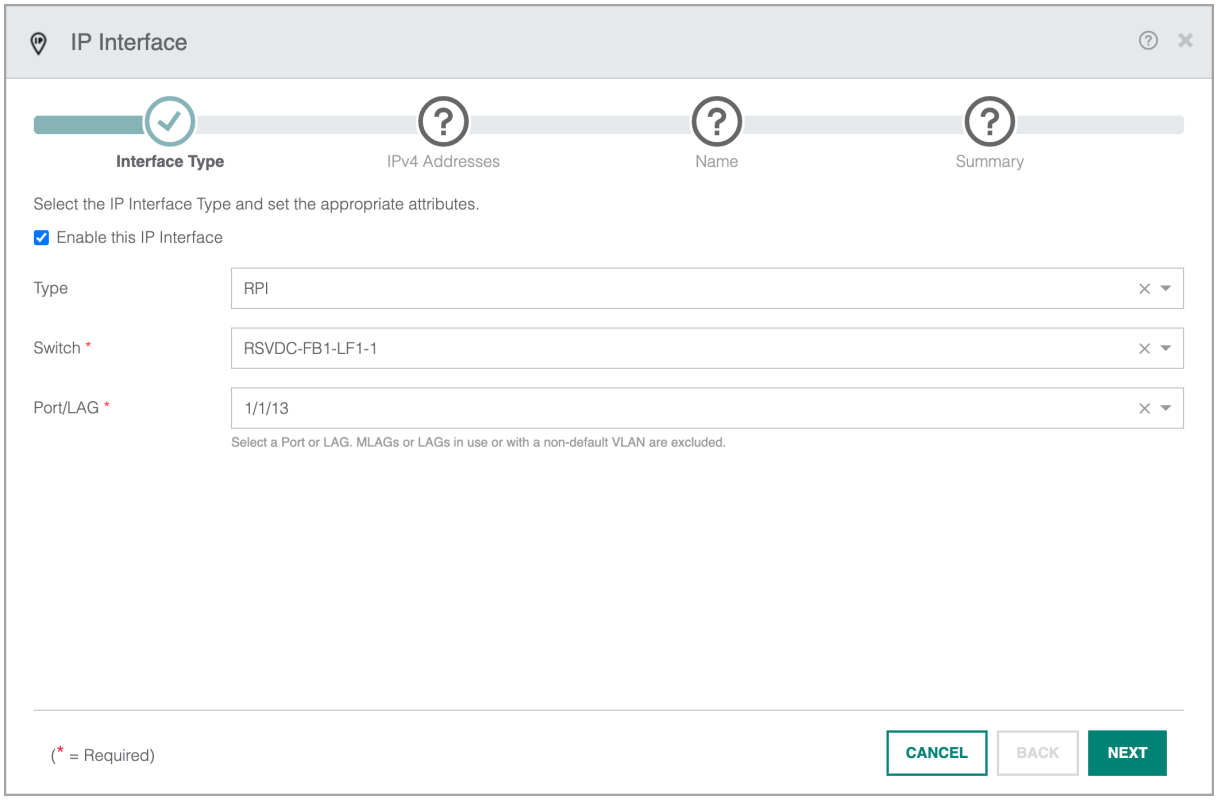


**Step 5** On the **IPv4 Addresses** page, enter the WAN IP address for the **Primary IPv4 Network Address**. Click **NEXT**.

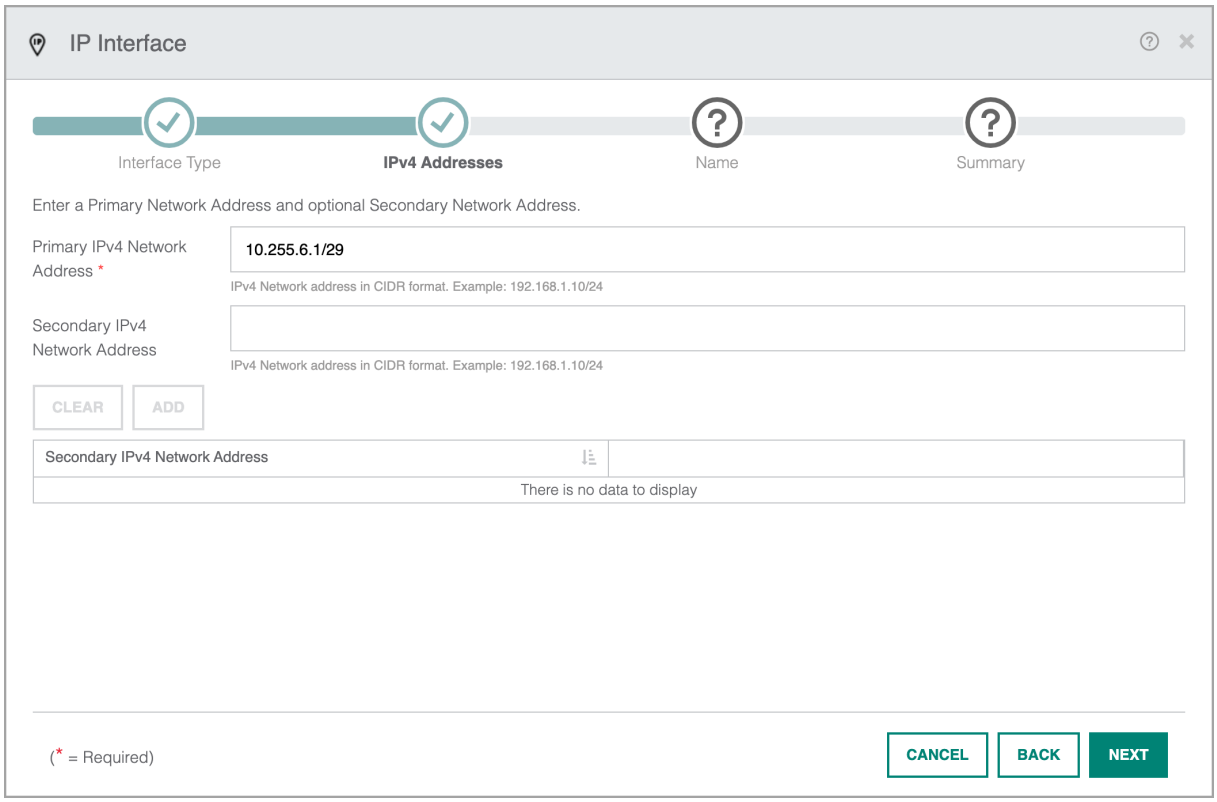


**Step 6** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

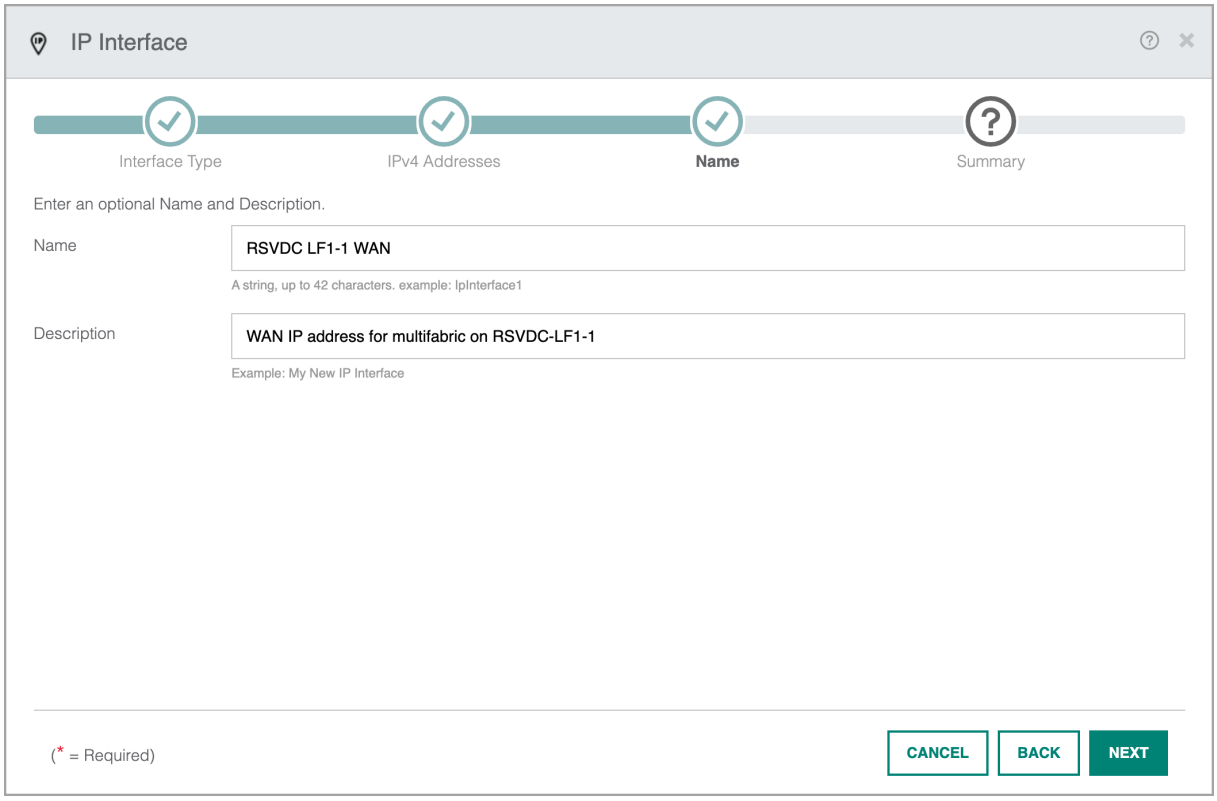


**Step 7** On the **Summary** page, verify the interface settings and click **APPLY**.

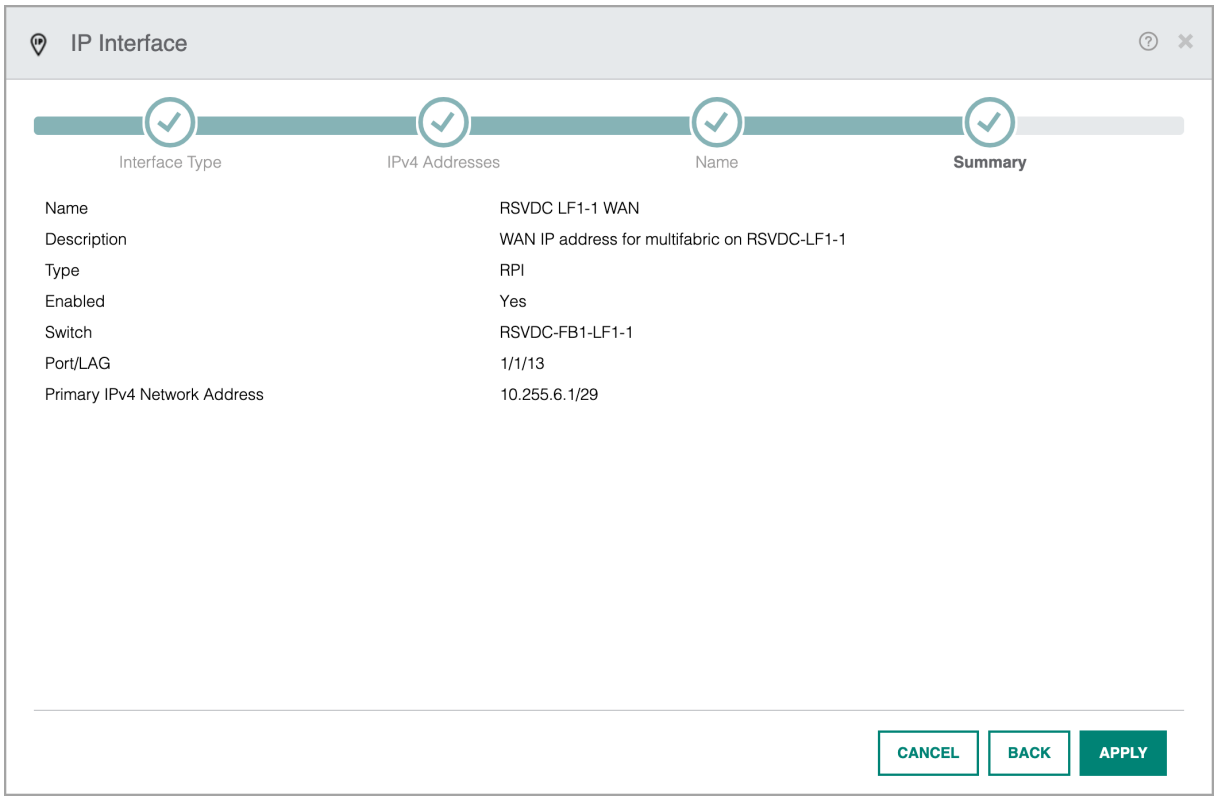


**Step 8** Repeat the procedure to assign IP addresses to the physical WAN interface of each border leaf.

| Switch | Type | Port/LAG | Primary IPv4 Network Address | Name | Description |
|---|---|---|---|---|---|
| RSVDC-LF1-2 WAN | RPI | 1/1/13 | 10.255.6.2/29 | RSVDC LF1-2 WAN | WAN IP address for multifabric on RSVDC-LF1-2 |
| RSVCO-LF1-1 WAN | RPI | 1/1/13 | 10.255.6.3/29 | RSVCO LF1-1 WAN | WAN IP address for multifabric on RSVCO-LF1-1 |
| RSVCO-LF1-2 WAN | RPI | 1/1/13 | 10.255.6.4/29 | RSVCO LF1-2 WAN | WAN IP address for multifabric on RSVCO-LF1-2 |

**NOTE:**

Click **VRF** in the current **Configuration / Routing / VRF / default** display in the upper left or in the left navigation pane to return to the **VRF** window. Select **RSVCO-FB2** in the **Fabric** menu to assign IP addresses in the second fabric.

## Configure Underlay Routing

External BGP (eBGP) using the IPv4 address-family is used to share IP loopback and VTEP reachability. The diagram below illustrates the eBGP IPv4 sessions established to share loopback and VTEP reachability information.

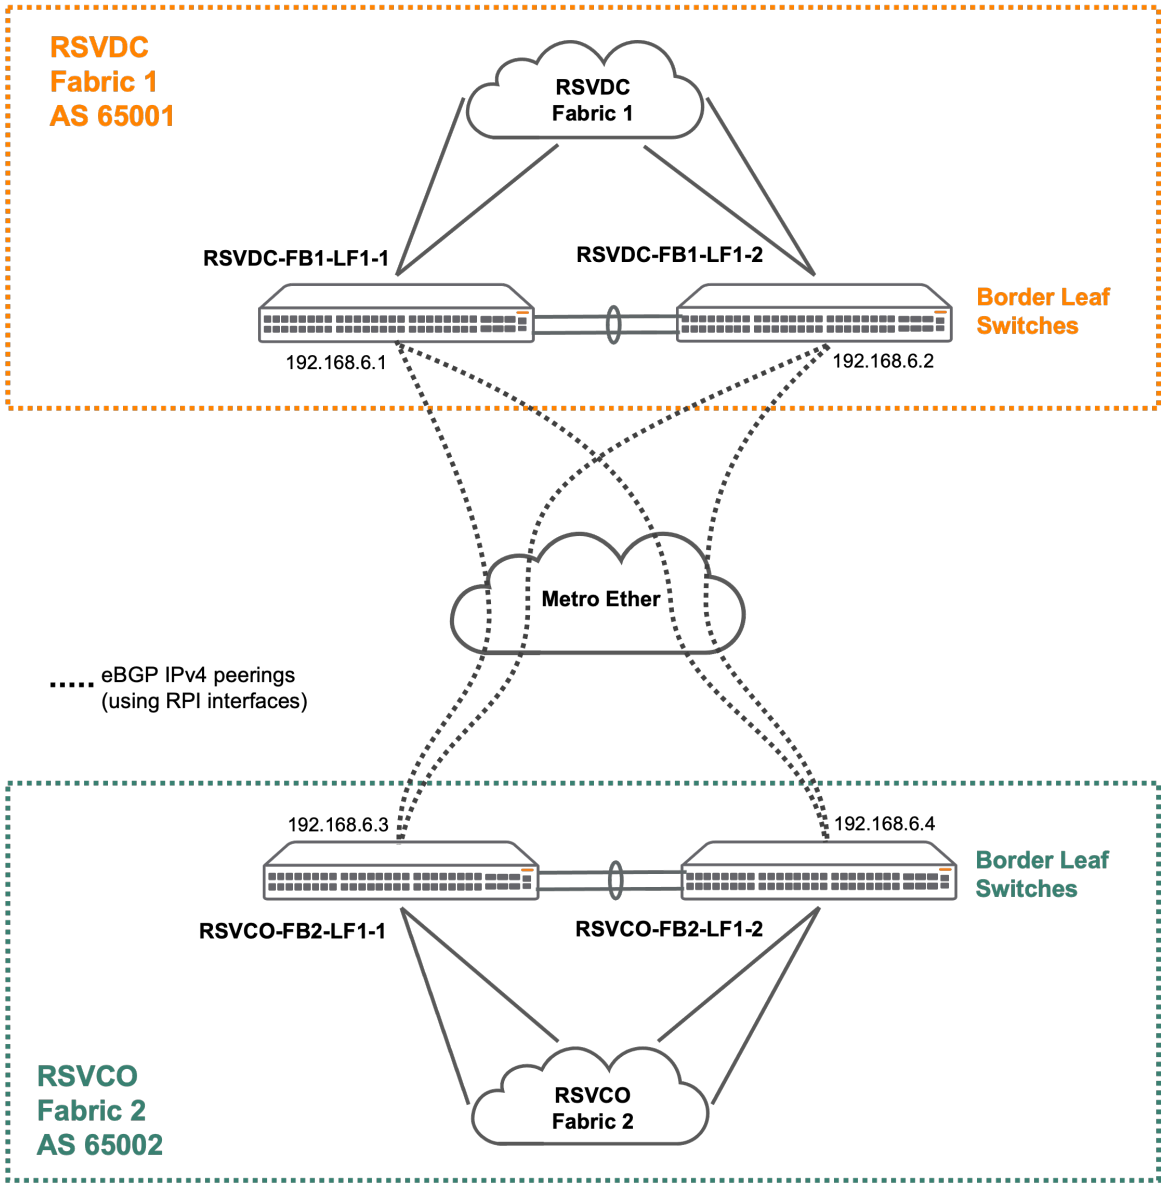


**Step 1** Select **Configuration > Routing > BGP** on the top menu.

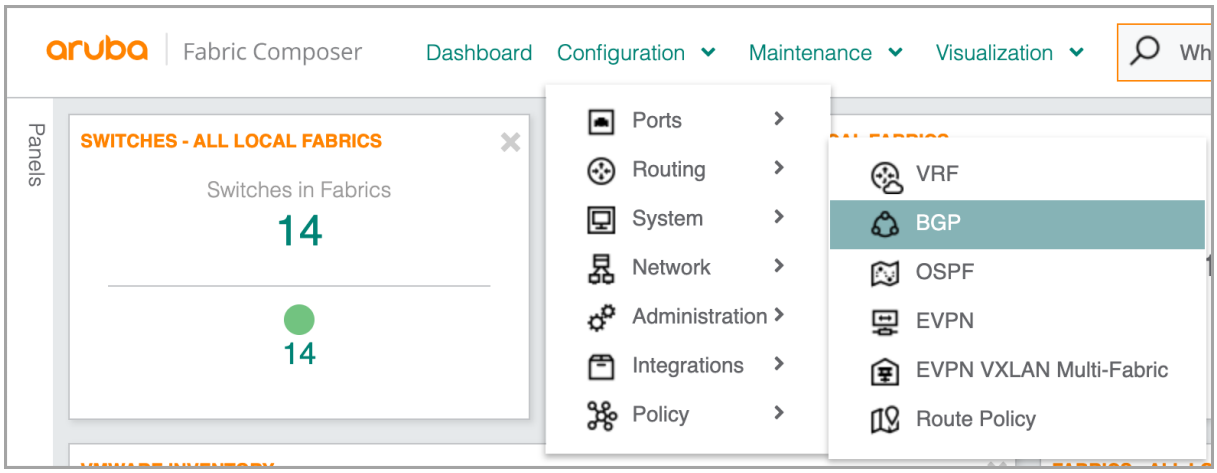


**Step 2** Select **RSVDC-FB1** in the **Fabric** menu. Click the • • • symbol next to **default** and select **Switches**.

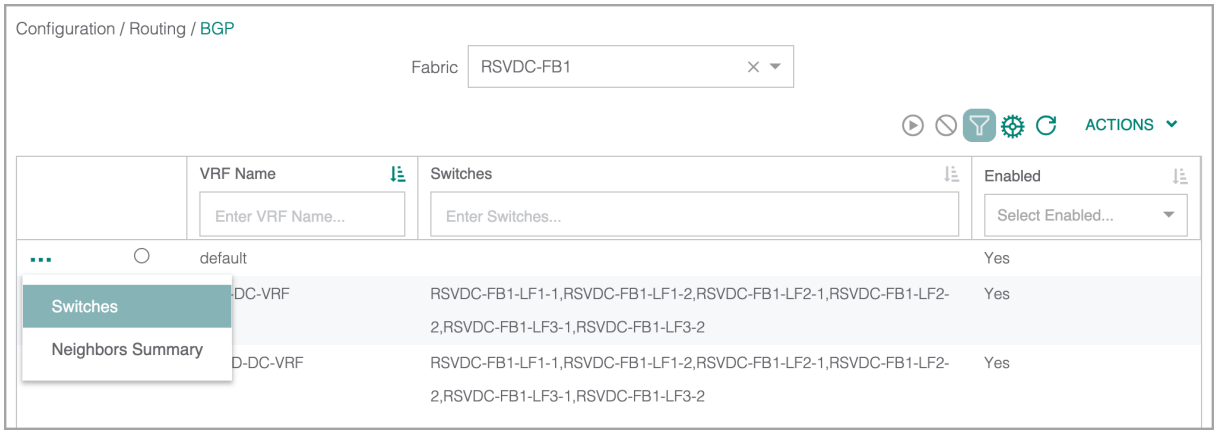


**Step 3** Click the • • • symbol next to **RSVDC-FB1-LF1-1** and select **Neighbors**.

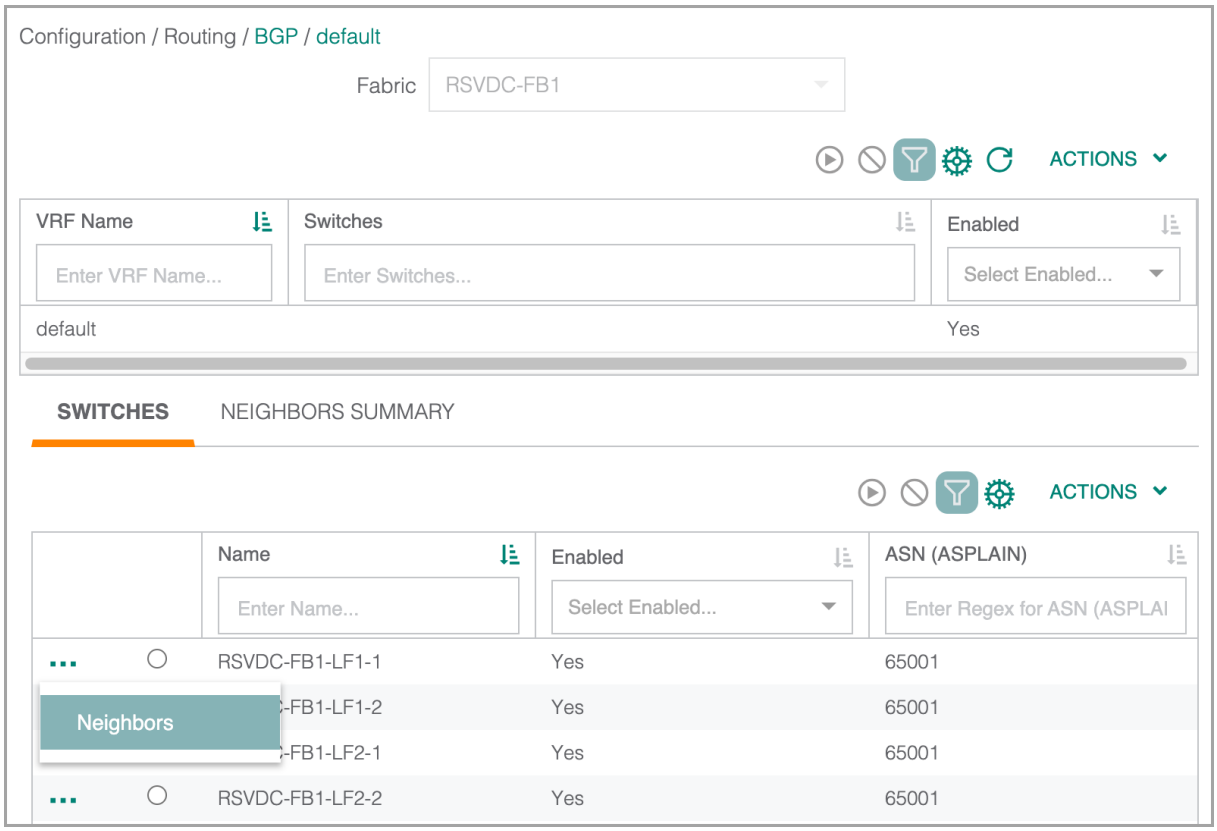


**Step 4** On the lower **ACTIONS** menu, select **Add**.

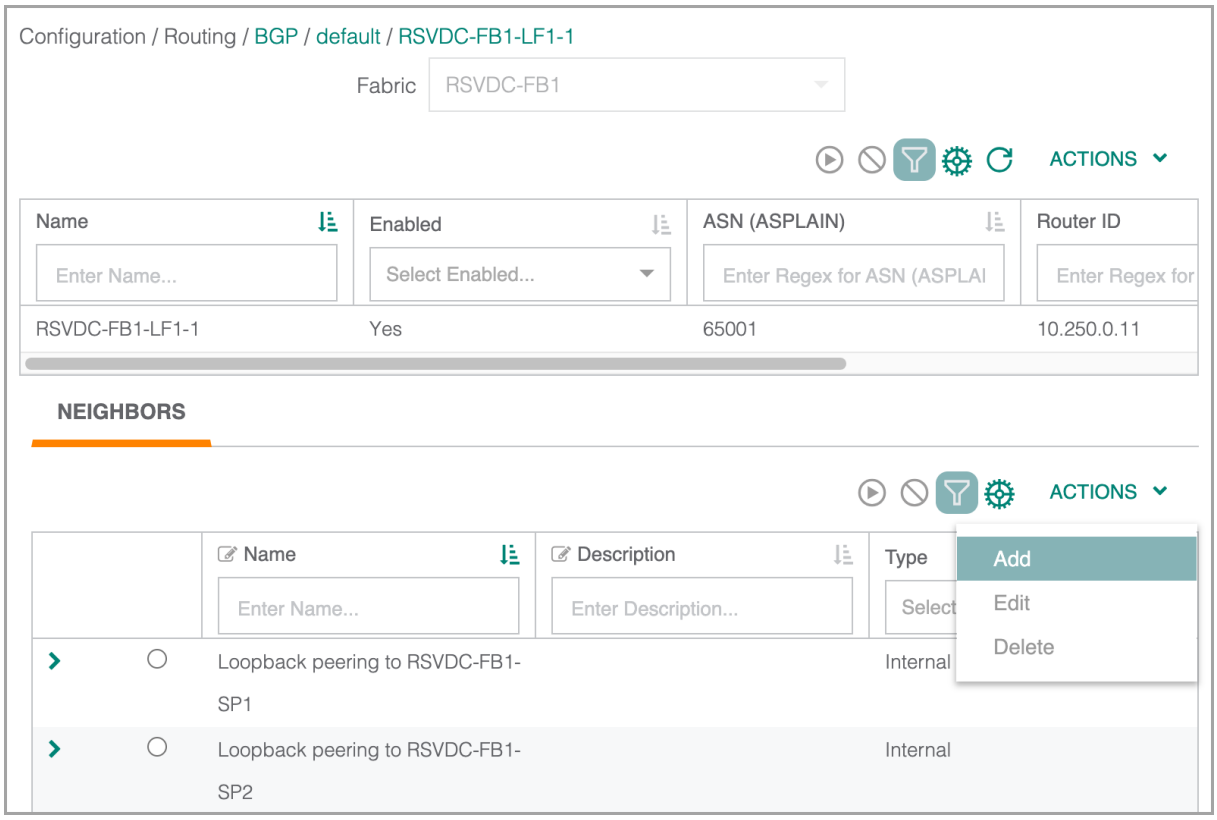


**Step 5** On the **Settings** page, enter the following non-default values and click **NEXT**.

- **Neighbor AS Number:** *65002*

- **IP Address:** *10.255.6.3*
- **Enable Bidirectional Forwarding Detection (BFD) Fall Over**: *< checked >*

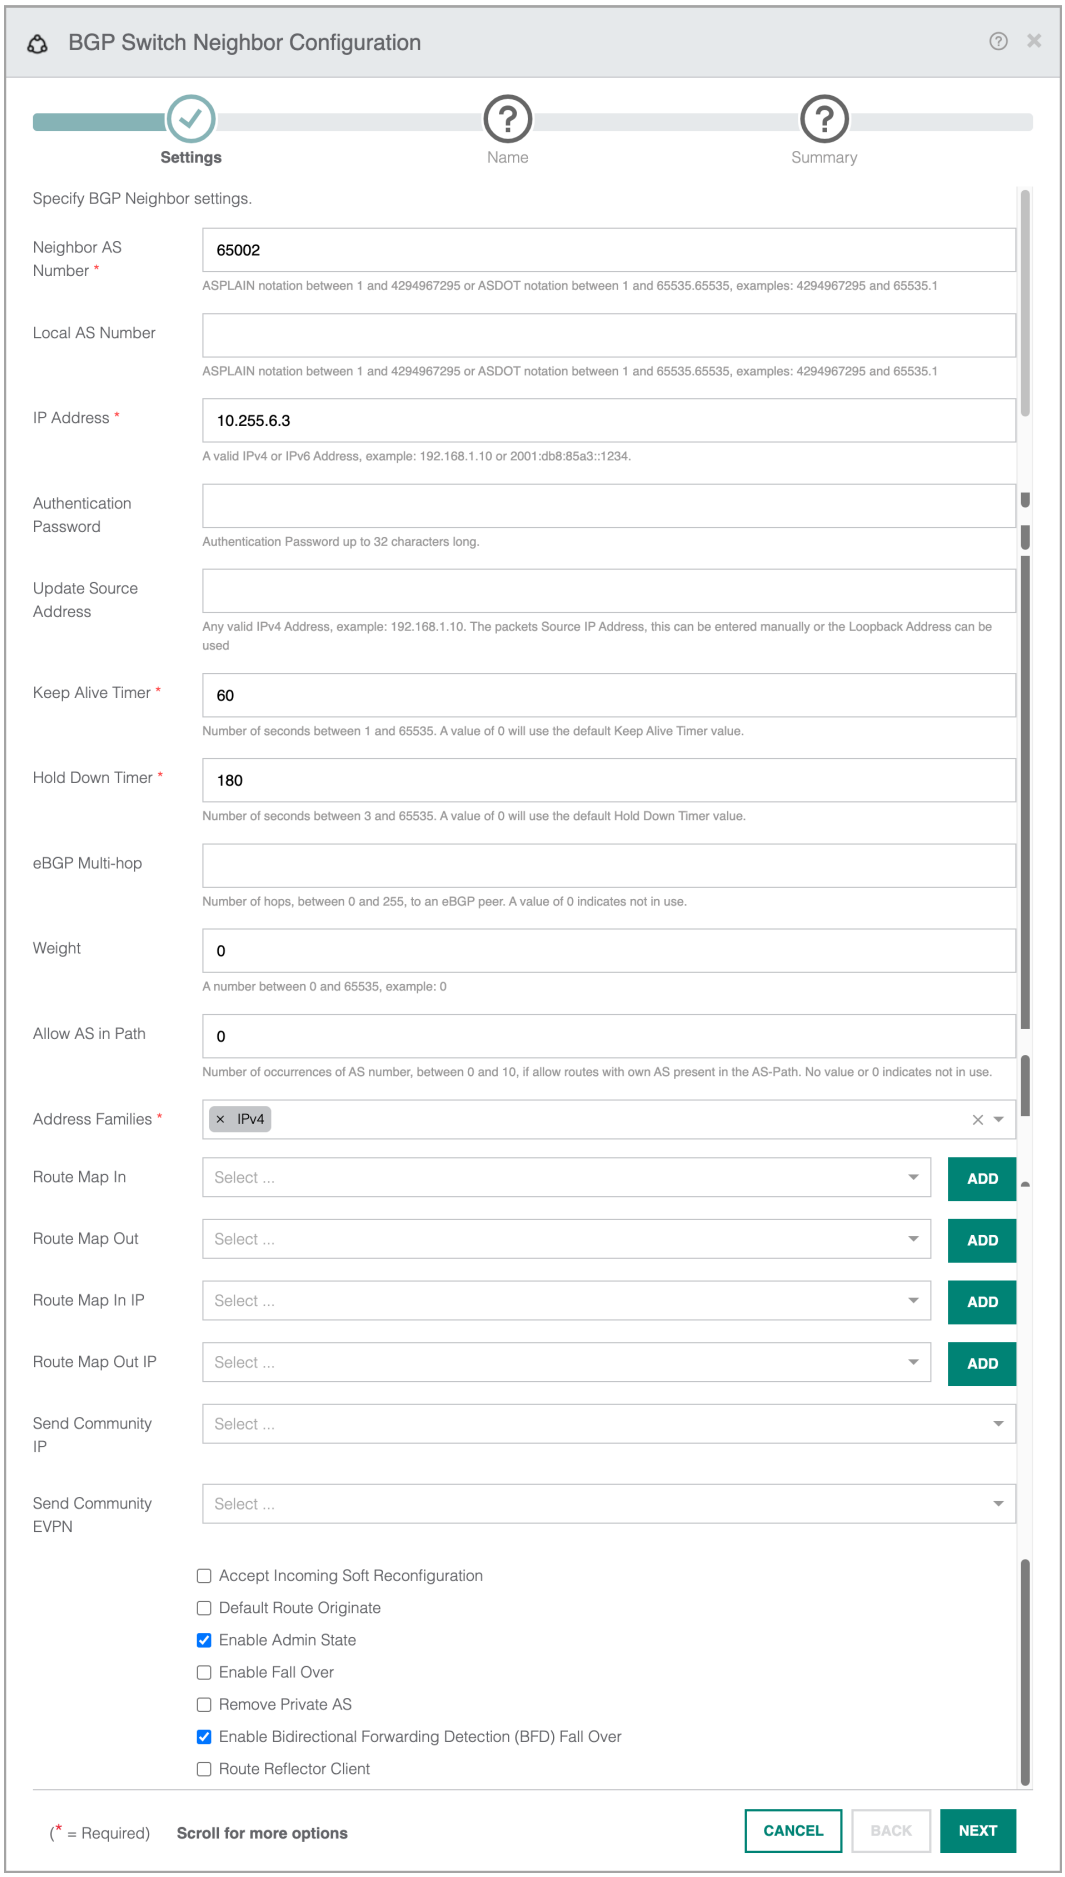
BGP Switch Neighbor Configuration
Settings
Name
Summary
Specify BGP Neighbor settings.
Neighbor AS Number *
65002
ASPLAIN notation between 1 and 4294967295 or ASDOT notation between 1 and 65535.65535, examples: 4294967295 and 65535.1
Local AS Number
ASPLAIN notation between 1 and 4294967295 or ASDOT notation between 1 and 65535.65535, examples: 4294967295 and 65535.1
IP Address *
10.255.6.3
A valid IPv4 or IPv6 Address, example: 192.168.1.10 or 2001:db8:85a3::1234.
Authentication Password
Authentication Password up to 32 characters long.
Update Source Address
Any valid IPv4 Address, example: 192.168.1.10. The packets Source IP Address, this can be entered manually or the Loopback Address can be used
Keep Alive Timer *
60
Number of seconds between 1 and 65535. A value of 0 will use the default Keep Alive Timer value.
Hold Down Timer *
180
Number of seconds between 3 and 65535. A value of 0 will use the default Hold Down Timer value.
eBGP Multi-hop
Number of hops, between 0 and 255, to an eBGP peer. A value of 0 indicates not in use.
Weight
0
A number between 0 and 65535, example: 0
Allow AS in Path
0
Number of occurrences of AS number, between 0 and 10, if allow routes with own AS present in the AS-Path. No value or 0 indicates not in use.
Address Families *
IPv4
Route Map In
Select ...
ADD
Route Map Out
Select ...
ADD
Route Map In IP
Select ...
ADD
Route Map Out IP
Select ...
ADD
Send Community IP
Select ...
Send Community EVPN
Select ...
Accept Incoming Soft Reconfiguration
Default Route Originate
Enable Admin State
Enable Fall Over
Remove Private AS
Enable Bidirectional Forwarding Detection (BFD) Fall Over
Route Reflector Client
(* = Required)
Scroll for more options
CANCEL
BACK
NEXT

**Step 6** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

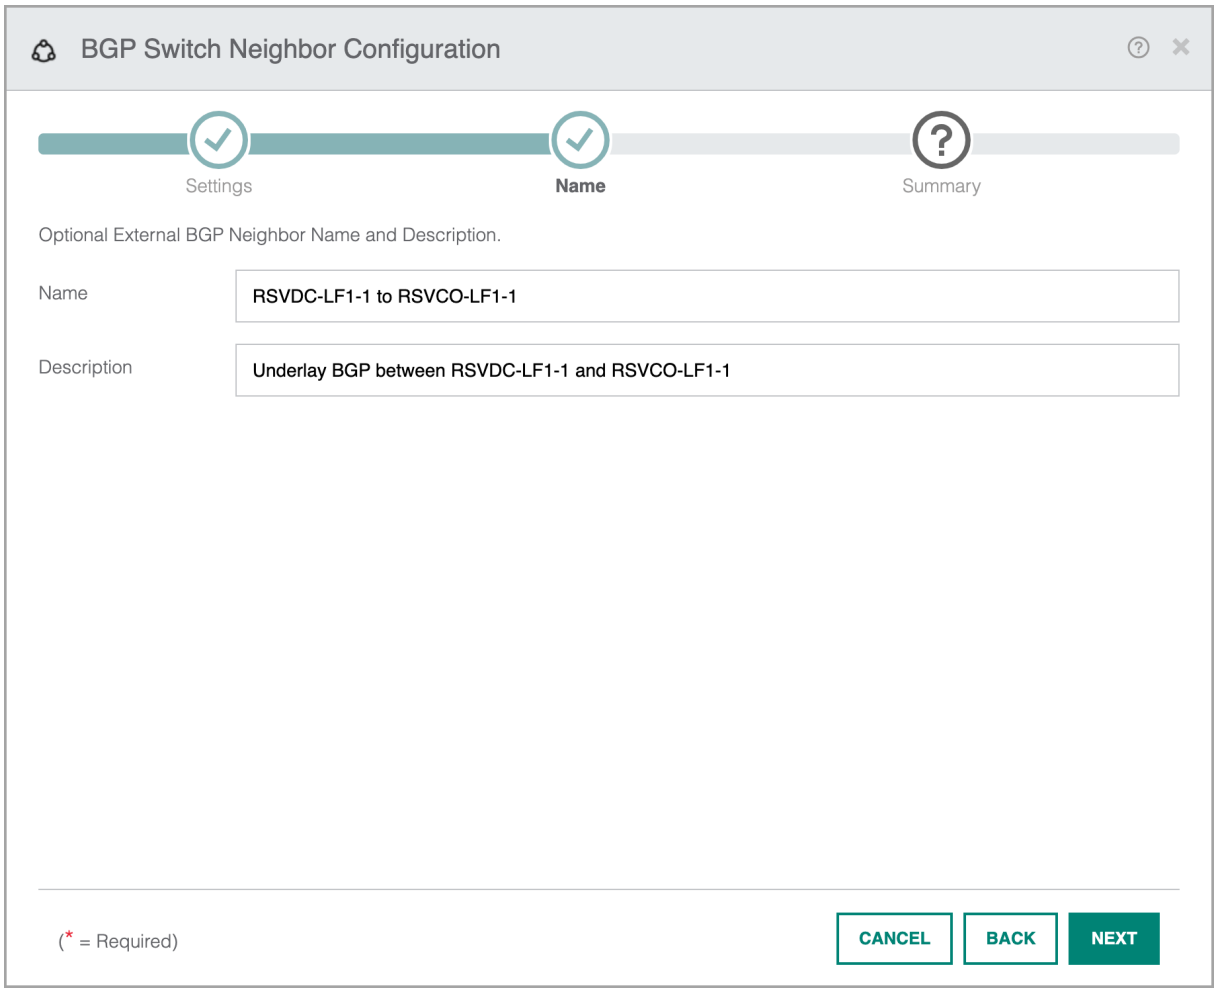


**Step 7** On the **Summary** page, verify the BGP settings and click **APPLY**.

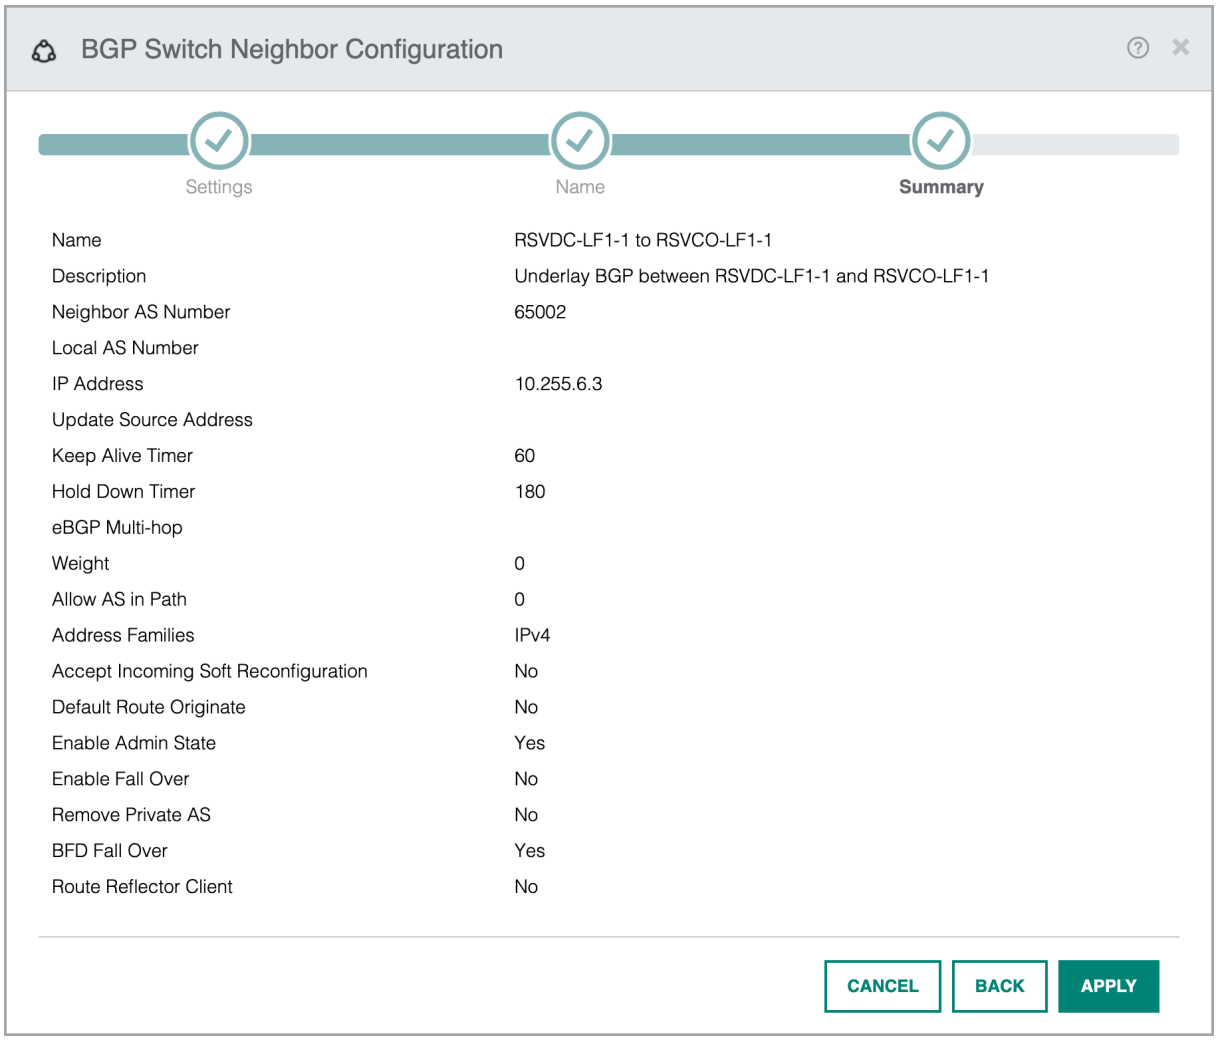


**Step 8** Repeat the procedure to configure underlay eBGP peerings on the remaining border leaf switches in the RSVDC-FB1 fabric.

| Name | Description | Neighbor ASN | IP Addresses |
|---|---|---|---|
| RSVDC-LF1-1 to RSVCO-LF1-2 | Underlay BGP between RSVDC-LF1-1 and RSVCO-LF1-2 | 65002 | 10.255.6.4 |
| RSVDC-LF1-2 to RSVCO-LF1-1 | Underlay BGP between RSVDC-LF1-2 and RSVCO-LF1-1 | 65002 | 10.255.6.3 |
| RSVDC-LF1-2 to RSVCO-LF1-2 | Underlay BGP between RSVDC-LF1-2 and RSVCO-LF1-2 | 65002 | 10.255.6.4 |

**Step 9** Repeat the procedure to configure underlay eBGP peerings on the border leaf switches in the RSVCO-FB2 fabric. In step 2, Select **RSVCO-FB2** in the **Fabric** menu to configure BGP peerings in the second fabric.

| Name | Description | Neighbor ASN | IP Addresses |
|---|---|---|---|
| RSVCO-LF1-1 to RSVDC-LF1-1 | Underlay BGP between RSVCO-LF1-1 and RSVDC-LF1-1 | 65001 | 10.255.6.1 |
| RSVCO-LF1-1 to RSVDC-LF1-2 | Underlay BGP between RSVCO-LF1-1 and RSVDC-LF1-2 | 65001 | 10.255.6.2 |
| RSVCO-LF1-2 to RSVDC-LF1-1 | Underlay BGP between RSVCO-LF1-2 and RSVDC-LF1-1 | 65001 | 10.255.6.1 |
| RSVCO-LF1-2 to RSVDC-LF1-2 | Underlay BGP between RSVCO-LF1-2 and RSVDC-LF1-2 | 65001 | 10.255.6.2 |

**Step 10** Click **BGP** in the left navigation pane and select **RSVDC-FB1** for the **Fabric** field.

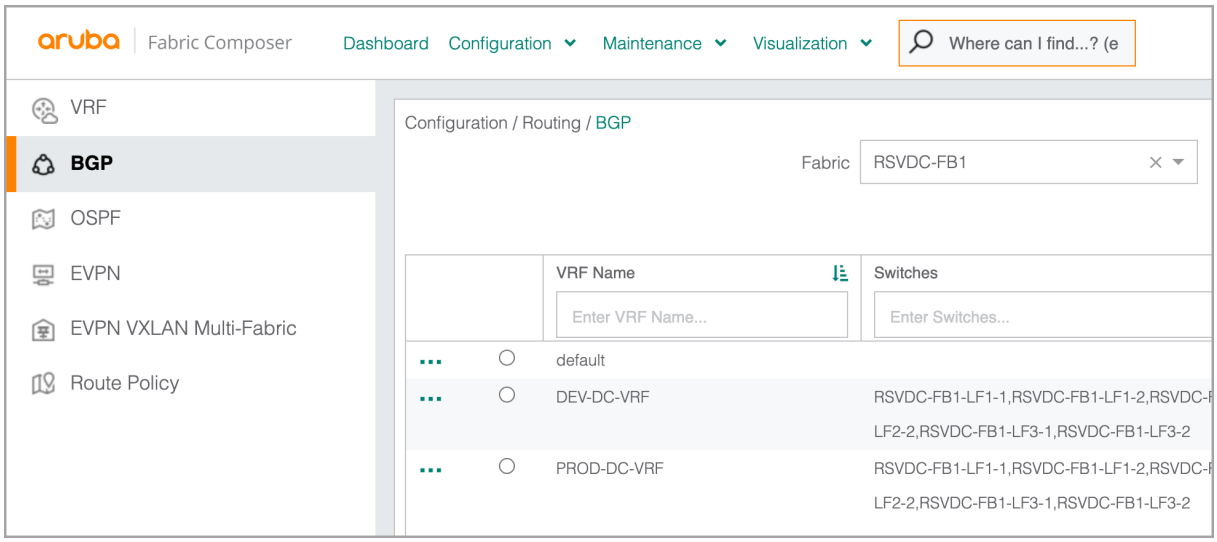


**Step 11** Click the • • • symbol next to **default** and select **Neighbors Summary**.

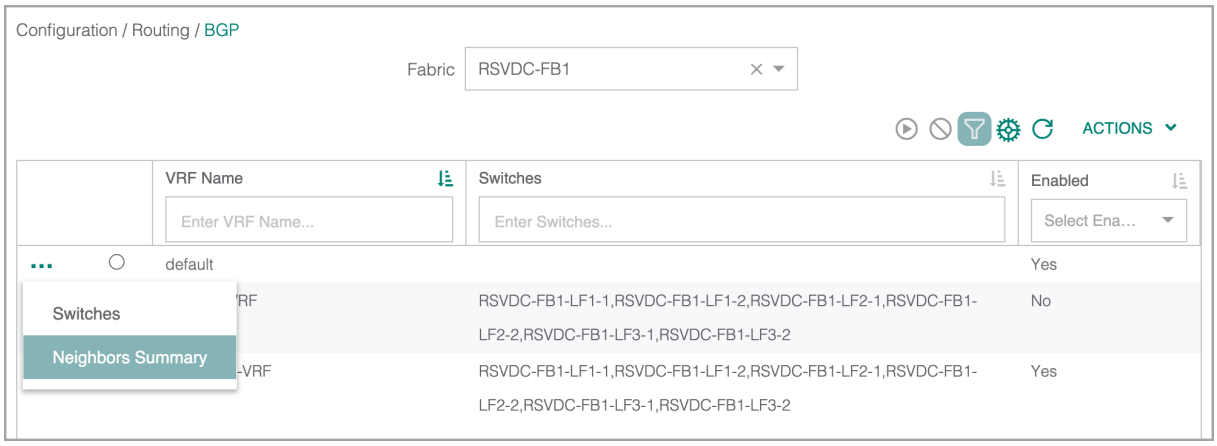


**Step 12** In the **Address Family** column filter, select **IPv4 Unicast**. Click the **Apply table filters** (arrow) icon.

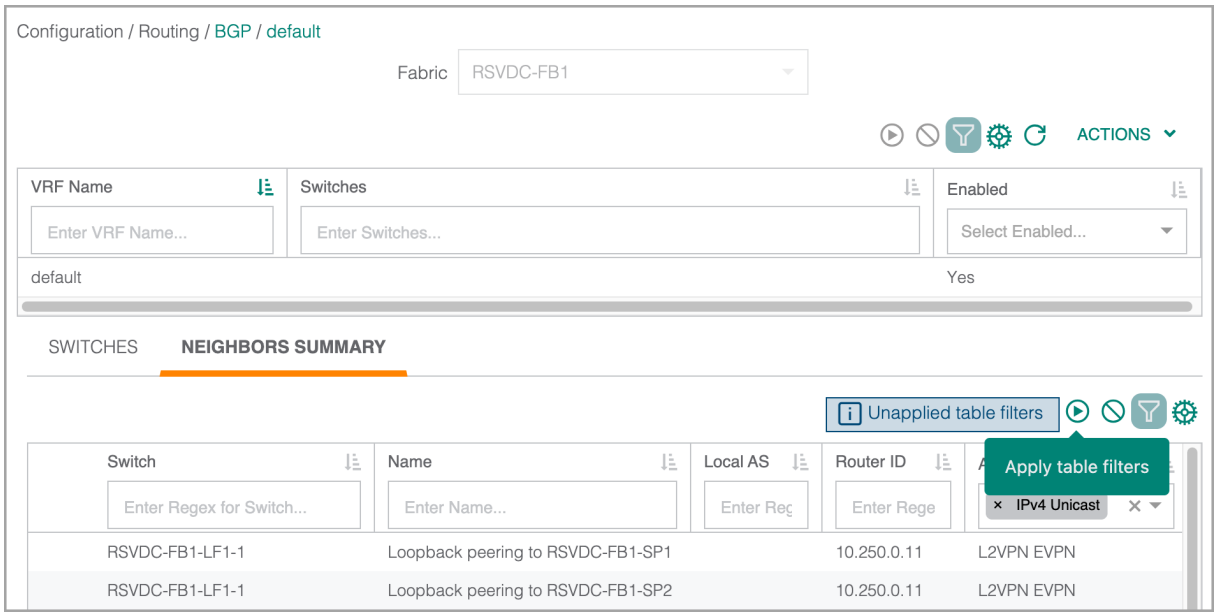


**Step 13** Verify that each BGP session displays **Established** in the **State** column.

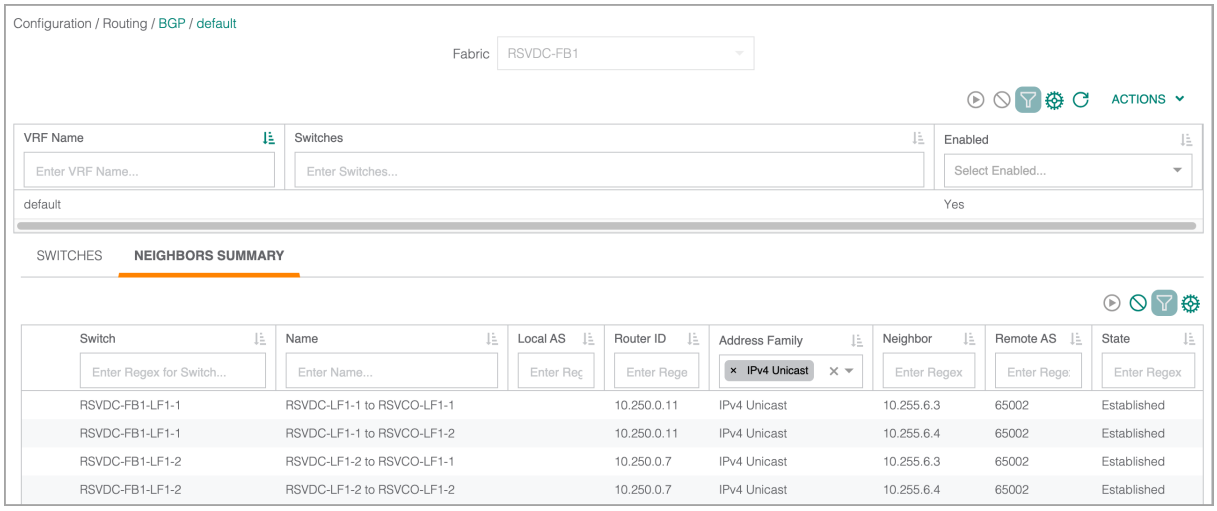


# Configure Overlay Control Plane

The overlay control plane uses MP-BGP EVPN advertisements to share host MAC and IP reachability across both fabrics. The diagram below illustrates the eBGP EVPN sessions established to share overlay reachability information.

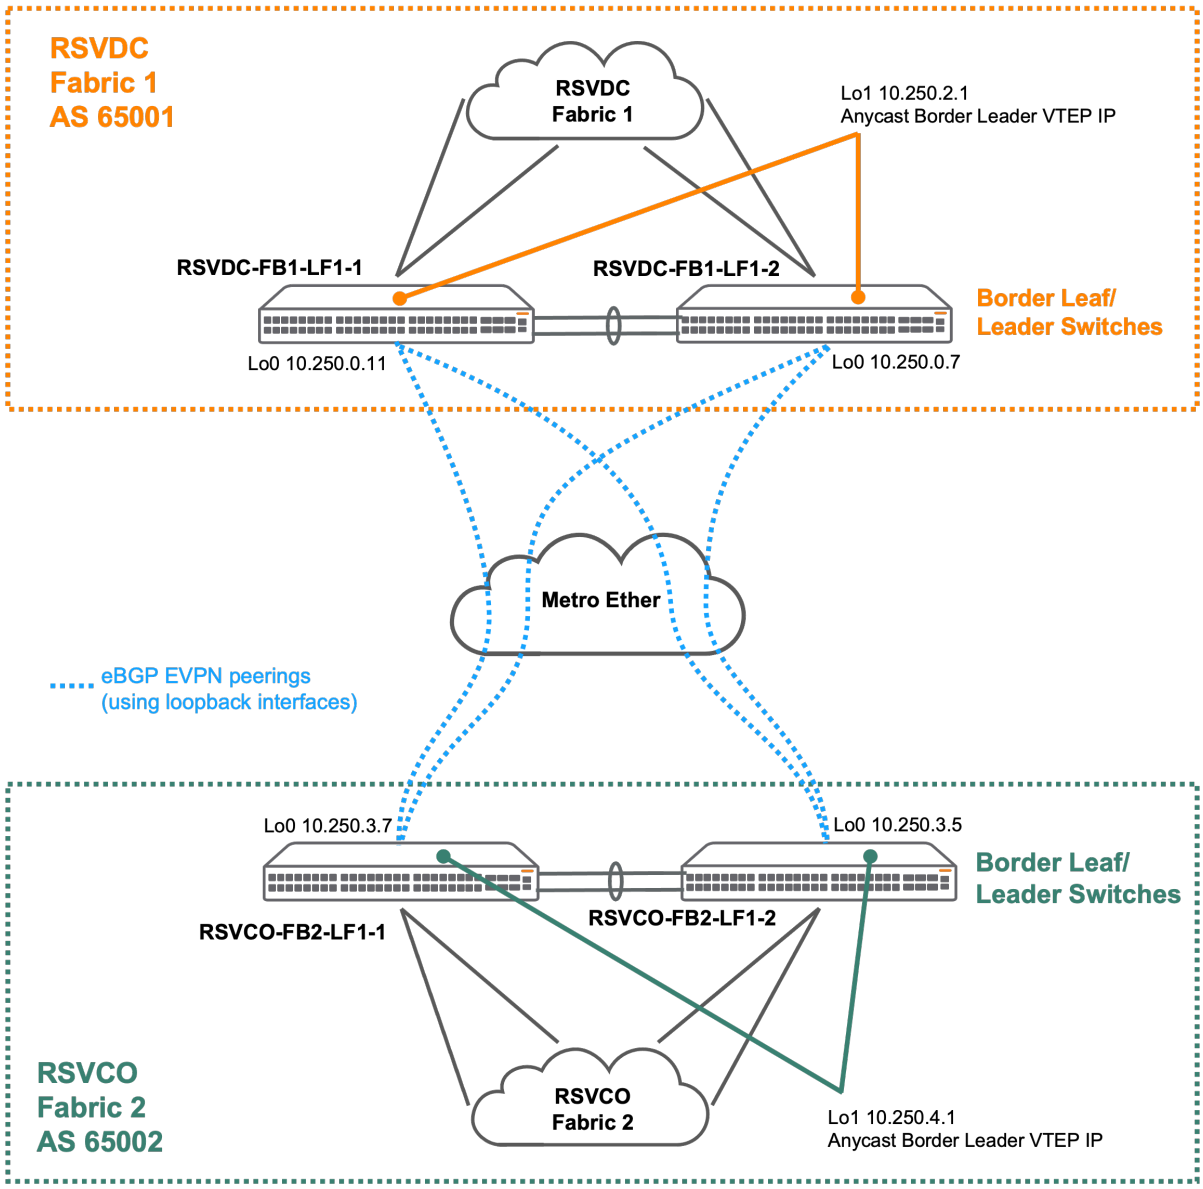


Each site contains one border leader VTEP. Border leader VTEP switches establish eBGP EVPN address-family peerings with border leaders in other sites. A full mesh of peerings is established between sites.

When more than one fabric is present in a single site, the border leader VTEP switches also establish eBGP EVPN address-family peerings with border leaf switches of each fabric in the same site.

**NOTE:**

The border leader is a control plane role that optimizes the number of eBGP sessions required to share EVPN reachability information between sites, when any individual site contains multiple fabrics. The multifabric VXLAN data plane still establishes a full mesh of VXLAN tunnels between between border leaf VTEPs.

## Multifabric VXLAN Tunnel Requirements

Overlay host traffic within a fabric only requires encapsulation in a single VXLAN tunnel, because a full mesh of tunnels is established between VTEPs. In a multifabric environment, a full mesh of tunnels between all leaf switches is not present. Traffic between fabrics is enabled using a VXLAN tunnel between the border leaf switches of each fabric. In this model, overlay host traffic between fabrics may traverse multiple VXLAN tunnels: VXLAN tunnels internal to a fabric and the inter-fabric VXLAN tunnel.

## Join Fabric Overlay Control Planes Together

The AFC **EVPN VXLAN Multi-Fabric** wizard configures eBGP EVPN peerings between the border leader switches of each fabric. Supporting prefix lists, route maps, OSPF redistributions, and BGP redistributions also are configured. VXLAN forwarding between iBGP and eBGP learned VXLAN tunnels is also enabled. These peerings share overlay host and prefix reachability between fabrics and are established between border leader loopback addresses.

AFC adds configuration to share the complete set of VTEP IPs in both fabrics to assist with troubleshooting.

**Step 1** On the left navigation pane, click **EVPN VXLAN Multi-Fabric**.

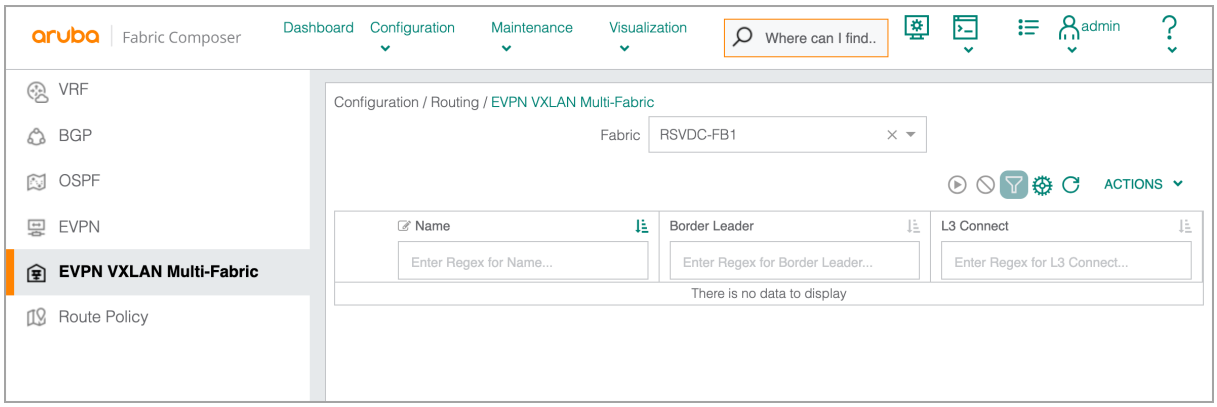


**NOTE:**

If the left pane no longer displays routing options, select **Configuration > Routing > EVPN VXLAN Multi-Fabric** on the top menu.

**Step 2** Verify that **RSVDC-FB1** is selected in the **Fabric** menu. On the **ACTIONS** menu, select **Add**.

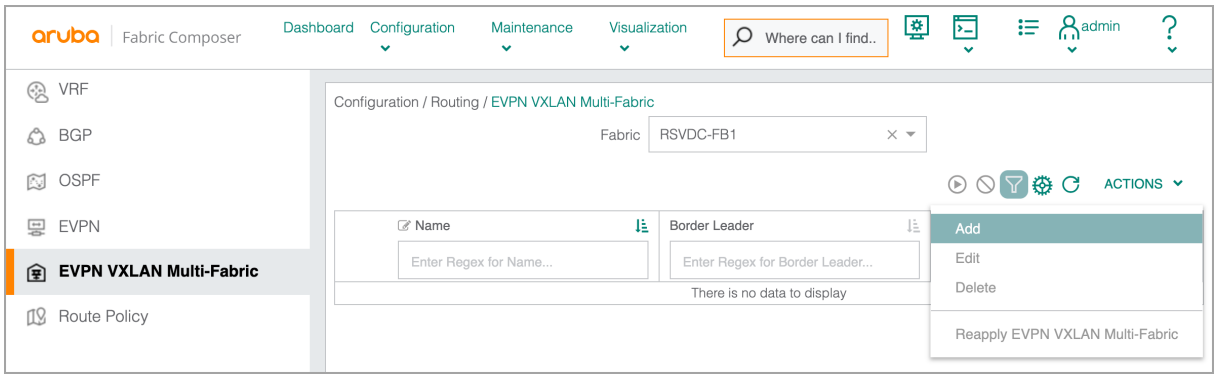


**Step 3** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

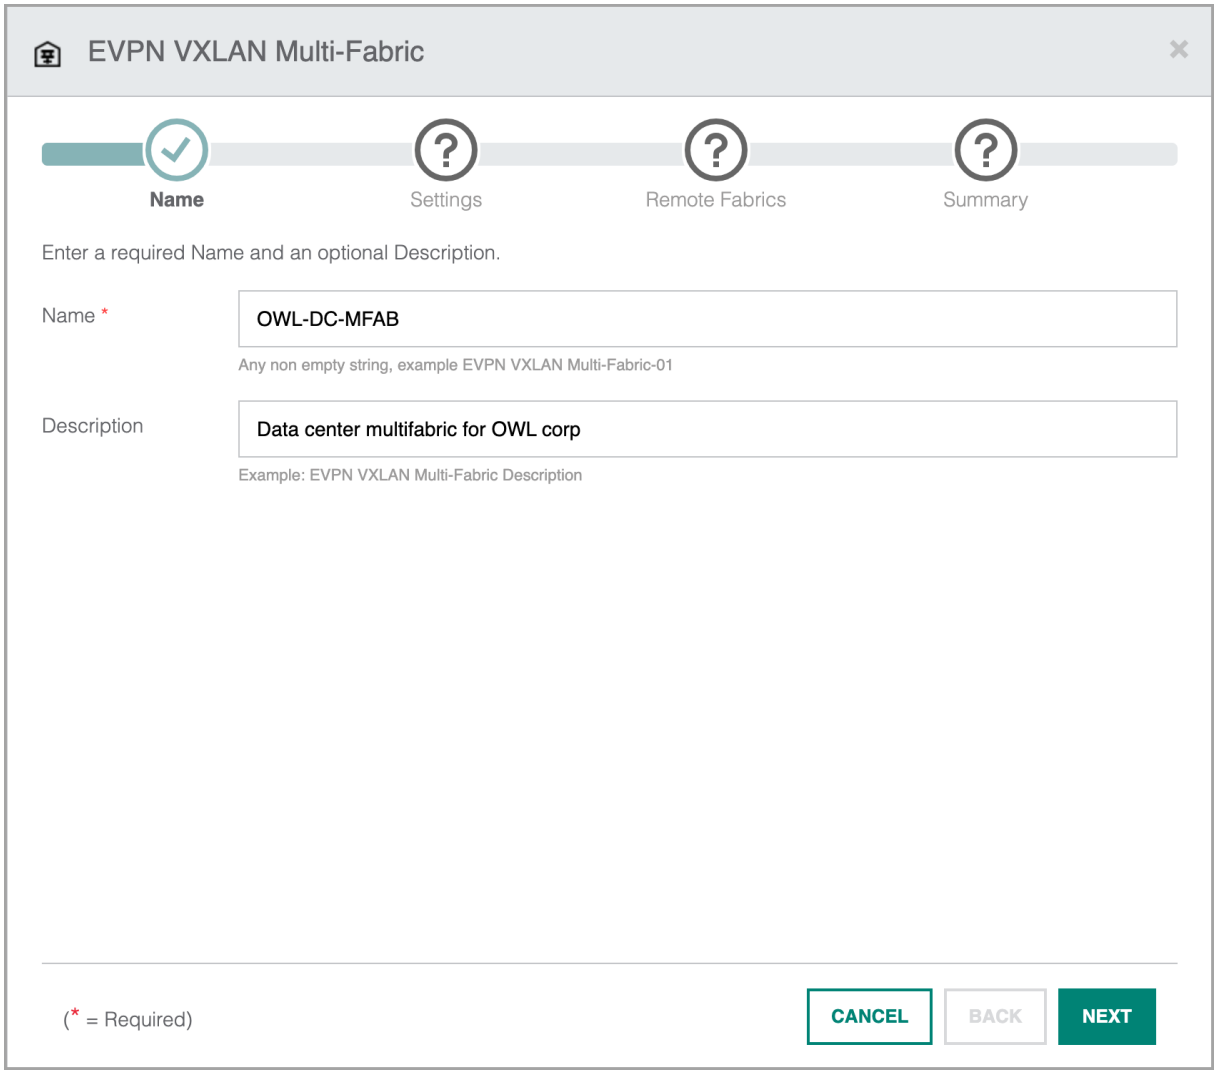


**Step 4** On the **Settings** page, select the border leader VSX pair in the **Border Leader** field, verify that both border leader switches are listed in the **L3 Connect** field, and click **NEXT**.

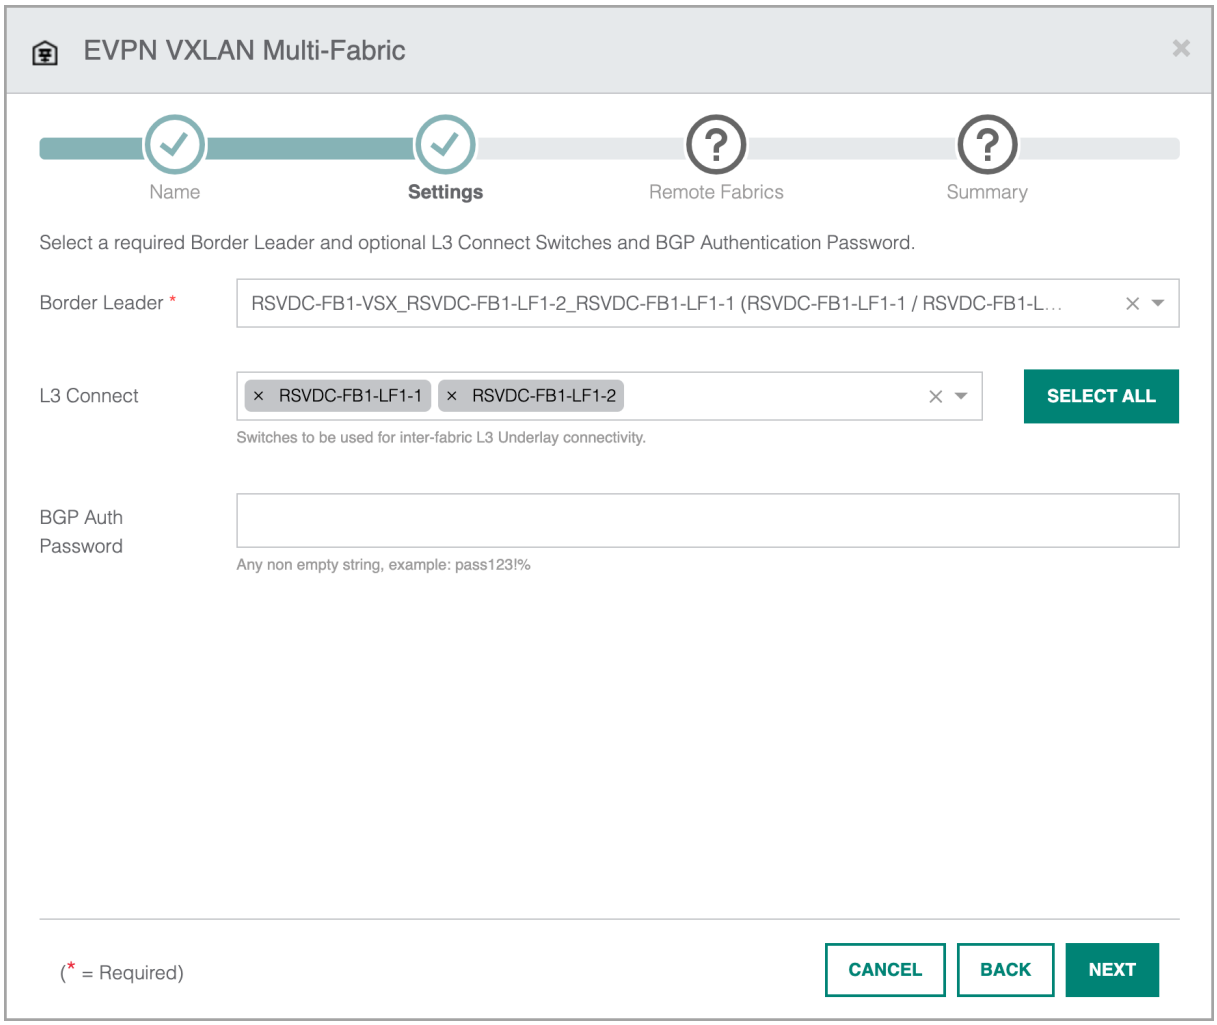


**Step 5** On the **Remote Fabrics** page, select **Local AFC** in the **AFC Site** field and **RSVCO-FB2** in the **Fabric** field. Verify that the auto-populated values for **AS Number**, **Remote Border Leader Address**, and **Secondary Remote Border Leader Address** are correct. Click **Add**.

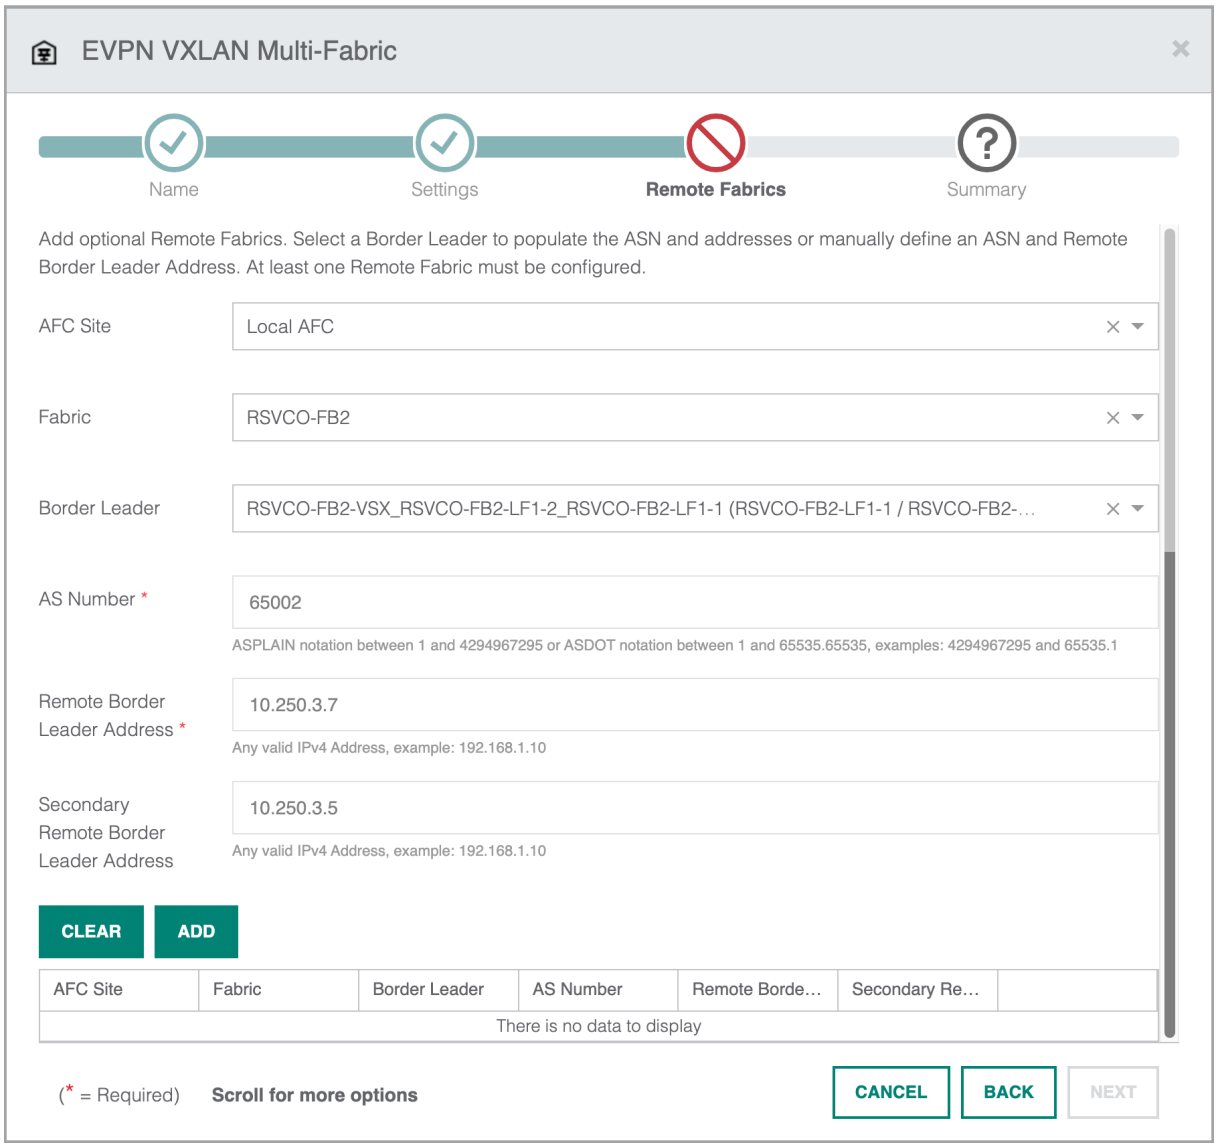


**NOTE:**

More than one fabric can be added on this page. Fabrics managed by remote AFC instances can be selected when the remote AFC instance is configured as an AFC Site on the local AFC application.

**Step 6** Click **NEXT**.

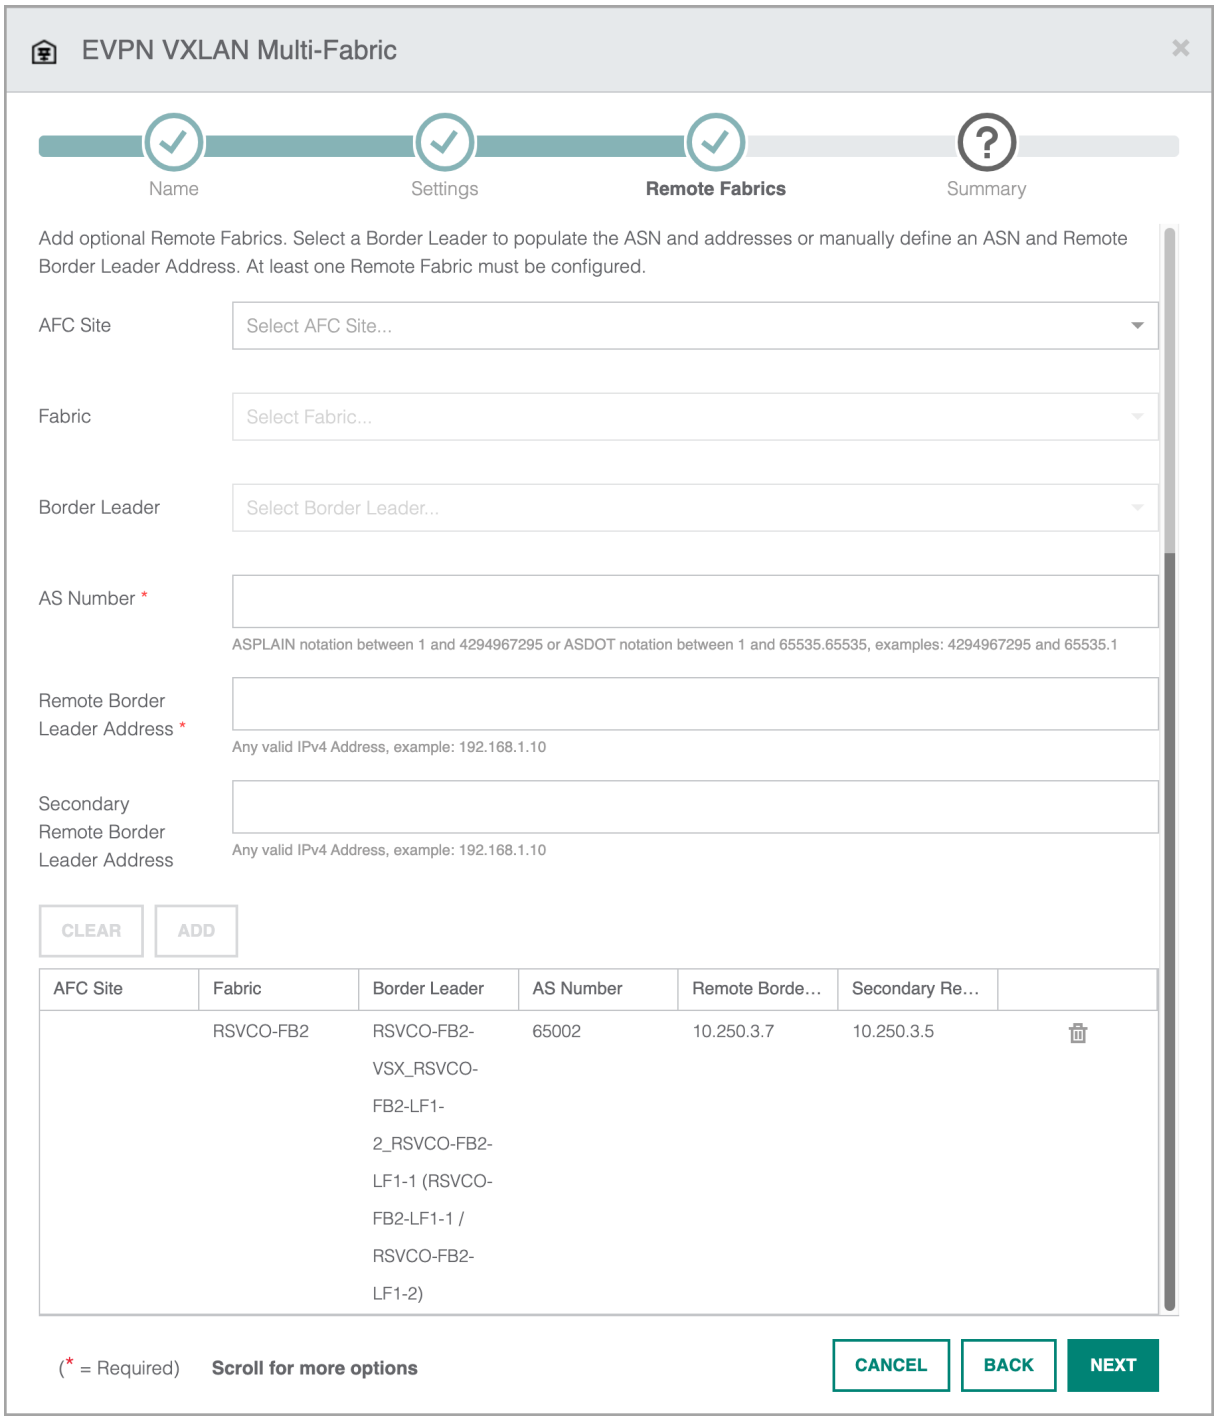


**Step 7** On the **Summary** page, verify the multifabric BGP settings and click **APPLY**.

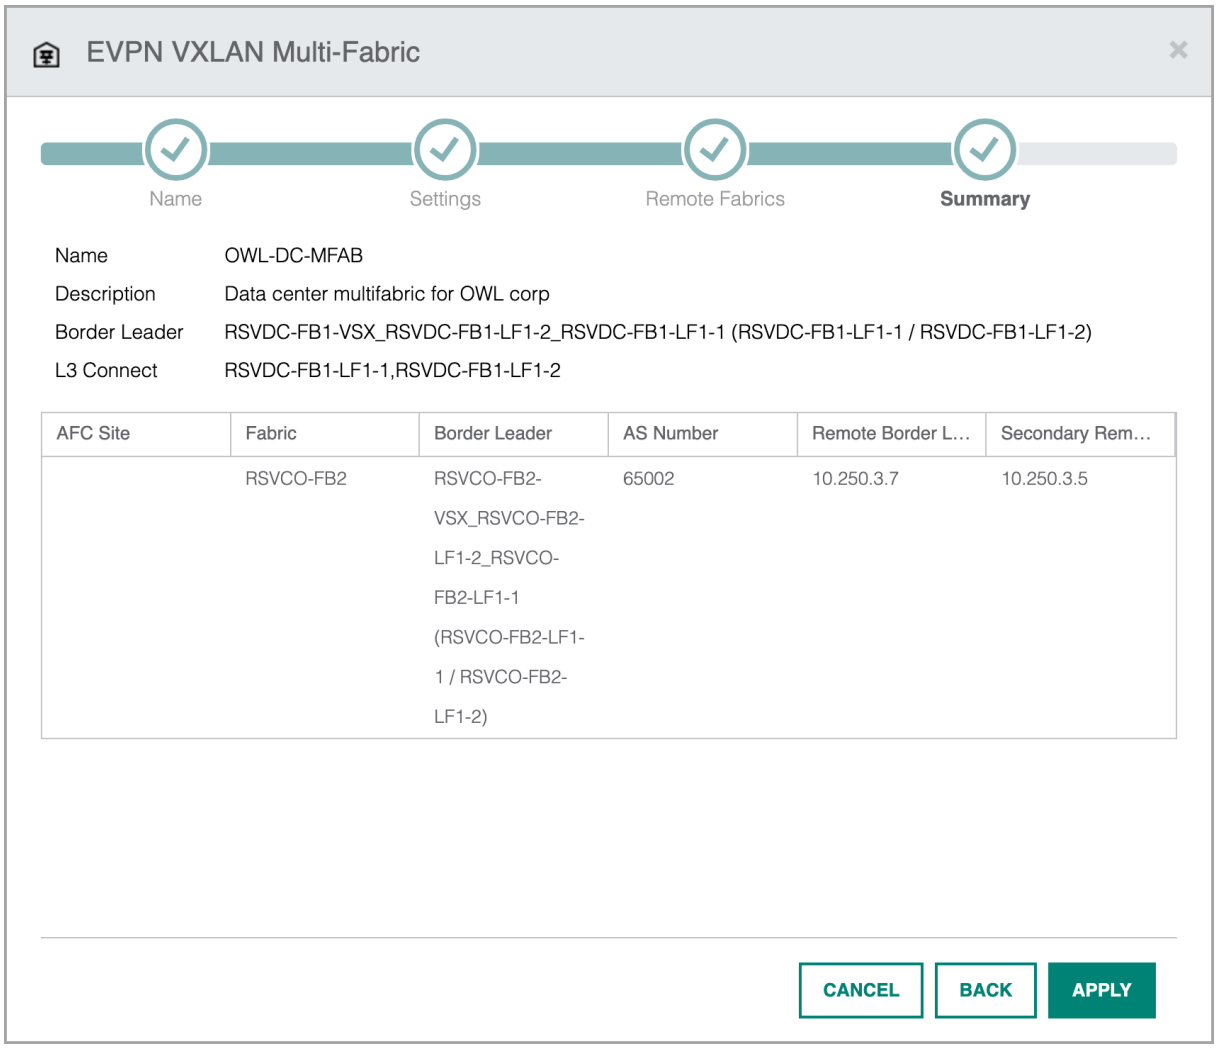


**Step 8** Repeat the procedure to configure eBGP EVPN peerings in the second fabric. At step 2, select **RSVCO-FB2** on the **Fabric** menu.

**Step 9** In the menu bar at the top right of the AFC display, click the **CLI Commands** icon and select **Show Commands**.

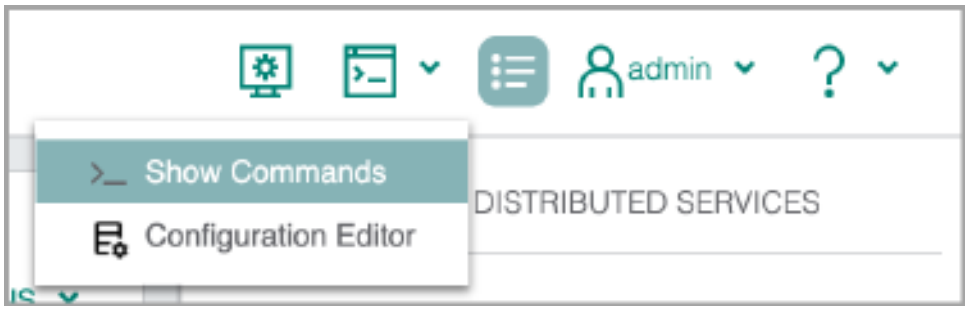


**Step 10** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** *RSVDC-FB1-LF1-1*, *RSVDC-FB1-LF1-1*, *RSVCO-FB2-LF1-1*, *RSVCO-FB2-LF1-2*
- **Commands:** *show bgp l2vpn evpn summary*

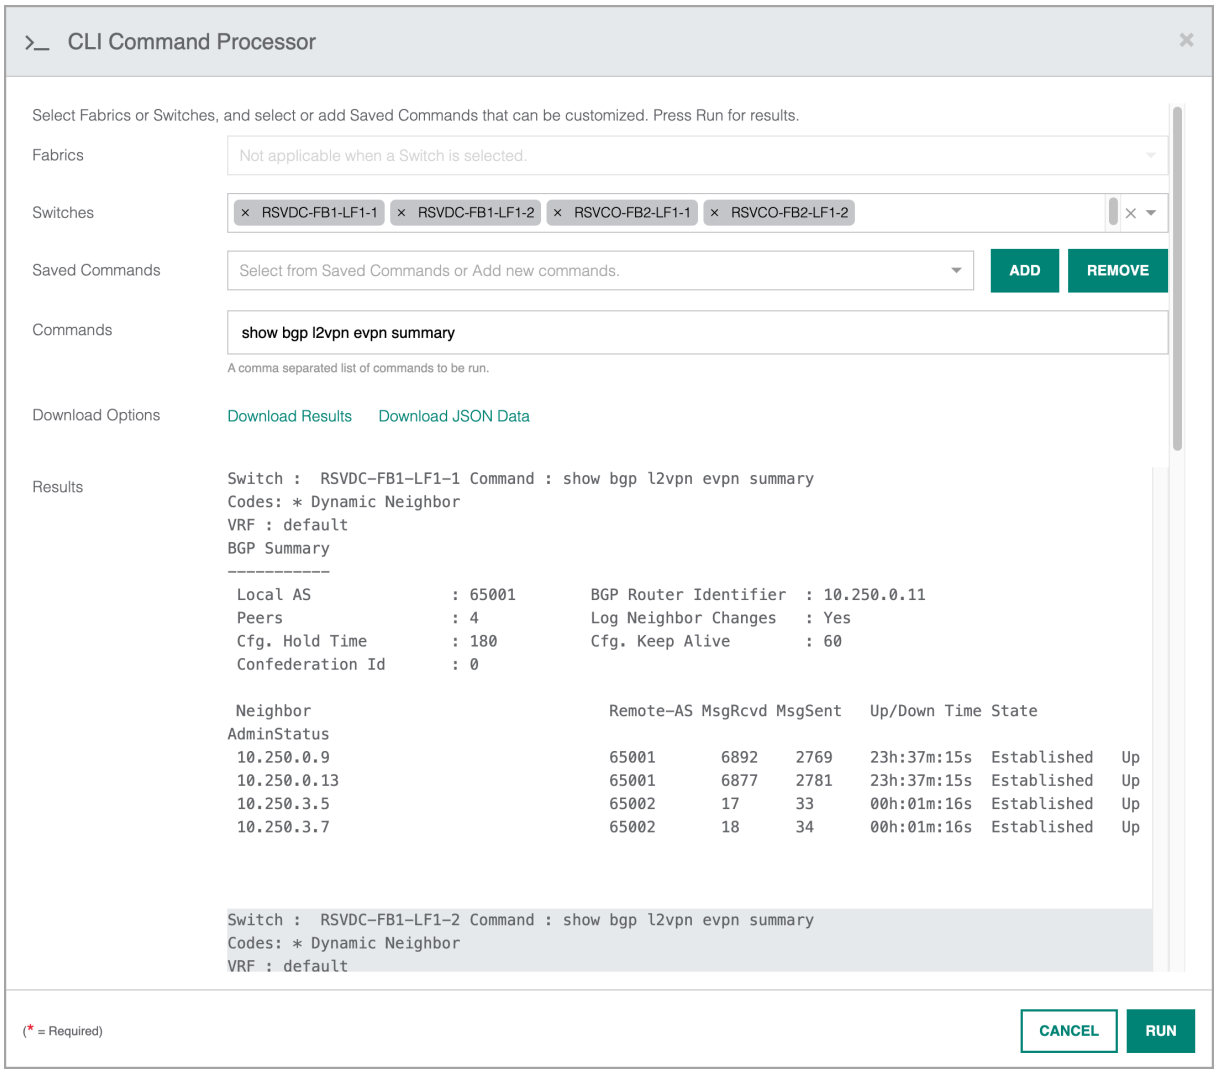


**Step 11** Verify that the output for each switch displays the newly configured peers with an *Established* state.

The eBGP EVPN route peerings establish the control plane mechanism to share overlay reachability between fabrics. Additional steps are necessary to install IP prefixes and MAC addresses between fabrics.

## Extend Layer 3 Reachability Across Fabrics

When using EVPN, MP-BGP EVPN route-type 5 advertisements share IP prefix reachability for a VRF. Route targets included in these advertisements control how Layer 3 IP addresses are installed in IP forwarding tables. Route targets defined during the creation of each fabric are typically used only to share information locally. New VRF route-targets are defined to share IP reachability between fabrics.

**Step 1** On the left navigation pane, click **VRF**.

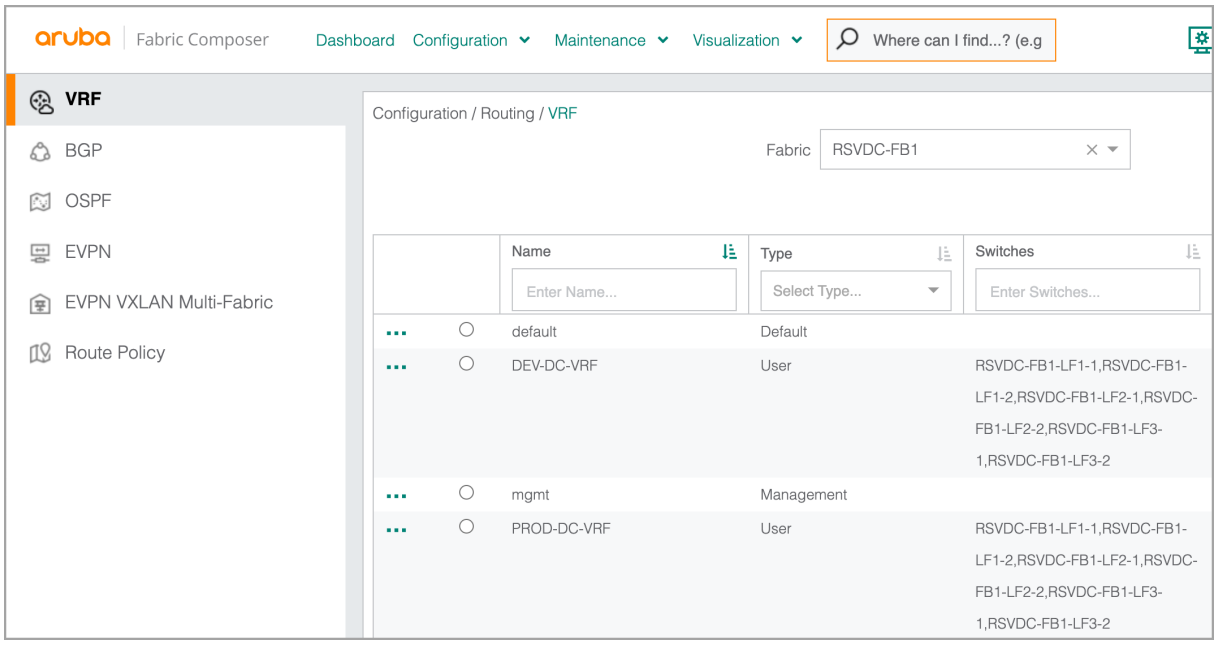


**NOTE:**

If the left pane no longer displays routing options, select **Configuration > Routing > VRF** on the top menu.

**Step 2** Select **RSVDC-FB1** in the **Fabric** field. Click the radio button next to **PROD-DC-VRF**. On the **ACTIONS** menu, select **Edit**.

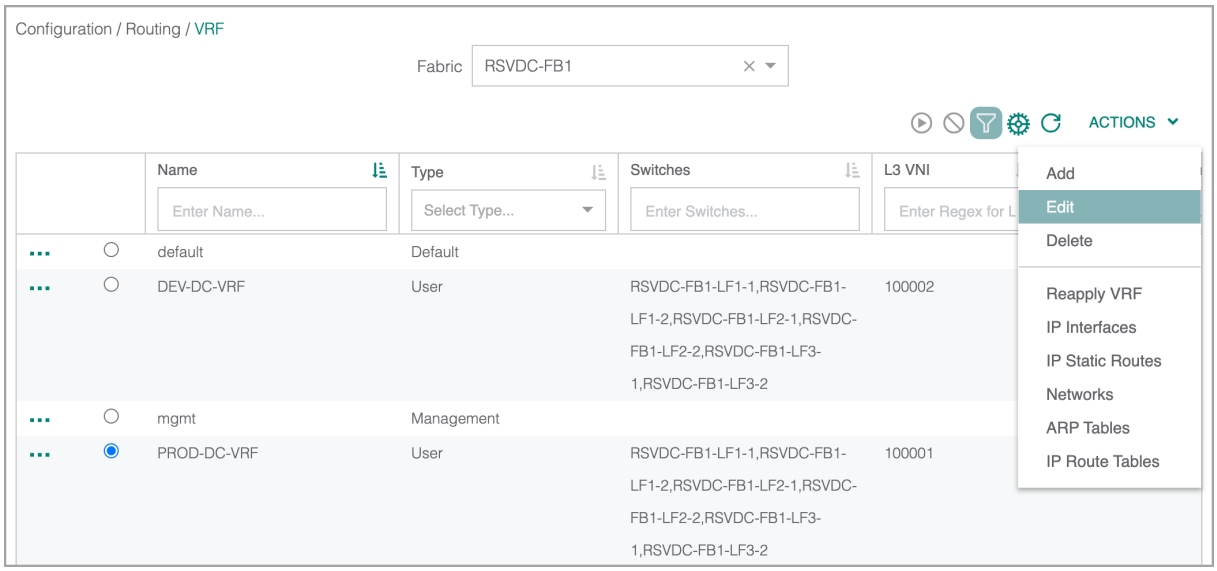


**Step 3** Click the the **ROUTE TARGETS** page heading. On the page, enter the values below and click **ADD**.

- **Route Target Mode:** *Both*
- **Route Target Ext-Community:** *1:100001*
- **Address Family:** *EVPN*

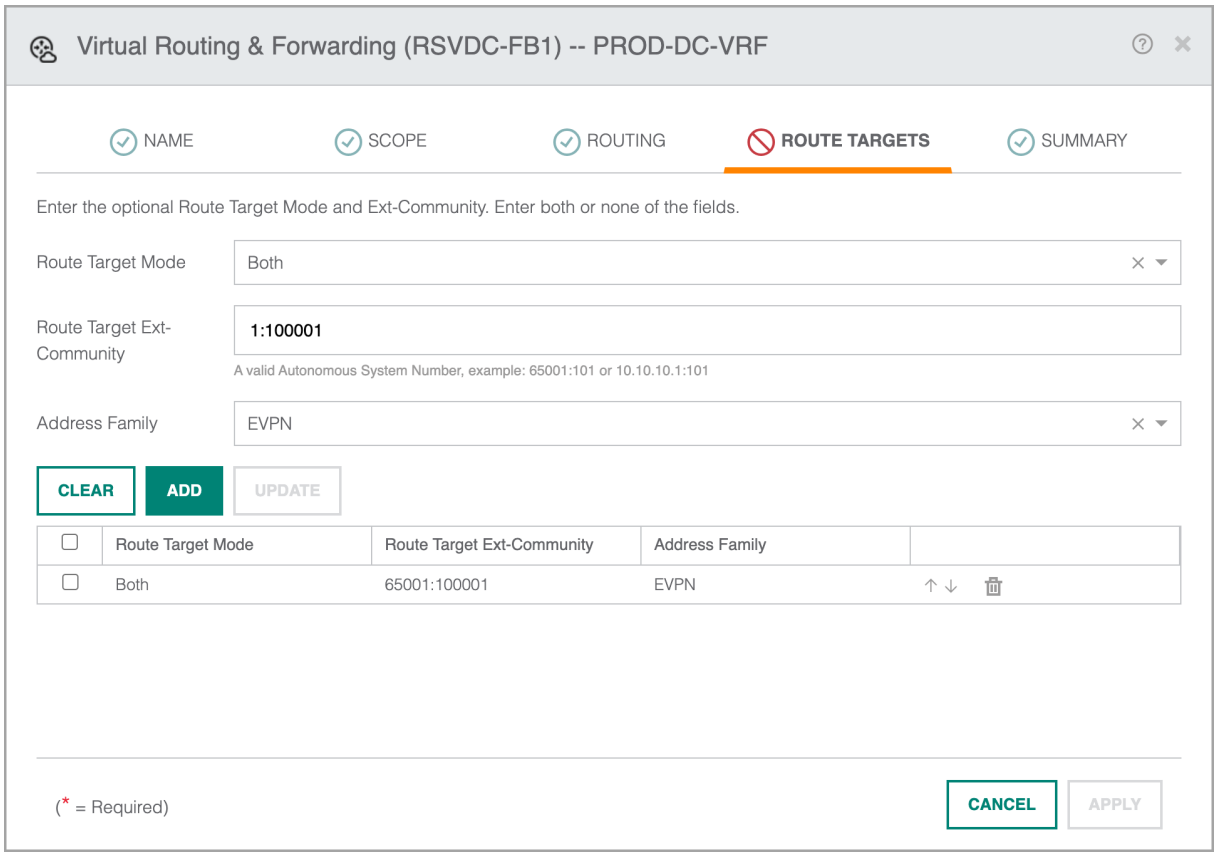


**NOTE:**

It is best practice to use a route target between fabrics that is distinct from the route target within a fabric to provide flexibility in controlling the installation of overlay IP prefixes.

**Step 4** Click **Apply**.

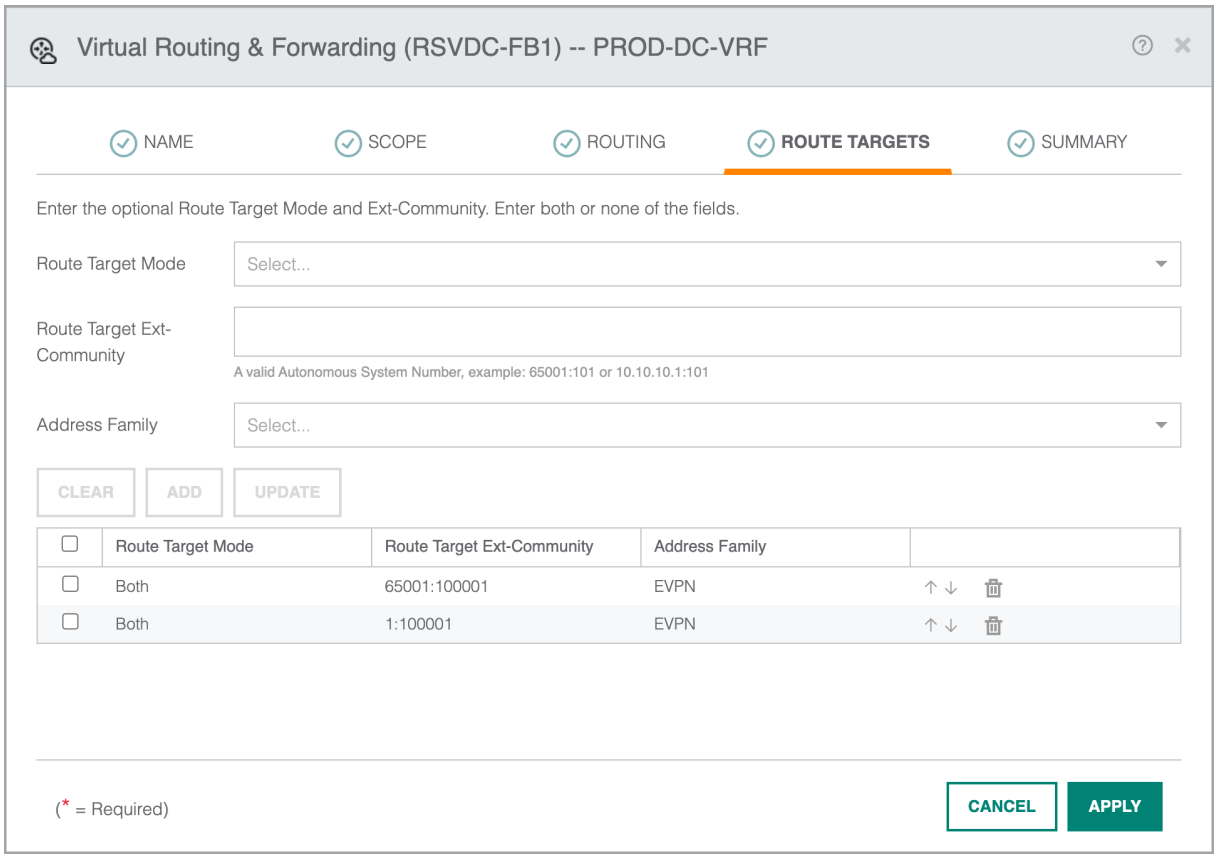


**Step 5** Repeat the procedure to assign the new route target to the RSVCO-FB2 PROD-DC VRF. Select **RSVCO-FB2** in the **Fabric** field to begin.

## Enable External Layer 3 Multifabric Advertisements

The AFC default configuration applies a route map that permits only local fabric prefixes to be shared across fabrics. The AFC created route map named *to-border-leaders* must be modified, if IP prefixes learned in the overlay from outside the fabric must be shared.

In this sample deployment, a campus summary prefix learned in PROD-DC-VRF of the RSVDC-FB1 fabric is shared with PROD-DC-VRF in the RSVCO-FB2 fabric. The default route learned in the same VRF also is shared as a backup default route if Internet connectivity in the colocation facility fails.

A new route map is created for the RSVDC-FB1 border leaders and applied to its eBGP EVPN peerings to allow advertising of routes learned from outside the fabric.

### *Create Route Map*

A new route map is created to allow local fabric reachability and enable advertising of campus-learned IP prefixes from the RSVDC-FB1 border leaders to the RSVCO-FB2 fabric. The route map **ALLOWED-EXT-AS** defined in the initial EVPN-VXLAN configuration is re-used in this procedure.

**Step 1** Click the **ROUTE MAPS** tab. On the **ACTIONS** menu, select **Add**.

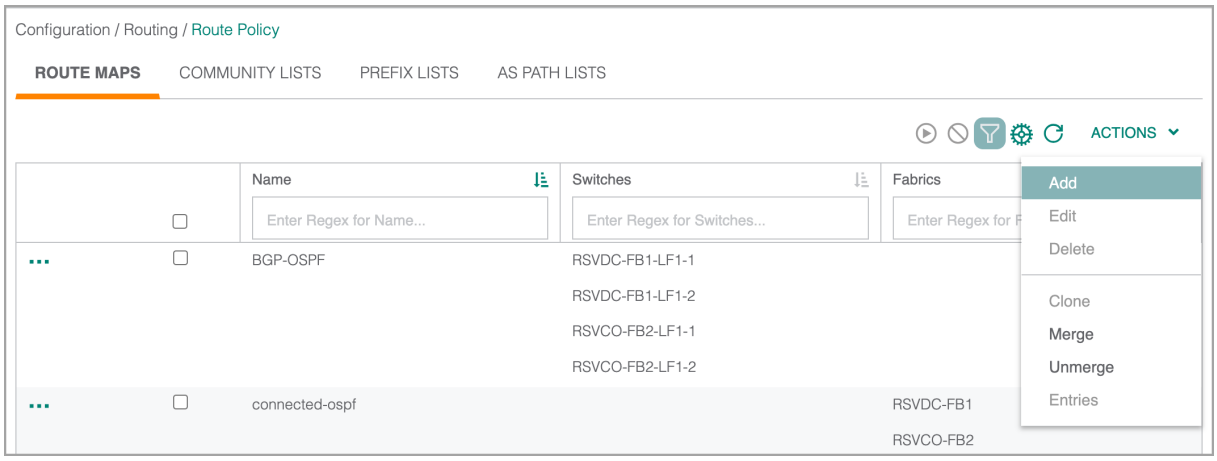


**Step 2** On the **Name** page, enter a **Name** and **Description**, then click **NEXT**.

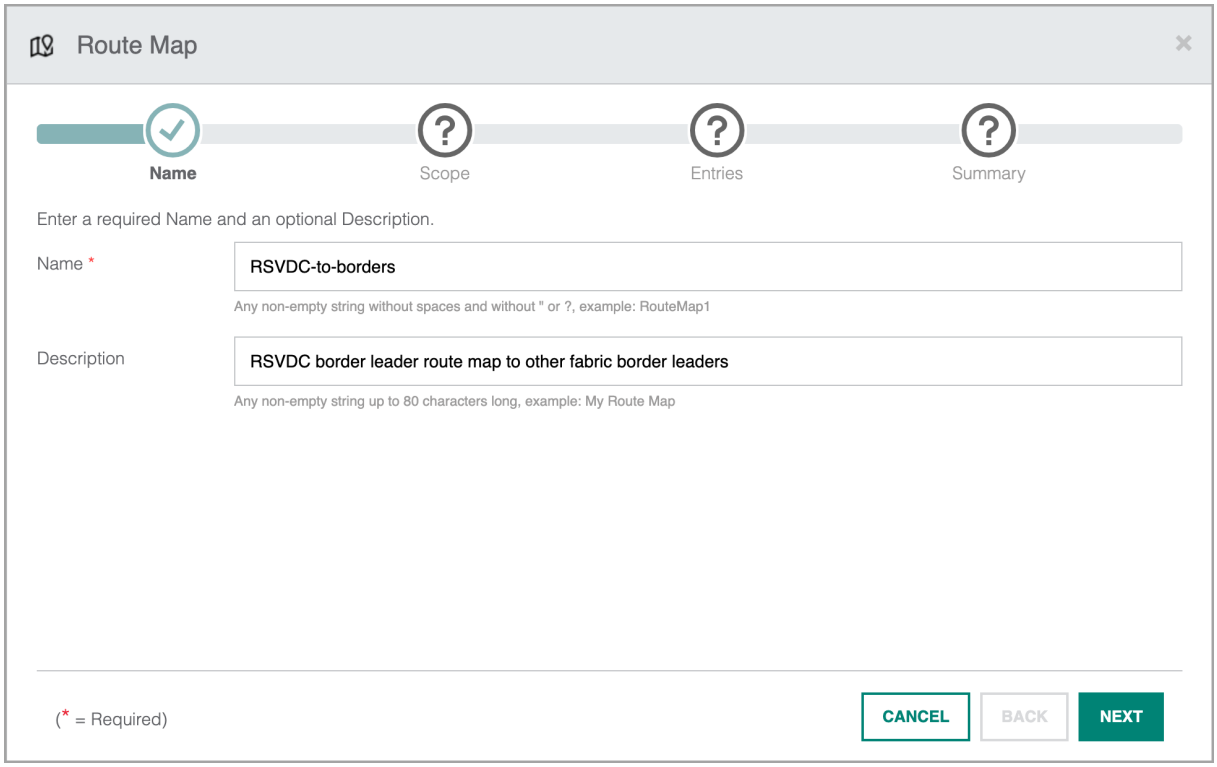


**Step 3** On the **Scope** page, select the RSVDC-FB1 border leaders in the **Switches** field and click **NEXT**.

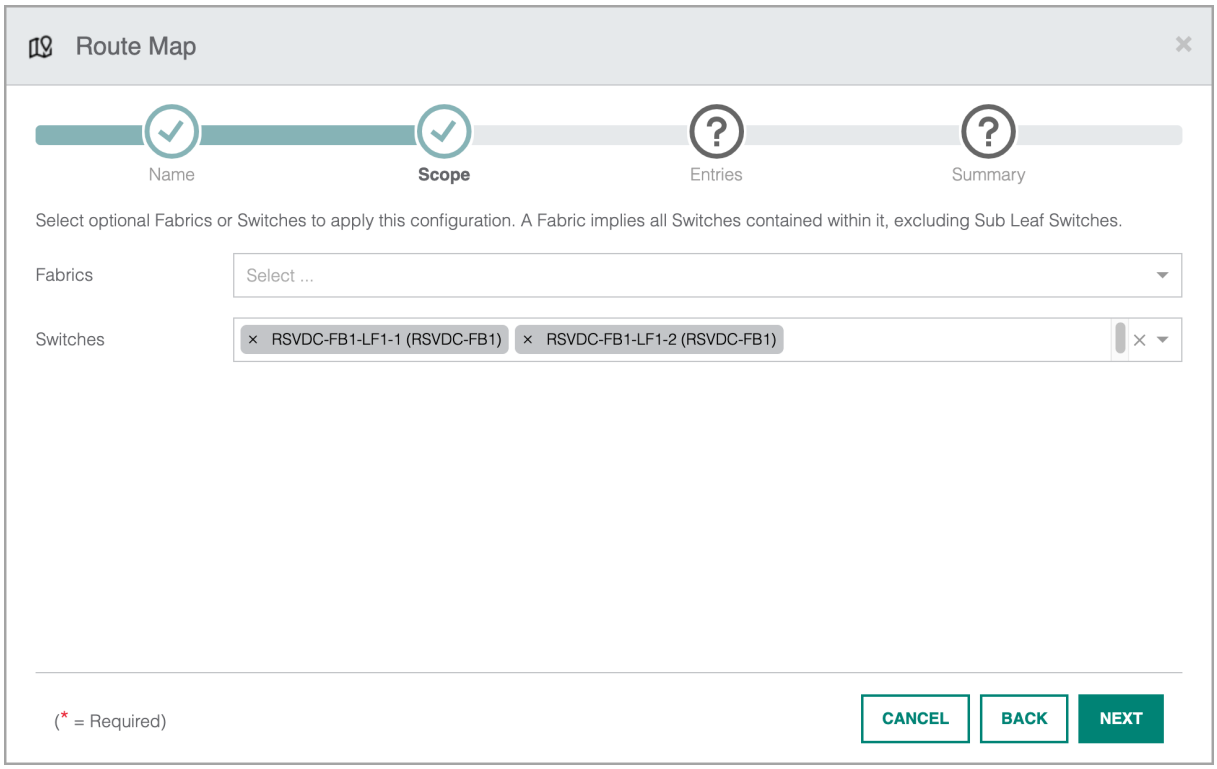


**Step 4** On the **ACTIONS** menu, select **Add**.

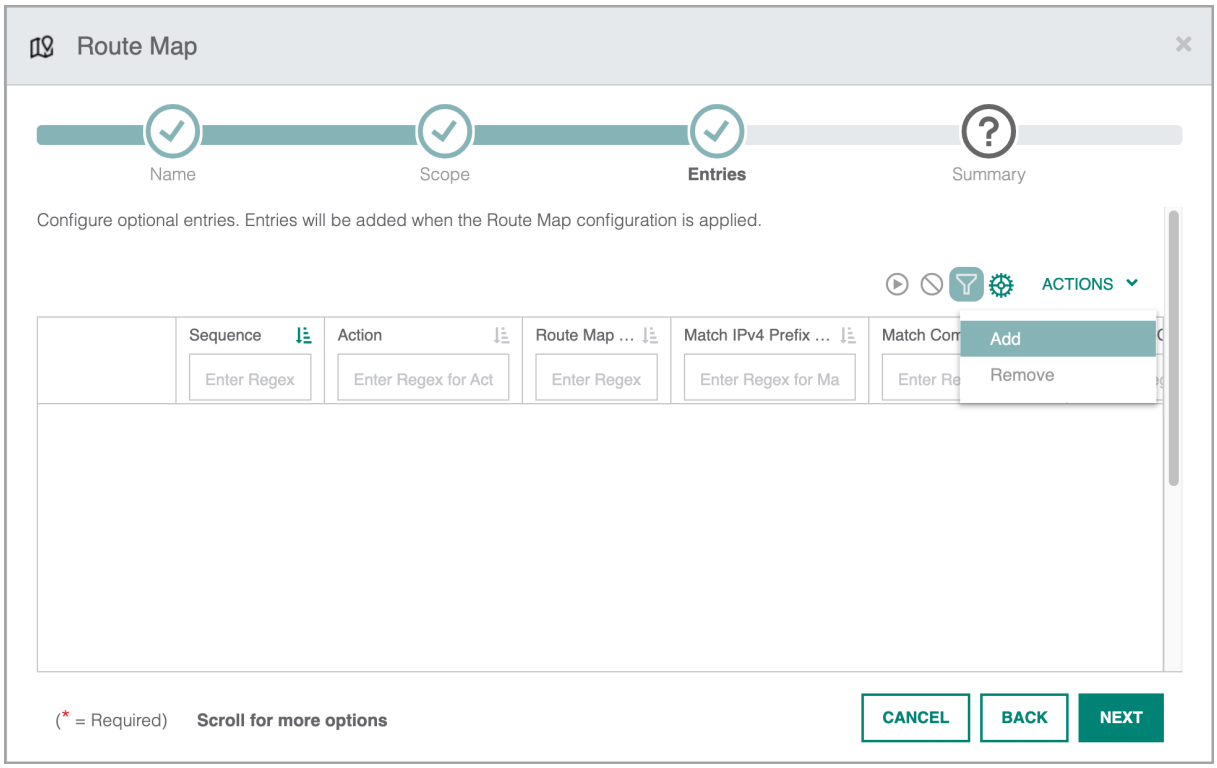


**Step 5** On the **Settings** page, enter the following values and click **NEXT**

- **Sequence**: *10*
- **Description**: *permit local overlay advertisements*
- **Action**: *Permit*

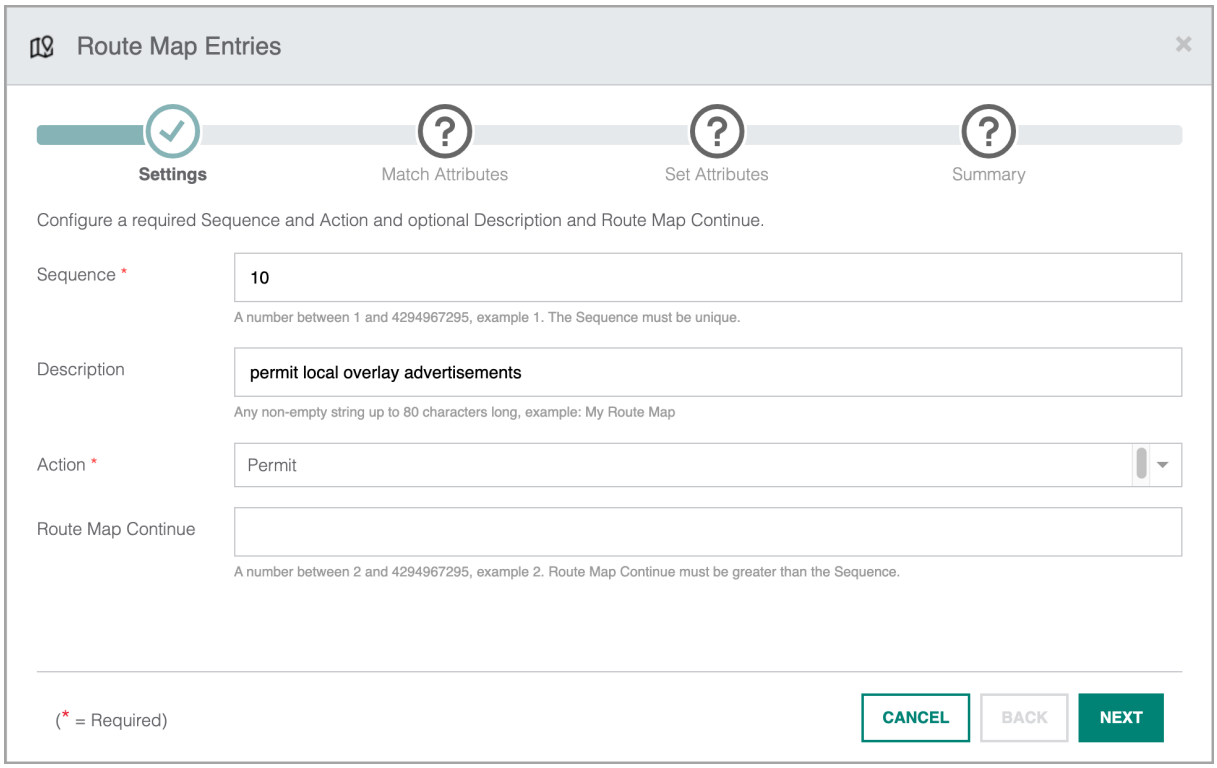


**Step 6** On the **Match Attributes** page, enter the following values and click **NEXT**.

- **Attributes**: *Match AS Path List*
- **Match AS Path List**: *local-fabric*

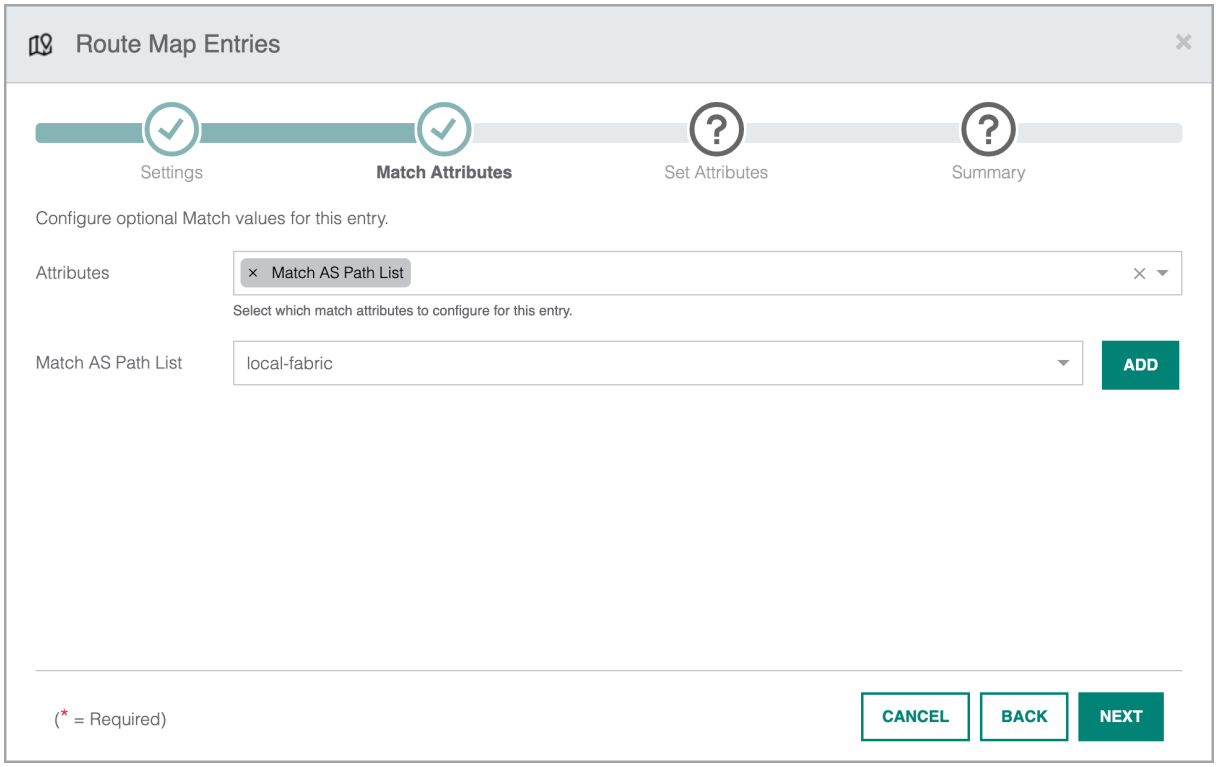


**Step 7** On the **Set Attributes** page, click **NEXT**.

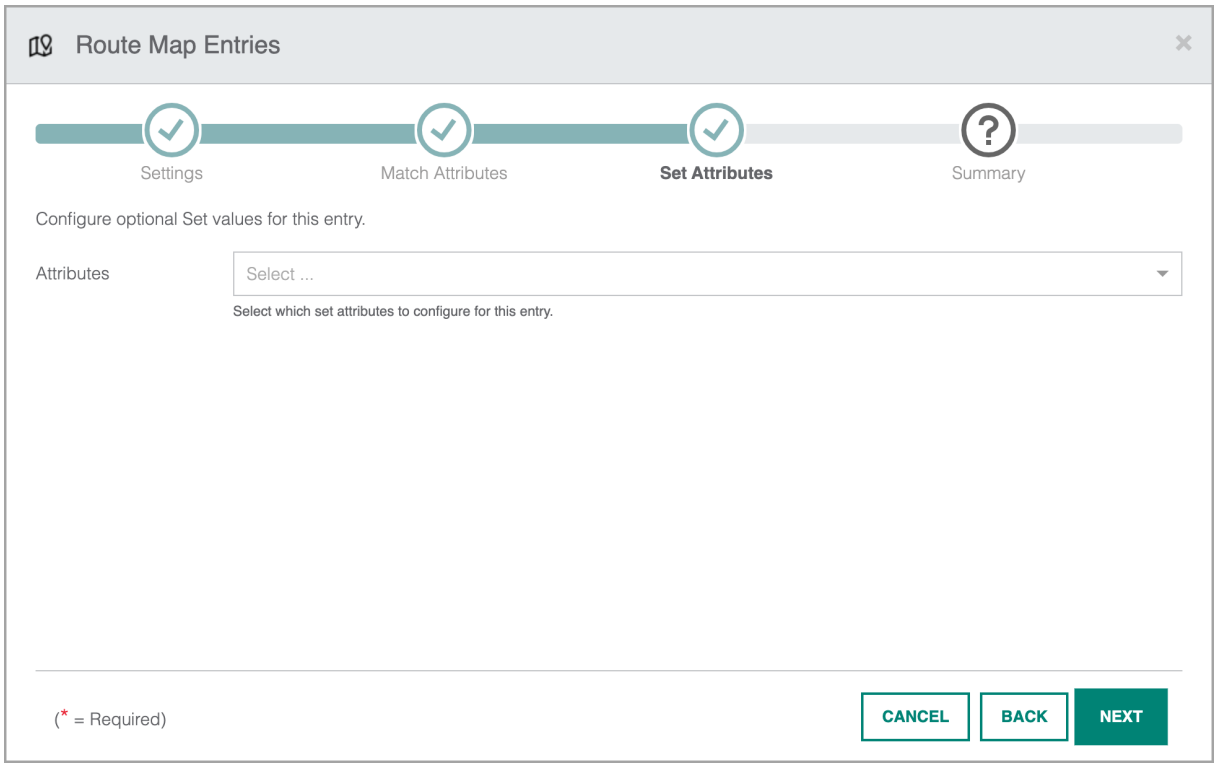


**Step 8** On the **Summary** page, verify the multifabric BGP settings and click **APPLY**.

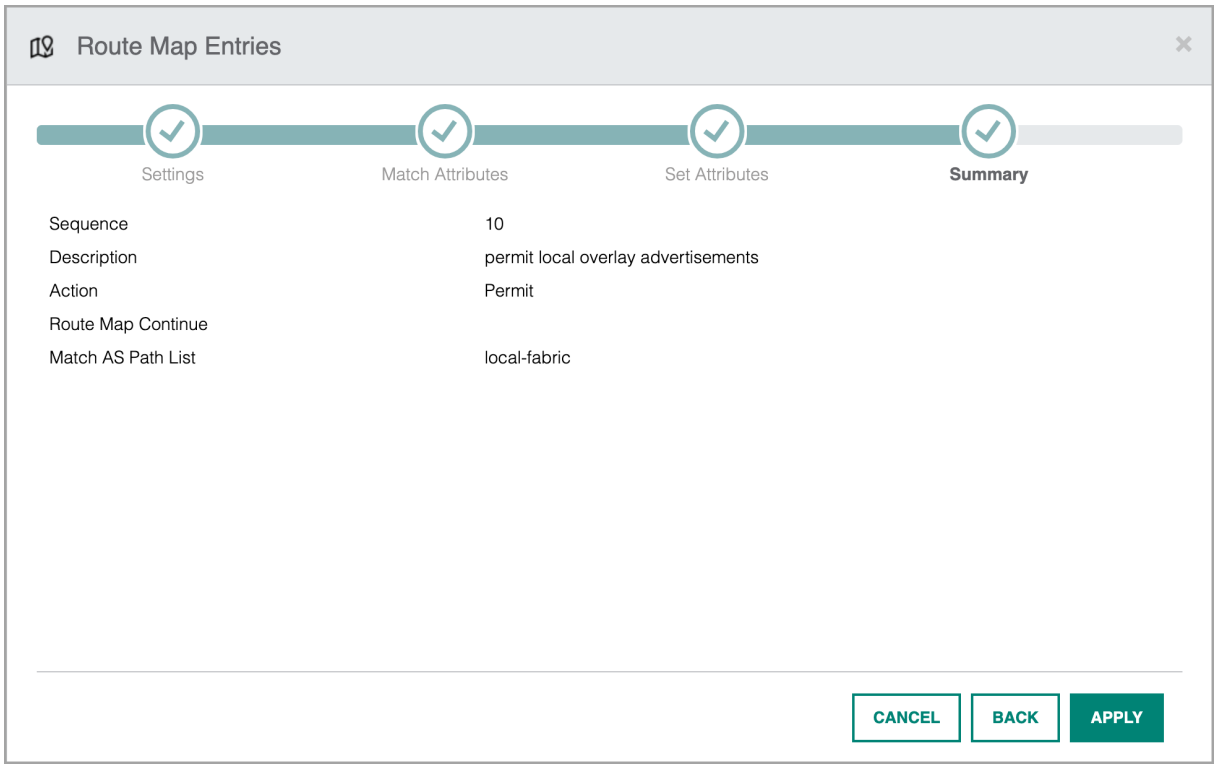


**Step 9** Repeat steps 11 to 15 to add a second route map entry with the values below.

- **Sequence**: *20*
- **Description**: *permit campus/firewall prefixes*
- **Action**: *Permit*
- **Attributes**: *Match AS Path List*

- **Match AS Path List**: *ALLOWED-EXT-AS*

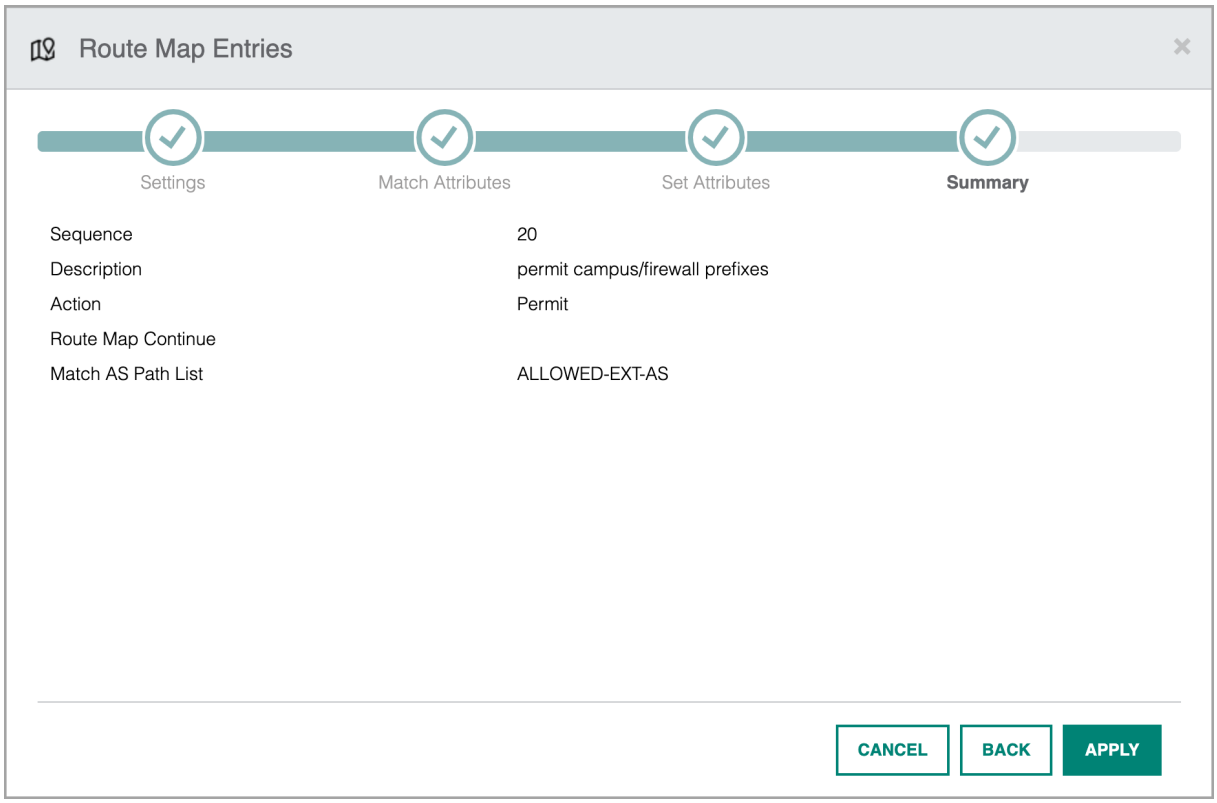


**Step 10** Click **NEXT**.

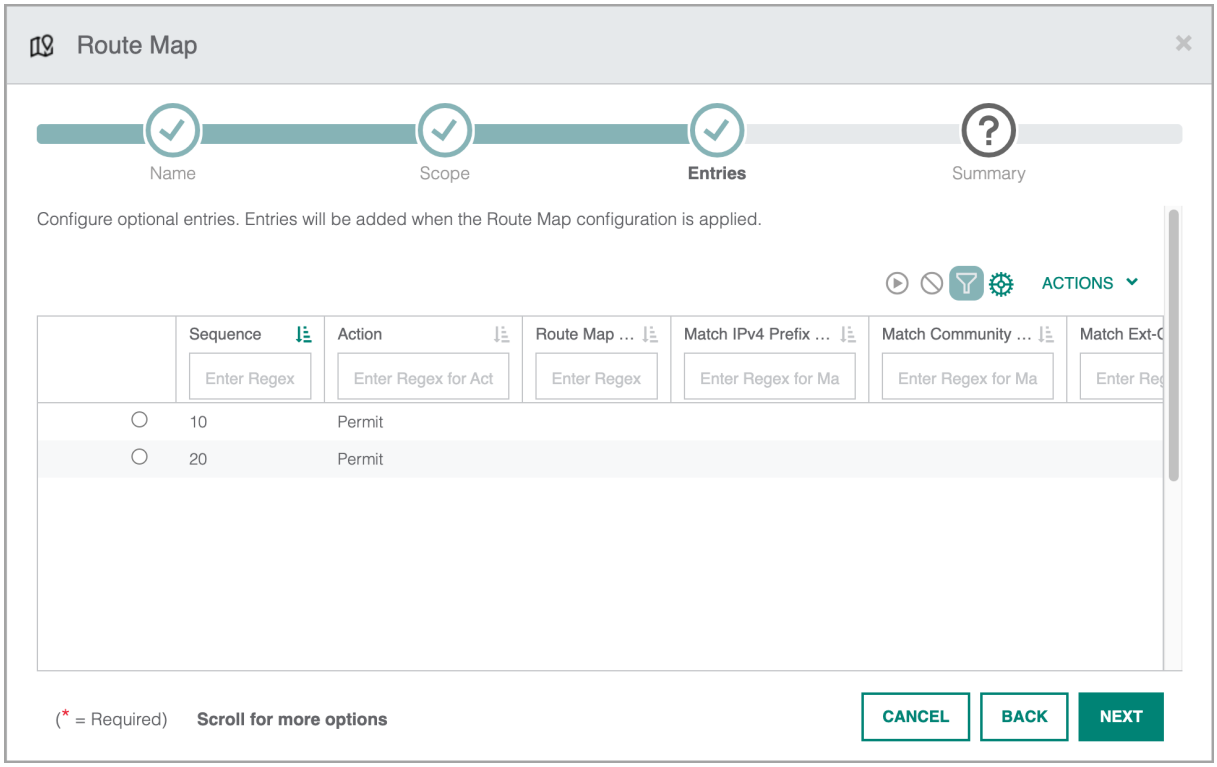


**Step 11** On the **Summary** page, verify the route map settings and click **APPLY**.

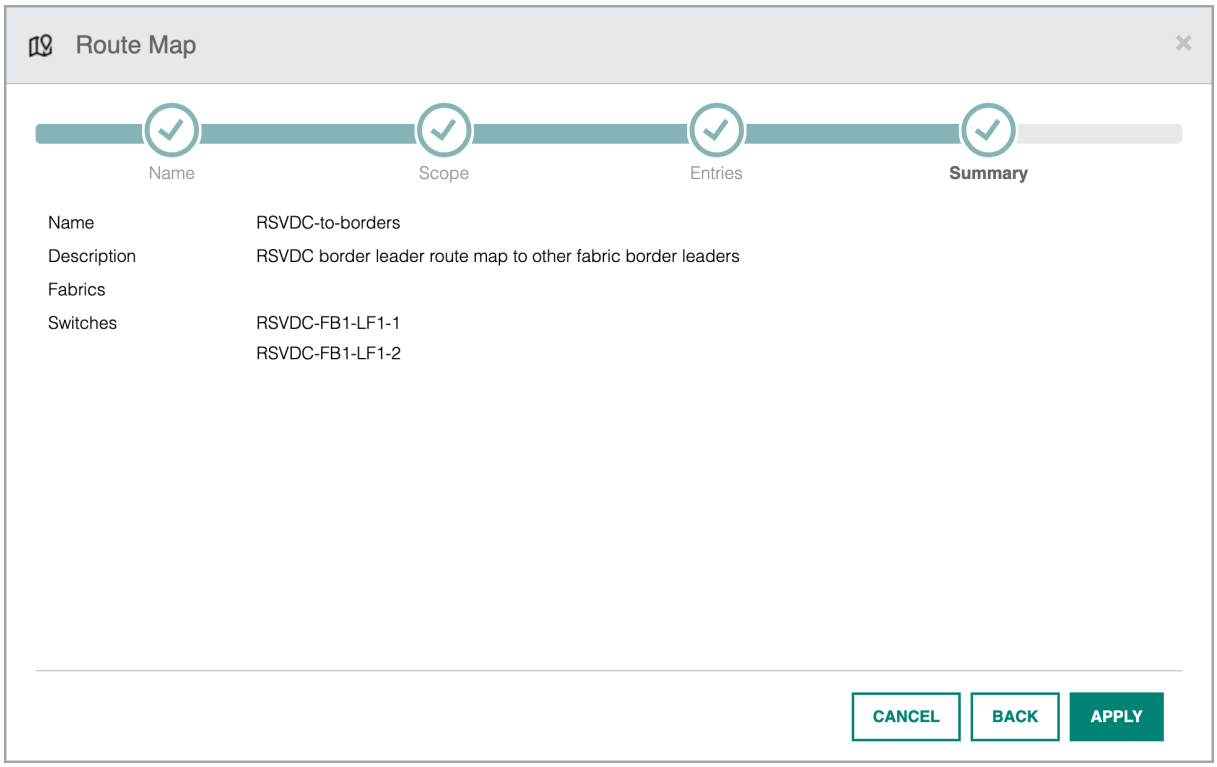


### *Apply Route Map to EVPN Peering*

**Step 1** On the menu bar at the top right of the AFC display, click the **CLI Commands** icon and select **Show Commands**.

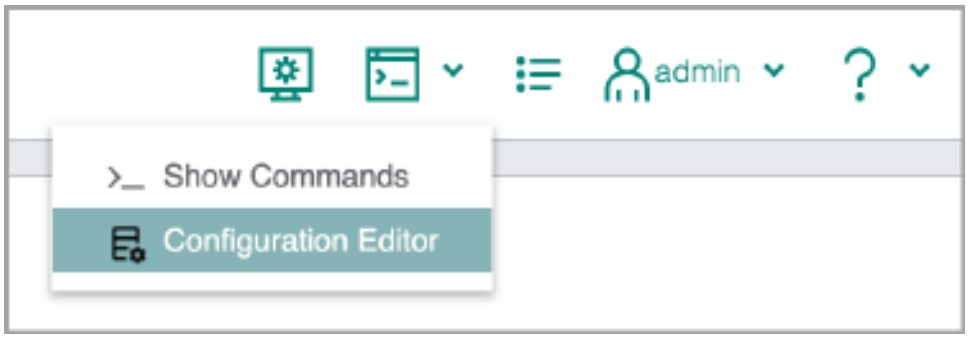


**Step 2** Select **RSVDC-FB1** in the **Fabric** field and both border leaf switches in the **Switch** field.

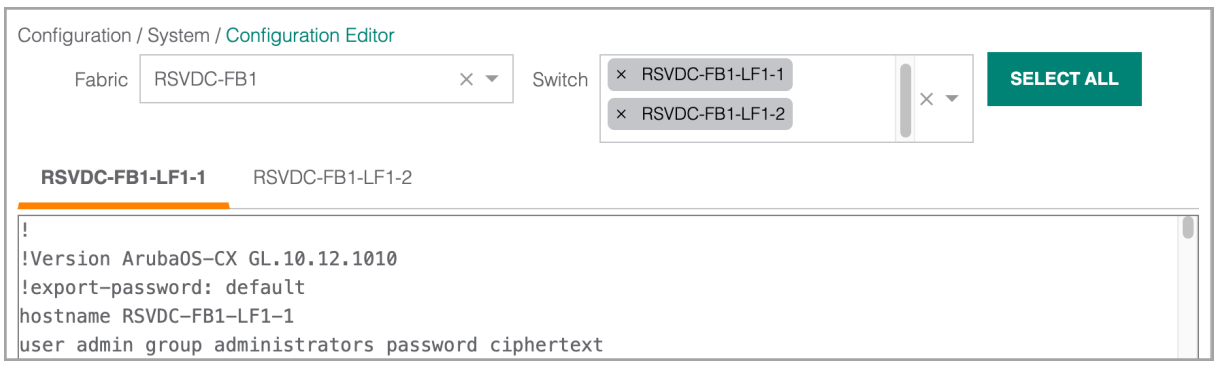

**Step 3** In the switch configuration window, scroll to the **address-family l2vpn evpn** stanza in the BGP configuration section. Set the **border-leaders** peer group's outbound route map to **RSVDC-to-borders**.

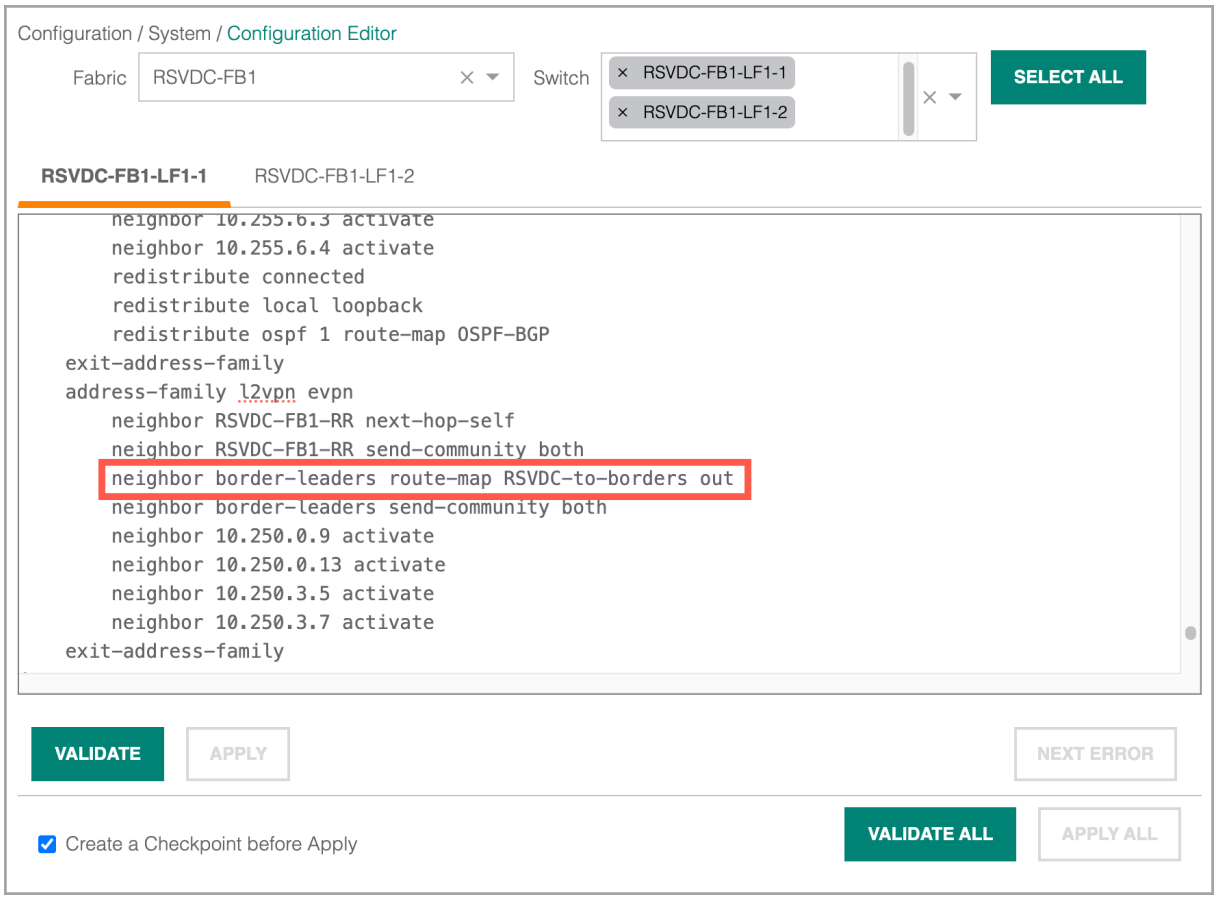


**Step 4** Click the **RSVDC-FB1-LF1-2** tab. Set the **border-leaders** peer group's outbound route map to **RSVDC-to-borders**. Click **VALIDATE ALL**.

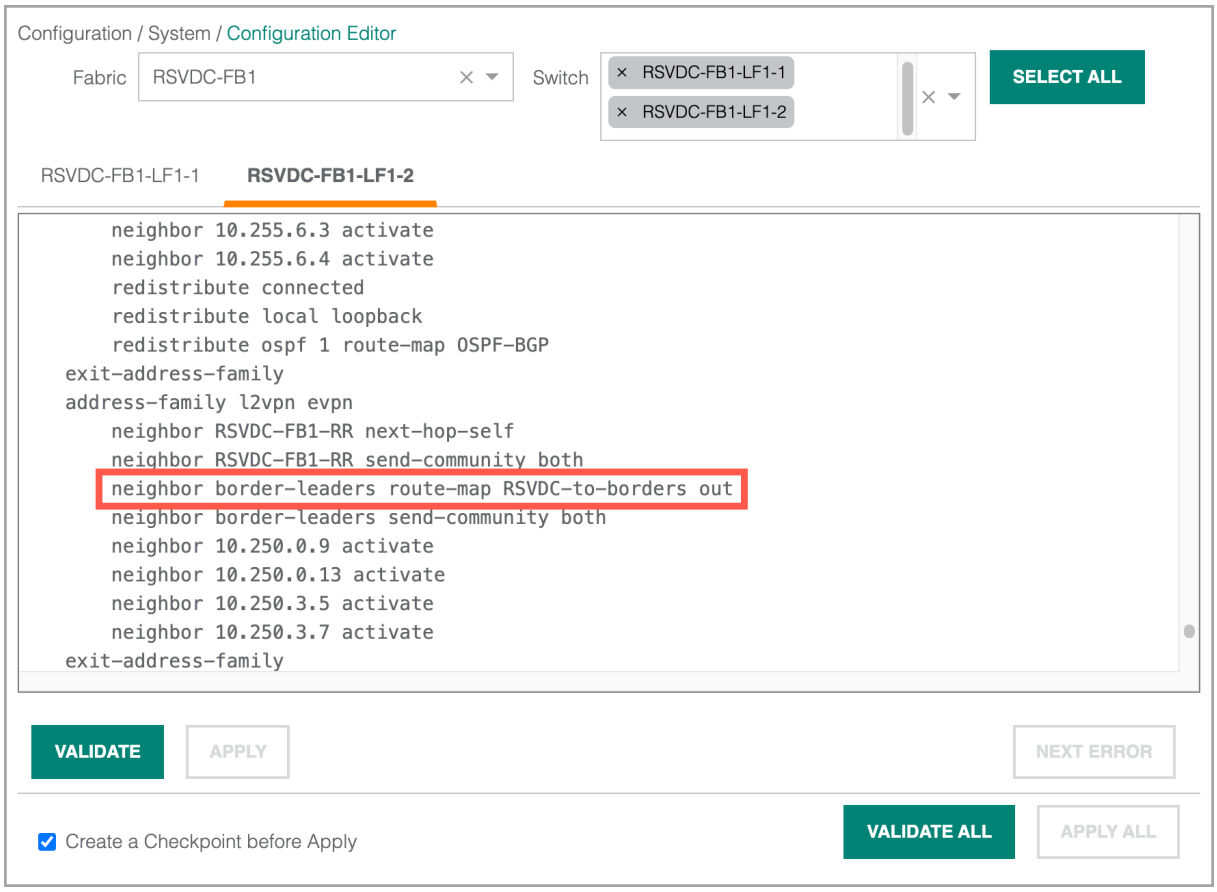


**Step 5** A success message verifies the configuration is valid. Click **APPLY ALL**.


**NOTE:**

If configuration errors are present on a switch, a red error icon appears on the switch tab and the configuration errors are highlighted in red. Correct the errors and validate the configuration again. It is possible an incomplete spanning-tree configuration was unintentionally created in the AFC guided setup process. If **spanning-tree config-name** is in the config with no name parameter and no other spanning-tree configuration present, it is safe to delete the line.

**Step 6** Success messages verify a configuration checkpoint was created and the configuration changes were applied.

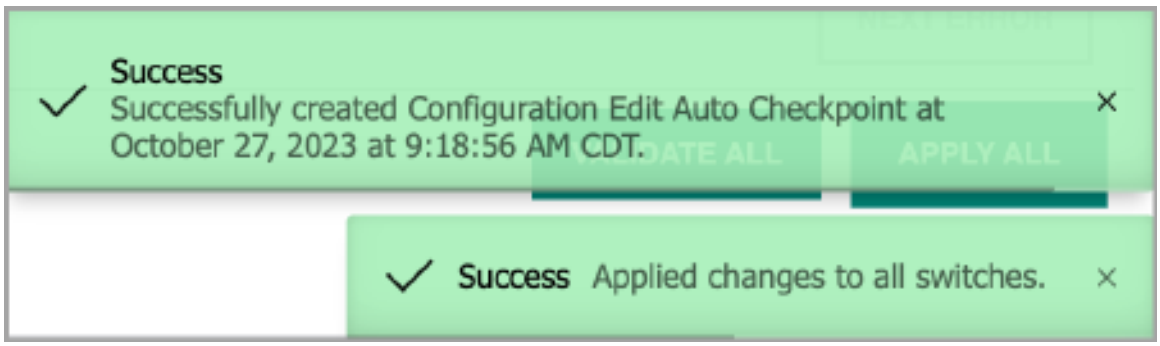

**Step 7** On the RSVDC-FB1-LF1-1 switch CLI, clear the EVPN BGP sessions to the second fabric to apply the new route map policy.

*clear bgp 10.250.3.5 clear bgp 10.250.7*

```
RSVDC-FB1-LF1-1# show bgp l2vpn evpn summary
Codes: * Dynamic Neighbor
VRF : default
BGP Summary
-----------
 Local AS               : 65001        BGP Router Identifier  : 10.250.0.11
 Peers                  : 4            Log Neighbor Changes   : Yes
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60
 Confederation Id       : 0

 Neighbor                                Remote-AS MsgRcvd MsgSent   Up/Down Time State        AdminStatus
 10.250.0.9                              65001       14663   6782    01d:19h:09m  Established   Up
 10.250.0.13                             65001       14684   6792    01d:19h:09m  Established   Up
 10.250.3.5                              65002       543     876     03h:44m:35s  Established   Up
 10.250.3.7                              65002       541     873     03h:44m:33s  Established   Up

RSVDC-FB1-LF1-1# clear bgp 10.250.3.5
RSVDC-FB1-LF1-1# clear bgp 10.250.3.7
RSVDC-FB1-LF1-1#
```

**Step 8** Repeat step 7 on the RSVDC-FB1-LF1-2 switch.

**Step 9** On the menu bar at the top right of the AFC display, click the **CLI Commands** icon and select **Show Commands**.

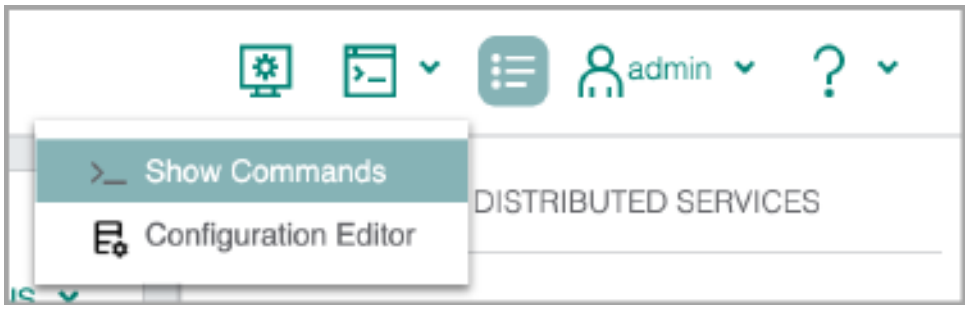


**Step 10** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Switches:** *RSVCO-FB2-LF1-1*, *RSVCO-FB2-LF1-1*, *RSVCO-FB2-LF2*, *RSVCO-FB2-LF3*
- **Commands:** *show ip route vrf PROD-DC-VRF*

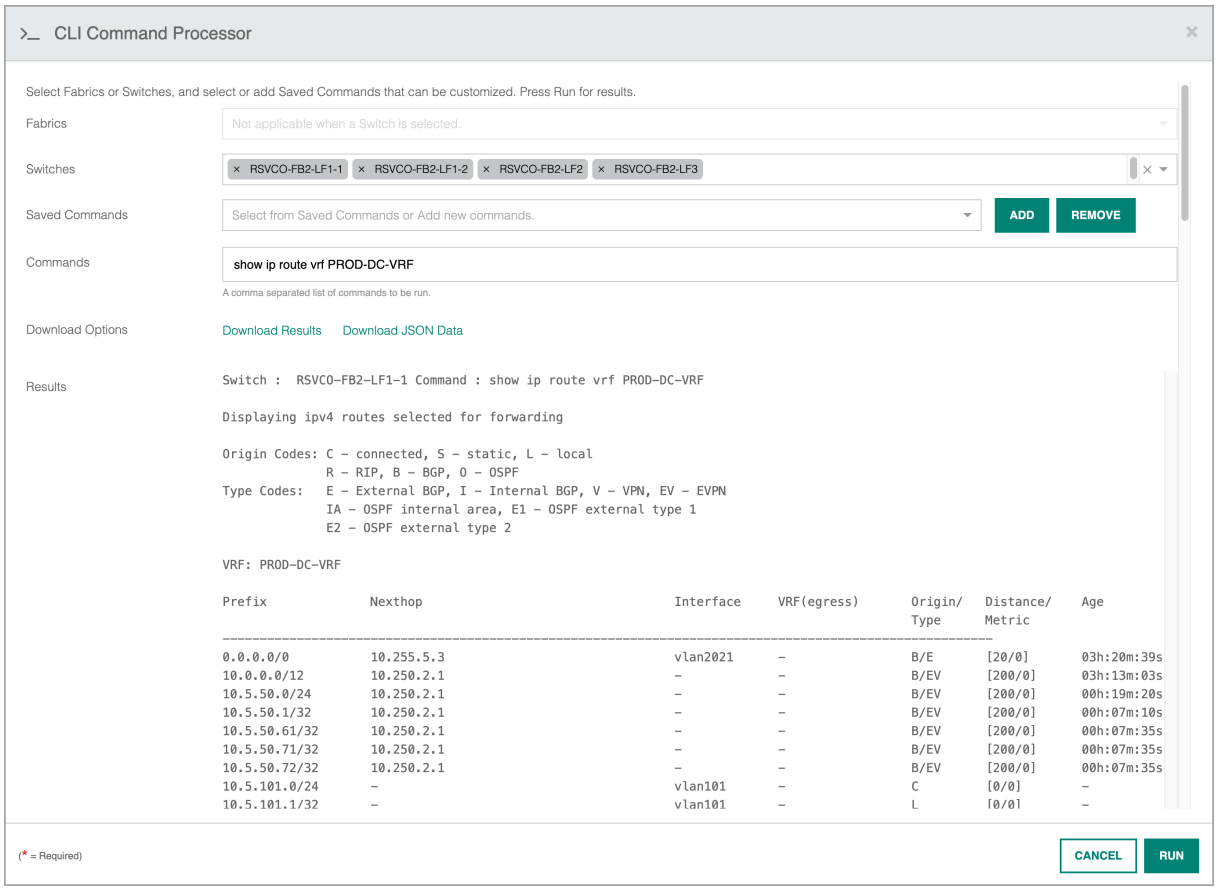


**Step 11** Verify that each leaf switch in the RSVCO fabric learns the 10.0.0.0/12 campus summary route and the 10.5.50.0/24 prefixes that are present only on CX 10000 switches in the RSVDC-FB1 fabric.

## Extend Layer 2 Reachability Across Fabrics

When using EVPN, host reachability advertisements include a route-target to inform remote VTEPs of the VLAN associated with the host MAC address advertisement. VLAN route-targets can be generated automatically within a fabric using an iBGP EVPN peering. The eBGP EVPN peering between fabrics requires an explicitly defined route target for each VLAN. This additional route target controls when MAC advertisements shared between fabrics are installed in local MAC address tables.

**Step 1** On the **Configuration** menu, select **Routing > EVPN**, then click the **EVPN MULTI SITE** tab.

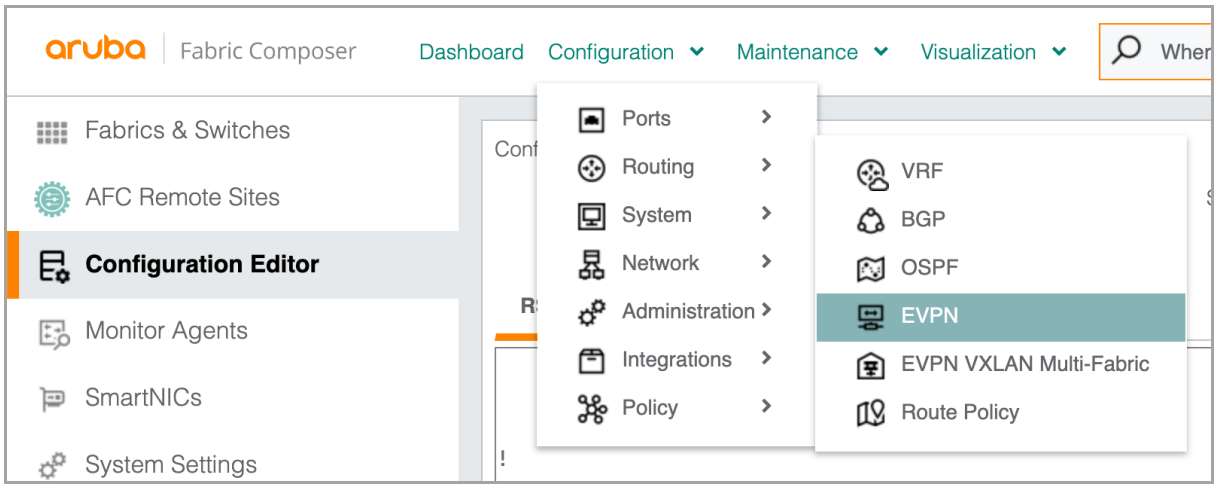


**NOTE:**

If the left pane no longer displays routing options, select **Configuration > Routing > EVPN** on the top menu.

**Step 2** With **RSVDC-FB1** selected in the **Fabric** menu, click the **ACTIONS** menu and select **Add**.

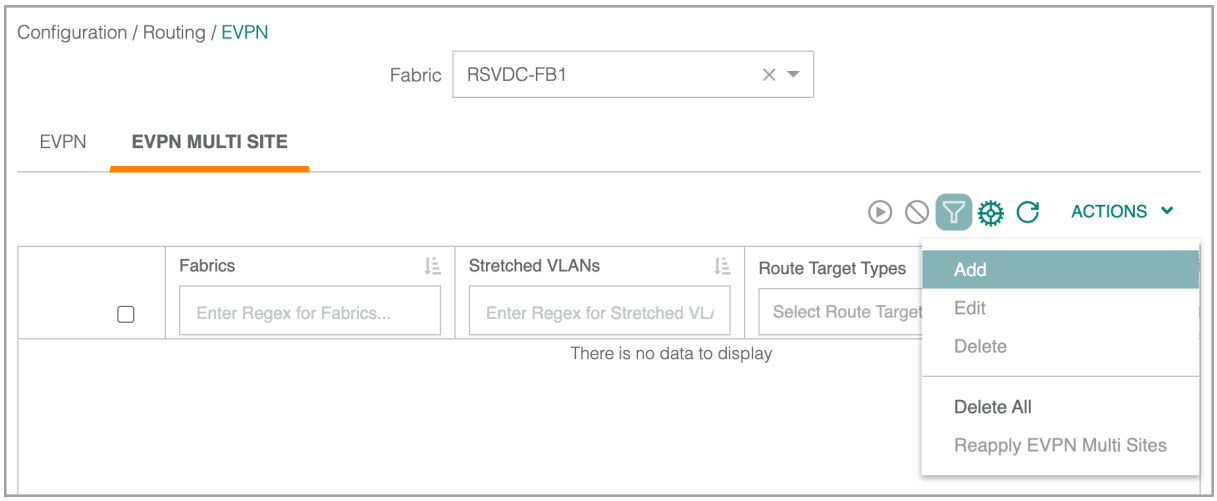


**Step 3** On the **Fabrics** page, select both fabrics and click **NEXT**.

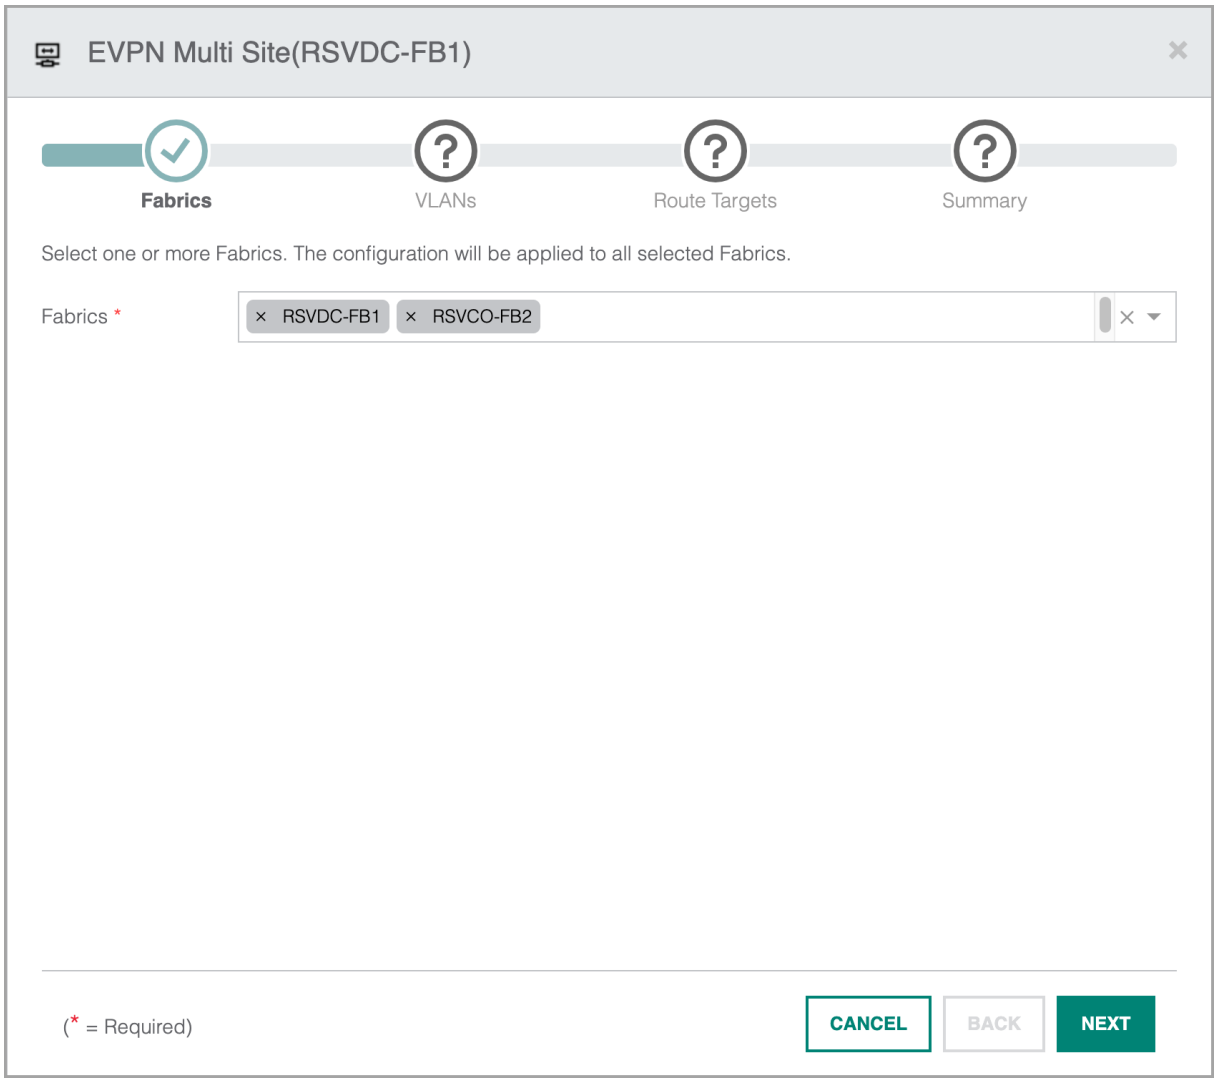


**Step 4** On the **VLANs** page, enter the VLAN IDs that share Layer 2 reachability information across both fabrics.

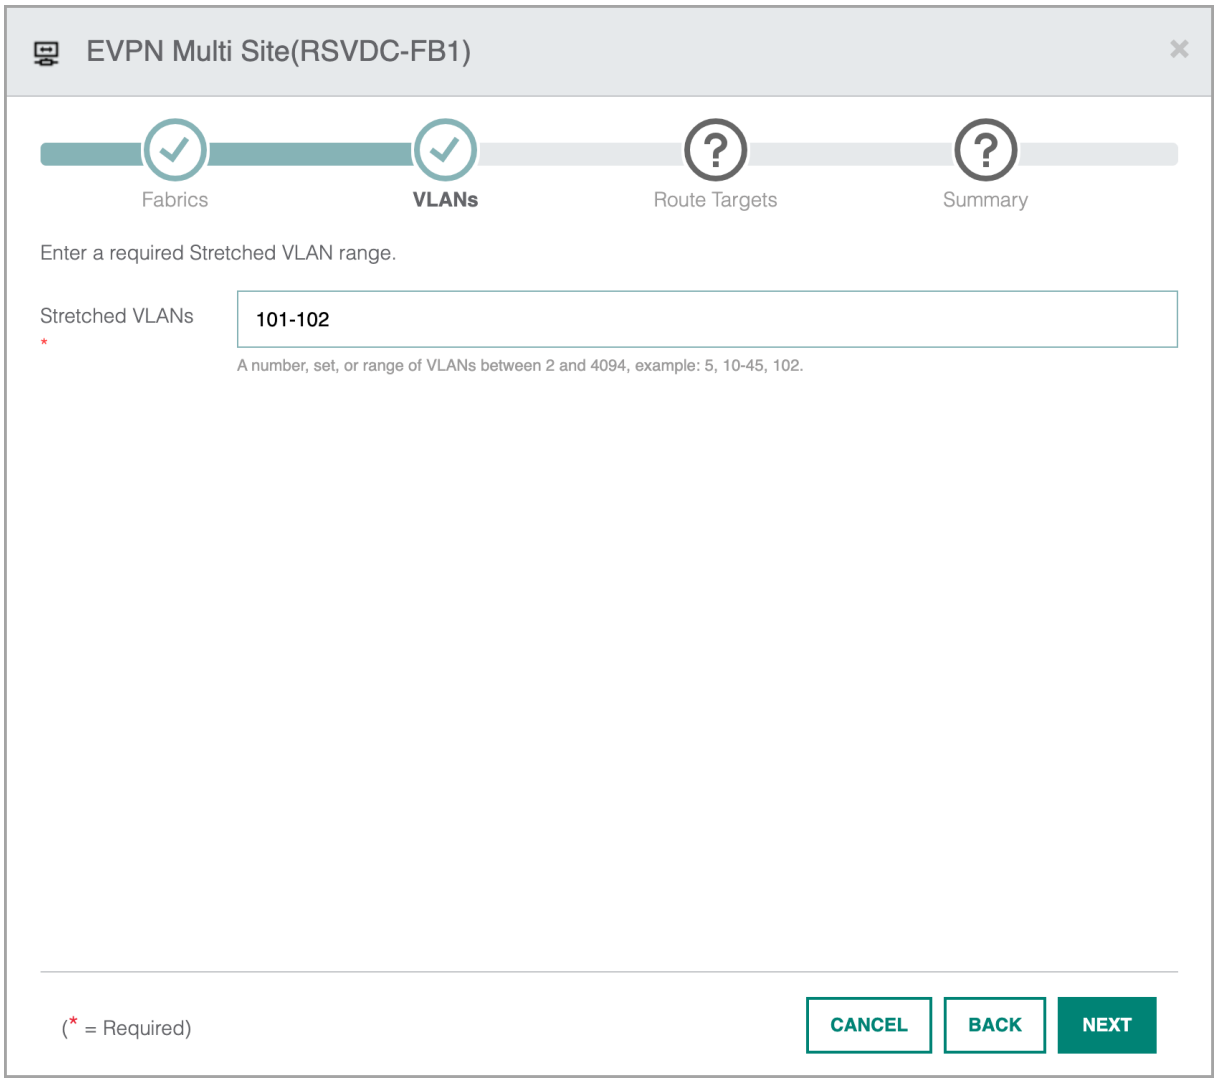


**Step 5** On the **Route Targets** page, enter the following values and click **ADD**.

- **Route Target Type**: *NN:VNI*
- **Administrative Number**: *1*

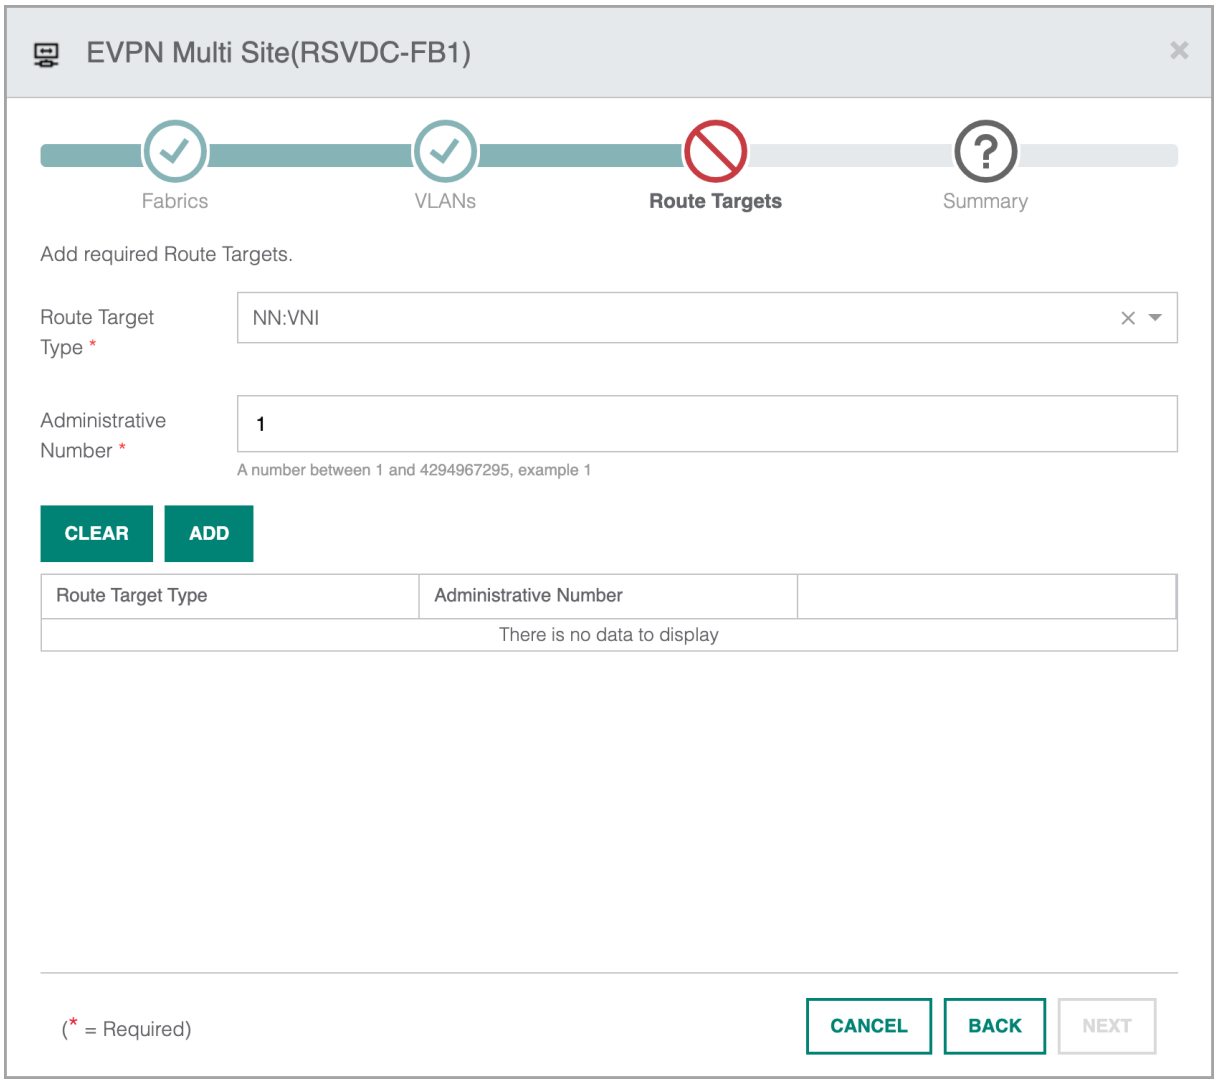


**Step 6** Click **NEXT**.

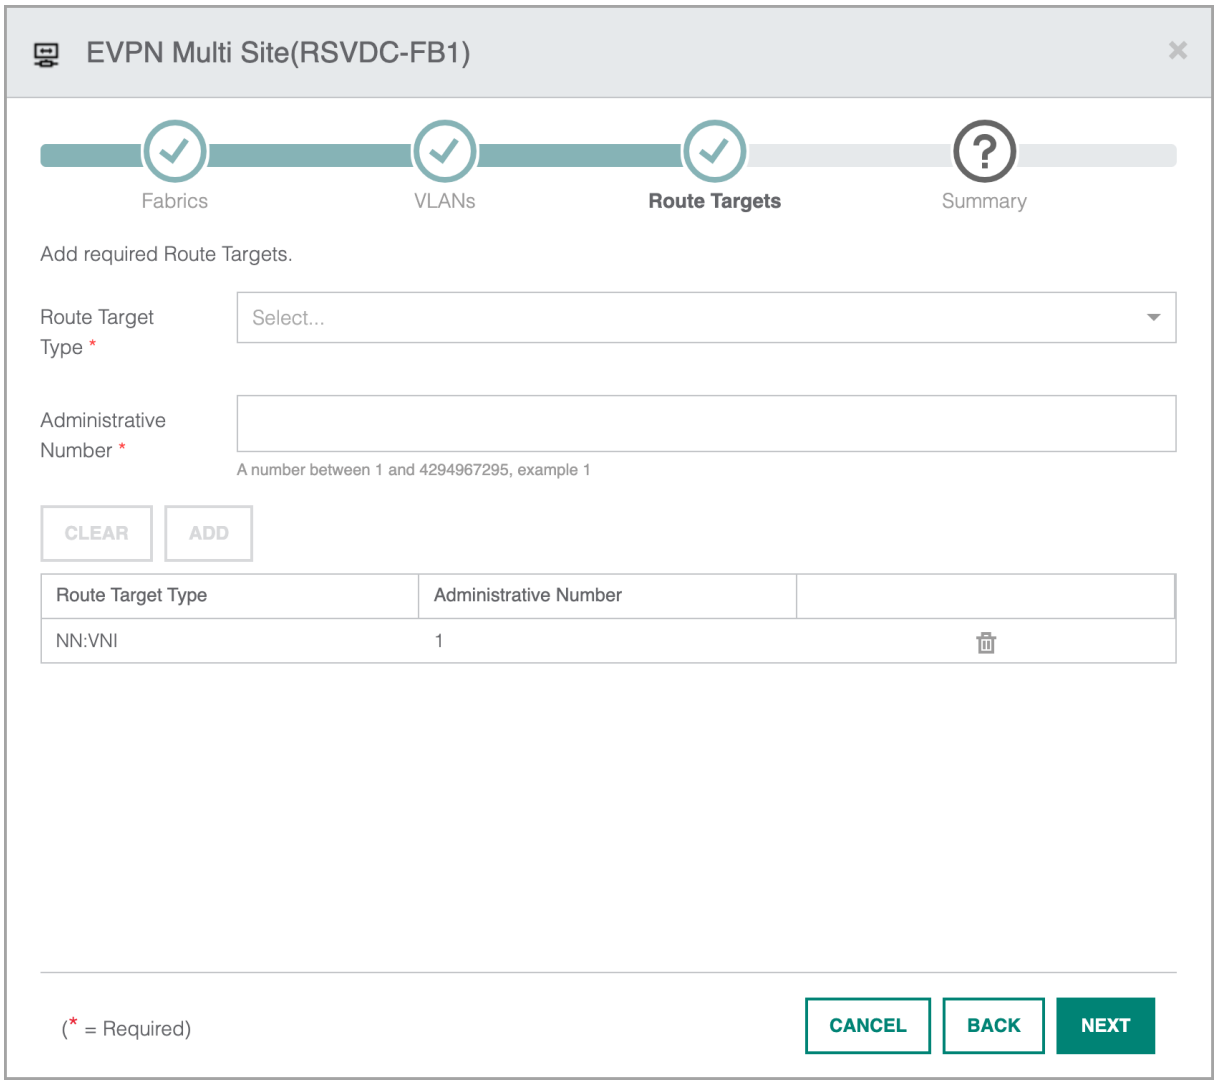


**Step 7** On the **Summary** page, verify the route target settings and click **APPLY**.

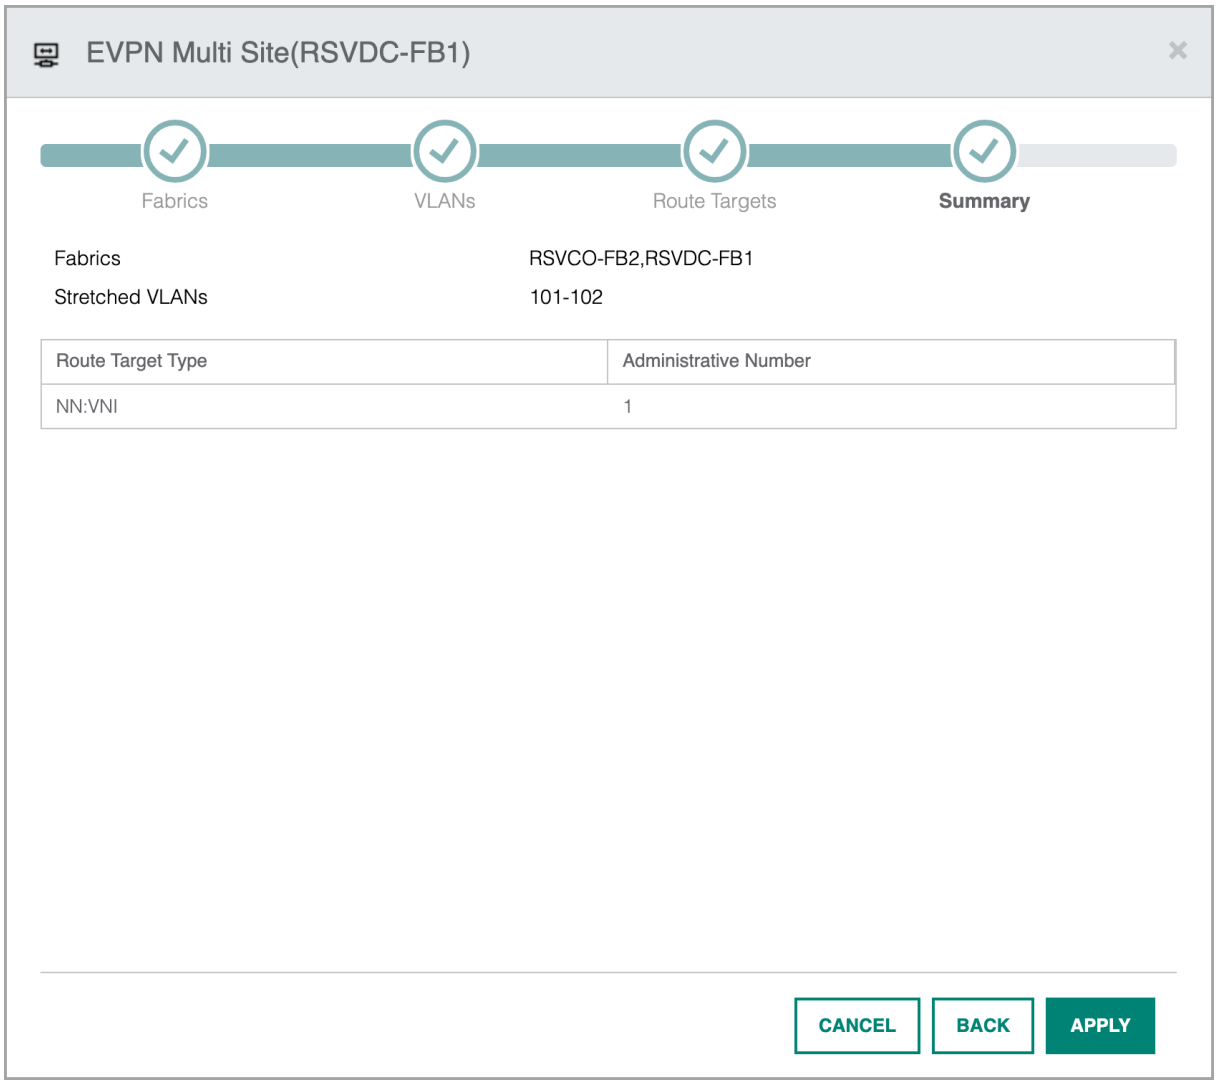


**NOTE:**

If the fabrics are managed by different AFC instances, the **EVPN** wizard must be run for each instance.

**Step 8** On the menu bar at the top right of the AFC display, click the **CLI Commands** icon and select **Show Commands**.

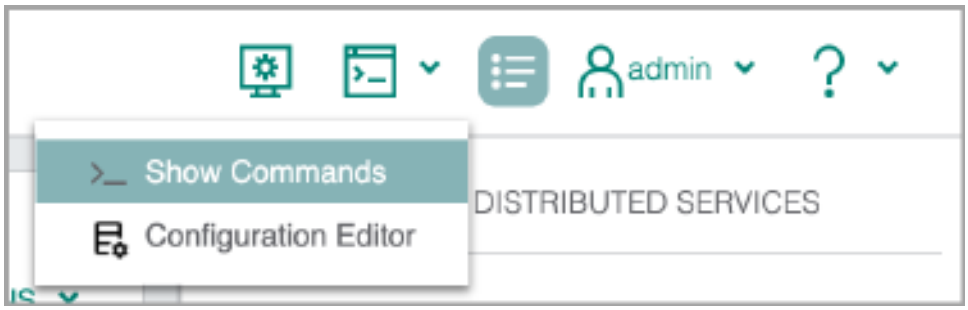


**Step 9** On the **CLI Command Processor** page, enter the following values, then click **RUN**.

- **Fabrics:** *RSVDC-FB1*, *RSVCO-FB2*
- **Commands:** *show mac-address-table*

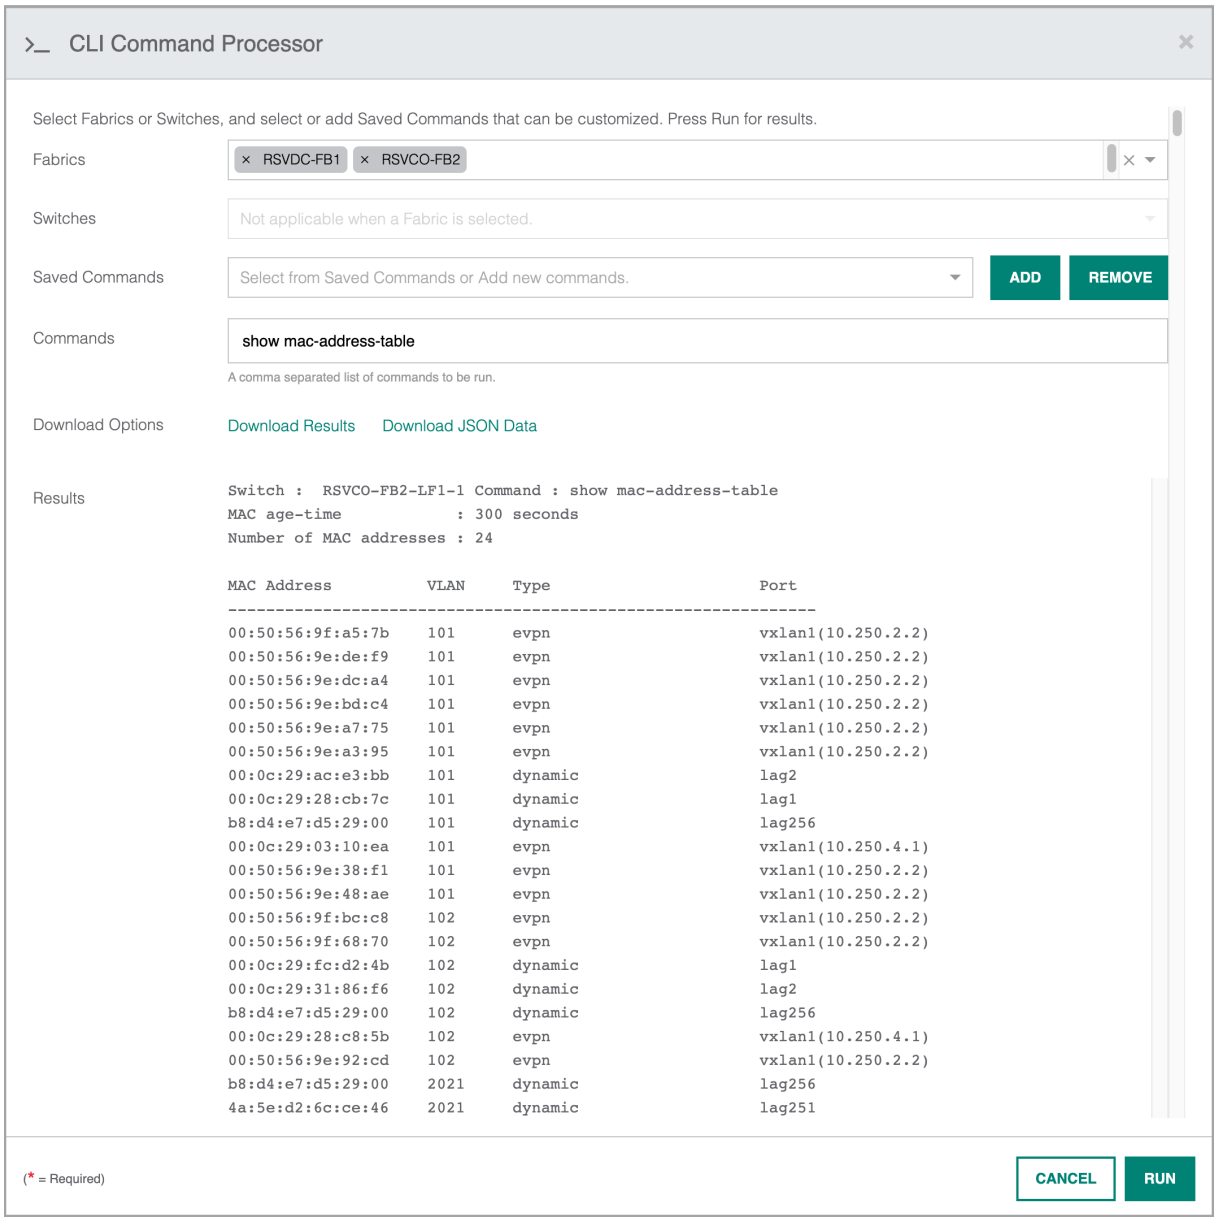


**NOTE:**

VXLAN-learned MAC address entries include the VTEP IP in parentheses in the **Port** column.

**Step 10** Verify that each switch in both fabrics learns MAC addresses from the other fabric (the MAC address entires are displayed with a VTEP IP of the border leader in the other fabric).

At the completion of this procedure, a multifabric EVPN is established between the RSVDC-FB1 and RSVCO-FB2 fabrics.

# Aruba Central Two-Tier Data Center

The Aruba ESP Two-Tier Data Center can be configured using Aruba Central or Aruba Fabric Composer (AFC). The Two-Tier architecture uses Layer 2 multi-chassis links between a VSX pair of core switches and a set of server access switches. The following deployment example uses Aruba Central, enabling a single management platform for both campus and data center networks that provides advanced troubleshooting features and performance feedback.

## Overview

The Aruba CX switching portfolio includes a range of products for data center core and access layers. Aruba Central documentation contains a list of supported AOS-CX switches.

Aruba ESP Two-Tier data centers meet the requirements for small- and medium-size data centers. It provides network resiliency by using multi-chassis link aggregations (MC-LAGs) at both switch tiers.

**Two-Tier Data Center Topology**

The diagram below summarizes the physical topology configured in this deployment guide and the relationship between components.

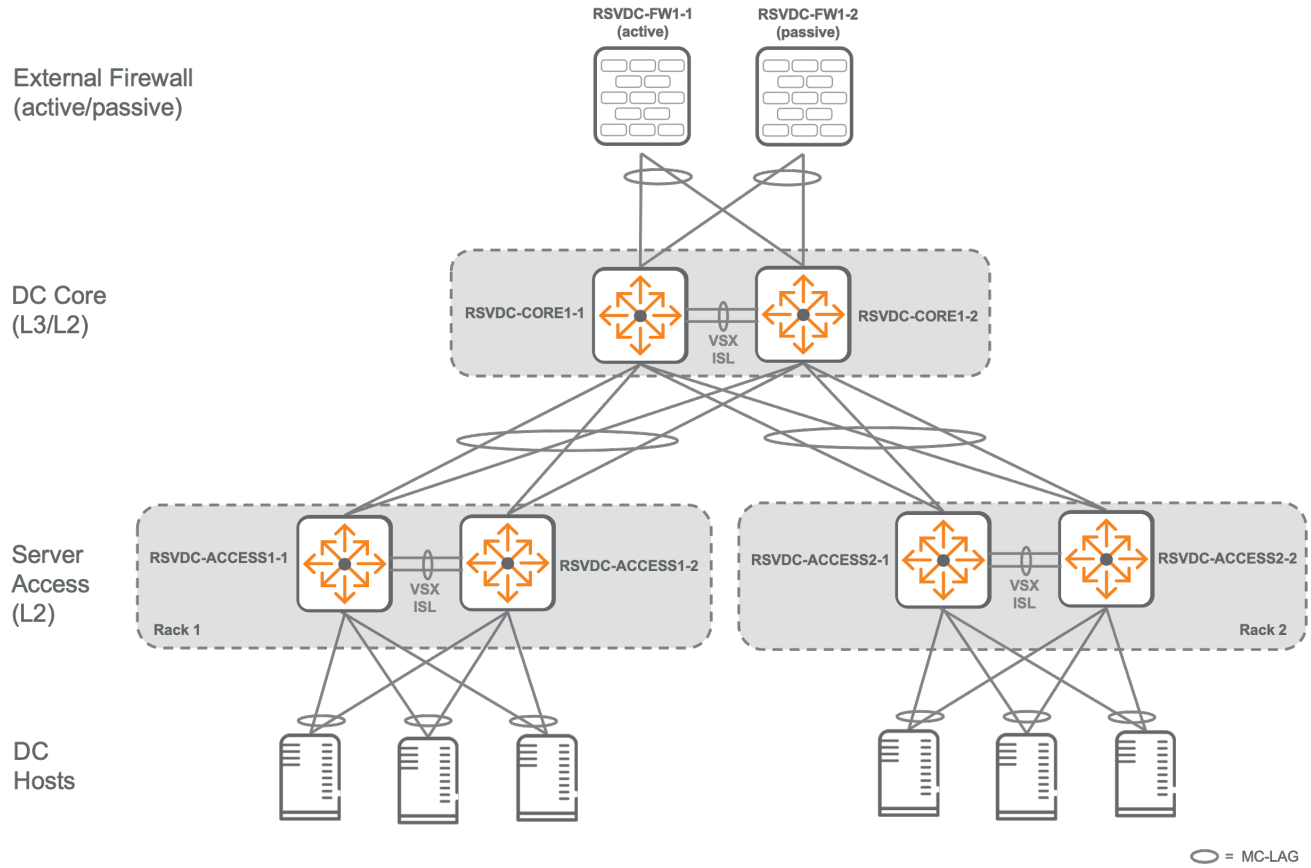

### Two-Tier Core Layer

The core layer provides redundant Layer 2 connectivity to downstream access switches. A VSX pair of core switches is configured with an MC-LAG to each downstream rack. All links from the core layer to the access layer for a single rack are members of the same MC-LAG, whether the rack is populated with a single switch or with a VSX-pair of access switches. MC-LAG provides network resiliency and load-balancing. It also mitigates the need for loop avoidance mechanisms between the core and access layer switches.

Layer 3 services for the data center are provided by the core layer. VLAN switched virtual interfaces (SVIs) are defined on core switches that route packets between data center subnets and provide redundant IP gateways to data center hosts. The core layer also provides redundant IP connectivity to upstream external networks. Typically, firewalls are placed between a data center and external networks for policy enforcement. The redundancy models between the data center core and external networks can vary, depending on device feature sets and organizational requirements. In this guide, a traditional active/passive redundant pair of firewalls is connected to the core switch pair using MC-LAGs.

### Two-Tier Access Layer

The access layer provides Layer 2 connectivity to downstream data center hosts.

When a single access switch is at the top-of-rack (ToR) position, the access layer connects to the core layer using a standard LAG. A single ToR switch can provide physical link redundancy using a standard LAG, but host connectivity is lost when performing firmware upgrades or when the ToR switch fails.

When using a VSX pair of ToR switches, the access layer provides physical switch redundancy to directly attached hosts. This model supports uninterrupted host connectivity, even when one of the ToR switches fails or a firmware upgrade is performed. Each access layer switch also is connected to each core switch. All core links across redundant access switches are members of the same MC-LAG for redundancy and loop avoidance.

## Planning the Deployment

This section provides sample values and rationale for naming and numbering schemes. Adjust values and formats as needed to accommodate specific requirements. Using a consistent approach in the physical and logical configurations improves the management and troubleshooting characteristics of a network.

### Naming Conventions

Establish a switch naming convention that indicates the switch type, role, and location to simplify identification and increase operating efficiency.

Example values used in this guide:

| Switch Name | Role | Description |
|---|---|---|
| RSVDC-CORE1-1 | Core | Roseville Data Center Core Switch, VSX Pair Member 1 (primary) |
| RSVDC-CORE1-2 | Core | Roseville Data Center Core Switch, VSX Pair Member 2 (secondary) |
| RSVDC-ACCESS1-1 | Access | Top-of-Rack Access Switch in Rack 1, VSX Pair Member #1 (primary) |
| RSVDC-ACCESS1-2 | Access | Top-of-Rack Access Switch in Rack #1, VSX Pair Member 2 (secondary) |

## Aruba Central Groups

Aruba Central organizes devices in groups with common configuration elements. Two functional roles in the two-tier data center architecture share configuration elements: the data center core and access layers. An Aruba Central group should be created for each layer.

Example Aruba Central groups used in this guide are:

- **DC-RSVCORE**
- **DC-RSVACESS**

## Aruba Central Sites

In addition to group membership, a device can be associated with a site that represents a physical location. Sites can be used to aggregate visibility, statistics, and troubleshooting tools across switches that are members of different groups.

In this guide, all data center switches are assigned to a site named **RSVDC**.

## IP Address Planning

Plan a consistent IP numbering scheme with values that can accommodate the current deployment size and leave room for growth. Define a range that can represent loopback addresses, IP addresses used in support protocols, and a range for data center hosts. It is beneficial to assign data center host subnets from a larger range of maskable IP addresses that summarizes all host subnets in the data center.

Example IP address ranges used in this guide:

| Subnet | Functional Description |
|---|---|
| 10.255.12.0/24 | Routed interface IP addresses |

| Subnet | Functional Description |
|---|---|
| 10.250.12.0/24 | Loopback IP addresses |
| 10.12.0.0/16 | Summary range of all data center host subnets |
| 10.12.101.0/24 | Example of a specific data center host subnet |

## MAC Address Planning

A Locally Administered Address (LAA) should be used when defining virtual MAC addresses for VSX and active gateway functions. This is required when configuring an Active Gateway for an SVI on a VSX pair and when configuring the system MAC address of VSX. An LAA is a MAC in one of the four formats shown below:

```
x2-xx-xx-xx-xx-xx
x6-xx-xx-xx-xx-xx
xA-xx-xx-xx-xx-xx
xE-xx-xx-xx-xx-xx
```

The *x* positions can contain any valid hex value. It is helpful to create a hexadecimal representation of the associated IP address or VLAN ID using the hex positions. For more details on the LAA format, see the IEEE tutorial guide.

In this guide, VSX system MAC addresses are set to 02:00:00:00:10:xx, where xx is replaced with the rack number of the VSX pair and the core switches use a value of 00.

Active Gateway MAC address are set to 02:00:xx:xx:xx:xx, where the last four octets are assigned a hexadecimal representation of the Active Gateway IP address. For example, the IP address 10.1.101.1 results in a MAC address of 02:00:0a:01:65:01. This simple method ensures MAC address uniqueness associated with Active Gateway IP addresses for troubleshooting purposes.

# Prepare Switches

After initial unboxing and inventory, the next step in deploying a data center network requires physical installation of network switches.

## Switch Installation

Verify the airflow configuration for the products to be installed to ensure that they support the cooling design for the data center. If required, an optional air duct kit is available for Aruba data center top-of-rack (ToR) switches to redirect hot air away from servers inside the rack.

Before installing switches, download the Aruba Installation Guide for the specific models. Review the Installation Guide before installing and deploying the switches. Carefully review requirements for power, cooling, and mounting to ensure that the data center environment is outfitted adequately for safe, secure operation.

**Step 1** Open a web browser and navigate to the Aruba Support Portal at https://asp.arubanetworks.com/.

**Step 2** On the **Support Portal** page, select the **Software & Documents** tab.

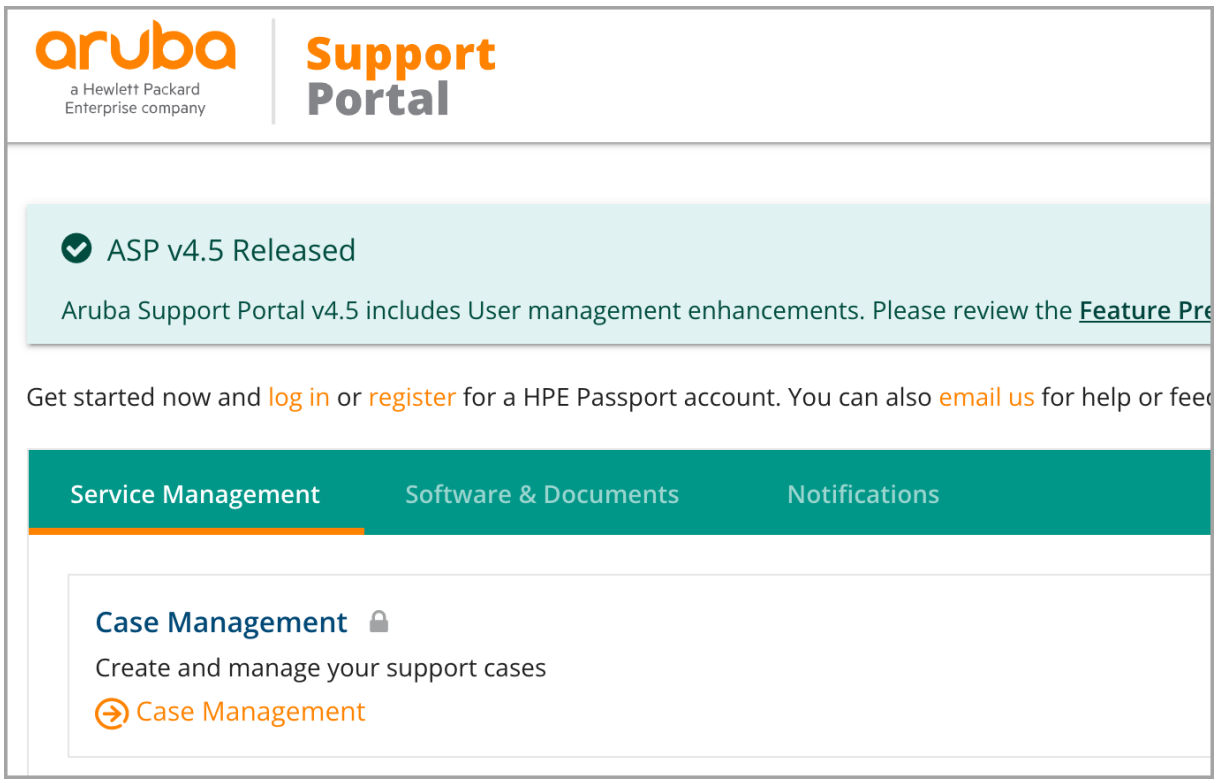


**Step 3** On the **Software & Documents** tab, select **Switches**.

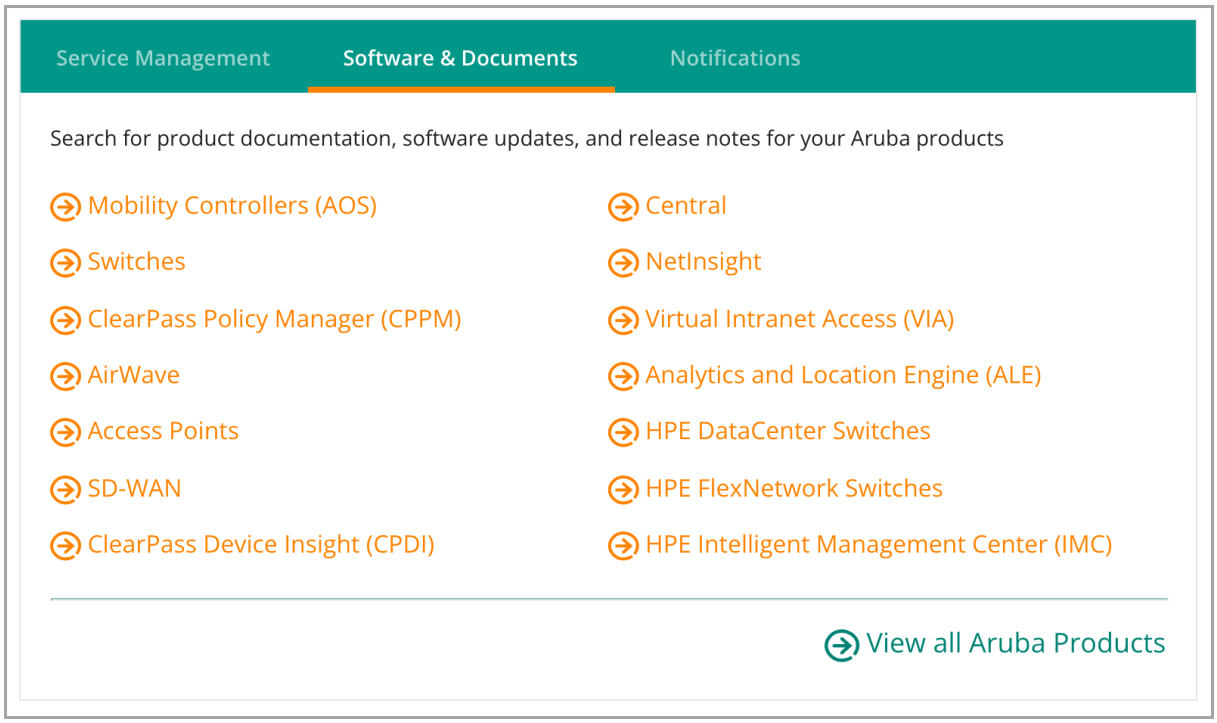


**Step 4** Select the filter options on the left.

- **File Type:** *Document*
- **Product:** *Aruba Switches*

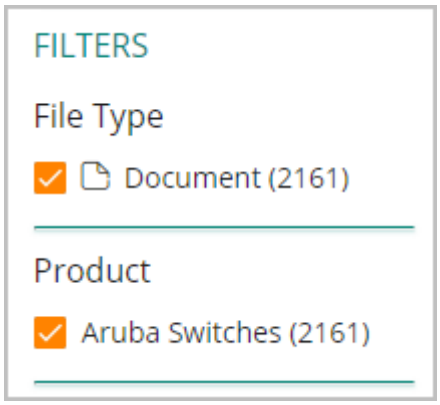


- **File Category:** *Installation Guide*

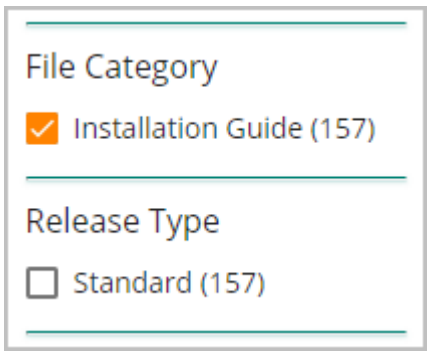


**Step 5** Download the Installation Guide version for the switch model to be installed.

**Step 6** Complete the physical installation of switches in the racks.

**NOTE:**

Core switches should be installed in a central location that meets cable distance requirements for the media used between core and access switches. Access switches should be installed at top-of-rack in high-density environments or middle-of-row in low-density environments.

# Physical Cabling

Consistent port selection for core and access switches increases the ease of configuration management, monitoring, reporting, and troubleshooting tasks in the data center.

Document all connections.

Ensure that distance limitations are observed for the chosen host connection media and between switches.

Refer to the “Data Center Design” section for guidance on cabling design options for the installation.

## Top of Rack Cabling

In a redundant ToR configuration, the first two uplink ports should be allocated to interconnect redundant peers (ports 49-50 on 8325-48Y8C and 10000-48Y6C switches), which provides physical link redundancy and sufficient bandwidth to accommodate a core uplink failure on one of the switches.

Two links between redundant peers are sufficient for most deployments unless the data center host implementation may result in high-traffic use of the inter-switch links under normal operating conditions, such as when many hosts in a rack are single-homed to only one of the redundant switches, or when the VM guest redundancy model uses only a single link to transmit traffic.

The heartbeat between a ToR redundant pair should be configured to use the out-of-band management port. Alternatively, the highest numbered non-uplink port can be used as a heartbeat link.

Before deploying ToR configurations that require server connectivity at multiple speeds, review the switch guide to determine if adjacent ports are affected by changing port speeds.

## Core-to-Access Cabling

The illustration below shows the physical port configuration on an 8325 32-port core switch.

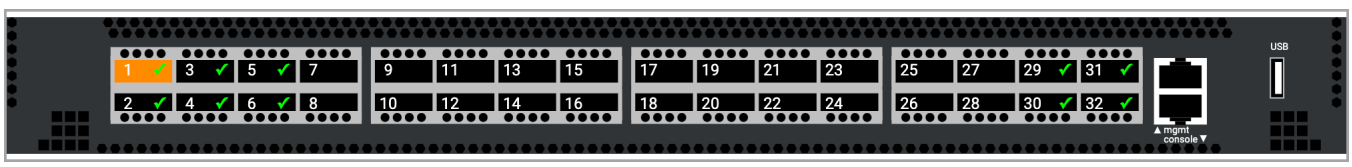


**Figure 8:** Spine switch

The core switches in a two-tier data center operate as a redundant VSX pair. The last two ports in a 1U core switch should be allocated to the inter-switch link (ISL) between them (ports 31-32 on an 8325-32C). When using a chassis-based switch model, the inter-switch links should use the last port on two different line cards to add line card diversity, enabling the ISL to continue to function in case of a single line card failure.

Connections from the core to access switches should begin with port 1. In a dual ToR configuration, each core switch must be connected to each ToR redundant switch. A 32-port core switch supports up to 14 racks in this design, after considering the inter-switch and external connectivity links. Use the same port number on each core switch to connect to the same access switch to simplify switch management and documentation. For example, assign port 1 of each core switch to connect to the same access switch.

# Management

## Out-of-Band Management

For an Aruba ESP data center, using a dedicated management LAN is strongly recommended. A dedicated management LAN on separate physical infrastructure ensures reliable connectivity to data center infrastructure for automation, orchestration, and management access. CX data center switches should be connected to the management network using their physical out-of-band (OOB) management port.

Deploy management LAN switches at the top-of-rack position. Plan for an IP subnet with enough capacity to support all switch and host OOB management ports in the data center. DNS and NTP services for the data center should be reachable from the out-of-band management network. The management LAN also must allow outbound connectivity to the Aruba Central cloud management platform.

Configuration steps for the management LAN are not covered in this guide.

## DHCP

A new switch must receive an IP address, DNS server address, and a default gateway via DHCP in order to enable successful Zero Touch Provisioning.

In a data center, it is recommended to reserve an IP address for each switch on a DHCP server. This ensures a predictable IP address for local management connections, such as SSH, while also enabling the switch to contact Aruba Central on boot.

When switches are unpacked and prepared for bench configuration or rack mounting, access the orange luggage tag on the switch and record the base MAC address. The MAC address of the management interface is the base MAC + 1. For example, A0:A0:01:00:00:00 becomes A0:A0:01:00:00:01. Use this management MAC address to create a DHCP reservation.

## Static Addressing

Static IP and gateway address assignments can be used when Zero Touch Provisioning is not required. A DNS server also must be assigned to enable reachability to Aruba Central.

## Central On-Premise

When organizational policy requires on-premise management of data center infrastructure, Aruba Central On-Premise (CoP) enables customers to run the Aruba Central management platform on local infrastructure. CoP supports CX 8xxx and 6400 series switches.

This guide does not cover the installation or use of CoP. Refer to the *Aruba Central On-Premises Supported Devices Reference Guide*, *Aruba Central (on-premises) User Guide*, and *Aruba Central (on-premises) Release Notes* in the Aruba Support Portal for additional information.

# Aruba Central Initial Configuration

This guide demonstrates initial configuration of Aruba Central UI *groups* and *sites* and the assignment of switches to both configurations.

Switches with a common set of configuration elements should be assigned to the same Central UI group. In a Two-Tier data center, one group is defined for core switches and a second group for access switches.

Configuration common to both groups is defined once for the core group, then cloned to create the access group.

Features configured at the group level include:

- the hostname,
- admin account password,
- Network Time Protocol (NTP) servers,
- Domain Name System (DNS) servers,
- VLANs,
- Spanning-Tree,
- Terminal Access Controller Access Control System (TACACS) servers,
- Authentication, Authorization, and Accounting (AAA) servers.

Best practice is to use as few groups as necessary to provide logical organization for the network and consistent configuration among devices. Configuration is not shared among groups.

Switches in the same physical location can be assigned to the same Central Site to aggregate statistics, reporting, and troubleshooting tools.

## Create Core Switch Group

**Step 1** Open HPE GreenLake in a web browser and login with administrator credentials.

**Step 2** Locate the account associated with the data center switches and click **Go to Account**.

**Step 3** Click **Launch** in the **Aruba Central** tile.

**Step 4** On the left **Aruba Central** navigation menu, click **Global**.

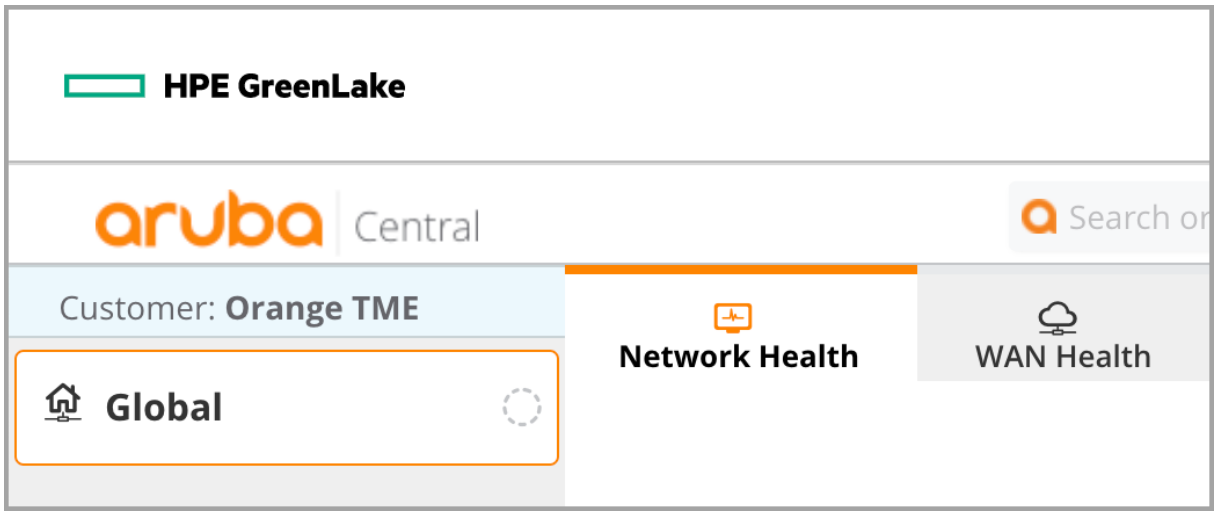


**Step 5** On the dropdown, click the **Groups** column title.

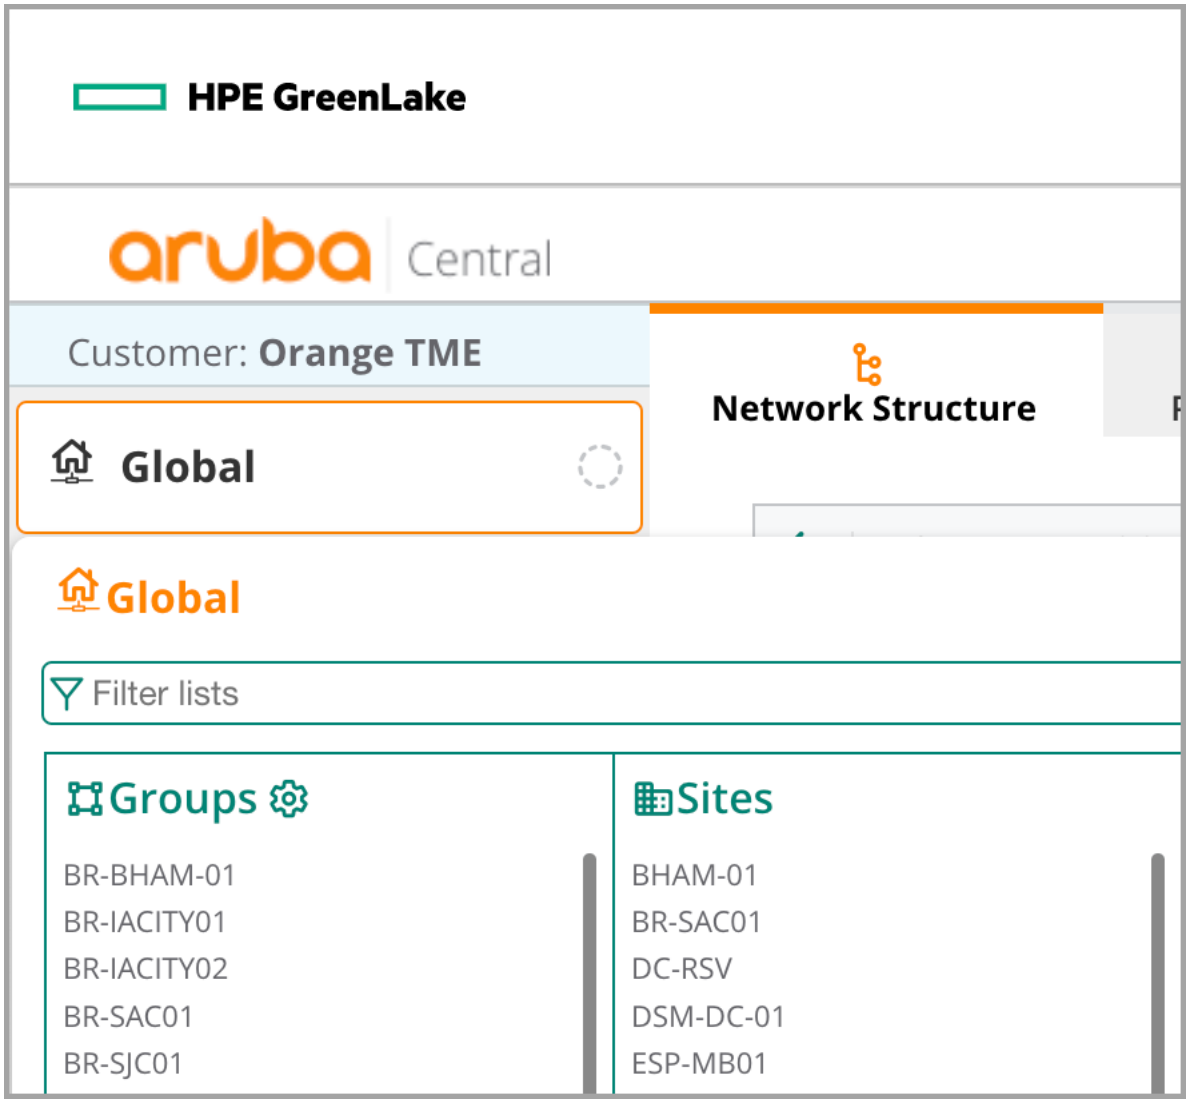


**Step 6** At the upper right corner of the **Groups** table, click the **+** (plus sign).

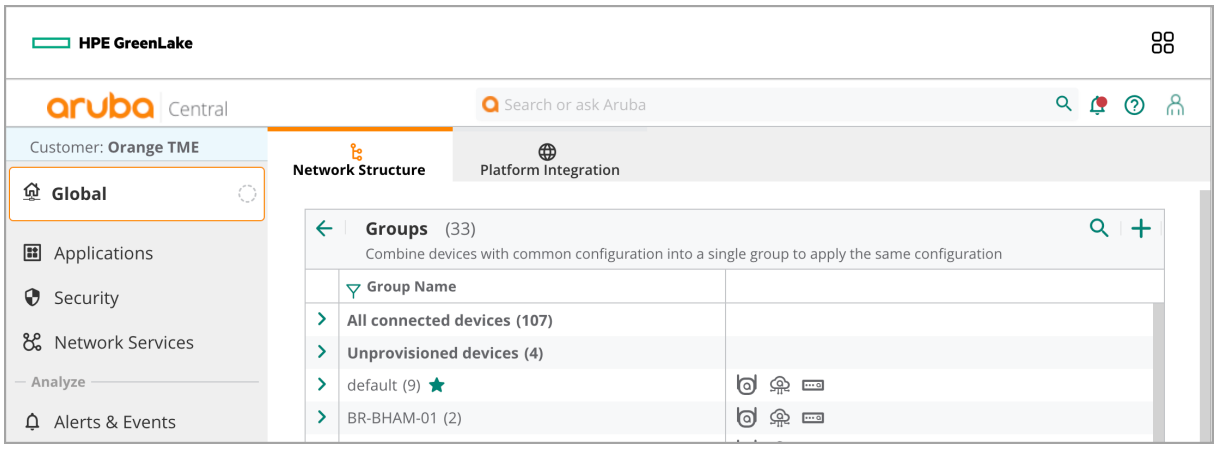


**Step 7** In the **Add Group** wizard, enter a UI group name for the data center core switches in the **Name** field, click the **Switches** checkbox, then click **Next**.

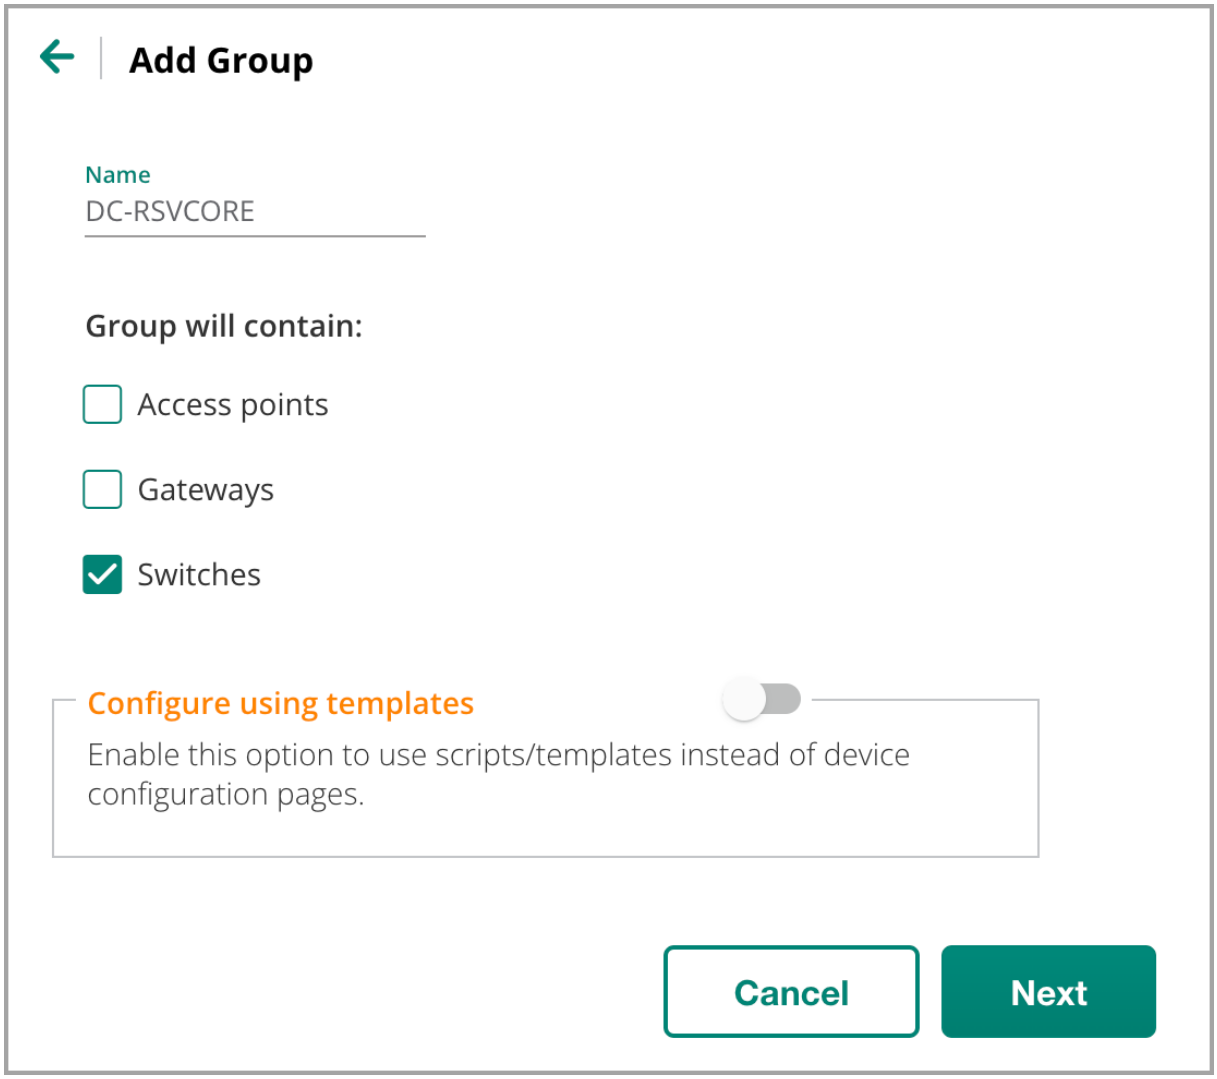


**Step 8** Click the **AOS-CX only** radio button, then click **Add**.

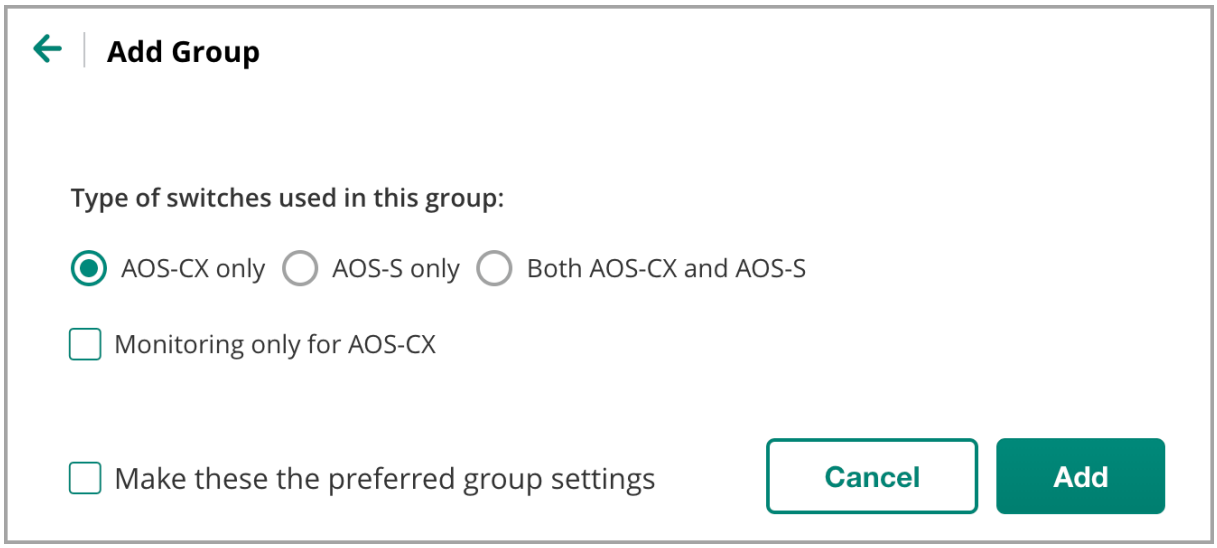


## Configure Core Switch Group Settings

Device configuration can be performed at the UI group level or at an individual device level. In this section, the core switch group is defined with initial configuration settings. Configuration defined at the group level is applied to all member switches, but exceptions can be made at the individual device level.

The access group is cloned from the core group, inheriting all group level settings at the time of cloning. In this initial configuration of the core switch group, only values shared with the access group are defined.

**Step 1** On the left navigation menu, click **Global**.

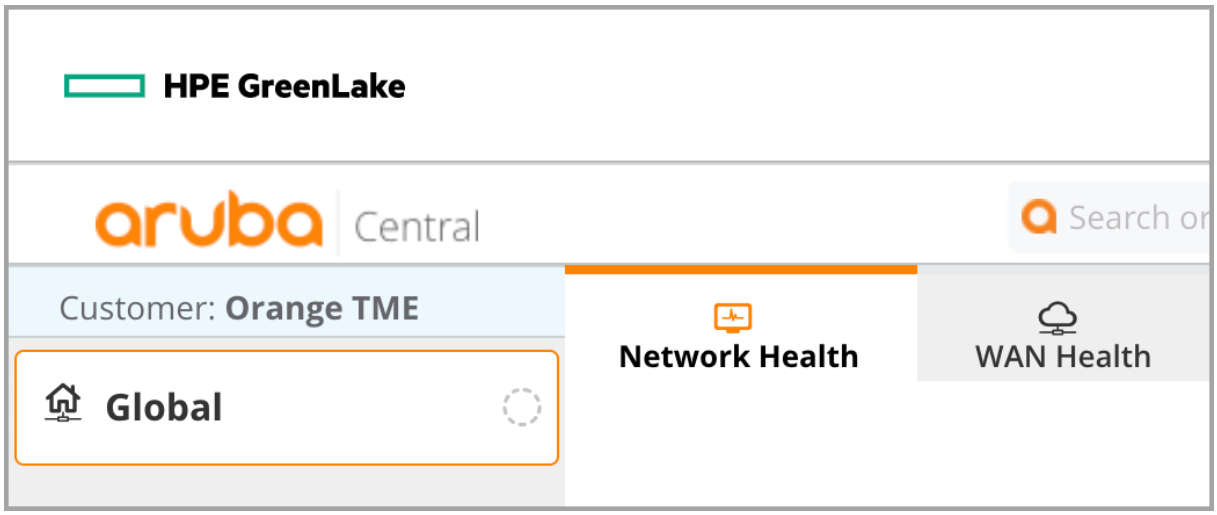


**Step 2** In the **Groups** column, click the data center core switch group name.

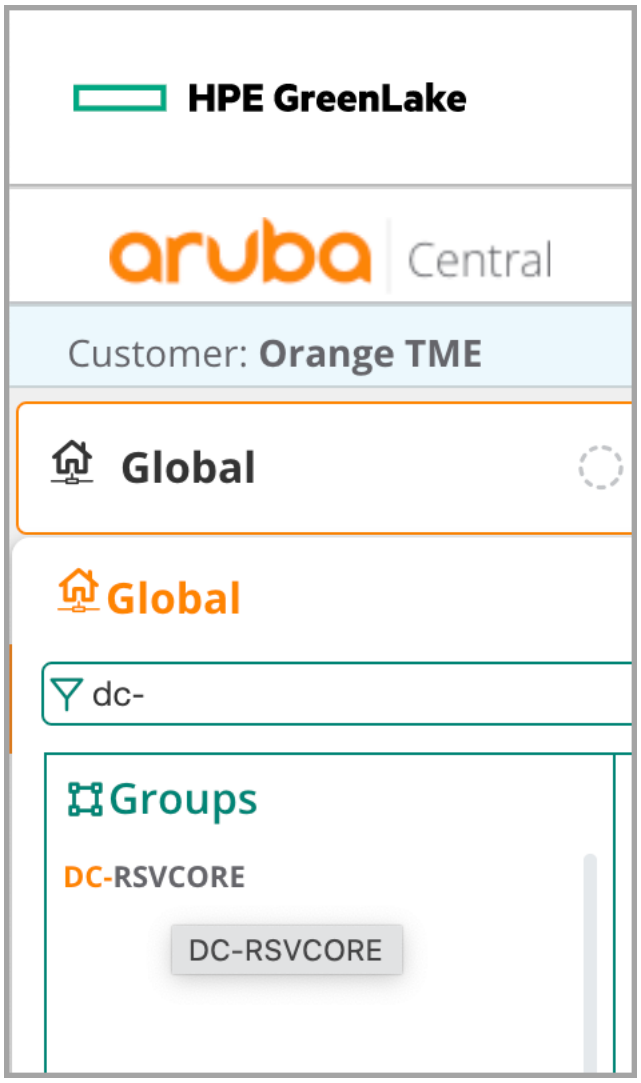


**NOTE:**

After clicking the group name, text can be entered immediately in the **Filter lists** field. Enter a portion of the switch group name in the **Filter lists** field to display items containing the string. The filter applies simultaneously to the groups, sites, and labels columns.

**Step 3** On the left navigation menu, click **Devices**.

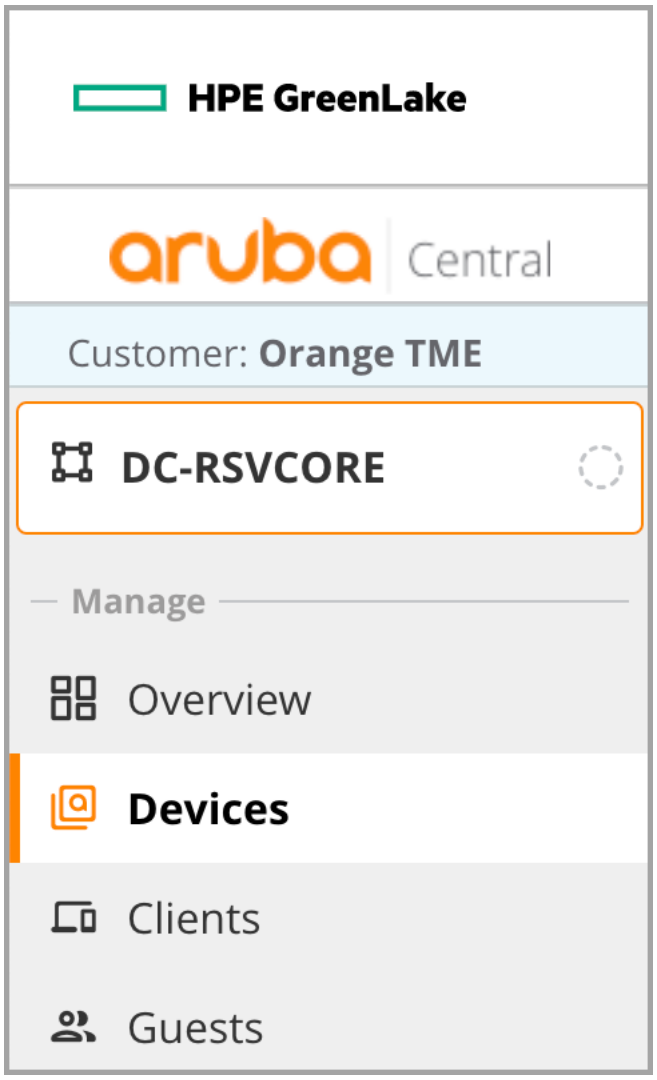


**Step 4** At the upper right of the **Switches** pane, click **Config**.

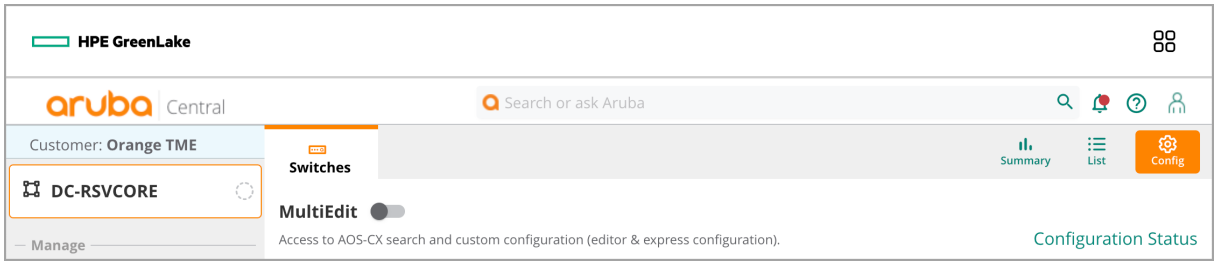


**Step 5** In the **Set Device Password** window, enter the switch **Administrator password** value, then click **SAVE**.

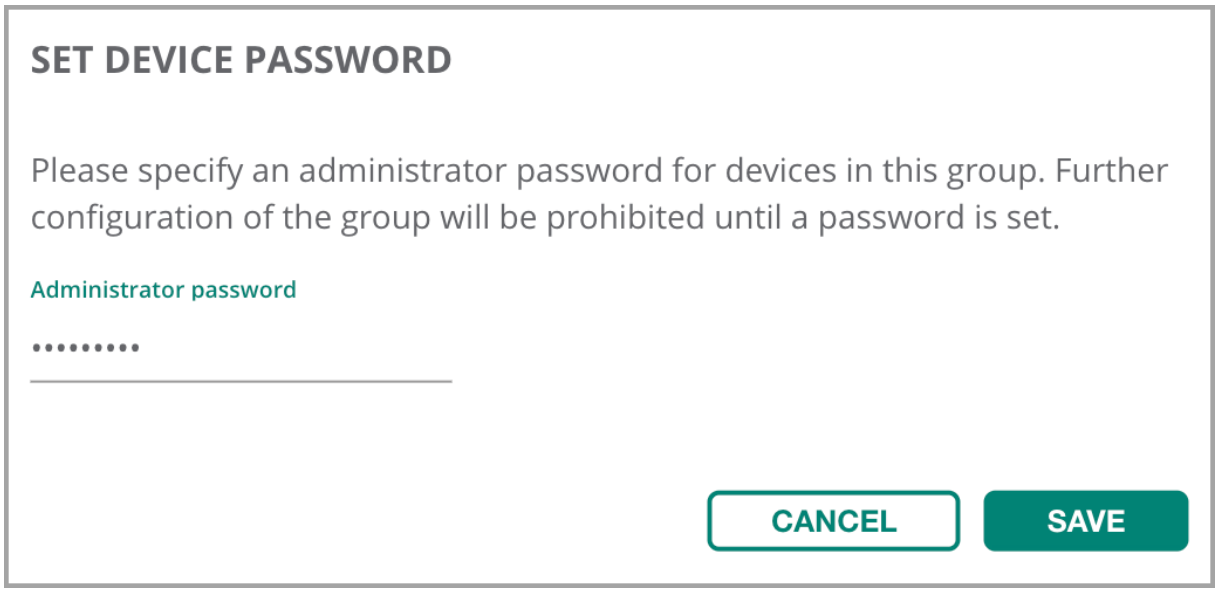


**Step 6** In the **System** tile, click **Properties**.

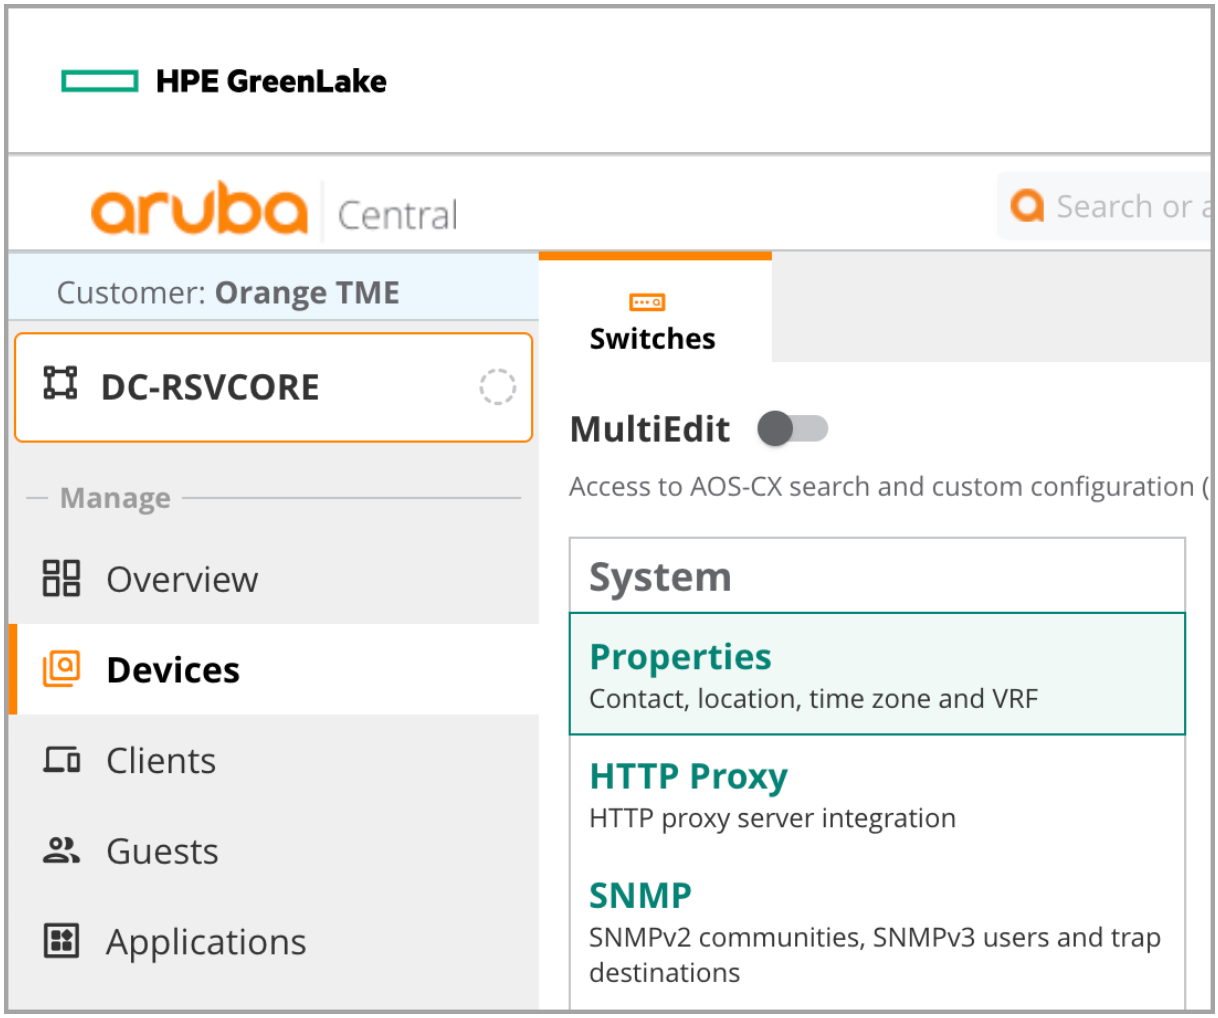


**Step 7** In the **Edit Properties** window, assign the following values, then click **SAVE**. - **Contact:** *ne-tadmin@orangetme.local* - **Location:** *DC01, Roseville, CA* - **Timezone:** *Los Angeles (UTC-8:00)* - **VRF:** *Management* - **DNS servers:** *10.2.120.98, 10.2.120.99* - **NTP servers:** *10.2.120.98, 10.2.120.99*

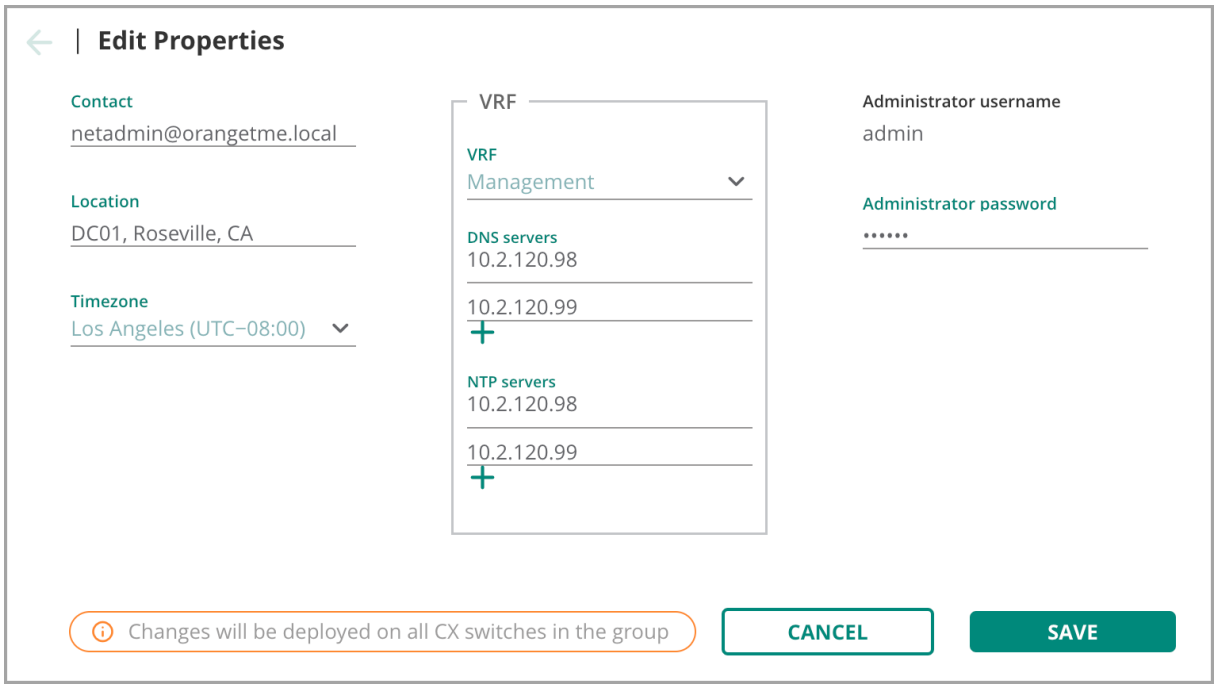


**NOTE:**

Set the VRF to the network where DNS and NTP will be reachable. When using a dedicated out-of-band management network connected to the mgmt interface, the *Management* value should be selected for **VRF**.Enter a complete IP address for DNS and NTP servers to make the **+** (plus sign) appear for entering additional servers.

**Step 8** In the **Security** tile, select **Authentication Servers**.

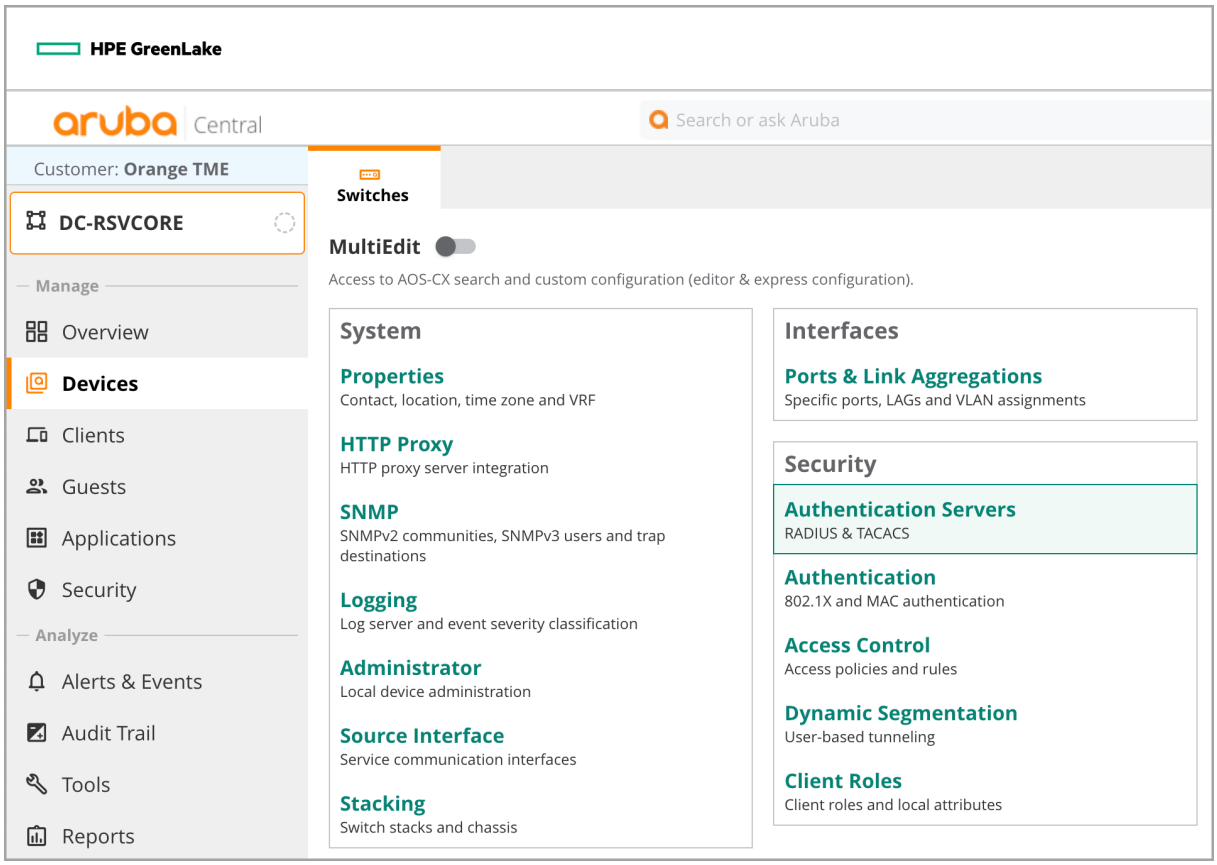


**Step 9** Mouse-over the **TACACS** row. At the far right of the **TACACS** row, click the **edit** icon (pencil).

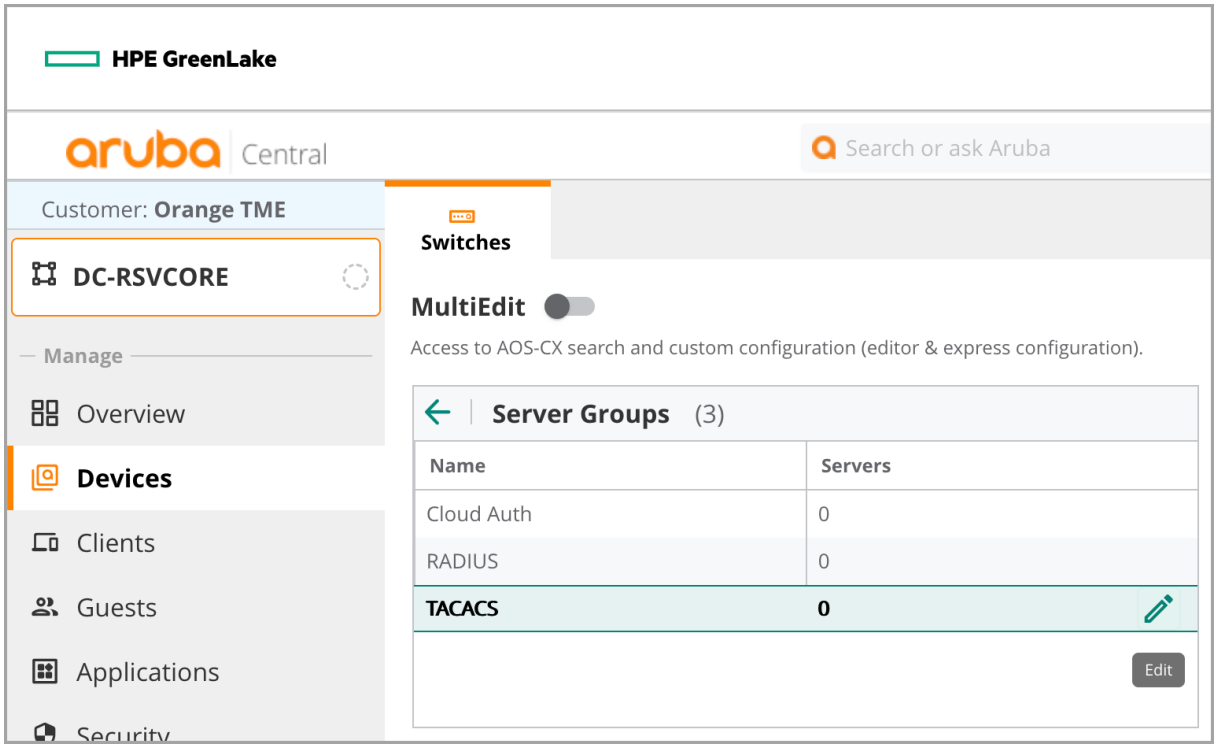


**Step 10** At the top right corner of the **TACACS Servers** table, click the **+** (plus sign).

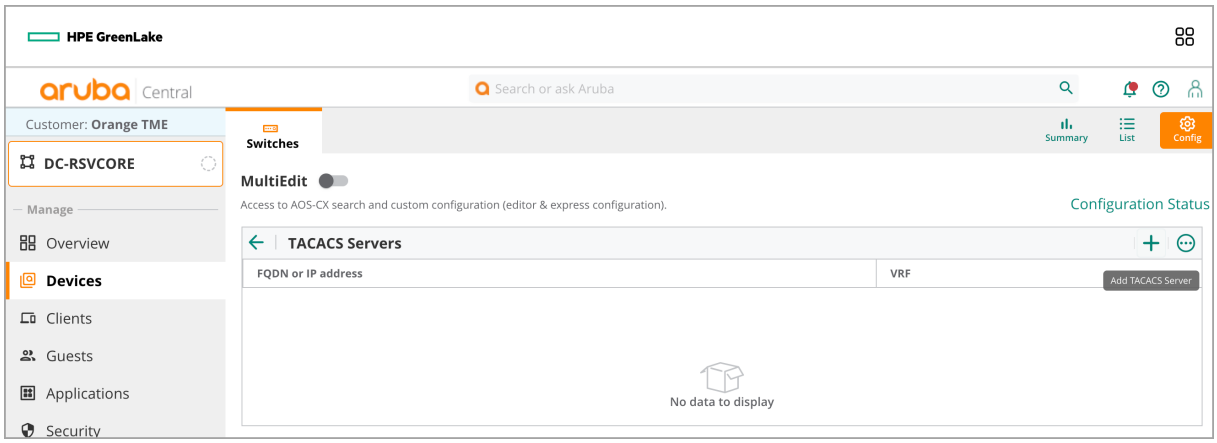


**Step 11** Assign the following settings on the **Add TACACS** page, then click **Apply**. - **FQDN or IP address:** *10.2.120.94*

- **Authentication Port:** *49* (default)
- **VRF:** *Management*
- **Shared secret:** *< shared secret >*
- **Timeout (secs):** *5* (default)

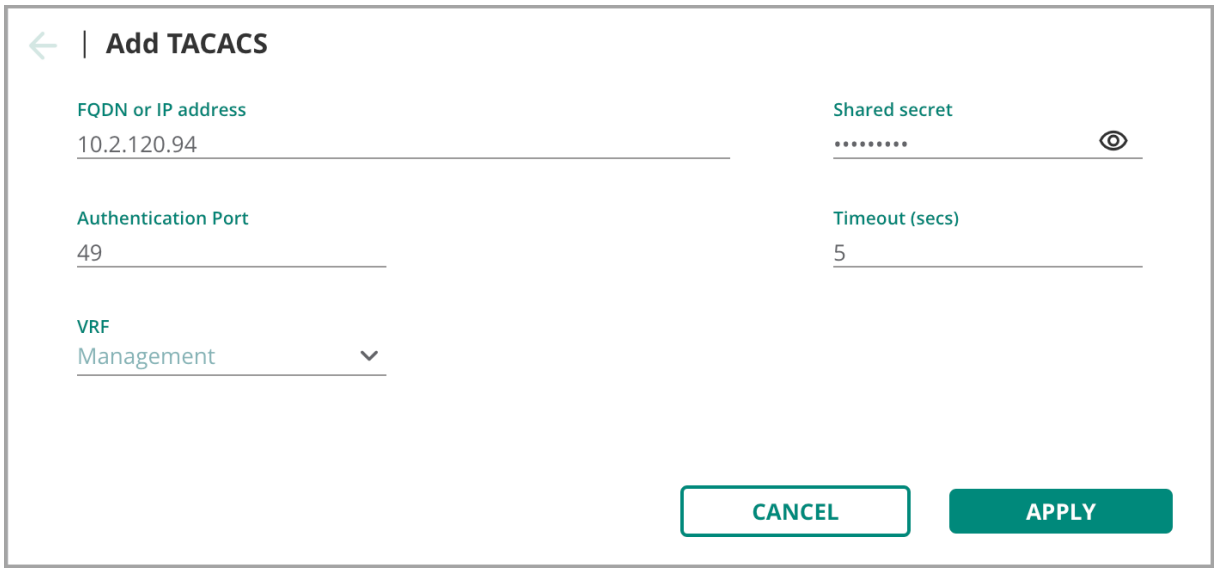


**Step 12** Assign additional servers by clicking the **+** (plus sign) at the top right corner of the **TACACS Servers** table.

**Step 13** After all servers are added, click **SAVE**.

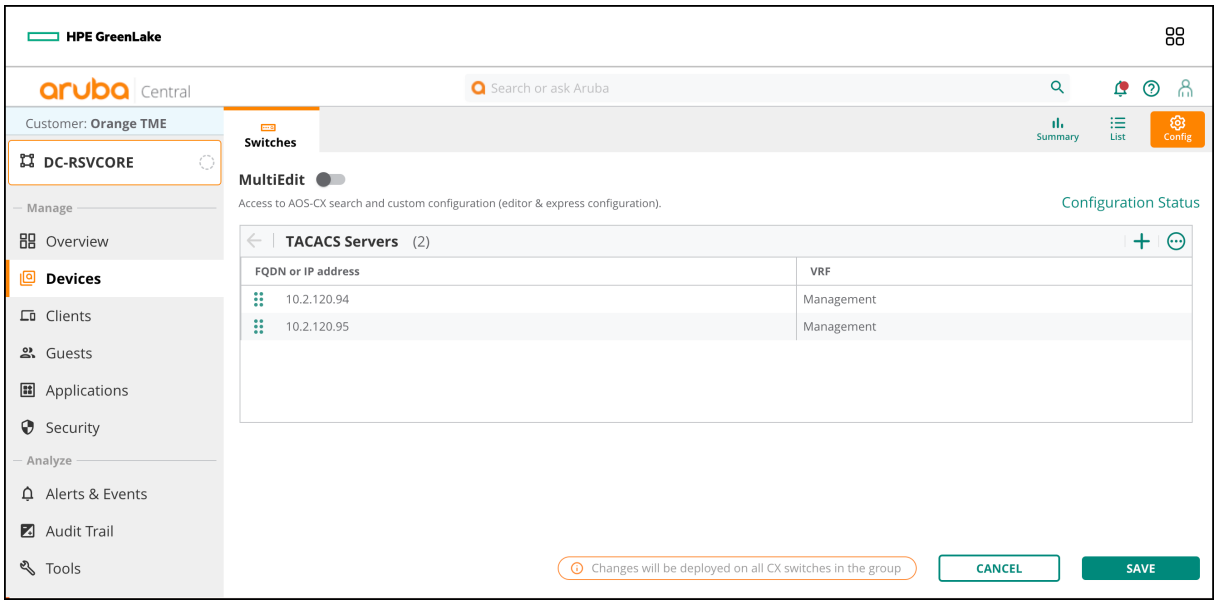


**Step 14** At the top left of the **Server Groups** table, click ← (left arrow).

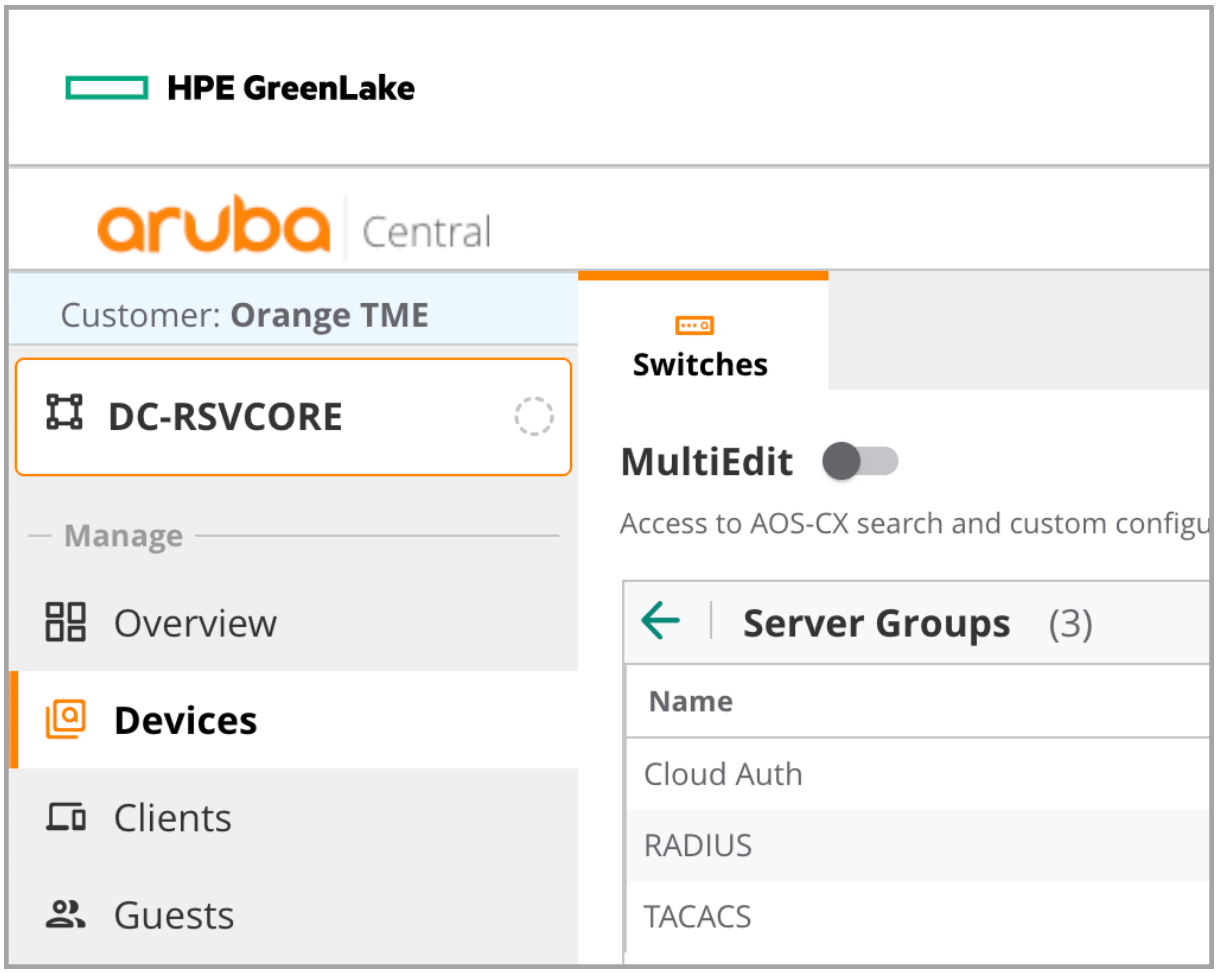

# Configure Group VLANs

Define the data center host VLANs shared by both core and access switches. Additional VLANs required only at the core level are defined in subsequent steps.

**Step 1** In the **Bridging** tile, click **VLANs**.

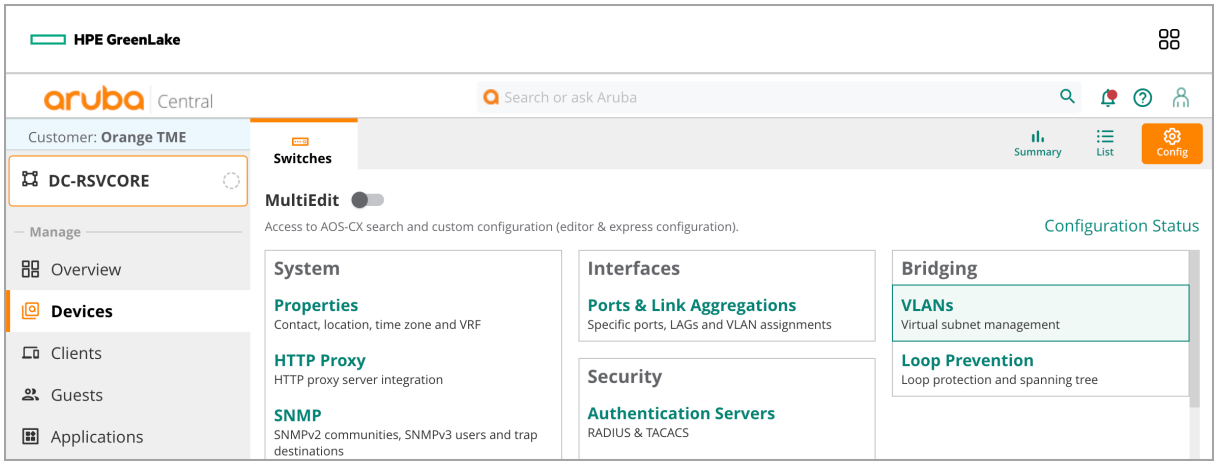


**Step 2** At the top right of the **VLANs** table, click the **+** (plus sign).

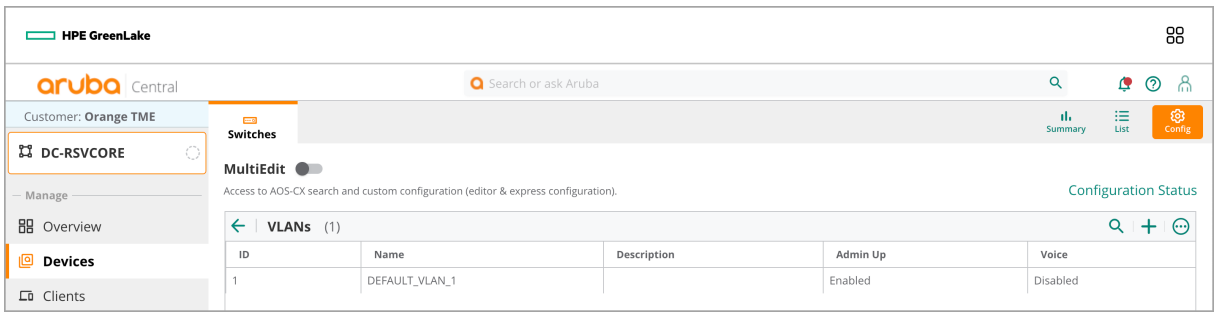


**Step 3** On the **Add VLAN** page, enter the following field values, then click **ADD**.

- **ID:** *101*
- **Name:** *PROD-WEB*
- **Description:** *< no value >*
- **Admin Up:** *< checked >*
- **Voice:** *< unchecked >*

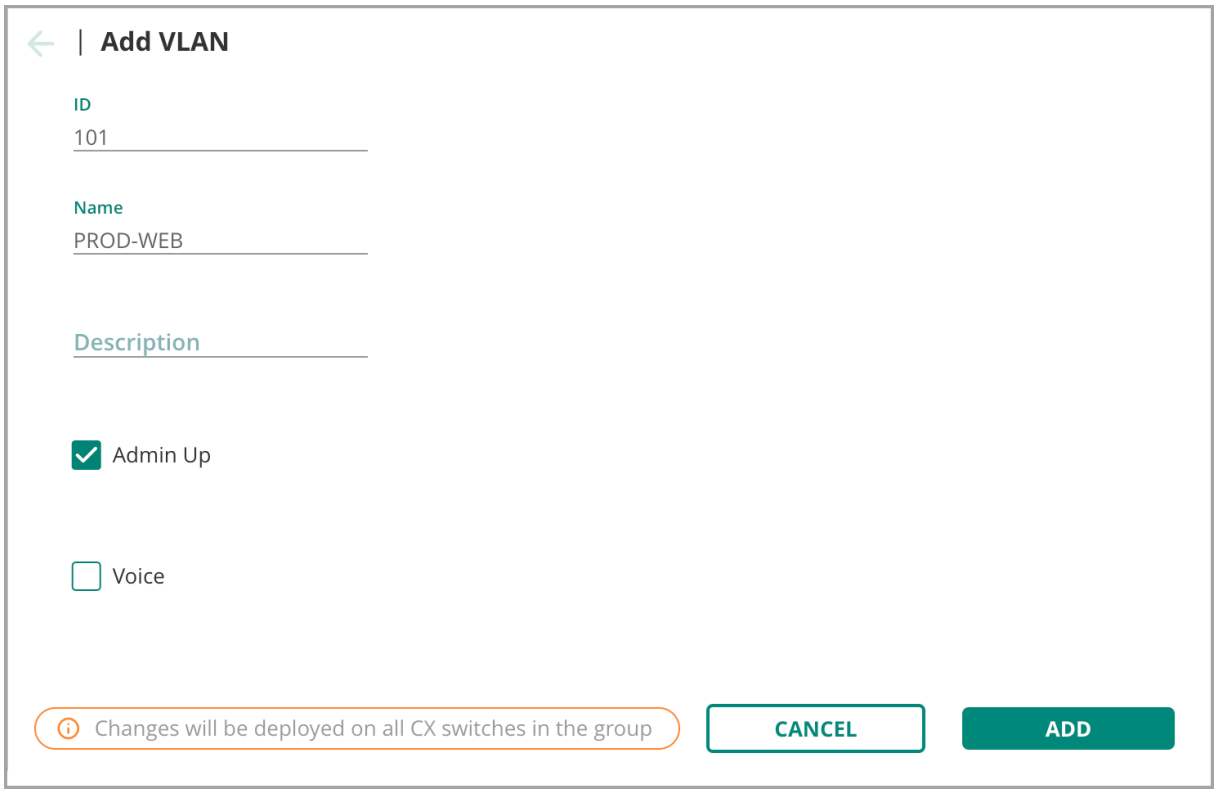


**Step 4** Repeat this process to create additional data center host VLANs.

- **ID:** *102*
- **Name:** *PROD-DB*
- **Description:** < *no value* >
- **Admin Up:** < *checked* >
- **Voice:** < *unchecked* >

## Clone Core Group for Access Switches

Aruba Central supports cloning group settings to a new UI group. Follow the steps below to create an access switch group that contains the same initial settings configured for the core switch group.

**Step 1** On the left navigation menu, click **DC-RSVCORE**.

**Step 2** On the dropdown, click the **Groups** column heading.

**Step 3** Enter search text in the **Group Name** column to filter displayed groups.

**Step 4** Move the cursor to the right of the name, and click the **Clone Group** icon.

**Step 5** On the **Clone Group** page, enter a name for the data center access switch group in the **Name** field, and click **Clone**.

**Step 6** Verify that the new access switch group is populated in the list of groups.

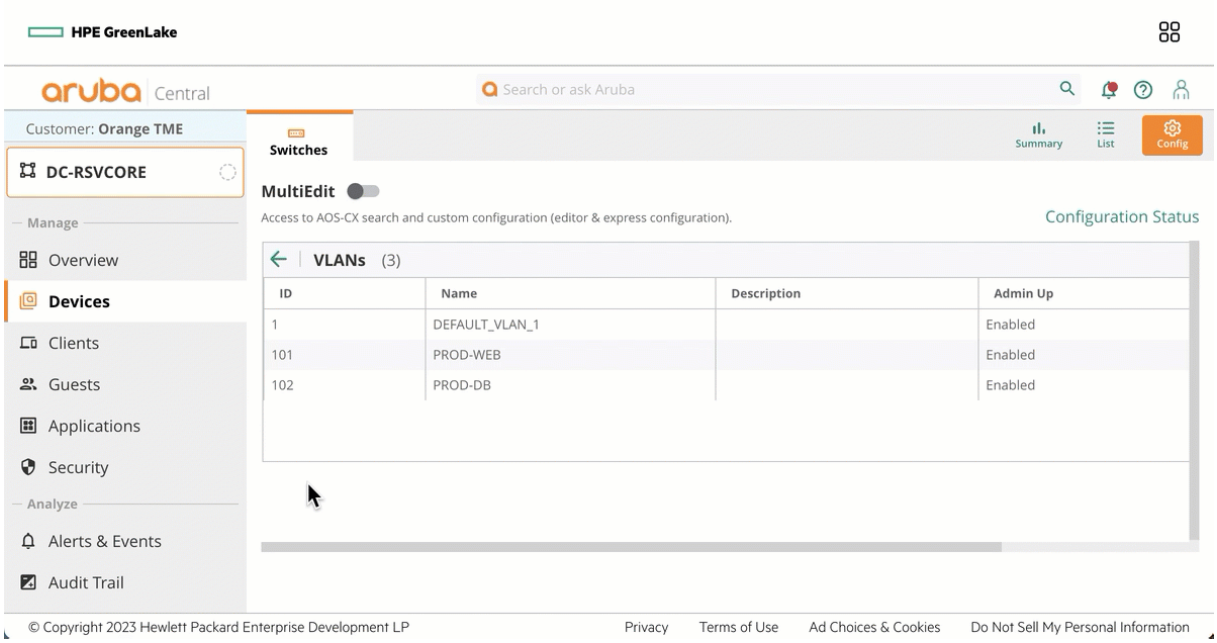


## Assign Switches to a Group

Use this procedure to assign switches to groups and synchronize initial configuration.

**Step 1** Expand the **Unprovisioned devices** group by clicking **>** (greater than) next to its name.

**NOTE:**

Only switches that were not previously configured in Central appear in the **Unprovisioned devices** group. If switches were removed from a different Central group for use in the data center, they appear in the **default** group.Switches new to Central must be identified by the unique serial number or MAC address.

**Step 2** Click both core switches.

**Step 3** In the lower right corner of the **Unprovisioned devices** group table, click the **Move** button.

**Step 4** On the **Move Devices** page, select the appropriate destination group for the selected switches and click **Move**.

**Step 5** Expand the list of switches for the destination group to verify that the switches moved.

**NOTE:**

The search filter in the top right corner of the Groups window can filter content to display devices matching the search criteria and the groups that contain those devices. The search criteria can match text in the **Name**, **Type**, **Serial Number**, and **MAC address** fields.Displayed group names also can be filtered by clicking the **Group Name** column heading and entering match criteria.

**Step 6** Repeat this procedure to assign access switches to the access switch group.

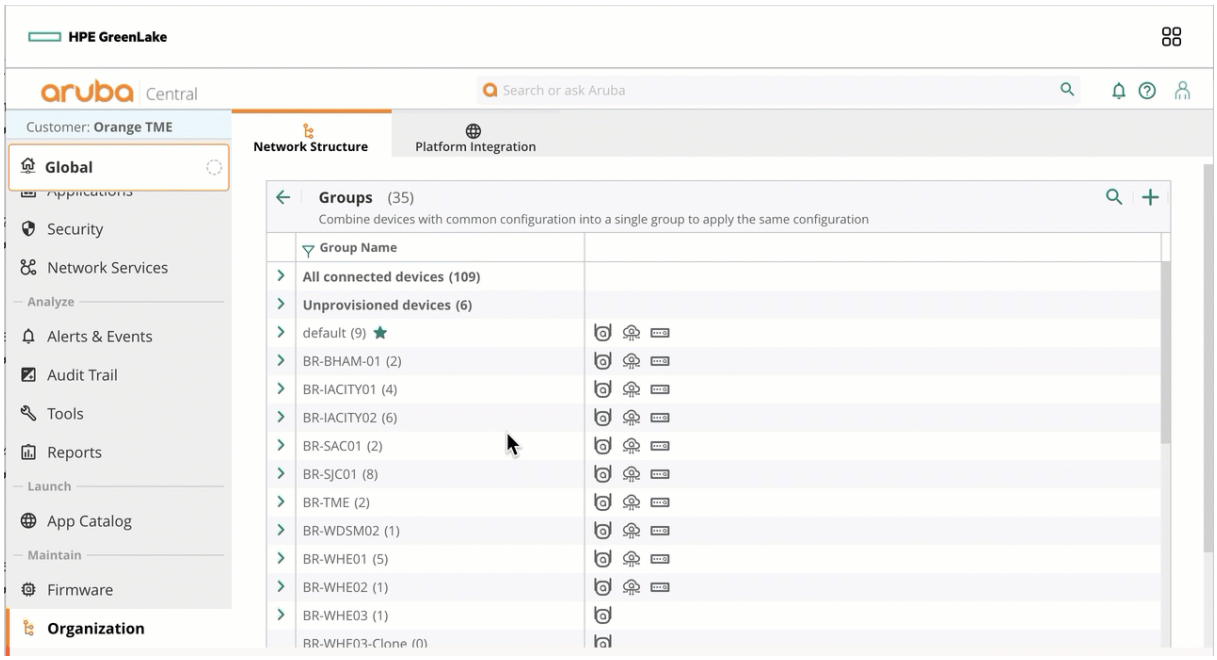


# Configure Switch Hostname

**Step 1** On the left navigation menu, click **Global** and select the data center core switch group.

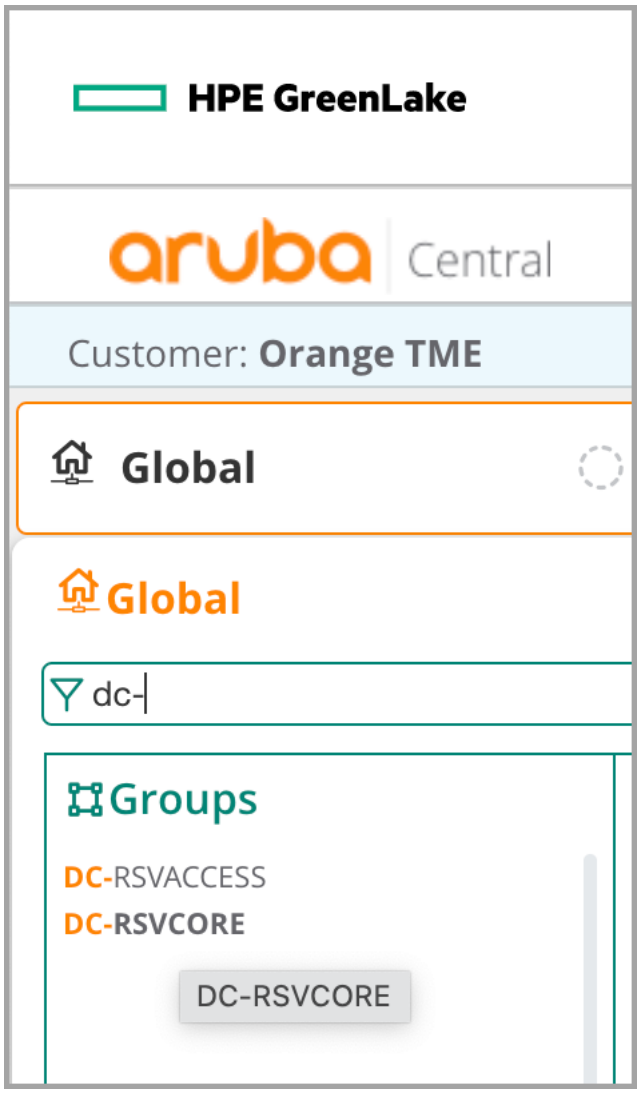


**Step 2** On the left navigation menu, select **Devices**.

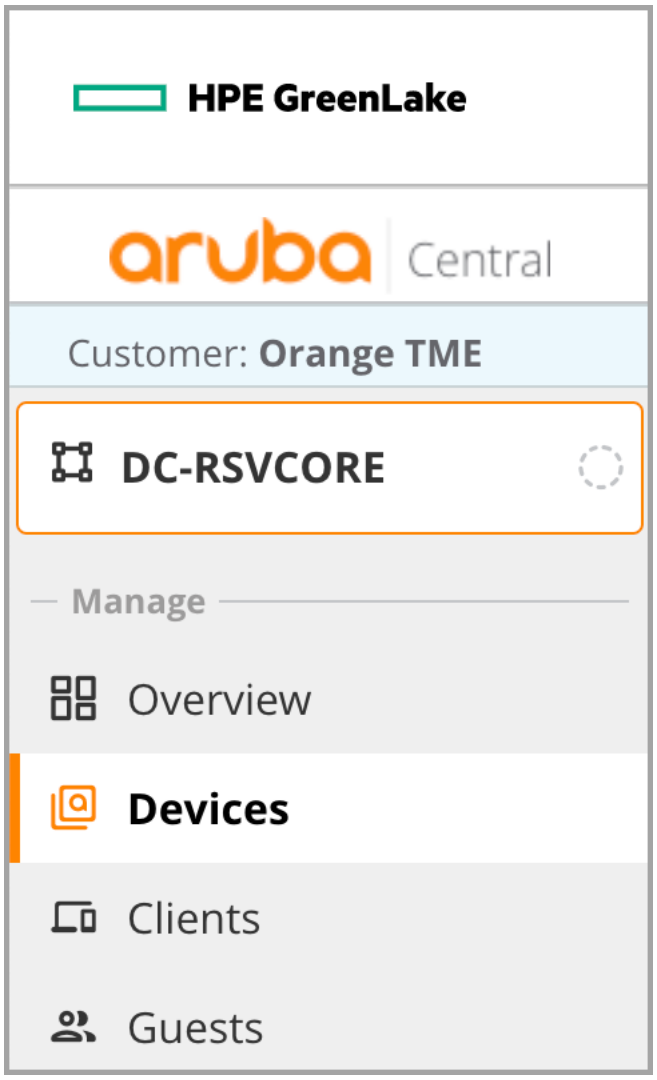


**Step 3** Scroll horizontally to view the identifying serial number and click the name of the recently added switch.

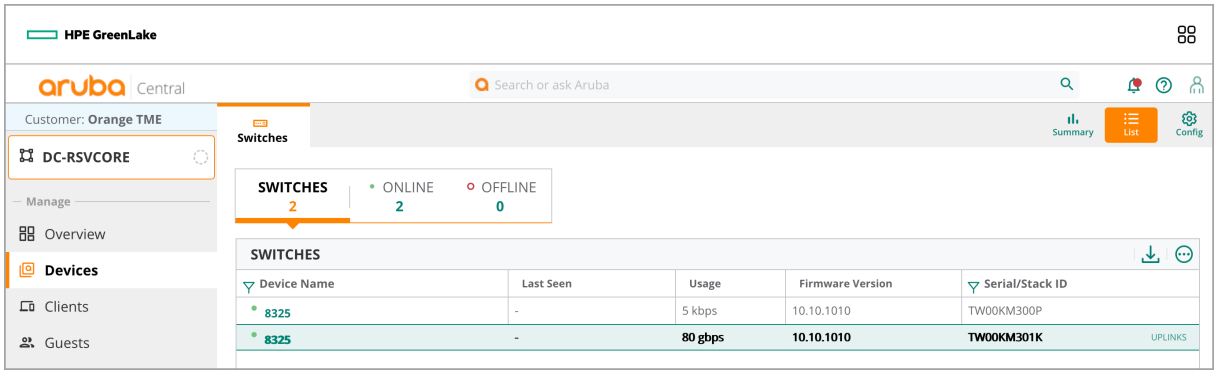


**Step 4** On the left navigation menu, click **Device**.

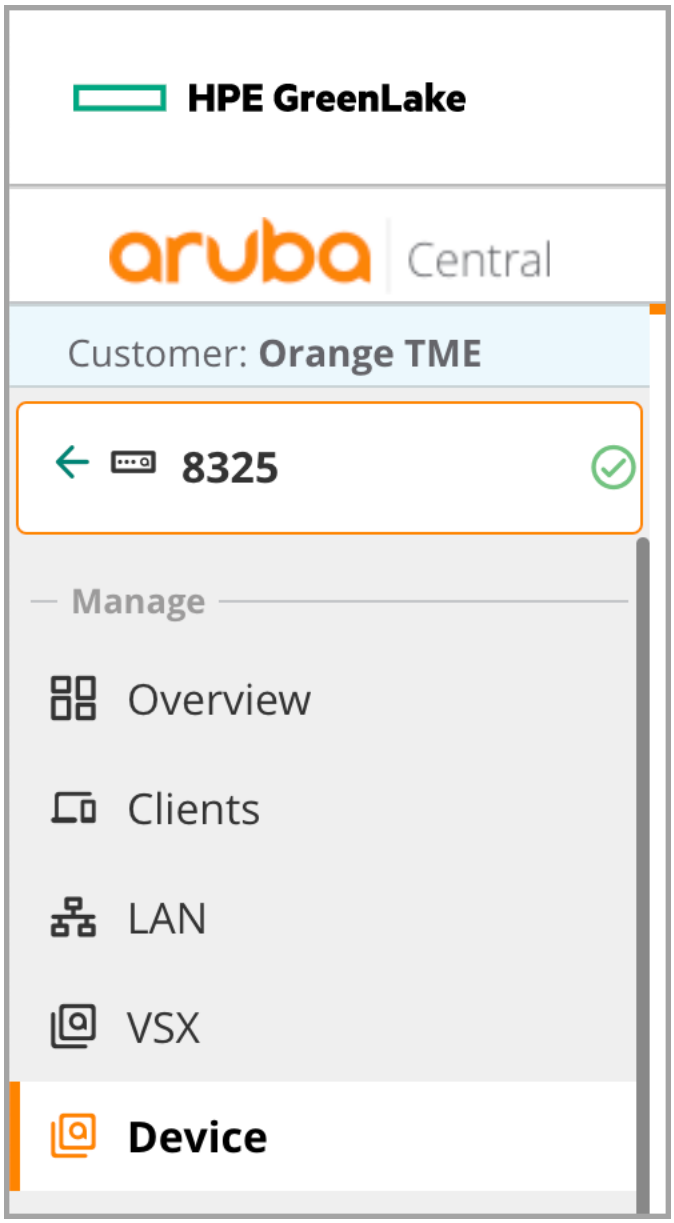


**Step 5** In the **System** tile, click **Properties**.

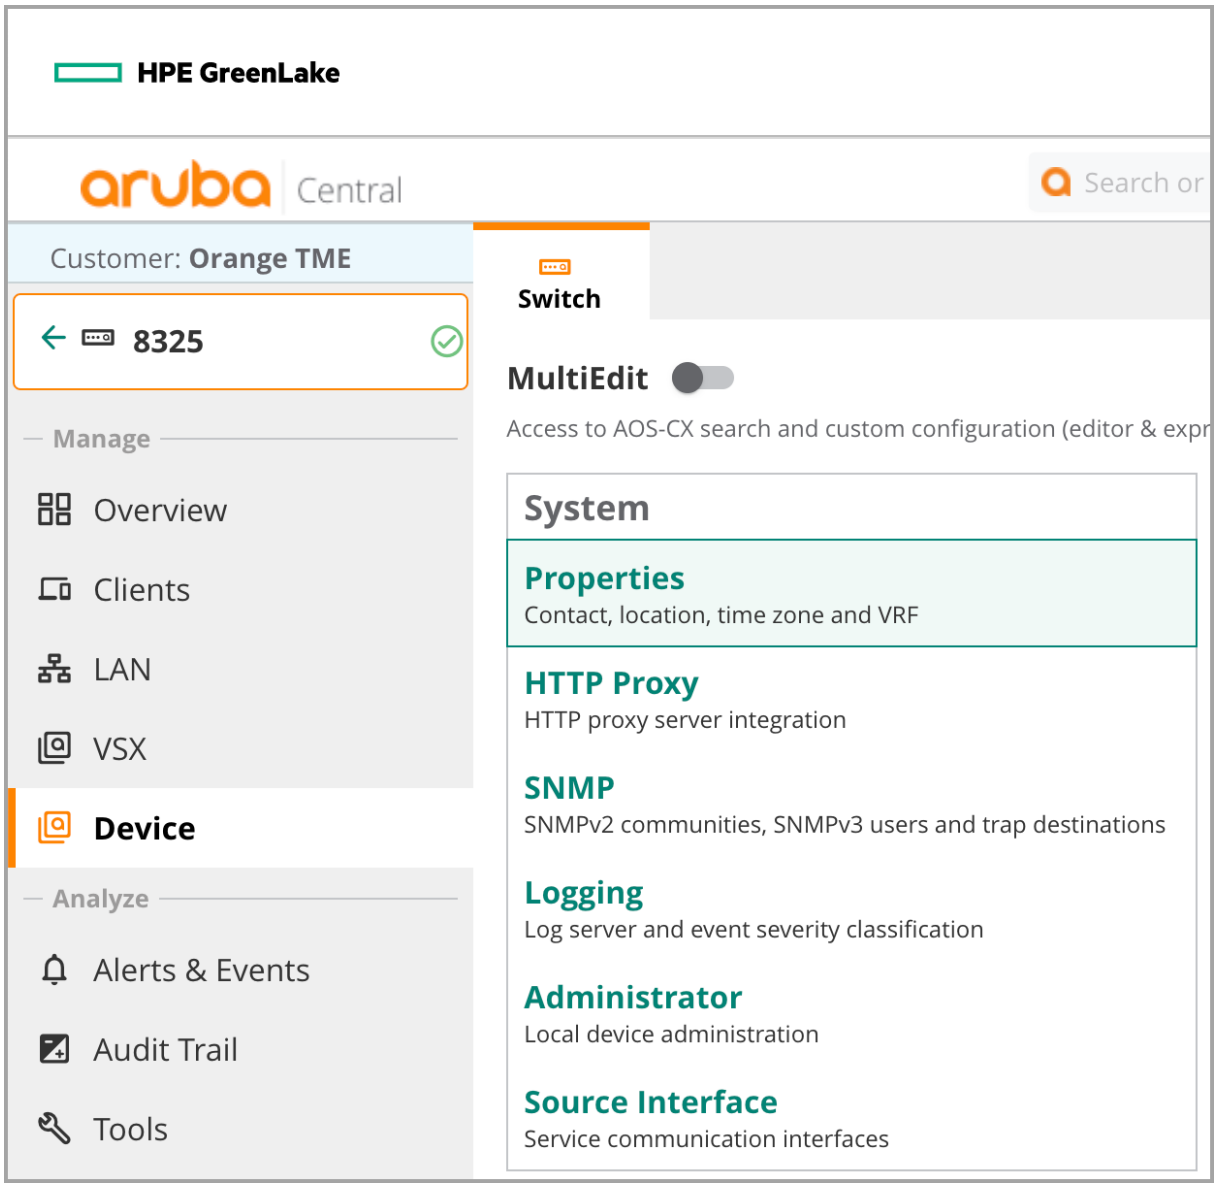


**Step 6** On the **Edit Properties** page, enter a hostname for the switch in the **Name** field, then click **SAVE**.

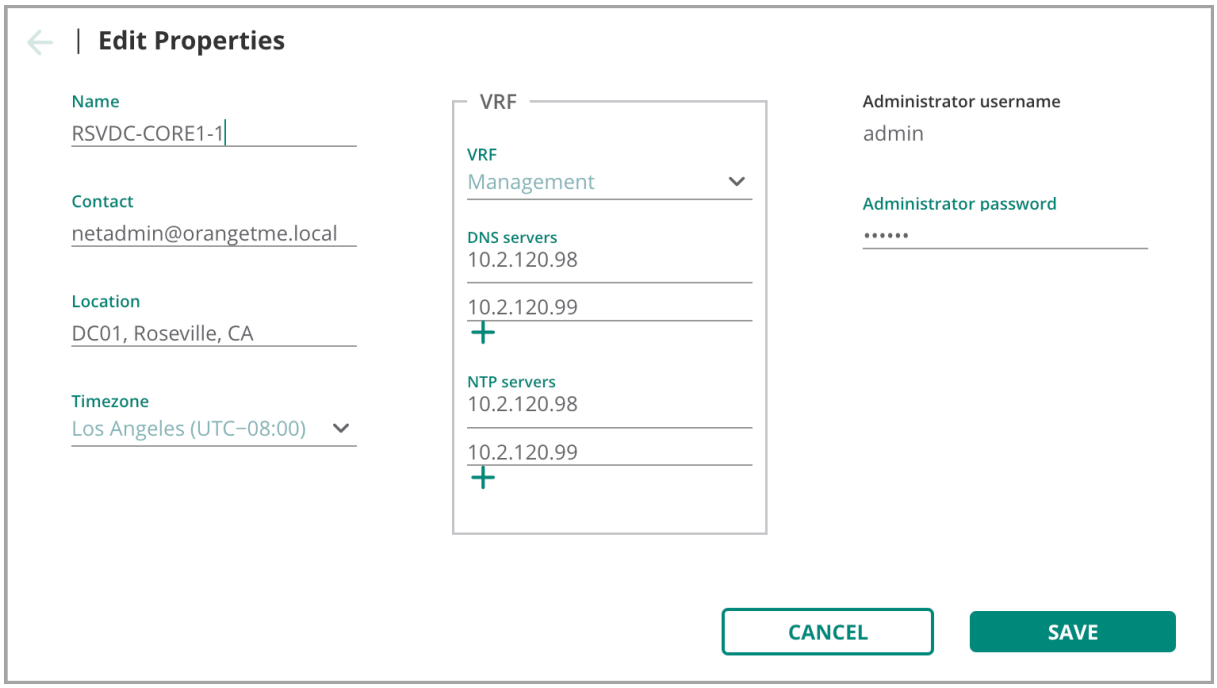


**Step 7** To return to the device list, at the top of the left navigation menu, click ← (left arrow) next to the switch name.

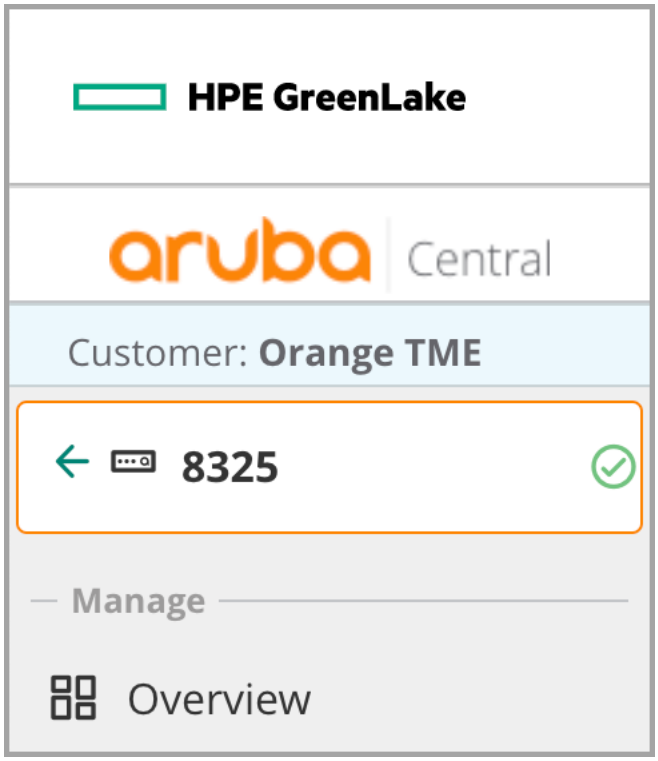


**Step 8** Repeat this procedure to assign a hostname for each switch in both data center groups.

**NOTE:**

It may take several minutes for a newly assigned name to display in the Central device list.

## Configure Data Center Site

**Step 1** At the top of the left navigation menu, click the current switch, then click the **Sites** column heading.

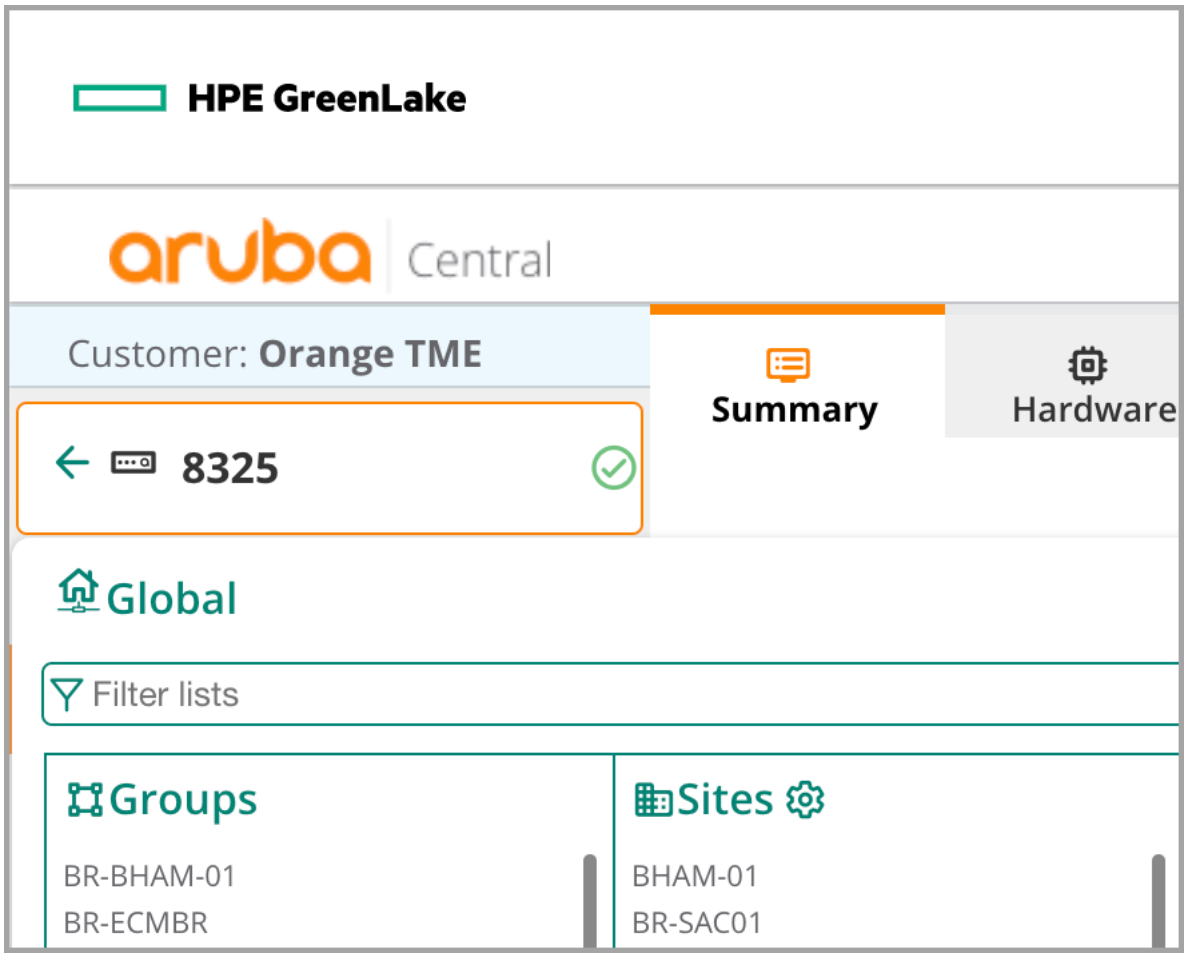


**Step 2** At the bottom left of the **Network Structure** pane, click **New Site**.

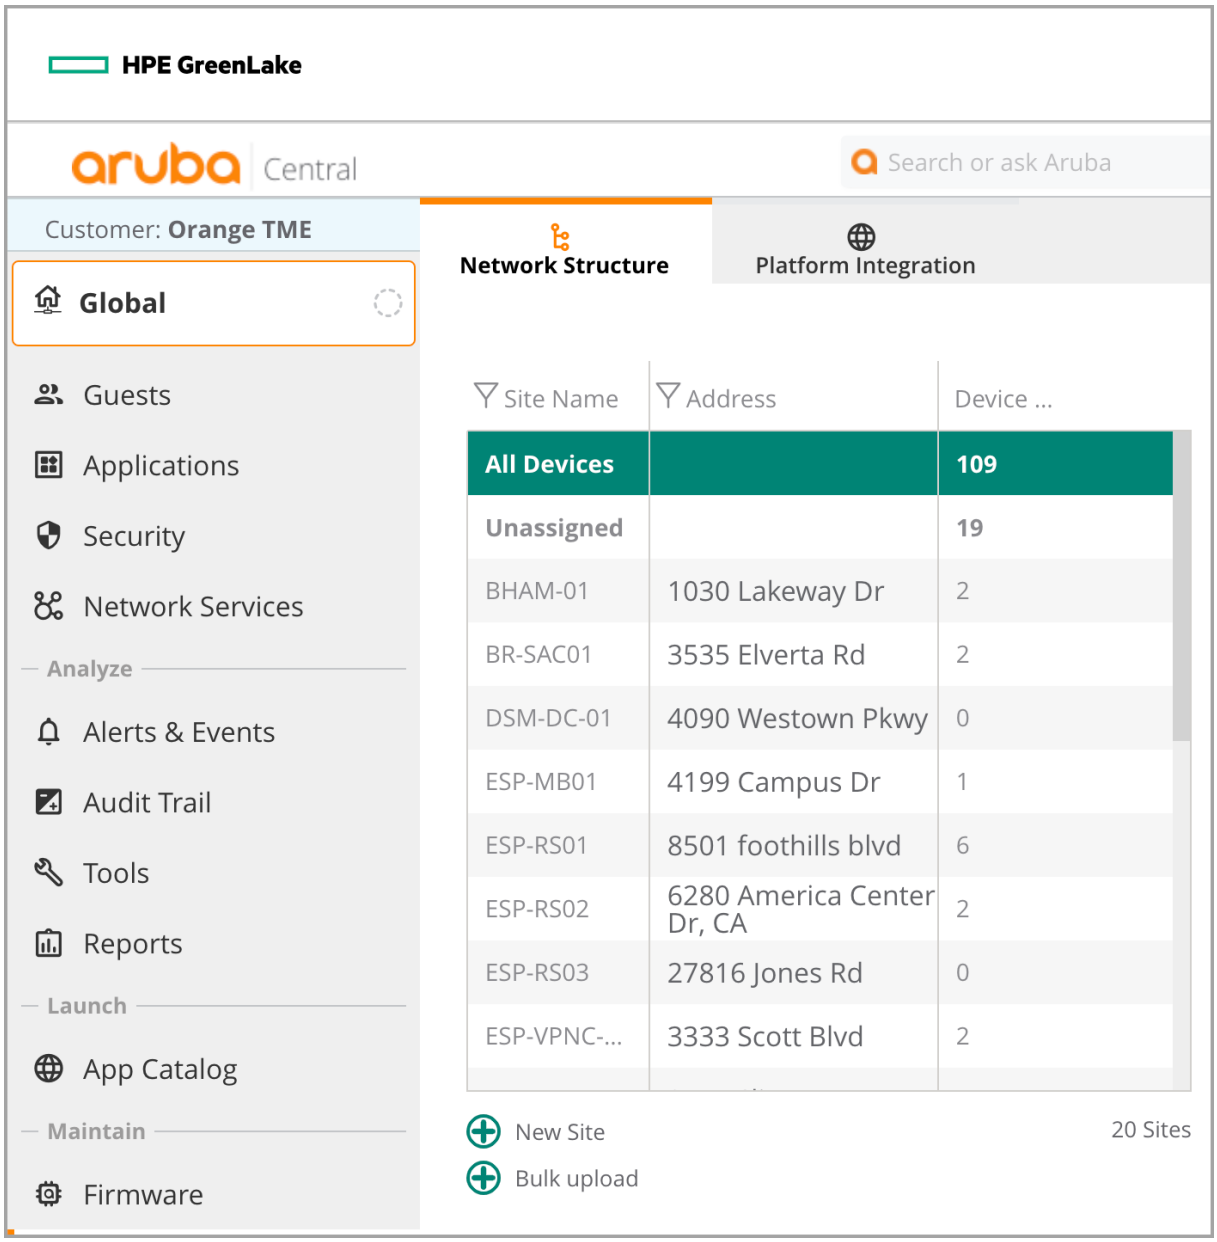


**Step 3** In the **CREATE NEW SITE** window, enter the site location information, then click **Add**.

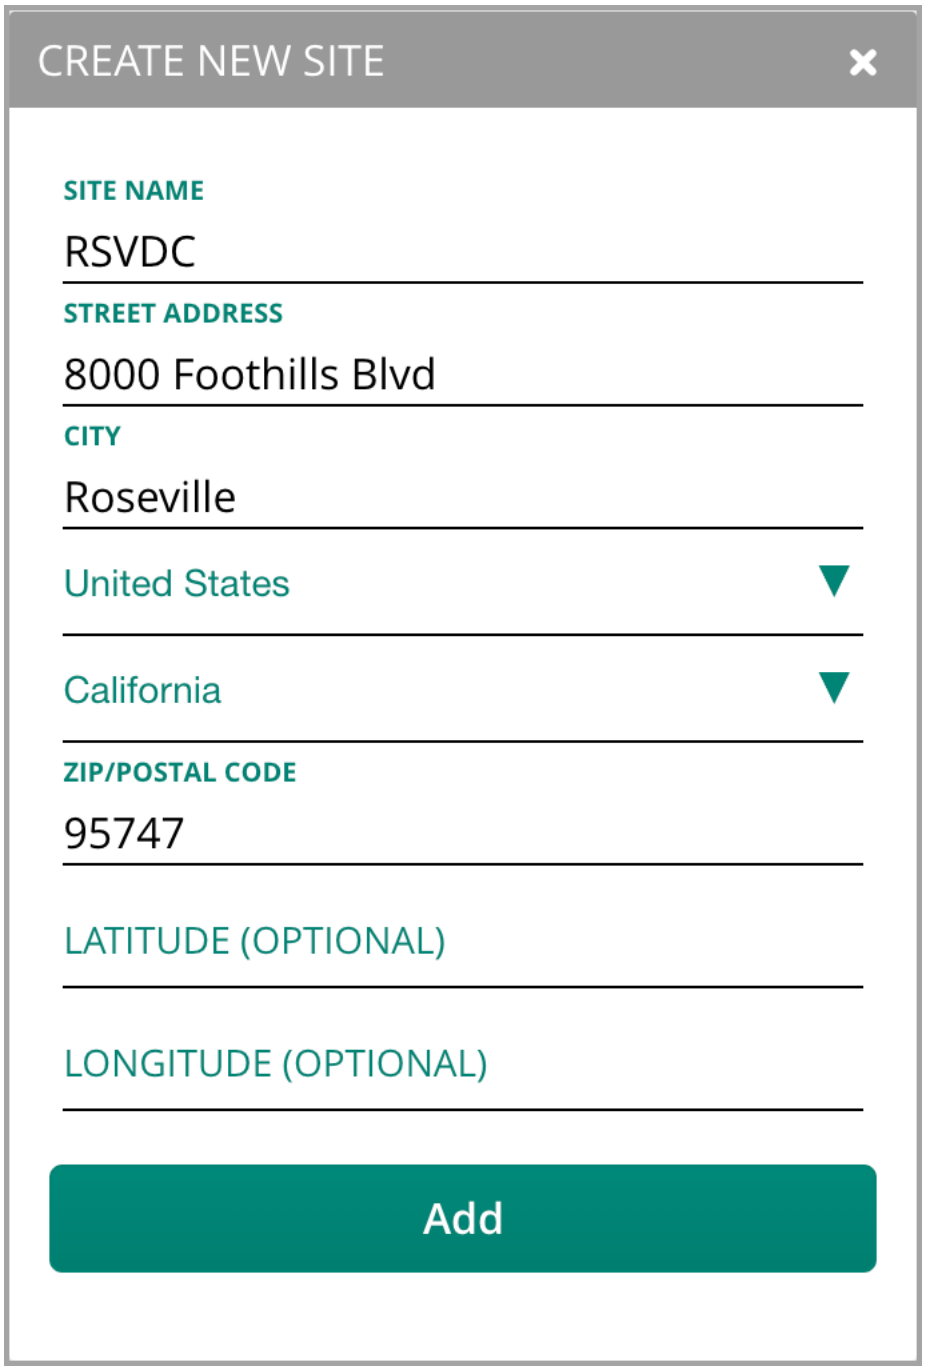


**Step 4** Click the **Site Name** column heading in the left table and enter search criteria to limit the site names displayed.

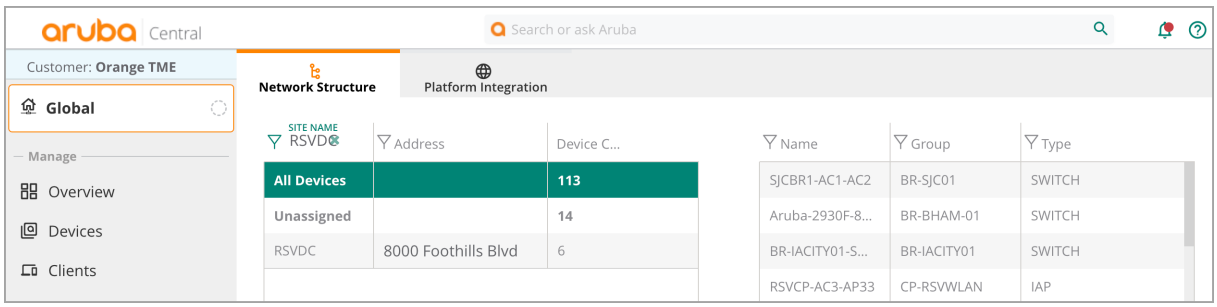

**Step 5** Click the **Name** column heading in the right table and enter search criteria to limit the switch names displayed.

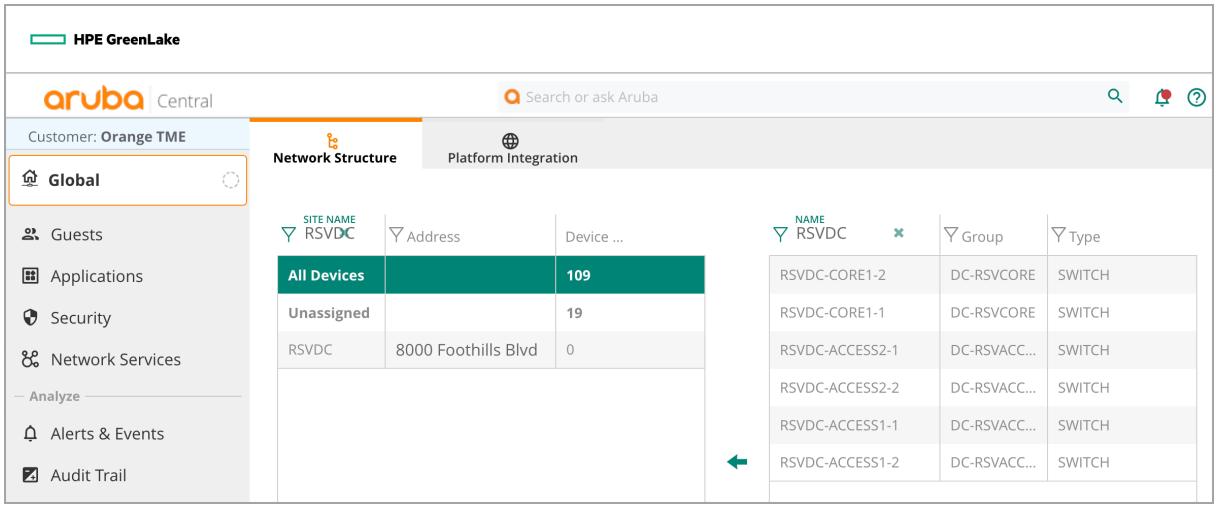


**Step 6** In the device list, select all data center core and access switches.

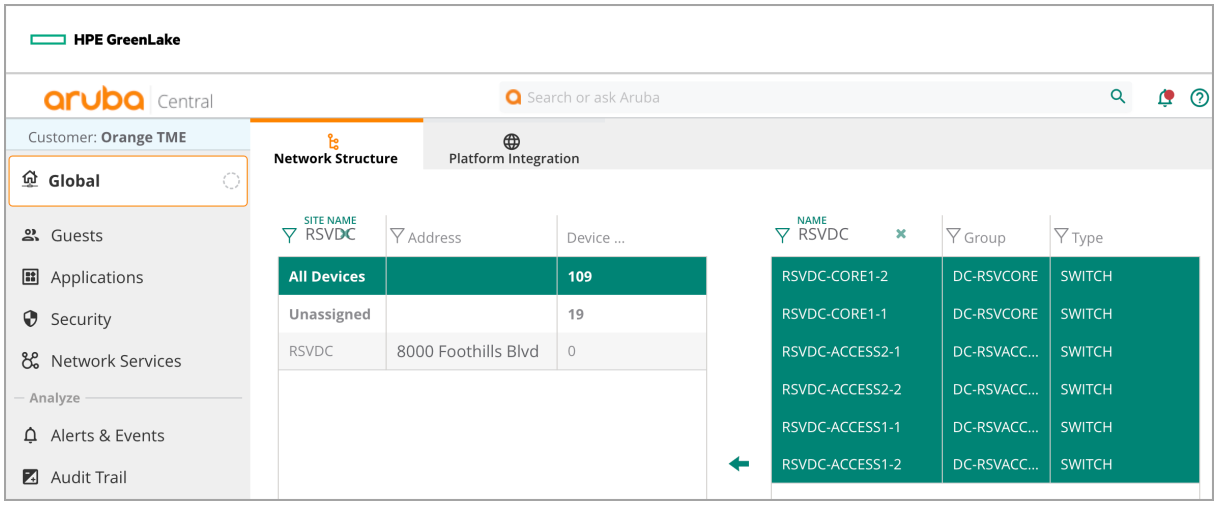


**Step 7** Click and hold on one device name and drag it to the group name on the left. When the group name is highlighted in blue, release the mouse button.

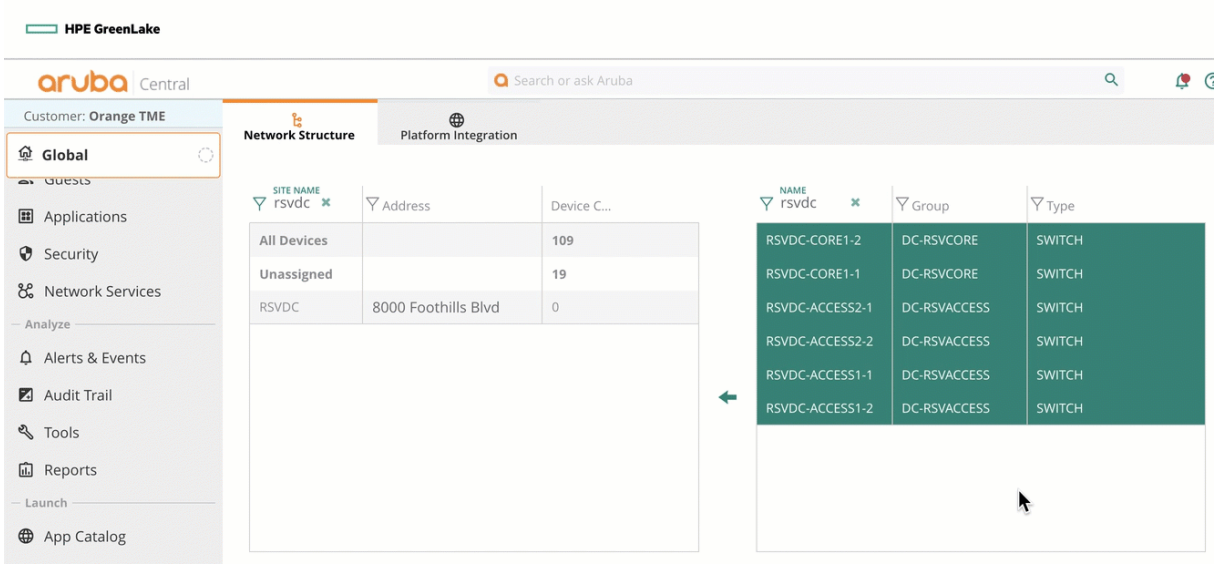

**Step 8** In the **CONFIRMATION ACTION** window, click **Yes**.

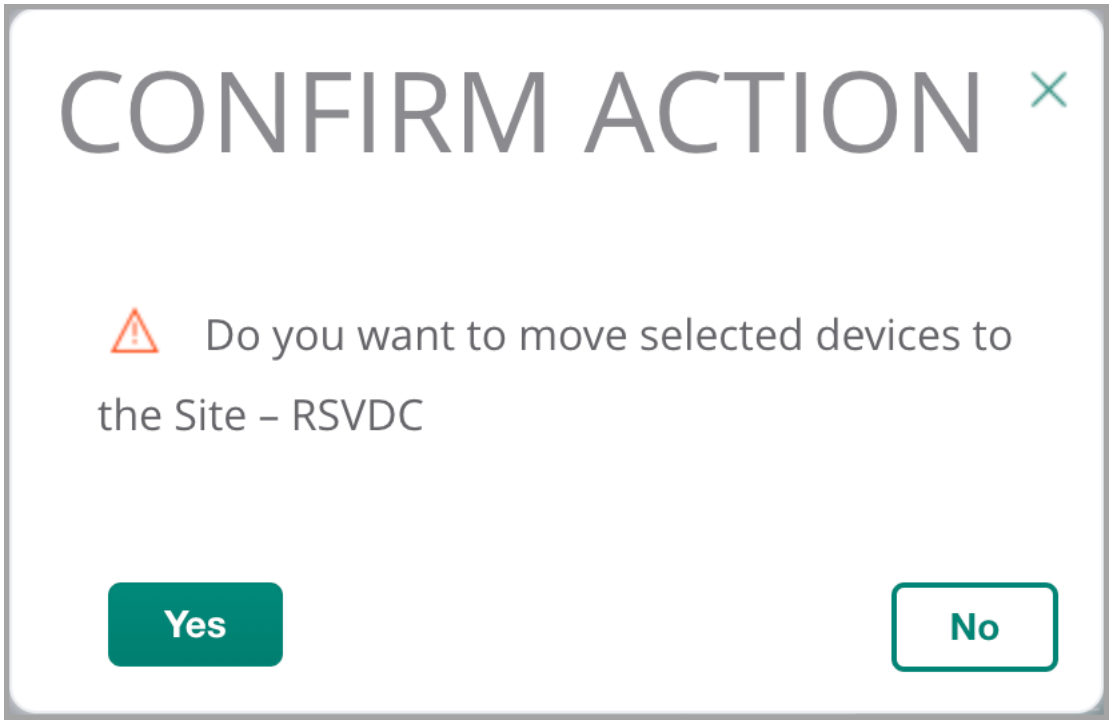


**Step 9** In the site list, click the name of the data center site.

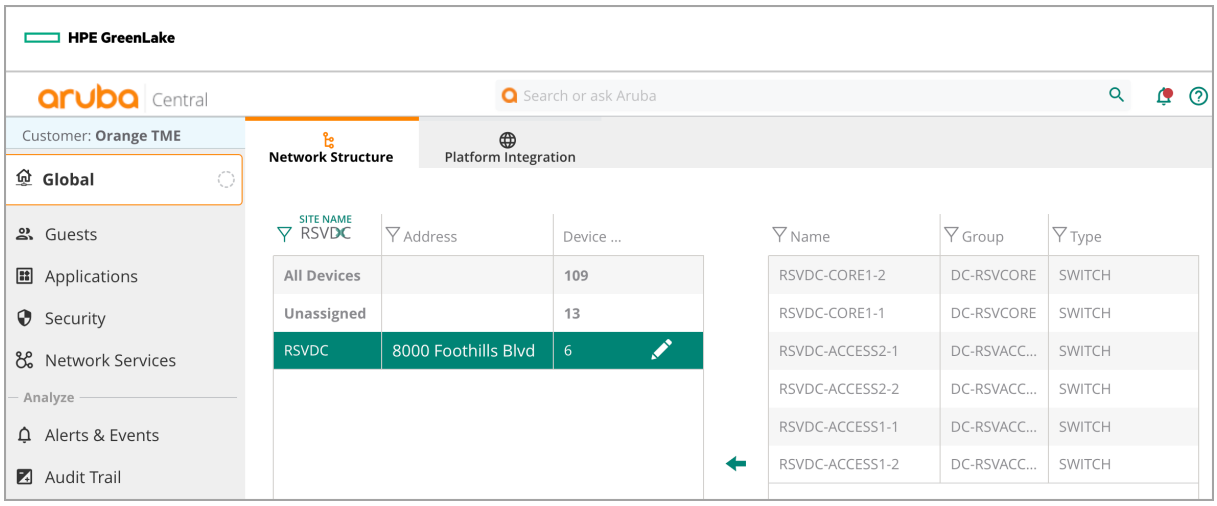


**Step 10** Verify that the complete list of data center switches appears in the device list on the right.

## Verify Data Center Cabling

**Step 1** At the top of the left navigation menu, click **Global**, then click the data center site name in the **Sites** column.

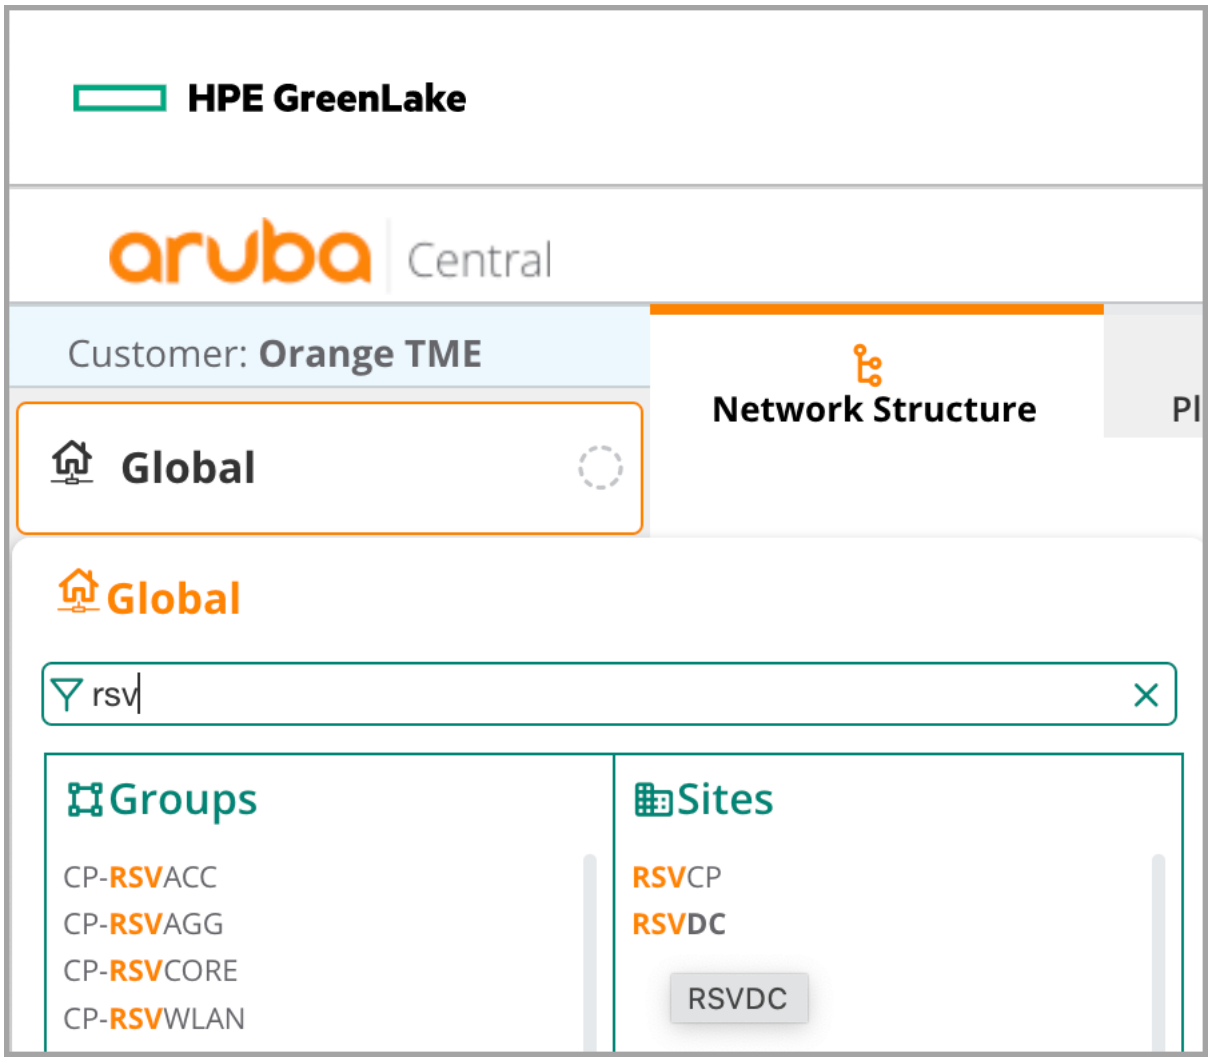


**Step 2** On the left navigation menu, click **Tools**.

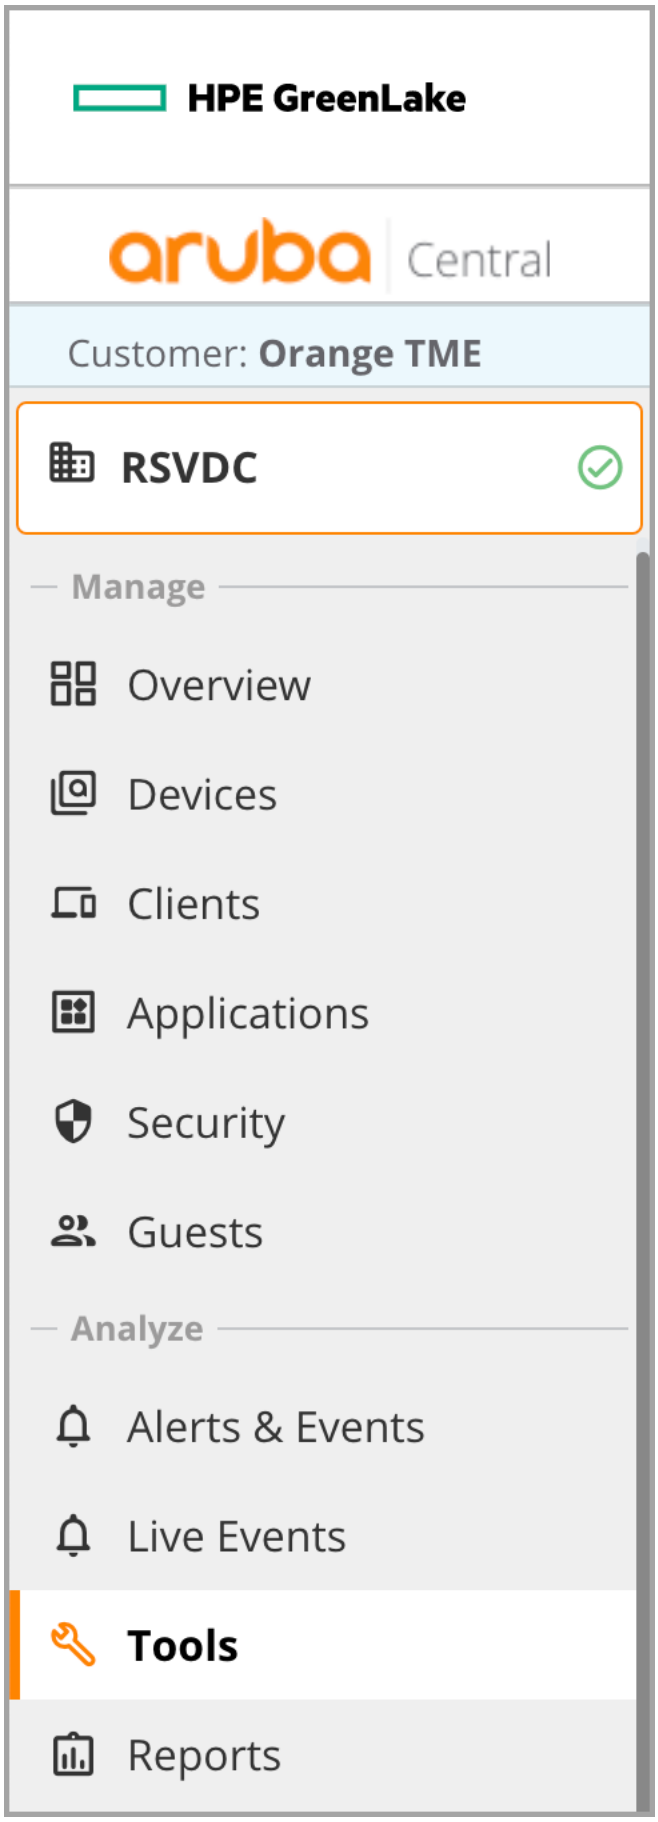
HPE GreenLake
aruba Central
Customer: Orange TME
RSVDC
Manage
Overview
Devices
Clients
Applications
Security
Guests
Analyze
Alerts & Events
Live Events
Tools
Reports

**Step 3** On the **Tools** menu at the top, click the **Commands** tab.

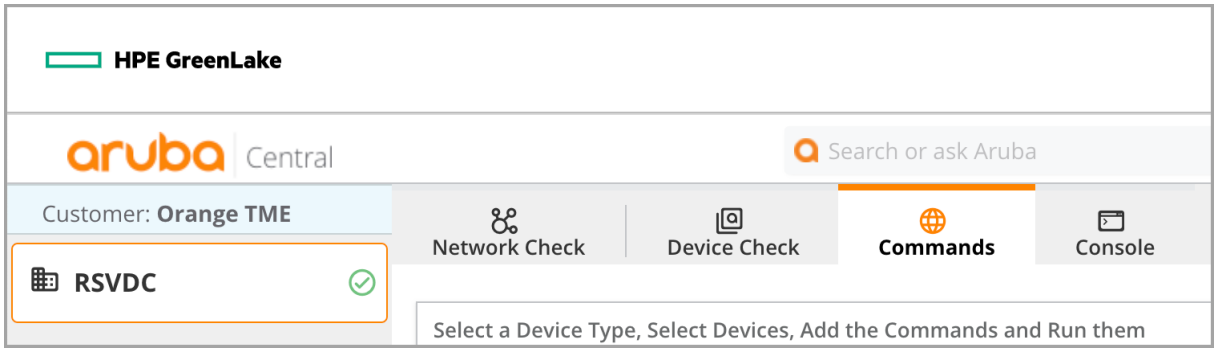


**Step 4** Click the **Available Devices** dropdown, then select all data center switches.

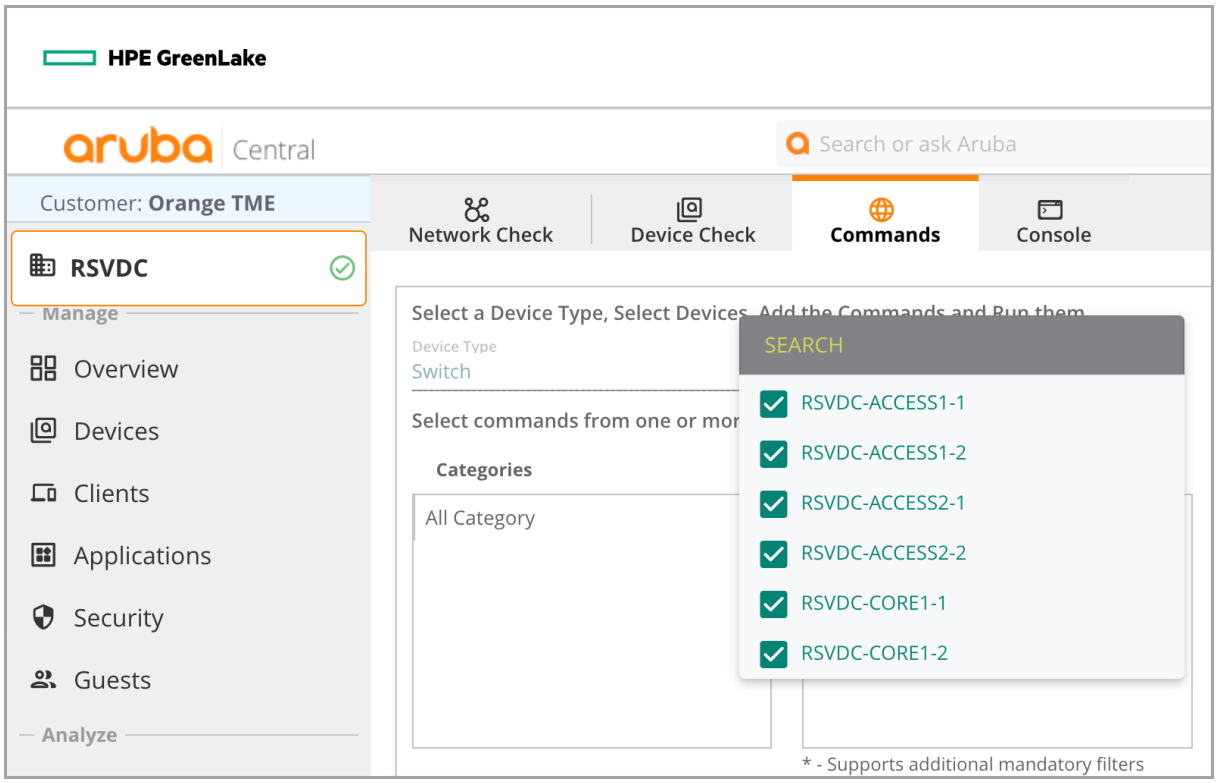


**Step 5** In the **Categories** list, click **All Category**, enter *lldp* in the commands list filter, click **show lldp neighbor**, and click **Add >**.

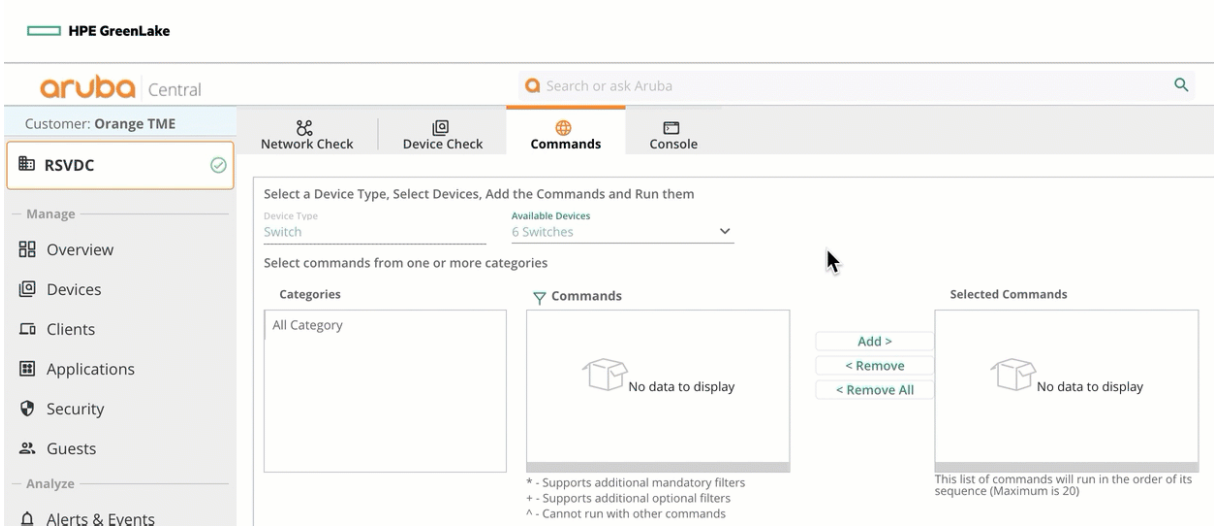


**Step 6** At the lower left of the **Commands** pane, click **RUN**.

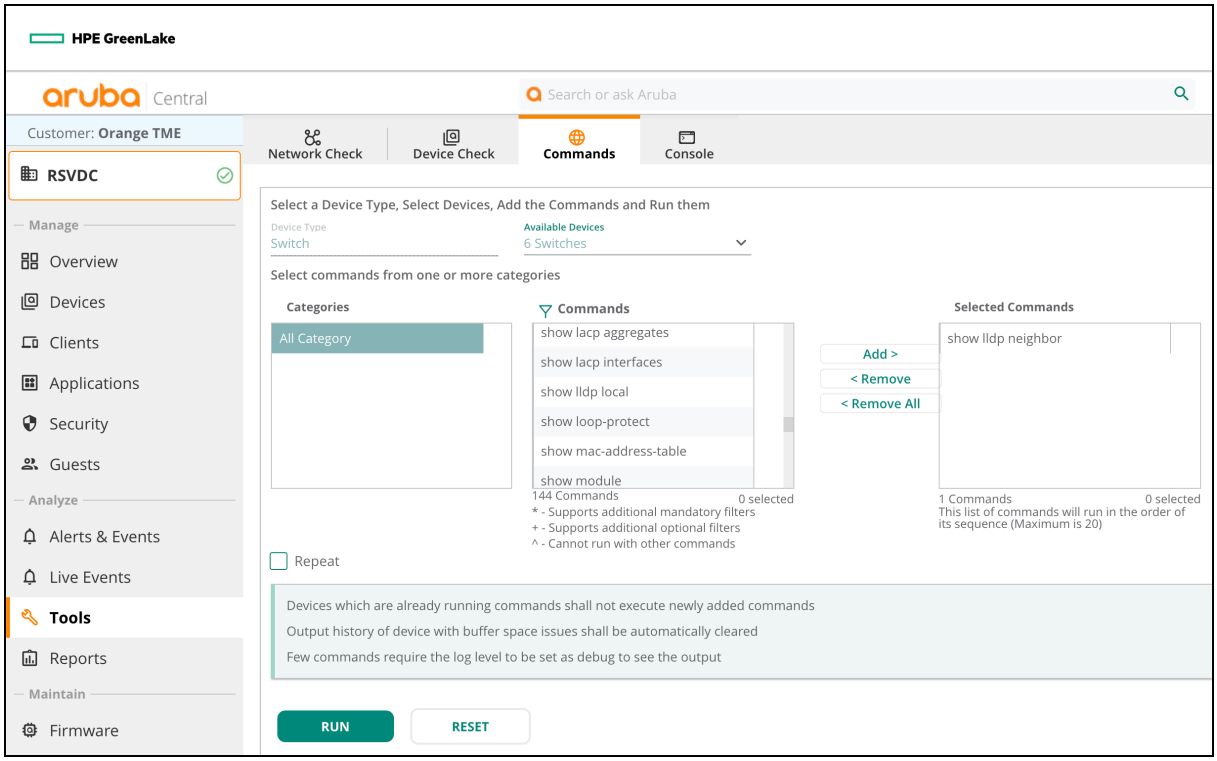


**Step 7** Scroll down to the command output. Review the results for each switch to ensure that LLDP neighbor relationships are consistent with planned cabling.

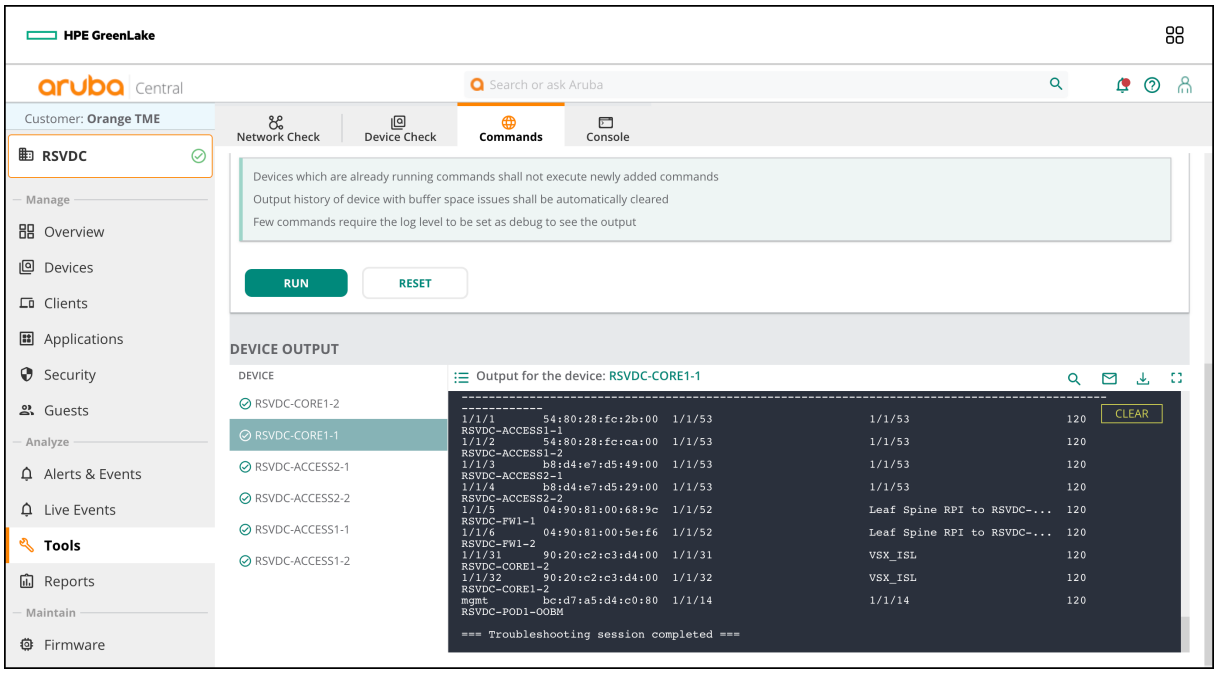


**NOTE:**

To reduce line wrapping, click the three dots/dashes icon to the left of *"Output for the device:"* to toggle **Device** column visibility and expand the the output window. The output also can be expanded to fill the web browser page by clicking the fill screen icon (**[ ]**) at the upper right of the window header.

# Two-Tier Core

Configure Two-Tier core switches as a VSX pair for Layer 2 aggregation of the data center access switches, IP data center services, and routing to the main campus.

## Configure Core VSX ISL Interface

To establish a VSX relationship between the core switches, create a link aggregation (LAG) interface for assignment as the VSX data plane's inter-switch link (ISL). The LAG can be defined at the Central UI group level when using the same ports for the VSX ISL on both core switches.

**Step 1** On the left **Aruba Central** menu, click the current context, then click the data center core switch group name in the **Groups** column.

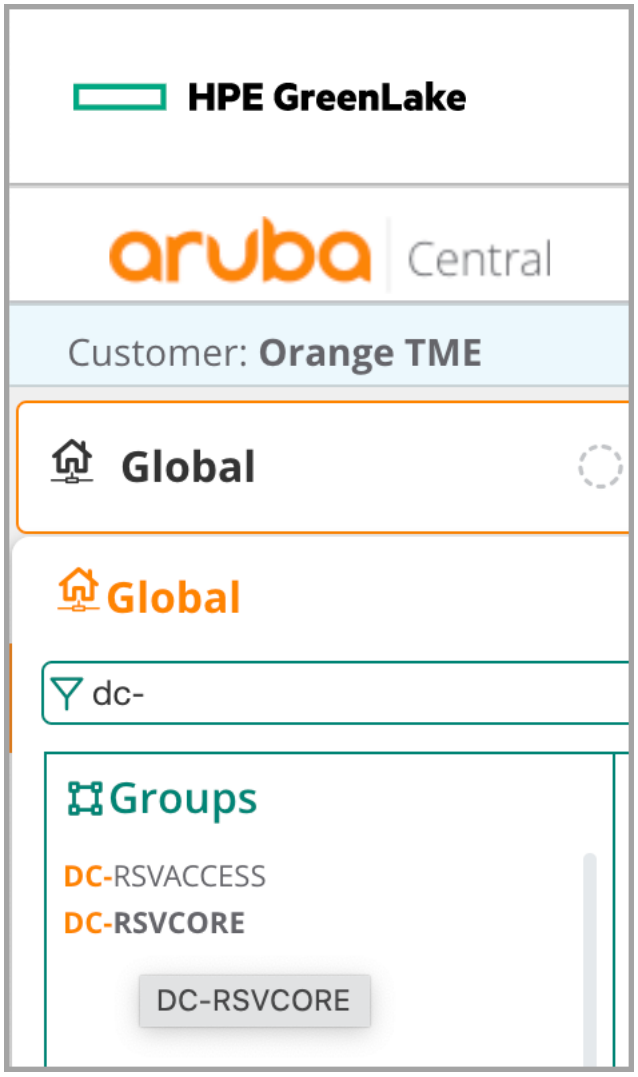


**NOTE:**

The current context in the screenshot above is **Global**.

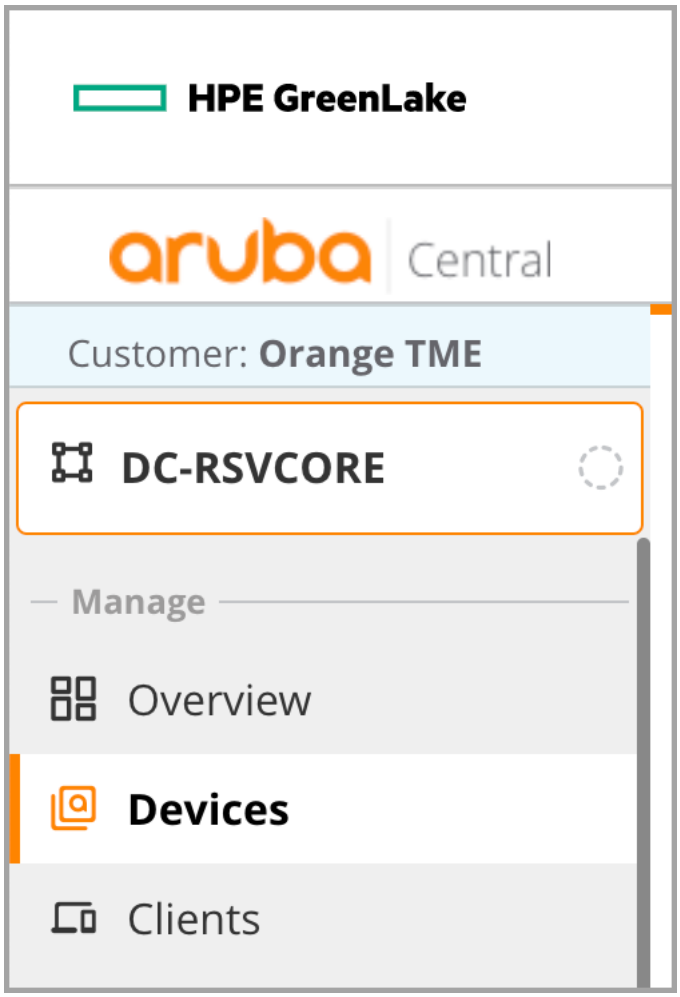


**Step 2** On the left navigation menu, click **Devices**.

**Step 3** At the upper right of the **Switches** pane, click **Config**.

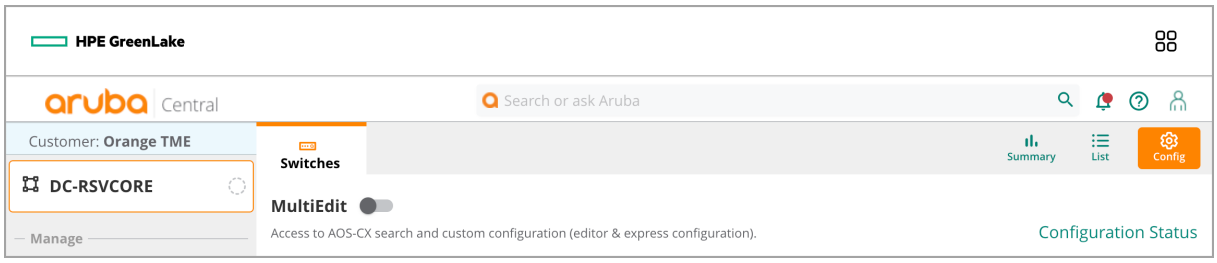


**Step 4** In the **Interfaces** tile, click **Ports & Link Aggregations**

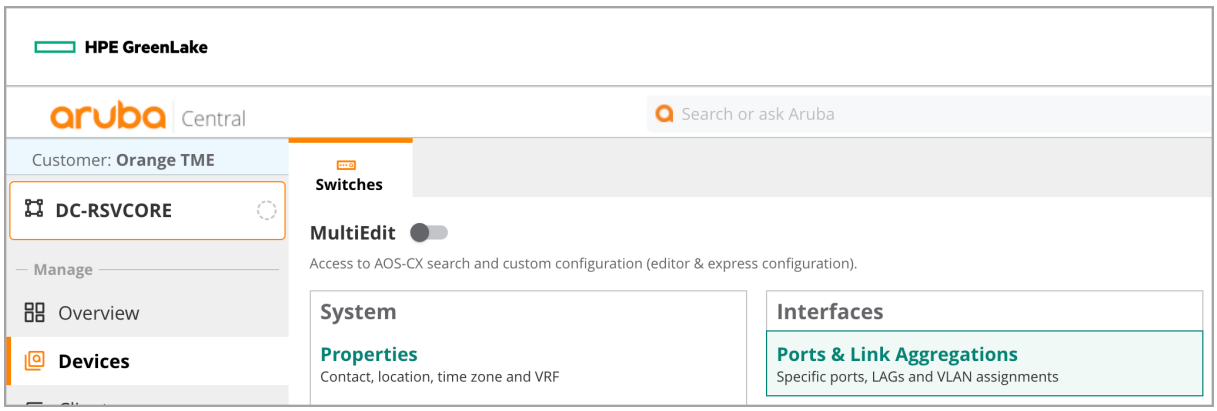

**Step 5** Scroll to the right of the **Ports & Link Aggregations** table, and click the **+** (plus sign) at the upper right.

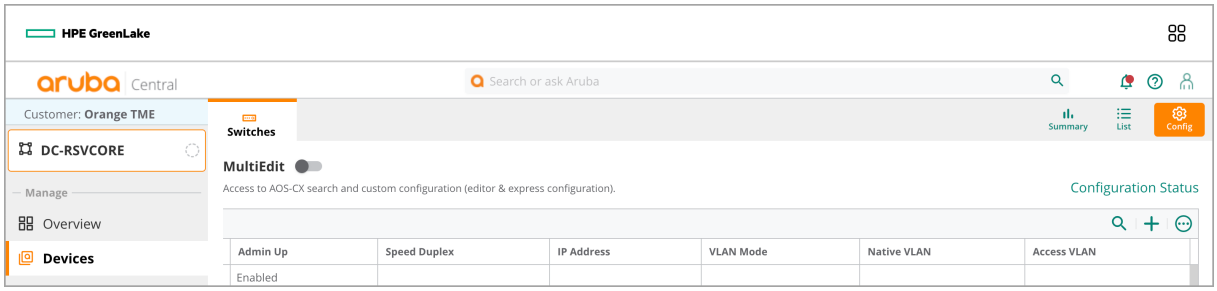


**Step 6** On the **Add LAG** page, enter the following values and click **ADD**:

- **Name:** *lag256*
- **Description:** *VSX-ISL-LAG*
- **Port Members:** *1/1/31, 1/1/32*
- **Speed Duplex:** *<no value>* (default)
- **VLAN Mode:** *trunk*
- **Native VLAN:** *1* (default)
- **Allowed VLANs:** *<no value>* (default)
- **Admin Up:** *checked*
- **Aggregation Mode:** *LACP Active*

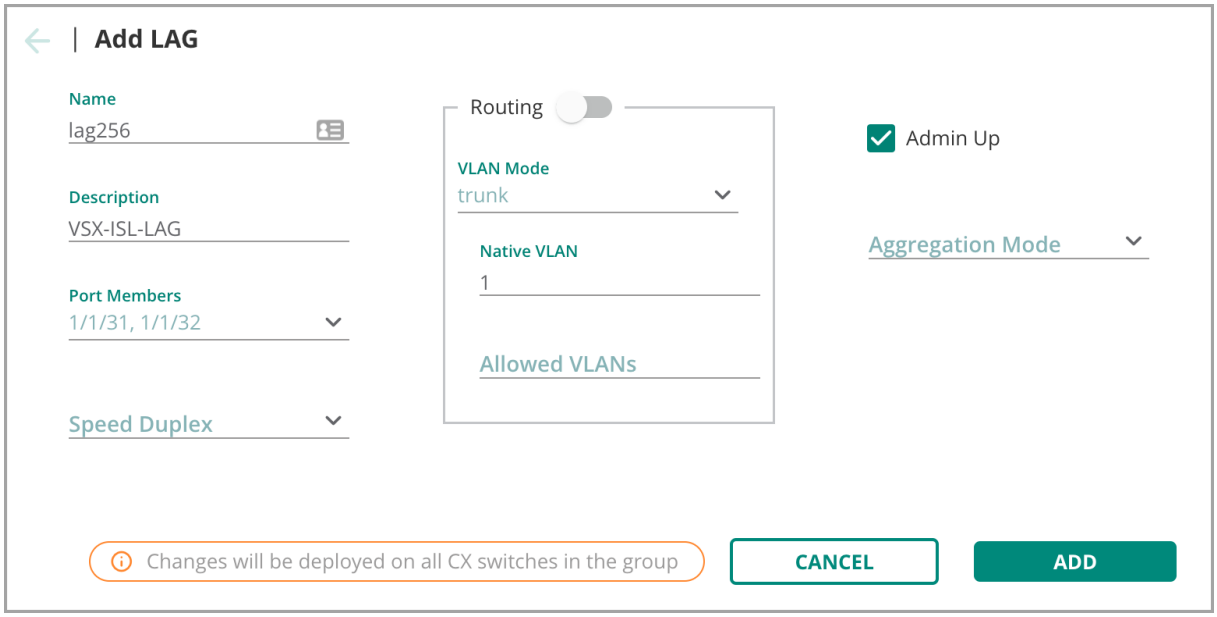


**Step 7** In the **Ports & Link Aggregations** table's title row, click ← (left arrow) to return to the main configuration page.

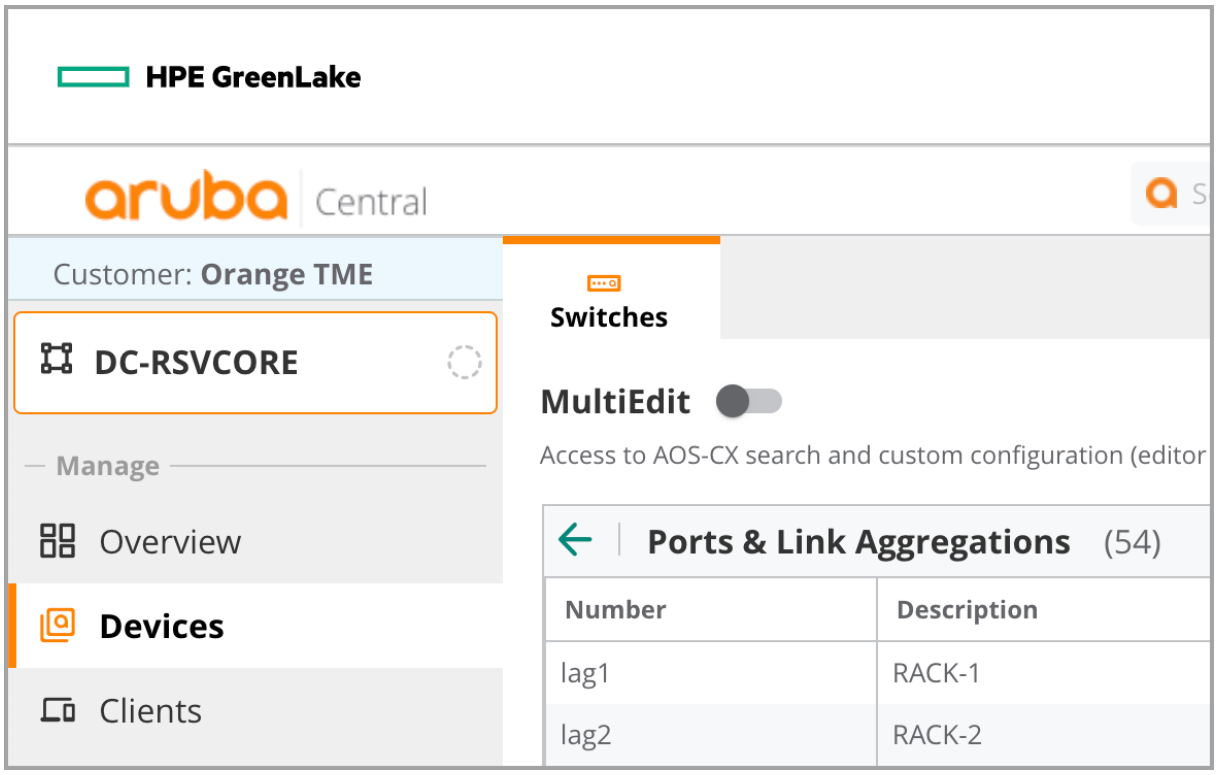


## Configure Routing VLAN

In this topology, a VLAN is created for routing traffic to upstream external networks. The same VLAN is used to establish a routed transit path between the core switches using the VSX ISL. If one of the core switches loses its external network peering, external reachability information is learned from the ISL and external hosts are still reachable from the other core switch.

**NOTE:**

When more than one VRF is present, a VLAN per VRF is created. This sample topology uses only the default VRF, so only one VLAN is created.

**Step 1** In the **Bridging** tile, click **VLANs**.

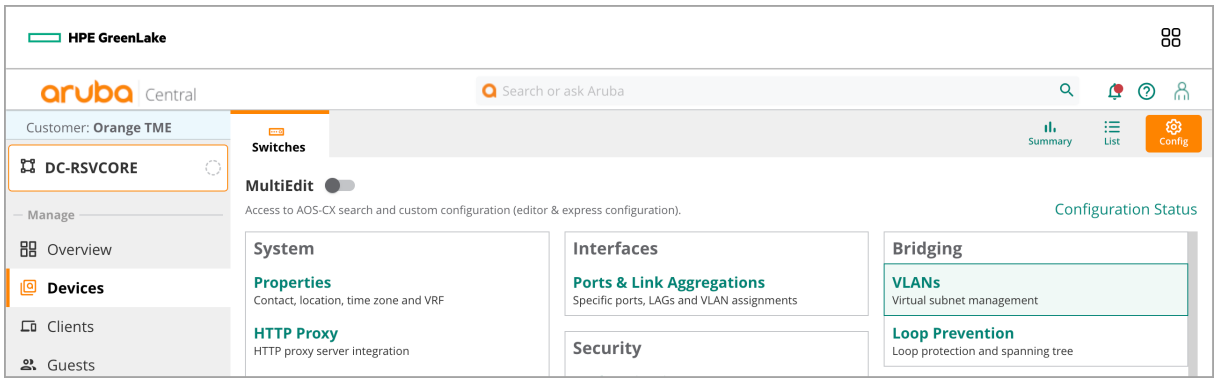


**Step 2** In the upper right of the **VLANs** table, click the **+** (plus sign).

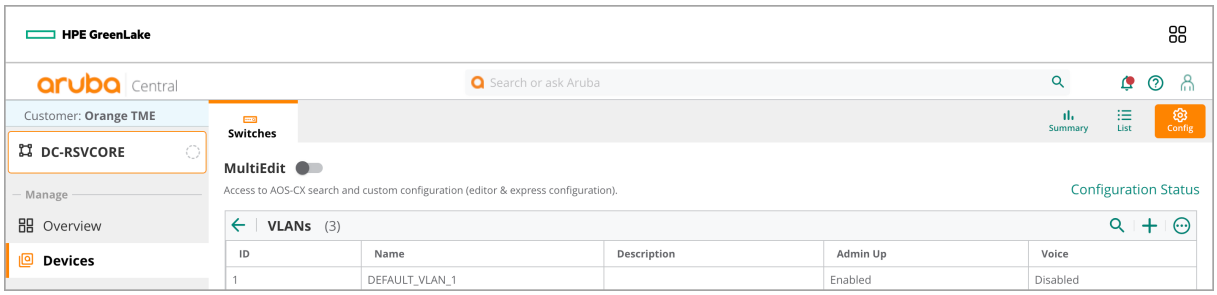


**Step 3** On the **Add VLAN** page, enter the following field values, then click **ADD**.

- **ID:** *4000*
- **Name:** *CORE-ROUTING*
- **Description:** *<no value>* (default)
- **Admin Up:** *checked* (default)
- **Voice:** *unchecked* (default)

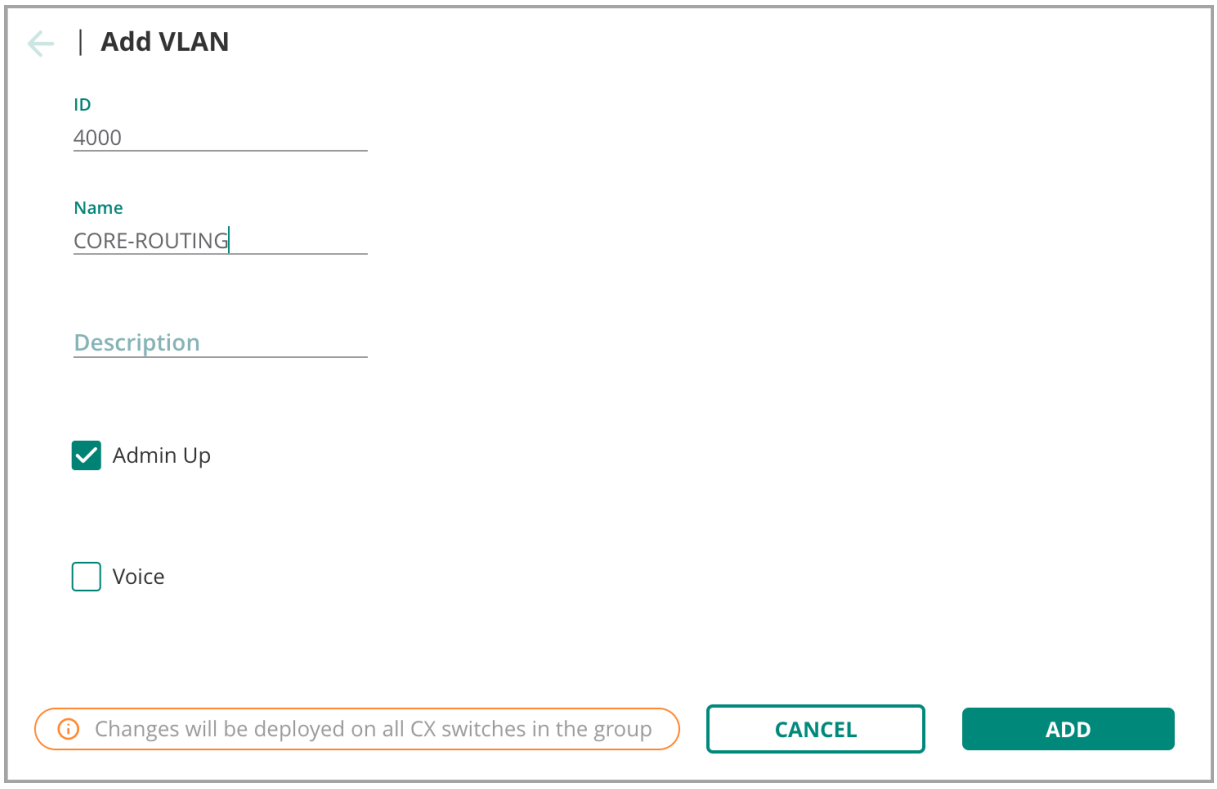


**Step 4** In the **VLANs** table's title row, click ← (left arrow) to return to the main configuration page.

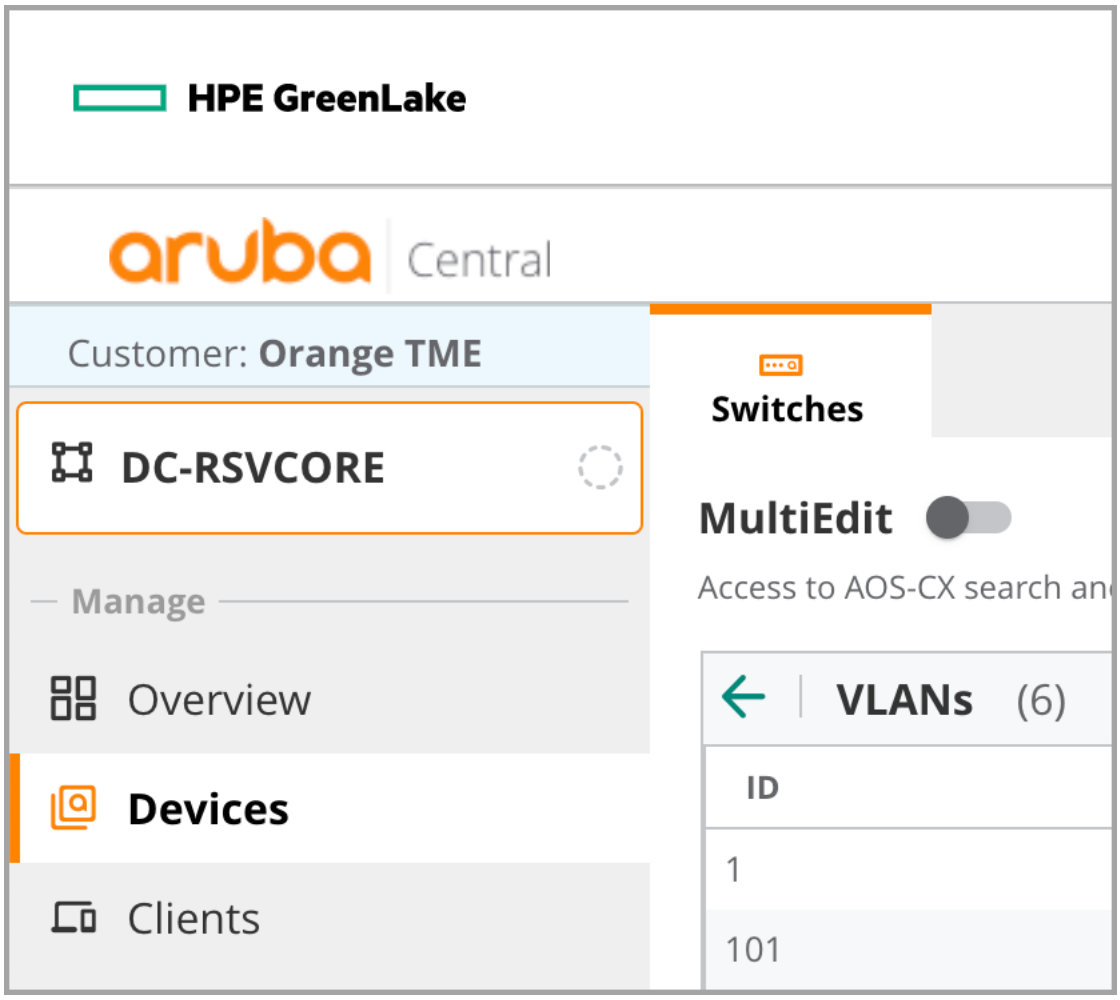


## Spanning Tree

Multi-chassis link aggregations (MC-LAGs) provide the primary loop prevention mechanism in a Two-Tier architecture. When configured on both core and access switches, MC-LAGs allow loop-free forwarding on all inter-switch links simultaneously in both directions, .

MC-LAGs provide efficient, hash-based load balancing with better performance than individually mapped VLANs to Multiple Spanning-Tree (MST) instances.

Spanning-tree (STP) is configured as a backup loop prevention mechanism in case of operator cabling errors when connecting hosts to top of rack switches.

Setting the spanning-tree priority to 0 ensures that the core VSX pair of switches is the STP root.

**Step 1** In the **Bridging** tile, click **Loop Prevention**.

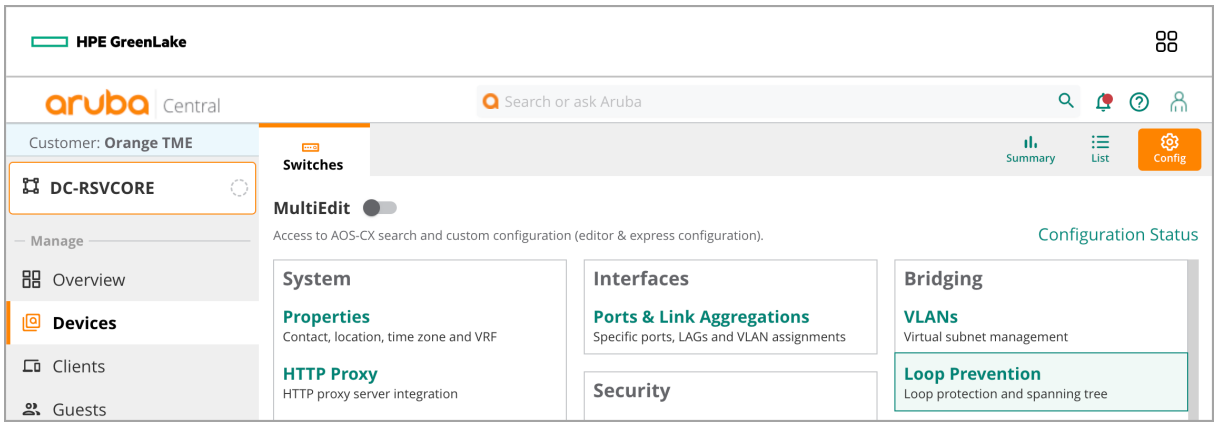


**Step 2** In the **Loop Prevention** window, set the following **Spanning Tree** values, then click **SAVE**. - **Priority:** *0* - **Region:** *RSVDC*

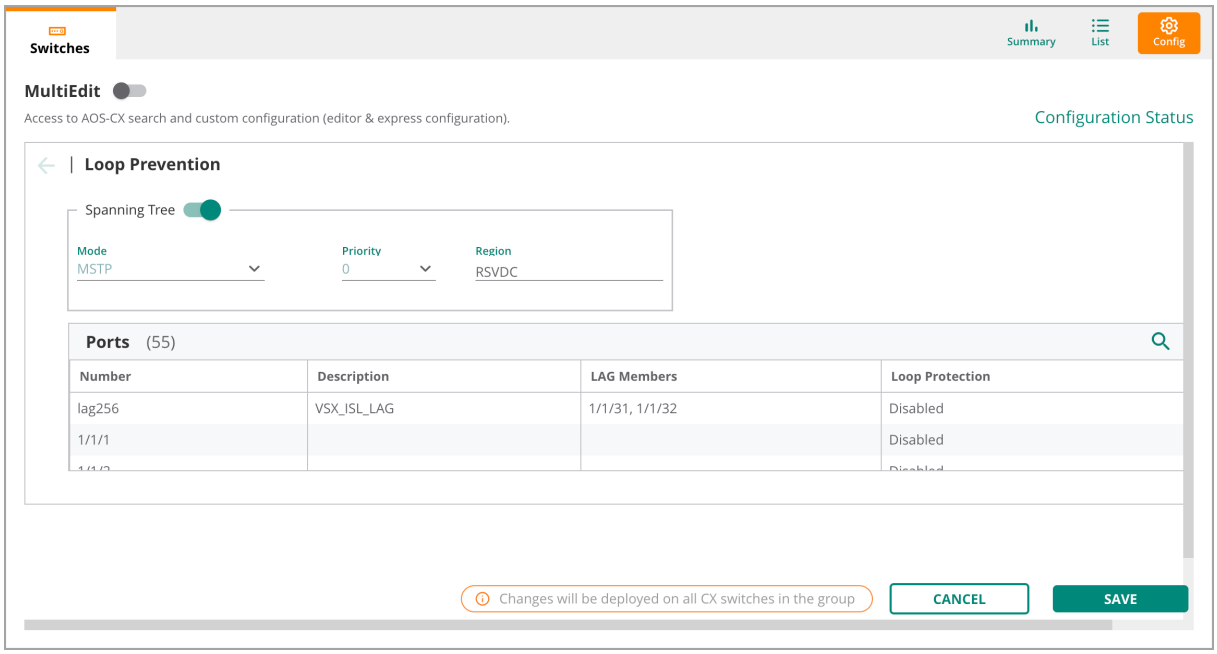


# Enter MultiEdit Configuration

The Central UI interface provides simplified access to most common switch configuration features. MultiEdit is a tool in the Central UI for CX switches that enables configuration of any CX feature using CLI syntax. MultiEdit provides syntax checking, colorization, and command completion.

For complete details on using MultiEdit, refer to the Editing Configuration on AOS-CX section of Central online help.

The text configuration snippets in the following steps are intended for copying and pasting into MultiEdit. To prevent potential copy/paste errors, scroll to the bottom of the configuration, create a new line, then paste the new configuration lines.  MultiEdit automatically positions new lines in the correct configuration context.

**Step 1** At the upper left of the **Switches** pane, click the **MultiEdit** enable slider.

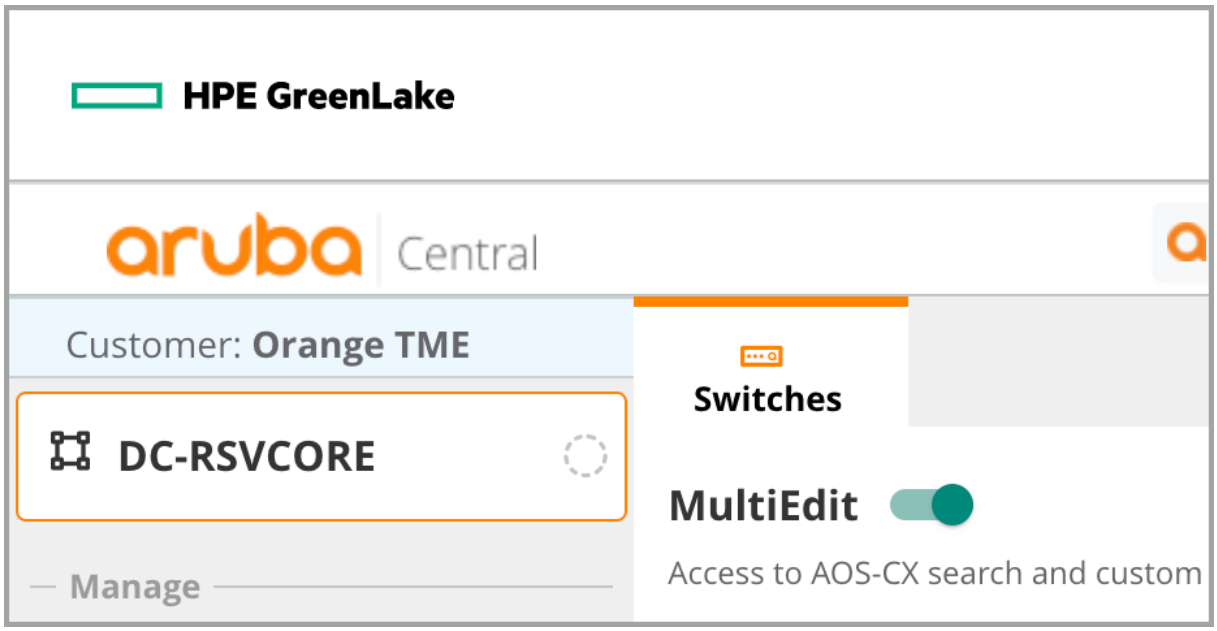


**Step 2** Click both core switches in the **Devices** lists, then click **EDIT CONFIG**.

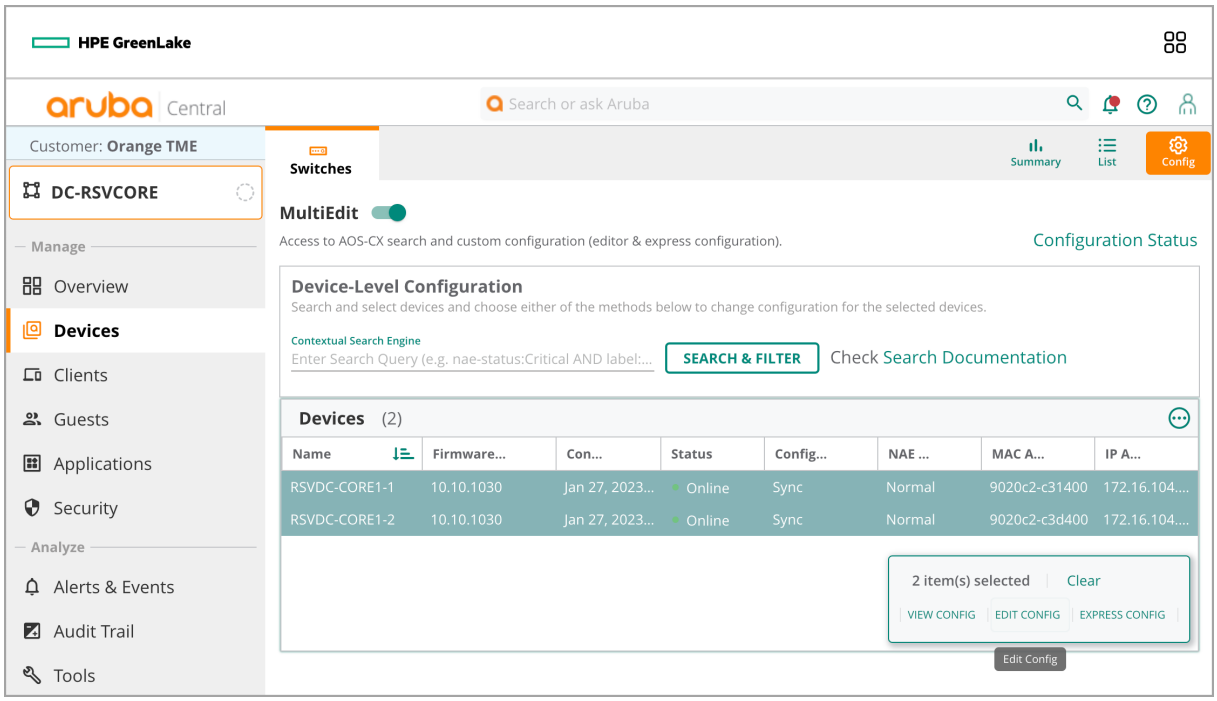


**NOTE:**

When using the Central MultiEditor, it is beneficial to save small sets of configuration at a time. This reduces the volume of configuration that must be inspected when errors occur and makes troubleshooting configuration elements faster.

# Configure Core Switch VSX

The core switches are configured as a VSX pair to support Layer 2 multi-chassis link aggregation (MC-LAG) to the access layer switches. The previously defined LAG is assigned as the VSX data path inter-switch link (ISL). The out-of-band (OOB) *mgmt* interface is used for VSX keepalives to maximize the number of ports available to connect access switches.

**Step 1** Enter the initial VSX configuration.

```
vsx
    system-mac 02:00:00:00:10:00
    inter-switch-link lag 256
    role primary
    keepalive peer 172.16.117.102 source 172.16.117.101 vrf mgmt
```

**NOTE:**

When the *mgmt* vrf is specified, the keepalive peer addresses are the IPs assigned to the out-of-band management interfaces. When using DHCP IP address assignments on the OOB management network, DHCP reservations must be created for VSX paired switches to avoid future keepalive failures.

**Step 2** Mouse-over the *role* value of **primary** to display the values for each individual switch, then right-click.

**Step 3** In the **Modify Parameters** window, click **primary** under *RSVDC-CORE1-2*, select **secondary** from the menu, then click **SAVE CHANGES**.

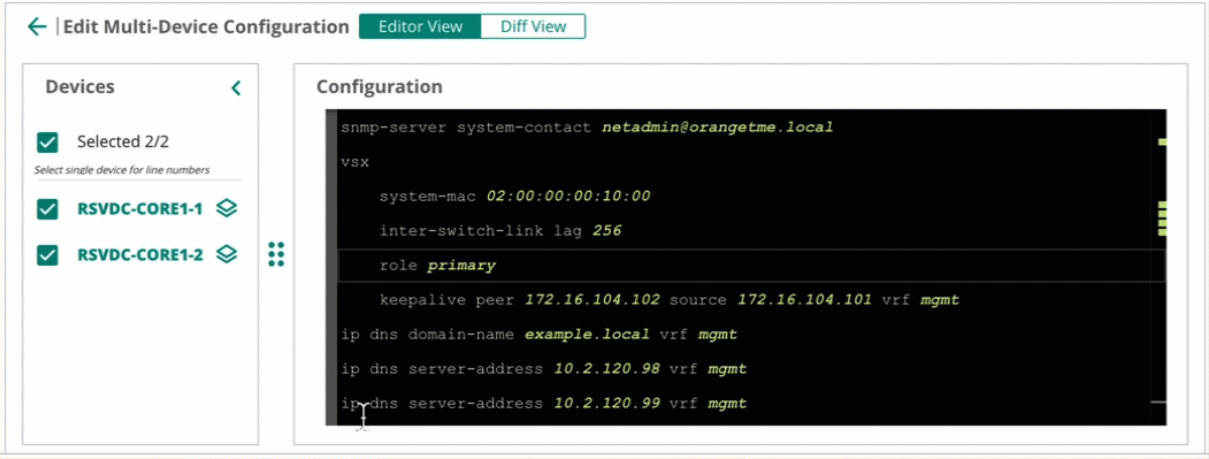


**NOTE:**

Hover the mouse over the per-switch values to display a switch's assigned value.

**Step 4** Modify the VSX keepalive **peer** and **source** parameters by right-clicking on the values.

| Switch | peer | source |
|---|---|---|
| RSVDC-CORE1-1 | 172.16.117.102 | 172.16.117.101 |

| Switch | peer | source |
|---|---|---|
| RSVDC-CORE1-2 | 172.16.117.101 | 172.16.117.102 |

**Step 5** Assign a description and maximum MTU value to the VSX ISL physical interfaces.

```
interface 1/1/31
    description VSX-ISL
    mtu 9198
interface 1/1/32
    description VSX-ISL
    mtu 9198
```

## Configure Core Switch MC-LAGs

**Step 1** Create MC-LAG interfaces for connecting to redundant top-of-rack access switches and upstream firewalls.

```
interface lag 1 multi-chassis
    description RACK-1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
    lacp fallback
    spanning-tree root-guard
interface lag 2 multi-chassis
    description RACK-2
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
    lacp fallback
    spanning-tree root-guard
interface lag 101 multi-chassis
    description EXT-FW1-1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 4000
    lacp mode active
    lacp fallback
    spanning-tree root-guard
interface lag 102 multi-chassis
    description EXT-FW1-2
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 4000
    lacp mode active
    lacp fallback
    spanning-tree root-guard
```

**NOTE:**

MC-LAG interfaces can scope trunked VLANs only to those required for a specific downstream rack. Tagging all VLANs on all core-to-access MC-LAGs supports ubiquitous host mobility across all racks within the Two-Tier structure and reduces the administrative overhead of maintaining VLAN assignments per rack.

**Step 2** Assign physical interfaces to the MC-LAGs..

```
interface 1/1/1
    description RSVDC-ACCESS1-1
    no shutdown
    mtu 9198
    lag 1
interface 1/1/2
    description RSVDC-ACCESS1-2
    no shutdown
    mtu 9198
    lag 1
interface 1/1/3
    description RSVDC-ACCESS2-1
    no shutdown
    mtu 9198
    lag 2
interface 1/1/4
    description RSVDC-ACCESS2-2
    no shutdown
    mtu 9198
    lag 2
interface 1/1/29
    description EXT-FW1-1
    no shutdown
    mtu 9198
    lag 101
interface 1/1/30
    description EXT-FW1-2
    no shutdown
    mtu 9198
    lag 102
```

**Step 3** Remove the following configuration line from interfaces configured above: **vlan access 1**

**NOTE:**

Interface LAG assignments and VLAN access statements cannot be assigned to an interface simultaneously. An error occurs when saving the MultiEdit configuration if the **vlan access 1** statement is not removed.

## Configure Routing Services

In a Two-Tier architecture, the core switches provide IP gateways to downstream hosts, route traffic between data center hosts in different subnets, and route traffic between the data center and external networks.

## Configure Routing Components

This sample deployment connects to a campus network via an active/passive pair of upstream firewalls. It is common to place firewalls between the data center and a campus network for granular policy enforcement. When using an active/passive redundant firewall pair, IP assignments are used only by the currently active firewall. Access control lists (ACLs) also can be used to augment firewall policy or in place of a firewall, when policy complexity does not require a dedicated appliance.

OSPF is the most common protocol used between Two-Tier data center core switches and external networks. Alternatively, BGP can be configured.

In the diagram below, OSPF adjacencies are formed between the two core switches and between each core switch and RSVDC-FW1-1, which is the active member of the firewall cluster pair. Under normal operating conditions, RSVDC-FW1-2 does not participate in routing or traffic forwarding. When RSVDC-FW1-2 becomes the active member of the firewall cluster, the OSPF sessions are moved and traffic is forwarded to RSVDC-FW1-2.

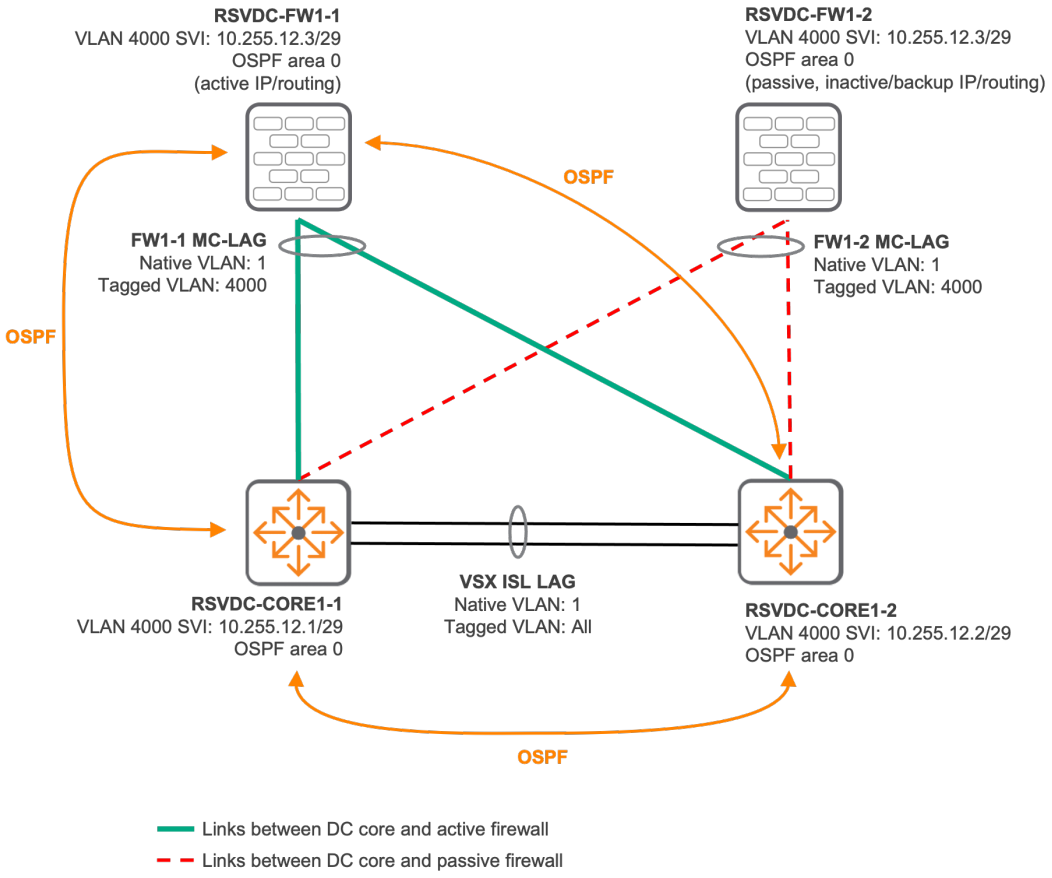


**NOTE:**

When more than one VRF is present, each VRF maintains its own set of OSPF peerings over a unique VLAN on the same physical links.

VSX combines two separate switches into a single, logical Layer 2 switch. Layer 3 functions remain independent. If an external link fails between one of the core switches and the upstream firewall, the same VLAN configured for external network reachability provides a routed transit path over the VSX ISL. An OSPF adjacency is configured between the core switches, which shares external network reachability between them.

**Step 1** Set the switch profile.

```
profile l3-agg
```

**NOTE:**

The available profile options are platform dependent. Selecting a profile optimizes switch hardware resources for its role in the network. It is recommended to assign the *l3-agg* profile to CX 8325 and CX 10000 core switches. CX 8360 switches should use their default *aggregation-leaf* profile. CX 9300 switches should use their default *leaf* profile.

**Step 2** Create the OSPF process.

```
router ospf 1
    router-id 10.250.12.1
    passive-interface default
    area 0.0.0.0
```

**Step 3** Mouse-over the OSPF router ID values *10.250.12.1*, right-click to set per switch values, set the **router-id** of RSVDC-CORE-1-2 to *10.250.12.2*, and click **SAVE CHANGES**.

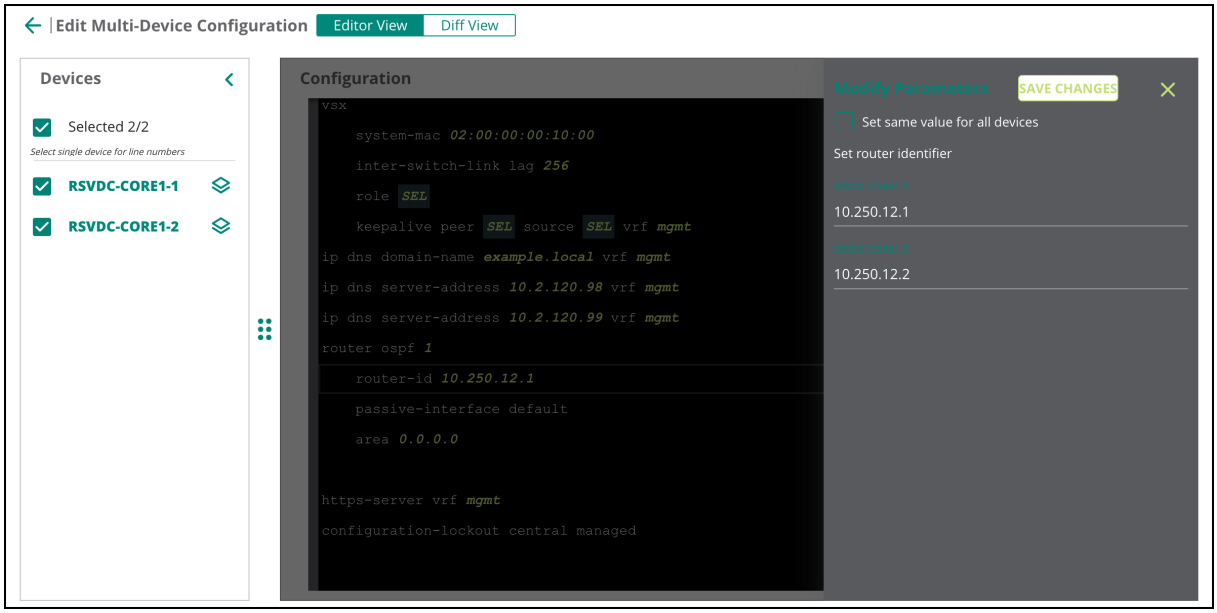


**Step 4** Create core switch loopback interfaces. The loopback IP should be the same value assigned to the OSPF router-id.

```
interface loopback 0
    ip address 10.250.12.1/32
    ip ospf 1 area 0.0.0.0
```

**Step 5** Mouse-over the loopback 0 **ip address** value of *10.250.12.1*, right-click to set per switch values, set the **ip address** of RSVDC-CORE1-2 to *10.250.12.2/32*, and click **SAVE CHANGES**.

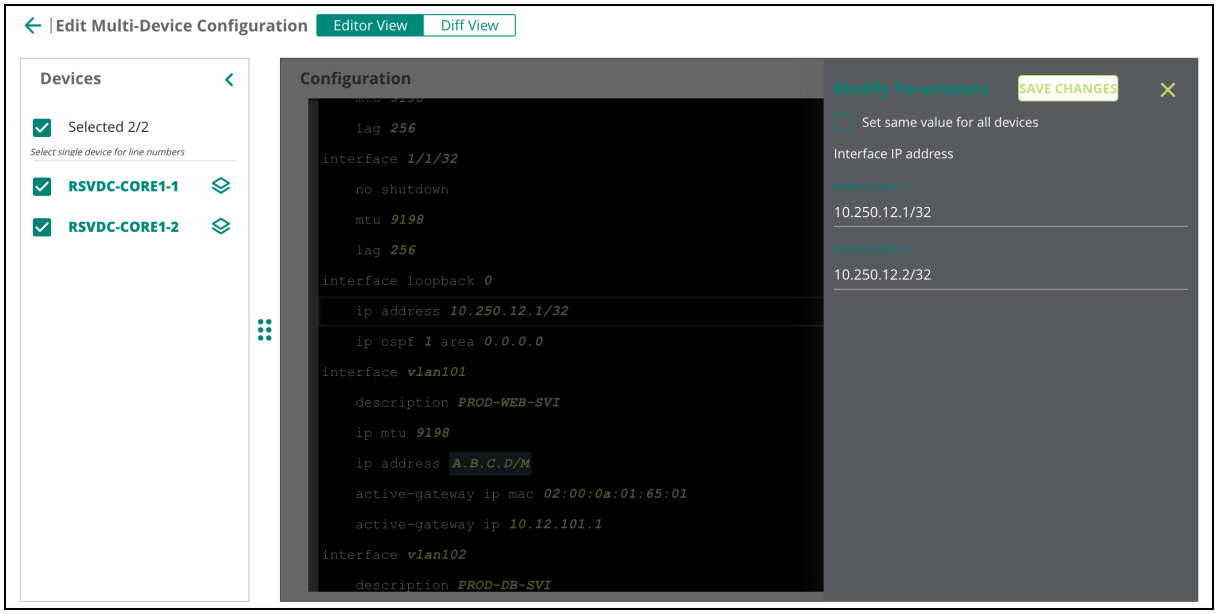


**Step 6** Configure the external/transit VLAN SVI.

```
interface vlan4000
    description CORE-ROUTING-SVI
    ip mtu 9198
    ip address 10.255.12.1/29
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
```

**Step 7** Mouse-over the transit VLAN **ip address** value of *10.255.12.1*, right-click to set per switch values, set the **ip address** of RSVDC-CORE1-2 to *10.255.12.2/29*, and click **SAVE CHANGES**.

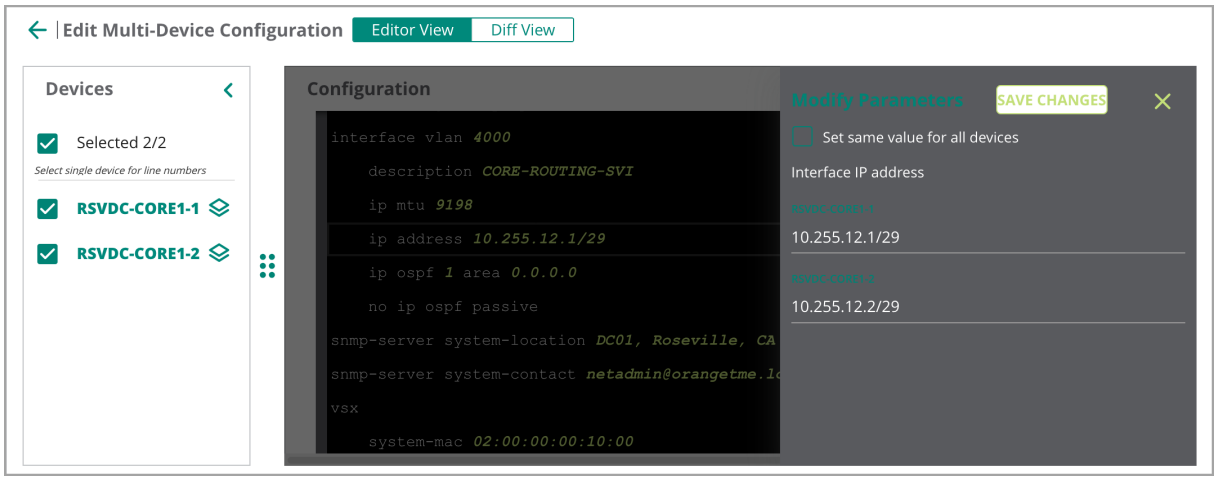


## Configure Host VLAN SVIs

**Step 1** Configure VLAN switched virtual interfaces (SVIs) for data center host VLANs. Core switches provide the default gateway to downstream data center hosts. An active gateway IP and MAC address are configured for each VLAN to allow both core switches to represent the same IP gateway.

```
interface vlan101
    description PROD-WEB-SVI
    ip mtu 9198
    ip address 10.12.101.2/24
    active-gateway ip mac 02:00:0a:01:65:01
    active-gateway ip 10.12.101.1
    ip ospf 1 area 0.0.0.0
interface vlan102
    description PROD-DB-SVI
    ip mtu 9198
    ip address 10.12.102.2/24
    active-gateway ip mac 02:00:0a:01:65:01
    active-gateway ip 10.12.102.1
    ip ospf 1 area 0.0.0.0
```

**NOTE:**

The sample active gateway MAC address associated with the virtual IP sets the locally administered bit to “1” and embeds a hexadecimal representation of the active gateway IP in the last four octets.

**Step 2** Mouse-over the VLAN 101 **ip address** value of *10.12.101.2*, right-click to set per switch values, set the **ip address** of RSVDC-CORE1-2 to *10.12.101.3*, and click **SAVE CHANGES**.

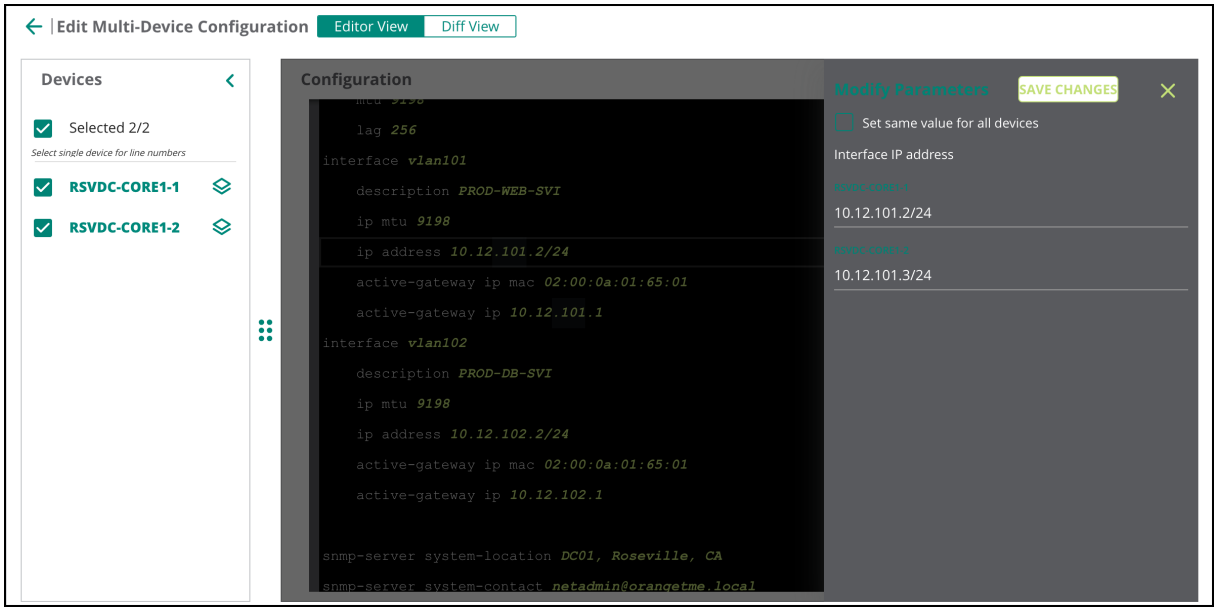


**Step 3** Mouse-over the VLAN 102 **ip address** value of *10.12.102.2*, right-click to set per switch values, set the **ip address** of RSVDC-CORE1-2 to *10.12.102.3*, and click **SAVE CHANGES**.

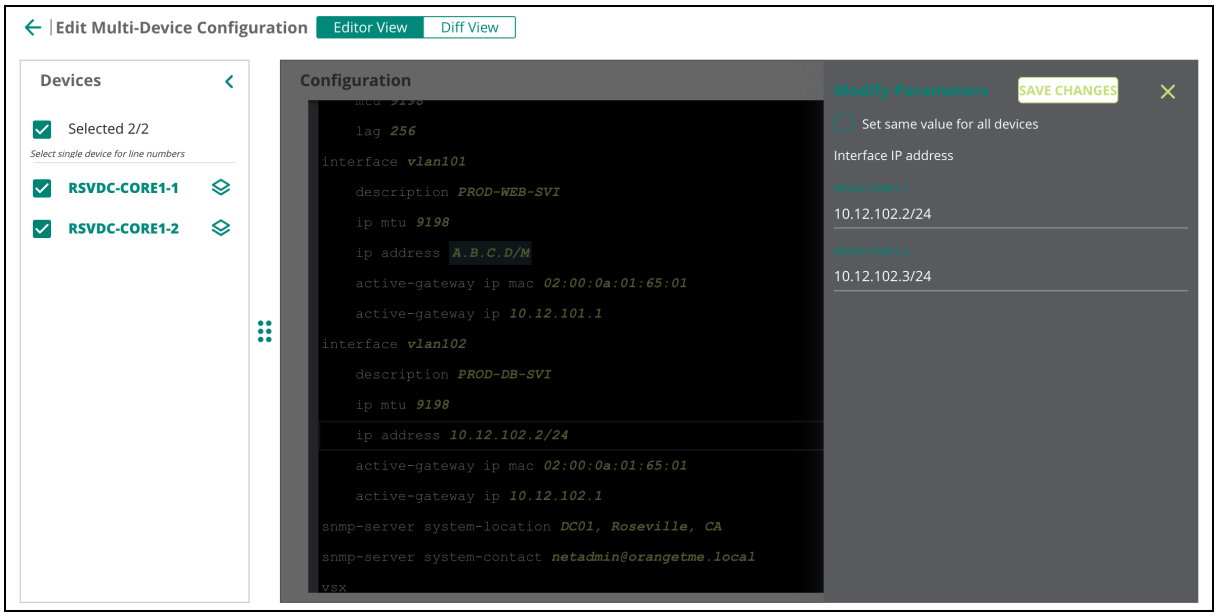


**Step 4** At the bottom right of the MultiEdit Configuration window, click **SAVE**.

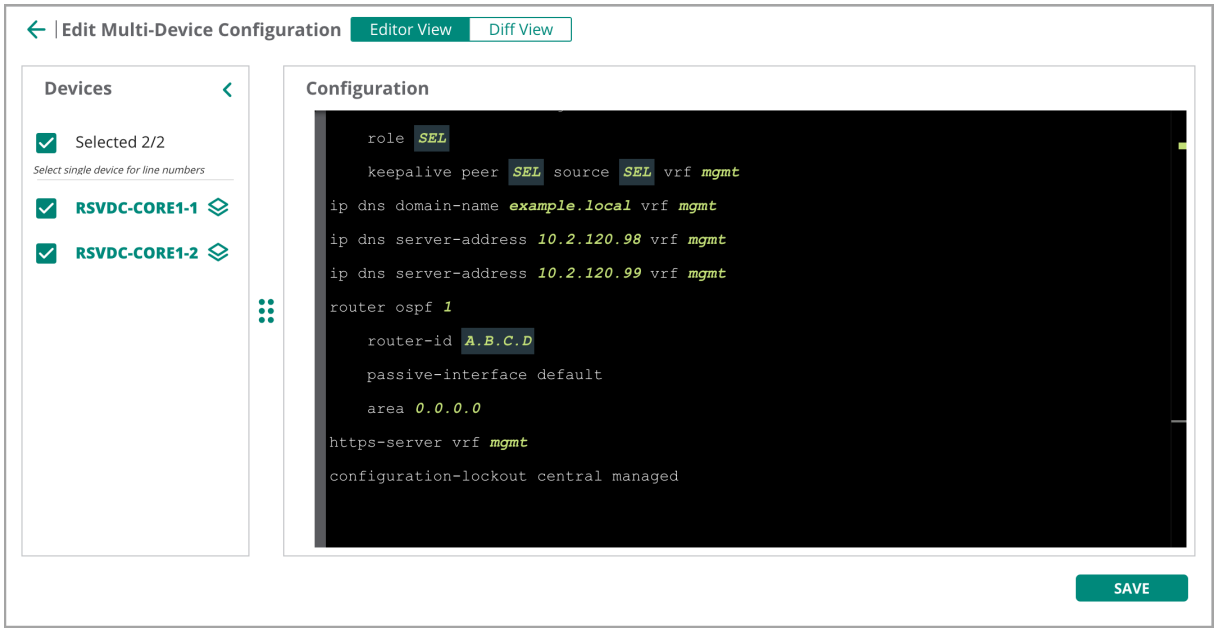


## Two-Tier Multicast

The ESP Two-Tier Data Center uses Protocol Independent Multicast—Sparse-Mode (PIM-SM) to distribute multicast source information and establish interface forwarding state. A centralized Rendezvous Point (RP) registers and distributes multicast sources throughout the network. The data center core switches perform multicast routing and typically point to an RP already established for campus usage to advertise data center sources and learn about campus sources. The Bootstrap Router (BSR) mechanism for PIM elects and learns the active RP.

Internet Group Messaging Protocol (IGMP) is enabled on data center host interfaces (routed-only interfaces and VLAN/SVI interfaces) to identify multicast listeners. When a host is interested in received traffic for a multicast group, it sends an IGMP join message.

**Step 1** Enable PIM routing on the core switches.

```
router pim
    enable
    active-active
```

**Step 2** Enable PIM-SM on external and data center host VLANs.

```
interface vlan4000
    ip pim-sparse enable
interface vlan101
    ip pim-sparse enable
interface vlan102
    ip pim-sparse enable
```

**Step 3** Enable IGMP on VLAN SVI interfaces for data center hosts.

```
interface vlan 101
    ip igmp enable
interface vlan 102
    ip igmp enable
```

## Verify Operational State

**Step 1** On the left navigation menu, click **Tools**.

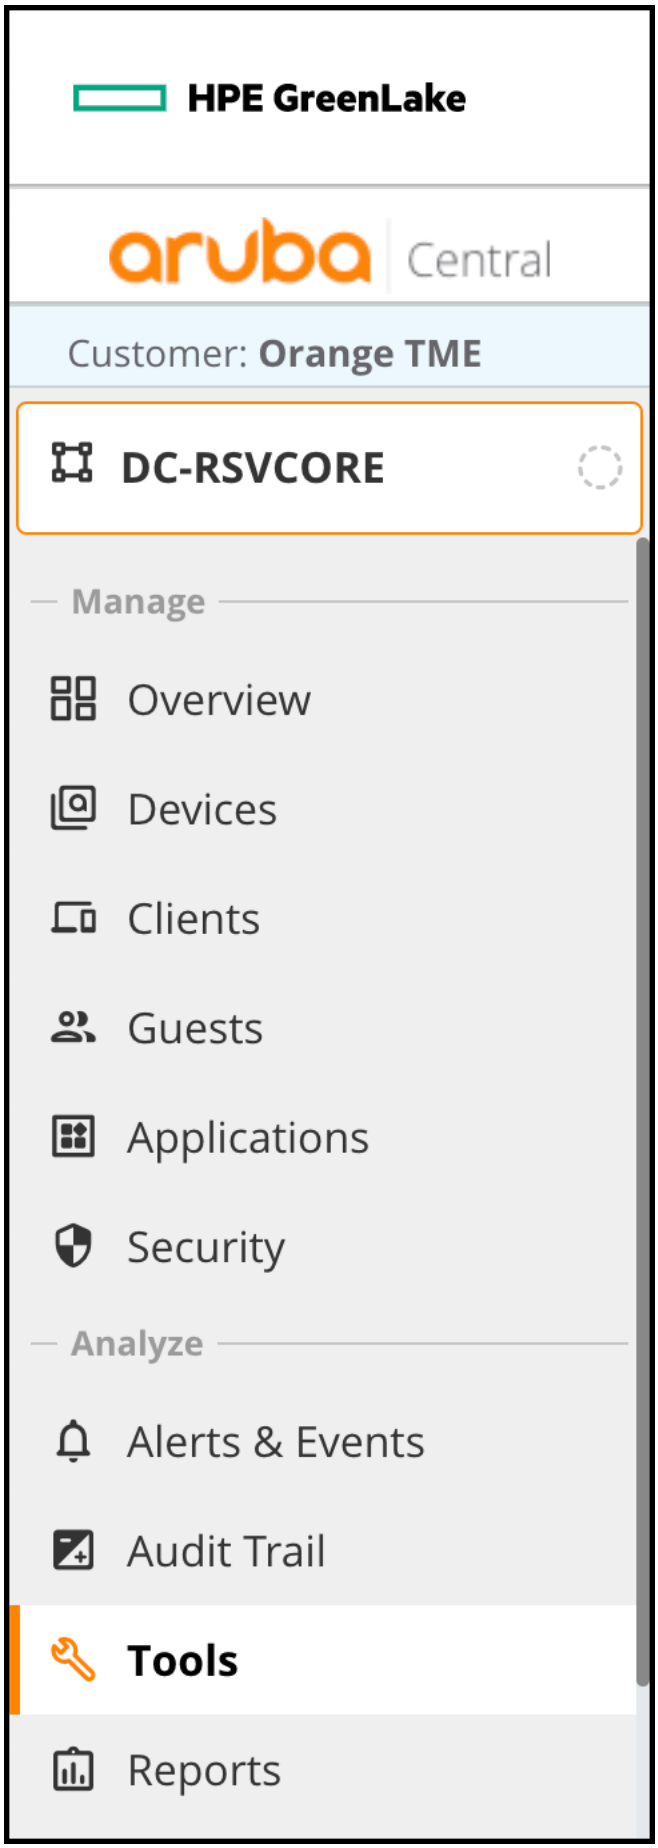
HPE GreenLake
aruba Central
Customer: Orange TME
DC-RSVCORE
Manage
Overview
Devices
Clients
Guests
Applications
Security
Analyze
Alerts & Events
Audit Trail
Tools
Reports

**Step 2** On the **Tools** menu at the top, click the **Commands** tab.

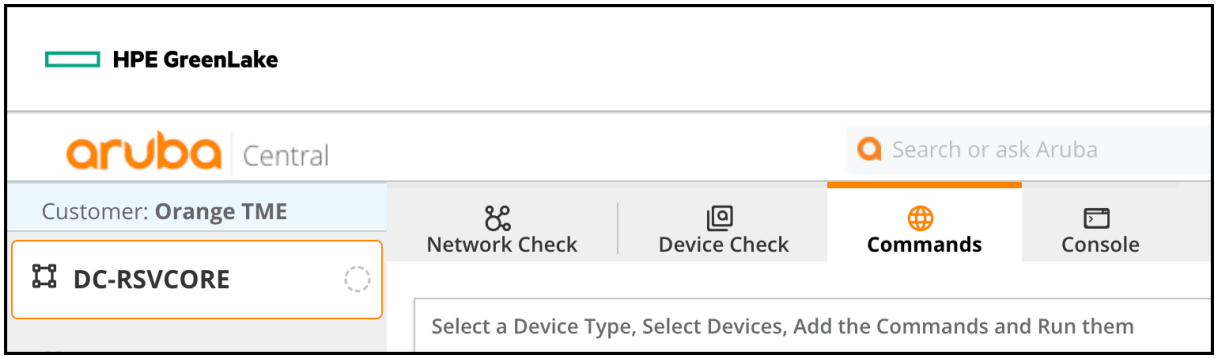


**Step 3** Click the **Available Devices** dropdown, select both data center core switches, then click elsewhere on the page.

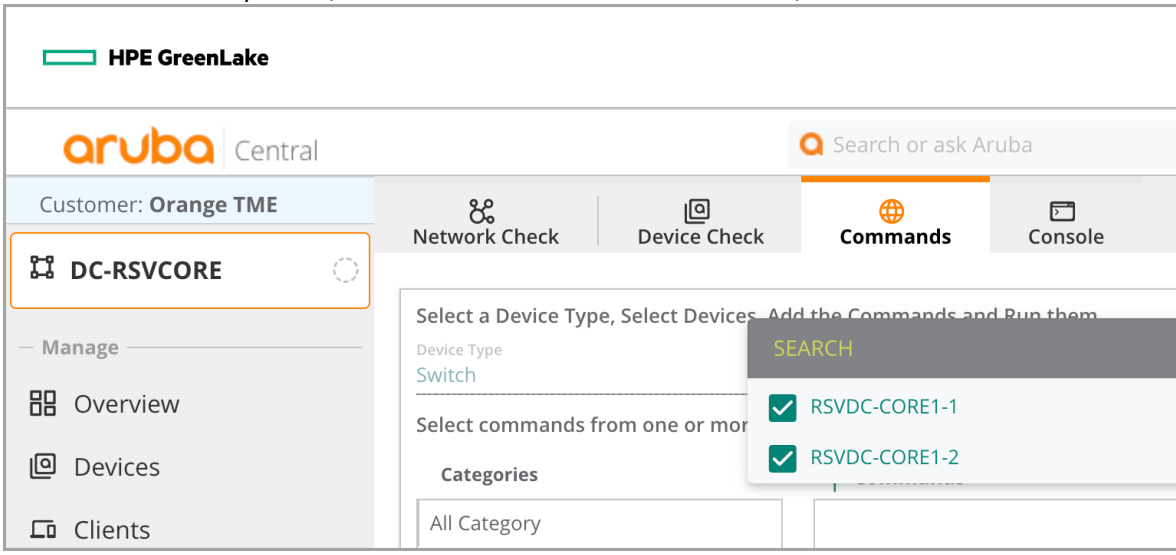


**Step 4** In the **Categories** list, click **All Category**. Enter *vsx* in the commands list filter, click **show vsx status**, then click **Add >**.

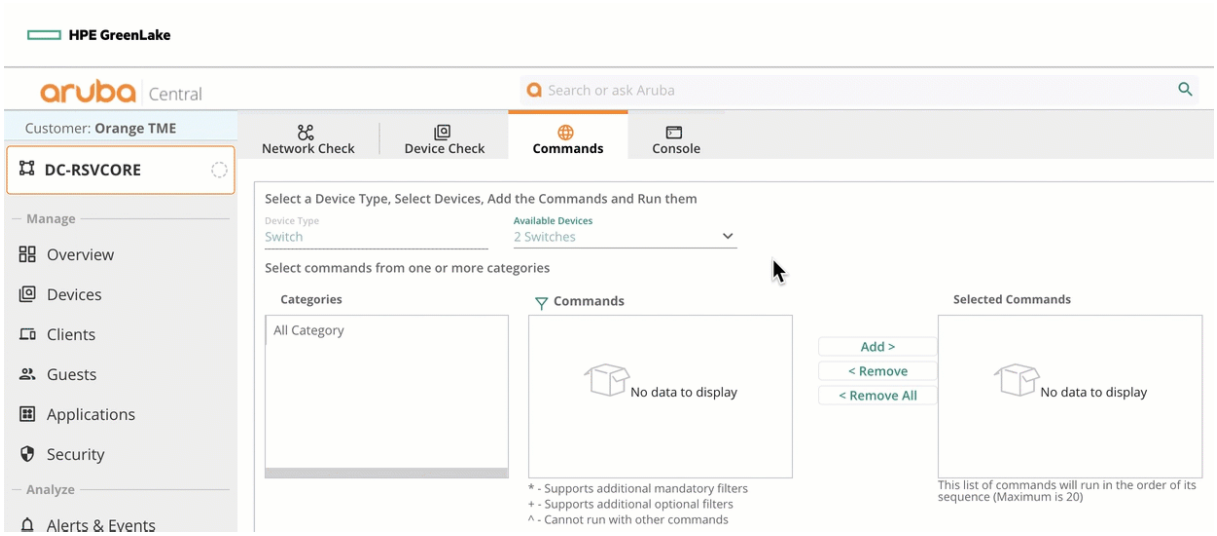


**Step 5** Add the following additional commands to the **Selected Commands** list.

- *show lacp interfaces*
- *show ip ospf interface all-vrfs*

- *show ip route all-vrfs*
- *show spanning-tree mst detail*
- *show ip pim*
- *show ip igmp*
- *show ntp status*

**Step 6** At the lower left of the **Commands** pane, click **RUN**.

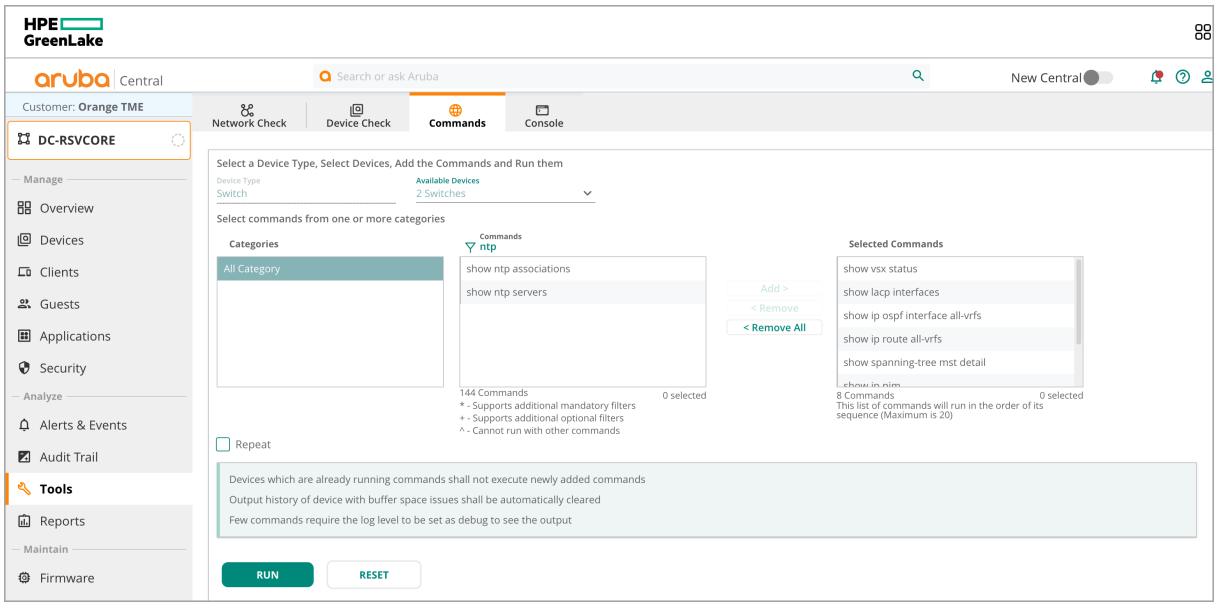


**Step 7** Scroll down to review the CLI command output for each switch. Verify key results for each command.

- *show vsx status*
  - **ISL channel**: *In-Sync*
  - **ISL mgmt channel**: *operational*
  - **Config Sync Status**: *In-Sync*
  - **Device Role**: set to *primary* and *secondary* on corresponding switches
  - Other VSX attributes display equal values for both VSX members

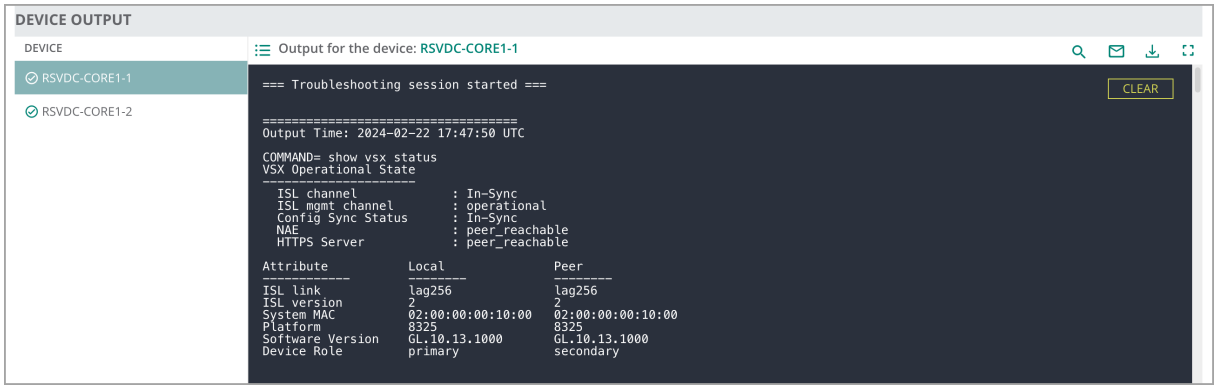


- *show lacp interfaces*
  - Both Actor and Partner have corresponding interfaces for each MC-LAG.

- All Actor interfaces have a **Forwarding State** of “up” for all host facing MC-LAGs and the upstream core switch facing MC-LAGs.
- All Actor and Partner interfaces have a state of “ALFNCD”.

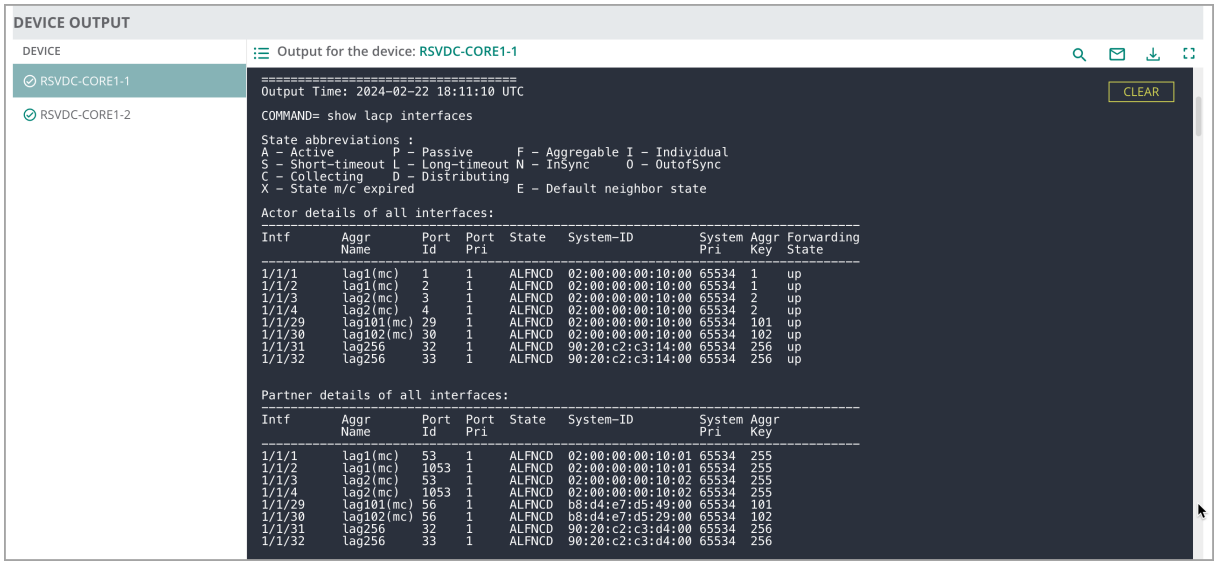


- *show ip ospf interface all-vrfs*
  - All interfaces display **Area** “0.0.0.0” and **Process** “1”.
  - VLAN 4000 **State/Type** is set to a valid value of DR, BDR or DR-other.
  - VLAN 4000 **DR** and **BDR** values are populated with IP addresses on the link.

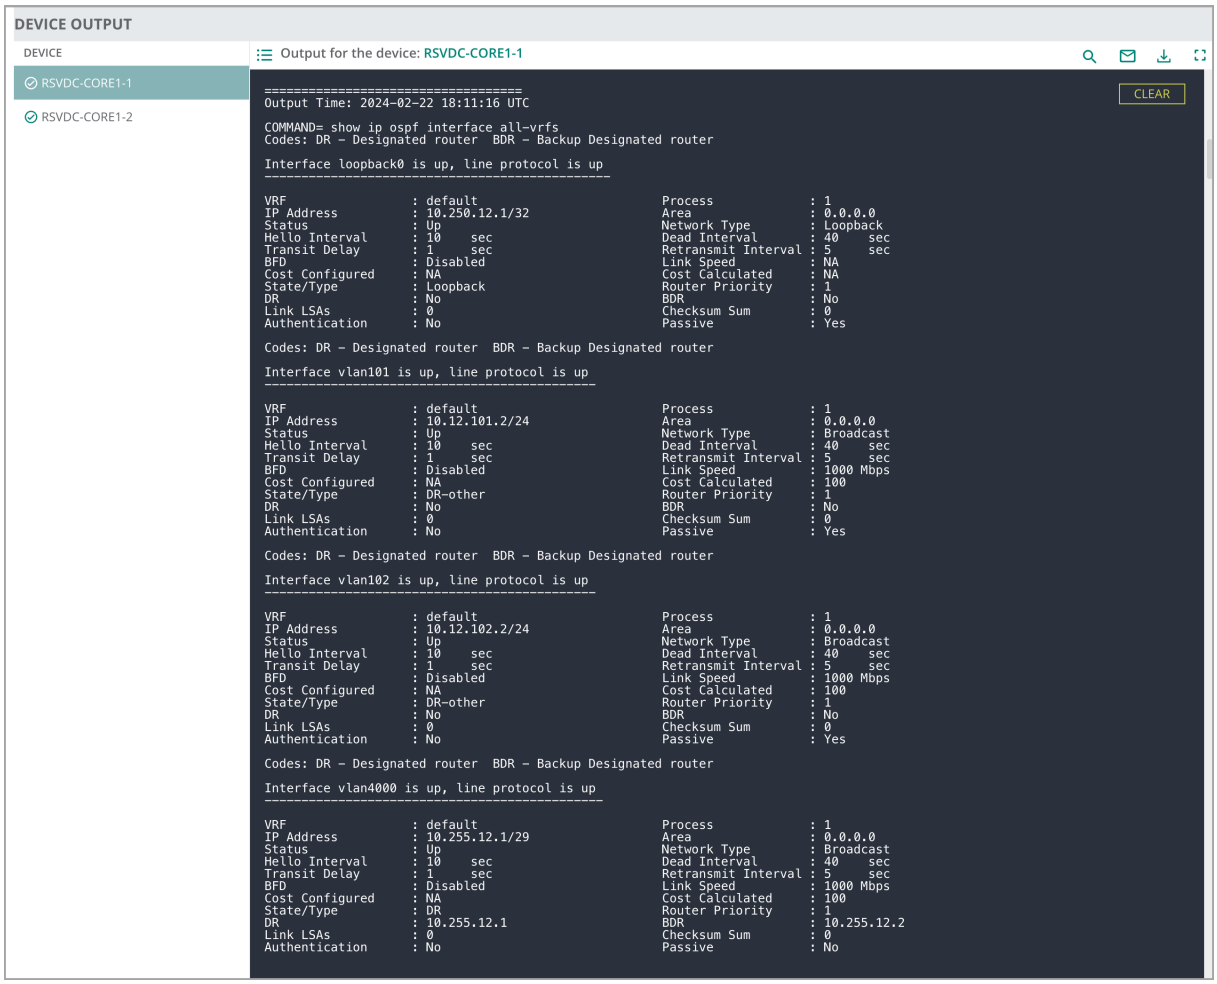


- *show ip route all-vrfs*
  - Verify that a default route is learned from the OSPF protocol and installed in the route table with the upstream firewall as the next hop.
  - Verify that valid campus routes are learned from the OSPF protocol.

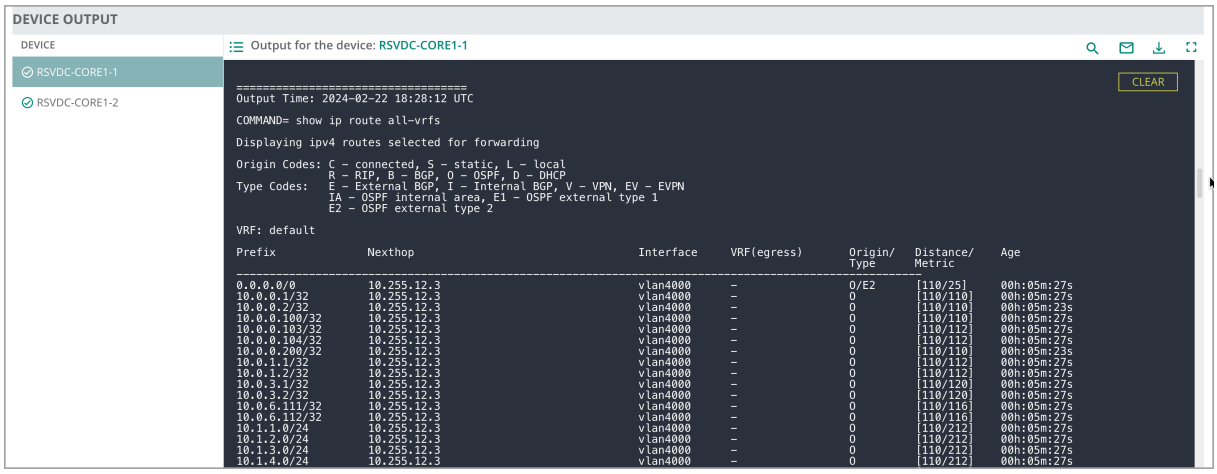


- *show spanning-tree mst detail*
  - Verify that the **Bridge Address** and **Root Address** values are the same.
  - Verify that all LAG interfaces have a **Role** of “Designated” and **State** of “Forwarding”.

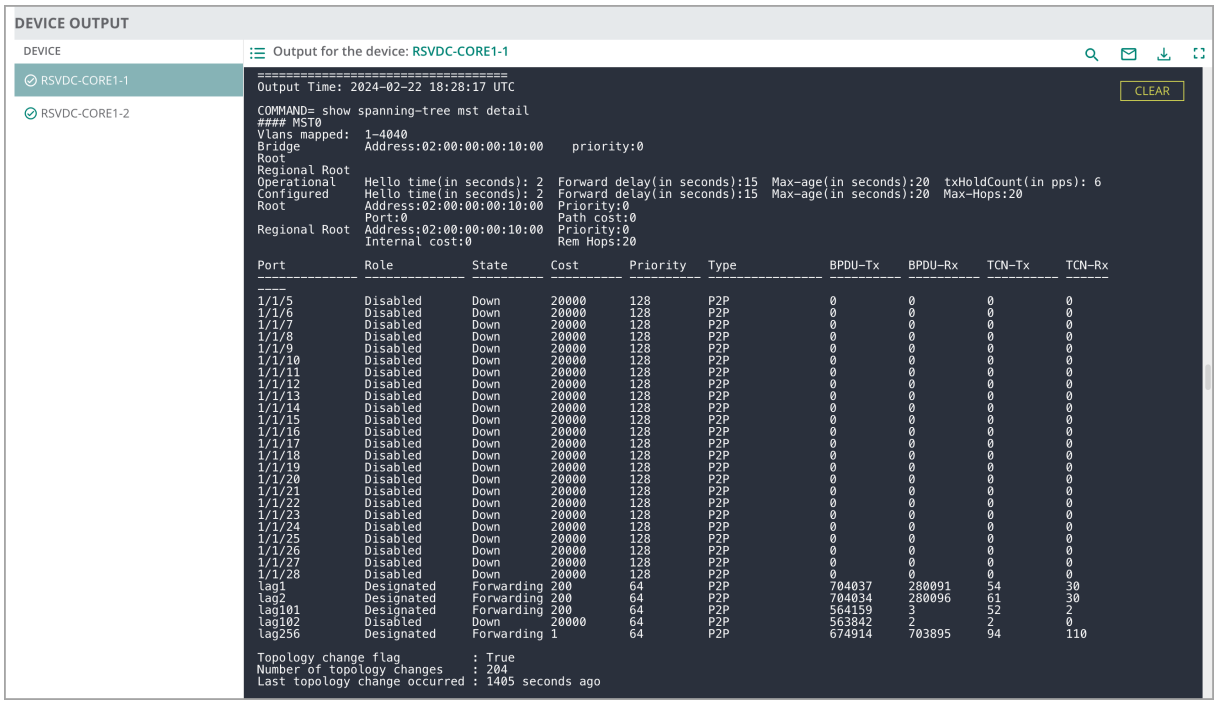


- *show ip pim*
  - Verify that **PIM Status** is “Enabled”.

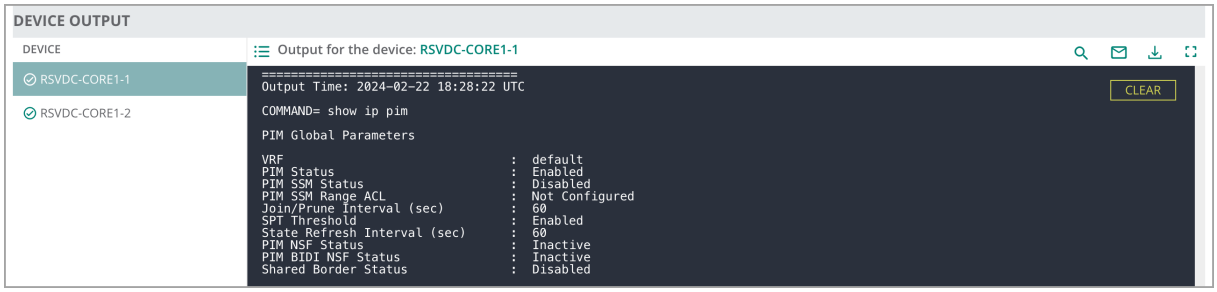


- *show ip igmp*
  - Verify that each host VLAN has an interface stanza with the **Querier IP** field populated with an IP address of one of the corresponding core switch VLAN SVIs.

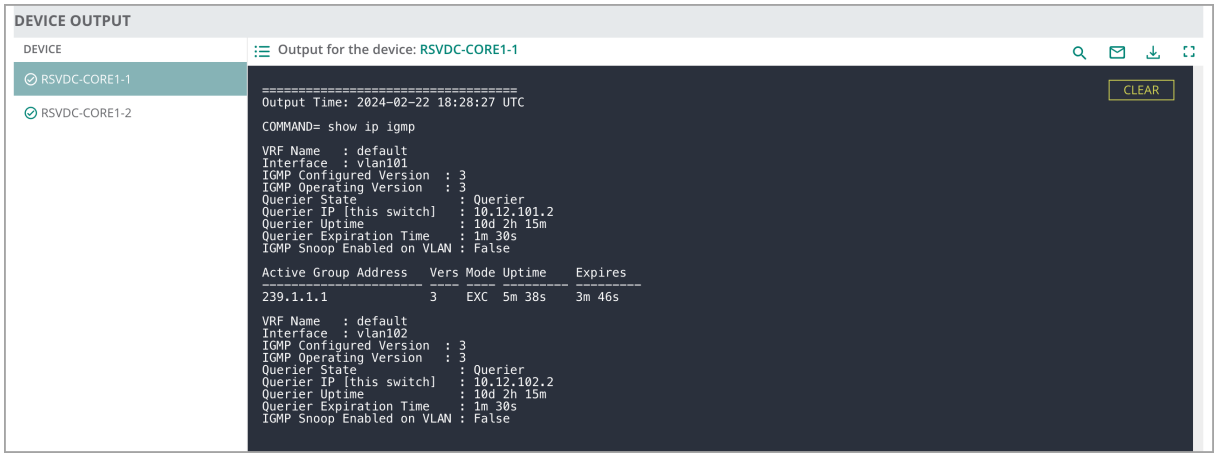


- *show ntp status*

- Verify that **NTP Server** is populated with a configured NTP server IP address.
- Verify that the **Time Accuracy** field is populated.

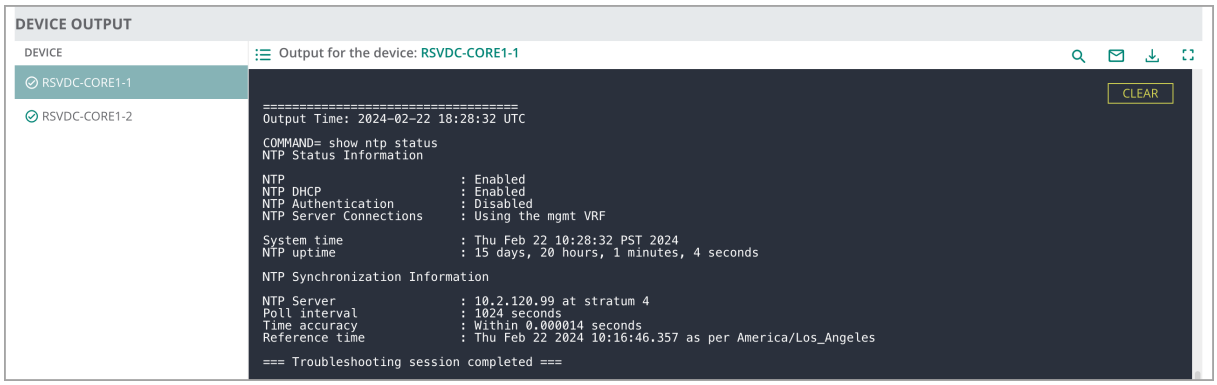

# Two-Tier Server Access

Configure Two-Tier access switches as VSX pairs for redundant multi-chassis link aggregation (MC-LAG) connections to the core and downstream data center hosts.

## Configure Access VSX ISL Interface

To establish a VSX relationship between each pair of access switches, a link aggregation (LAG) interface must be created for assignment as the VSX data plane's inter-switch link (ISL). Standardizing the ToR model enables configuring the same ports on all access switches for the VSX ISL link at the UI group level.

**Step 1** On the left navigation menu, click **DC-RSVCORE**, then click the data center access switch group name in the **Groups** column.

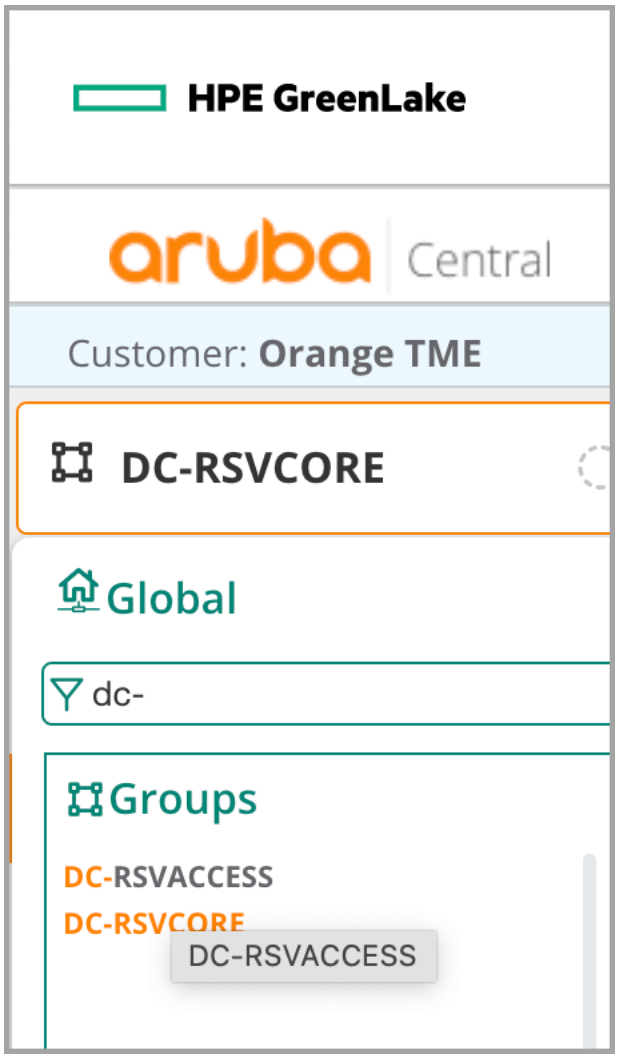

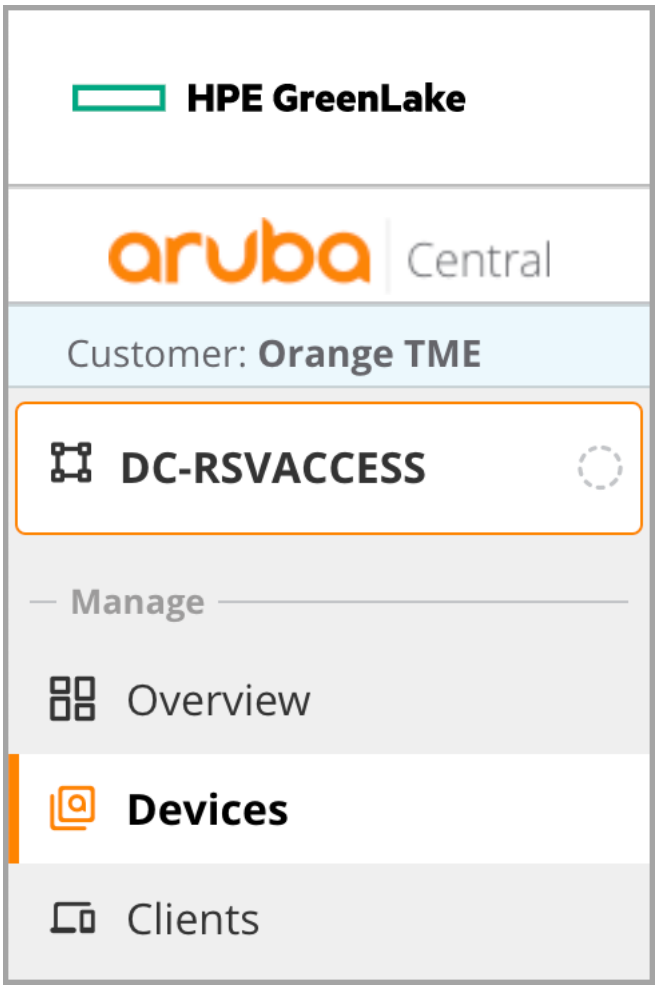


**Step 2** On the left navigation menu, click **Devices**.

**Step 3** At the upper right of the **Switches** pane, click **Config**.

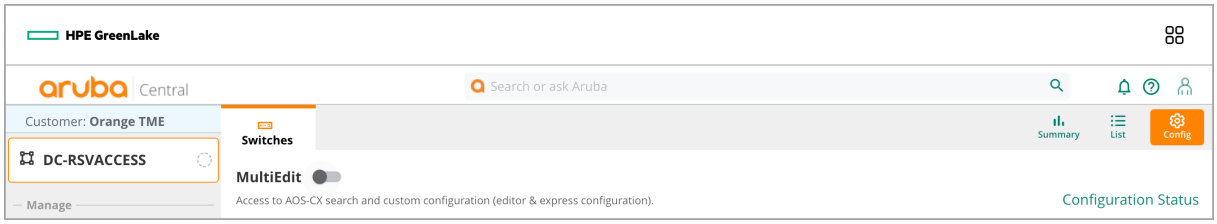


**Step 4** In the **Interfaces** tile, click **Ports & Link Aggregations**

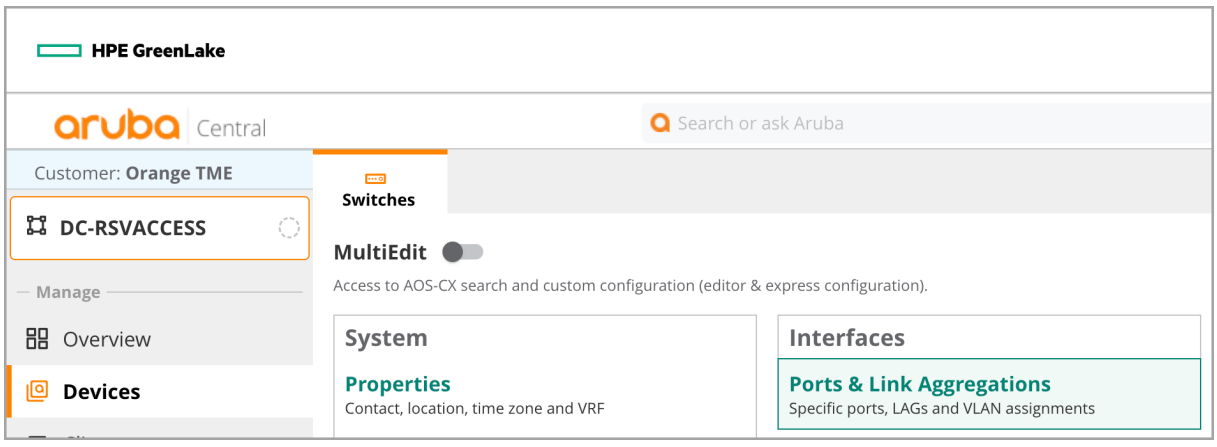

**Step 5** Scroll to the right of the **Ports & Link Aggregations** table, and click the **+** (plus sign) in the upper right.

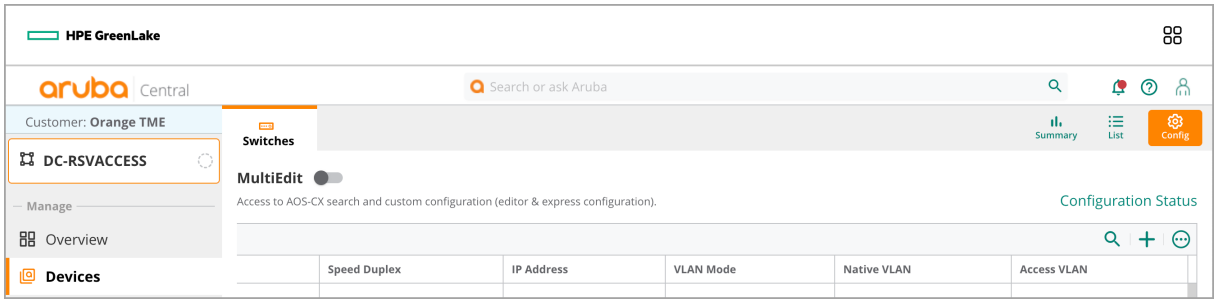


**Step 6** On the **Add LAG** page, assign the following values:

- **Name:** *lag256*
- **Description:** *VSX_ISL_LAG*
- **Port Members:** *1/1/49, 1/1/50*
- **Speed Duplex:** *<no value>* (default)
- **VLAN Mode:** *trunk*
- **Native VLAN:** *1* (default)
- **Allowed VLANs:** *<no value>* (default)
- **Admin Up:** *checked*
- **Aggregation Mode:** *LACP Active*

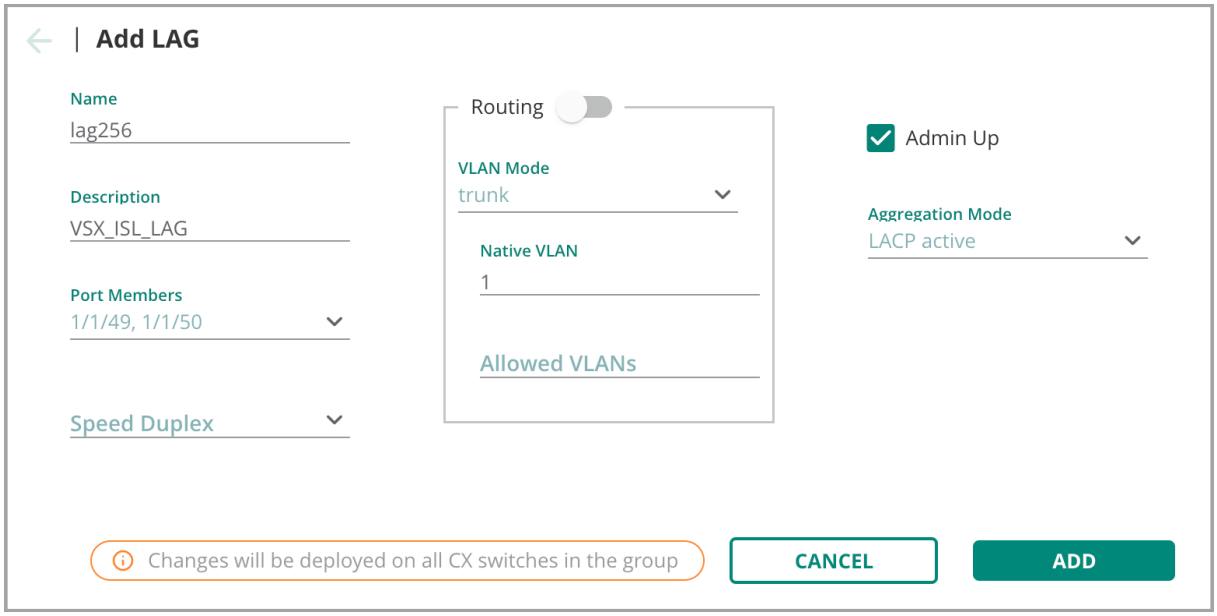


**Step 7** In the **Ports & Link Aggregations** table's title row, click ← (left arrow) to return to the main configuration page.

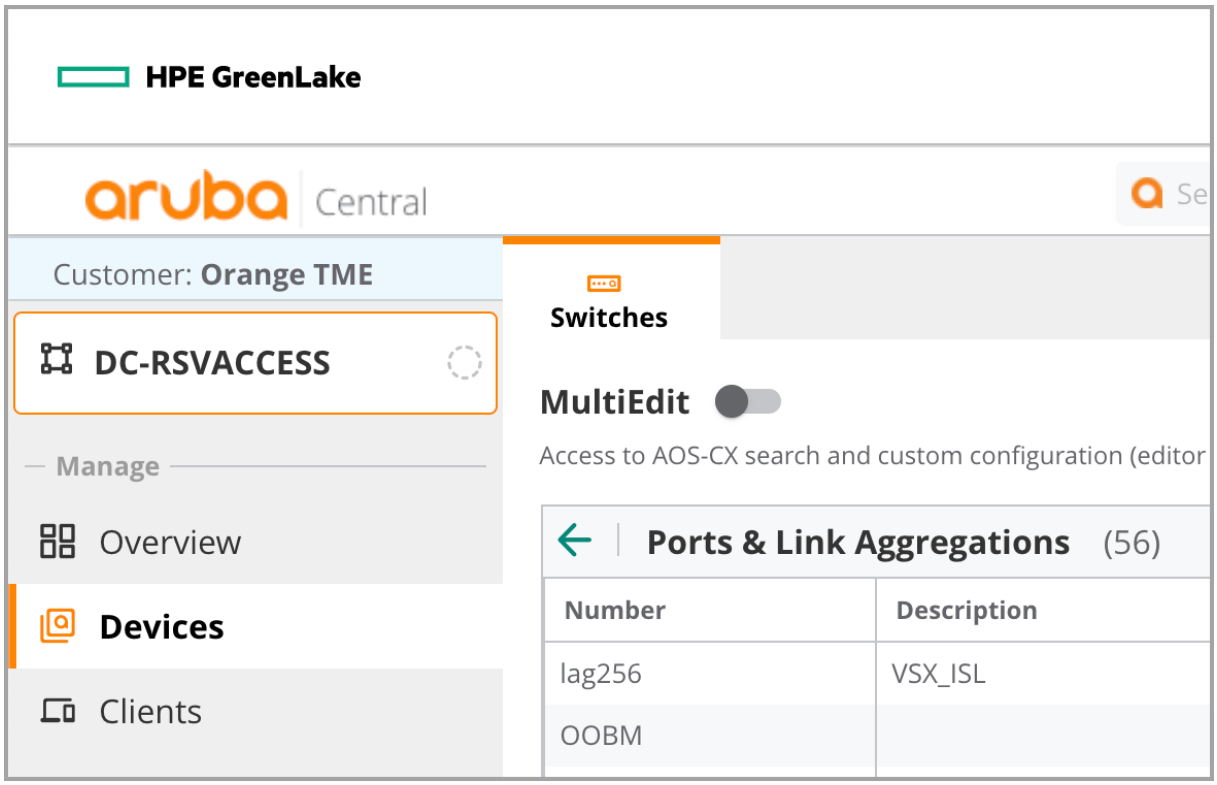


## Spanning Tree

MC-LAGs provide loop prevention in a Two-Tier architecture. Spanning-tree (STP) is configured as a backup loop prevention mechanism in case of host cabling errors to ToR switches.

**Step 1** In the **Bridging** tile, click **Loop Prevention**.

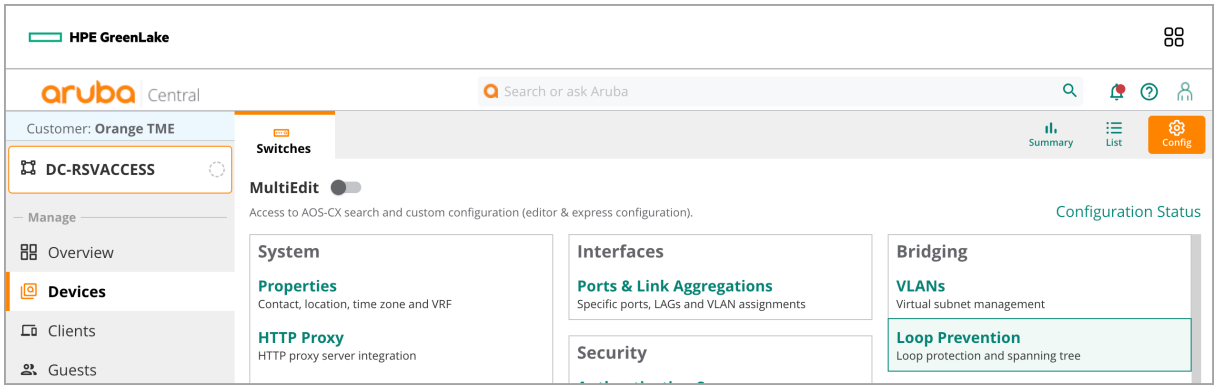


**Step 2** In the **Loop Prevention** window, set the Spanning Tree **Region** to **RSVDC**, leave all other settings at their default, then click **SAVE**.

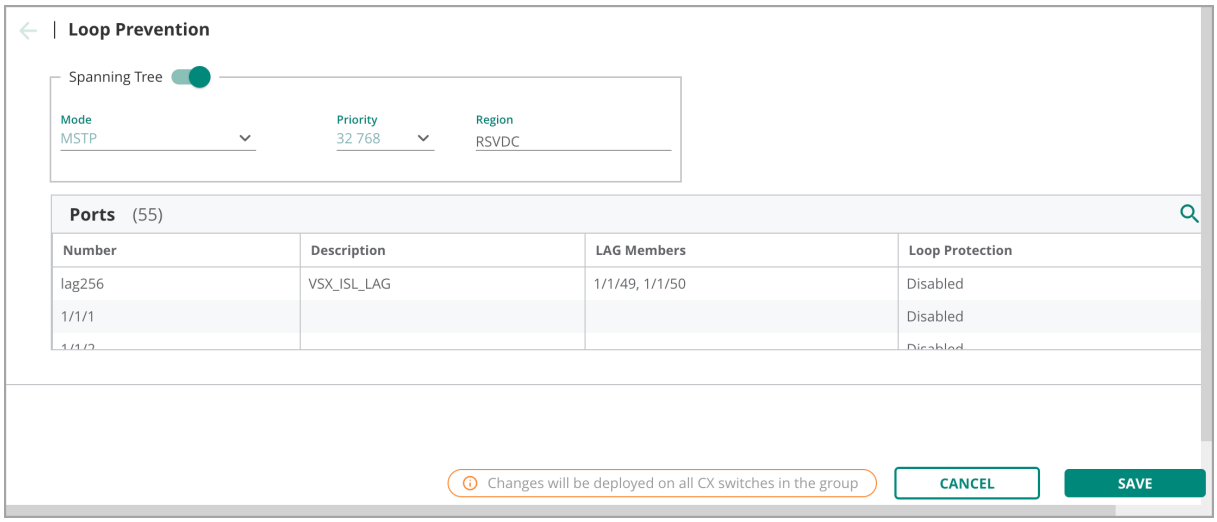


# Enter MultiEdit Configuration

**Step 1** At the upper left of the **Switches** pane, click the **MultiEdit** enable slider.

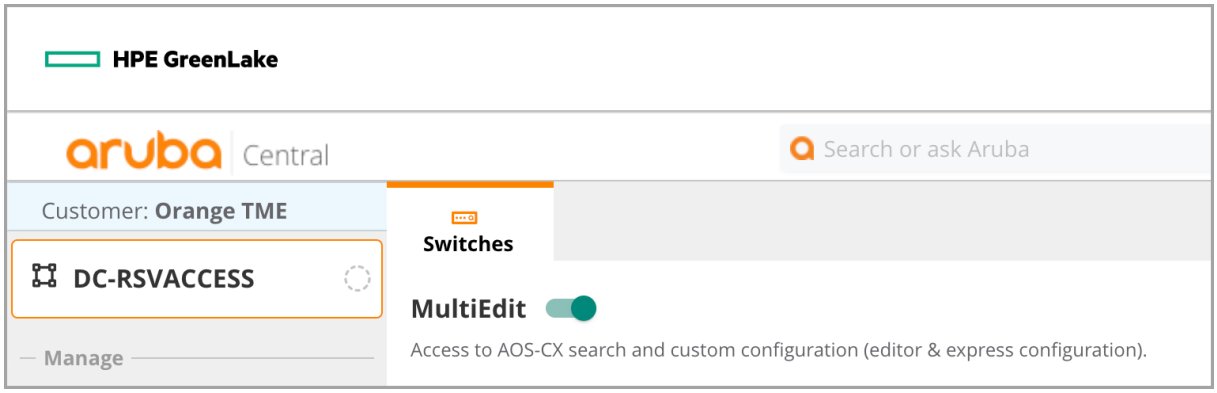


**Step 2** Select all access switches in the **Devices** lists, then click **EDIT CONFIG**.

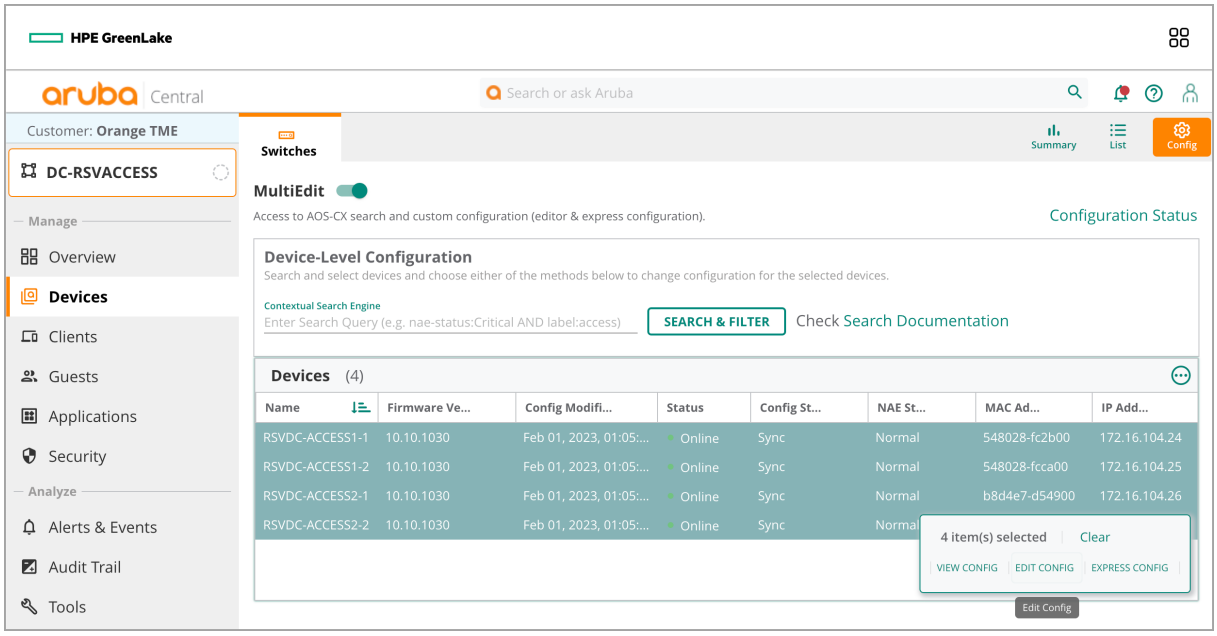

# Configure Access Switch VSX Pairs

The access switches are configured as VSX pairs to support Layer 2 multi-chassis link aggregation to the core layer and downstream data center hosts. A two-port link aggregation is configured and assigned as the VSX data path inter-switch link (ISL). The out-of-band *mgmt* interface is used for VSX keepalives to maximize the number of ports available for connecting access switches.

**Step 1** Enter the initial VSX configuration.

```
vsx
    system-mac 02:00:00:00:10:01
    inter-switch-link lag 256
    role primary
    keepalive peer 172.16.104.25 source 172.16.104.24 vrf mgmt
```

**Step 2** Mouse-over the field values in the table column headings below, right-click, and set the appropriate values for each switch.

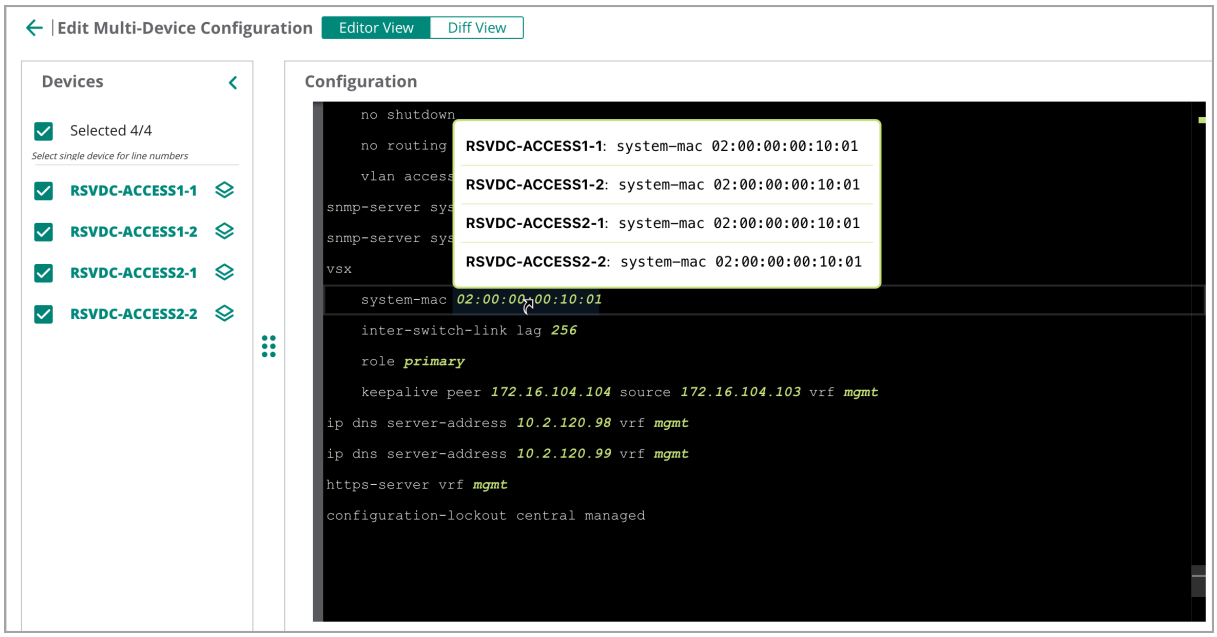


| Switch | system-mac | role | peer | source |
|---|---|---|---|---|
| RSVDC-ACCESS1-2 | 02:00:00:00:10:0 change] | secondary | 172.16.104.103 | 172.16.104.104 |
| RSVDC-ACCESS2-1 | 02:00:00:00:10:02 | primary[no- change] | 172.16.104.106 | 172.16.104.105 |
| RSVDC-ACCESS1-2 | 02:00:00:00:10:0 | secondary | 172.16.104.105 | 172.16.104.106 |

**Step 3** Assign a description and maximum MTU value for the VSX ISL physical interfaces.

```
interface 1/1/49
    description VSX-ISL
    mtu 9198
interface 1/1/50
    description VSX-ISL
    mtu 9198
```

## Configure Access to Core MC-LAGs

**Step 1** Create the core-facing MC-LAG interface.

```
interface lag 255 multi-chassis
    no shutdown
    description DC-CORE
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
```

**NOTE:**

Tag all VLANs on all inter-switch MC-LAGs to support ubiquitous host mobility across all racks in the Two-Tier structure.

**Step 2** Assign physical interfaces to the core MC-LAG interface.

```
interface 1/1/53
    no shutdown
    mtu 9198
    description RSVDC-CORE1-1
    lag 255
interface 1/1/54
    no shutdown
    mtu 9198
    description RSVDC-CORE1-2
    lag 255
```

**NOTE:**

The same physical interface on each access switch in the data center should connect to the same upstream core switch. For example, interface 1/1/53 on every ToR access switch can be configured to connect to the primary switch in the VSX core pair. This creates a consistent configuration that is easy to troubleshoot.

**Step 3** Remove the following configuration line from interfaces 1/1/53 and 1/1/54: *vlan access 1*

## Configure Access Switch to Host MC-LAGS

**Step 1** Configure the host-facing MC-LAG interface.

```
interface lag 1 multi-chassis
    no shutdown
    description ESXi-01
    no routing
    vlan trunk native 1
    vlan trunk allowed 101-102
    lacp mode active
    spanning-tree root-guard
```

**NOTE:**

MC-LAG interfaces to downstream hosts should scope allowed VLANs only to those required for host.

**Step 2** Mouse-over the description field, right-click, then modify values appropriately for each switch.

| Switch | description |
|---|---|
| RSVDC-ACCESS1-2 | ESXi-01[no-change] |
| RSVDC-ACCESS2-1 | ESXi-02 |
| RSVDC-ACCESS2-2 | ESXi-02 |

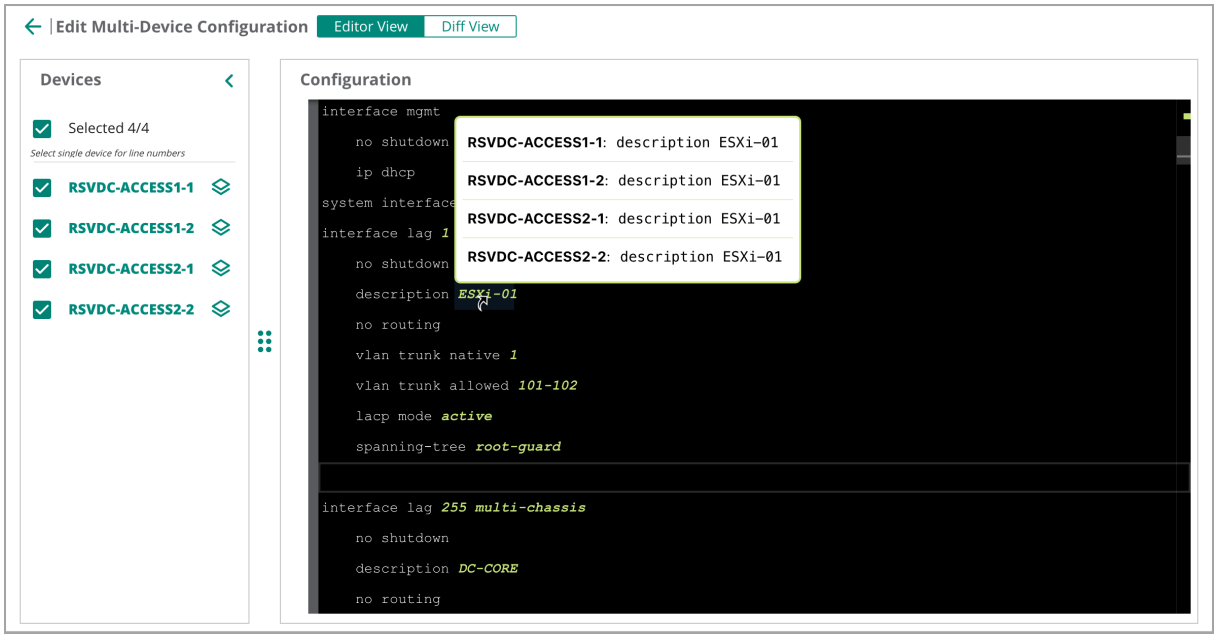


**NOTE:**

Additional field configuration values, such as allowed VLANs on the trunk, also can be modified, when appropriate.

**Step 3** Associate physical interfaces with the MC-LAG. The following configuration assigns an interface MTU and associates the interfaces with the previously created MC-LAG interface.

```
interface 1/1/1
    no shutdown
    mtu 9198
    lag 1
```

**NOTE:**

Standardize the association of LAG index values to physical interfaces across all access switches for efficient configuration of switch interfaces. The example above assigns LAG index 1 to interface 1/1/1 on all selected switches in MultiEdit.

**Step 4** Remove the following configuration line from interfaces 1/1/1: *vlan access 1*

**Step 5** Repeat this process for each host facing MC-LAG.

## Configure Multicast

Server Access switches do not perform multicast routing or client services, but IGMP snooping observes IGMP requests from hosts to optimize multicast forwarding at Layer 2. IGMP populates Multicast MAC addresses corresponding to IP multicast groups in the MAC table. This conserves bandwidth by ensuring that multicast traffic is forwarded only to interested receivers.

**Step 1** Enable IGMP snooping on all VLAN interfaces.

```
vlan 101
    ip igmp snooping enable
vlan 102
    ip igmp snooping enable
```

**Step 2** At the bottom right of the MultiEdit **Configuration** window, click **SAVE**.

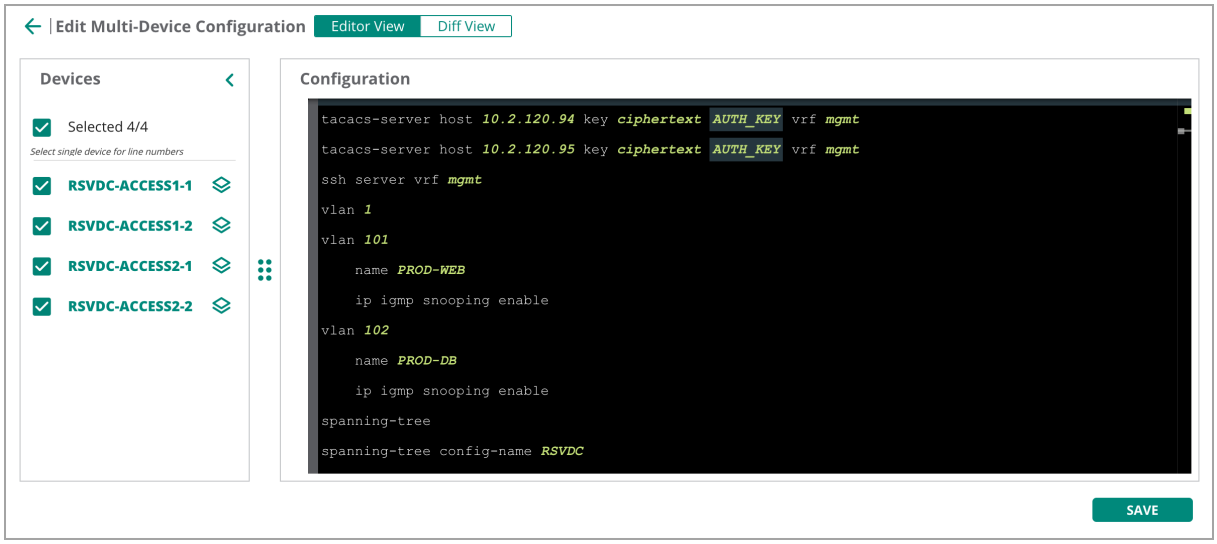

# Configure Physical Port Speeds

The default port speed on a switch may be different than the supported speed of a connected device. When attached host speeds are not common across racks, MultiEdit can be used to select only the ToR VSX pair of switches to be modified.

Aruba CX 8325 and CX 10000 switches set physical interface speeds in groups. Every non-uplink interface is associated with an interface group. All members of an interface group use the same operational speed. The size of the group depends on the switch model. This sample topology uses a CX 8325, which groups sets of 12 non-uplink physical interfaces to four distinct interface groups.

**Step 1** Select switches that require interface group speed settings, then click **EDIT CONFIG**.

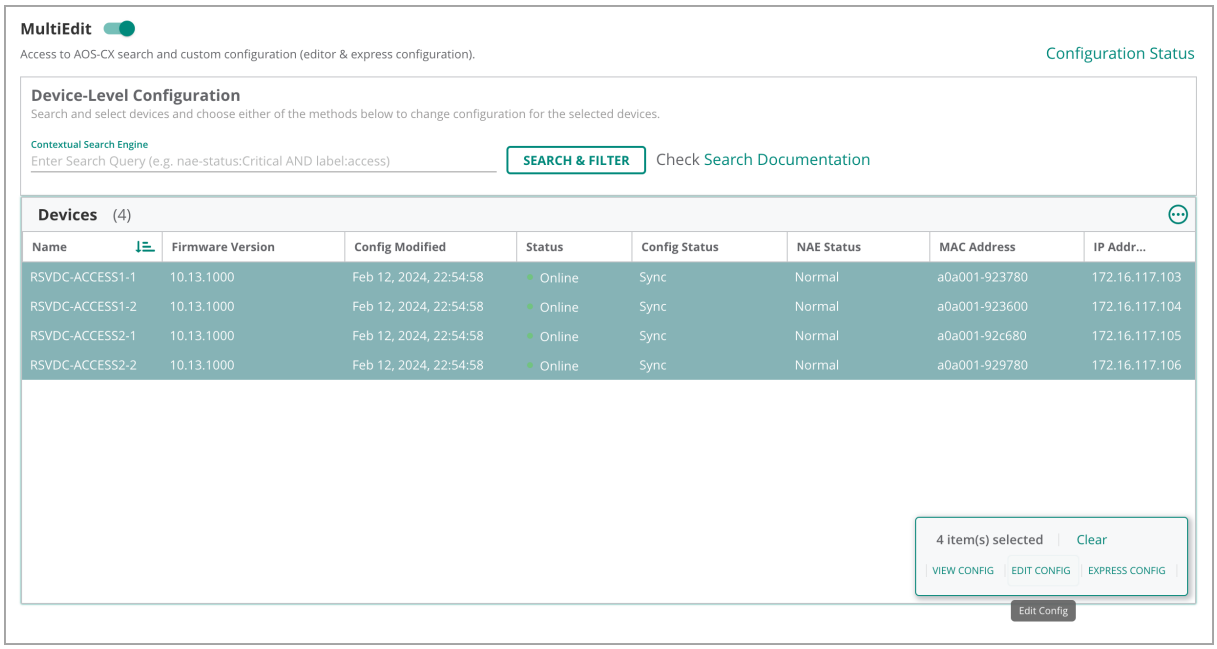


**Step 2** Set the interface group port speed to *10Gbps*.

```
system interface-group 1 speed 10g
```

**NOTE:**

The command above sets physical ports 1/1/1–1/1/4 on an 8360 and ports 1/1/1–1/1/12 on an 8325 to operate at 10 Gbps.

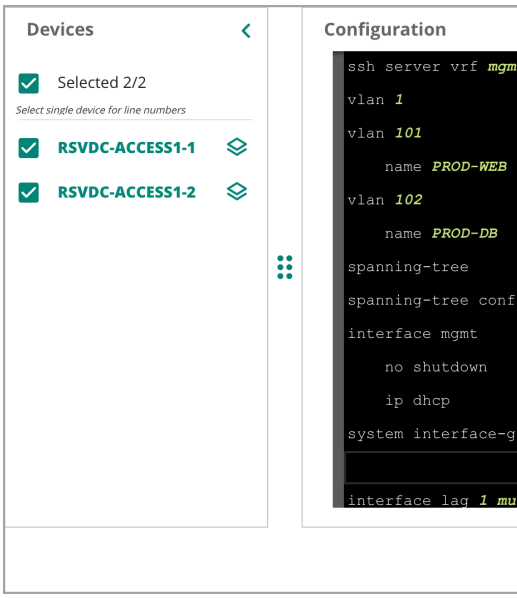


**Step 3** At the lower right of the MultiEdit **Configuration** window, click **SAVE**.

## Verify Configuration

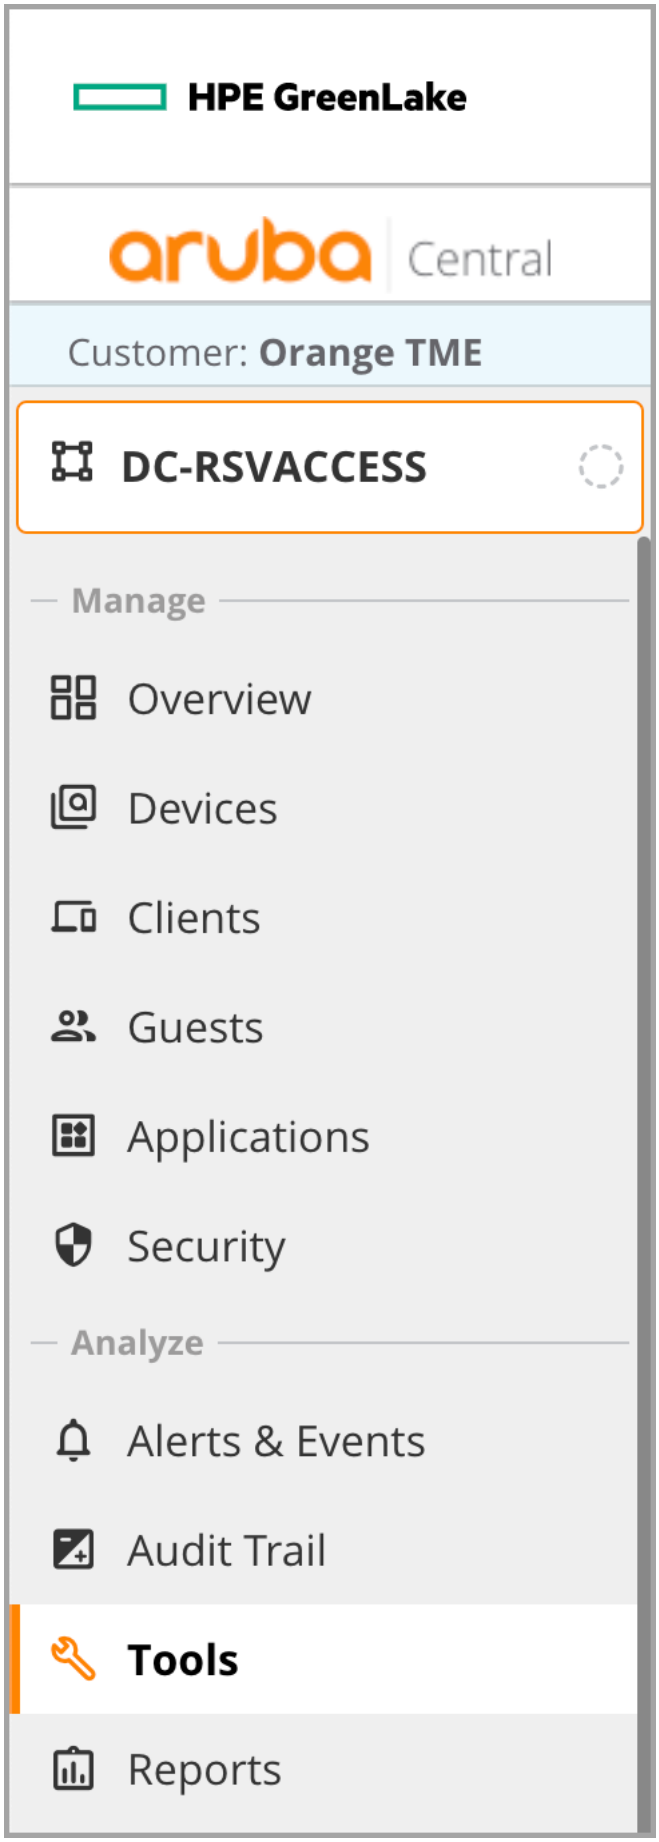


**Step 1** On the left navigation menu, click **Tools**.

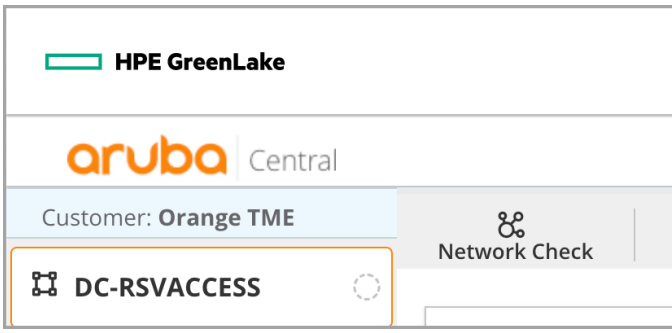


**Step 2** On the **Tools** menu at the top, click the **Commands** tab.

**Step 3** Click the **Available Devices** dropdown, select all access switches, then click elsewhere on the page.

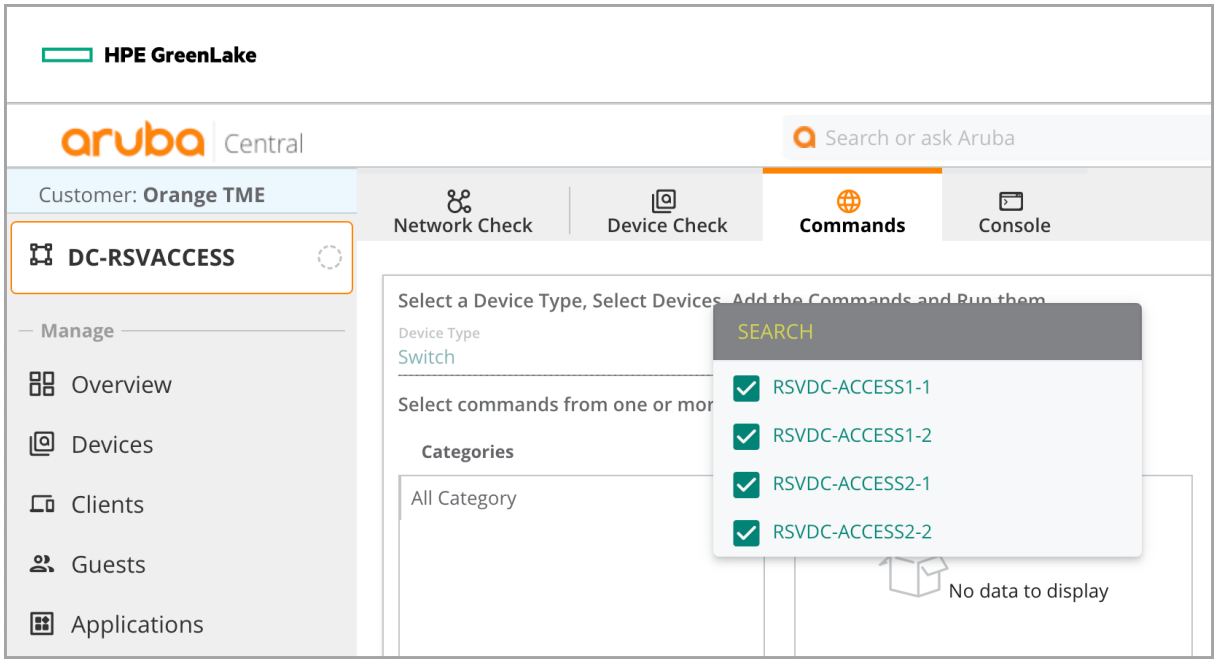


**Step 4** In the **Categories** list, click **All Category**. Enter *vsx* in the commands list filter, click **show vsx status**, then click **Add >**.

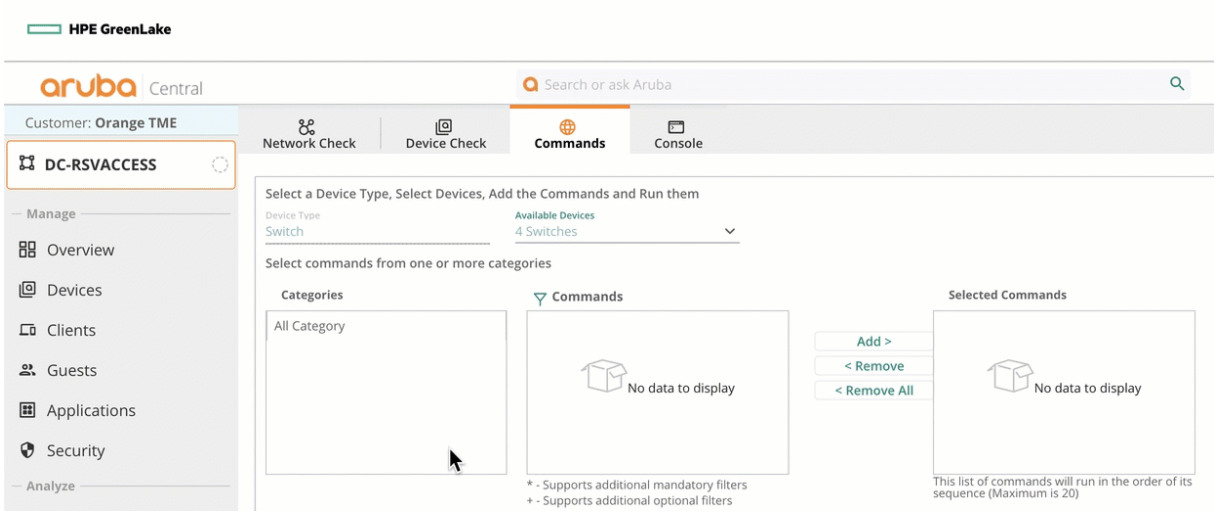


**Step 5** Add the following additional commands to the **Selected Commands** list.

- *show lacp interfaces*

- *show spanning-tree mst detail*
- *show ntp status*

**Step 6** At the lower left of the **Commands** pane, click **RUN**.

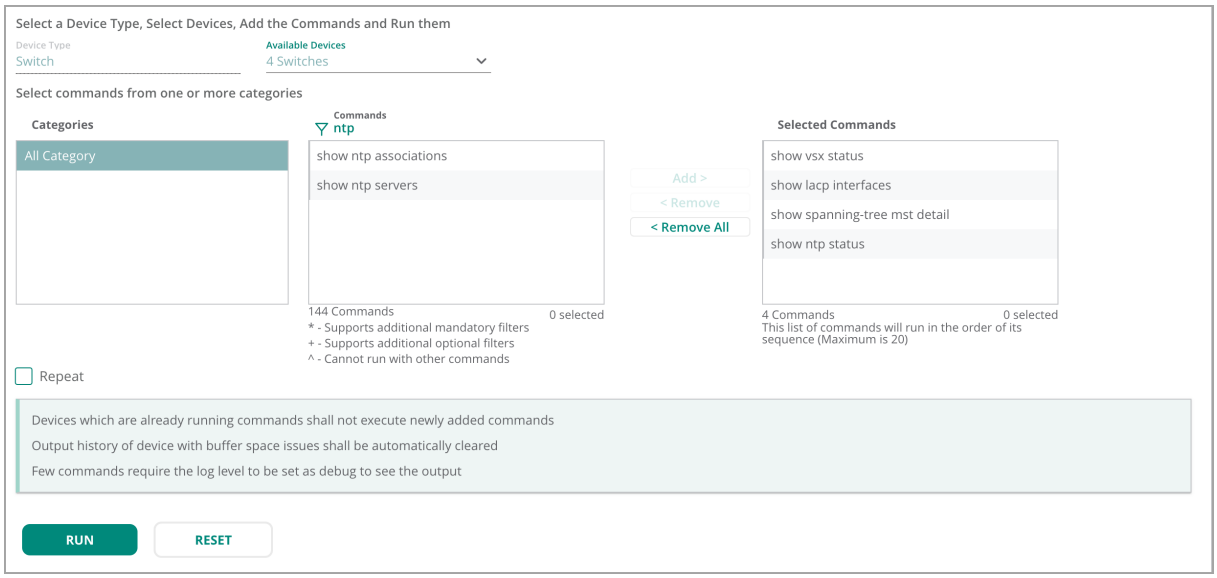


**Step 7** Scroll down to review the CLI command output for each switch. Verify key result data for each command.

- *show vsx status*
  - **ISL channel**: *In-Sync*
  - **ISL mgmt channel**: *operational*
  - **Config Sync Status**: *In-Sync*
  - **Device Role**: set to *primary* and *secondary* on corresponding switches
  - Other VSX attributes display equal values for both VSX members

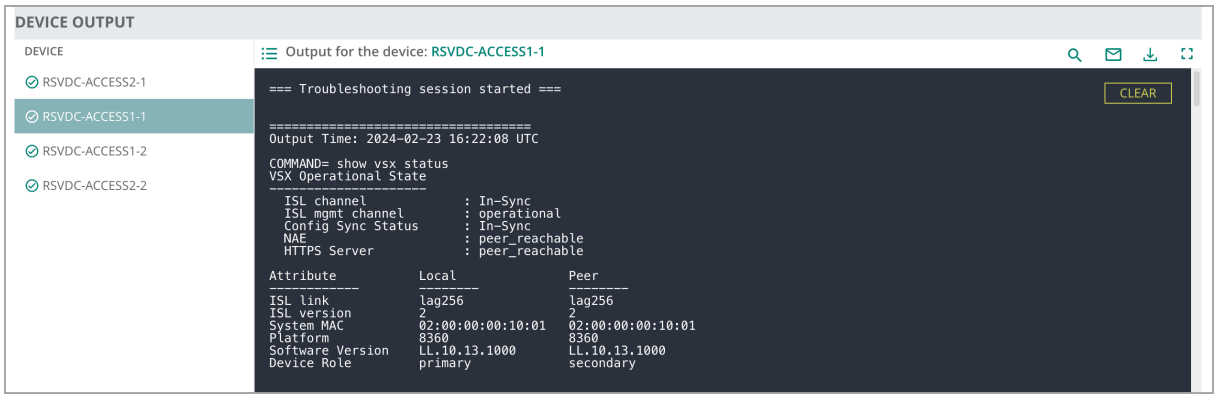


- *show lacp interfaces*
  - Both Actor and Partner have a corresponding interface for each MC-LAG.
  - All Actor interfaces have a **State** of “ALFNCD”.
  - All Actor interfaces have a **Forwarding State** of “up” for all host facing MC-LAGs and the upstream core switch facing MC-LAGs.
  - All Partner interfaces have a state of “PLFNCD” or “ALFNCD”.

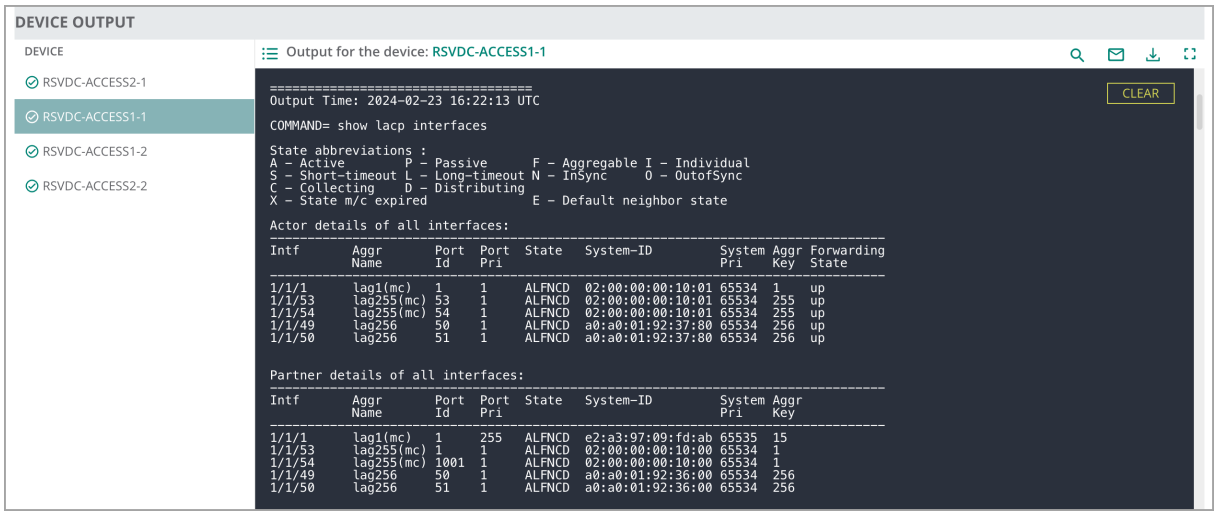


**NOTE:**

"(mc)" in the **Aggr Name** column indicates an MC-LAG. The switch on which the *show lacp interfaces* command is run is considered the Actor. The other VSX member switch is considered the Partner.

- *show spanning-tree mst detail*
  - Verify that the **Root Address** value is the virtual VSX MAC address on the core switches.
  - Verify that the **Role** for LAG 255 connected to the core switches is "Root" with a **State** of "Forwarding".
  - Verify that the **Role** for all other LAGs and ports with connections is "Designated" with a **State** of "Forwarding".

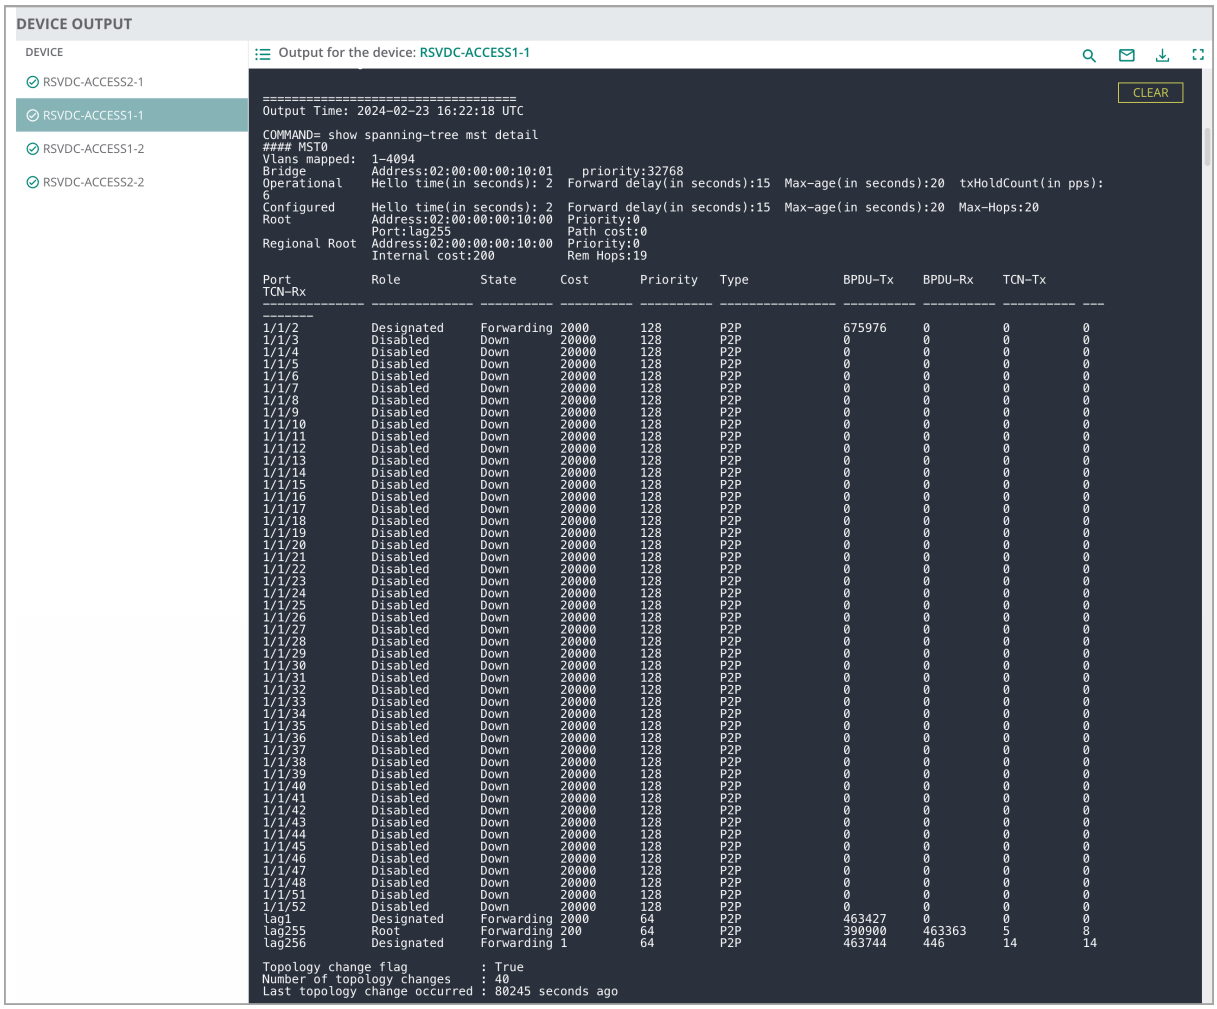


- *show ntp status*
  - Verify that **NTP Server** is populated with a configured NTP server IP address
  - Verify that the **Time Accuracy** field is populated.

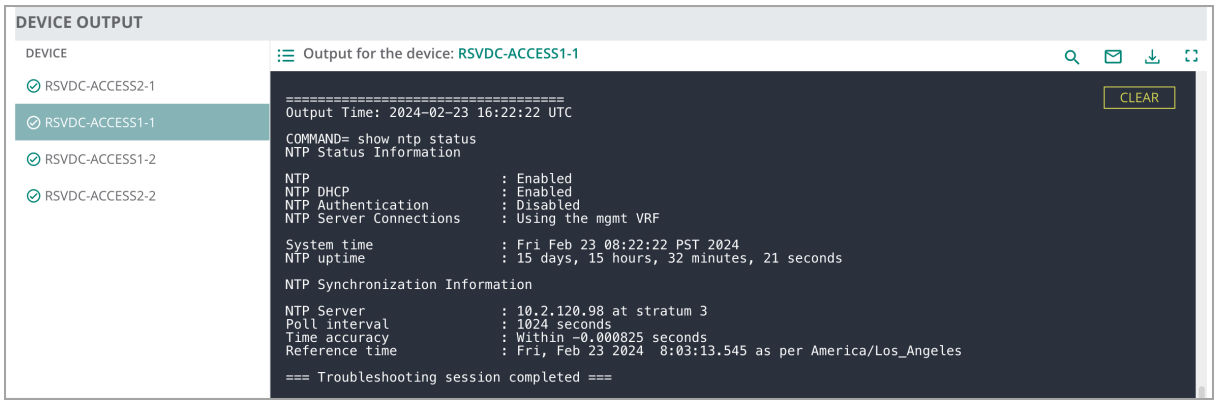

# Ansible Two-Tier Data Center

HPE Aruba is committed to providing effective, flexible network automation strategies tailored to customer needs. In addition to workflow-based automations provided by Aruba Central and Aruba Fabric Composer, the HPE Aruba Networking Developer Hub provides comprehensive tooling to support CX switch configuration using Ansible.

## Overview

Ansible is an open-source orchestration framework maintained by Red Hat. It automates provisioning, configuration management, and application deployment.

An Ansible playbook automates CX switches using the AOS-CX Ansible Collection that configures switches using multiple REST API calls and CLI commands via SSH.

The Ansible workflow in this guide provides turnkey automation of an AOS-CX Two-Tier Data Center. The zipped version of the project can be downloaded from the Github repository into your Ansible control machine using the following **git clone** command:

```
$ git clone https://github.com/aruba/aoscx-ansible-dcn-workflows.git
```

**NOTE:**

HPE Aruba's Getting Started with Ansible and AOS-CX guide provides additional information on how to use the AOS-CX Ansible Collection.

## Ansible Project Prerequisites

This project assumes a working knowledge of Ansible. If you are new to Ansible automation, please review HPE Aruba's Getting Started with Ansible and AOS-CX guide on the Developer Hub.

An automation server or VM in the networking environment with SSH reachability to the IP addresses assigned to the Aruba CX out-of-band management interfaces is required.

The Ansible control machine requires **Python3.5+** and **Ansible 2.13.1+**, which can be installed using Ansible's Installing Ansible guide.

This project requires HPE Aruba's AOS-CX Ansible Collection, which can be installed by executing the **ansible-galaxy** command using the requirements.yml file in the HPE Ansible data center repository.

```
$ cd aoscx-ansible-dcn-workflows
$ ansible-galaxy install -r requirements.yml
```

The following Python libraries are required for this project.

- **Jinja2** 2.10+

- **paramiko** 2.1.1+
- **pip** 6.0+
- **requests** 2.2.0+
- **netaddr** 0.7.5+
- **pyaoscx** 2.5.1+
- **openpyxl**

The Python libraries can be installed by running **pip** with the requirements.txt file in the HPE Ansible data center repository.

```
$ cd aoscx-ansible-dcn-workflows
$ pip install -r requirements.txt
```

## Ansible Project Structure

The Ansible files for the Two-Tier Data Center project are maintained in a general data center workflow repository on Github. The repository also contains files for additional projects. Visit the AOS-CX Data Center Automation with Ansible Developer Hub for details on the other workflows hosted within the data center repository.

The files necessary for the Two-Tier Data Center workflow are listed in the repository structure below.

```
configs                                 # Directory for generated configurations
|- sample_configs                       # Sample Final Configurations for all
    workflows

templates                               # Place to hold Jinja templates for config
    generation
|- 2TierV2                              # Jinja2 configuration Templates for Two-
    Tier DCN V2
|   |- access.j2                        # Access switch Jinja2 template for
    Architecture II version 2
|   |- core.j2                          # Core switch Jinja2 template for
    Architecture II version 2

deploy_2tierv2_dcn.yml                  # Playbook for Architecture II version 2
inventory_2tierv2_dcn.yml               # Inventory for Architecture II version 2
requirements.txt                        # Python library requirements for project
requirements.yml                        # Galaxy collection requirements for
    project
```

The inventory and template files are critical for running the Two-Tier Data Center playbook, as described in separate chapters in this guide.

## Two-Tier Data Center Topology

The Ansible workflow deploys the same sample topology used in the Aruba Central Two-Tier Data Center guide. Some information from the Aruba Central guide is repeated here for reference and readability.

HPE Aruba Two-Tier data centers meet the requirements of small- and medium-size data centers. For network resiliency, multi-chassis link aggregations (MC-LAGs) are used at both switch tiers. The diagram below summarizes the physical topology configured in this deployment guide and the relationship between components.

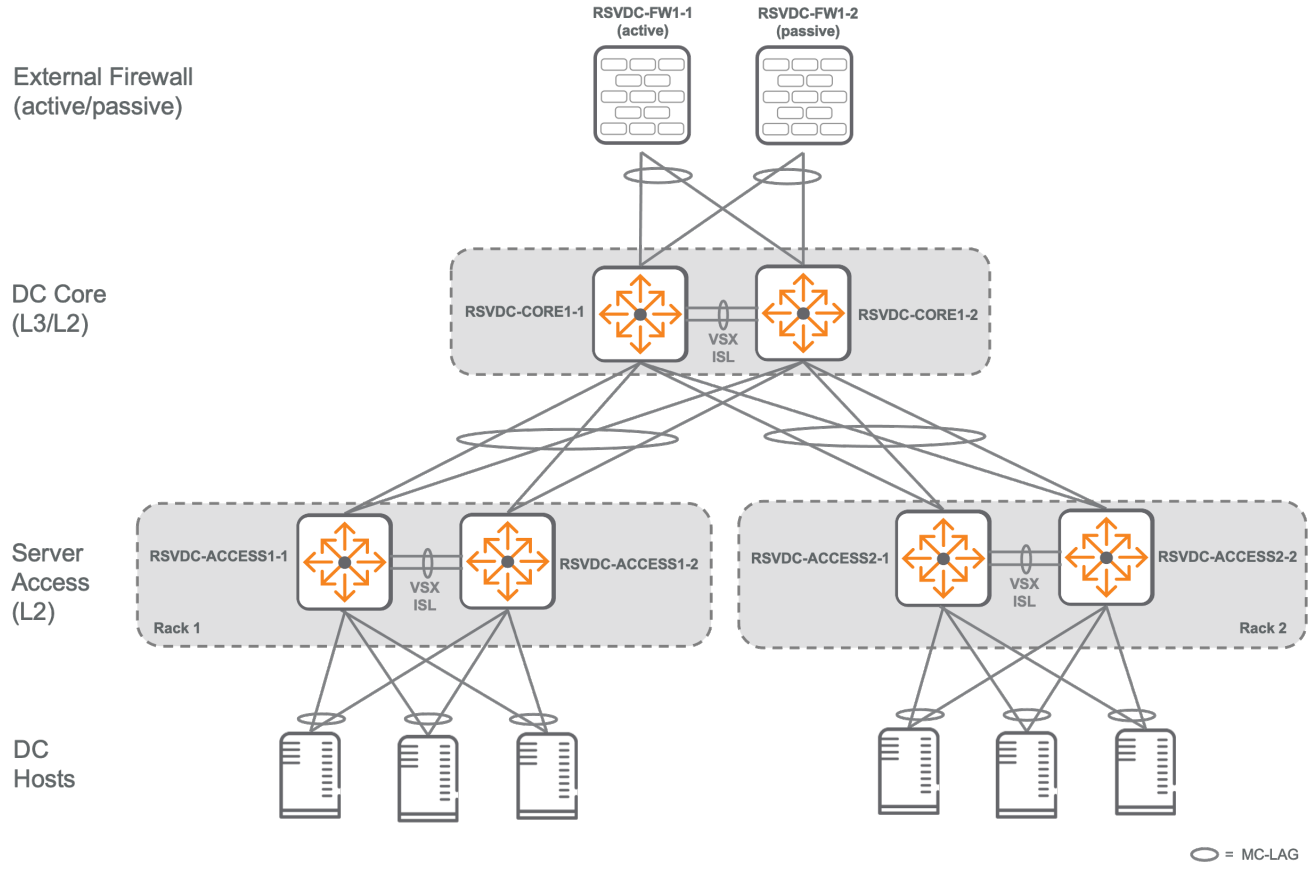


## Two-Tier Core Layer

The core layer provides redundant Layer 2 connectivity to downstream access switches. A VSX pair of core switches is configured with an MC-LAG to each downstream rack. All links from the core layer to the access layer for a single rack are members of the same MC-LAG, whether the rack is populated with a single switch or with a VSX-pair of access switches. MC-LAG provides network resiliency and load-balancing. It also mitigates the need for loop avoidance mechanisms between the core and access layer switches.

Layer 3 services for the data center are provided by the core layer. VLAN switched virtual interfaces (SVIs) define data center subnets, and Aruba Active Gateway provides redundant IP gateways to data center hosts. The core layer also provides redundant IP connectivity to upstream external networks. Typically, firewalls are placed between a data center and external networks for policy enforcement. The redundancy strategies between the data center core and external networks can vary, depending on device features and organizational requirements. In this guide, a traditional active/passive redundant pair of firewalls is connected to the core switch pair using MC-LAGs.

## Two-Tier Access Layer

The access layer provides Layer 2 connectivity to downstream data center hosts.

When a single access switch is at the top-of-rack (ToR) position, the access layer connects to the core layer using a standard LAG. A single ToR switch can provide physical link redundancy using a standard LAG, but host connectivity is lost when performing firmware upgrades or when the ToR switch fails.

This example deployment uses a VSX pair of ToR switches at the access layer, which provides physical switch redundancy to directly attached hosts. This model supports uninterrupted host connectivity, even when one of the ToR switches fails or a firmware upgrade is performed. Each access layer switch also is connected to each core switch using an MC-LAG for redundancy, load balancing, and loop avoidance.

# Ansible Two-Tier Inventory

An Ansible inventory file serves two primary functions for AOS-CX automation: it defines the set of switches to apply automated configuration, and it defines a set of variables that provide configuration values for those switches.

## YAML Inventory File

Ansible inventory files are formatted using YAML data structures. The Two-Tier data center inventory file provides a reference for formatting host, group, and variable data in an Ansible automation.

Red Hat's How to build your inventory guide provides a general reference on creating inventory files.

The example inventory variable values must be modified to align with the organization's environment. This includes customizing the host IP addresses, DNS servers, port assignments, etc. This chapter provides background information and the variable names needed to modify the example inventory file successfully for individual needs.

### Inventory Hosts

Hosts are the targets for an Ansible playbook automation. They are defined in one or more **hosts** sections in the inventory file.

An **ansible_host** variable is assigned to each switch with a fully-qualified DNS name (FQDN), IPv4 address, or IPv6 address value. The playbook uses this value for SSH and API communication with switch management interfaces.

Additional variables are associated with each host that represent the desired configuration state of the target switch. Variable values can be assigned directly to an individual host or inherited through the group structure described in this chapter.

### Inventory Variables

Variables defined in the Ansible inventory file serve two general purposes: they specify switch configuration and Ansible playbook behavior. Most variables in the Ansible inventory file represent values assigned to switch configuration statements when the playbook builds a switch configuration file.

All the variables in the example inventory file are necessary for the Two-Tier data center playbook to run properly.

#### *Single Value Assignments*

Simple variables assign a single value to the right of the variable name with a colon separator.

```
mtu: 9198
vsx_role: primary
vsx_system_mac: 02:00:00:00:10:01
```

### List and Dictionary Assignments

It can be advantageous to assign multiple values to a single variable for automating switch configuration.

Lists can be defined using two methods in the inventory file. The first method assigns a set of comma delimited values in square brackets to a variable name. The second method assigns each list member on a new line preceded by a dash.

The second method provides an easily understandable visual structure for creating dictionaries (multidimensional arrays), where each dictionary member contains its own set of nested variables. A dictionary is analogous to an employee record, where each employee is assigned a set of values such as name, address, and phone number.

| Array Definition Method | Array Syntax |
|---|---|
| Method 1: comma delimited simple list in square brackets | variable_name: [value1, value2, value3] |
| Method 2: simple list with dash-delimited lines | variable_name: - value1 - value2 - value3 |
| Method 2: dictionary entries containinga subset list of variable assignments | dictionary_variable_name: - dict1_var1: dict1_var1_value dict1_var2: dict1_var2_value dict1_var3: dict1_var3_value - dict2_var1: dict2_var1_value dict2_var2: dict2_var2_value dict2_var3: dict2_var3_value - dict3_var1: dict3_var1_value dict3_var2: dict3_var2_value dict3_var3: dict3_var3_value |

Example lists from the inventory file are shown below:

```
ntp_servers: [10.2.120.98, 10.2.120.99]
dns_servers: [10.2.120.98, 10.2.120.99]
vsx_isl_ports: ['1/1/49', '1/1/50']
```

The following example illustrates how to assign sets of variables in dictionary form to the **host_vlans** variable. Grouping sets of values in this manner facilitates AOS-CX config automation in the template file.

```
host_vlans:
- id: 101
  name: PROD-WEB
  ip_address: 10.12.101.2
  active_gateway_mac: 02:00:0a:01:65:01
  active_gateway_ip: 10.12.101.1
- id: 102
  name: PROD-DB
  ip_address: 10.12.102.2
  active_gateway_mac: 02:00:0a:01:65:01
  active_gateway_ip: 10.12.102.1
```

### *AOS-CX Ansible Collection Variables*

AOS-CX Ansible Collection variables are reserved names that modify playbook behavior and connectivity to automated hosts.

Some AOS-CX Ansible Collection variables require a specific value, such as the following connection variables.

```
ansible_connection: arubanetworks.aoscx.aoscx   # DO NOT CHANGE
ansible_network_os: arubanetworks.aoscx.aoscx   # DO NOT CHANGE
ansible_httpapi_use_ssl: True   # DO NOT CHANGE
```

## Inventory Groups

Groups provide a hierarchical structure for the inventory file. Ansible hosts and variables can be defined at any group level in the hierarchy.

In the example deployment, six groups are defined to organize the inventory file into a four-level hierarchy.

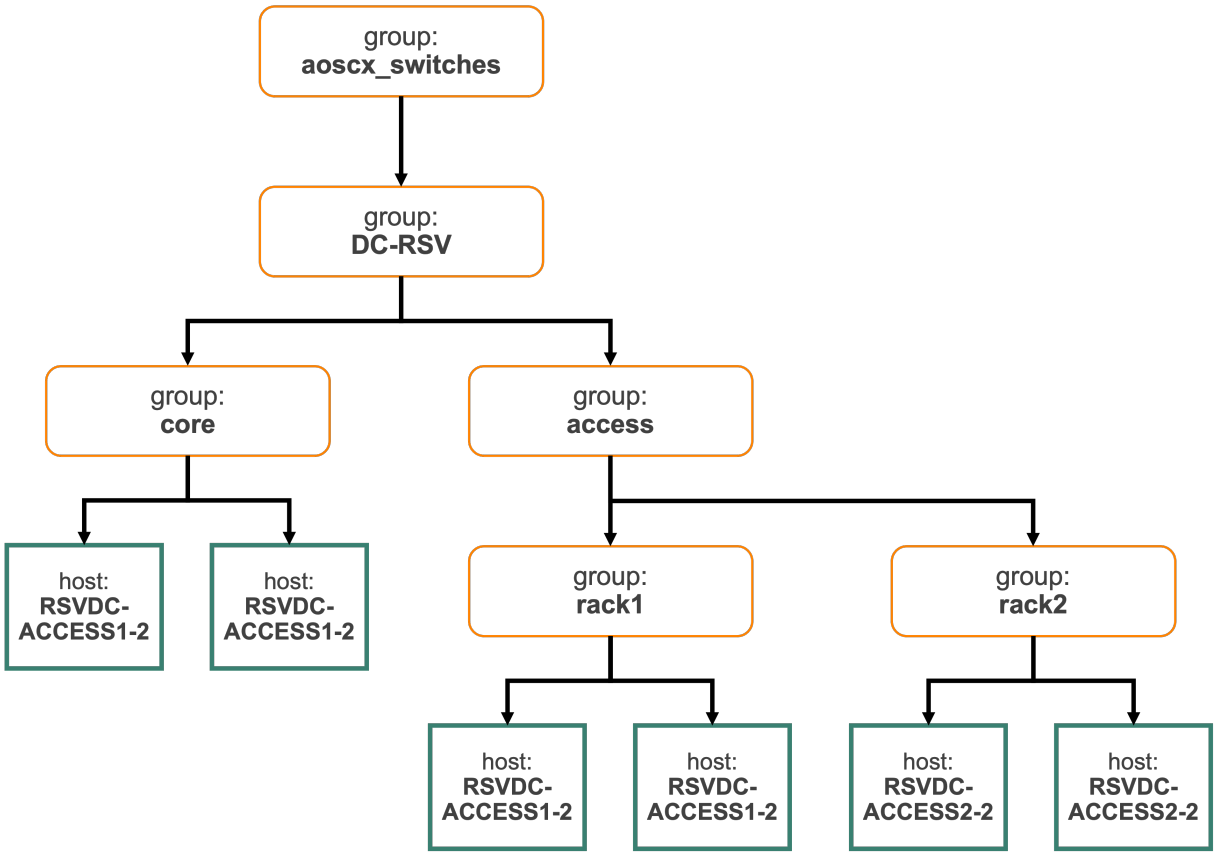


Groups enable efficient organization of common configuration elements, similar to Aruba Central. Each host inherits the variables assigned to groups in its path hierarchy. This enables assigning common configuration elements on multiple affinities, such as data center location, functional role, or rack.

Core switch hosts are defined in the **core** group. Server access switches are defined in their respective rack group to accommodate configuration common to each VSX pair. For example, server access switches in rack 1 are defined in the **rack1** group.

AOS-CX Ansible Collection variables that apply to all switches are defined once in the **aoscx_switches** group. Configuration values common to all data center switches are defined once in the **DC-RSV** group. The **core** and **access** groups assign configuration common to switches in those respective roles. Configuration values common to access switches in the same rack are defined at the rack level. The rack groups are defined primarily to apply VSX and MC-LAG configuration to redundant top-of-rack switches. Variables that have unique values at each host are defined at the host level for both core and access switches.

The following diagram illustrates a sample variable at each group level inherited by a host and a unique host variable assignment:

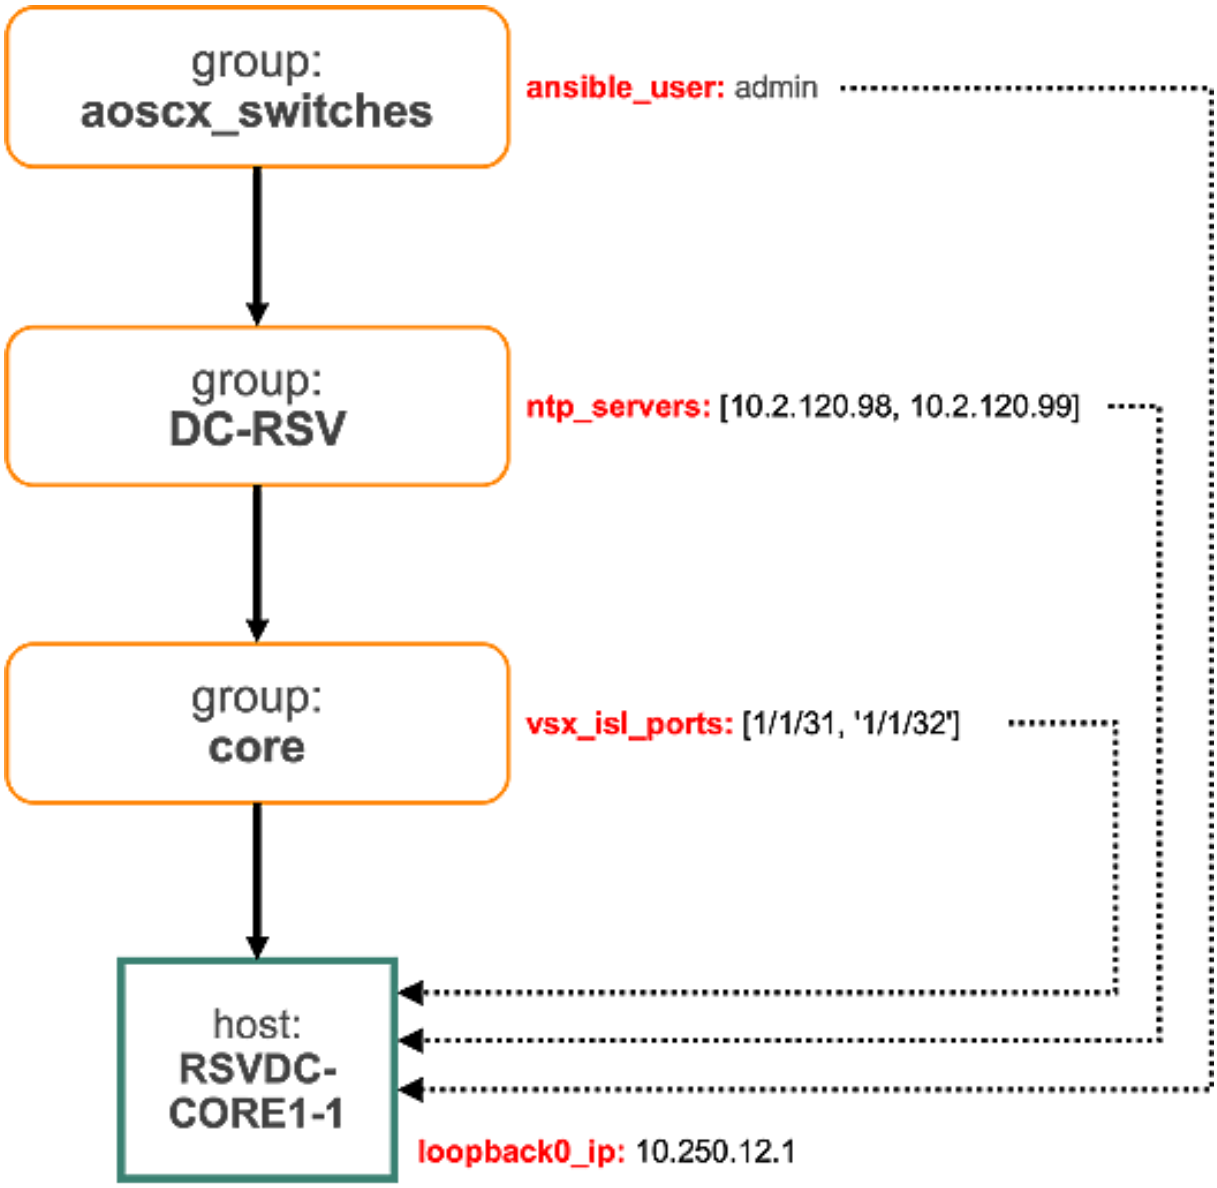


When the same variable exists in multiple contexts, the assignment nearest the host in the hierarchy is used by the Ansible playbook. A host variable assignment takes precedence over all group assignments.

The example Ansible group organization is just one approach for categorizing switches with common configuration elements. Organizations should use a grouping methodology that best meets their needs.

## Specify a Template File

The playbook uses inventory variables in conjunction with a configuration template to build AOS-CX configuration statements.

When automating more than one switch model or role, some configuration elements may be unique to the model or role. Two examples include: - Assigning switch profiles to optimize hardware resources with a switch's network functions. Switch profile names are not consistent across all switch models. - In the Two-Tier data center, OSPF routing is configured on core switches, but not on server access switches.

A separate template file for core and access switches is used in our example deployment. Standardizing a consistent switch platform for each role can simplify template creation, although additional scripting logic can be applied in a template file to accommodate model differences.

It is best practice to use different configuration templates for each switch role for efficient template file administration.

The configuration templates are assigned to a host with the **config_template** inventory variable. In this project, the **config_template** variable is assigned in the **core** and **access** group contexts.

```
config_template: templates/2TierV2/access.j2
```

Configuration templates are described further in the Ansible Template chapter.

## Planning Inventory File Values

This section provides sample values and rationale for naming and numbering schemes. Adjust values and formats as needed to accommodate specific requirements. Use a consistent approach in the physical and logical configurations to improve network management and troubleshooting.

### Naming Conventions

Establish a switch naming convention that indicates the switch type, role, and location to simplify identification and increase operating efficiency.

Example values used in this guide:

| Switch Name | Role | Description |
|---|---|---|
| RSVDC-CORE1-1 | Core | Roseville Data Center Core Switch, VSX Pair Member 1 (primary) |
| RSVDC-CORE1-2 | Core | Roseville Data Center Core Switch, VSX Pair Member 2 (secondary) |
| RSVDC-ACCESS1-1 | Acces | Top-of-Rack Access Switch in Rack 1, VSX Pair Member #1 (primary) |
| RSVDC-ACCESS1-2 | Access | Top-of-Rack Access Switch in Rack #1, VSX Pair Member 2 (secondary) |

In the Ansible inventory file, the switch name is defined as an Ansible inventory name in the **hosts** stanza. For example, the inventory name **RSVDC-ACCESS-1-1** is created in the hosts stanza in the **DC-RSV > access > rack1** group context. Each link in the table above references the location in the inventory file for a specific switch's inventory name.

When executing the playbook, the hostname in the configuration file is set based on the current iterated switch inventory name.

## IP Address Planning

Plan a consistent IP numbering scheme with values that can accommodate the current deployment size and leave room for growth. Define a range that can represent loopback addresses, IP addresses used in supporting protocols, and a range for data center hosts. It is beneficial to assign data center host subnets from a larger range of maskable IP addresses that summarizes all host subnets in the data center.

Example IP address ranges used in this guide:

| Subnet | Functional Description |
|---|---|
| 10.255.12.0/24 | Routed interface IP addresses |
| 10.250.12.0/24 | Loopback IP addresses |
| 10.12.0.0/16 | Summary range of all data center host subnets |
| 10.12.101.0/24 | Example of a specific data center host subnet |

The Ansible inventory statically defines all the IP addresses needed, including loopback addresses and routed interfaces. Each link in the table above references the location in the inventory file of an example IP address assigned from the pool. Some values and variables are repeated for devices. Be sure to review all entries in the inventory file before execution.

## MAC Address Planning

A Locally Administered Address (LAA) should be used when defining virtual MAC addresses for VSX and active gateway functions. This is required when configuring an Active Gateway for an SVI on a VSX pair and when configuring the system MAC address of VSX. An LAA is a MAC in one of the four formats shown below:

```
x2-xx-xx-xx-xx-xx
x6-xx-xx-xx-xx-xx
xA-xx-xx-xx-xx-xx
xE-xx-xx-xx-xx-xx
```

The *x* positions can contain any valid hex value. It is helpful to create a hexadecimal representation of the associated IP address or VLAN ID using the hex positions. For more details on the LAA format, see the IEEE tutorial guide.

In this guide, VSX system MAC addresses are set to 02:00:00:00:10:xx, where xx is replaced with the rack number of the VSX pair and the core switches use a value of 00.

Active Gateway MAC addresses are set in the format 02:00:0a:xx:xx:xx.

The **host_vlans** variable defined in the **RSV-DC** group is redefined for each core switch at the host level. Be mindful to use consistent **vlan_id** and **name** values in the inventory file at both levels.

**NOTE:**

A variable's value may be reassigned at a different level in the inventory file's hierarchy for several reasons. In our example, the dictionary values assigned to the **host_vlans** variable in the **DC-RSV** group are inherited by all access switches. The core switches require additional variables in the dictionary for their roles in the network, where some values are unique to the host. Redefining the **host_vlans** variable at each core switch's host level overwrites the dictionary values defined in the higher group, so it is necessary to include the previously defined **vlan_id** and **name** variables in the new dictionary list assignment at the host level.

## Inventory Variable Reference

This section describes inventory variables included in this example Ansible AOS-CX playbook automation. The example inventory file provides a complete reference for the Two-Tier data center automation.

The top level **aoscx_switches** group assigns variables associated to all switch hosts focused on script execution.

| Variable Name | Value | Notes |
|---|---|---|
| ansible_user | [switch username] | Login user for switch. This value can be changed. It is best practice to define an automation user account. |
| ansible_password | [switch password] | Password for the ansible_user account on the switch. |
| ansible_connection | arubanetworks.ac | This value is used by the AOS-CX Ansible Collection and must not be changed. |
| ansible_network_os | arubanetworks.aoscx.aoscx | This value is used by the AOS-CX Ansible Collection and must not be changed. |

| Variable Name | Value | Notes |
|---|---|---|
| ansible_httpapi_use_ssl | True | This value enables HTTPS communication and should not be changed. |
| ansible_httpapi_validate_certs | False | Change this value to **True** to perform certificate validation using the host's certificate store. |
| ansible_acx_no_proxy | True | |
| ansible_aoscx_validate_certs | False | |
| ansible_aoscx_use_proxy | False | |

The **DC-RSV** group assigns AOS-CX configuration variables common to all data center switches.

| Variable Name | Sample Value | Notes |
|---|---|---|
| hostname | "{{inventory_name}}" | The playbook iterates over each host entry in the inventory file, and the *inventory name* is the name of the host entry currently being iterated. |
| group | RSVDC | This value is referenced by another variable to assign the MST instance name. |
| config_path | "configs/" | The directory that playbook generated configurations will be stored. |
| timezone | "america/los_angeles" | Value used to assign the time zone for all switches at this data center location. |
| mtu | 9198 | Value used to assign the Layer 2 interface MTU. AOS-CX defines the Layer 2 MTU value as the acceptable Layer 3 payload size, not the Ethernet frame size directly. |
| stp_config_name | "{{group}}" | Value used to assign the MST instance name to the value of the **group** variable. |

| Variable Name | Sample Value | Notes |
|---|---|---|
| ntp_servers | [10.2.120.98, 10.2.120.99] | List of NTP servers used to assign multiple NTP servers. |
| ntp_vrf | mgmt | Value used to assign source VRF for NTP queries. |
| dns_servers | [10.2.120.98, 10.2.120.99] | List of DNS servers used to assign multiple DNS servers. |
| dns_domain | example.local | Value used to assign DNS domain. |
| system_location | DC01, Roseville, CA | Value used to assign SNMP server system location |
| system_contact | netadmin@orangetme.local | Value used to assign SNMP server system contact |
| opsf_area | 0.0.0.0 | Value used to assign the OSPF area on core switches. The server access switches do not participate in OSPF. Although the ospf_area inventory variable is associated with access switches by its inclusion in the **DC-RSV** group, it is not applied to server access switch configurations, because it is not referenced by the access switch configuration template when building configuration files. |
| tacacs_servers | | This defines the start of a dictionary containing multiple variables for each TACACS server entry. |
| tacacs_servers.host | 10.2.120.94 | Value used to assign the IP address of an individual TACACS server. Each tacacs_server dictionary entry has its own unique IP assignment. |
| tacacs_server.ciphertext | [shared secret] | Value used to assign the encrypted shared key for an individual TACACs server. |

| Variable Name | Sample Value | Notes |
|---|---|---|
| host_vlans | | This defines the start of a dictionary containing multiple variables for each server VLAN. The host_vlans dictionary defined at the **DC-RSV** level is inherited by server access switches, but it is overwritten by host-level variables for each core switch. |
| host_vlans.id | 101 | Value used to assign VLAN ID for each VLAN in the data center. A unique value is assigned to each dictionary entry. |
| host_vlans.name | PROD-WEB | Value used to assign VLAN name. A unique value is assigned to each dictionary entry. |
| vsx_isl_lagid | 256 | Value used to assign the LAG ID used as the ISL in a VSX pair. |
| vsx_keepalive_vrf | mgmt | Value used to assign the VRF used for VSX keepalive messages. |

The **core** group assigns AOS-CX configuration variables common to data center core switches.

| Variable Name | Sample Value | Notes |
|---|---|---|
| config_template | templates/2TierV | Assigns the core switch configuration template to both core switches. The playbook uses the configuration template and inventory variables in combination to generate switch configuration files. |
| vsx_system_mac | 02:00:00:00:10:00 | Value used to assign VSX system MAC address to both VSX core switches. |

| Variable Name | Sample Value | Notes |
|---|---|---|
| vsx_keepalive_ip_primary | 172.16.117.101 | Value used to identify the VSX primary switch IP address for VSX keepalive configuration, but not assignment to an interface. The switches used in this deployment example use DHCP to obtain a **mgmt** interface IP assignment. The IP address assignments for the VSX primary and secondary switches are identified through the use of DHCP reservations. |
| vsx_keepalive_ip_secondary | 172.16.117.102 | Value used to identify the VSX secondary switch IP address for VSX keepalive configuration. |
| vsx_isl_ports | [1/1/31, “1/1/32”] | Value used to identify the physical interfaces assigned to the VSX ISL. The same interfaces are used on both switches. |
| mclags | | This defines the start of a dictionary containing multiple variables for each switch VSX/MC-LAG. |
| mclags.id | 1 | Value used to assign MC-LAG interface IDs. A unique value is assigned to each dictionary entry. |
| mclags.interfaces | [“1/1/1”, “1/1/2”] | Value used to assign physical interfaces for a specific MC-LAG ID. Unique values are assigned to each dictionary entry. |
| mclags.allowed_vlans | [“all”] | Value used to assign VLAN IDs for a specific MC-LAG ID. This example is a simple list containing only one element. The **all** keyword allows all VLANs between core and access switches. |
| mclags.mtu | 9198 | Value used to assign the Layer 2 MTU for MC-LAG member interfaces. It is best practice to set the Layer 2 MTU to its highest value. |

| Variable Name | Sample Value | Notes |
|---|---|---|
| mclags.description | RACK-1 | Value used to assign the description for a specific MC-LAG ID. Unique values are assigned to each dictionary entry. |

Each core switch is assigned a set of variables at the host level. The variables and lists assigned at the host level contain unique values not shared by other switches. The following table shows sample values for the RSVDC-CORE1-1 switch.

| Variable Name | Sample Values | Notes |
|---|---|---|
| ansible_host | 172.16.117.101 | Value used by the AOS-CX Ansible playbook to connect to this specific switch host. |
| vsx_role | primary | Value used to assign the VSX role. The template file uses this variable in a conditional statement to generate proper VSX keepalive IP configuration. |
| loopback0_ip | 10.250.12.1 | Value used to assign an IP address to the loopback0 interface. |
| routing_vlans | | This defines the start of a dictionary for VLANs used to route to external networks. In this example deployment, only a single external routing VLAN is defined with a single dictionary entry. This dictionary is assigned at the host level for each core switch, because unique IP address assignments exist per host that cannot be inherited from a higher level group. |
| routing_vlans.id | 4000 | Value used to assign VLAN ID for the VLAN used to connect to external networks. |
| routing_vlans.name | CORE-ROUTING-SVI | Value used to assign VLAN name. |
| routing_vlans.ip_address | 10.255.12.1 | Value used to assign a unique VLAN SVI for this specific switch. |

| Variable Name | Sample Values | Notes |
|---|---|---|
| host_vlans | | This defines the start of a dictionary containing multiple variables for each server VLAN. This dictionary is assigned at the host level for each core switch, because unique IP address assignments exist per host that cannot be inherited from a higher level group. |
| host_vlans.id | 101 | Value used to assign VLAN ID for each VLAN in the data center. A unique value is assigned to each dictionary entry. |
| host_vlans.name | PROD-WEB | Value used to assign VLAN name. A unique value is assigned to each dictionary entry. |
| host_vlans.ip_address | 10.12.101.2 | Value used to assign unique VLAN SVI for this specific switch. |
| host_vlans.active_gateway_mac | 02:00:0a:01:65:01 | Value used to assign the Active Gateway MAC address shared by the VLAN on both VSX members. |
| host_vlans.active_gateway_ip | 10.12.101.1 | Value used to assign the Active Gateway IP address shared by the VLAN on both VSX members. |

The **access** group assigns AOS-CX configuration variables common to data center access switches.

| Variable Name | Sample Values | Notes |
|---|---|---|
| config_template | templates/2TierV2, | Assigns the access switch configuration template to all server access switches. The playbook uses the configuration template and inventory variables in combination to generate switch configuration files. It may be necessary to assign the **config_template** variable at the host level, when differences between hardware models do not support the same features. |

| Variable Name | Sample Values | Notes |
|---|---|---|
| speed_interface_group_10g | [1] | A list of switch interface groups to set to 10 Gbps operation. The example list contains only a single element. |
| vsx_isl_ports | [“1/1/49”, “1/1/50”] | Value used to identify the physical interfaces assigned to the VSX ISL. The same interfaces are used on both switches. The same VSX interfaces can be assigned to all server access switches, because the same switch model and port layout are used for all access switches. |

A rack group assigns AOS-CX configuration variables common to a VSX redundant pair of rack switches. In this example deployment, MC-LAGS are defined at the rack level to accommodate unique description and VLAN assignments per LAG. It is possible to define MC-LAGs for all host-facing ports just once in the **access** group, when assigning all VLANs to the MC-LAGs without unique descriptions.

| Variable Name | Sample Values | Notes |
|---|---|---|
| vsx_system_mac | 02:00:00:00: | Value used to assign VSX system MAC address to both VSX core switches. |
| vsx_keepalive_ip_primary | 172.16.117.103 | Value used to identify the VSX primary switch IP address for VSX keepalive configuration, but not assignment to an interface. The switches used in this deployment example use DHCP to obtain a **mgmt** interface IP assignment. The IP address assignments for the VSX primary and secondary switches in all redundant top-of-rack server access pairs are identified through the use of DHCP reservations. |
| vsx_keepalive_ip_secondary | 172.16.117.10⁴ | Value used to identify the VSX secondary switch IP address for VSX keepalive configuration. |
| mclags | | This defines the start of a dictionary containing multiple variables for each server access switch VSX/MC-LAG to physical servers. |

| Variable Name | Sample Values | Notes |
|---|---|---|
| mclags.id | 1 | Value used to assign MC-LAG interface IDs. A unique value is assigned to each dictionary entry. |
| mclags.interfaces | [1/1/1] | Value used to assign physical interfaces for a specific MC-LAG ID. Unique values are assigned to each dictionary entry. |
| mclags.allowed_vlans | [101, 102] | Value used to assign VLAN IDs for a specific MC-LAG ID. This example is a simple list containing two allowed VLANs. |
| mclags.mtu | 9198 | Value used to assign the Layer 2 MTU for MC-LAG member interfaces. It is best practice to set the Layer 2 MTU to its highest value. |
| mclags.description | ESXi-01 | Value used to assign the description for a specific MC-LAG ID. Unique values are assigned to each dictionary entry. |

Each access switch is assigned a set of variables at the host level. The variables and lists assigned at the host level contain unique values that are not shared by other switches. The following table shows sample values for the RSVDC-ACCESS1-1 switch.

| Variable Name | Sample Values | Notes |
|---|---|---|
| ansible_host | 172.16.117.10 | Value used by the AOS-CX Ansible playbook to connect to this specific switch host. |
| vsx_role | primary | Value used to assign VSX role. The template file uses this variable in a conditional statement to generate proper VSX keepalive IP configuration. |

# Ansible Two-Tier Template

When executing an AOS-CX Ansible playbook, Jinja2-based template files are used to create AOS-CX switch configurations.

## Overview

Template files are called by the playbook to build a complete AOS-CX configuration file for each target switch. The template file's Jinja2 syntax in combination with variables defined in the inventory file support the building of dynamic and complex switch configurations.

Multiple template files can be used to build configurations based on switch roles or other criteria. The playbook selects the template assigned to a switch host in the inventory file using the **config_template** variable's value. The variable can be assigned directly to a host or it can be inherited from a group.

## Template Syntax

Jinja2 syntax supports the use of variable replacement values, loops, and conditional statements.

This chapter describes basic formatting of the Jinja2-based template file. The sample template can be modified or used as a reference to build new templates.

The core and access template files in this example deployment can be used unmodified when deploying CX 8325-32C core switches and any 83xx-series CX model as access switches.

### Standard CLI Statements

Any standard AOS-CX switch configuration statement can be added to a template file. The configuration line is copied to the switch exactly as it appears in the file. Strict adherence to AOS-CX syntax is required, and it is best practice to use a full and complete configuration statement. Configuration statements should be placed in the correct context in the file. AOS-CX indentation uses four spaces.

The following example uses standard AOS-CX CLI configuration statements in a template file:

```
profile l3-agg
router pim
    enable
    active-active
https-server vrf mgmt
```

### Simple Variable Substitution

The template file can substitute a single variable value in place of a static value in a configuration statement. Variable values are assigned in the inventory file. To use a variable's value in the template file, place the inventory file's variable name inside two curly braces: **{{variable_name}}**.

By substituting static values with a variable, the same template file can be used to configure multiple switches that require different values in config statements. In addition to applying unique values such as IP addresses, variables enable the same template to be used across different administrative and geographic boundaries. For example, inventory variables can be defined for timezones, SNMP information, DNS servers, NTP servers, and other configuration components that may have different values based on location.

The flexibility of a template file is increased as you increase the number of variables. Each administrator must balance the additional complexity of inventory and template files when increasing the number of variables for added flexibility.

Variable values can be assigned directly to a host in the inventory file, or they can be inherited from a parent group. The following example of AOS-CX configuration statements uses simple variable substitutions in a template file.

hostname {{hostname}} clock timezone {{timezone}} interface lag {{vsx_isl_lagid}}

**NOTE:**

A playbook error will occur if a template references a variable name that is not defined in the inventory file.

## Iteration Over a List of Values

When an inventory variable is assigned a list of values, the template file can iterate over each member of the list using a **for** loop to generate multiple AOS-CX configuration lines.

Iteration over a simple list is helpful when generating multiple lines of configuration that contain a single modified value. For example, generating multiple AOS-CX configuration statements to define a set of NTP or DNS servers. Iteration over a dictionary enables more complex configuration, such as creating a set of VLAN interfaces, where each interface requires a set of additional information to build the configuration, such as VLAN ID, IP address and, active gateway assignments.

Examples of a simple list and dictionary-based iteration are provided below.

### *Example 1: Iteration Over a Simple List*

The example inventory file defines a simple comma-delimited list for DNS servers at the **DC-RSV** group level as follows:

```
dns_servers: [10.2.120.98, 10.2.120.99]
```

The playbook calls upon the template file which iterates over this list and generates the proper AOS-CX config using the following syntax:

p{% for server in dns_servers %} ip dns server-address {{server}} vrf mgmt {% endfor %} p

In the above example, the **for** statement reads one list member at a time from the **dns_servers** variable, and assigns it to the local **server** variable. The **server** variable value is used to complete the AOS-CX configuration statement located between the beginning and end of the for loop, by replacing the variable name with its assigned value on each iteration. This process is repeated until all list members are read and each configuration line is written to the switch config file. In this example, the first list value of **10.2.120.98** is assigned to the **server** variable to generate the first configuration line from the loop. After the second list value of **10.2.120.99** is read and another configuration line is generated, the **for** loop is complete, since no additional list members are present.

The following AOS-CX configuration lines are generated from the above **for** loop:

```
ip dns server-address 10.2.120.98 vrf mgmt
ip dns server-address 10.2.120.99 vrf mgmt
```

### *Example 2: Iteration Over a Dictionary*

The example inventory file defines a dictionary at the RSVDC-CORE1-1 host level to assign VLAN interface values (**DC-RSV > core > RSVDC-CORE1-1**). In the example below, there are two entries in the **host_vlans** dictionary, and each entry contains a list of associated variable values that are required to define a VLAN interface such as **id**, **name**, **ip_address**, etc.

```
host_vlans:
- id: 101
  name: PROD-WEB
  ip_address: 10.12.101.2
  active_gateway_mac: 02:00:0a:01:65:01
  active_gateway_ip: 10.12.101.1
- id: 102
  name: PROD-DB
  ip_address: 10.12.102.2
  active_gateway_mac: 02:00:0a:01:65:01
  active_gateway_ip: 10.12.102.1
```

The playbook calls upon the template file to iterate over the host_vlans dictionary and generate the proper AOS-CX VLAN interface config using the following syntax:

{% for vlan in host_vlans %} interface vlan {{vlan.id}} {% endfor %}

In the above example, the **for** statement reads one dictionary member at a time from the **host_vlans** array. Each array entry is read as a set of variables and assigned to the local **vlan** variable. The individual sub-variable values are referenced by concatenating the parent variable (**vlan**) and one of the sub-variable names (i.e., **id**, **name**, **ip_address**, etc.) with a dot separator. For example, **vlan.name** references the VLAN name value of the current dictionary entry read into the **vlan** variable.

This method supports writing complex configuration, where multiple variable values are associated with a single logical AOS-CX configuration area.

When the first dictionary entry is read, the following variable value assignments are made: - **vlan.id**: *101* - **vlan.name**: *PROD-WEB* - **vlan.ip_address**: *10.12.101.2* - **vlan.active_gateway_mac**: *02:00:0a:01:65:01* - **vlan.active_gateway_ip**: *10.12.101.1*

**NOTE:**

The **for** loop contains additional logic, where if no **vlan.mask** variable is defined, the default value of **24** is substituted. It also contains reference to the **ospf_area** variable that is an inherited value of **0.0.0.0** for all hosts assigned at the **DC-RSV** group level.

The following AOS-CX configuration is generated by the template using the example **for** loop:

```
interface vlan 101
    description PROD-WEB
    ip mtu 9198
    ip address 10.12.101.2/24
    active-gateway ip mac 02:00:0a:01:65:01
    active-gateway ip 10.12.101.1
    ip ospf 1 area 0.0.0.0
    ip igmp enable
    ip pim-sparse enable
interface vlan 102
    description PROD-DB
    ip mtu 9198
    ip address 10.12.102.2/24
    active-gateway ip mac 02:00:0a:01:65:01
    active-gateway ip 10.12.102.1
    ip ospf 1 area 0.0.0.0
    ip igmp enable
    ip pim-sparse enable
```

## Conditional Statements

Jinja2 supports evaluating conditional if/then/else statements, which can be used when generating configuration statements.

The example below contains two conditional statements, a second nested inside the first:

{% if vsx_keepalive_int is defined %} interface {{vsx_keepalive_int}} {% if vsx_role == “primary” %} {% else %} {% endif %} {% endif %}

In the example above, the conditional statement **{% if vsx_keepalive_int is defined %}** checks if the **vsx_keepalive_int** variable is defined for the current host for which the configuration is being generated. If not, the entire stanza above is ignored and no VSX keepalive interface configuration is generated for the switch. If the variable is defined, the configuration is generated using variable substitution.

The second conditional evaluation is designed to assign the correct IP address to the VSX keepalive interface based on the role of the switch in the VSX pair. Both the primary and secondary keepalive IP addresses are defined in the inventory file. If the current switch processed by the playbook is the primary switch, then the **vsx_keepalive_ip_primary** variable value is assigned. If the switch is not the primary, then the **vsx_keepalive_ip_secondary** variable value is assigned.

# Run the Two-Tier Playbook

Execute the following command from the root level of the cloned repository:

```
ansible-playbook deploy_2tierv2_dcn.yml -i inventory_2tierv2_dcn.yml
```

The playbook performs the following actions on every device in the inventory file:

**Step 1** Generate switch configuration files. The switch configurations are based on inventory file variable values and Jinja2 configuration templates assigned in the inventory file.

**NOTE:**

The **core** group in the inventory file assigns the templates/2TierV2/core.j2 template file to the **config_template** variable. Both core switches are configured using this template.The **access** group in the inventory file is assigned the templates/2TierV2/access.j2 template file to the **config_template** variable. All access switches are configured using this template.

**Step 2** Push the generated configurations to each switch using the AOS-CX Ansible SSH module **aoscx_config**.

**Step 3** Enable 10g speed interface groups, if defined in the inventory file. In the example, the access switches have interface groups defined that operate at the non-default value of 10 Gbps .

# Aruba Validated Hardware and Software

The following hardware and software versions are used for the creation of this guide.

Features have been independently validated. Interoperability validation of the combined feature set in this guide is pending.

## EVPN-VXLAN Fabric

### Wired Switching

| Product Name | Software Version |
|---|---|
| Aruba CX 8325 | 10.11.1050 |
| Aruba CX 10000 | 10.11.1050 |
| Aruba CX 8360 | 10.11.1050 |
| Aruba CX 6300 | 10.11.1050 |
| Aruba CX 6400v2 | 10.10.1010 |

### Management and Orchestration

| Product Name | Software Version |
|---|---|
| Aruba Fabric Composer | 6.5.3-13069 |
| Pensando Policy and Services Manager | 1.54.5-T-2 |

## Layer 2 Two-Tier

### Switching Products

| Product Name | Software Version |
|---|---|
| Aruba CX 8325 | 10.13.1000 |
| Aruba CX 8360 | 10.13.1000 |

## Management and Orchestration

| Product Name | Software Version |
| --- | --- |
| Aruba Central | 2.5.8 |


www.arubanetworks.com

3333 Scott Blvd. Santa Clara, CA 95054
1.844.472.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550