aruba
a Hewlett Packard Enterprise company

Validated Solution Guide

ESP DATA CENTER VOLUME 2

Deployment Guide

TABLE OF CONTENTS

About This Guide

This document is from a family of technology guides called Aruba Validated Solution Guides (VSG). VSGs are cross-portfolio solution guides that cover multiple technology areas, including wired, wireless, data center, SD-WAN and security. They are validated by Aruba's Solution TME and Solution Quality Assurance teams on an ongoing basis using a rigorous process. A VSG provides prescriptive guidance focused on the Aruba recommended best practices specific to the solution being covered.

The goal is to describe a solution implementation which addresses the primary use cases for customer networks, while avoiding the corner cases. The intent is to enable partners and customers to efficiently install end-to-end solutions using Aruba technology in a consistent and repeatable manner. The result will be improved stability and supportability by limiting the number of deployment variations.

VSGs are categorized into volumes to differentiate each guide type from the others.

Volumes

- *Design:* Identify products and technologies to meet customer business requirements
- *Deploy:* Step-by-step set of procedures to build the solution
- *Operate:* Recommended procedures to maintain and optimize the solution

Document Conventions

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

Examples:

• the **show vsx status** command

• Navigate to **Configurations > Routing > VRF**

Italic text indicates the definition of important terminology.

Example:

• A *VXLAN Tunnel Endpoint (VTEP)* is the function within leaf switches that handles the origination and termination of point-to-point tunnels forming an overlay network.

Code blocks indicates a variable for which you should substitute a value appropriate for your environment.

Example:

• BGP router id: `10.0.5.1`

Introduction

The Aruba Networks ESP Data Center is built on a technology platform which provides the tools for transforming the data center into a modern, agile, services delivery platform satisfying the requirements of organizations large, small, distributed, and centralized. The Aruba AOS-CX operating system simplifies operations and maintenance with a common switch operating system across the campus, branch, and data center, managed from the cloud or on-premises, and backed by an artificial intelligence capability which provides best practices guidance throughout the operational lifecycle of a network.

Converged ethernet is changing the way compute hosts access storage in the modern data center. Dedicated storage area networks are no longer required. Lossless ethernet and bandwidth management protocols ensure timely reads and writes over a traditional IP LAN. The cost savings and operational simplicity of converged ethernet are major drivers for transformation in the data center today.

At the same time, network topologies have become virtualized. While this virtualization promotes the flexibility required to meet today's transformational data center requirements, it can lead to complexity during implementation and management. The Aruba ESP Data Center mitigates these challenges by leveraging automation in the management plane and capabilities of the Aruba AOS-CX operating system such as automated configuration backups and built-in alerts instrumented on critical network performance metrics.

As you begin the process of designing a new or transformed data center the first step is to understand your organization's cloud applications strategy. This will allow you to determine which applications will remain on-premises and what a right-sized data center looks like for your requirements. When establishing a new data center intended to grow and adapt into the future, plan to implement a spine and leaf underlay supporting software defined overlay networks. The Aruba Networks CX 83xx and 84xx switching platforms provide a best in class suite of products featuring a variety of high throughput port configurations and industry leading operating system modularity providing real-time analytics and always up maintenance.

Purpose of this Guide

This guide covers the Aruba Data Center Network deployment, including reference architectures along with their associated hardware and software components. It contains best practices recommendations for deploying a next generation spine and leaf data center fabric using VXLAN and BGP EVPN leveraging the orchestration capabilities of Aruba Fabric Composer. Please refer to volume one of this VSG for design guidance:

Aruba VSG: 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 can fill a variety of roles:

- Systems engineers who need a standard set of procedures for implementing 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 are changing rapidly. The most pressing challenge is to maintain operational stability and visibility while placing compute and storage resources where they need to be in order to best serve users. In addition, data center teams are being asked to support the rapid pace of DevOps environments including requirements to connect directly with public cloud infrastructure. Given the rapidly changing landscape for data center requirements it is critical that network and system engineers be provided with tools to simplify and automate complex infrastructure configurations.

Initializing the Network

The first step in deploying a data center involves the physical installation of the switches and compute hosts. Please verify the airflow configuration for the products to be installed to make sure 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 the servers inside the rack.

Switch Installation

Prior to installation of the switches, download the Aruba Installation Guide for the specific model being deployed. Review the Installation Guide prior to beginning to install the switches. Review requirements for power, cooling, and mounting, then ensure the required data center infrastructure is available.

Step 1: Open a web browser and navigate to the Aruba Support Portal at https://asp.arubanetworks.com/.

Step 2: From the Support Portal page, select the **Software & Documents** tab.

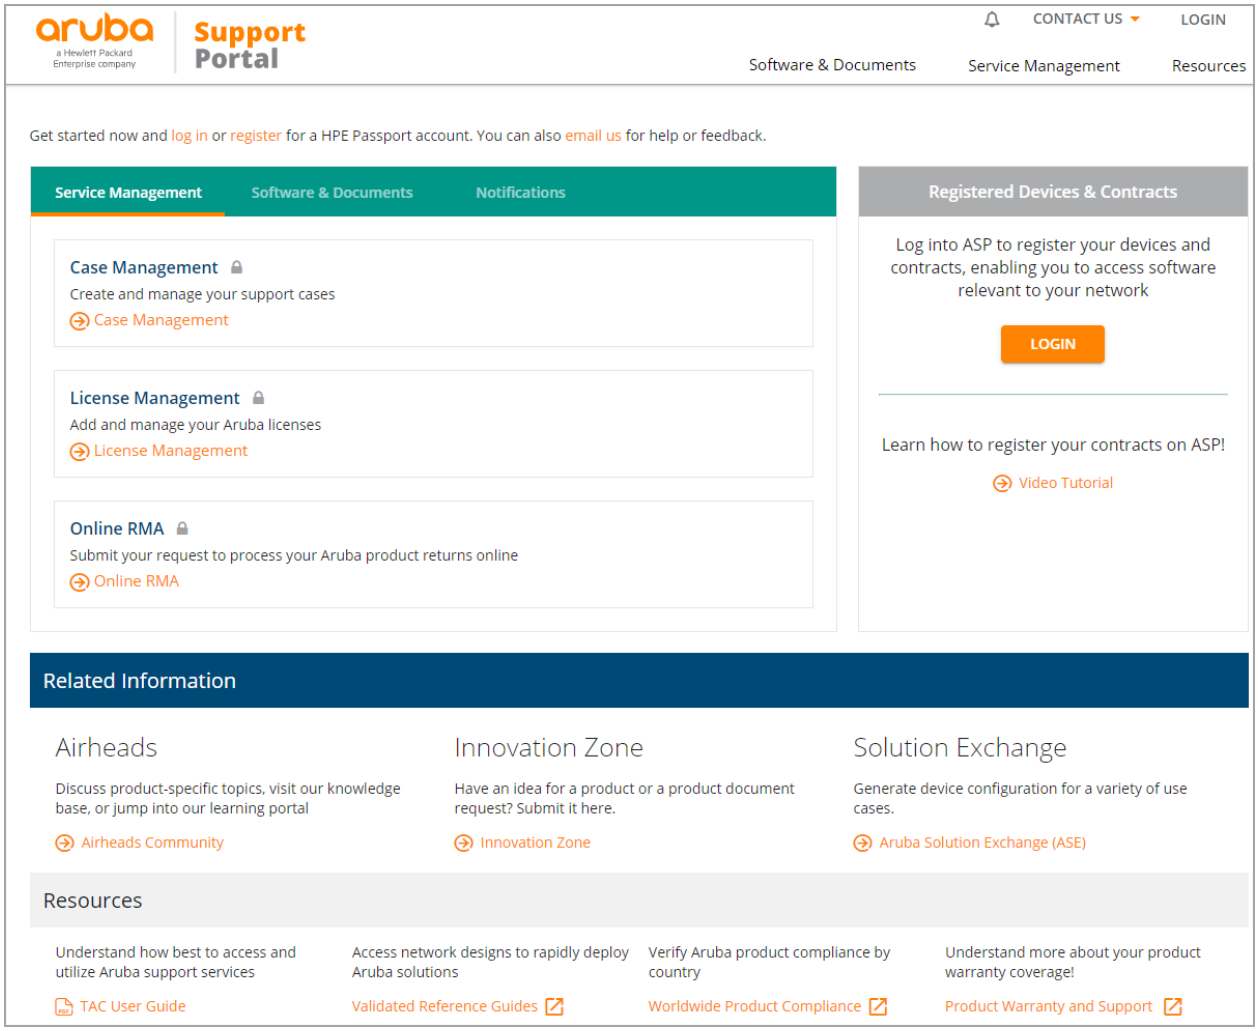

Step 3: From the Software & Documents tab, select **Switches**.

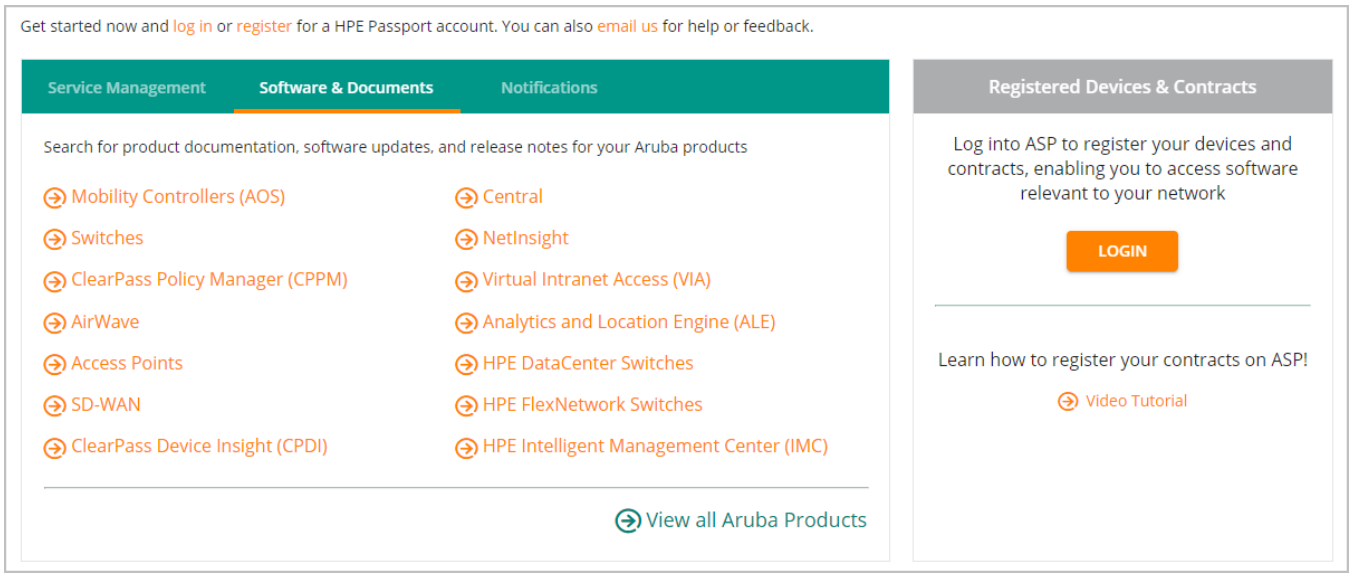

Step 4: Select filter options on the left.

File Type: *Document*

Product: *Aruba Switches*

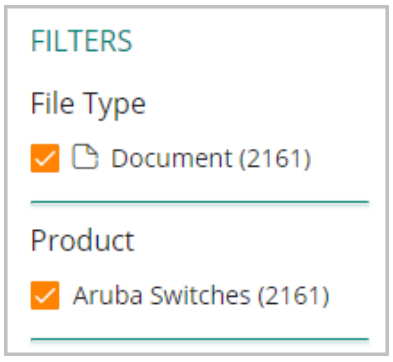

File Category: *Installation Guide*

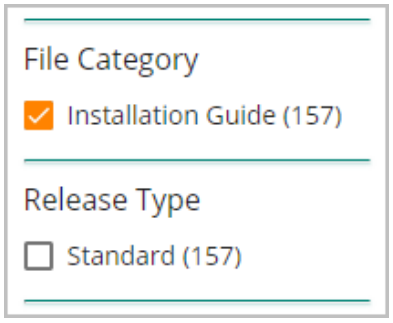

Step 5: Download the Installation Guide version for the switch model being installed.

Step 6: Complete the physical installation of switches into the racks.

NOTE:
Leaf switches should be installed ToR in high-density environments or middle-of-row in low-density environments. Spine switches should be installed at middle-of-row or end-of-row locations depending on cabling requirements and space availability. The key consideration is cable distance and media types used between leaf and spine switches.

Physical Cabling

Consistent port selection across racks and in the spine switches will result in increased ease of configuration management, monitoring, reporting and troubleshooting tasks within the data center. Document all connections. Ensure distance limitations are observed for your preferred host connection media and between switches. Please refer to Volume 1 for guidance related to cabling design options for your installation.

Top of Rack Cabling

The following picture illustrates the port configuration on an 8325 48-port ToR switch.

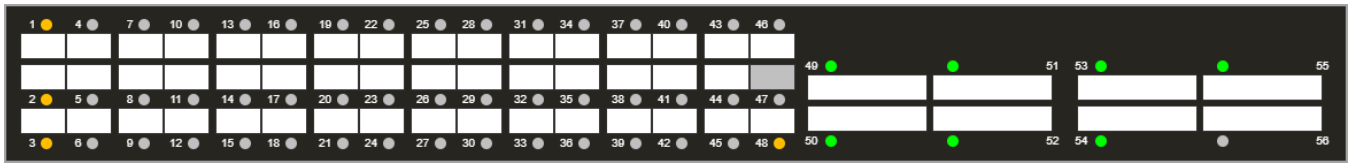

In a high-density, ToR configuration the first set of uplink ports (49-52) should be allocated for interconnecting to a redundant peer ToR switch while the second set of uplink ports (53-56) should be allocated for connecting to spine switches. The number of spine switches should match the number of leaf to spine links required on each ToR (providing a fully meshed, Clos switch topology) and the number of inter-switch links required between ToR peers (providing sufficient bandwidth between peers to accommodate a complete uplink failure on one). This port allocation approach will also ensure consistent port selection in the event that an additional spine layer is added in the future to increase capacity within the fabric.

A similar approach can be followed when using lower density ToR designs. When deploying ToR configurations that require server connectivity at multiple speeds, review the switch guide to determine if doing so will effect adjacent ports. Configuration steps for changing port speeds are covered later in this guide. Please refer to Volume 1 for design guidance regarding port speed groups across the different HW platforms.

Spine to Leaf Cabling

The following picture illustrates the port configuration on an 8325 32-port spine switch.

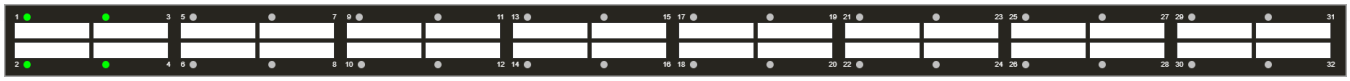

In a dual ToR configuration, a spine switch needs to be connected to each switch in the redundant ToR pair in each rack. A 32 port spine switch will support up to 16 racks in this design. Using the same port number on each spine switch to connect to the same leaf switch (e.g. port 1 of each spine switch connects to leaf switch 1) results in simplified switch management and documentation.

Border Leaf Cabling

In a VXLAN spine and leaf design a pair of leaf switches is the single entry and exit point to the data center. This is called the border leaf but it is not required to be dedicated to that function. Cabling of the border leaf can vary between deployments as it is dependent on how the external network is connected to the data center and if the border leaf contains hosts or service nodes such as firewalls and load balancers.

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

Out-of-Band Management

The Aruba ESP Data Center spine and leaf design strongly recommends a dedicated management LAN for the data center. 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 should also be the network to which Aruba Fabric Composer and Aruba NetEdit are connected. Ensure that the host infrastructure needed for those applications can also be connected to the management LAN.

Deploy management LAN switches at top-of-rack with switch and host management ports connected. Plan for an IP subnet with enough capacity to support all management ports 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. For design assistance, refer to the ESP Data Center Volume 1 Design Guide, as mentioned in the Purpose of this Guide section.

Switch Initialization

Using the Aruba Support Portal at https://asp.arubanetworks.com/ and following the same process as earlier, filter for the AOS-CX Fundamentals Guide which matches the version of the operating system you plan to run.

NOTE:
The version decision is made easier by referring to the operating system release notes and consulting with an Aruba Networks SE or TAC team member.

Under the Initial Configuration section of the Fundamentals Guide, there are 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 switch power supplies.

Step 2: Monitor the console window and observe the boot messages. The console window should look similar to the one below.

Step 3: Confirm that all switches in the fabric are running AOS-CX version 10.06 or later for compatibility with Aruba Fabric Composer.

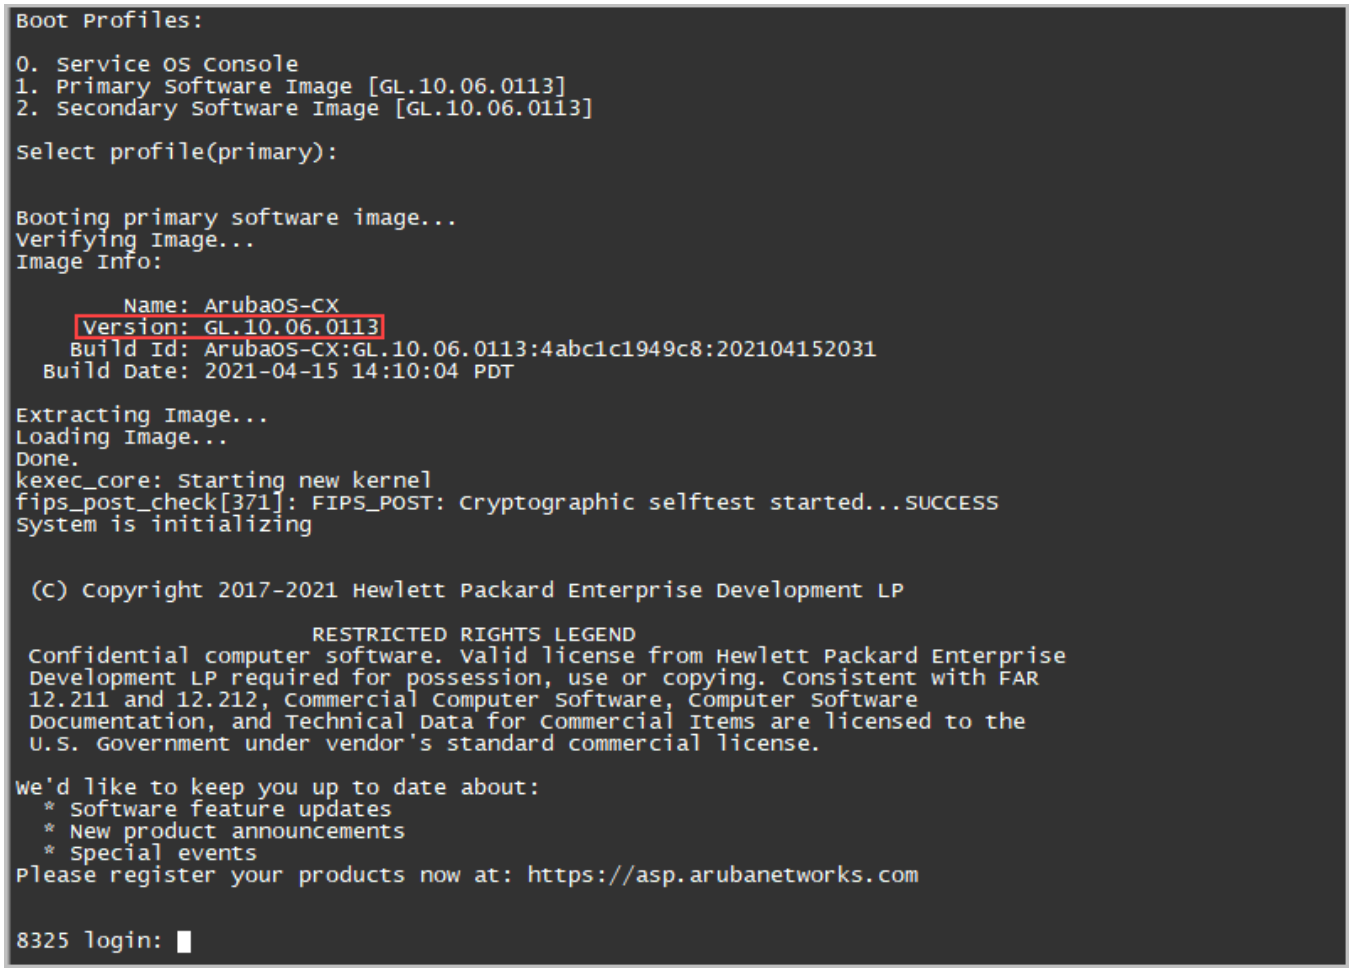

Step 4: Login with the username **admin** and an empty password.

Step 5: Enter a new password for the admin account.

NOTE:
As with the earlier step, under the Initial Configuration section of the Fundamentals Guide, you will find detailed instructions for logging into the switch for the first time.

Step 6: If the switch has been previously configured, reset to factory config. Aruba Fabric Composer requires a factory default configuration for orchestration of 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 the switch hostname.

```
hostname 8325-DC1-L1
```

NOTE:
It is important to use a canonical naming scheme which makes it easier to identify each switch functionally. See the above schemas an example to name the switch in a way that will make it easy to add it to the correct fabric and role when configuring the network from Aruba Fabric Composer.

Step 8: Configure the Switch Management Interface.

```
interface mgmt interface mgmt     no shutdown     ip static 172.16.20.232/24
default-gateway 172.16.20.1     nameserver 172.16.1.98
```

NOTES:
According to your existing IP address management process, determine a subnet to be used for your management LAN. The management LAN is where out-of-band (OOB) management ports on your switches are configured. This is also the LAN on which you should install Aruba Fabric Composer and Aruba NetEdit.
Under the Initial Configuration section of the Fundamentals Guide, you will find detailed instructions for configuring the management interface.

Download Aruba Fabric Composer

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

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

Step 3: From the menu on the left under Product, select **Show All**.

Step 4: In the Product popup, select Aruba Fabric Composer, and then click **Apply**.

Step 5: From the search results, choose the latest OVA version and download it to your computer.

Install Aruba Fabric Composer

On the second page of the Aruba Support Portal search results, find the Aruba Fabric Composer Install Guide. Review the installation considerations to ensure you have adequate host resources available.

Step 1: Select the OVA file using the **Deploy OVF Template** workflow within vCenter.

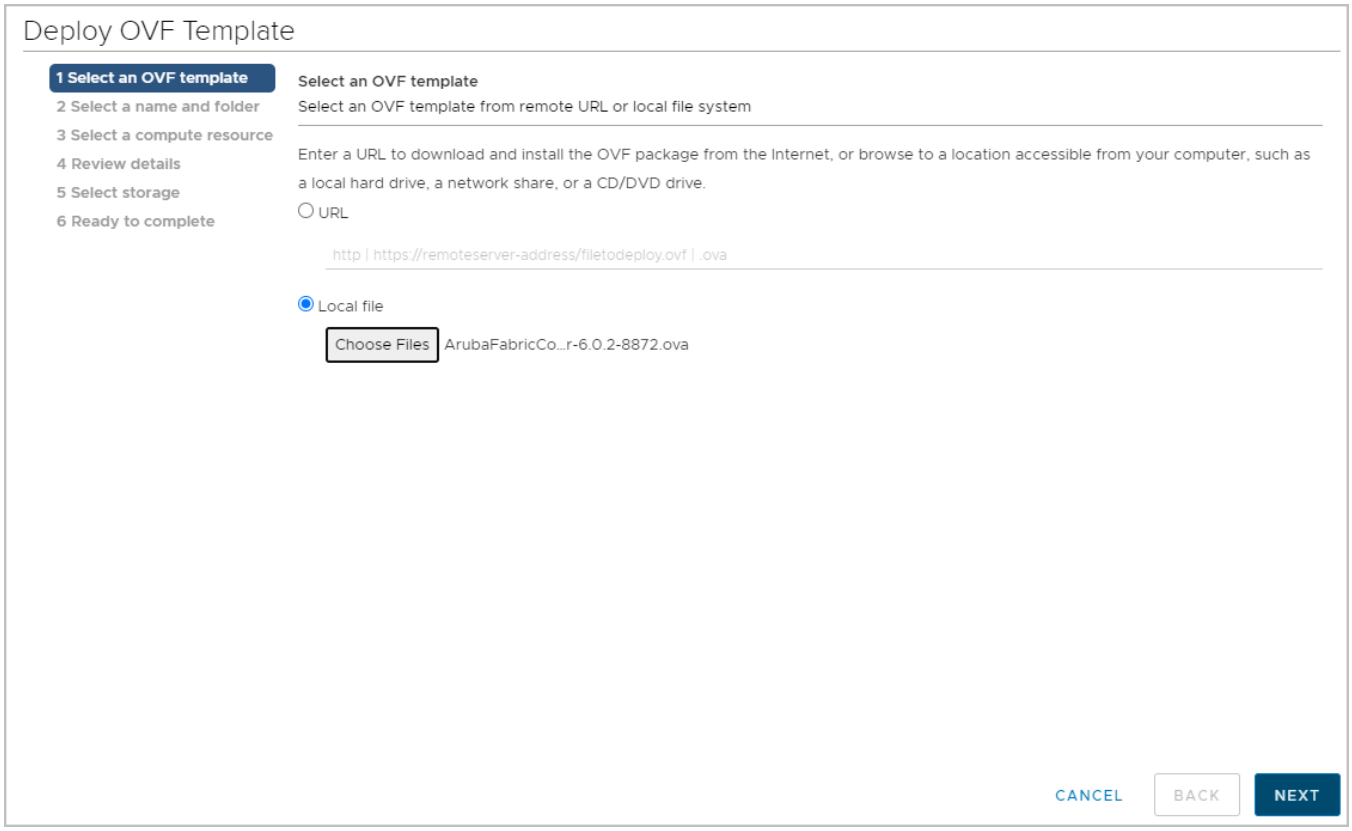

NOTE:
Refer to the Aruba Fabric Composer Release Notes available on the Aruba Support Portal for minimum host requirements.

Steps 2: Proceed through selection of appropriate vSphere resources for your environment and accept the license agreement.

Step 3: Complete the **Customize template** form in the 2nd to last step of the pre-deployment. See below for the types of information you will need to have ready.

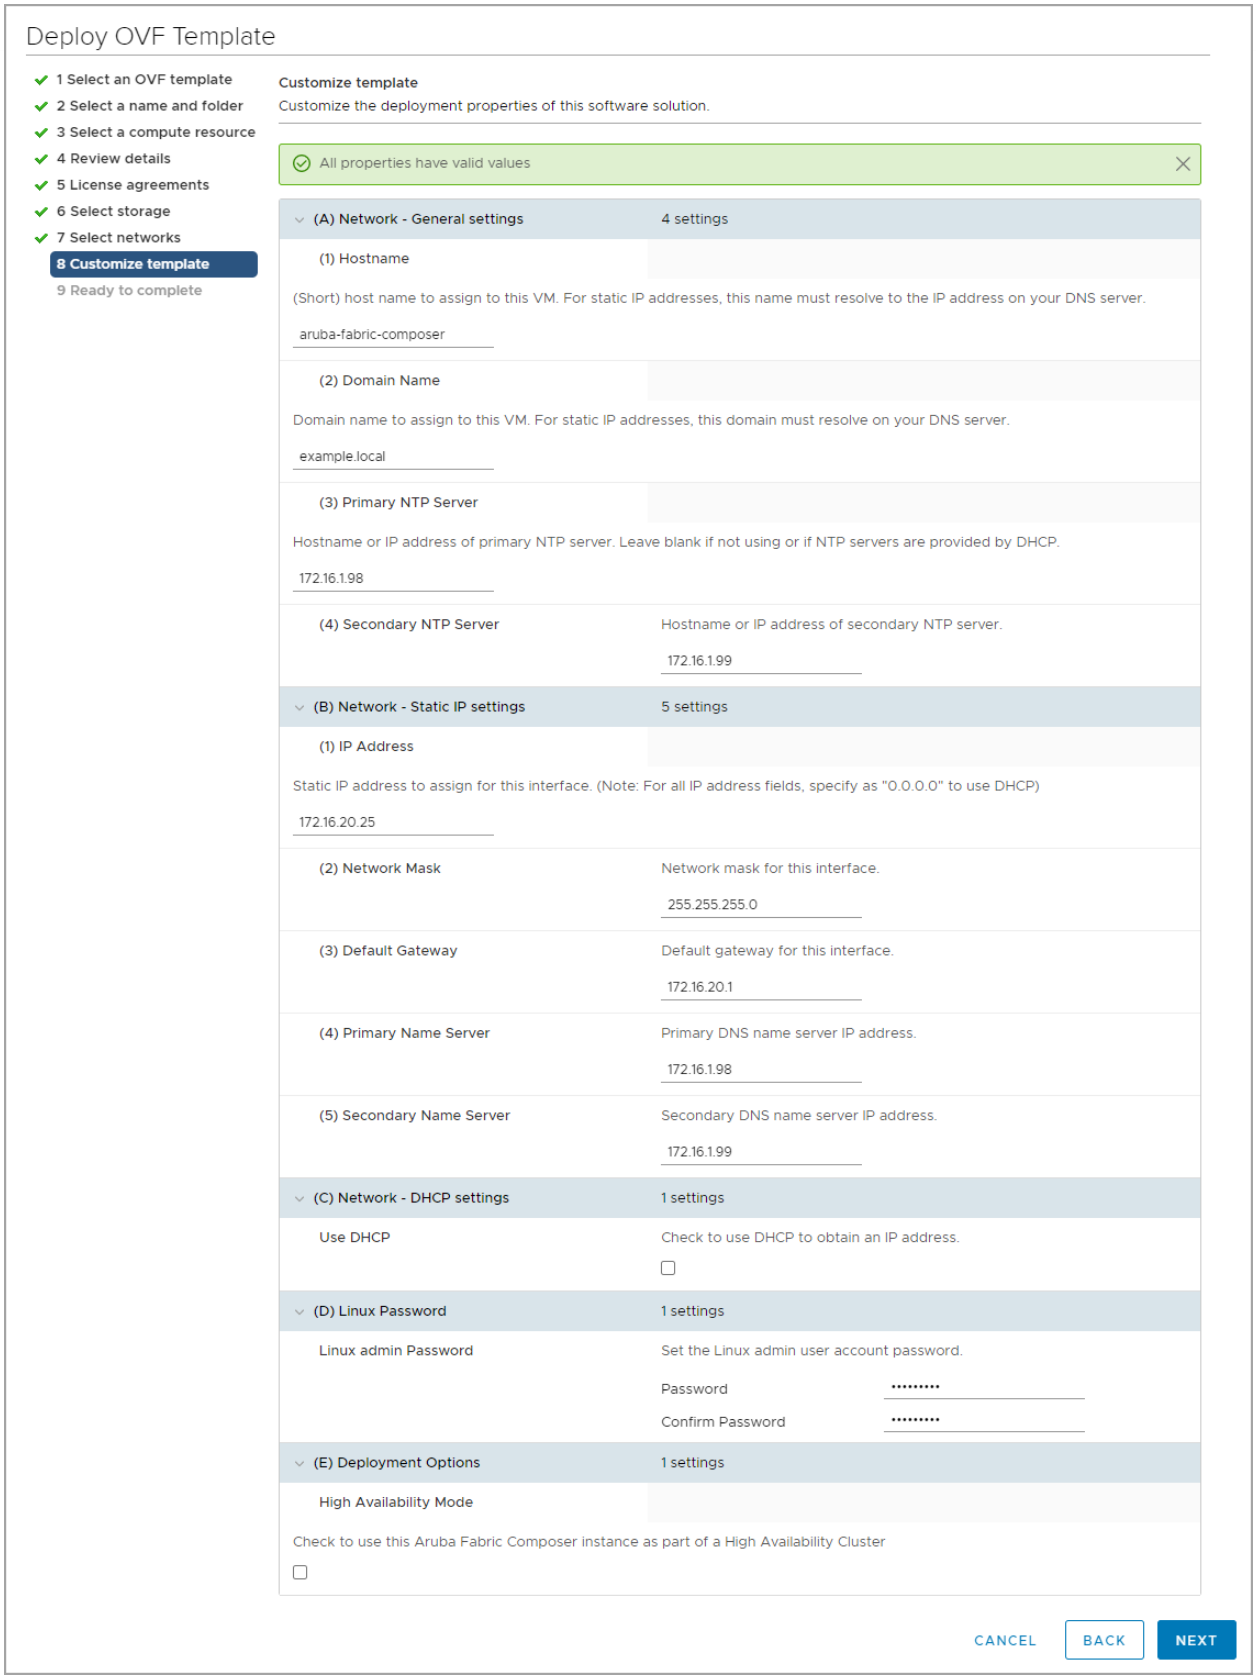

Step 4: Verify all settings and power on the new virtual machine. Wait several minutes for the system to initialize and for the application to become available.

Step 5: Open a web browser and connect to Aruba Fabric Composer at the IP address previously configured.

NOTE:
The software version will not be displayed while the system is initializing and will not allow login during that time.

Step 6: From the Fabric Composer page, enter the following default credentials, and then click **LOGIN**.

Username: *admin*

Password: *aruba*

Step 7: Change your password from the default, and then click **APPLY**.

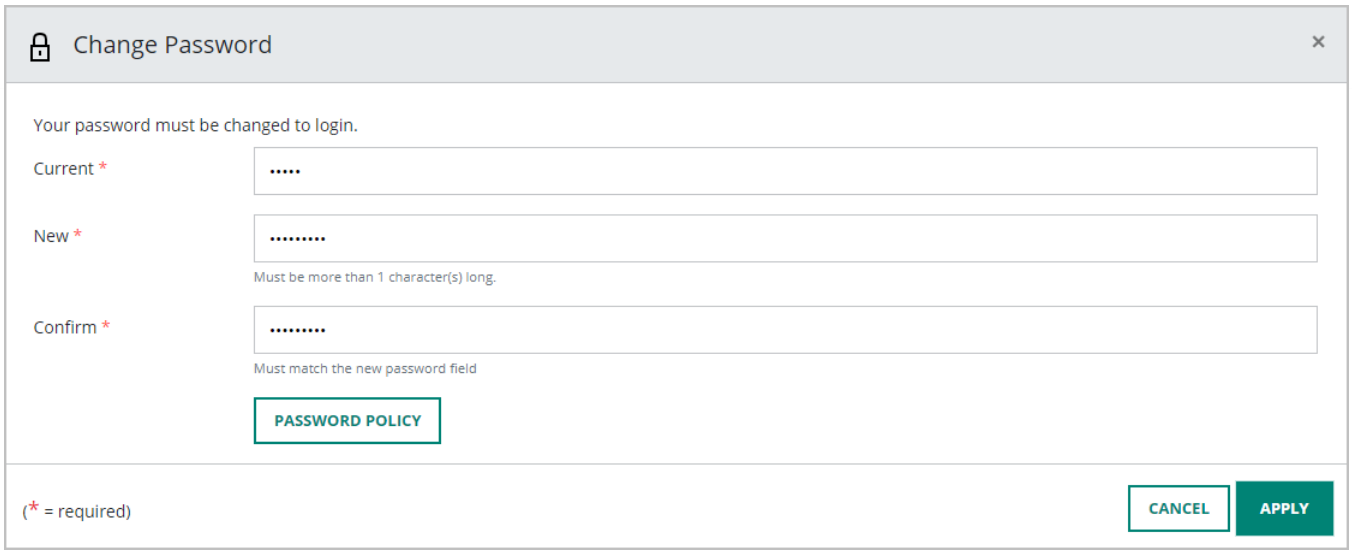

Step 8: Confirm that the Aruba Fabric Composer empty dashboard view is displayed.

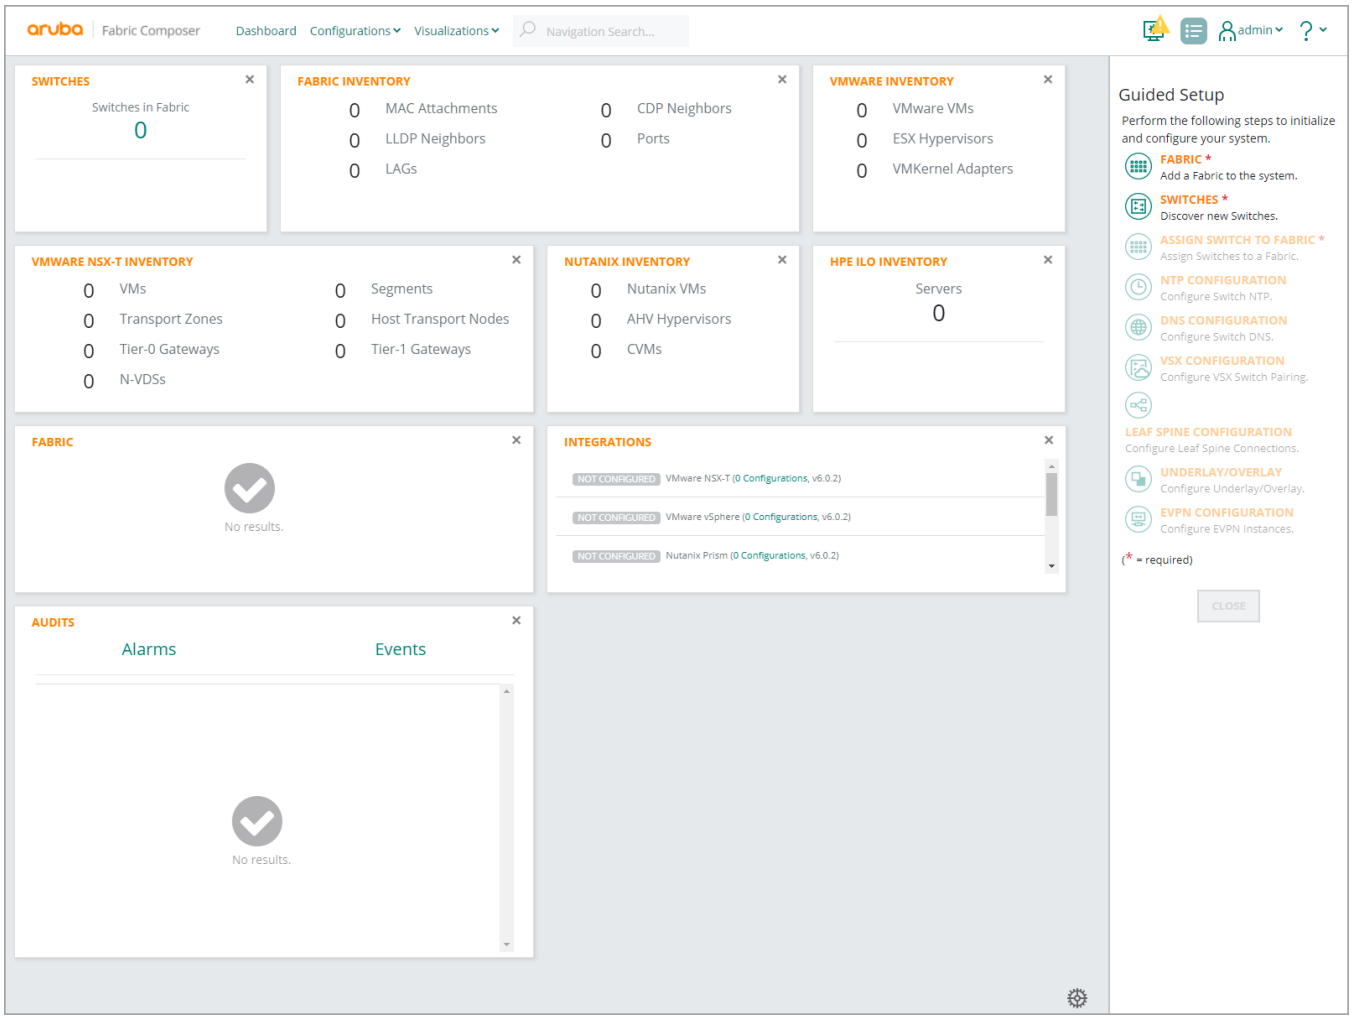

Installing Aruba Fabric Composer for High Availability

Refer to the Aruba Fabric Composer Installation Guide available on the Aruba Support Portal. In the section **Deploying High Availability for Aruba Fabric Composer**, review the installation requirements described and ensure you have adequate host resources available. Follow the steps provided for deployment of the HA cluster.

Deploying the fabric

Configuration of an Aruba ESP Data Center fabric using a spine and leaf topology is best performed using the Aruba Fabric Composer Guided Setup process. Guided Setup will take you through the following steps:

- **Fabric creation** - Definition of the logical construct within Aruba Fabric Composer to identify a fabric.
- **Switch discovery** - Discovery and inventory of switches within Aruba Fabric Composer
- **Switch assignment** - Switch and role assignments within a fabric.
- **NTP & DNS configuration** - Information about NTP and DNS servers for the fabric
- **VSX configuration** - Automated configuration of VSX for redundant ToR, leaf pairs.
- **Leaf/Spine configuration** - Automated configuration of IP links between leaf and spine switches.
- **Underlay/Overlay configuration** - Configuration of the OSPF underlay and BGP overlay.
- **EVPN configuration** - Establishment of EVPN instances for mapping VLANs into overlay VRFs.

Once the Guided Setup is completed, configuration details for compute host onboarding and external fabric connectivity are also required:

- Layer 3 services within overlays.
- Routing between the data center and campus.
- Multi-chassis LACP LAG configuration for host connectivity.

For additional details regarding the Guided Setup steps, find the Aruba Fabric Composer User Guide on the Aruba Support Portal. In the User Guide, refer to the section **Guided Setup > Guided Setup Configuration Options**.

Planning for Deployment

Before starting the guided setup it is important to plan ahead and identify values to be used that will ensure consistent numbering and addressing schemes that will accommodate your deployment size while leaving room for growth. A consistent approach in the physical and logical configurations will improve the management and troubleshooting characteristics of the fabric. This section provides example values and context regarding why those values were chosen. Some values may need to be adjusted to best accommodate the size of the fabric being deployed.

Naming Conventions

The Aruba Fabric Composer supports executing operations on a single switch or a selected group of switches. Establishing a switch naming convention that indicates the switch type, role and location will simplify and increase efficiency when operating production scale fabrics.

Example values used in this guide:

Name	Fabric Role	Description
8325-DC1-L1	Leaf	Fabric #1, Leaf #1 (Rack #1)
8325-DC1-L2	Leaf	Fabric #1, Leaf #2 (Rack #1)
8325-DC1-S1	Spine	Fabric #1, Spine #1

During the Guided Setup you will be prompted to include a Name Prefix for specific steps. These prefixes are logical names inside of Aruba Fabric Composer. Choosing a descriptive name will make it easy to monitor, edit and execute operations. During the detailed procedures you will see example names that can be used in your deployment.

Underlay Connectivity and Addressing

The point-to-point connections between spine and leaf switches are discovered and automatically configured for IP connectivity using /31 subnet addresses from a single network range. A subnet using a /24 mask will support up to 125 links inside a fabric. 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 the event of a host link failure.
- A /31point-to-point interface between ToR switch pairs is used to transmit keepalive messages for VSX peer loss detection.

Each of these subnet types is automatically created by Aruba Fabric Composer from a single network range provided during the VSX setup process. In the case that VSX is not used, the network range will be provided during the underlay configuration process.

Example values used in this guide:

Name	Description	Example
Leaf-Spine IPv4 Subnetwork Address	An IPv4 address block used for creating /31, point-to-point layer 3 links between leaf and spine switches.	172.18.105.0/24
VSX Keep Alive Interfaces IPv4 Subnetwork Address	An IPv4 address block used to allocate loopback addresses (/32) for each switch, for VSX keep-alive point-to-point connection (/31) and also as a transit routed VLAN between redundant ToRs (/31)	10.0.5.0/24

Overlay Connectivity and Addressing

The overlay network is created using VXLAN tunnels that are originated by Virtual Tunnel Endpoints (VTEPs) within the leaf switches in the fabric. A single logical VTEP per rack is required. This is implemented by creating a dedicated /32 loopback interface which is common to both ToR peer switches. The interfaces are automatically assigned from a single subnet scope provided during the Overlay guided setup.

Two BGP Autonomous System Number (ASN) are required for Guided Setup. One is used in the spine layer. The other ASN is used for the leaf layer. A future release of Aruba Fabric Composer will include support for configuring a single ASN across the fabric as described in Volume 1 of this guide.

A Virtual Network Identifier (VNI) is a numerical value used to identify network segments within the overlay topology of the fabric. A VNI is carried in the VXLAN header to enable switches in the fabric to identify which overlay each frame belongs to 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/L3 VNI. Plan your VNI numbering scheme in advance to ensure values do not overlap.

Example values used in this guide:

Name	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

MAC Address Best Practices

A Locally Administered Address (LAA) should be used anytime Aruba Fabric Composer requires a MAC address to be entered. This will primarily be required when configuring the Active Gateway for a Distributed SVI. An LAA is a MAC which looks like one of the four examples 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 be any valid hex value. It is helpful to create a binary representation of the associated VLAN ID using the hex positions. For more details on the LAA format, see the IEEE tutorial guide.

Fabric Initialization

Configuration of an Aruba ESP Data Center fabric using a spine and leaf topology is best performed using the Aruba Fabric Composer Guided Setup process. Any time that you need to return to Guided Setup, simply select it on the menu bar at the top of the Aruba Fabric Composer UI.

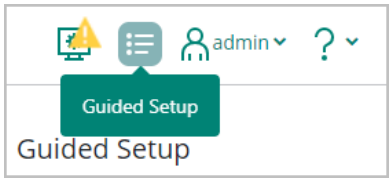

Create a Fabric

The first procedure will create a named fabric within the Aruba Fabric Composer system. The fabric name will be used in later steps as the target of fabric wide configuration settings. This name is internal to Aruba Fabric Composer operations and is not tied to any specific switch configuration. In the 6.0 release of Aruba Fabric Composer only a single fabric can be managed per software instance.

Step 1: At the top of the Guided Setup menu on the right side of the dashboard select **Fabric**.

Step 2: Define a unique logical name for your data center fabric, assign a location, and then click **APPLY**.

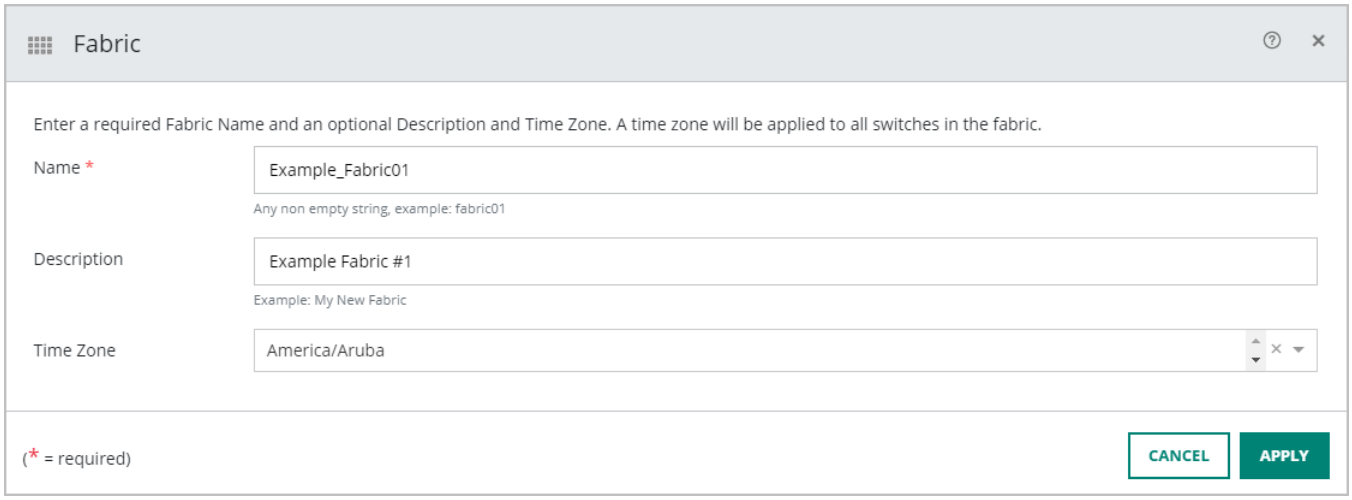

Discover Switches on the Network

Switches to be added to the fabric must first be discovered on the network and added to the device inventory. An orderly naming convention for switch host names should be implemented prior to continuing with this procedure in order to simplify switch selection in the following steps.

Step 1: From the Guided Setup menu, select **Switches**.

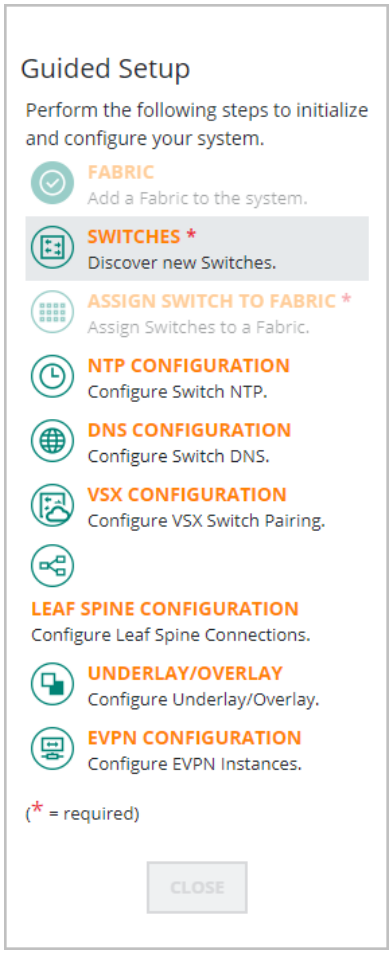

Step 2: From the Discover Switches page, enter the OOBM IP information for the fabric switches. Switch IP addresses can be entered in a comma separated list or in one or more ranges.

CAUTION:
If a range is provided, it must contain only valid switch addresses.

Step 3: Enter the **admin Switch Password** created during switch initialization step then confirm it in the next field.

Step 4: Enter a new password for the **afc_admin Account Password**, confirm it in the next field, and then click **APPLY**.

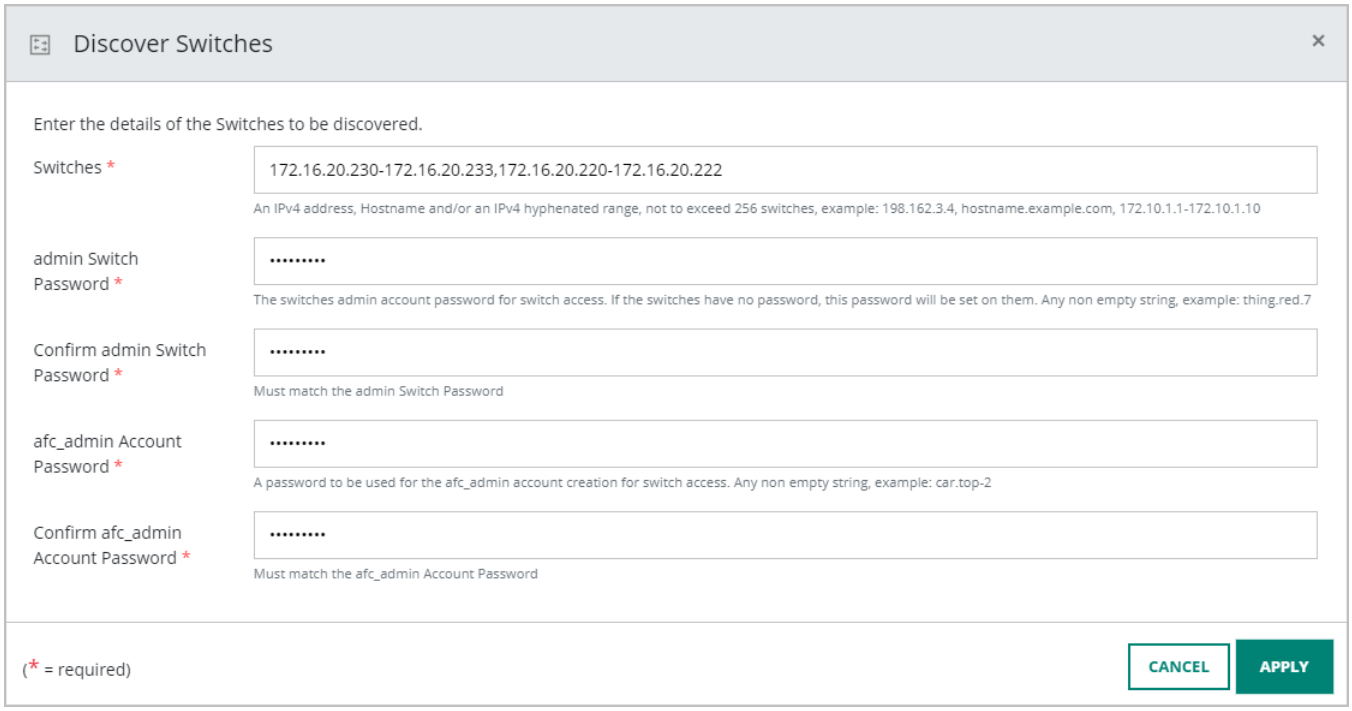

Step 5: Navigate to **Configuration > Maintenance > Switches** and verify the health status of the switches.

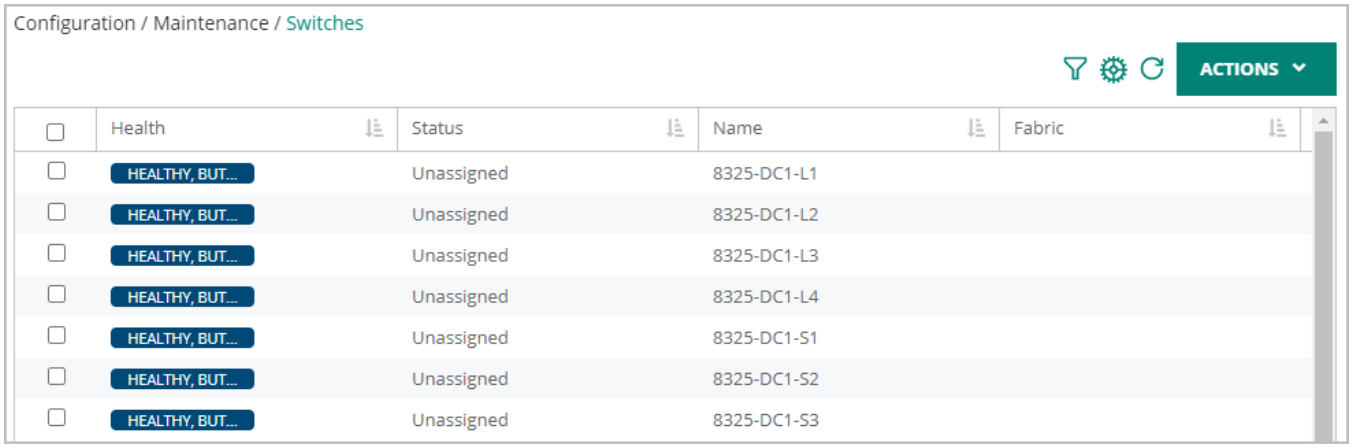

NOTE:
At the completion of the previous step, a new afc_admin account for API access from Aruba Fabric Composer will be created in all of the switches. The health status of the switches should indicate **HEALTHY, BUT NOT ADDED TO A FABRIC**. This state is indicated with the dark blue color.

Add Switches to the Fabric

Switches must be added to a fabric before they can be configured. When adding a switch to the 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: From the Guided Setup menu, select **ASSIGN SWITCH TO FABRIC**.

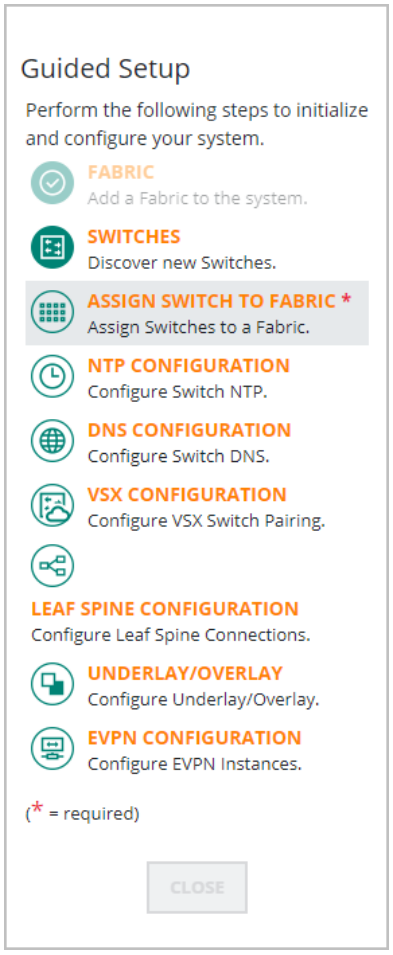

Step 2: Assign switches to the fabric grouped by role. Use the drop-down to expand a list of switches and roles available. For a large fabric the form can be used to filter on hostname. First select all spine switches, select **Spine** as Role, and then click **ADD**.

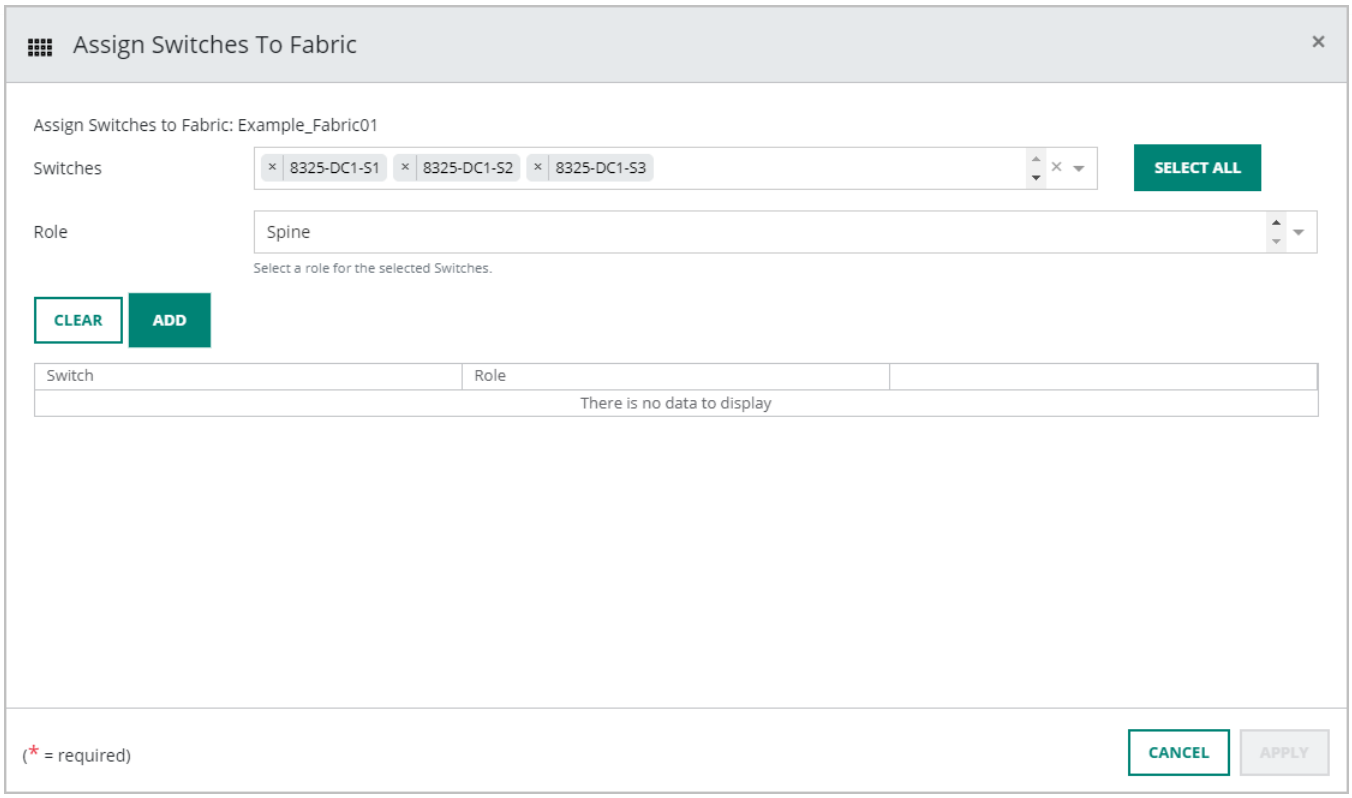

Step 3: Repeat the above steps for border leaf switches.

Step 4: Verify all switches are listed with the correct role before proceeding to the next step, and then click **ADD**.

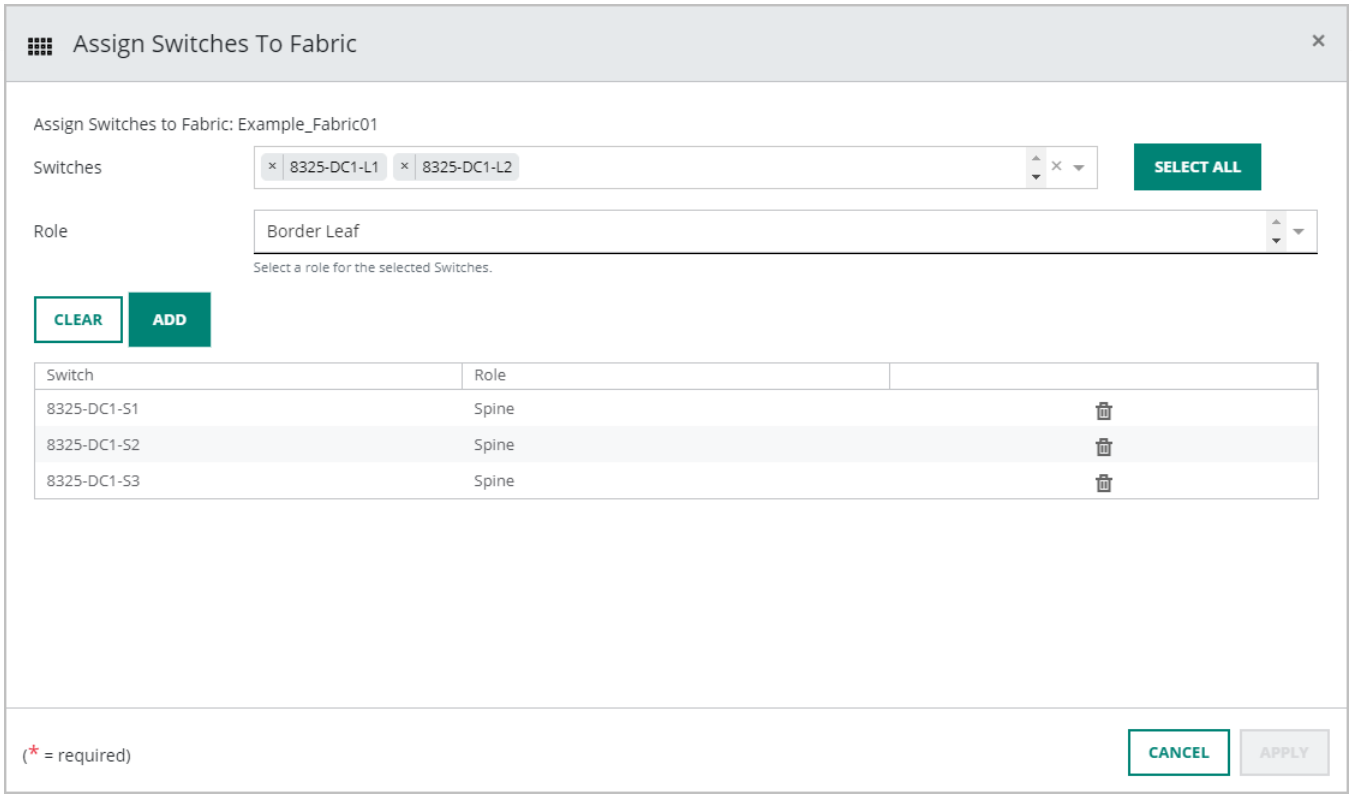

Step 5: Repeat the previous step for leaf switches.

NOTE:
SELECT ALL can be used to catch all the leaf switches. It is important to verify the role assignments to ensure successful configuration of the switches.

Step 6: Once all switches are added to the fabric with the correct role, click **APPLY**.

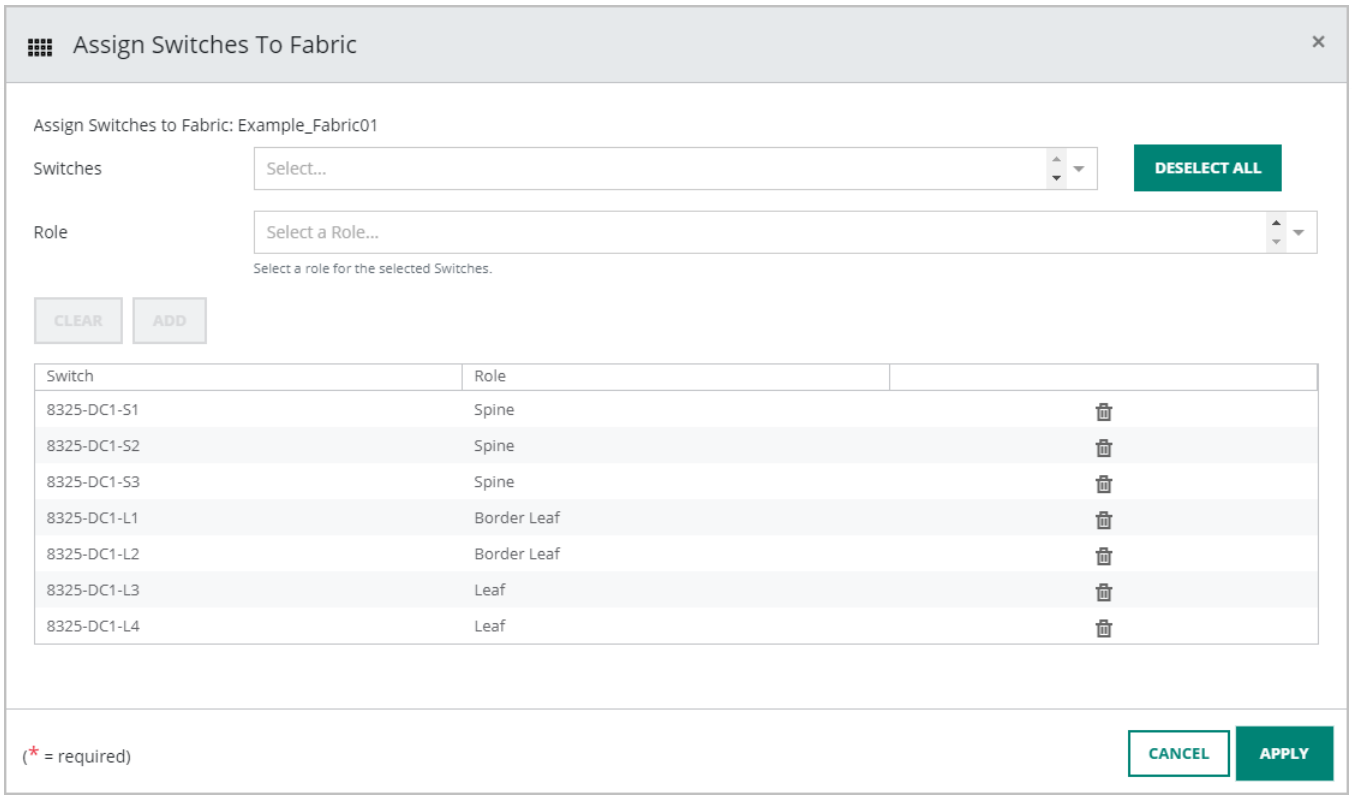

Step 7: Navigate to **Configuration > Maintenance > Switches** and verify the health status of the switches.

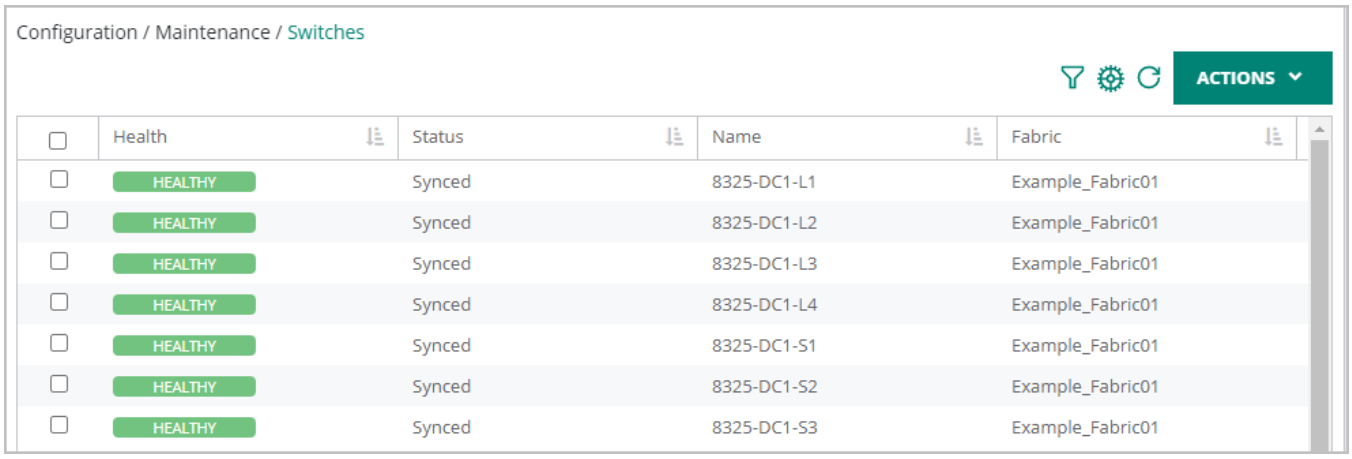

NOTE:
All switches should be in the fabric and displayed as **HEALTHY** in green. The MTU of the physical ports will be adjusted to 9198 in order to support jumbo frames and VXLAN encapsulation overhead.

Configure NTP for the Fabric

Modern networks require accurate, synchronized time. In the following steps, provide the host details for one or more NTP servers reachable from the management LAN of the data center network and associate them with your fabric. At the completion of this procedure the data center switches should have date and time synchronized with the NTP servers.

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

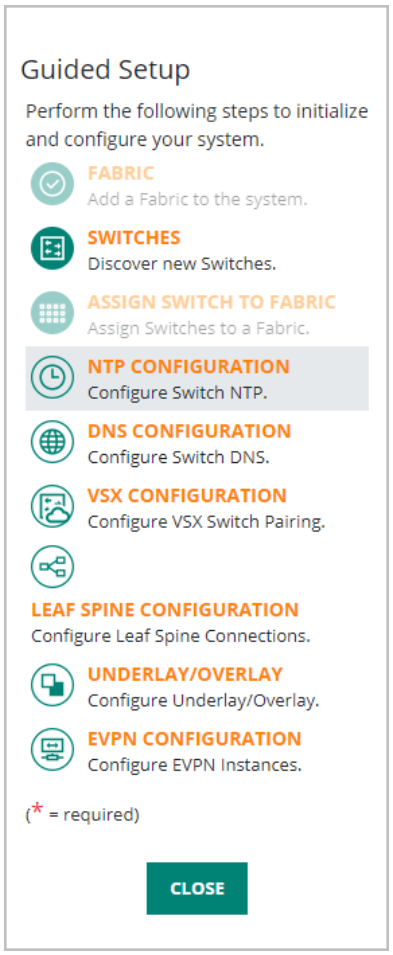

Step 2: From the NTP Configuration page, enter a name and description, and then click **NEXT**.

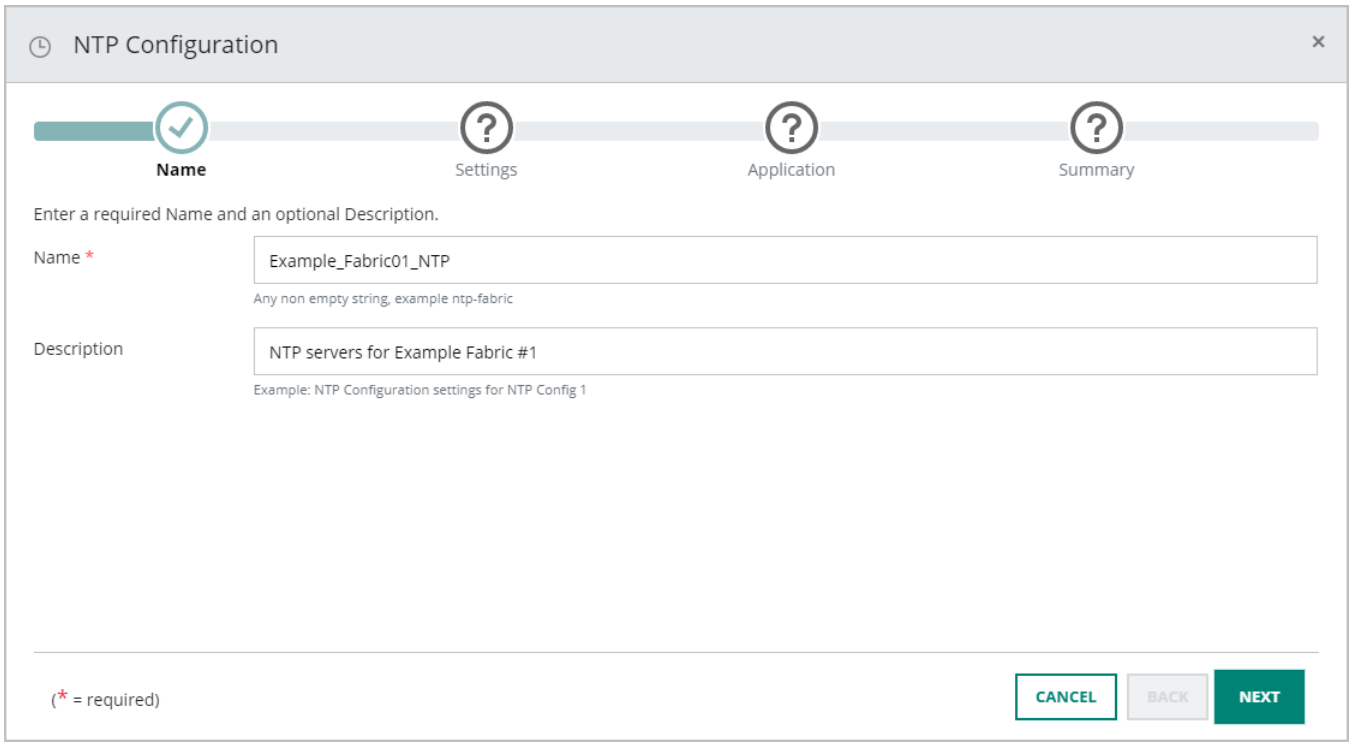

Step 3: Enter a valid NTP server host name or IP address. Use the tab or return key to complete the entry or the mechanism illustrated below. After at least one valid NTP server is entered, click **NEXT**.

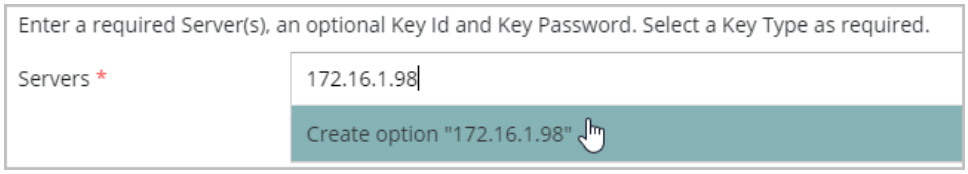

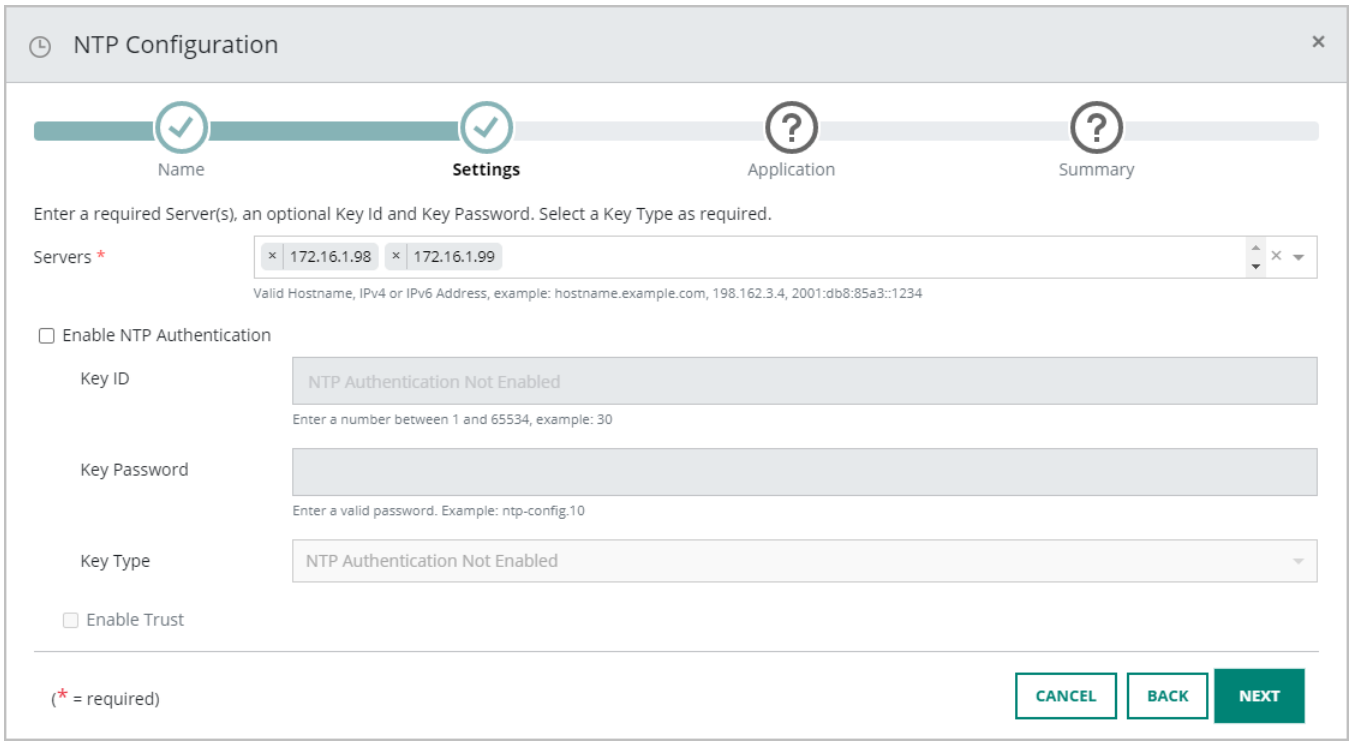

Step 4: In the **Fabric** field, select the name of the fabric, and then click **NEXT**.

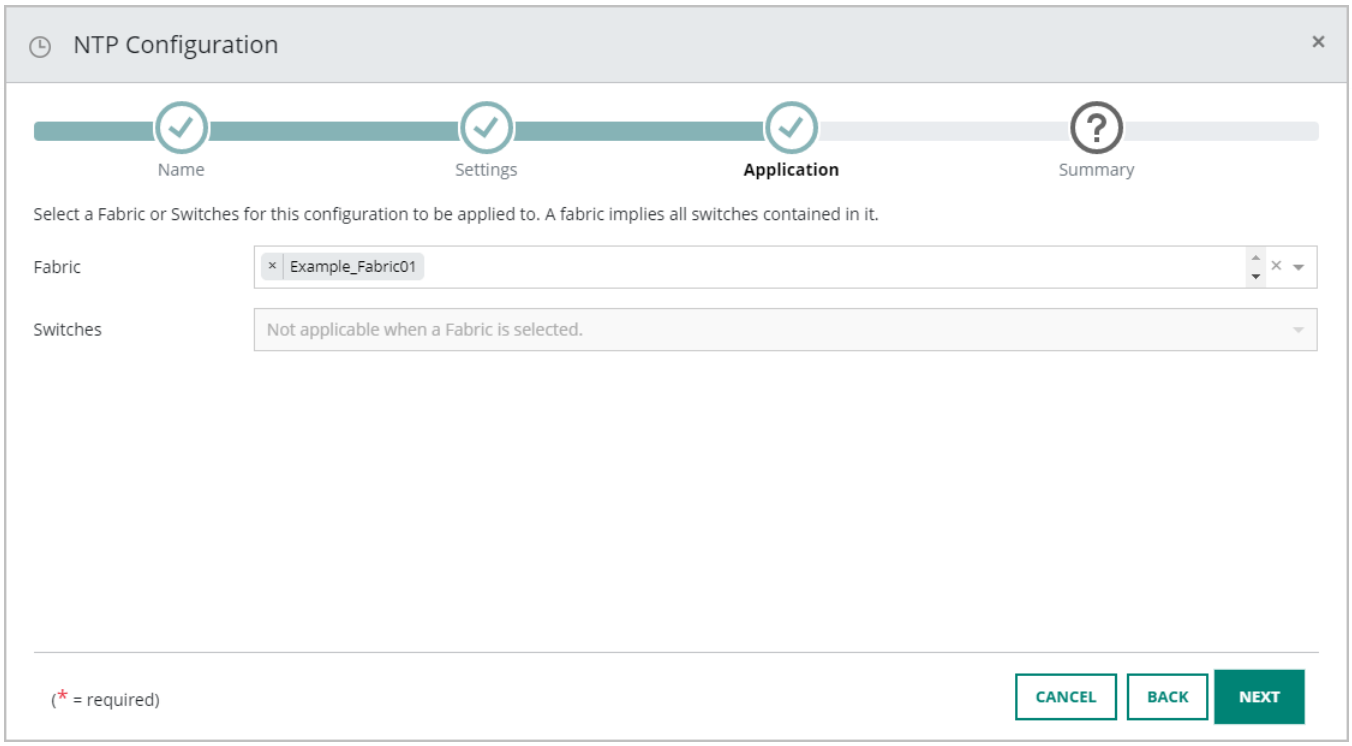

Step 5: Confirm the information was entered correctly and click **APPLY**.

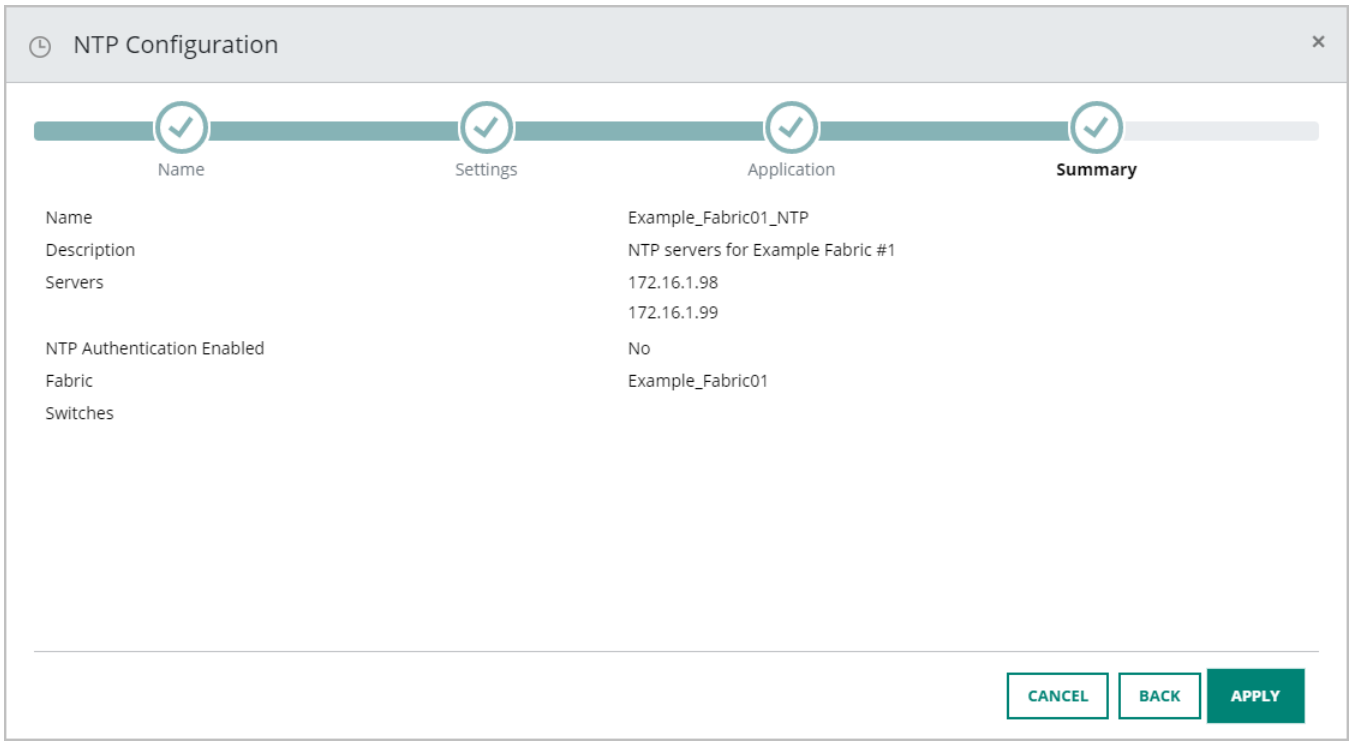

Step 6: Verify the switches have synchronized date and time with the **show ntp status** command.

```
8325-DC1-L1# show ntp status
NTP Status Information

NTP                          : Enabled
NTP Authentication           : Disabled
NTP Server Connections       : Using the mgmt VRF

System time                  : Thu Apr 29 10:29:49 AST 2021
NTP uptime                   : 16 hours, 12 minutes, 55 seconds

NTP Synchronization Information

NTP Server                   : 172.16.1.99 at stratum 3
Poll interval                : 64 seconds
Time accuracy                : Within 0.169670 seconds
Reference time               : Thu Apr 29 2021 10:16:24.115 as per America/Aruba
```

NOTE:
This could take a few minutes and the verification step will fail if a hostname is used until the next step is completed for defining DNS servers.

Configure DNS for the fabric

In most deployments, DNS is required for applications to function correctly. In the following steps, provide one or more DNS server IP addresses reachable from the management LAN of the data center network and associate them with your fabric.

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

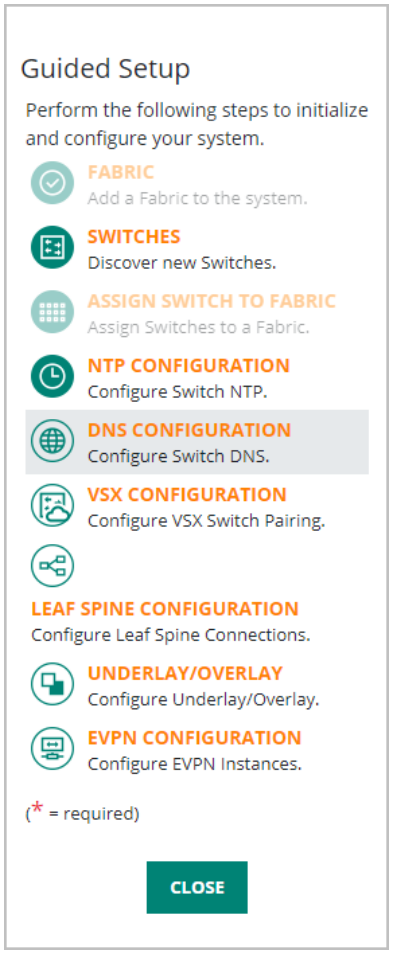

Step 2: From the DNS Configuration page, enter a name and description, and then click **NEXT**.

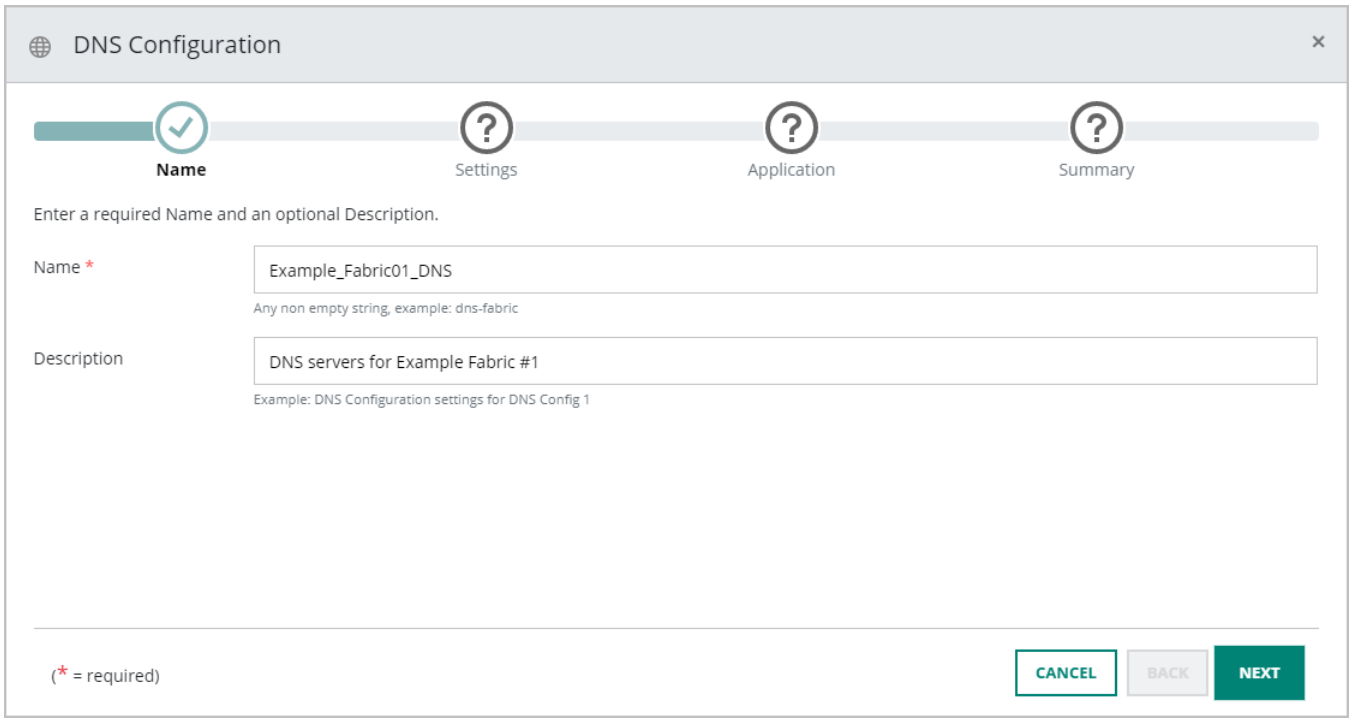

Step 3: Observe the mechanism for entering **Name Servers** as discussed above for NTP Servers, and then click **NEXT**.

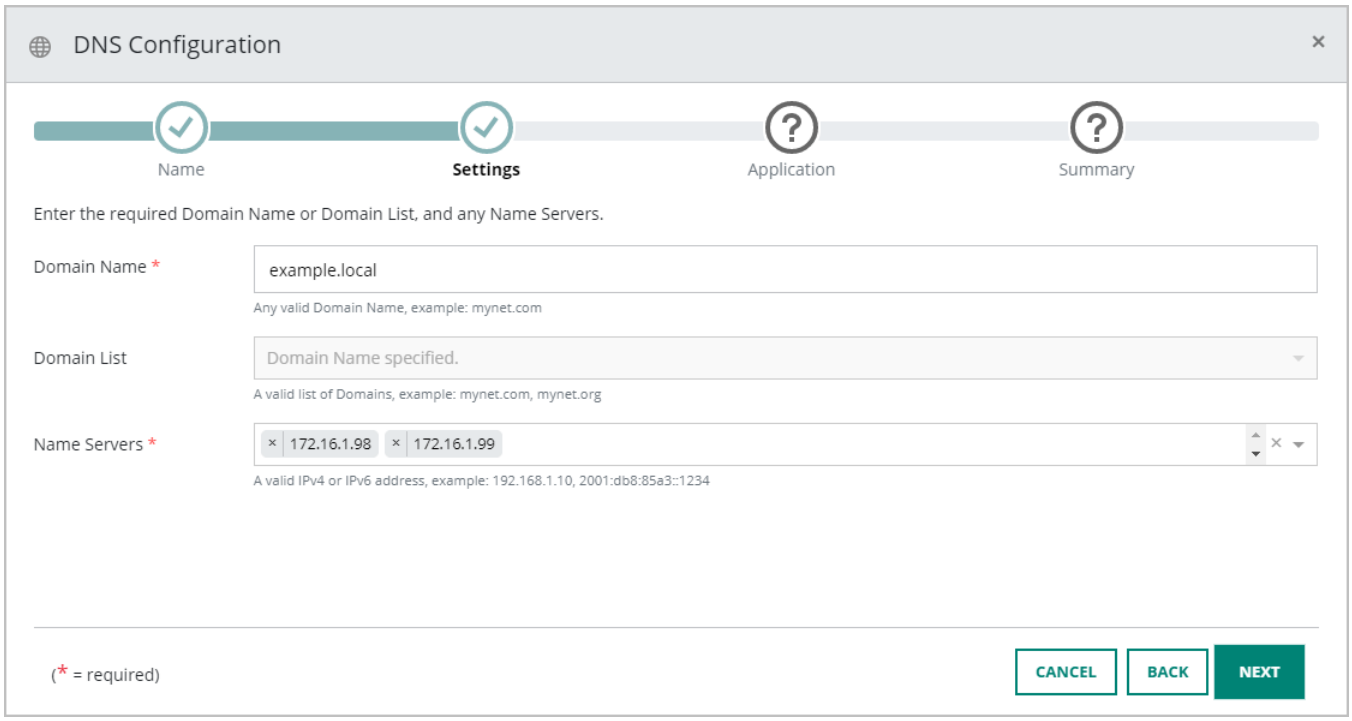

Step 4: In the **Fabrics** field, select the name of the fabric, and then click **NEXT**.

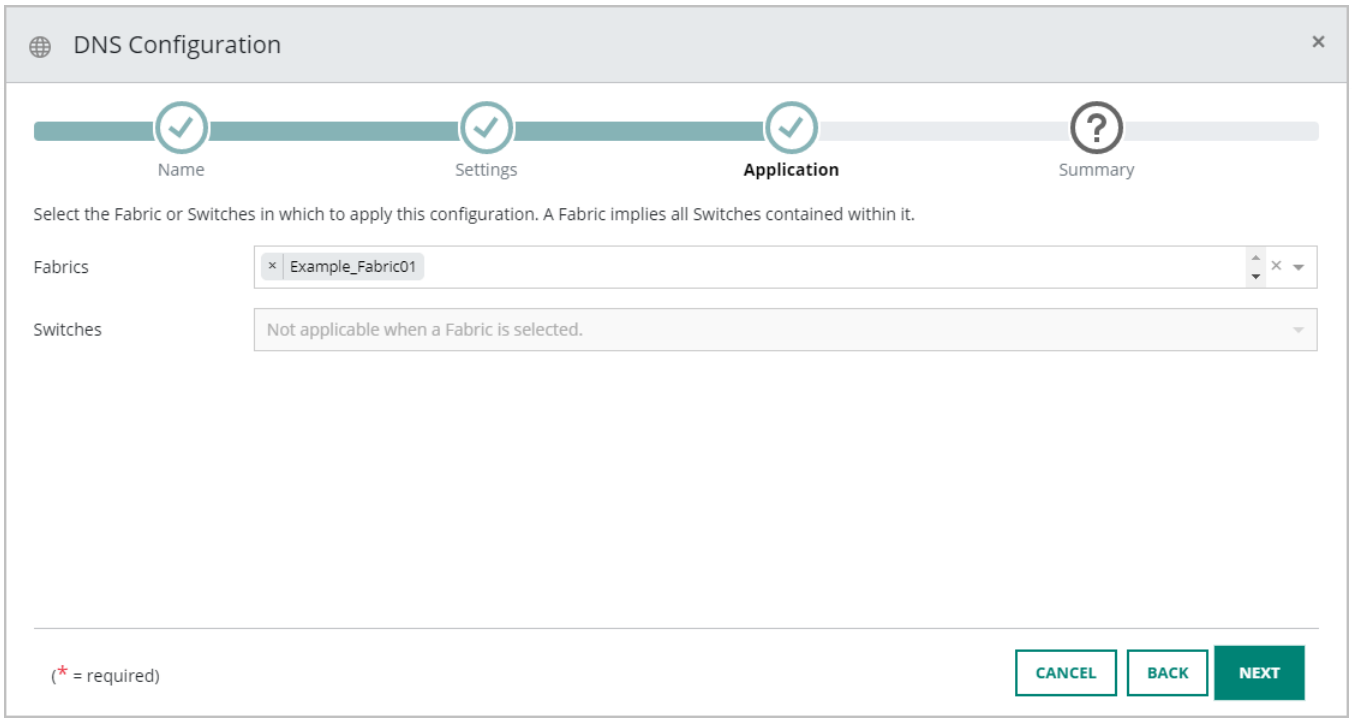

Step 5: Confirm the information was entered correctly and click **APPLY**.

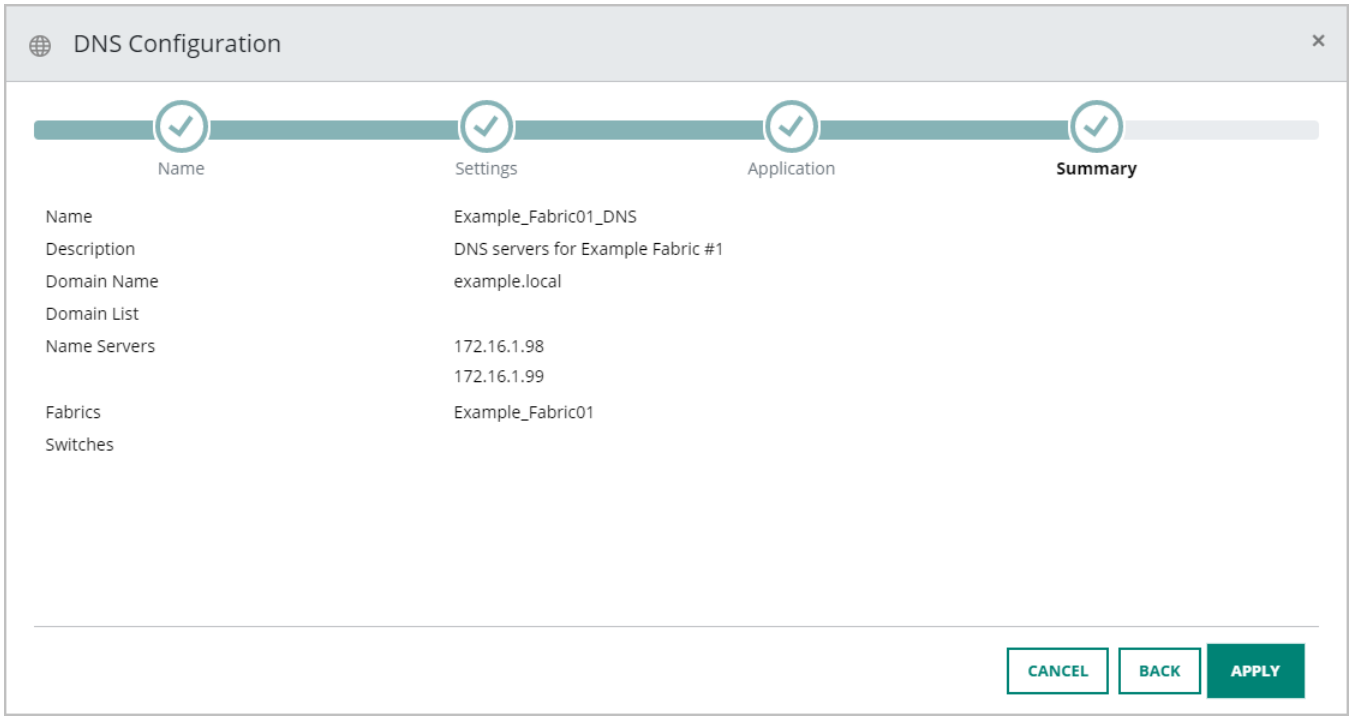

Configure VSX on Leaf Switches

Aruba Fabric Composer will automatically identify VSX switch pairs and configure them appropriately. Be sure to cable the redundant ToR leaf pairs fully and consistently across racks before proceeding to the next steps.

NOTE:
It is recommended that the number and speed of Inter-Switch Link (ISL) ports on the ToRs should equal the uplink port capacity. An additional link of any available speed can be used for keepalive messages between VSX peers.

Step 1: From the Guided Setup menu, select **VSX CONFIGURATION**

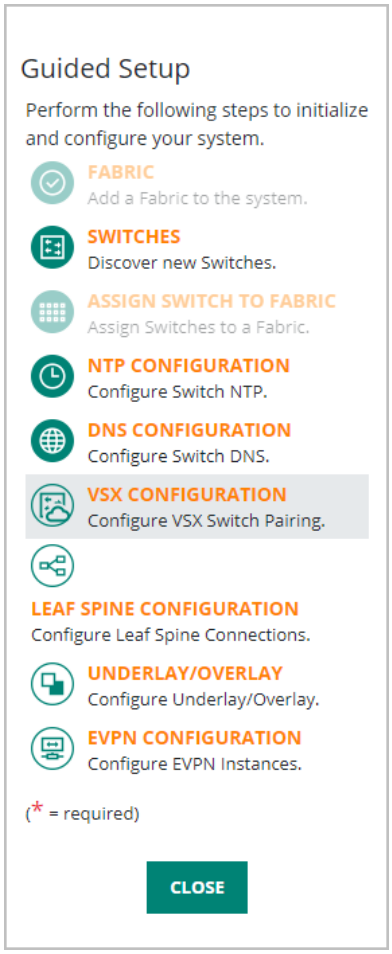

Step 2: From the VSX Create Mode page, select **Automatically generate VSX Pairs**, and then click **NEXT**.

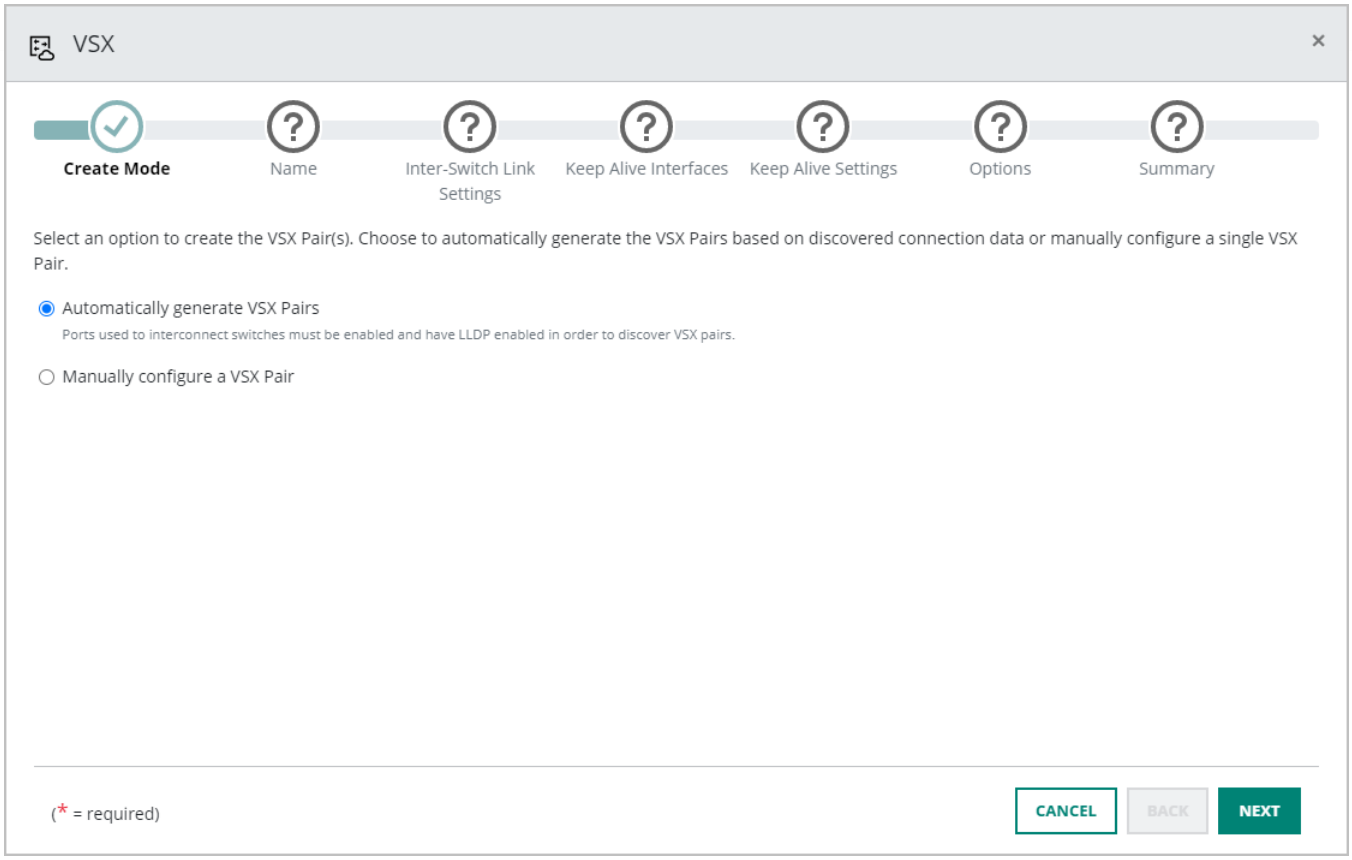

Step 3: From the VSX Name page, enter a name prefix and description, and then click **NEXT**.

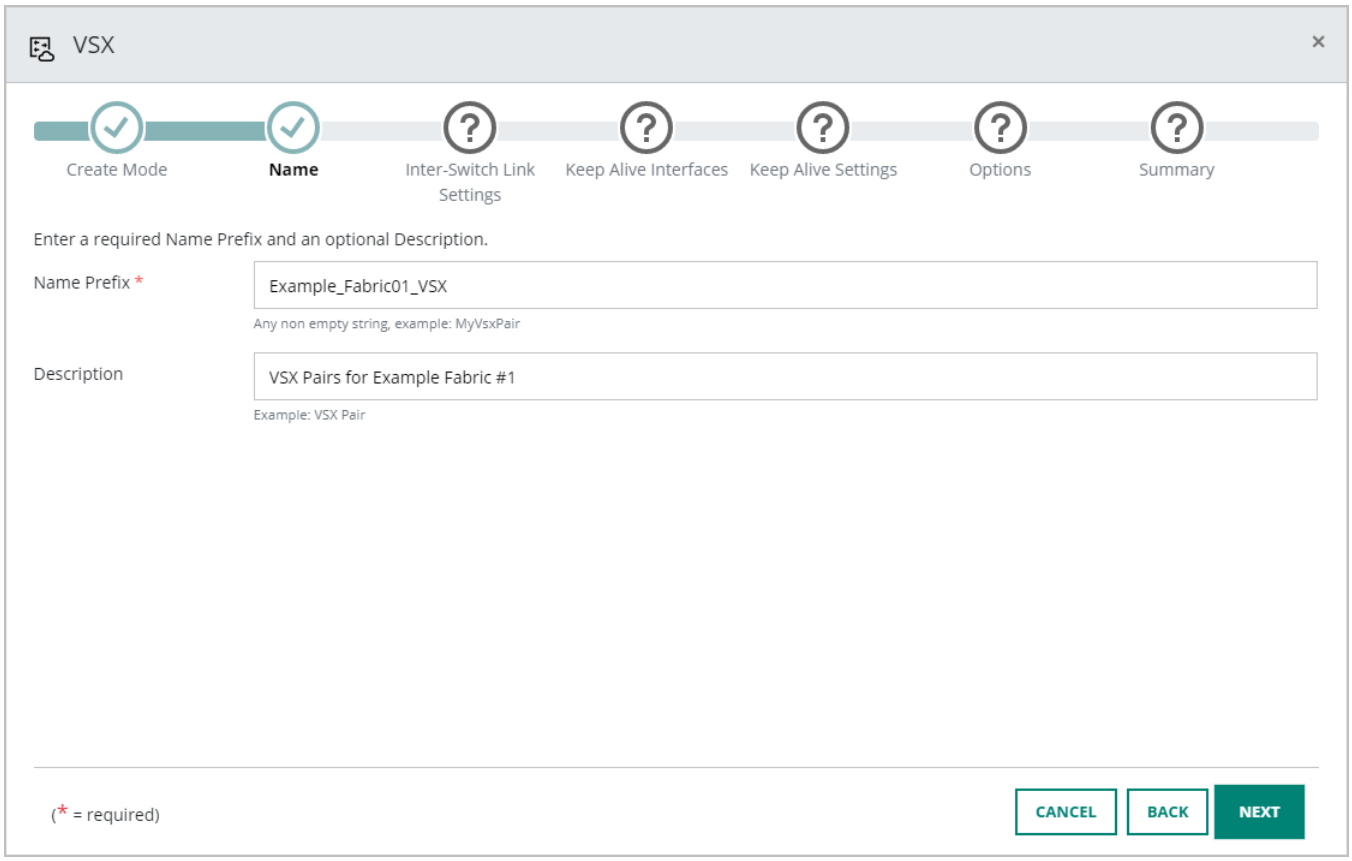

Step 4: From the VSX Inter-Switch Link Settings, leave the default settings, and then click **NEXT**.

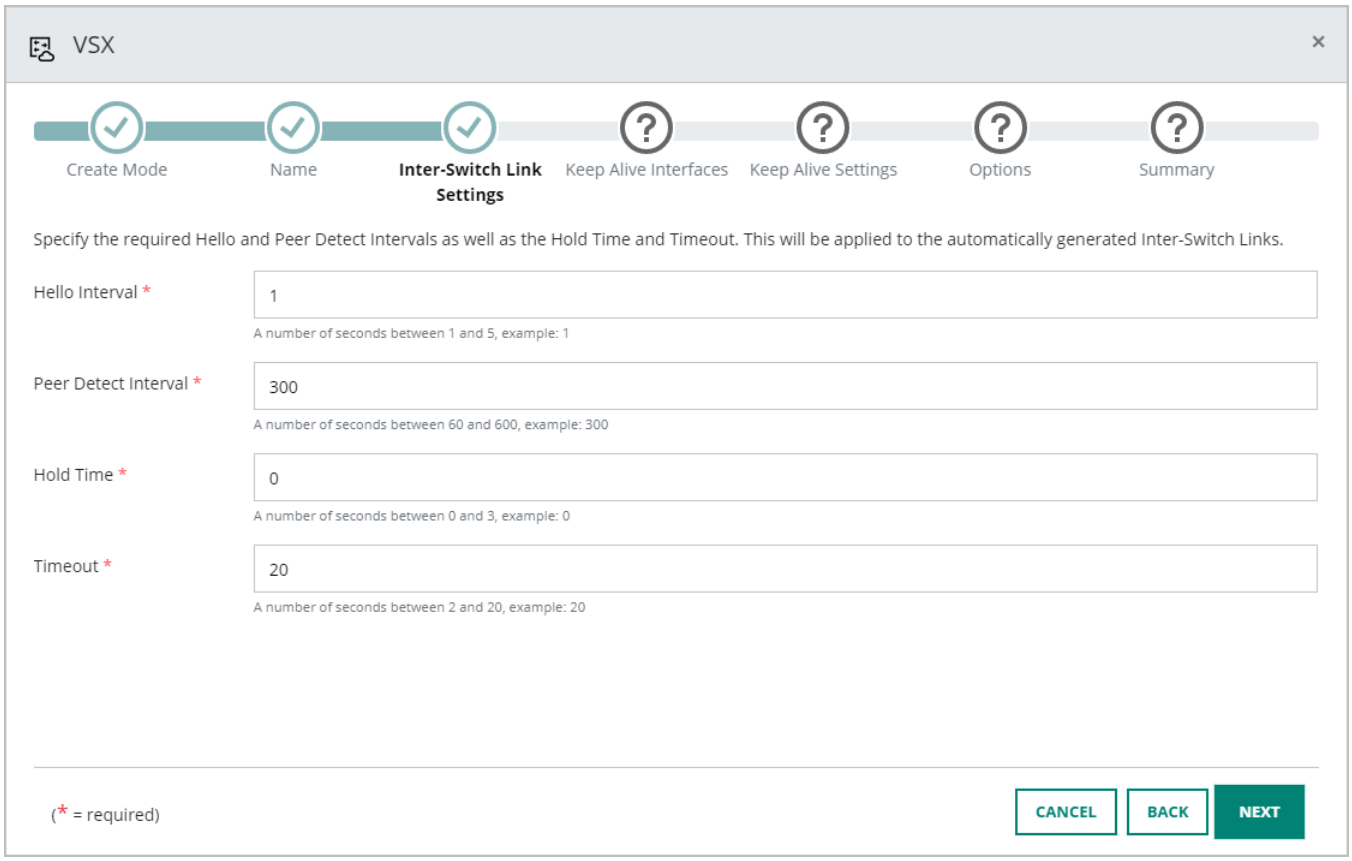

Step 5: From the VSX Keep Alive Interfaces page, implement the following setting, and then click **Next**.

- **Interface Mode:** *Point-to-Point*
- **IPv4 Subnetwork Address:** *10.0.5.0/24*

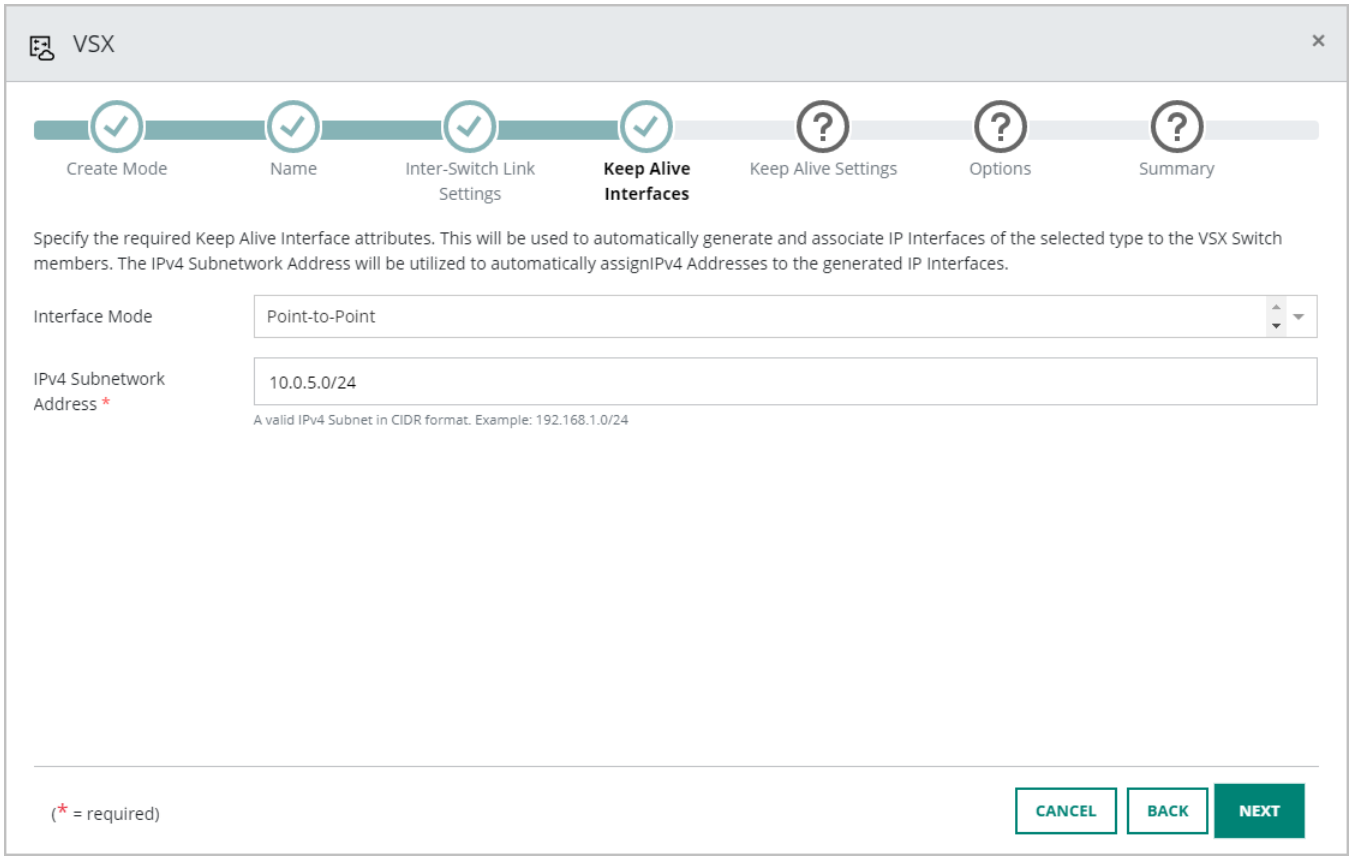

NOTE:

The IPv4 address block will be used to allocate loopback addresses (/32) for each switch, for VSX keep-alive point-to-point connection (/31) and also as a transit routed VLAN between redundant ToRs (/31). Please plan to use a block large enough to support the size of the fabric.

Step 6: From the VSX Keep Alive Settings page, leave the default settings, and then click **NEXT**.

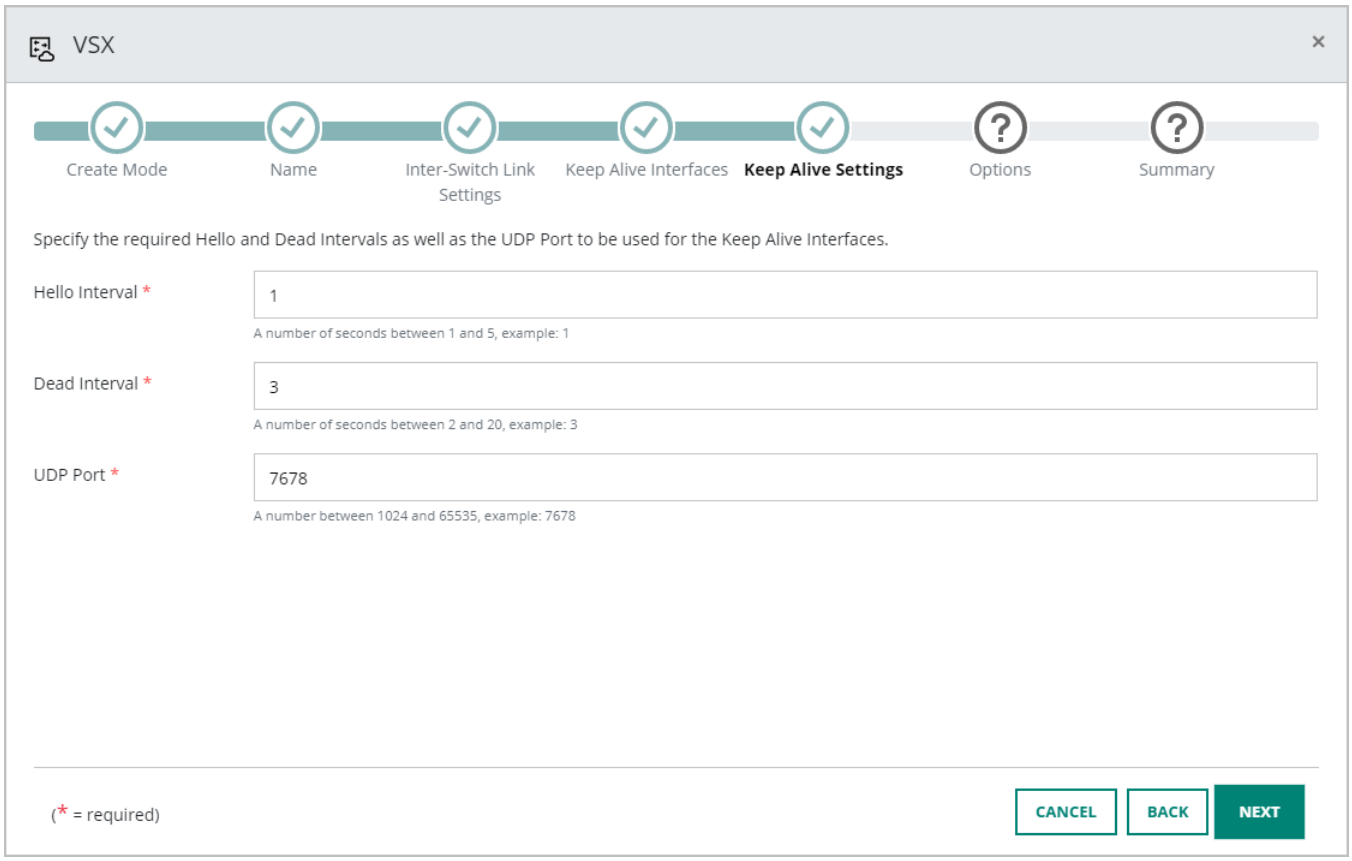

Step 7: From the VSX Options page, leave the default settings, and then click **NEXT**.

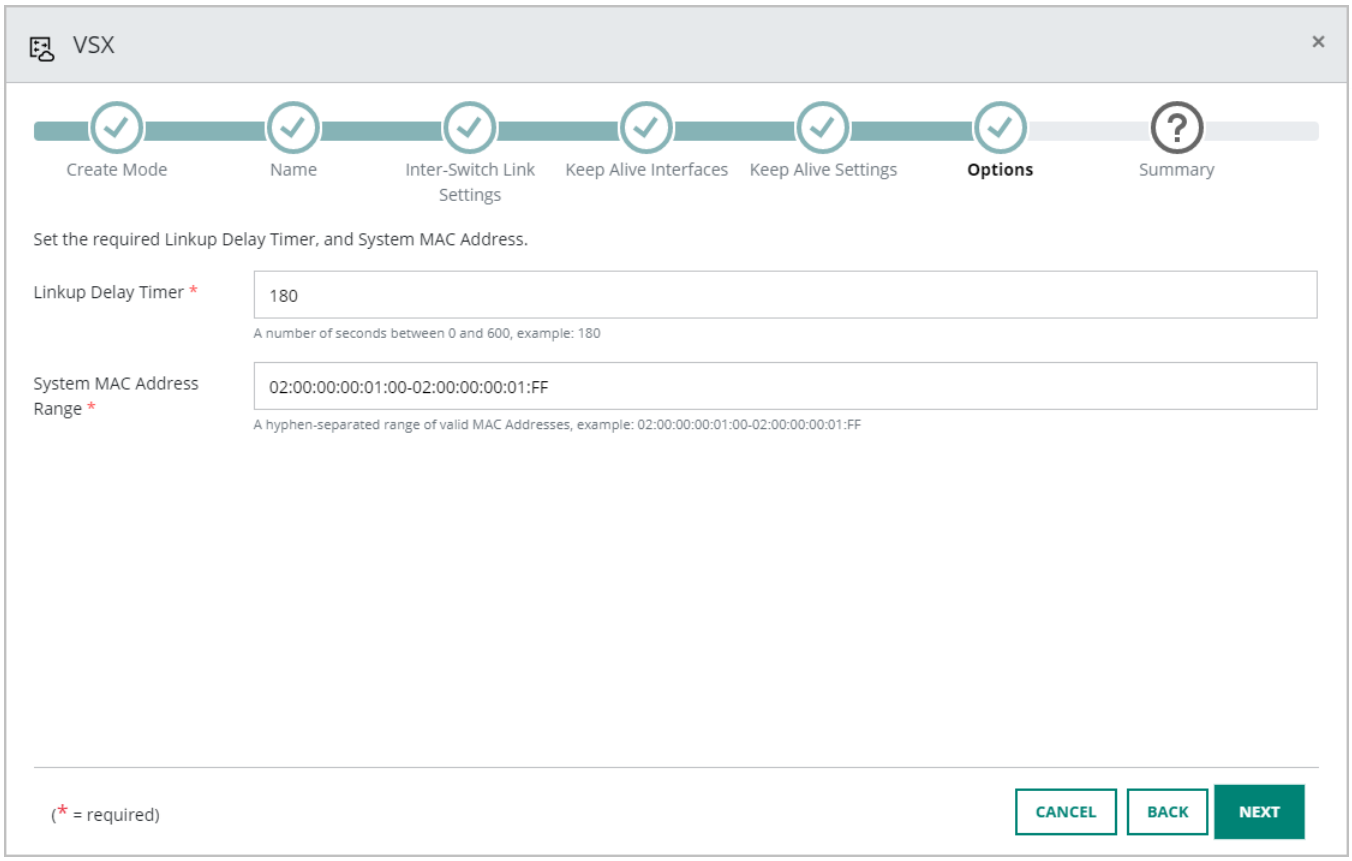

Step 8: Confirm the information was entered correctly and click **APPLY**.

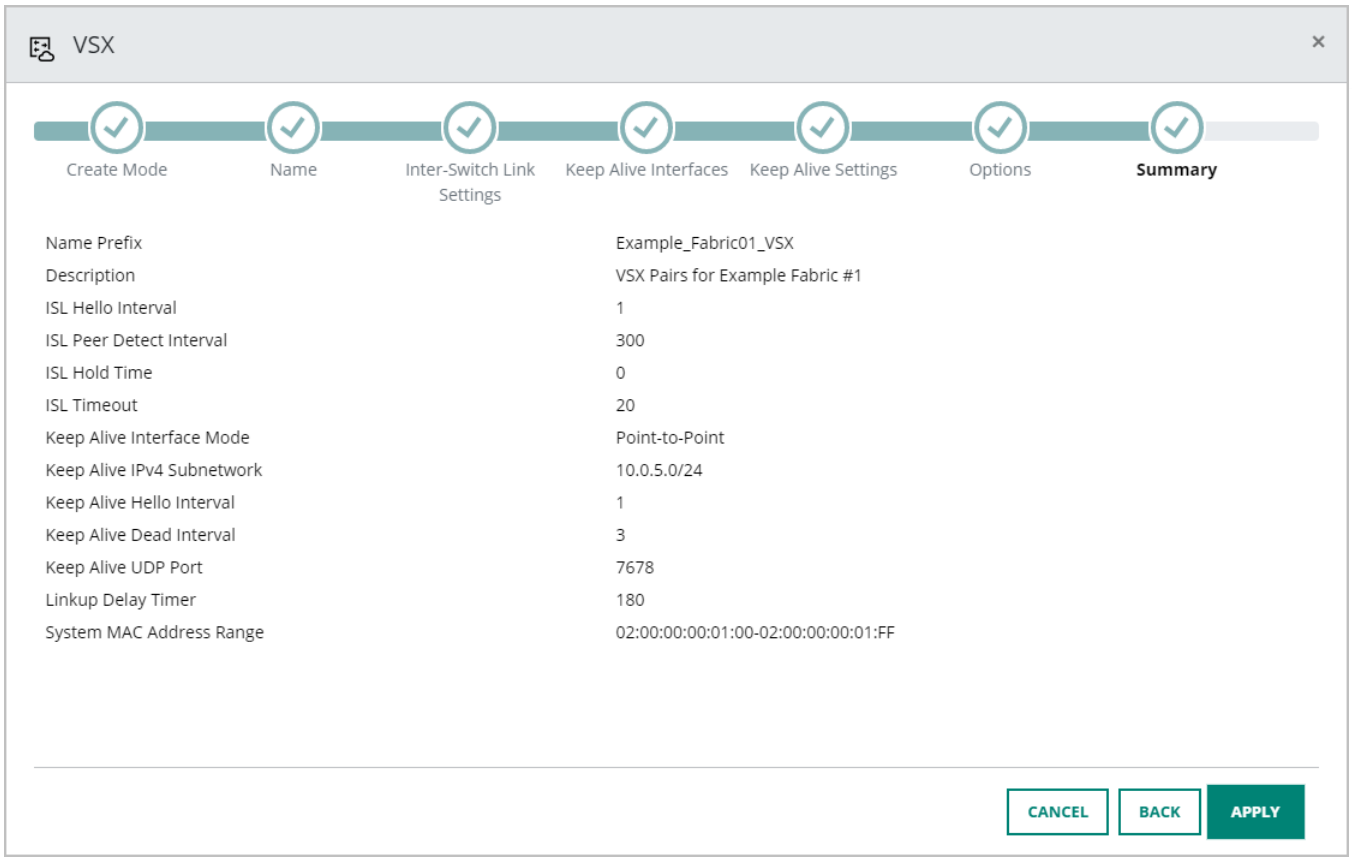

Configure Leaf to Spine Connections

Aruba Fabric Composer will automatically identify leaf to spine connections. Be sure to cable the links fully and consistently across the data center. This allows for automated configuration of the routed ports and related IP subnets between leaf and spine switches. In the following steps, an IP network address space is provided, additional default values are accepted, and Aruba Fabric Composer completes the configuration of routed links between leaf and spine switches.

Step 1: From the Guided Setup menu, select **LEAF SPINE CONFIGURATION**.

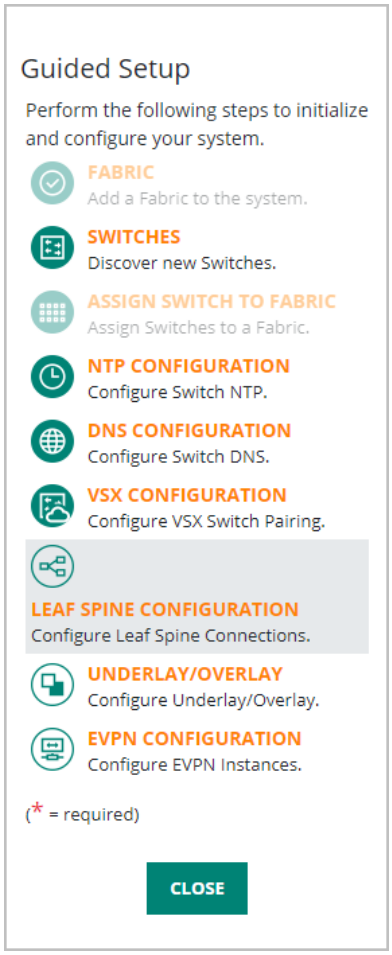

Step 2: From the Leaf-Spine Create Mode page, select **Automatically generate Leaf-Spine Pairs**, and then click **NEXT**.

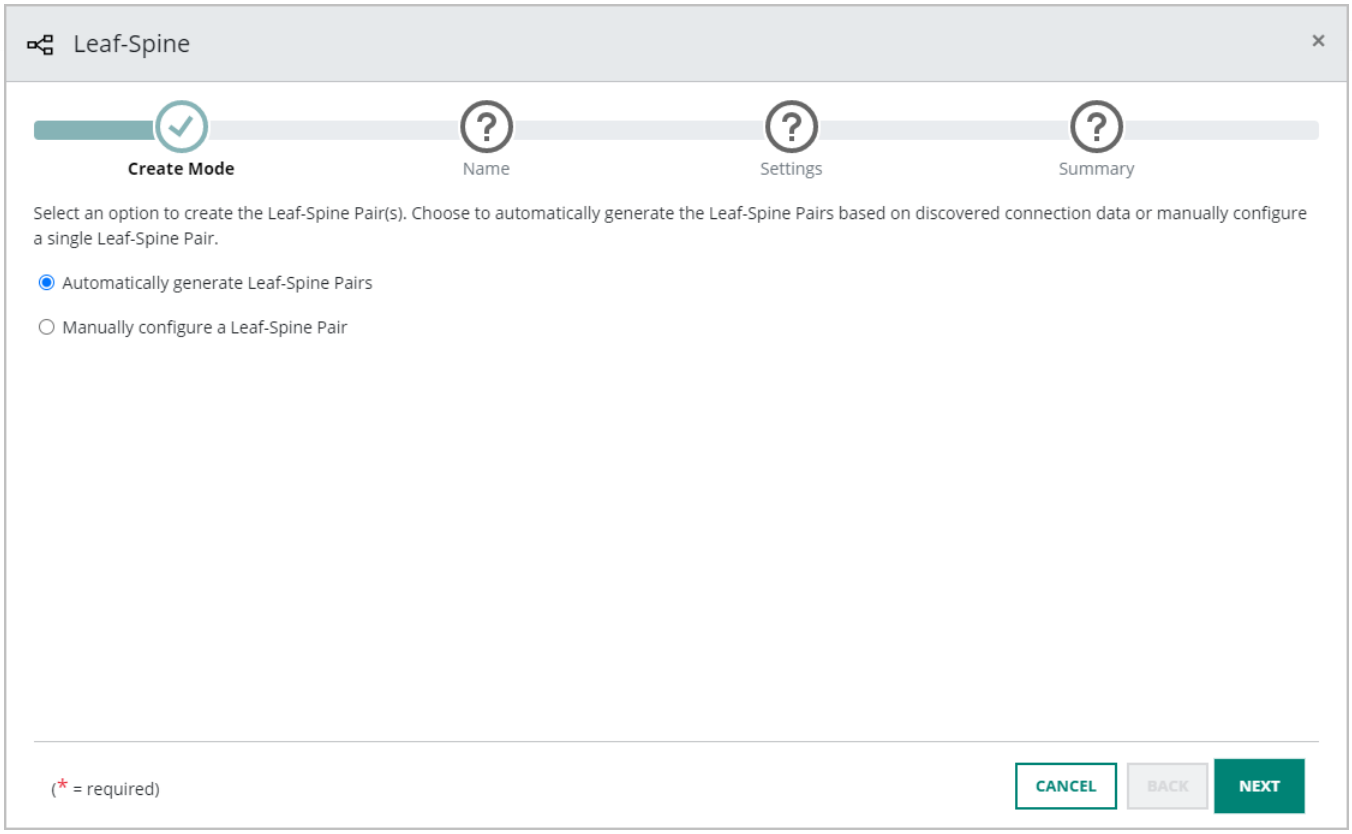

Step 3: From the Leaf-Spine Name page, enter a name prefix and description, and then click **NEXT**.

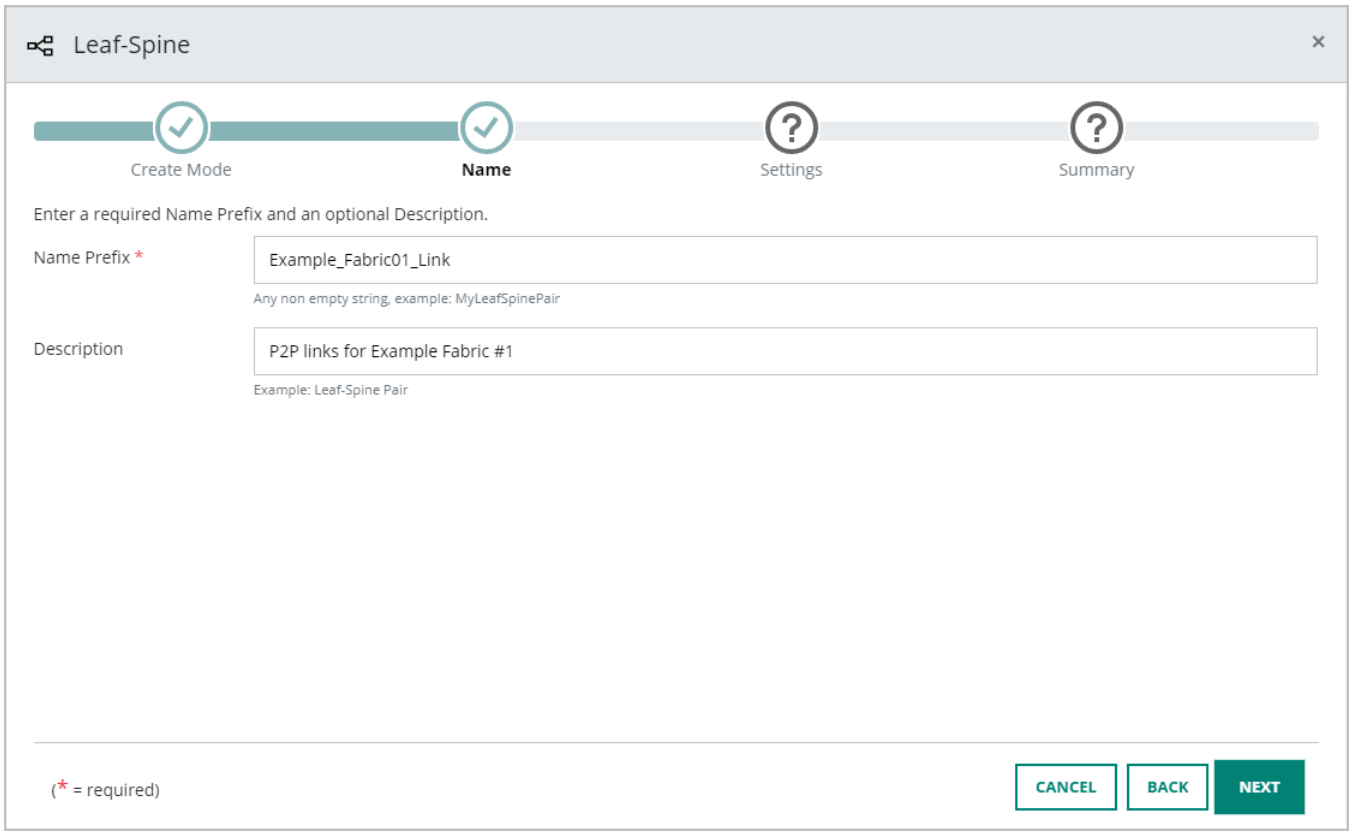

Step 4: From the Leaf-Spine Settings page, enter an IPv4 Subnetwork Address, and then click **NEXT**.

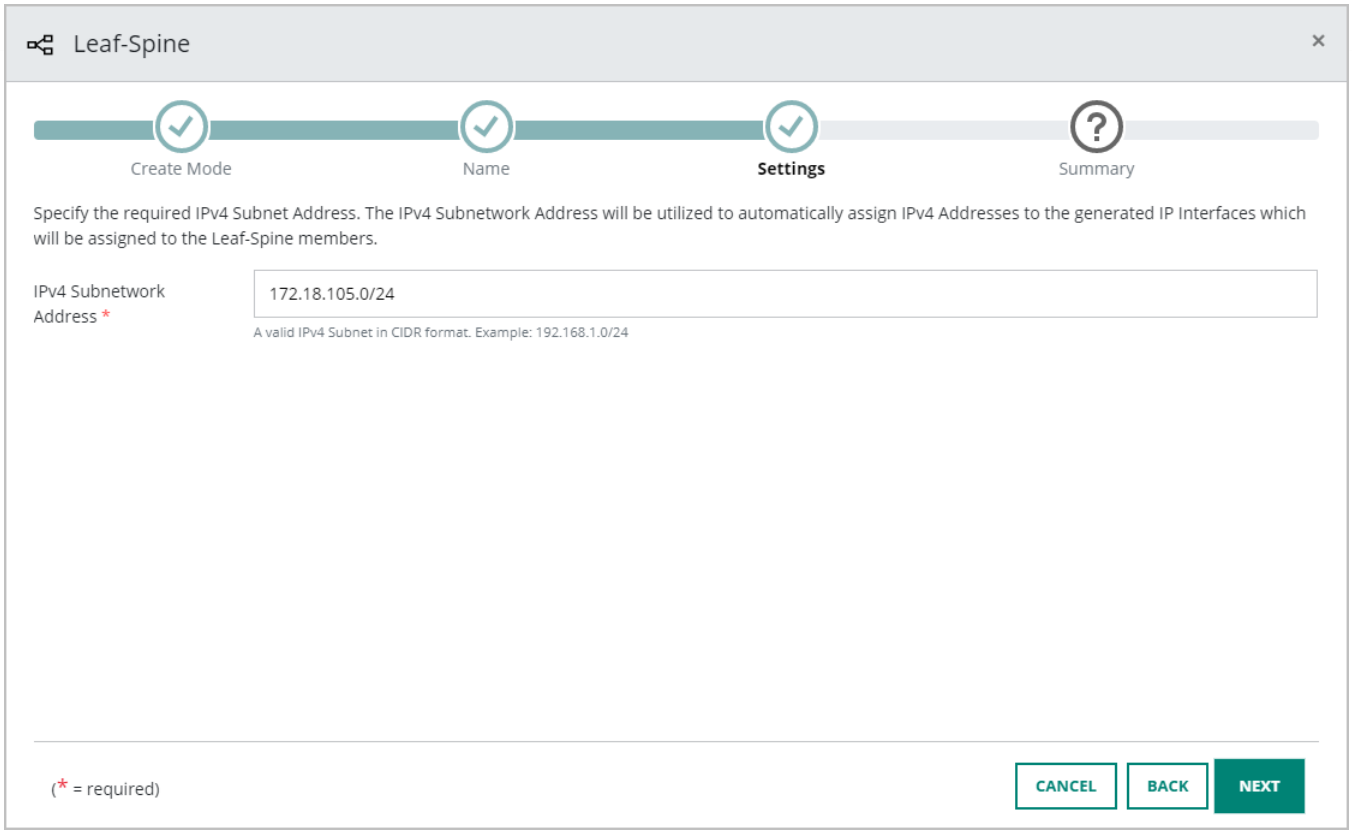

NOTE:
Use a subnet that is distinct from the subnets you plan to use in your overlay networks. The provided subnet will be used to configure routed ports on the fabric switches. In order to simplify route summarization between the campus and data center, ensure this range is unique to the underlay.

Step 5: Confirm the information was entered correctly and click **APPLY**.

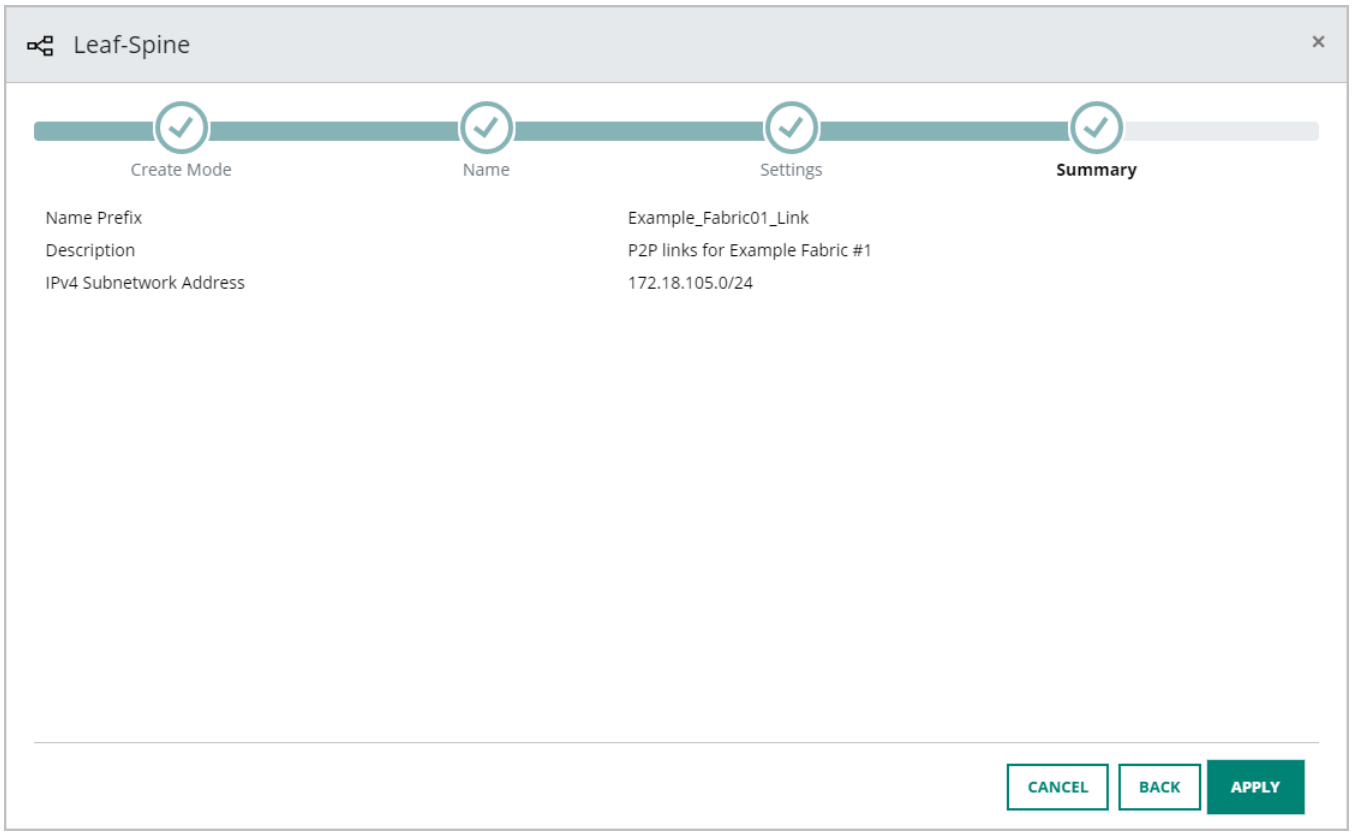

Configure Underlay Network Routing

The Aruba ESP Data Center spine and leaf design uses OSPF as the underlay routing protocol. In the following steps, a transit VLAN ID for redundant ToRs is provided, additional default values are accepted, and Aruba Fabric Composer completes the OSPF configuration.

Step 1: From the Guided Setup menu, select **UNDERLAY/OVERLAY**.

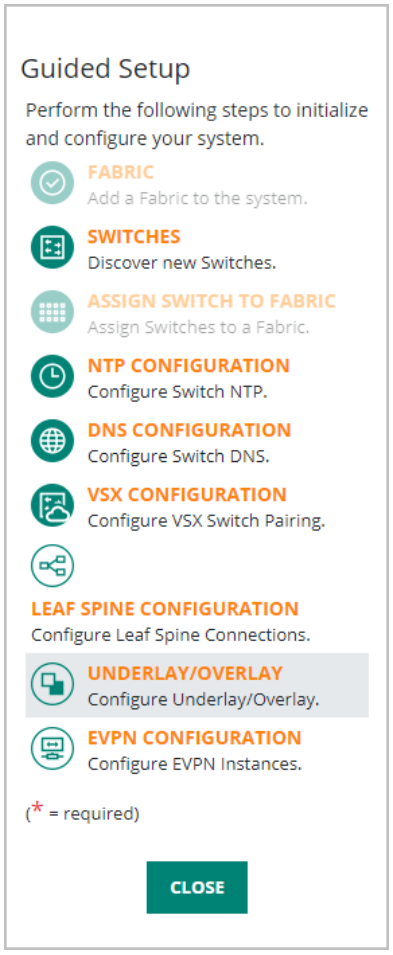

Step 2: Select the **UNDERLAYS** tab, and then from the **ACTIONS** menu on the right, click **Add**.

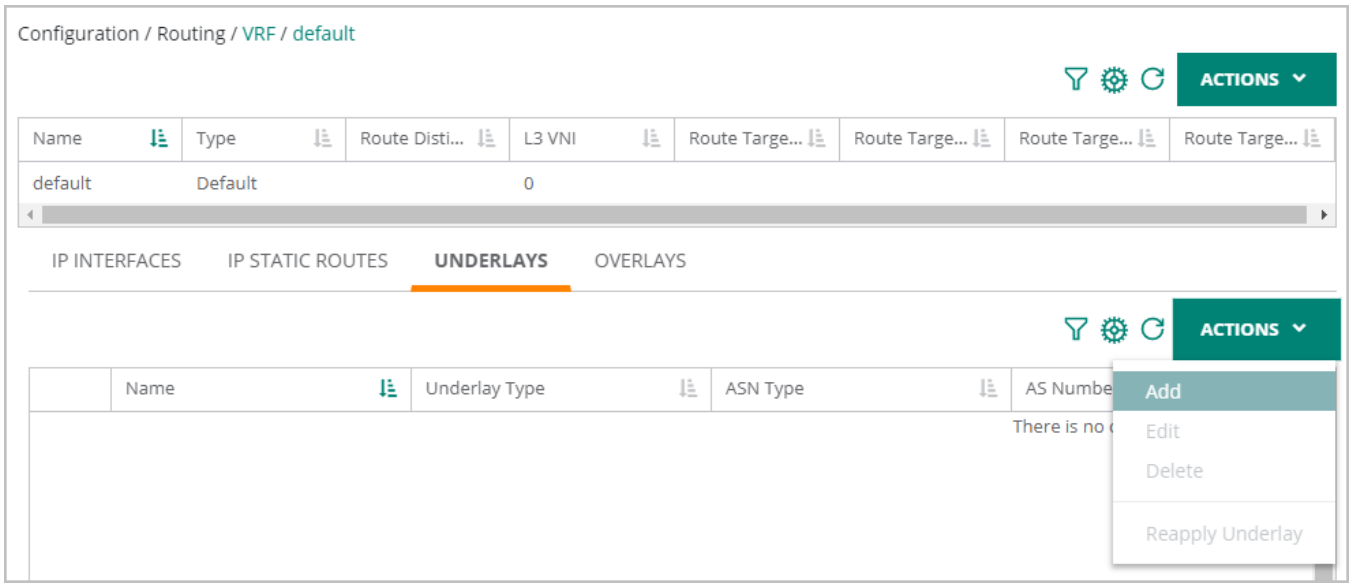

Step 3: From the Underlay Configuration Name page, enter a name and description, and then click **NEXT**.

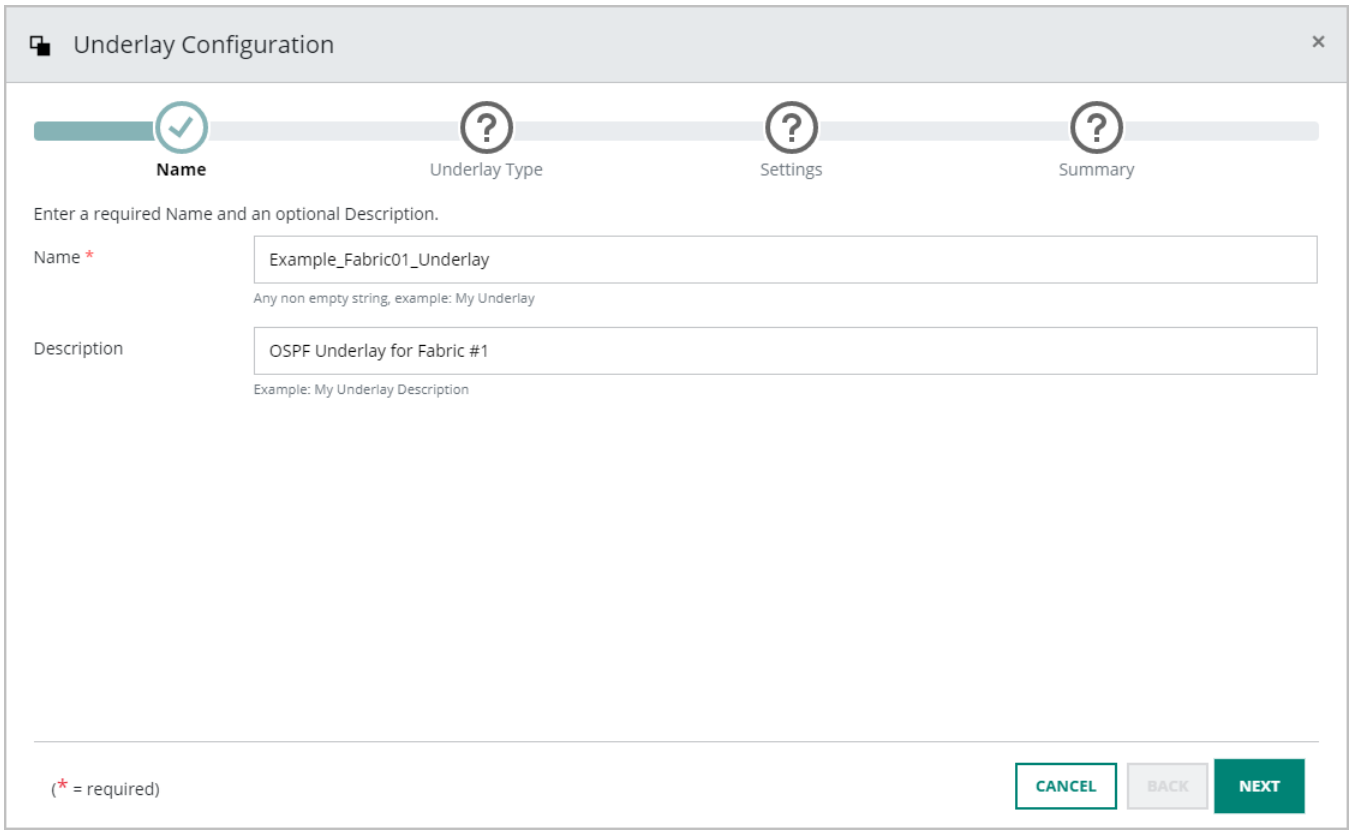

Step 4: From the Underlay Configuration Underlay Type page, leave the default **OSPF**, and then click **NEXT**.

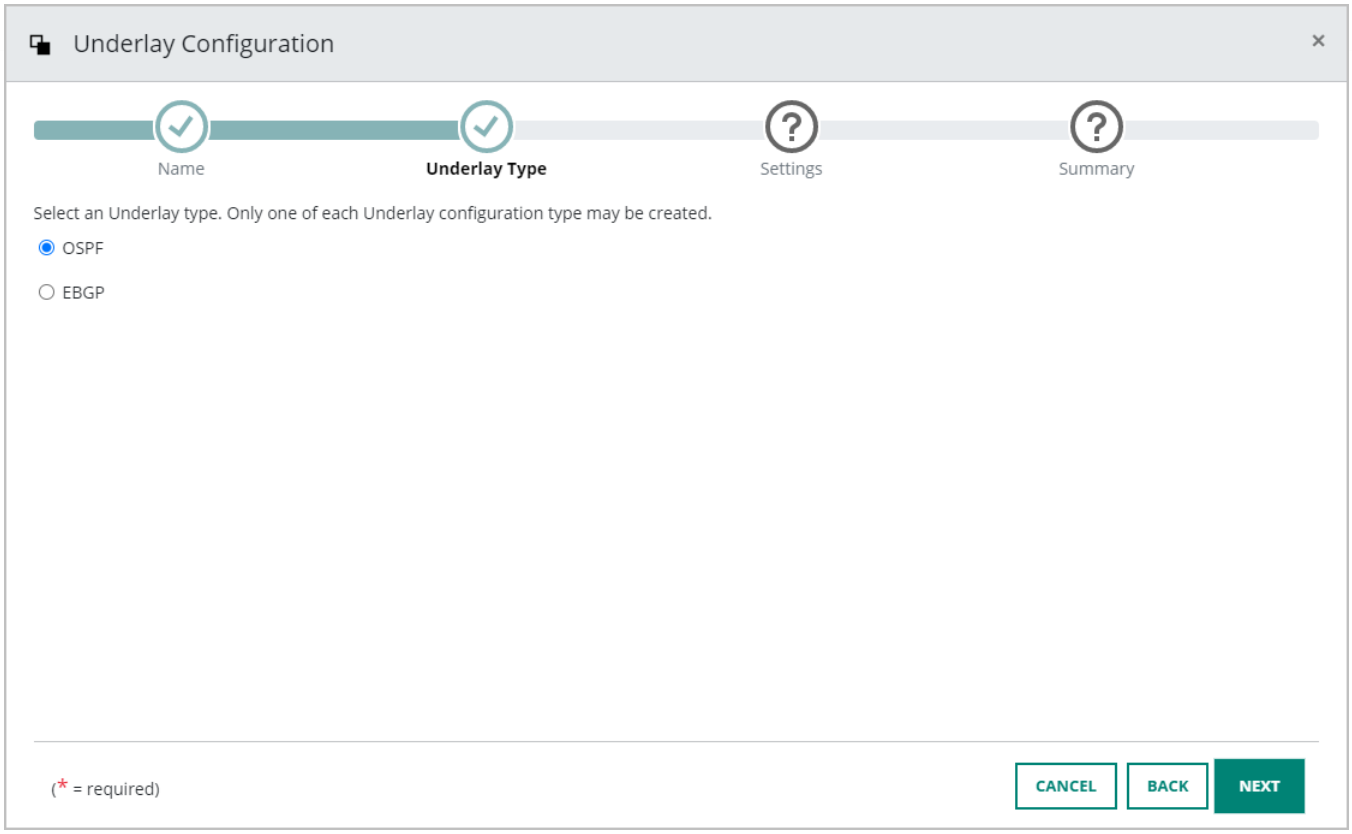

Step 5: From the Underlay Configuration Settings page, implement the following setting, and then click **Next**.

- **Transit VLAN:** *3966*
- **Bidirectional Forwarding Detection:** *Disabled*

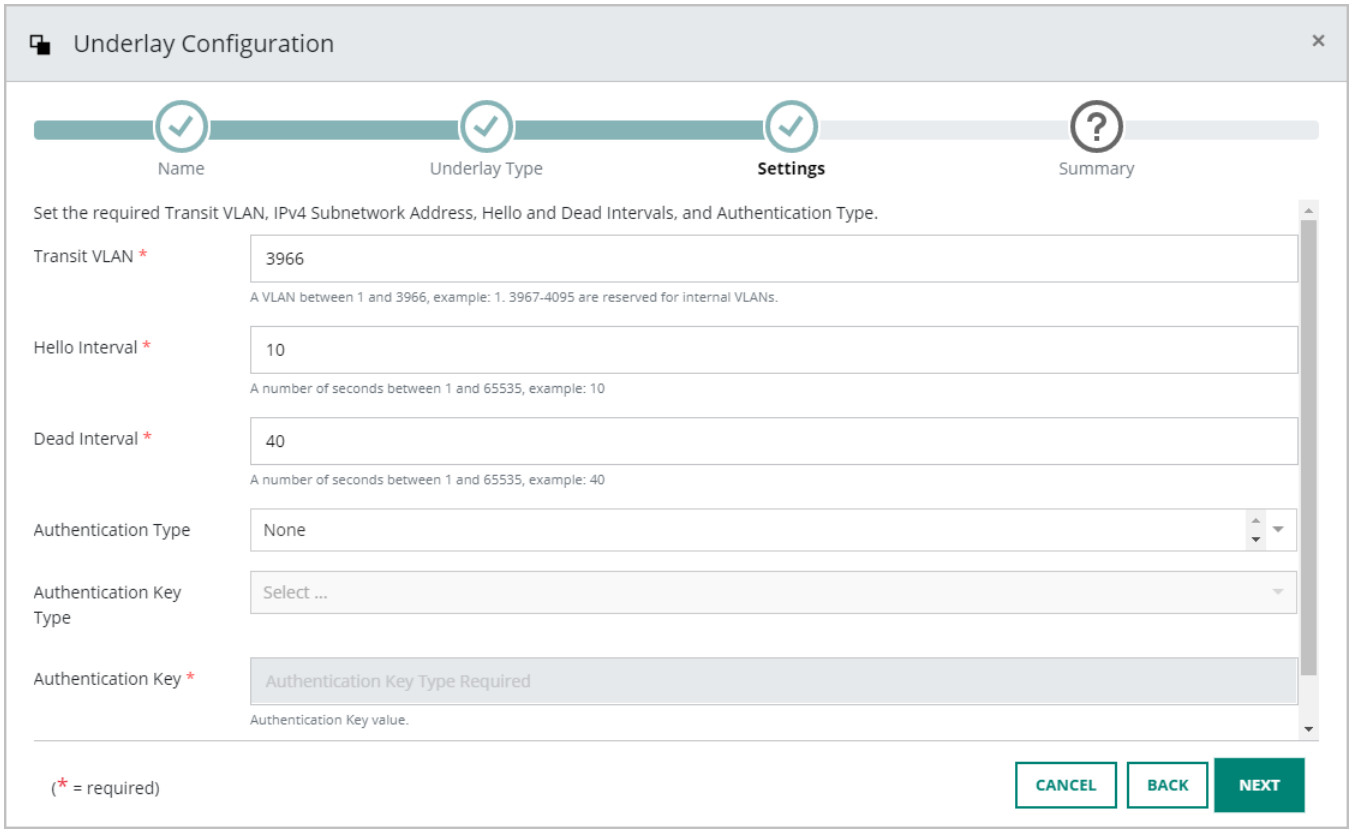

NOTE:
Enter a VLAN ID which will not be easily confused with other VLANs within the network. Scroll down to the bottom to see the **Bidirectional Forwarding Detection** option.

Step 6: Confirm the information was entered correctly and select **APPLY**.

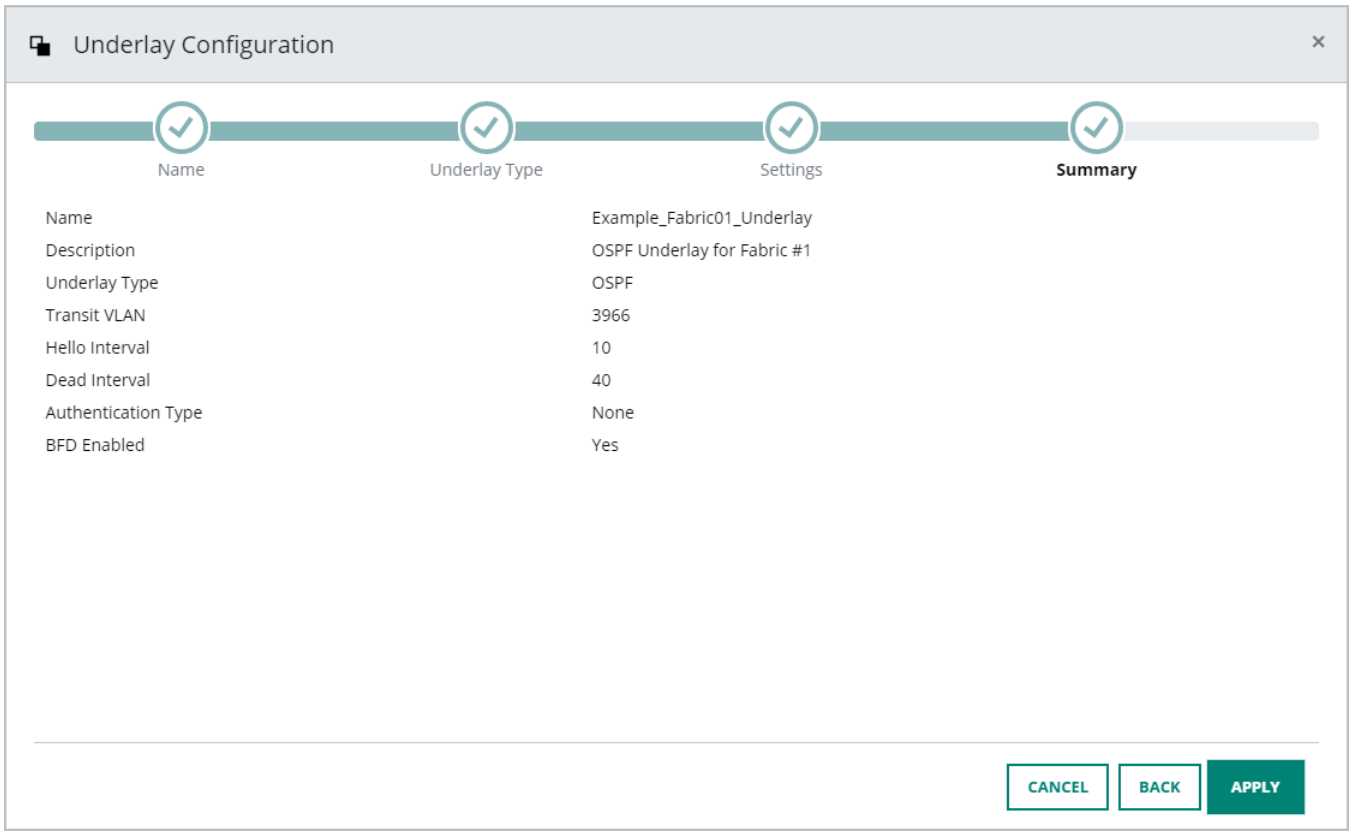

Step 7: Verify the spine switches have established **FULL** OSPF adjacencies to all leaf switches with the **show ip ospf neighbors** command.

```
8325-DC1-S1# sh ip ospf nei
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 4

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.0.5.9         n/a       FULL              172.18.105.11      1/1/1
10.0.5.10        n/a       FULL              172.18.105.13      1/1/2
10.0.5.8         n/a       FULL              172.18.105.15      1/1/3
10.0.5.6         n/a       FULL              172.18.105.17      1/1/4
```

Configure Overlay Network Routing

The first series of steps for deploying an overlay network is to configure BGP as the control plane for the fabric. The Aruba ESP Data Center spine and leaf design implements MP-BGP in the overlay to communicate host routes across the fabric using the EVPN address family. In the following steps, a default private ASN is accepted, an IP network address space for VTEPs is provided, additional default values are accepted, and Aruba Fabric Composer completes the router configuration.

Step 1: Navigate to **Configuration > Routing > VRF > default** and select the **Overlays** tab.

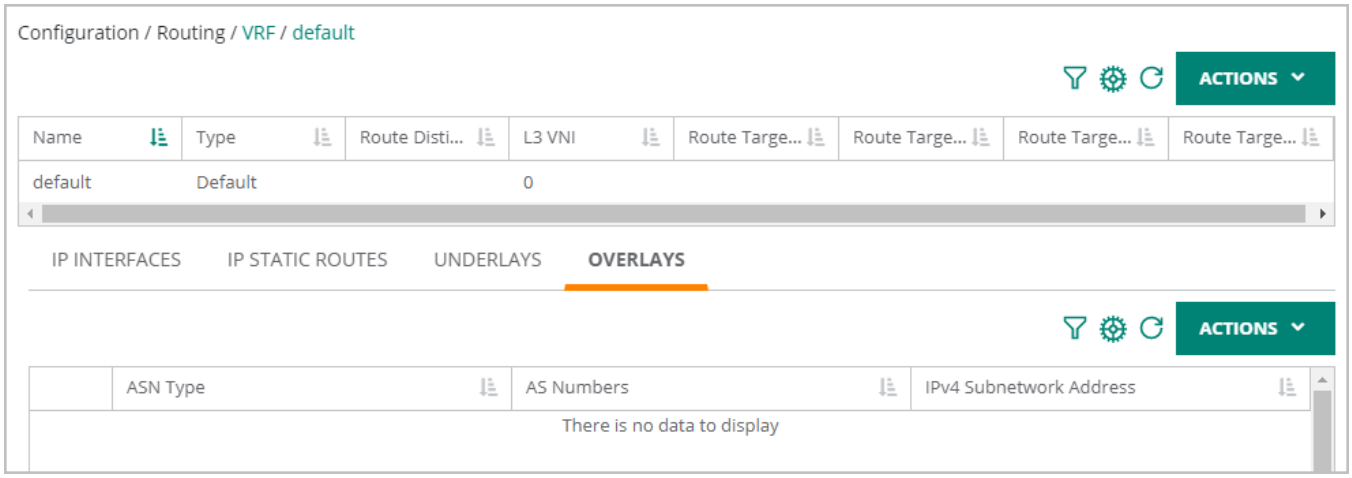

Step 2: From the **ACTIONS** menu on the right, select **Add**.

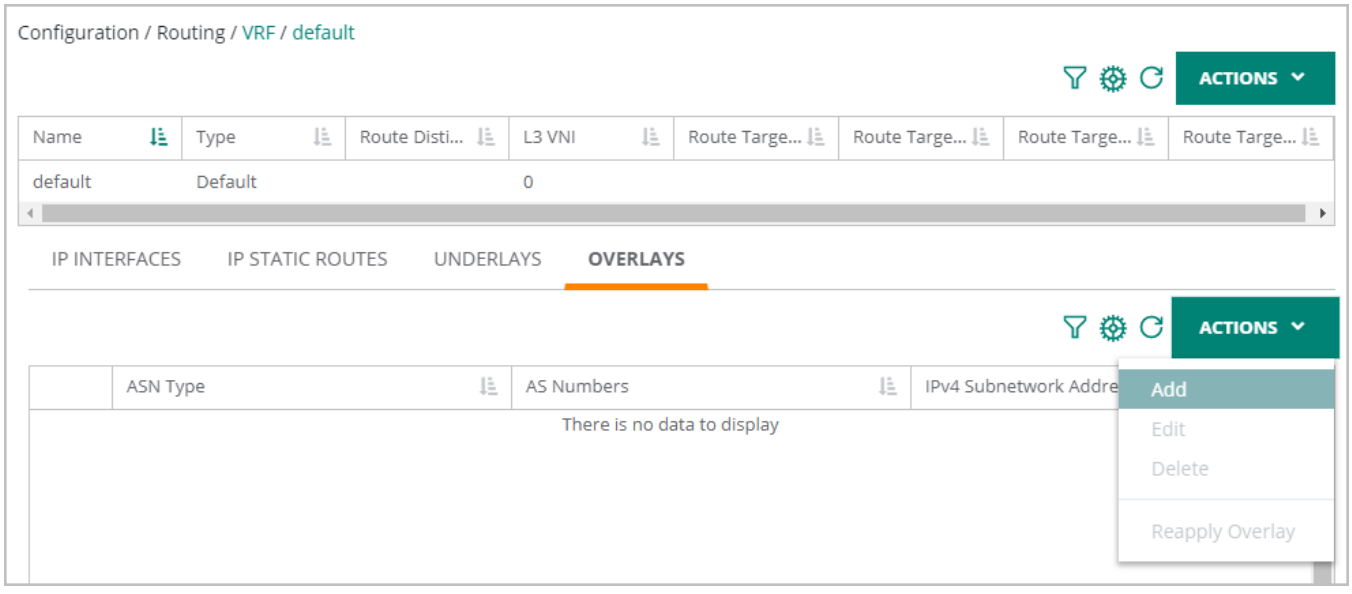

Step 3: From the Overlay Configuration AS Numbers page, implement the following settings, and then click **Next**.

- **Dual-ASN:** *Select*
- **Spine ASN:** *65001*
- **Leaf ASN:** *65000*

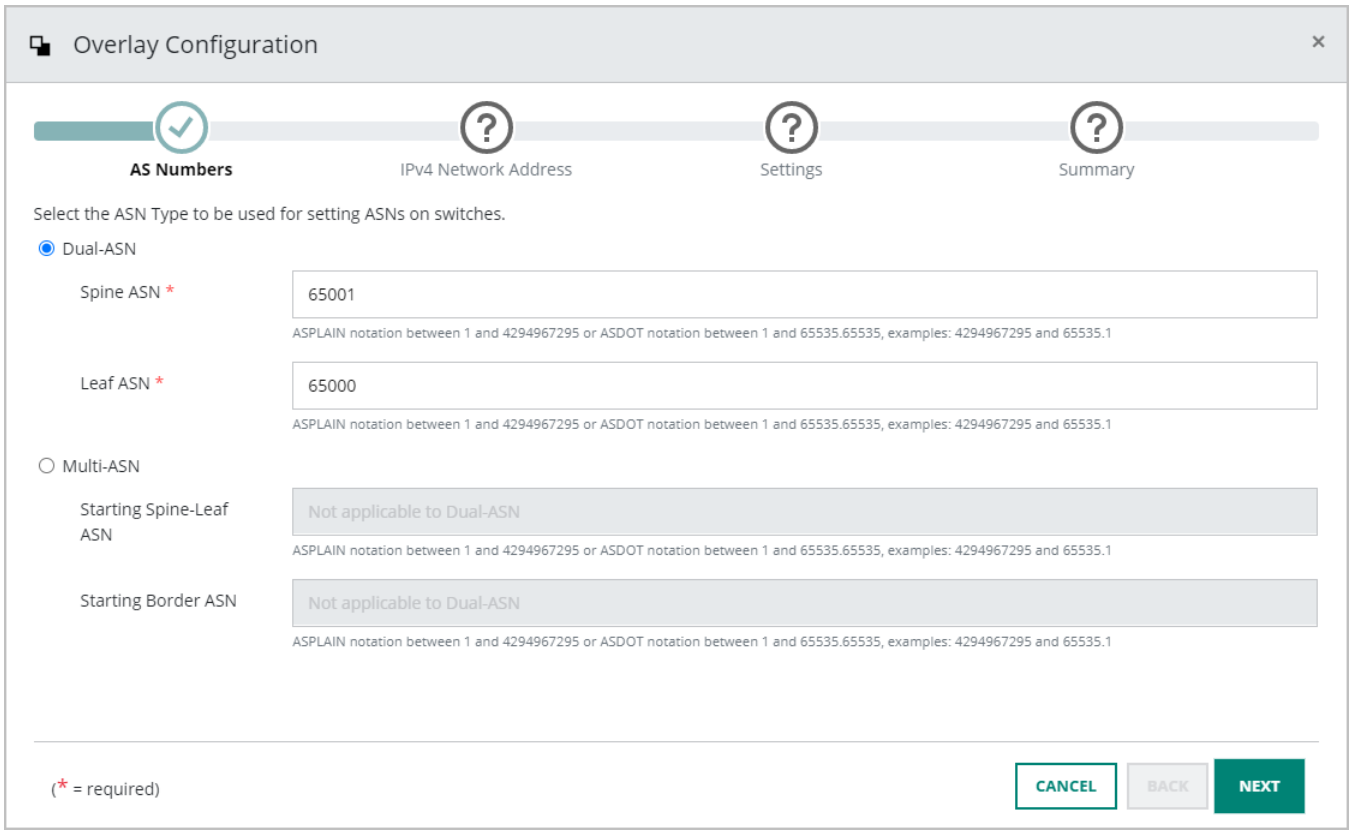

NOTE:
The use of a 2-byte private ASN (64512-65534) will make the configurations easier to parse.

Step 4: From the Overlay Configuration IPv4 Network Address page, enter an IPv4 Subnetnetwork Address, and then click **NEXT**.

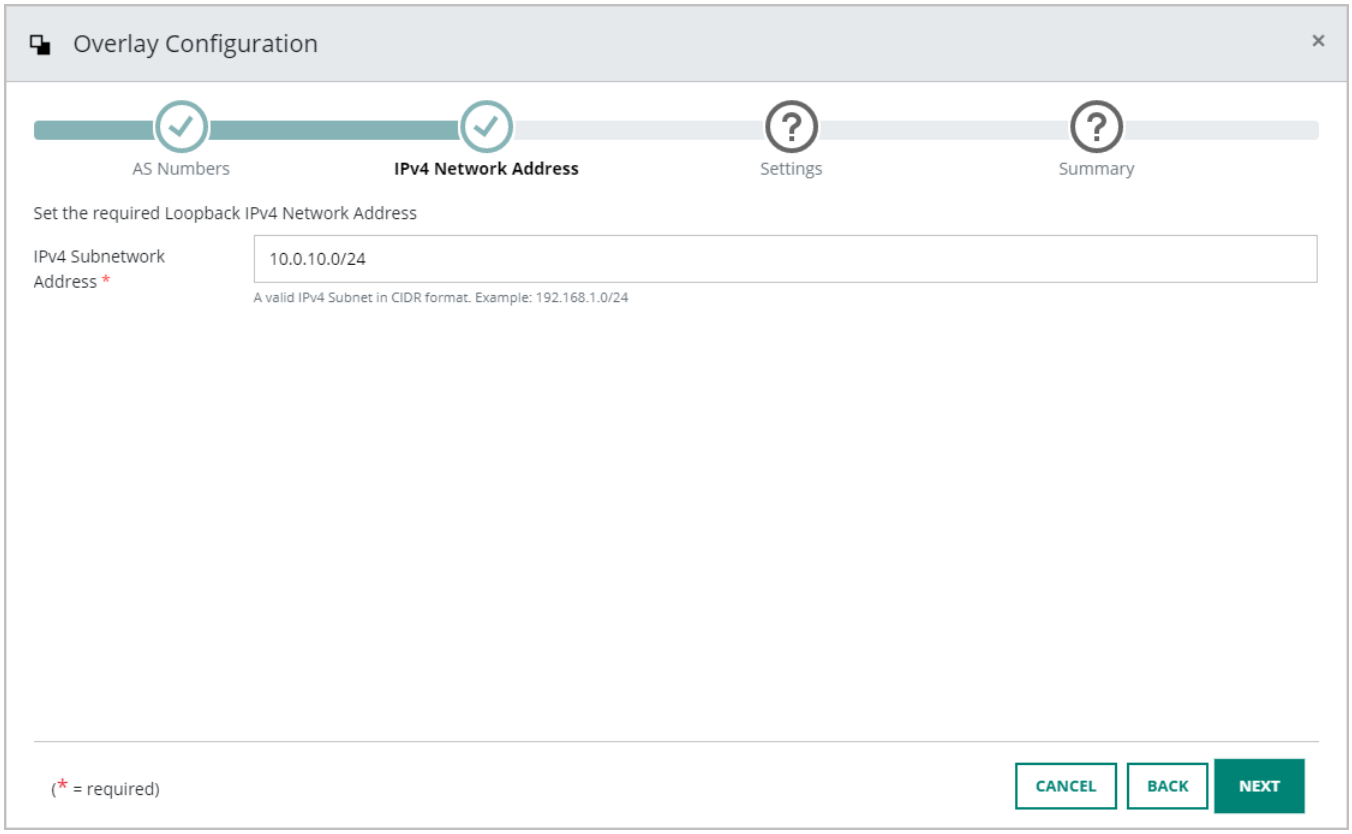

NOTE:
The IPv4 subnet address is used to configure VTEP addresses.

Step 5: From the Overlay Configuration Settings page, leave the default settings, and then click **NEXT**.

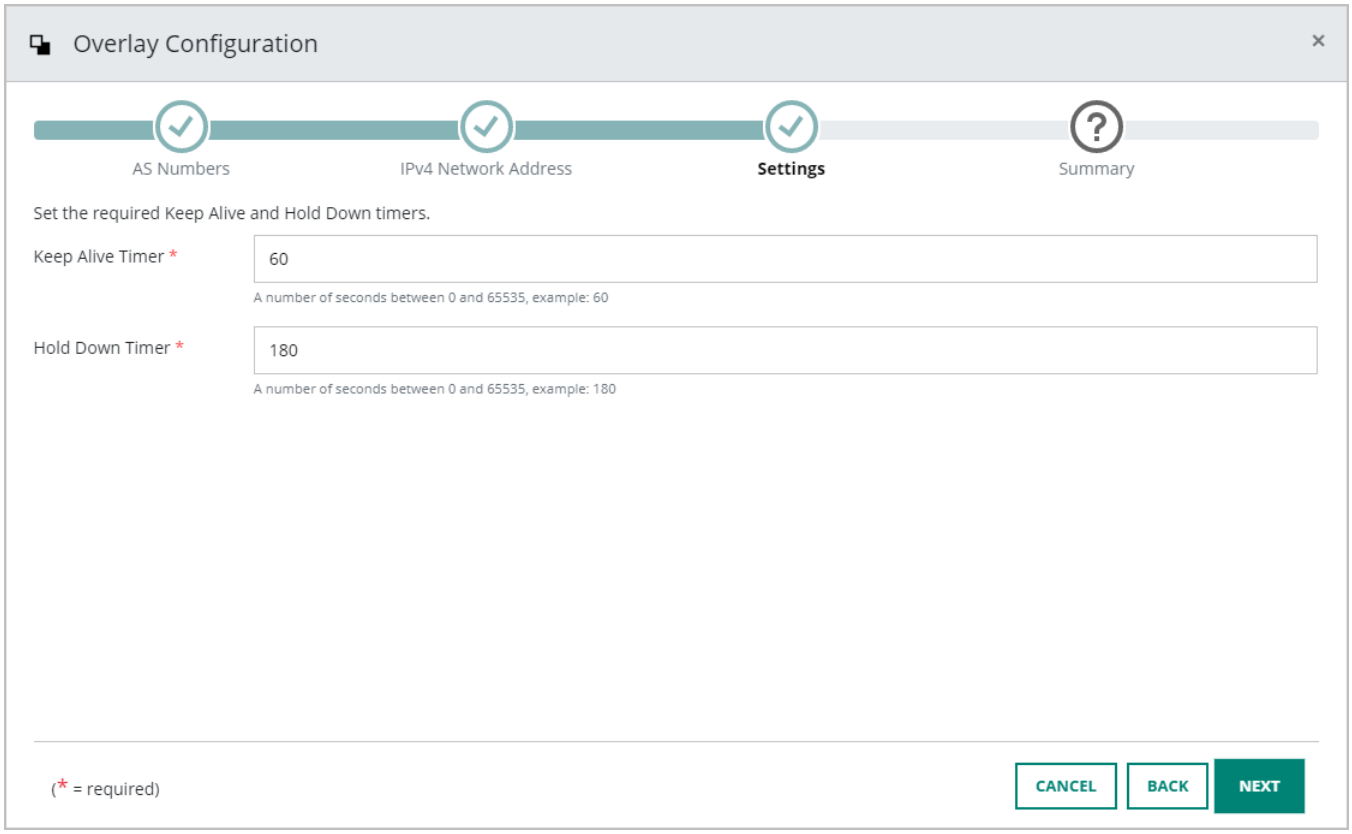

Step 6: Confirm the information was entered correctly and click **APPLY**.

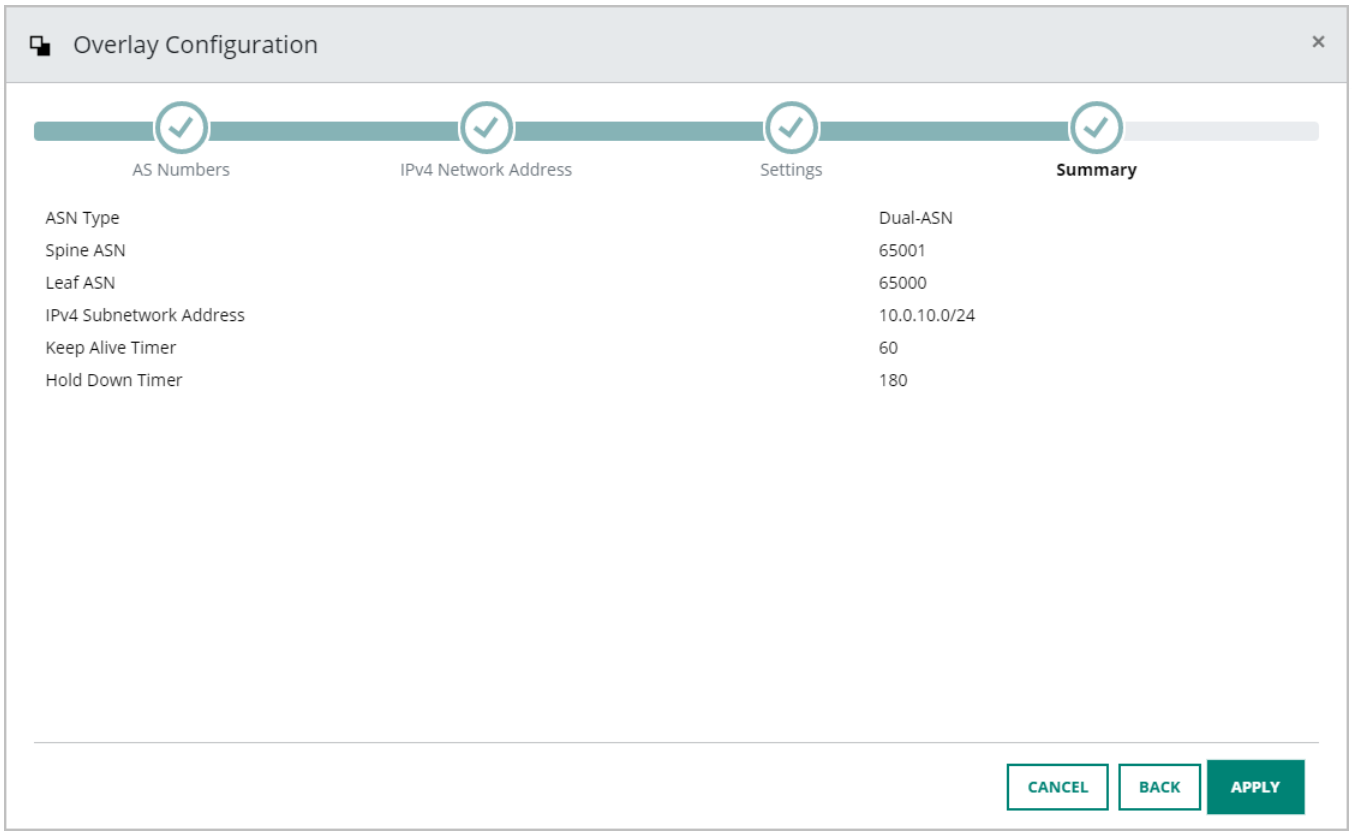

Step 7: Verify the spine switches have **Established** BGP neighbors to all leaf switches with the **show bgp all-vrf all summary** command.

```
8325-DC1-S1# sh bgp all-vrf all summary
VRF : default
BGP Summary
-----------
 Local AS               : 65001        BGP Router Identifier  : 10.0.5.11
 Peers                  : 4            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60
 Confederation Id       : 0

Address-family : IPv4 Unicast
-----------------------------

Address-family : IPv6 Unicast
-----------------------------

Address-family : L2VPN EVPN
-----------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State       AdminStatus
 10.0.5.6        65000       1325    1320    18h:51m:47s  Established   Up
 10.0.5.8        65000       1332    1327    18h:51m:32s  Established   Up
 10.0.5.9        65000       1329    1330    18h:51m:47s  Established   Up
 10.0.5.10       65000       1323    1324    18h:51m:47s  Established   Up
```

Configure Overlay VRFs

The next series of steps for deploying the fabric is to configure the overlay VRFs. An overlay VRF provides network virtualization and macro segmentation, necessary for flexible and secure data centers. In the Aruba ESP Data Center spine and leaf design each overlay VRF is associated to an L3VNI and one or more L2VNIs for host access. In the following steps, two overlay VRFs are created, an L3 VNI is created using a provided ID number, and an EVPN enabled route target is created. The VNIs and RTs should be unique to preserve separation and to be used as a filtering mechanism to enable route leaking if required.

Step 1: From the menu on the left, select **VRF**.

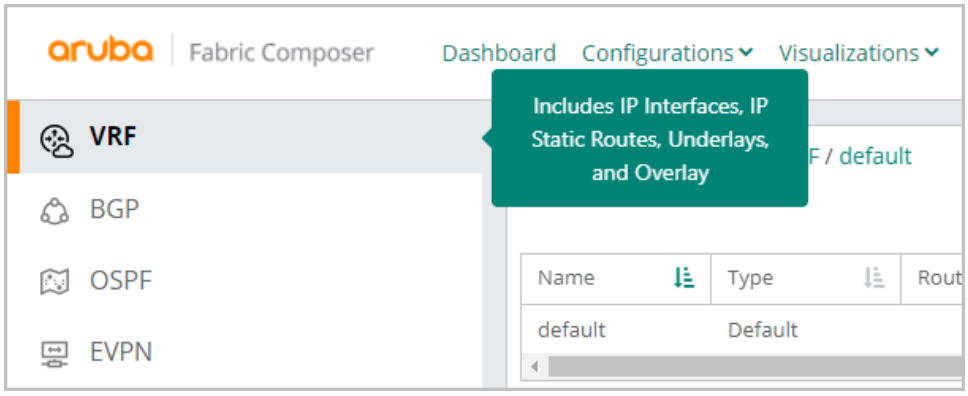

Step 2: Confirm that the view context is set to Configuration/Routing/VRF, and then from the **ACTIONS** menu on the right, select **Add**.

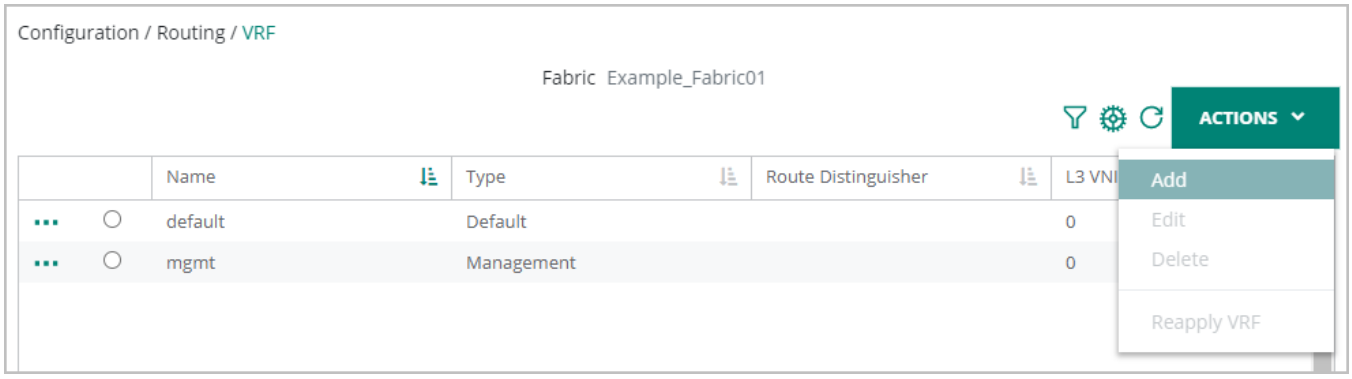

Step 3: From the Virtual Routing & Forwarding Name page, enter the name and description, and then click **NEXT**.

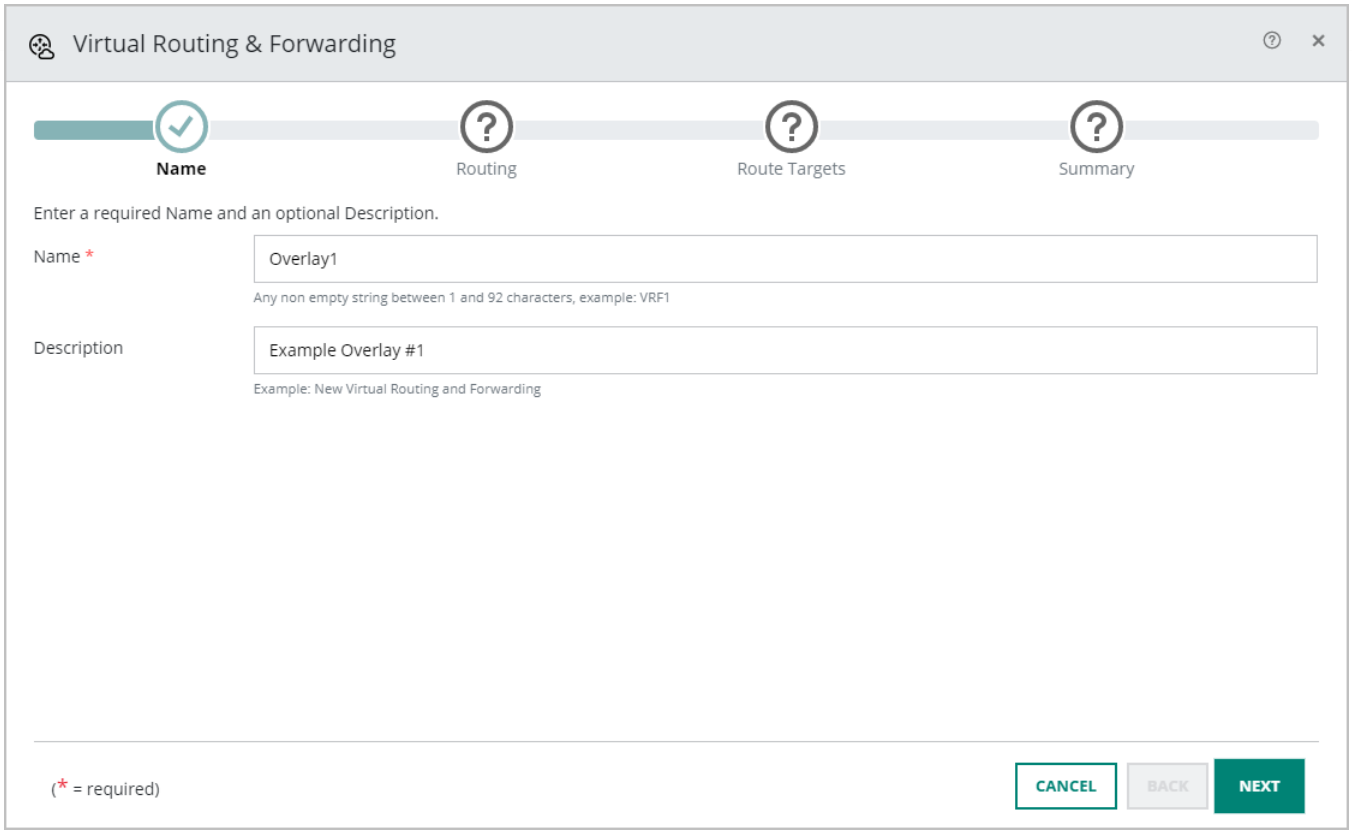

Step 4: From the Virtual Routing & Forwarding Routing page, enter an L3 VNI, and then click **NEXT**.

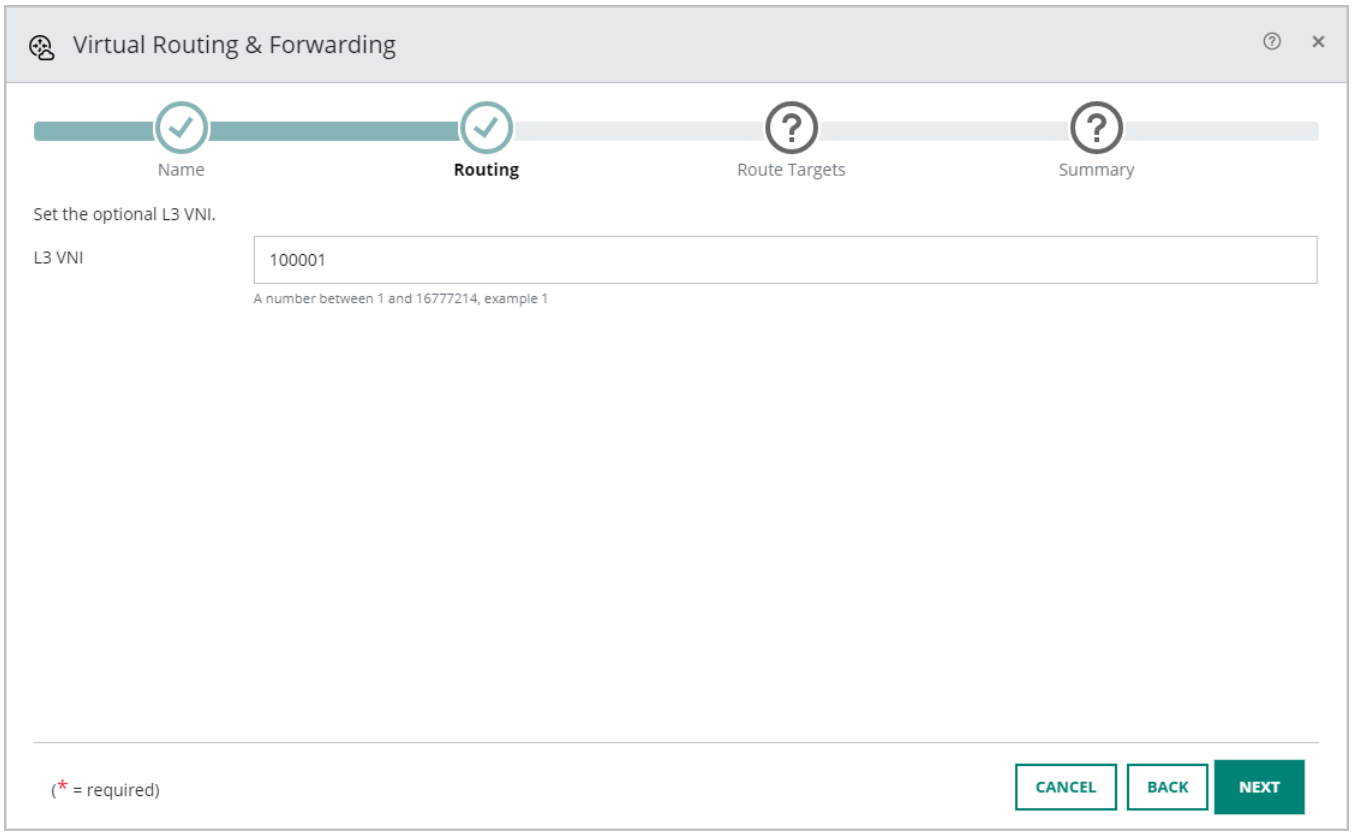

NOTE:
Refer back to Overlay Connectivity and Addressing section for a reference numbering approach.

Step 5: From the Virtual Routing & Forwarding Route Targets page, implement the following settings, and then click ADD

- **Route Target Mode:** *Both*
- **Route Target Ext-Community:** *65000:101*
- **Address Family:** *IPv4 Unicast*
- **Enable EVPN:** *checkmark*

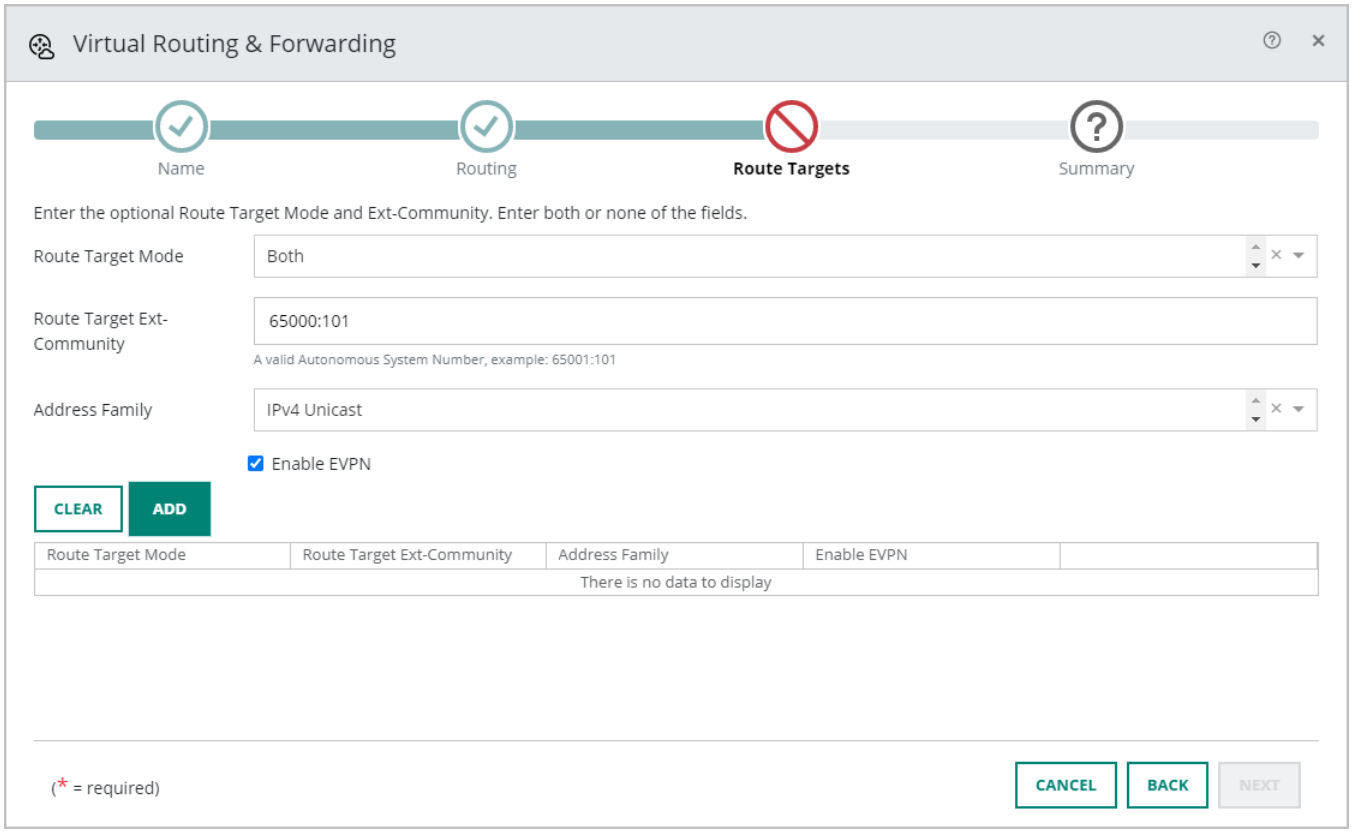

NOTE:
For the Route Target External Community, enter the private autonomous system number used in the leaf ASN and a unique ID you plan to use for the overlay separated by a colon.

Step 6: Verify the Route Target information is correct and click **NEXT**.

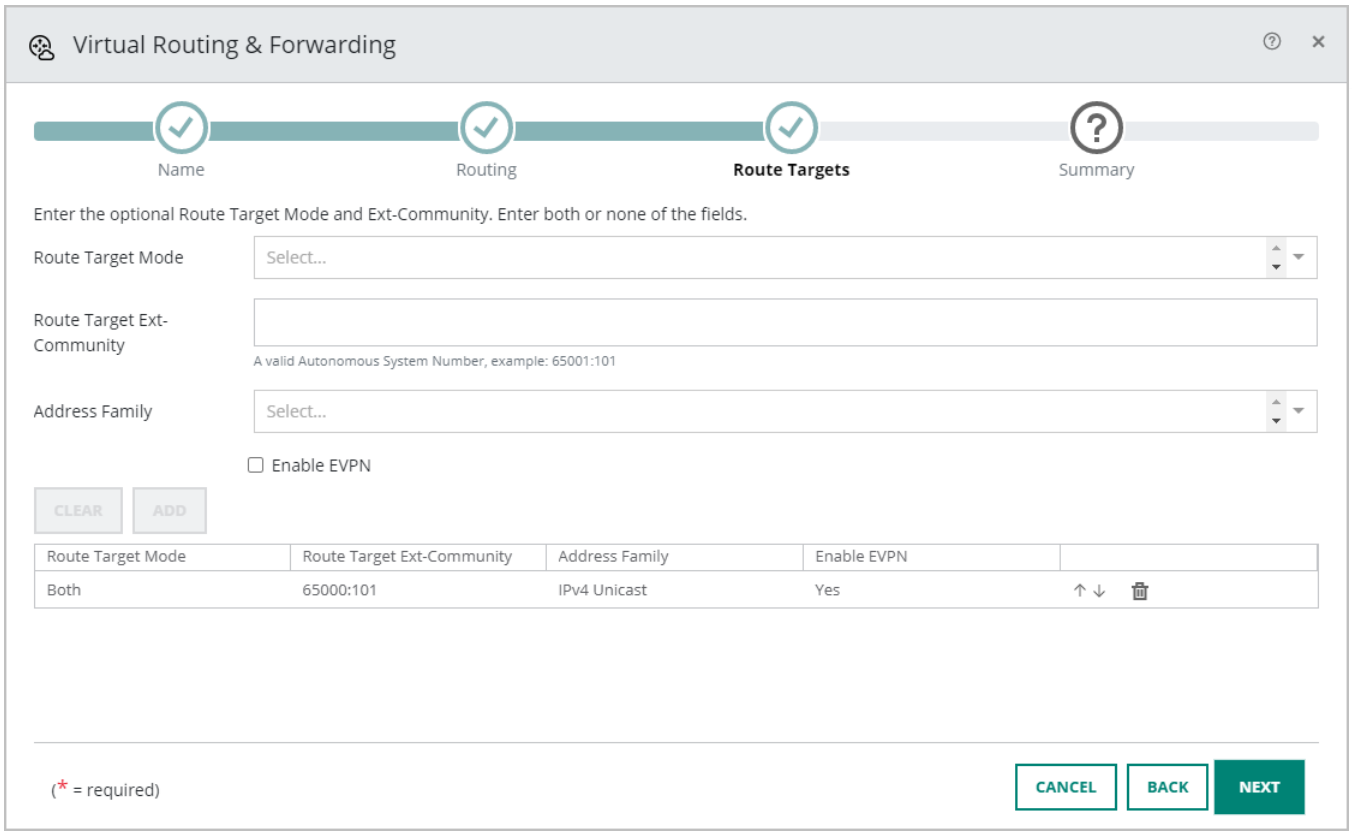

Step 7: Confirm the information was entered correctly and click **APPLY**.

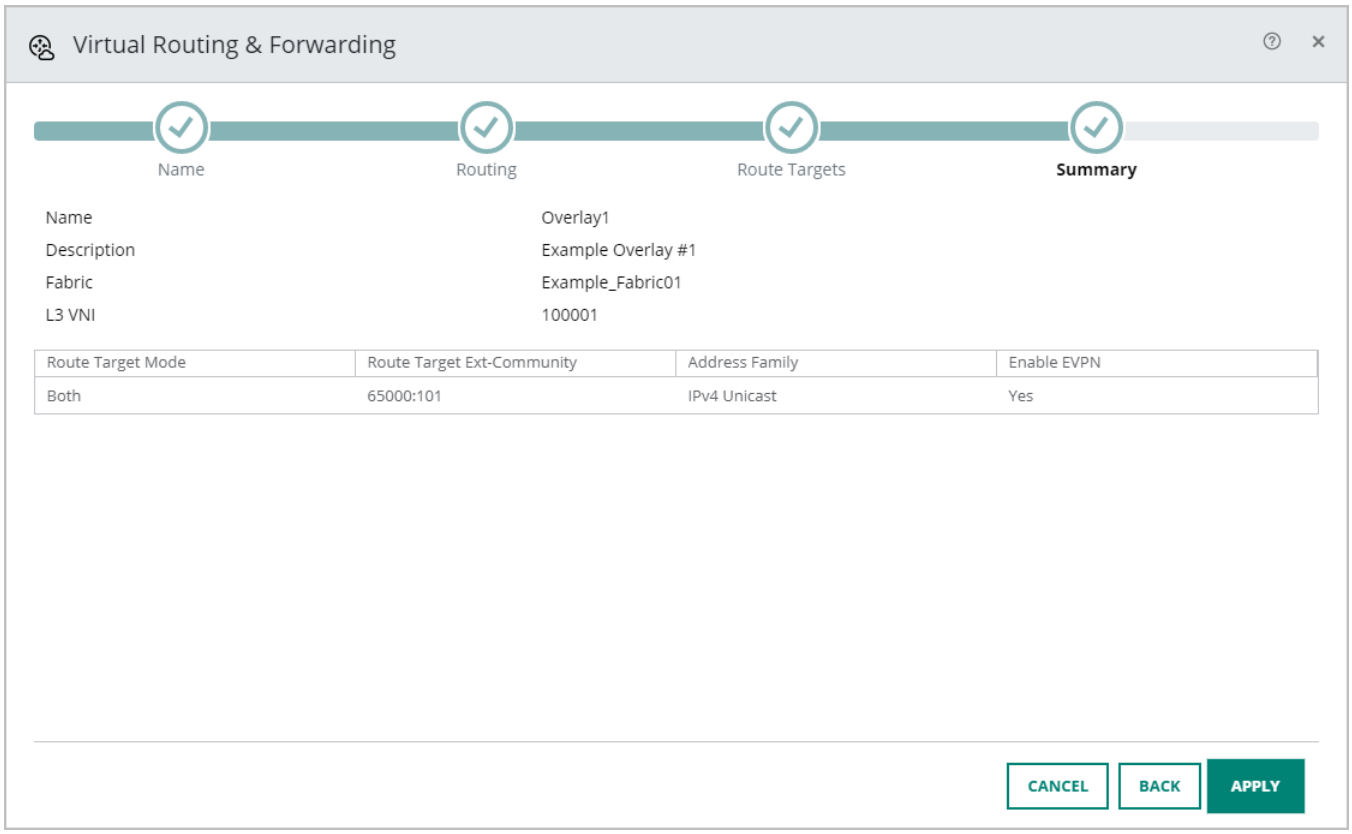

Step 8: Repeat this procedure for any additional overlay VRFs required.

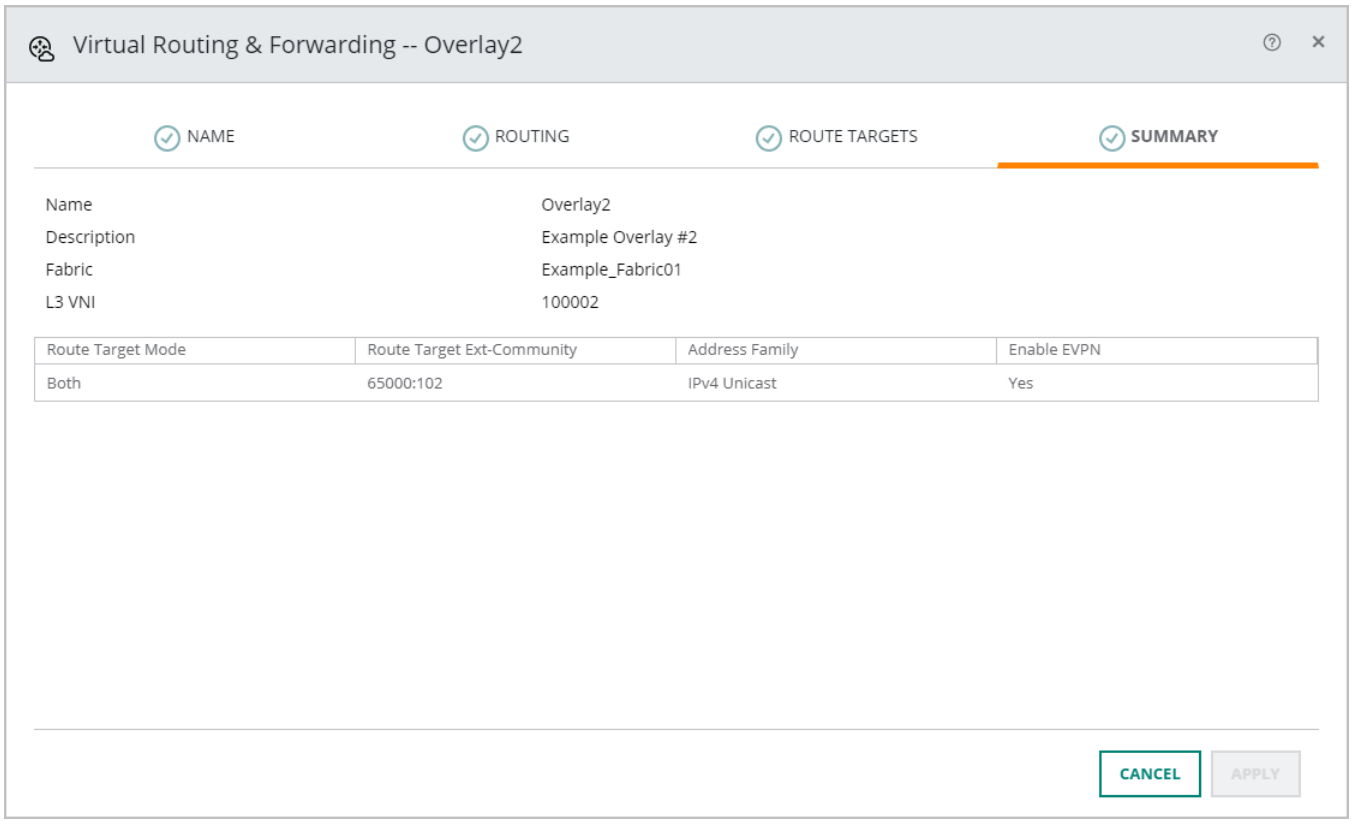

Configure Overlay VLANs and SVIs

The Aruba ESP Data Center uses one or more SVIs within each VRF to provide routed connectivity within the fabric. In the following steps, a Distributed SVI for an overlay VRF is created, the SVI is assigned an IP network range from which the active gateway and local SVI IP addresses are also assigned, and a locally administered MAC address is provided.

Step 1: Confirm that the view context is set to Configuration/Routing/VRF, and then from the **ACTIONS** menu on the right, select **IP interfaces**.

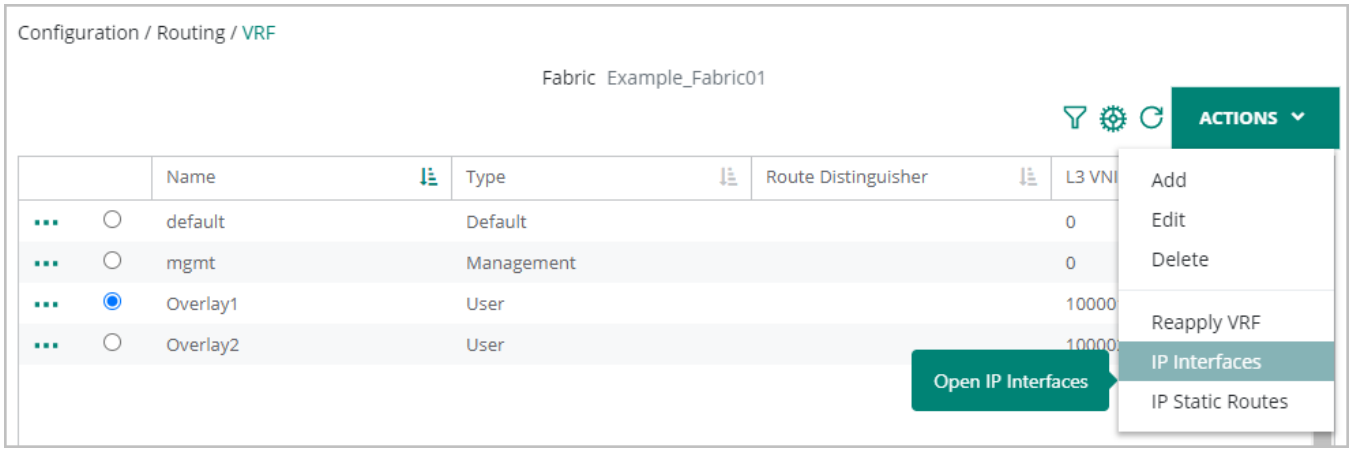

Step 2: From the Configuration/Routing/VRF/Overlay 1 page, select the **ACTIONS** menu and the right, and then click **Add**.

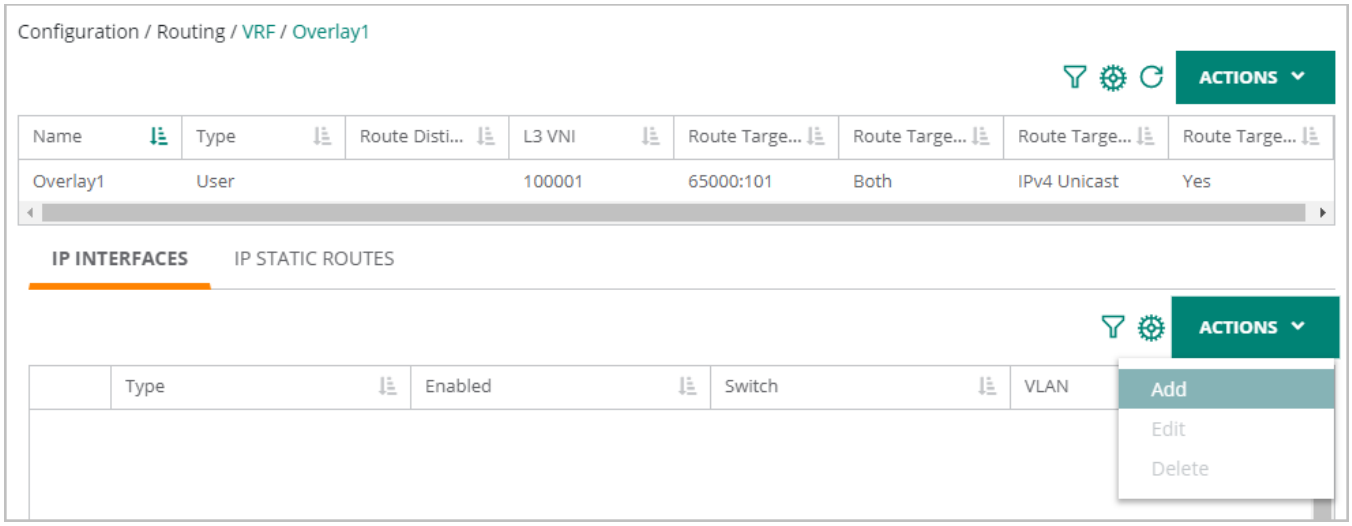

Step 3: From the IP Interface, Interface Type page, implement the following settings, and then click **NEXT**.

- **Type:** *Distributed SVI*
- **VLAN:** *100*
- **IPv4 Subnetwork Address:** *10.5.1.0/24*
- **Active Gateway IP Address:** *10.5.1.1*
- **Active Gateway MAC Address:** *02:00:10:05:01:01*
- **Switch Addresses:** *10.5.1.2-10.5.1.5*

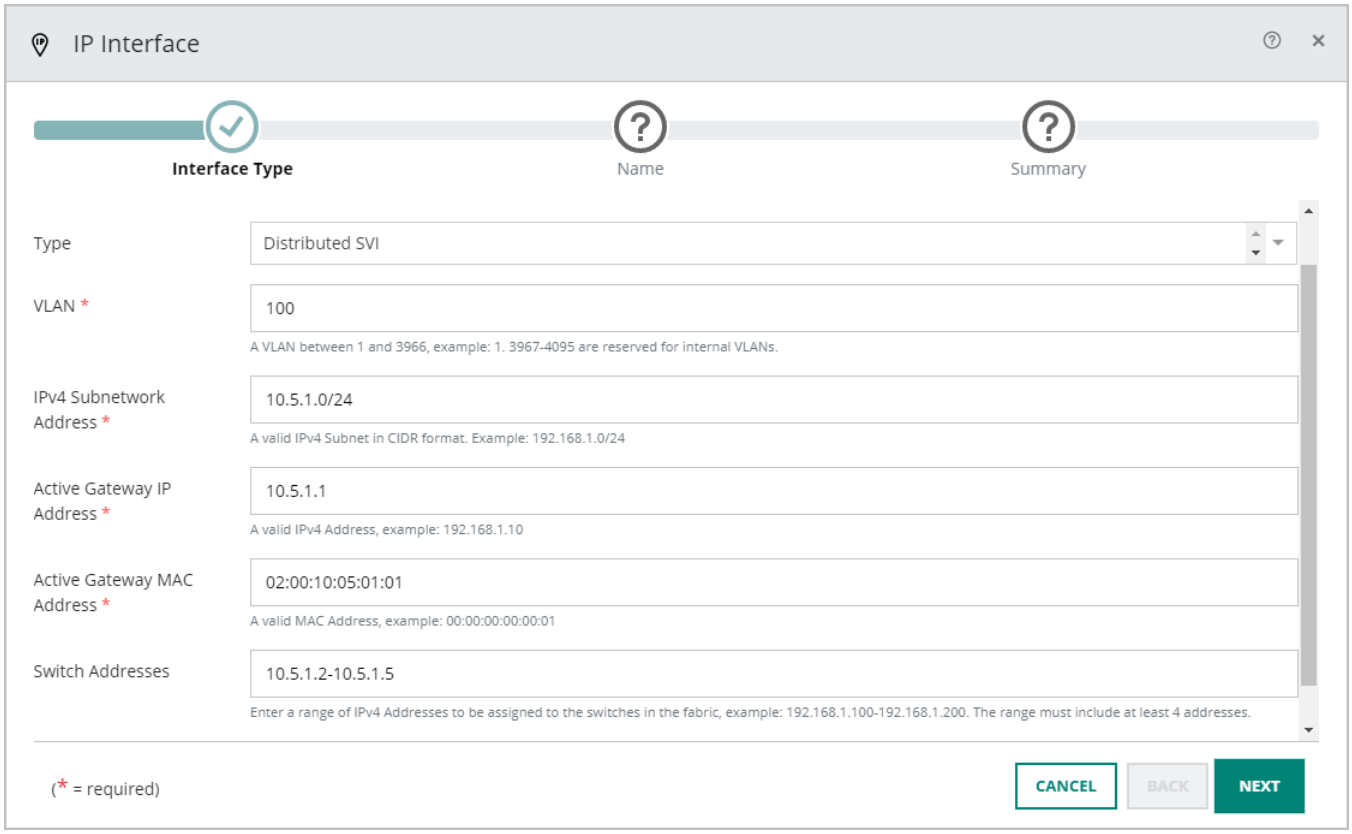

NOTES:
The range provided for **Active Gateway IP Address** and **Switch Addresses** must be from the same network range as the **IPv4 Subnetwork Address** provided initially above.
Active Gateway is the IP address which will be configured on each switch that will serve as default gateway for the subnet.
Active MAC is the MAC address which will be configured on each switch.
Switch addresses a range of addresses to be assigned as the switch unique IP for each SVI. Allocate a range that will support one IP address for every ToR.

Step 4: From the IP Interface Name page, enter the name and description, and then click **Next**.

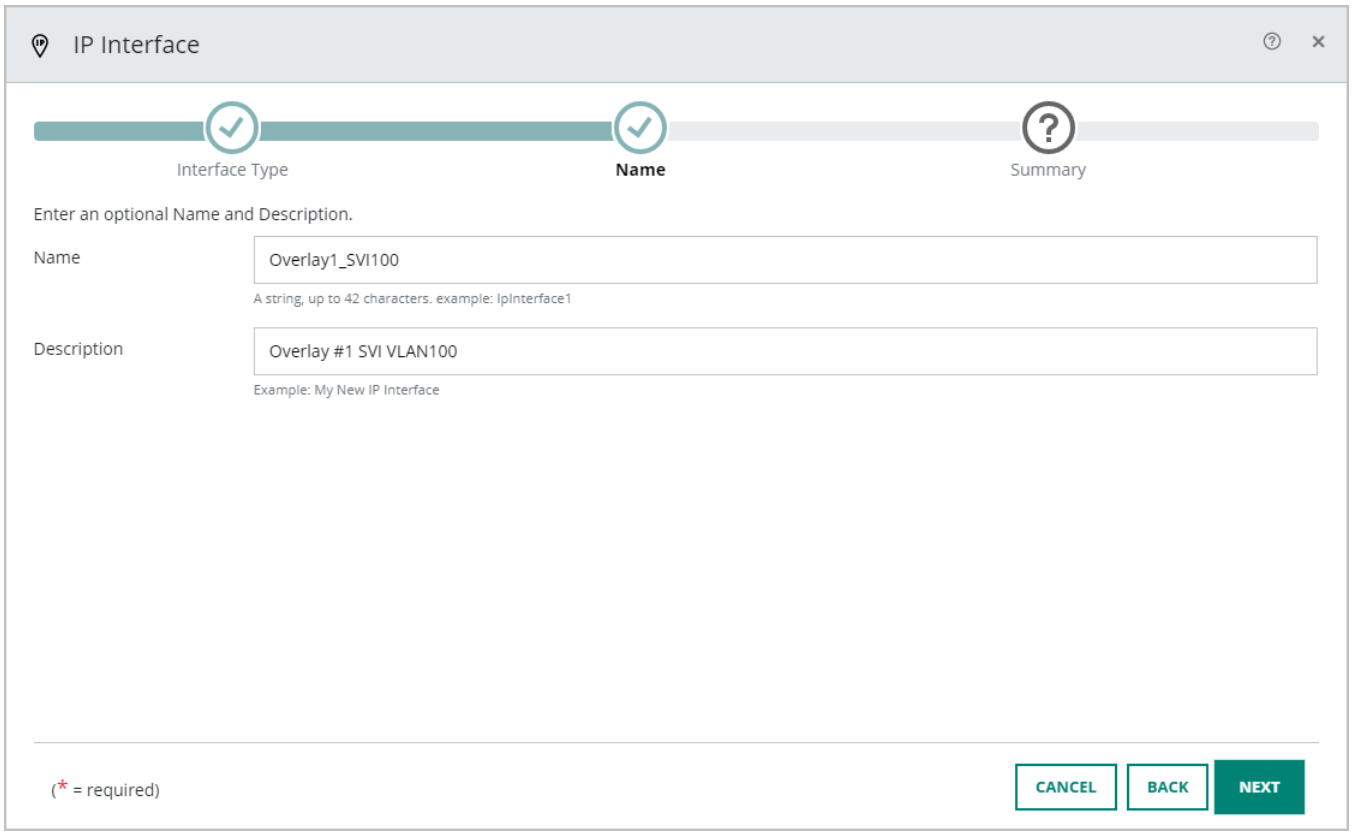

NOTE:
Including the VLAN ID in the Name can be helpful during management operations.

Step 5: Confirm the information was entered correctly and click **APPLY**.

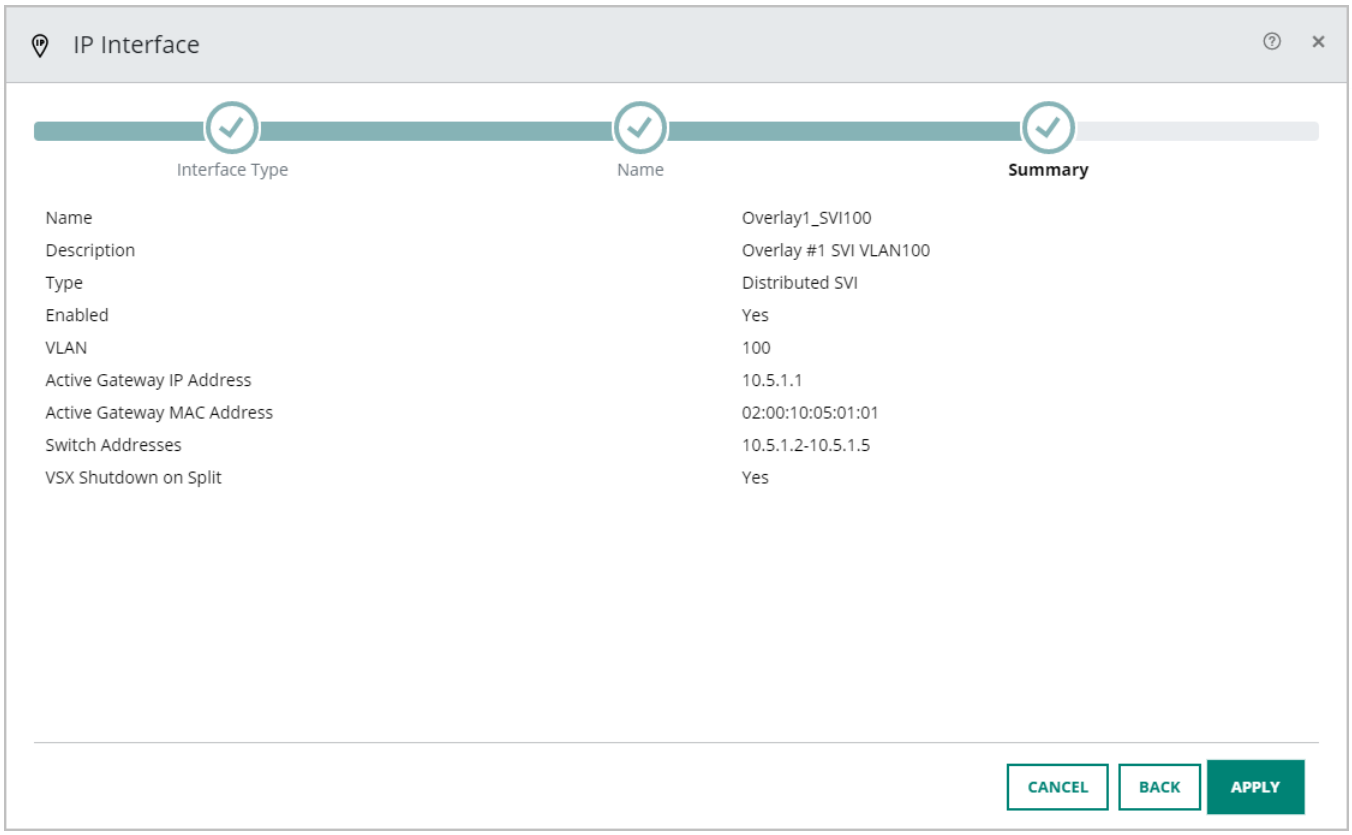

Step 6: Repeat this procedure for each additional overlay subnet required.

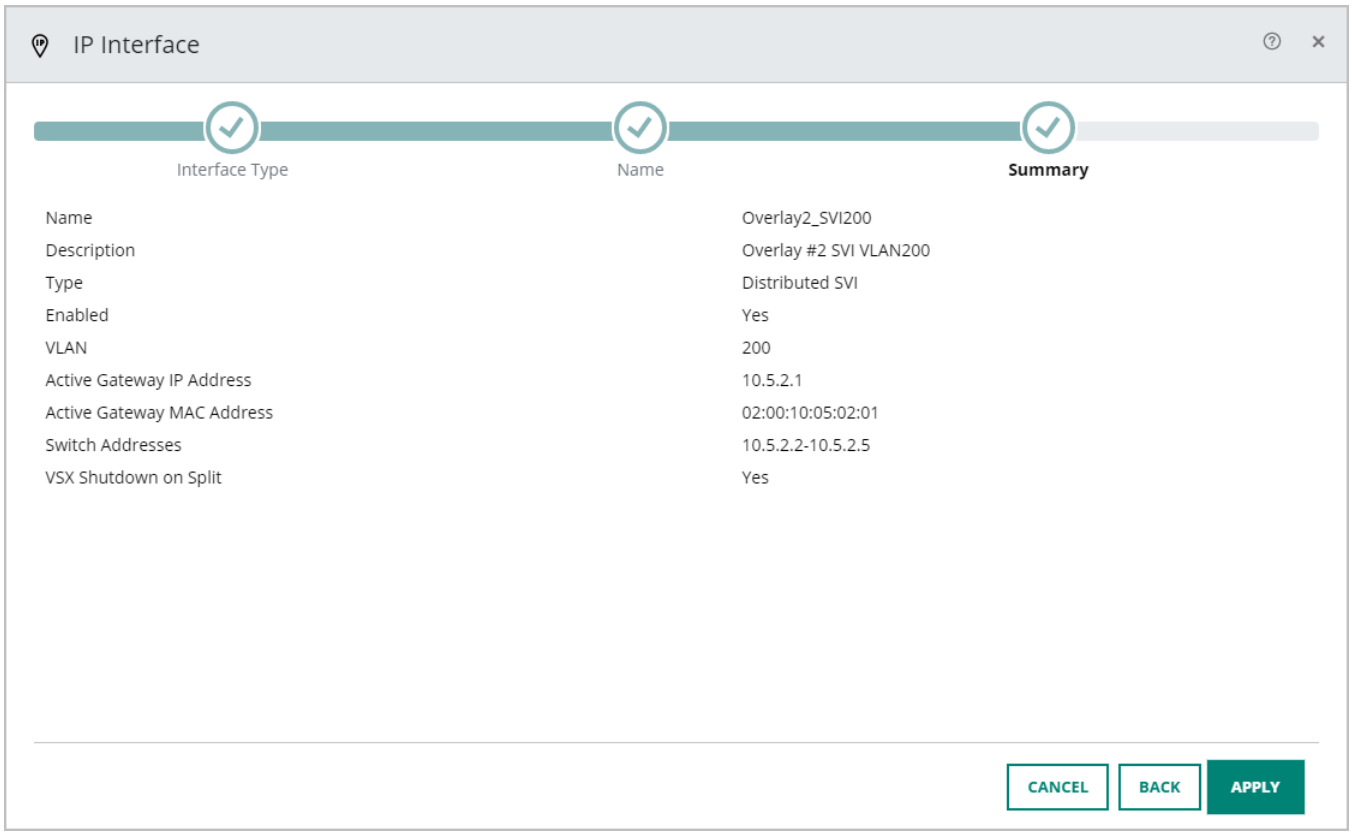

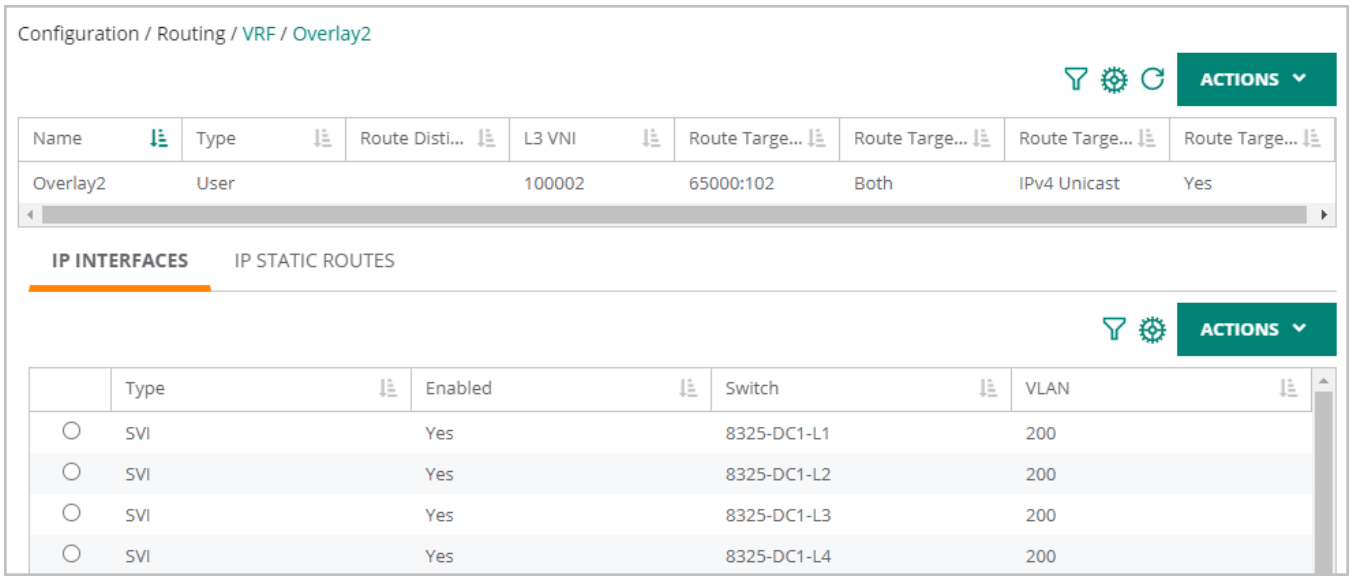

Configure EVPN Instances

In order to dynamically provide host reachability information across the overlay fabric, the Aruba ESP Data Center uses the EVPN address family of MP-BGP to communicate the location of an IP address. In the steps below, an EVPN instance is created, VLAN IDs are mapped to automatically generated L2 VNI values so they map to overlay networks. A prefix value is provided for automatic generation of route targets.

Step 1: From the Guided Setup menu, select **EVPN CONFIGURATION**.

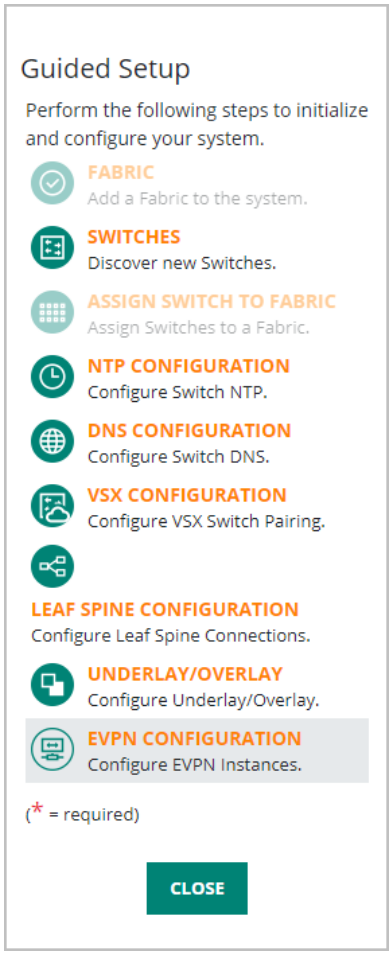

Step 2: Observe the provided guidance and click **NEXT**.

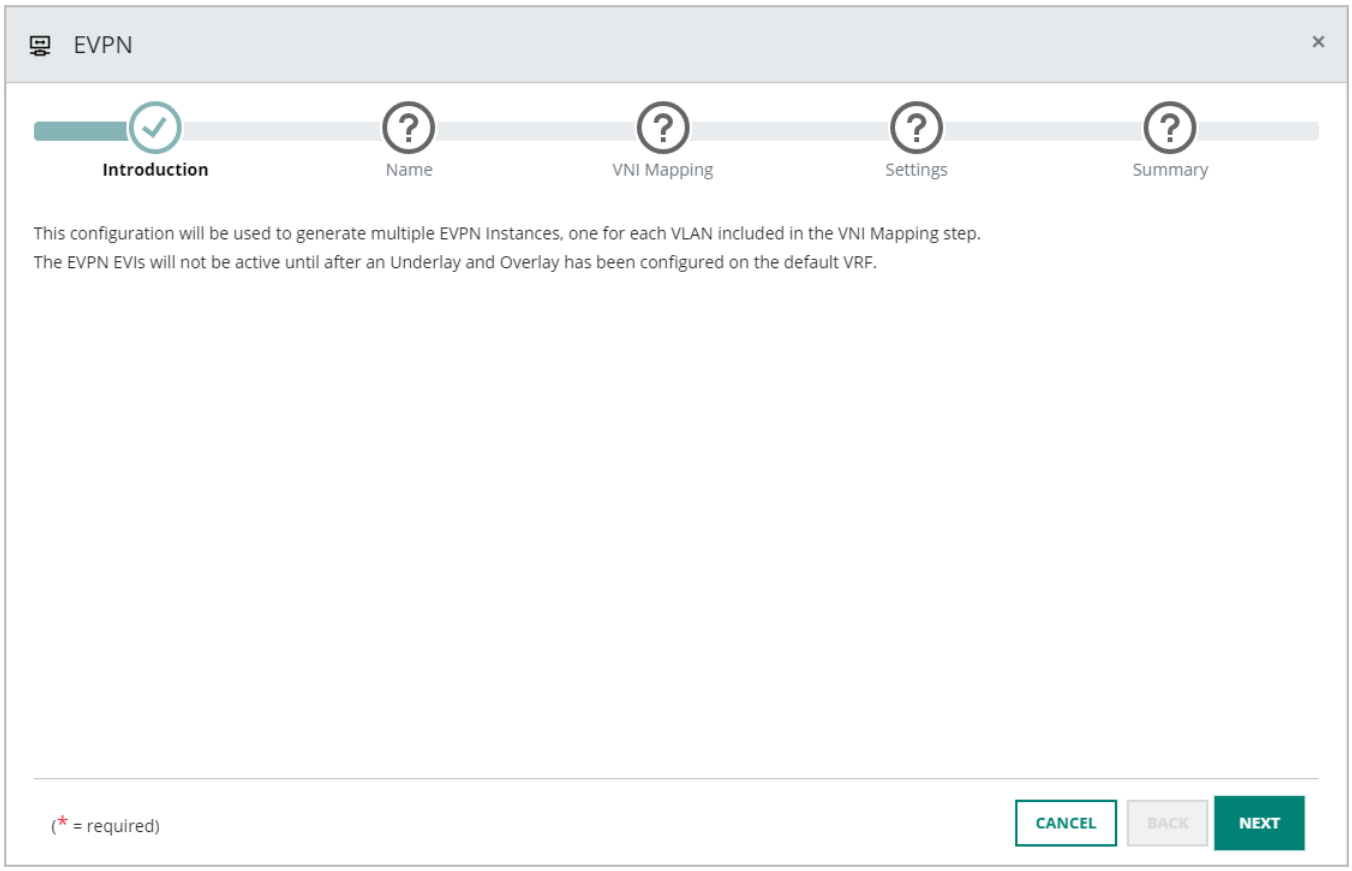

Step 3: From the EVPN Name page, enter a name prefix and description, and then click **NEXT**.

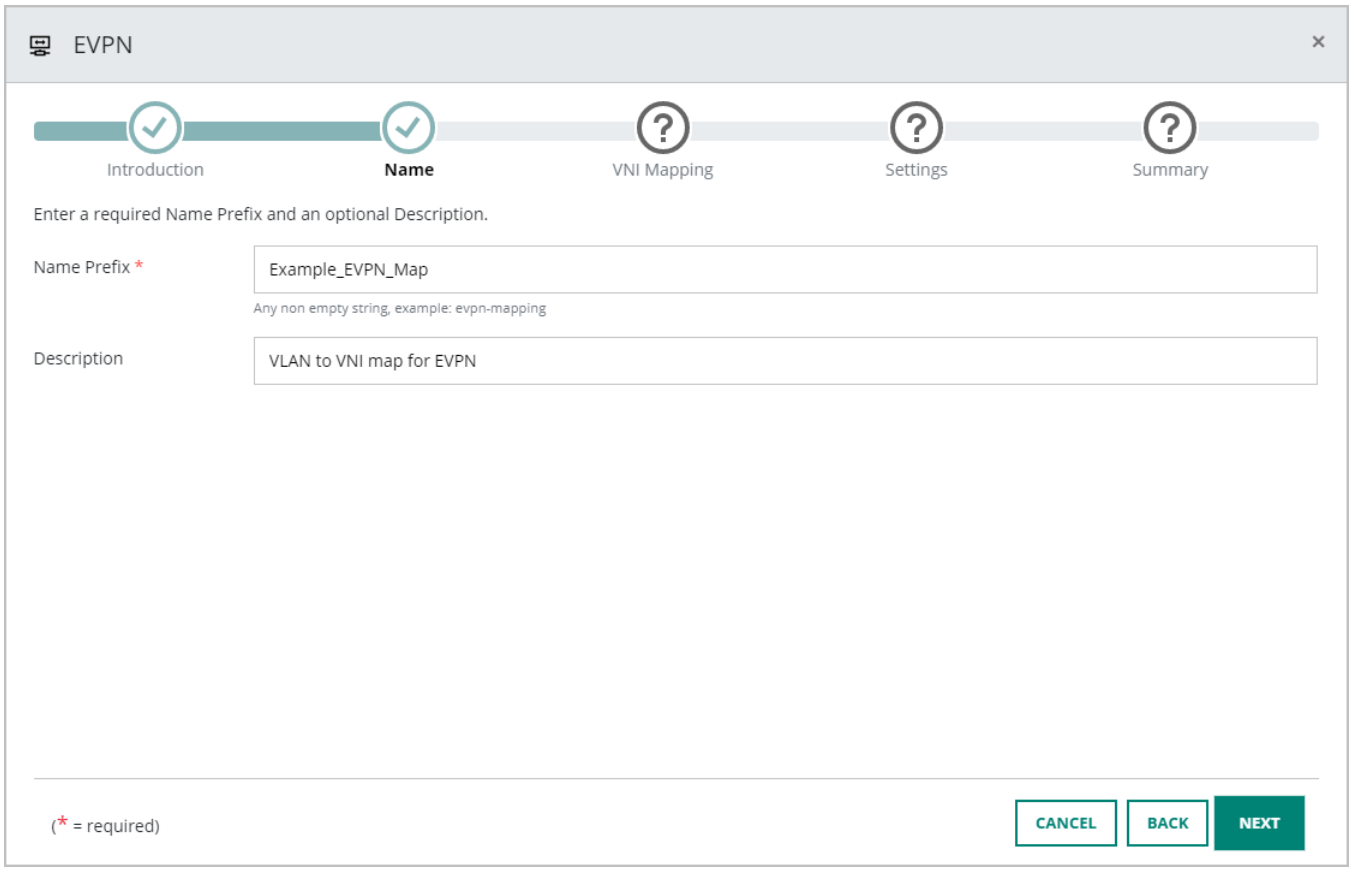

Step 4: From the EVPN VNI Mapping page, enter one or more VLANs and a Base L2VNI, and then click **NEXT**.

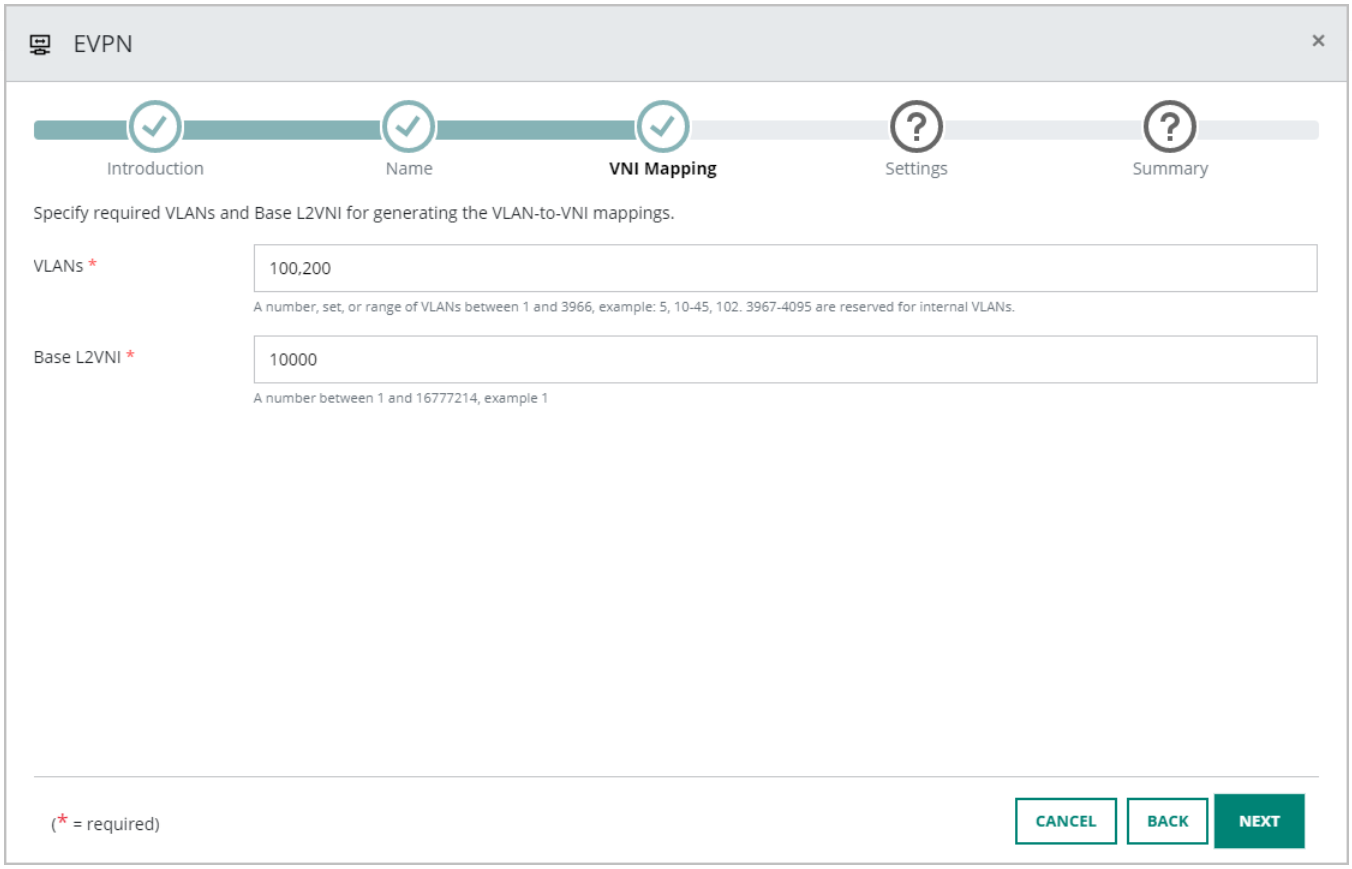

NOTE:
The **Base L2VNI** value which will be the base for automatically generated VNIs.

Step 5: From the EVPN Setting page, enter Route Target ASN Prefix, accept the default for System MAC Address Range, and then click **NEXT**.

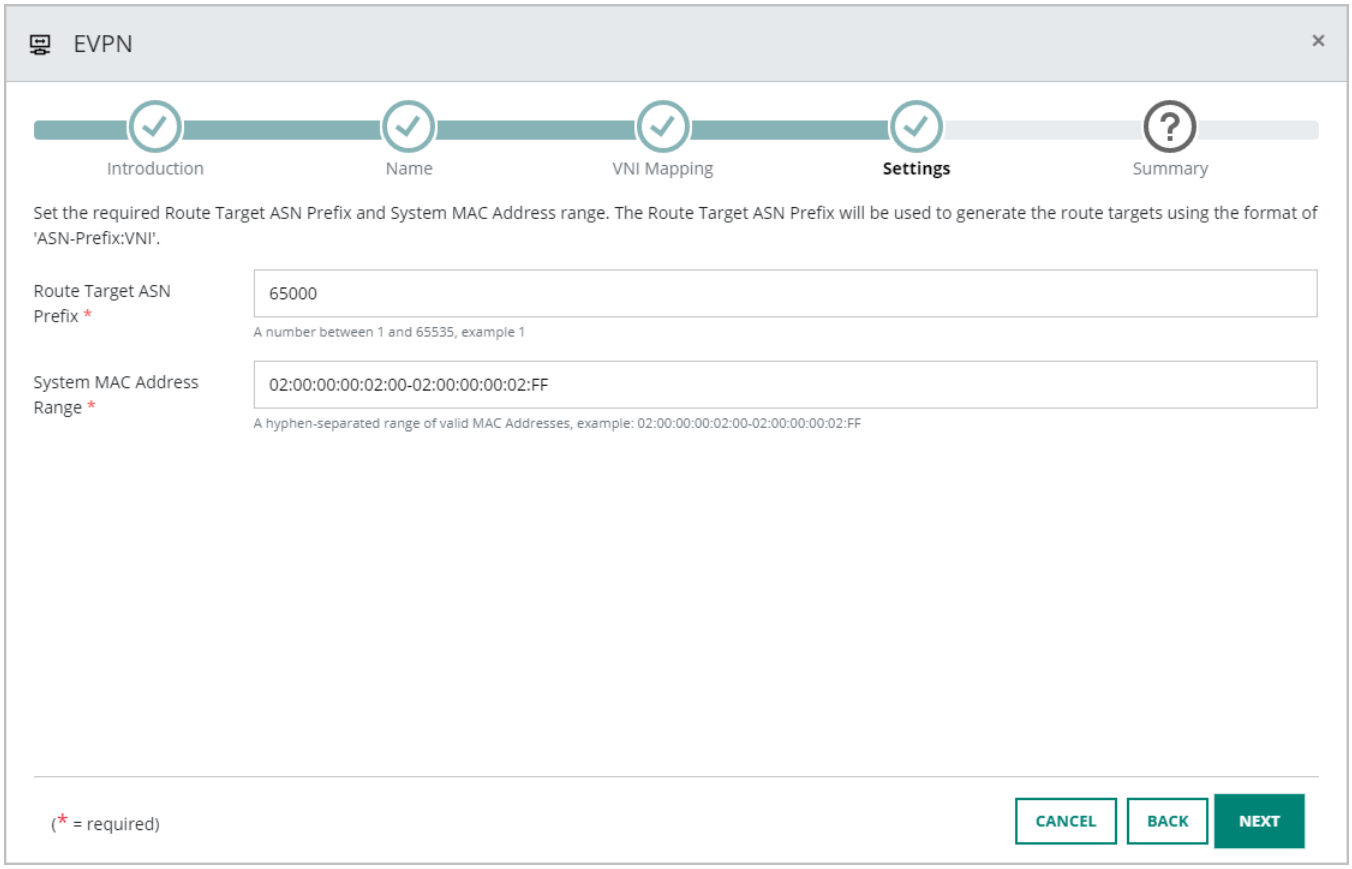

NOTE:
The Route Target ASN Prefix should be the same or similar to the ASN created during the BGP configuration.

Step 6: Confirm the information was entered correctly and click **APPLY**.

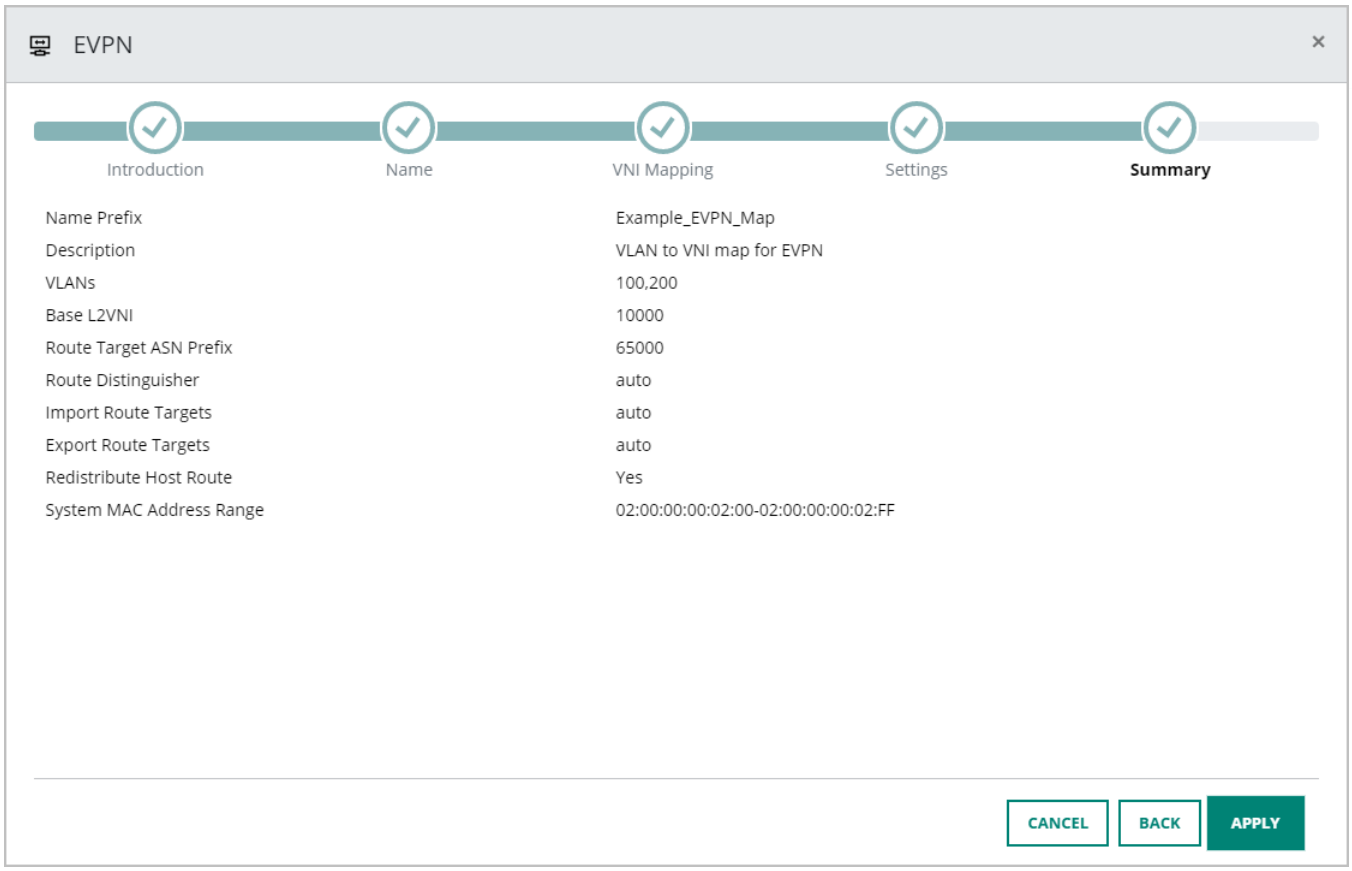

Step7: The Guided Setup is now complete.

Border Leaf Configuration

The border leaf is the ToR switch pair connecting the data center fabric to other networks such as a campus, WAN, or DMZ. Inside the data center, overlay networks provide multi-tenancy and segmentation while supporting overlapping IP address ranges.

When connecting overlay networks to external networks there are three fundamental design approaches to handling segmentation at the border:

- Keep the overlay network restricted to the data center, no external connectivity to campus and/or WAN required.
- Preserve overlay network isolation outside the data center using a 1:1 mapping of an overlay network to an external VRF and VLAN(s).
- Combine overlay networks outside the data center using a N:1 mapping of multiple overlay networks to a single IGP and/or BGP network for external connectivity.

The following diagram illustrates the topology orchestrated using Aruba Fabric Composer and the detailed connectivity between the border leaf and campus core switches.

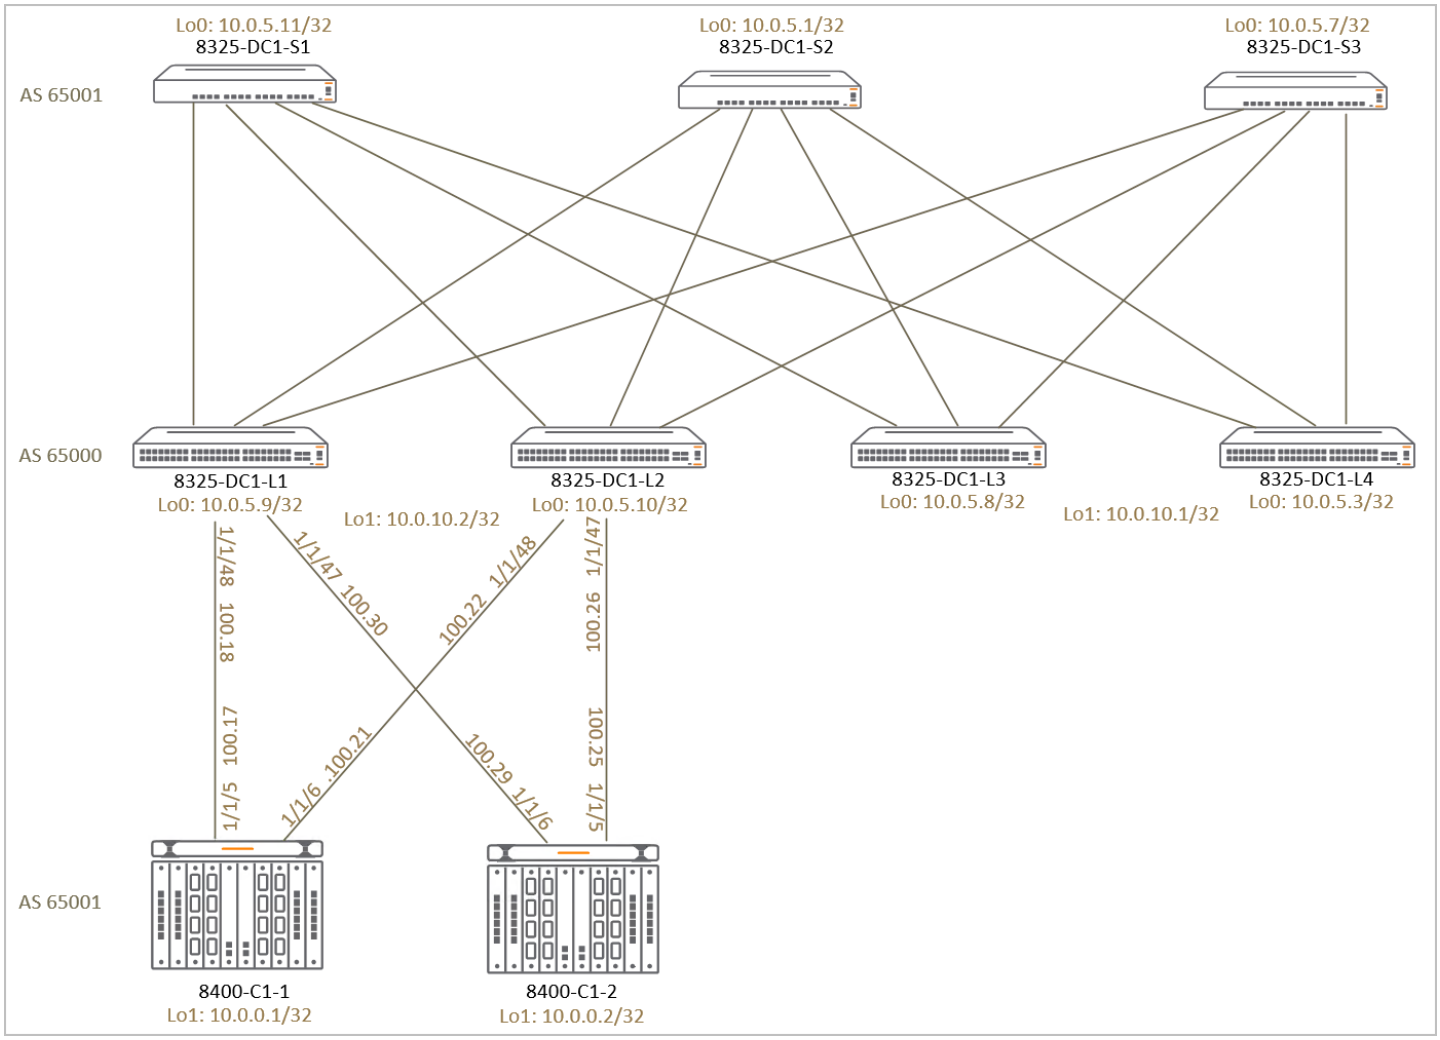

When using MP-BGP with EVPN address families, type-2 EVPN routes identify host reachability while type-5 routes identify network reachability inside the data center. EVPN host routes should be summarized when propagated outside of the data center. When combining multiple overlay VRFs into a single IGP and/or BGP network the use of Inter-VRF Route Forwarding (IVRF) is required. This allows routes to dynamically propagate from data center overlay VRFs into a common external VRF.

In the configuration steps below, a new VRF named Campus is created, eBGP routing is established between the campus and data center networks and route redistribution between Campus and overlay VRFs is established. This example configuration is for the design option in which the virtual networks in the data center are combined into the global routing table in the campus.

Step 1: Create a new VRF called Campus on each border leaf switch. Assign an RD value that is unique and use the same value on both border leaf switches.

```
vrf Campus
    rd 10.0.10.2:100
```

NOTE:
A Route Distinguisher (RD) is used to identify which routes belong to a VRF and must be a unique value assigned to each VRF. Aruba Fabric Composer uses the VTEP IP with a unique sequence number for assigning RD values to VRFs

Step 2: Redistribute routes between the Campus VRF and the Overlay VRFs on each border leaf switch. Notice the use of import RT values to dynamically redistribute prefixes from both overlays into the campus and vice versa. Overlay VRFs do not import RT values between each other to preserve separation of routing tables inside the data center.

```
vrf Campus
    address-family ipv4 unicast
        route-target export 65000:1
        route-target import 65000:101
        route-target import 65000:102
     exit-address-family


vrf Overlay1
    address-family ipv4 unicast
       route-target import 65000:1
    exit-address-family


vrf Overlay2
    address-family ipv4 unicast
       route-target import 65000:1
    exit-address-family
```

NOTE:
Route targets (RT) are BGP extended communities that are associated with prefixes inside a VRF and are used for establishing filtering rules.

Step 3: Configure the IP interface on the border leaf switches which will connect to the Campus core switches. Add the interfaces to the Campus VRF.

```
interface 1/1/47
    no shutdown
    mtu 9198
    vrf attach Campus
    ip mtu 9198
    ip address 172.18.100.30/30


interface 1/1/48
    no shutdown
    mtu 9198
    vrf attach Campus
    ip mtu 9198
    ip address 172.18.100.18/30
```

CAUTION:
Order of operations is important, the IP address must be configured after attaching the VRF to the interface.

Step 4: Configure the IP interface on the core switches which will connect to the Data Center border leaf switches.

```
interface 1/1/5
    no shutdown
    mtu 9198
    description 8325-DC1-1 1/1/1
    ip mtu 9198
    ip address 172.18.100.17/30


interface 1/1/6
    no shutdown
    mtu 9198
    description 8325-DC1-2 1/1/1
    ip mtu 9198
    ip address 172.18.100.21/30
```

NOTE:
Unlike the border leaf, there is no need to create and attach a VRF on the campus core in this non-virtualized campus example.

Step 5: Configure a prefix list on the border leaf switches that will match /32 subnet mask. This will be used in the next step to create a route map that will prevent EVPN type-2 host prefixes from getting propagated into the campus as host routes.

```
ip prefix-list PL_DC-Campus seq 10 permit 0.0.0.0/0 ge 32
```

Step 6: Configure a route-map on the border leaf switches that will deny /32 subnet mask. This will be used in the next step to attach to the BGP neighbor definition

```
route-map DC-Campus deny seq 10
     match ip address prefix-list PL_DC-Campus
route-map DC-Campus permit seq 20
```

Step 7: Configure eBGP peering from the border leaf switches to campus core routers.

```
router bgp 65000
    neighbor 172.18.100.17 remote-as 65001
    neighbor 172.18.100.29 remote-as 65001
!
    vrf Campus
        bgp router-id 10.0.5.10
        neighbor 172.18.100.17 remote-as 65001
        neighbor 172.18.100.29 remote-as 65001
        address-family ipv4 unicast
            neighbor 172.18.100.17 activate
            neighbor 172.18.100.17 route-map DC-Campus out
            neighbor 172.18.100.17 send-community extended
            neighbor 172.18.100.29 activate
            neighbor 172.18.100.29 route-map DC-Campus out
            neighbor 172.18.100.29 send-community extended
        exit-address-family
```

NOTE:
BGP router ID should be the unique loopback address of each switch. Use the IP addresses of the interfaces between border leaf switches and campus core for eBGP peering. Make sure to apply the route-map created in the previous step to the eBGP neighbors.

Step 8: Configure eBGP peering from the campus core to the border leaf switches.

```
router bgp 65001
    bgp router-id 10.0.0.1
    bgp fast-external-fallover
    neighbor 172.18.100.18 remote-as 65000
    neighbor 172.18.100.22 remote-as 65000
    address-family ipv4 unicast
        neighbor 172.18.100.18 activate
        neighbor 172.18.100.18 default-originate
        neighbor 172.18.100.18 send-community extended
        neighbor 172.18.100.22 activate
        neighbor 172.18.100.22 default-originate
        neighbor 172.18.100.22 send-community extended
    exit-address-family
```

NOTE:
In this example we use default-originate to advertise a default route from the campus into the overlay VRFs in the data center.

Step 9: Configure route redistribution between BGP and OSPF on the Campus core routers.

```
router ospf 1
    redistribute bgp
```

Step 10: From the campus core switches, verify the BGP neighbors are **Established** with the **show bgp all summary** command.

```
8400-C1-1# sh bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 65001        BGP Router Identifier  : 10.0.0.1
 Peers                  : 2            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60
 Confederation Id       : 0

Address-family : IPv4 Unicast
-----------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State        AdminStatus
 172.18.100.18   65000       15      15      00h:09m:56s  Established  Up
 172.18.100.22   65000       16      16      00h:09m:56s  Established  Up

Address-family : IPv6 Unicast
-----------------------------

Address-family : L2VPN EVPN
-----------------------------
```

Step 11: From the campus core switches, verify that data center subnets are present in the routing table with the **show ip route bgp** command.

```
8400-C1-1# show ip route bgp

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.5.1.0/24, vrf default
        via  172.18.100.18,  [20/0],  bgp
        via  172.18.100.22,  [20/0],  bgp
10.5.2.0/24, vrf default
        via  172.18.100.18,  [20/0],  bgp
        via  172.18.100.22,  [20/0],  bgp
```

Step 12: From the border leaf switches, verify default route pointing to campus is present in the overlay VRF routing tables with the **show ip route vrf Overlay1** command.

```
8325-DC1-L1# show ip route vrf Overlay1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf Overlay1
        via  172.18.100.17[vrf Campus],  [200/0],  bgp
        via  172.18.100.29[vrf Campus],  [200/0],  bgp
10.5.1.0/24, vrf Overlay1
        via  vlan100,  [0/0],  connected
10.5.1.4/32, vrf Overlay1
        via  vlan100,  [0/0],  local
10.5.1.20/32, vrf Overlay1
        via  10.0.10.1[vrf default],  [200/0],  bgp,  encap vxlan,  l3vni 100001

8325-DC1-L1# show ip route vrf Overlay2

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf Overlay2
        via  172.18.100.17[vrf Campus],  [200/0],  bgp
        via  172.18.100.29[vrf Campus],  [200/0],  bgp
10.5.2.0/24, vrf Overlay2
        via  vlan200,  [0/0],  connected
10.5.2.4/32, vrf Overlay2
        via  vlan200,  [0/0],  local
10.5.2.40/32, vrf Overlay2
        via  10.0.10.1[vrf default],  [200/0],  bgp,  encap vxlan,  l3vni 100002
10.5.2.50/32, vrf Overlay2
        via  10.0.10.1[vrf default],  [200/0],  bgp,  encap vxlan,  l3vni 100002
```

Host Port Configuration

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

Port Group Configuration

The SFP28 ports in the Aruba 8325-48Y8C switches (JL624A and JL625A) are organized into four groups of 12 ports each. The ports default to 25Gb/s speed and must be manually set if 10Gb/s is required.

For additional details, find the complete Aruba 8325 Switch Series Installation and Getting Started Guide on the Aruba Support Portal. Navigate to the section **Installing the switch > Install transceivers > Interface-Group operation**.

Step 1: From the **Configurations** menu, select **Ports > Ports**.

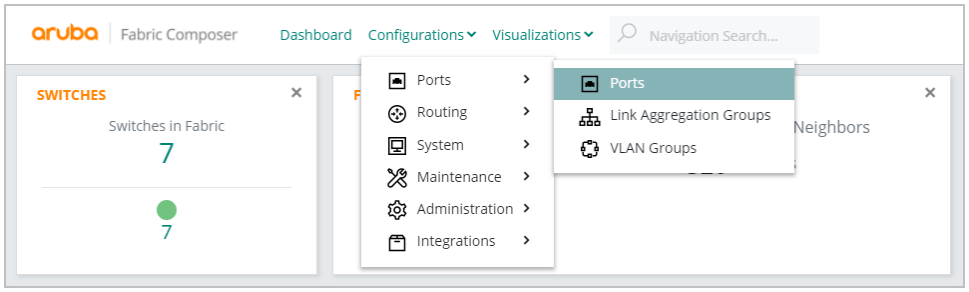

Step 2: Select the switch and single port in the group being changed, from the **ACTIONS** menu on the right, select **Edit**.

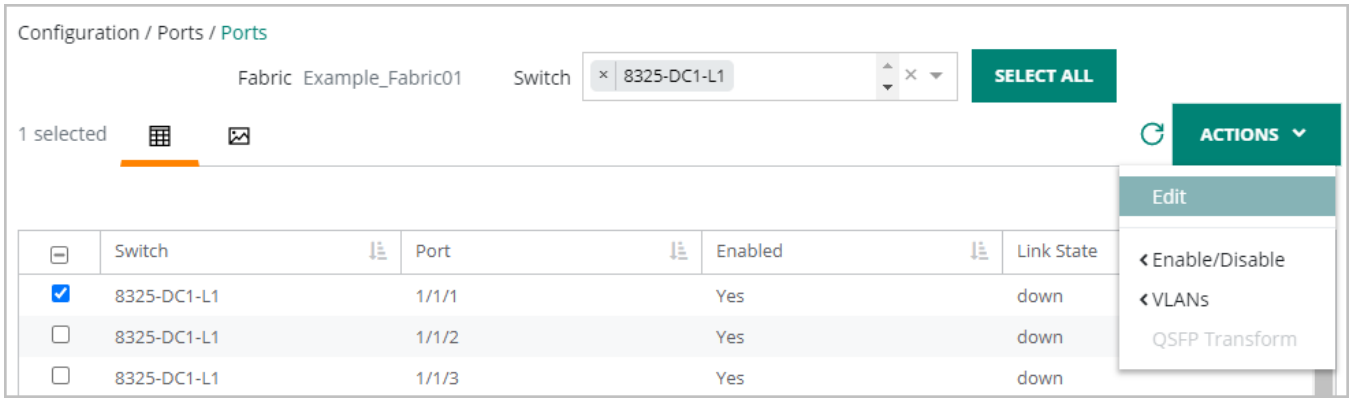

Step 3: From the Ports page, select the **Speed** tab, use the drop down menu to select the Speed, and then click **APPLY**.

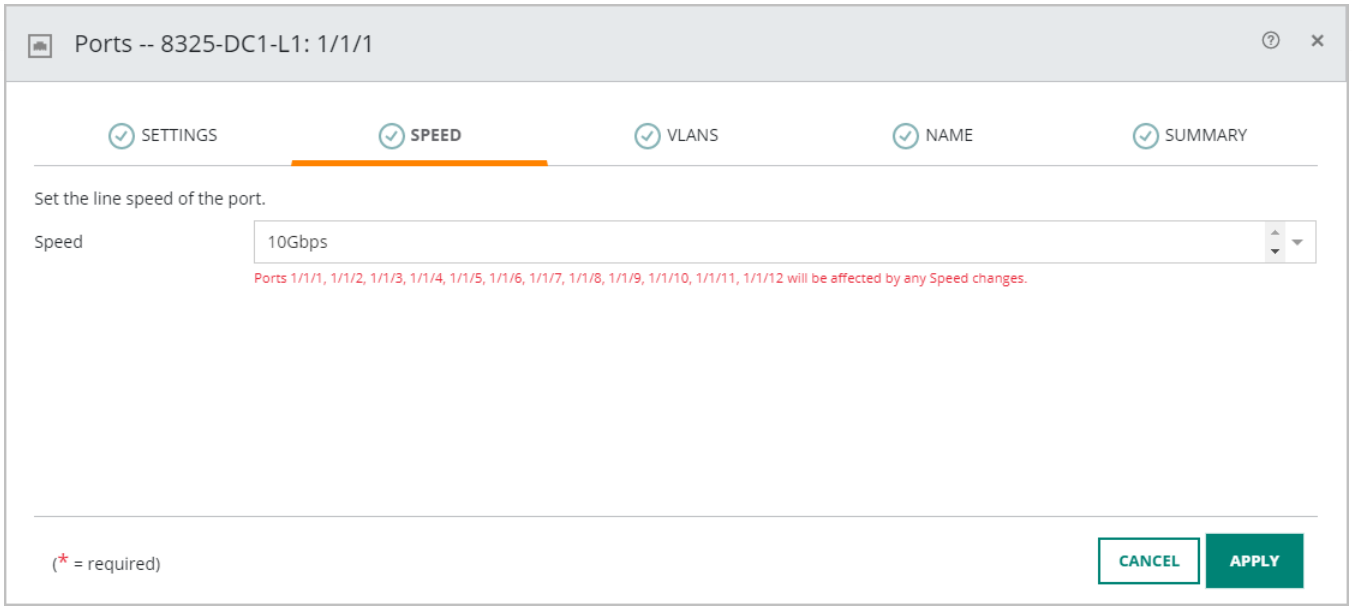

Note:
Observe the full list of ports which will be effected by the speed change. Ensure that this is the correct speed setting for all listed ports.

LACP Host LAG Configuration

The Aruba ESP Data Center uses LACP link aggregation groups to provide fault tolerance and efficient bandwidth utilization to physical hosts in the data center. The steps below show how to configure LAGs and LACP on the switches in the fabric. Configuration is also required on the connected hosts. That topic is not covered in this guide as it may vary between server platforms and operating systems.

Step 1: From the **Configurations** menu, select **Ports > Link Aggregation Groups**.

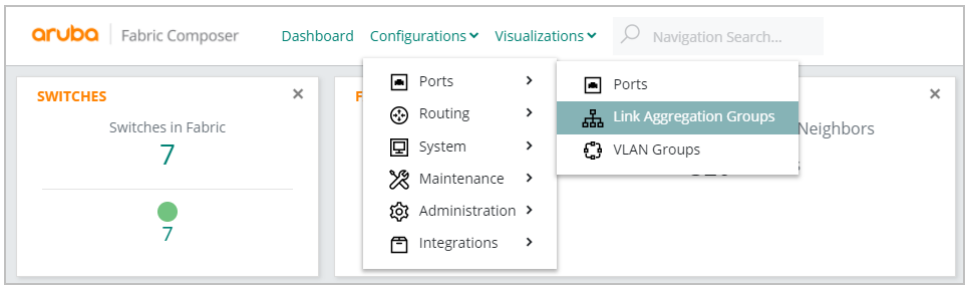

Step 2: From the **ACTIONS** menu on the right, select **Add**.

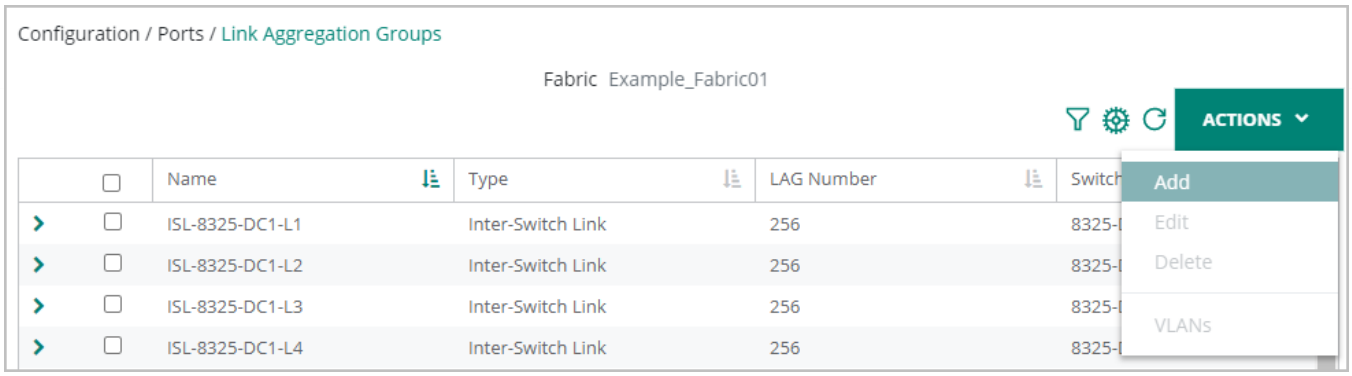

Step 3: From the Link Aggregation Group Settings page, enter a name, description and LAG number, and then click **NEXT**.

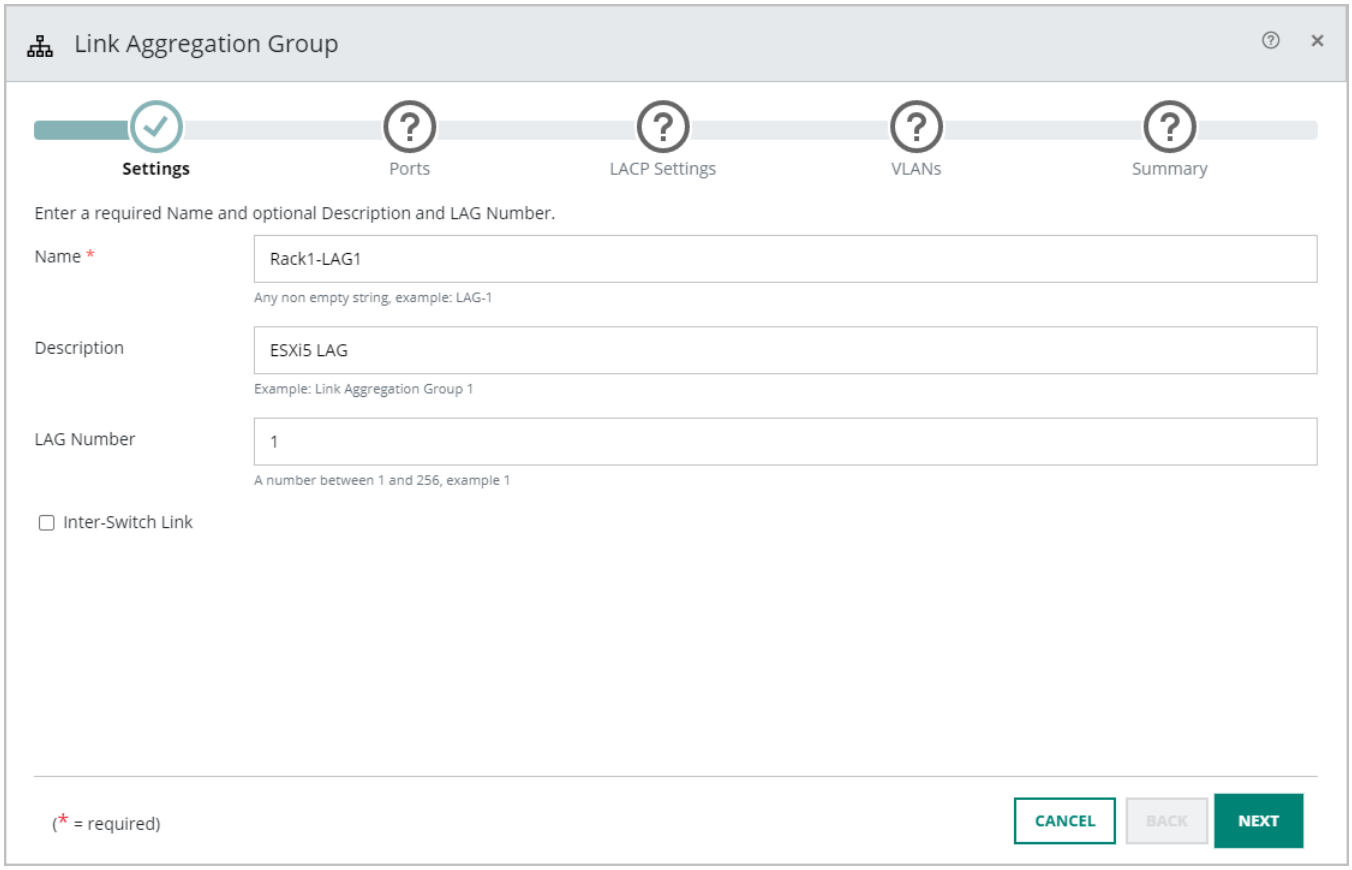

NOTE:
Consider using a **Name** and **LAG Number** which will help to identify the hosts where they are connected.

Step 4: From the Link Aggregation Group Ports page, select the **LAG Switch Member** dropdown menu, and then choose the switches participating in the multi-chassis VSX LAG.

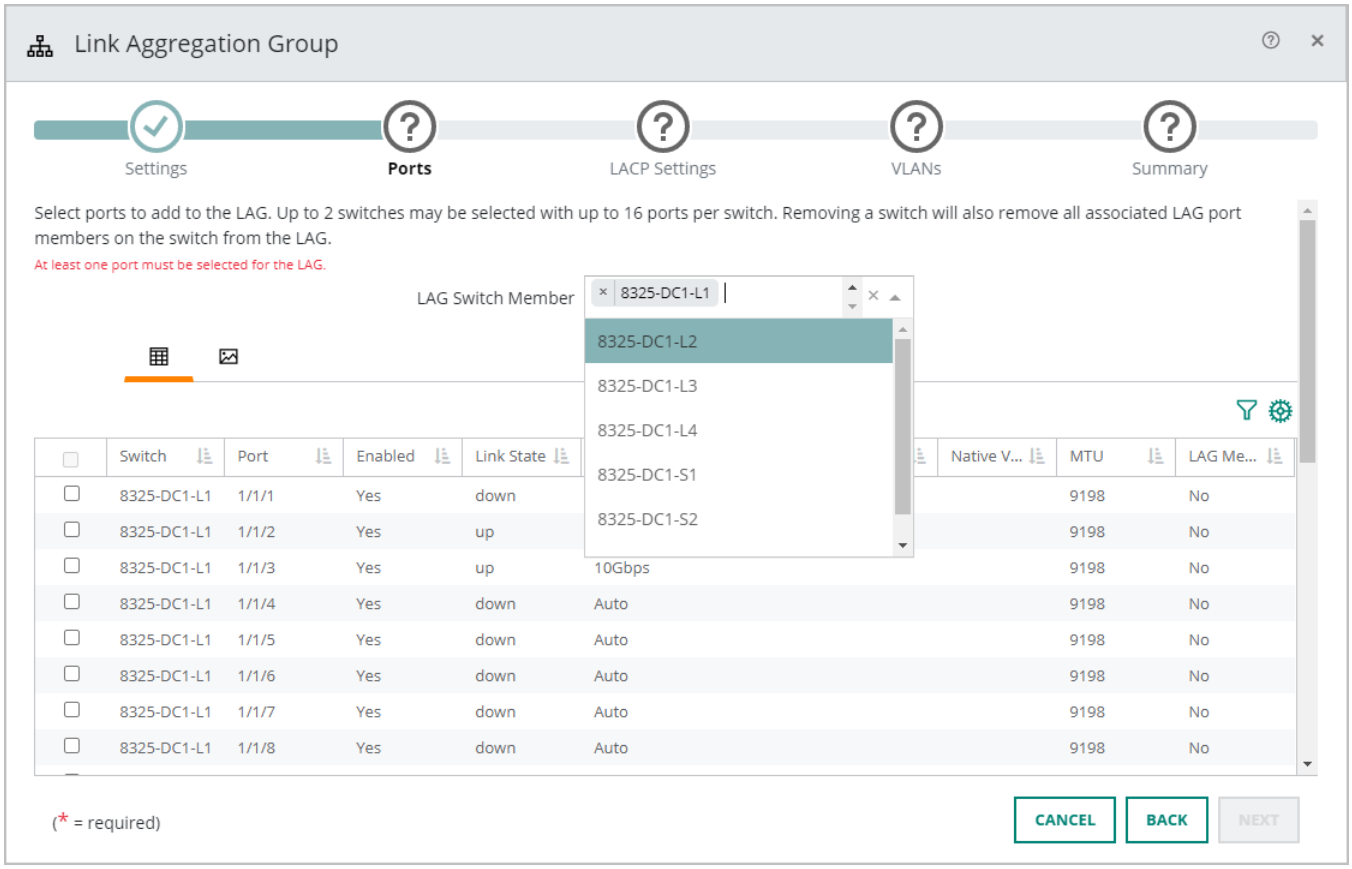

Step 5: From the Link Aggregation Group Ports page, change to **Switch View** mode to make it easier to identify the ports.

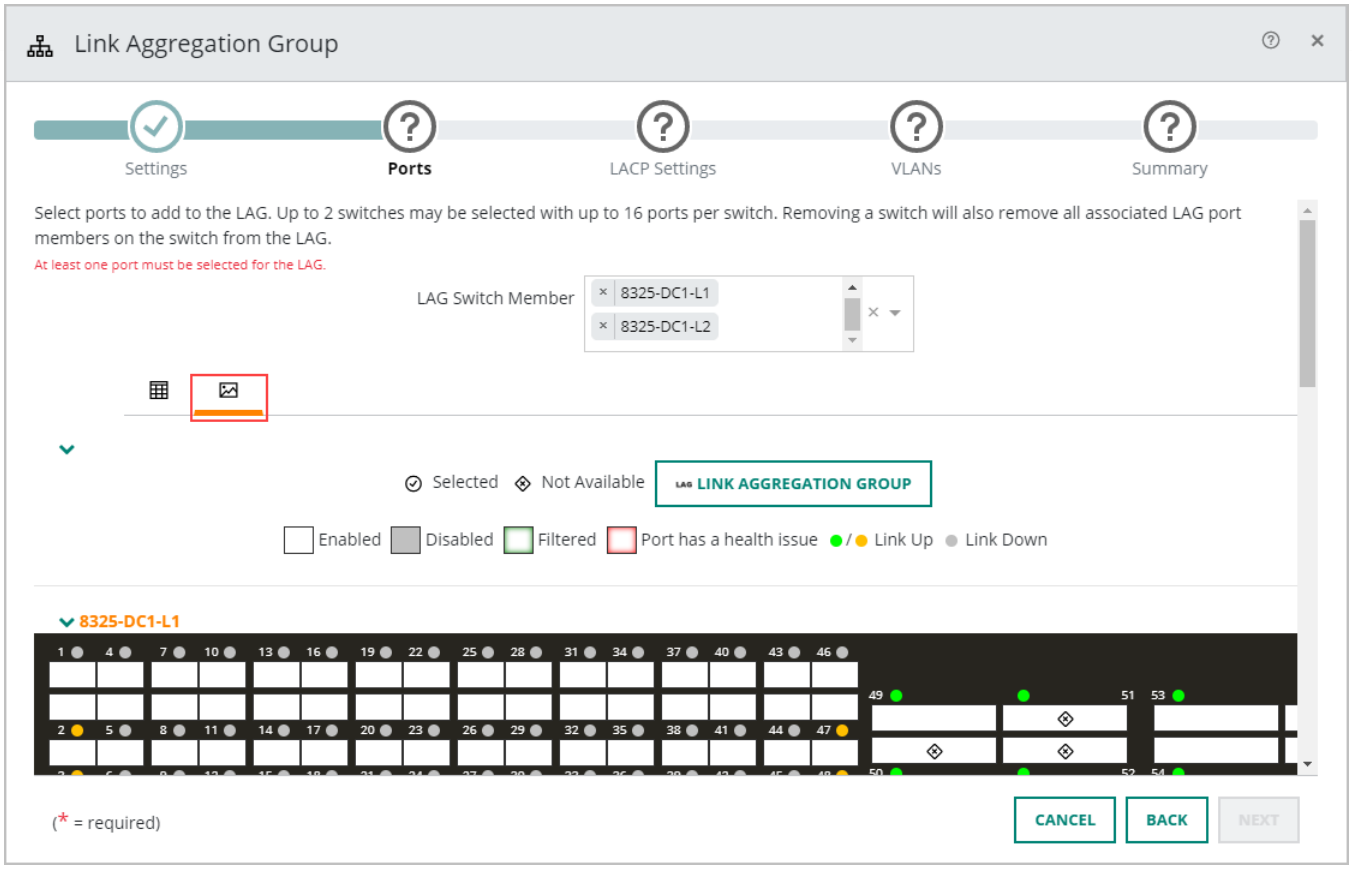

Step 6: On each switch, click the ports for the **LAG Group**, and then click **NEXT**.

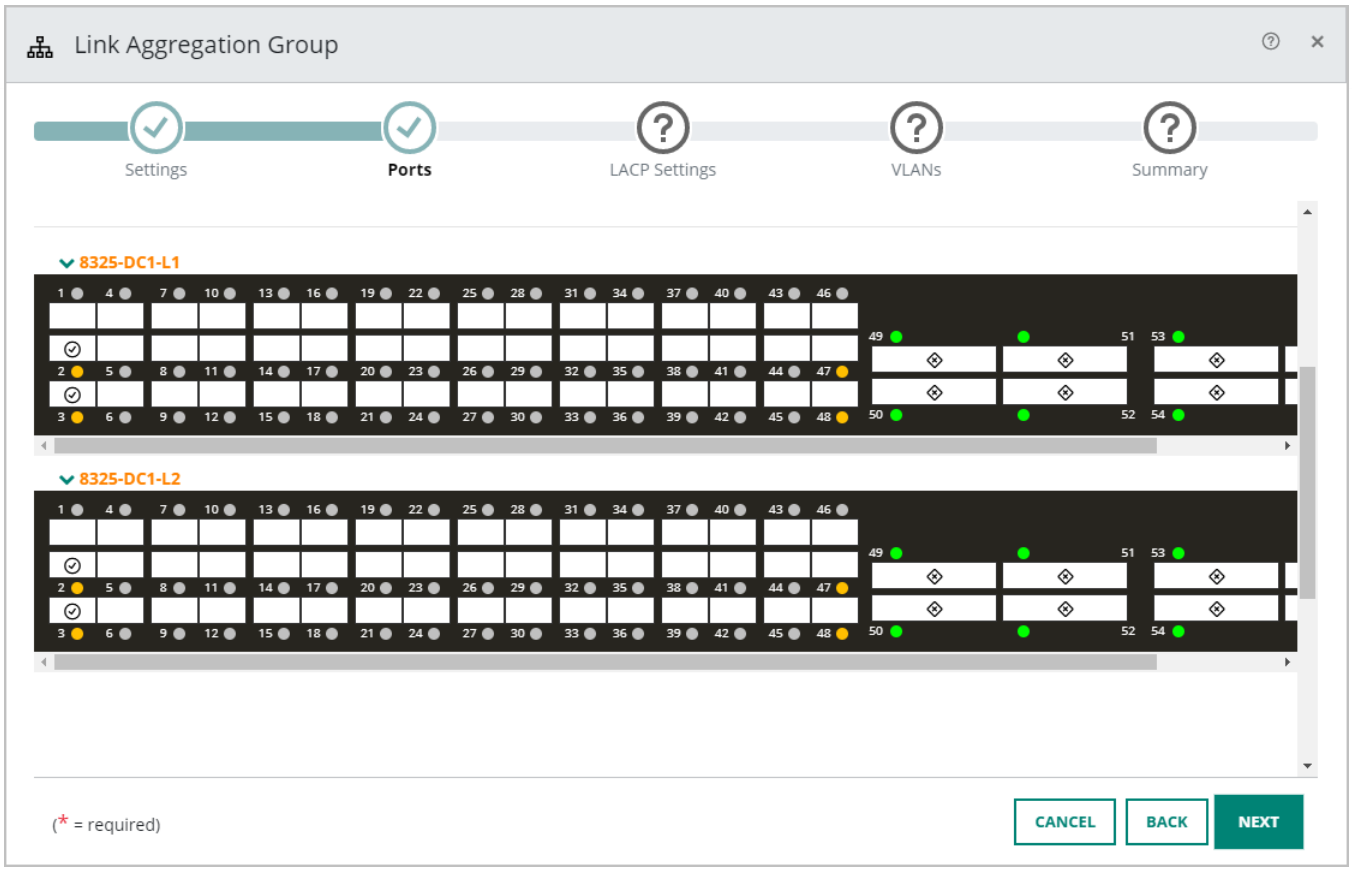

NOTE:
A checkmark will appear on the newly selected ports. The diamond icon appears on ports already configured for a LAG Group.

Step 7: From the Link Aggregation Group LACP Settings page, select **Enable LACP Fallback**, choose both switches in the table and alter the **LACP settings** as needed, and then click **NEXT**.

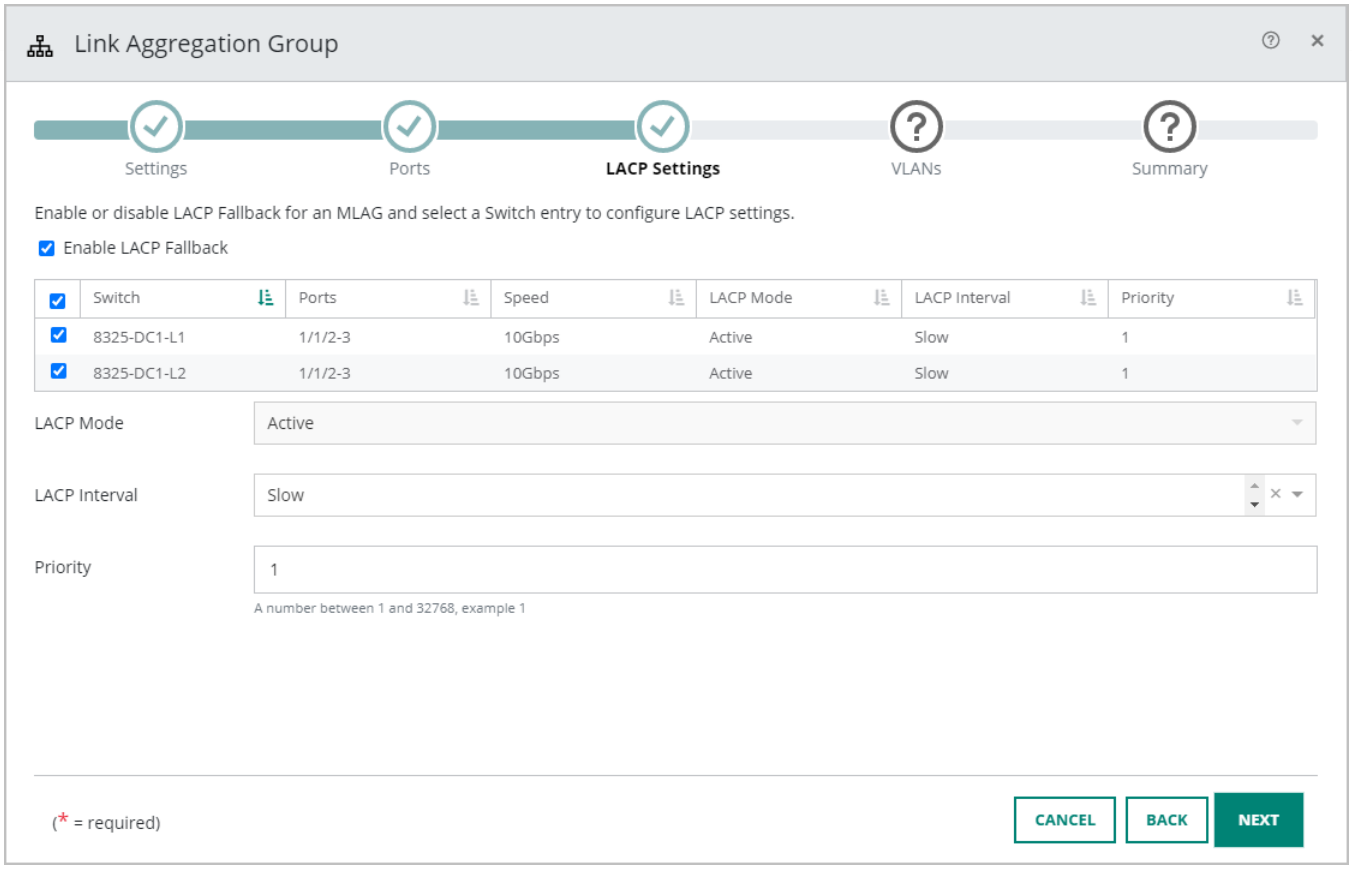

NOTE:
Aruba CX switches default to Active to ensure that LACP can be established regardless of the LACP configuration of the host platform. The default settings are recommended.

Step 8: From the Link Aggregation Group LACP Settings page, enter the Native VLAN or leave blank as appropriate to the data center design, enter the appropriate **VLANs**, and then click **NEXT**.

Note:
If the hosts will be connected to the Overlay network created previously, ensure the VLAN for that network is in the VLANs field.

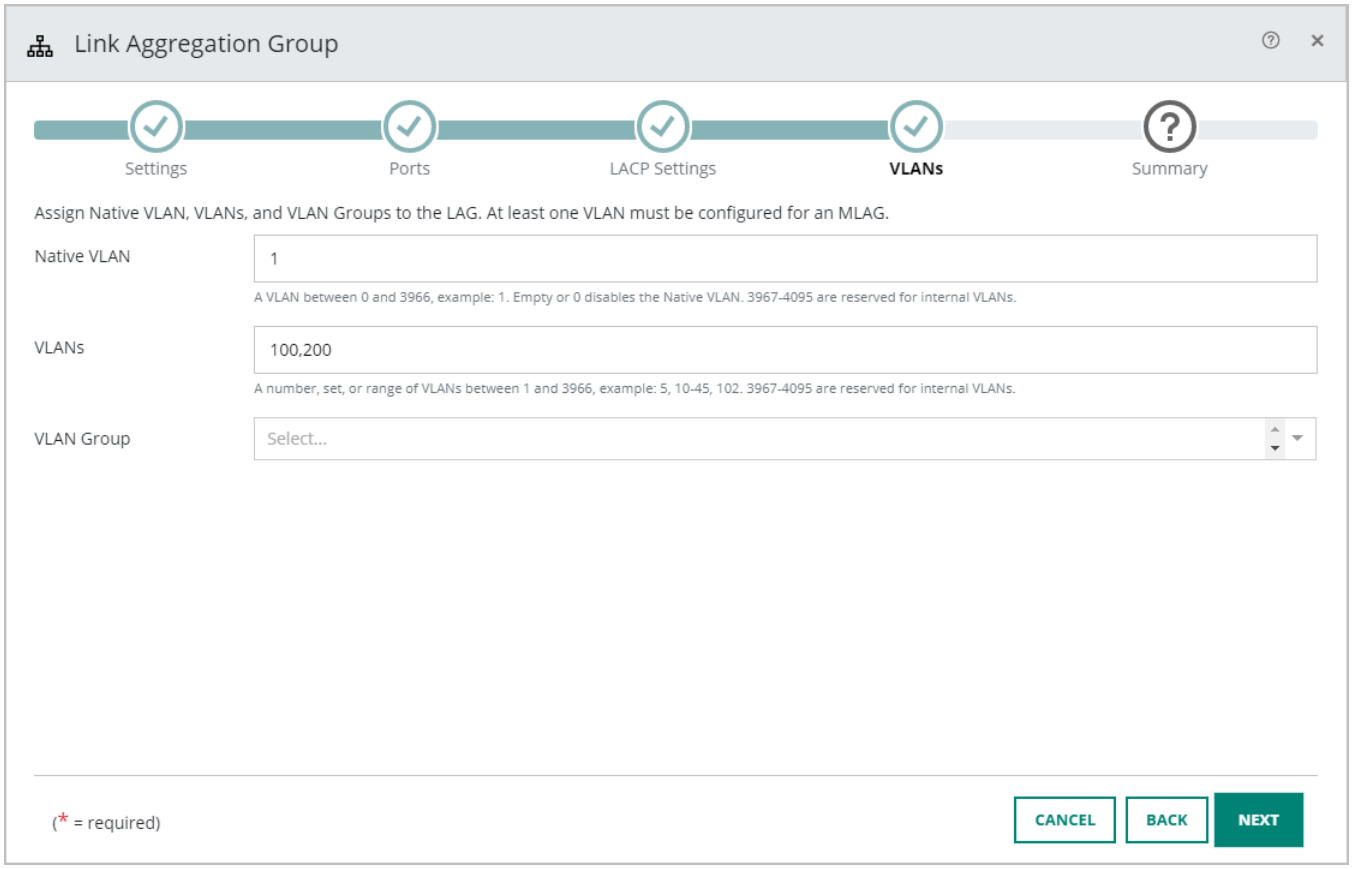

Step 9: Confirm the information was entered correctly and click **APPLY**.

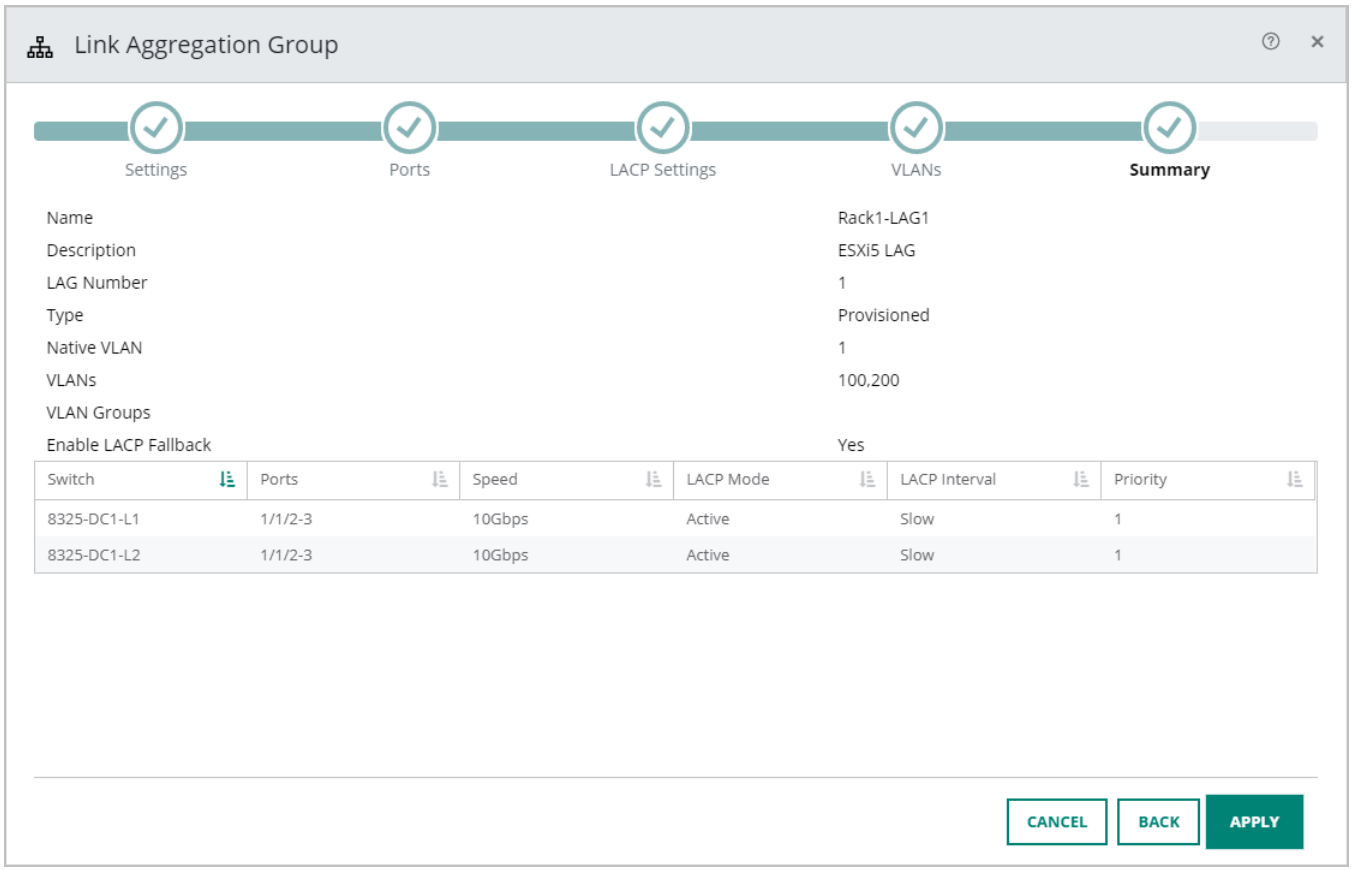

Step 10: Repeat the procedure for each host connected to the fabric.

Step 11: From the leaf switches, verify the LACP LAG connections are configured correctly with the **show lacp interfaces** command.

```
8325-DC1-L1# sh lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/2      lag1(mc)   2     1     ALFNCD  02:00:00:00:01:01 65534  1    up
1/1/3      lag1(mc)   3     1     ALFNCD  02:00:00:00:01:01 65534  1    up
1/1/50     lag256     51    1     ALFNCD  54:80:28:fc:2b:00 65534  256  up
1/1/51     lag256     52    1     ALFNCD  54:80:28:fc:2b:00 65534  256  up
1/1/52     lag256     53    1     ALFNCD  54:80:28:fc:2b:00 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/2      lag1(mc)   2     255   PLFNCD  b4:7a:f1:13:a2:02 65535  15
1/1/3      lag1(mc)   4     255   PLFNCD  b4:7a:f1:13:a2:02 65535  15
1/1/50     lag256     51    1     ALFNCD  54:80:28:fc:ca:00 65534  256
1/1/51     lag256     52    1     ALFNCD  54:80:28:fc:ca:00 65534  256
1/1/52     lag256     53    1     ALFNCD  54:80:28:fc:ca:00 65534  256
```

NOTE:

A combination of VSX peer LAG interfaces and VSX multi-chassis LAG interfaces to hosts are likely included in the command output. As seen above, the multi-chassis interfaces are denoted by **(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 above **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.

VMWare vSphere Integration

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

Step 1: From the **Configurations** menu, select **Integrations > VMware vSphere**.

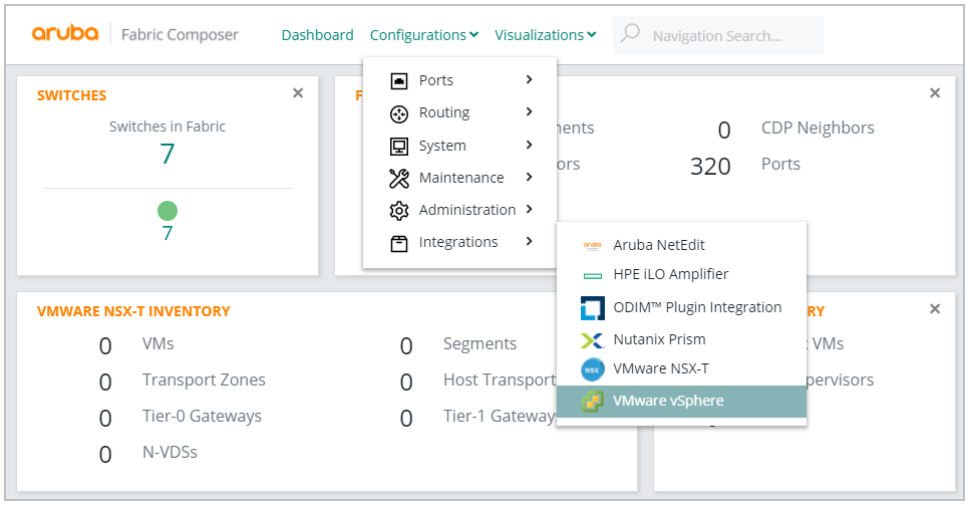

Step 2: From the **ACTIONS** menu on the right, click **Add**.

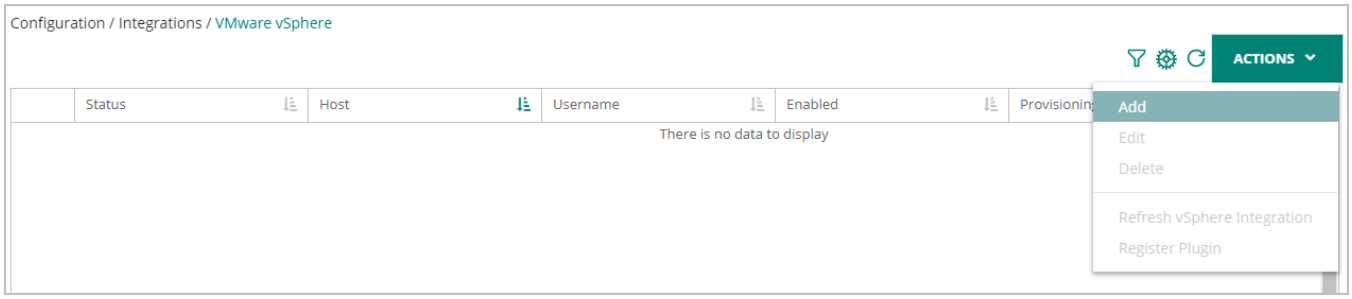

Step 3: From the VMware vSphere Host page and implement the following settings.

- **Name:** *Example-vSphere1*
- **Description:** *Example vSphere integration*
- **Host:** *172.16.1.212*
- **Username:** *administrator@example.local*
- **Password:** *password*
- **Validate SSL/TLS certificates for Aruba Fabric Composer:** *unchecked*
- **Enable this configuration:** *checkmark*

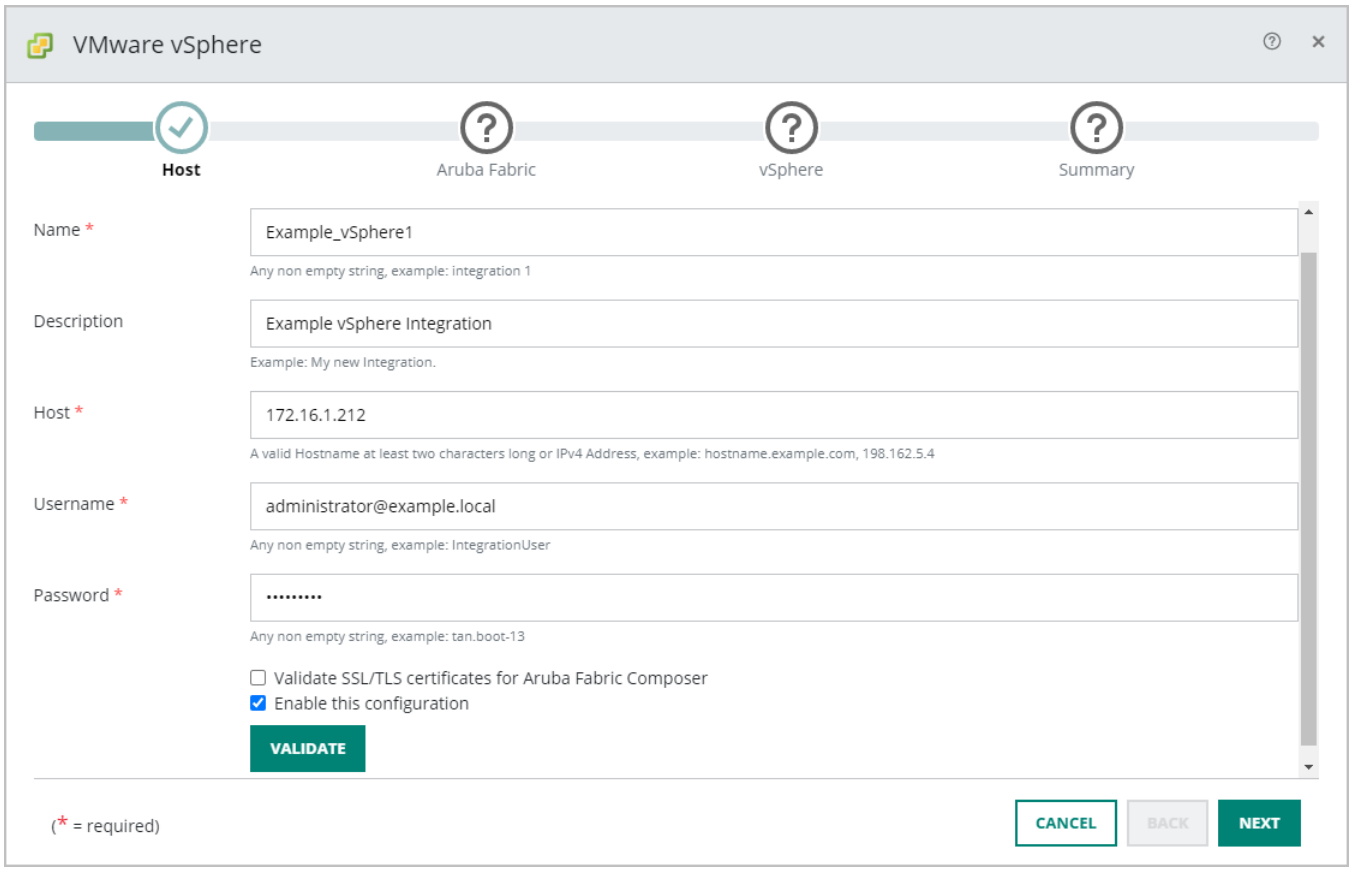

NOTES:
Host is the resolvable hostname or IP address of the vCenter server.
Username is the name of an administrator account on the vCenter server.
Password is the password for the administrator account on the vCenter server.

Step 4: Verify the provided credentials are correct by clicking **VALIDATE**, a green **Success** message should appear at the bottom right, and then click **NEXT**.

Step 5: From the VMware vSphere Aruba Fabric page, choose from the options below, enter a VLAN range, and 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, choose **Automated VLAN provisioning for ESX hosts connected through intermediate switches**.

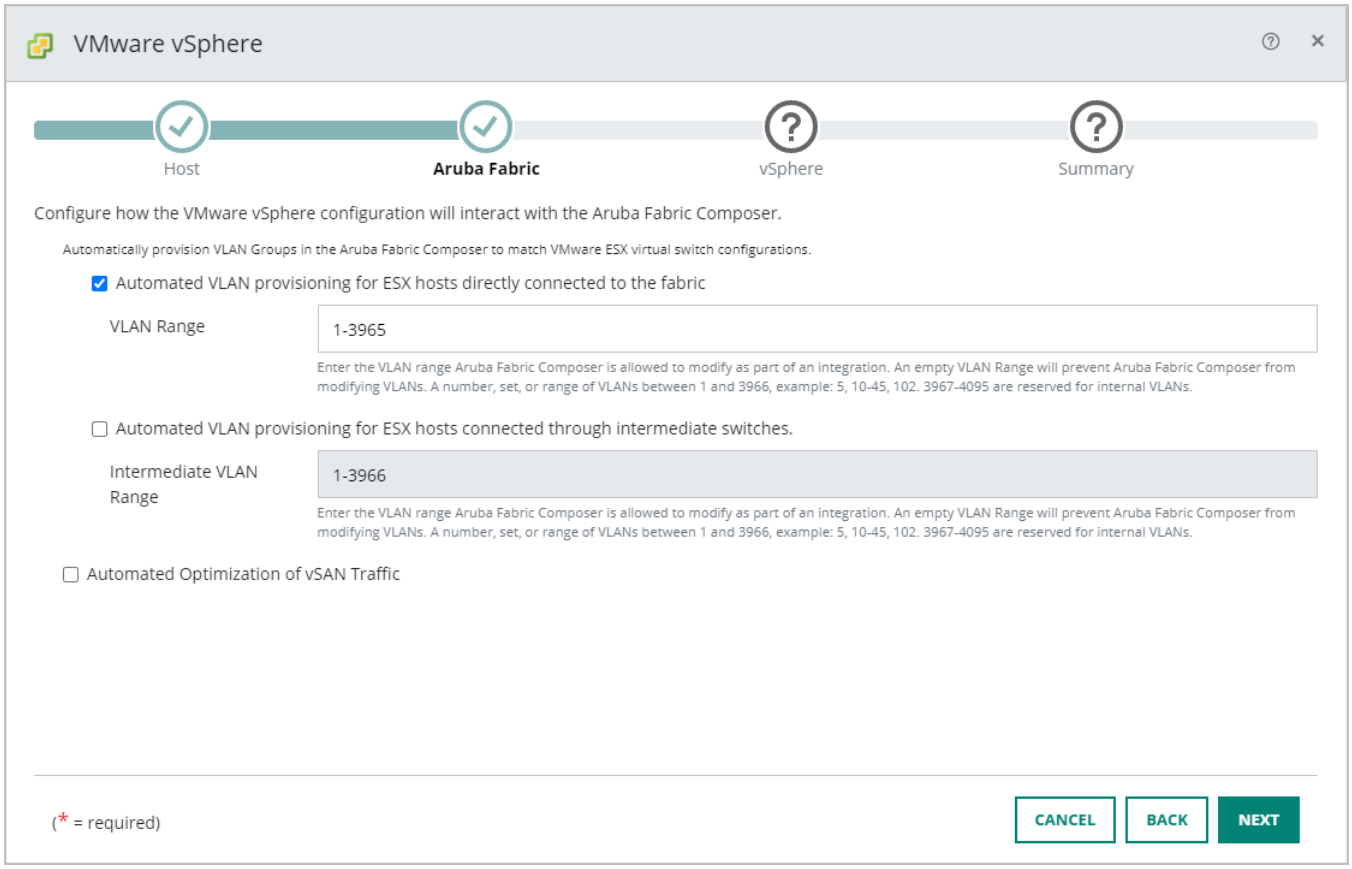

NOTE:
For additional details, please refer to the Aruba Fabric Composer User Guide.

Steps 6: From the VMware vSphere, vSphere page, select the checkbox for **Discovery protocols** and then click **NEXT**.

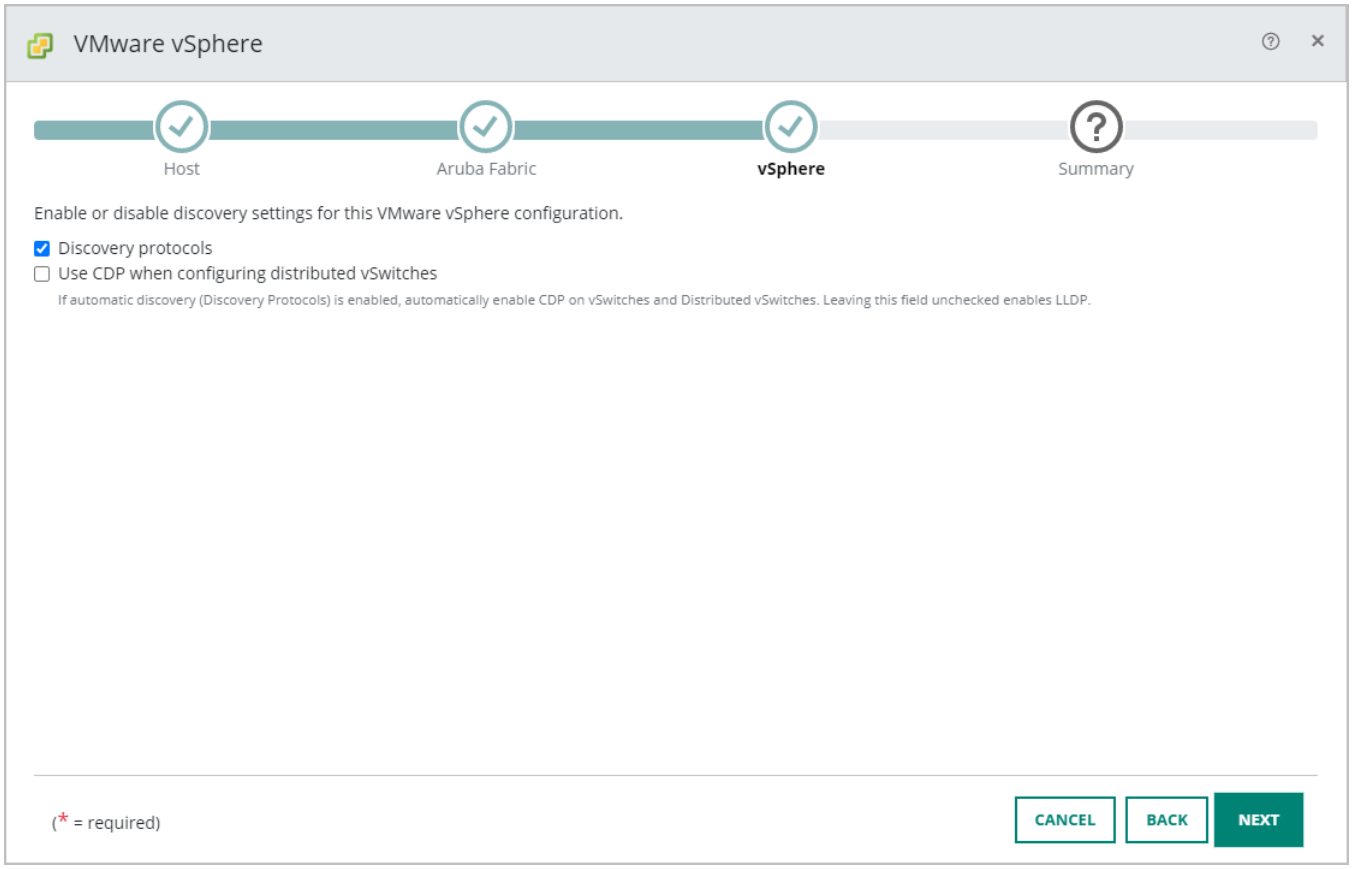

CAUTION:
If **Discovery protocols** is not enabled, the VMWare integration will not display virtual switches correctly.

Step 7: Confirm the information was entered correctly and click **APPLY**.

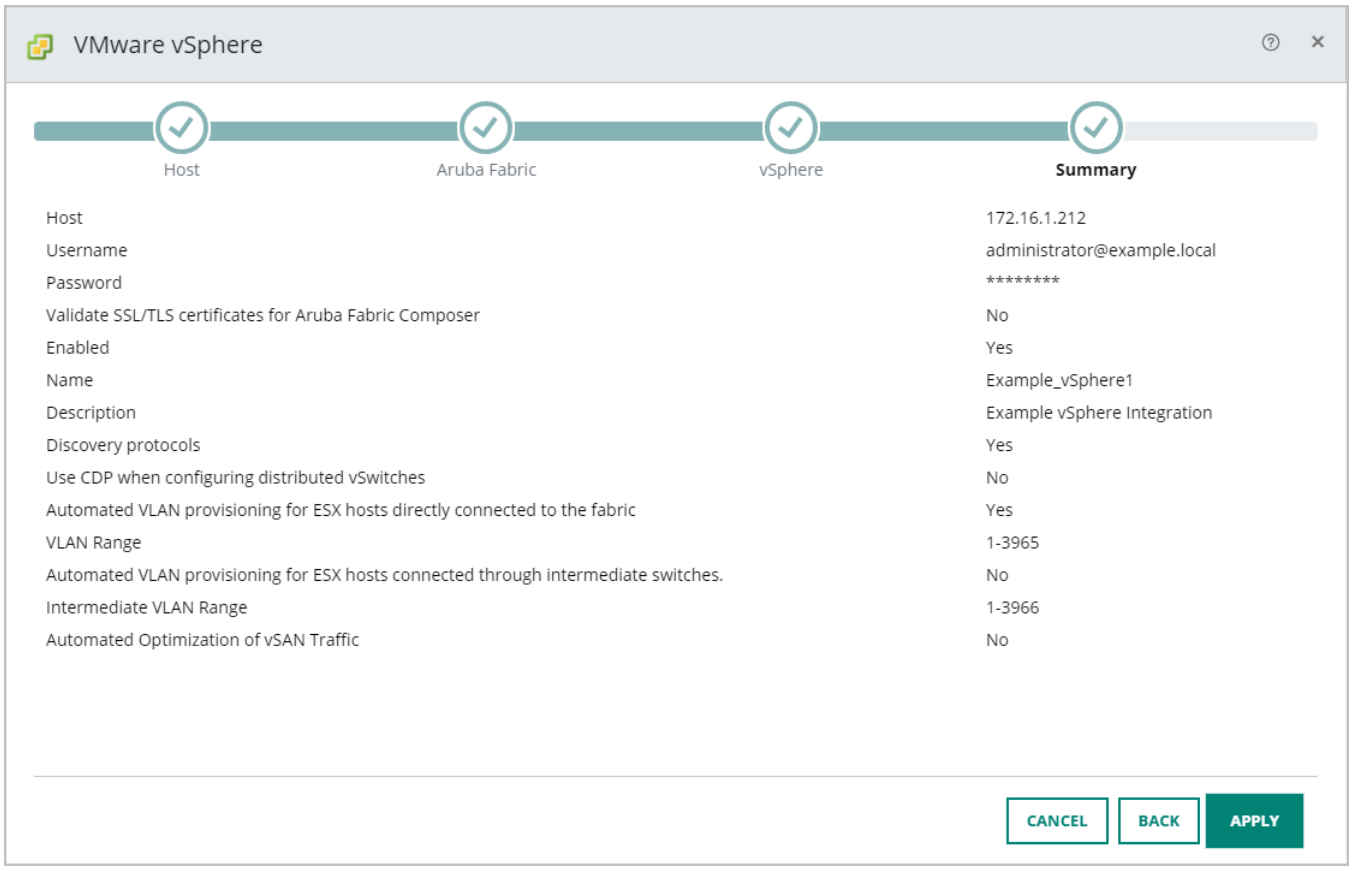

Step 8: Navigate to **Visualizations > Hosts**, verify connectivity from the hypervisor layer to the leaf switches.

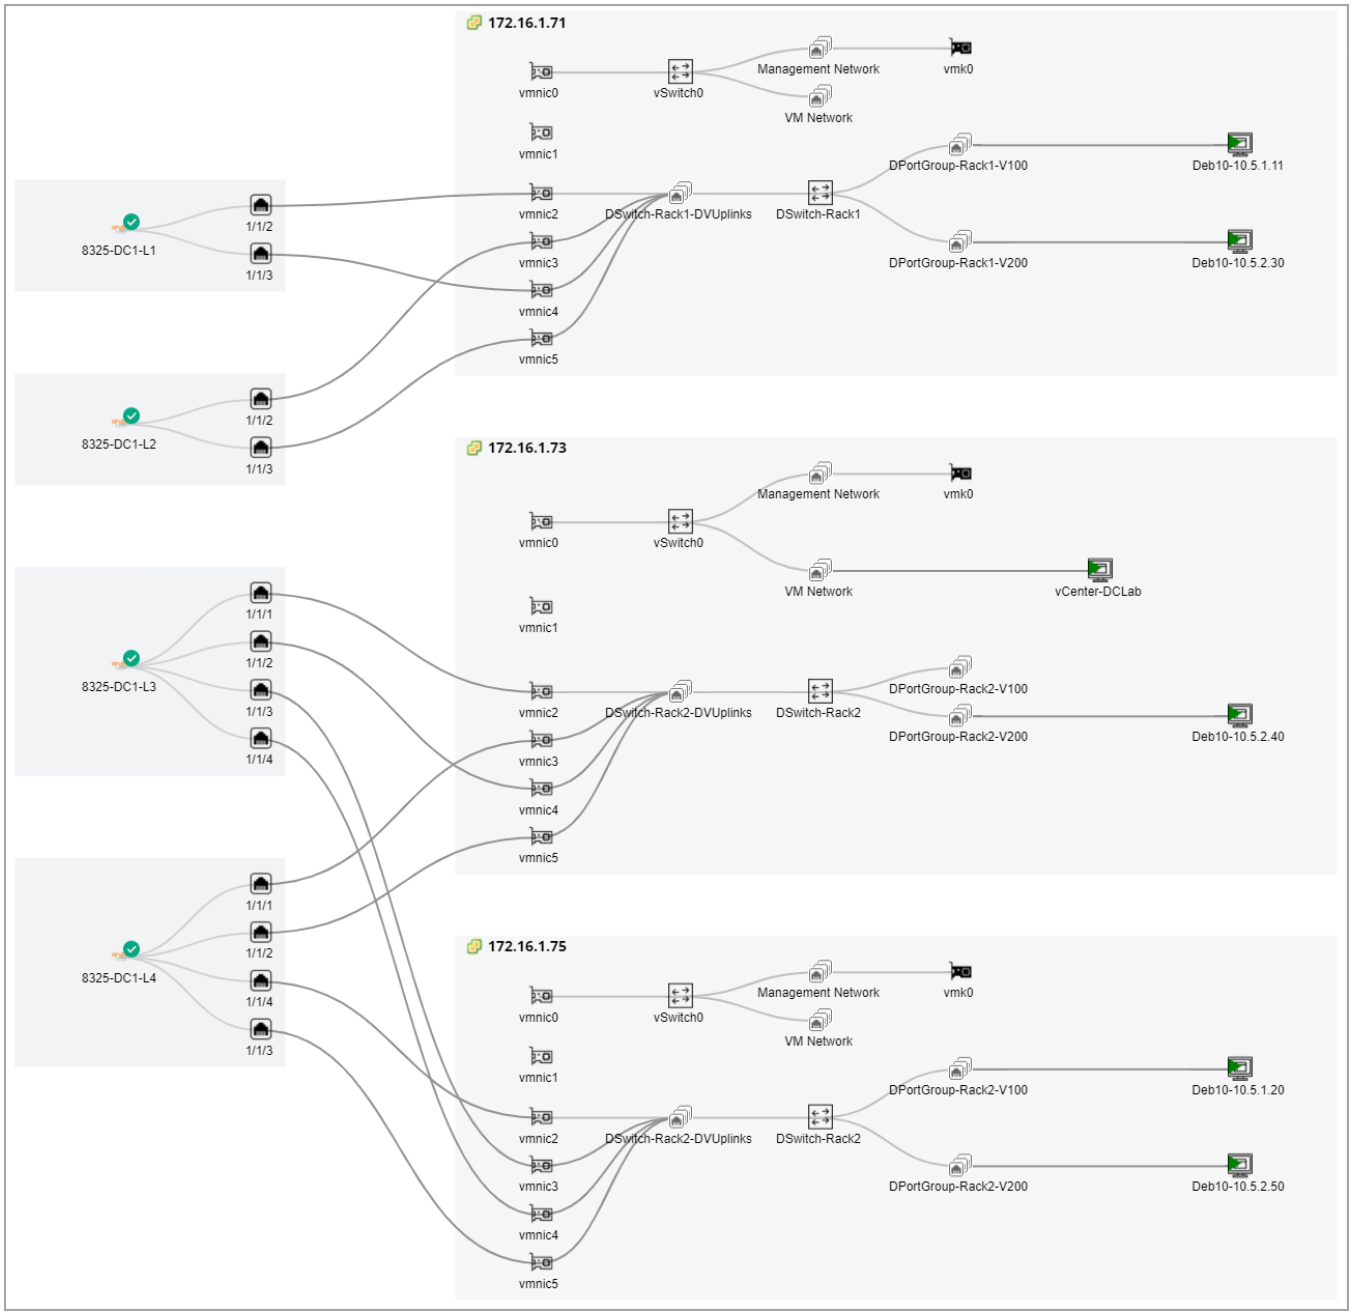

Validated Hardware and Software

The following hardware and software versions were validated for this guide. For compatibility, please upgrade to at least the versions listed below.

Switches

Product Name	Software Version
Aruba CX 8325	10.06.0113
Aruba CX 8400	10.06.0113

Management and Orchestration

Product Name	Software Version
Aruba Fabric Composer	6.0.2-8872

Reference Configuration Files

These full switch configurations were used to develop this guide:

Border Leaf #1 Configuration

Border Leaf #2 Configuration

Leaf #3 Configuration

Leaf #4 Configuration

Spine #1 Configuration

Spine #2 Configuration

Spine #3 Configuration

Campus Core #1 Configuration

Campus Core #2 Configuration

What’s New in This Version

We made the following changes since Aruba last published this guide:

- This is a new guide.