

NetConductor

User Guide



Hewlett Packard
Enterprise

Copyright Information

© Copyright 2024 Hewlett Packard Enterprise Development LP.

Open Source Code

This product includes code licensed under certain open source licenses which require source compliance. The corresponding source for these components is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, please check if the code is available in the HPE Software Center at <https://myenterpriselicense.hpe.com/cwp-ui/software> but, if not, send a written request for specific software version and product for which you want the open source code. Along with the request, please send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company
Attn: General Counsel
WW Corporate Headquarters
1701 E Mossy Oaks Rd, Spring, TX 77389
United States of America



Contents	3
About This Document	1
Intended Audience	1
Related Documents	1
Conventions	1
Contacting Support	3
NetConductor Overview	4
NetConductor Vocabulary	6
Deployment Scenarios	8
Distributed Campus-wide Fabric	8
Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch	9
Centralized Multi-site Fabric with Third-Party Networks	9
Supported Devices	11
Aruba Access Points	11
AOS-CX Switches	12
Aruba Gateways	13
Getting Started	15
Prerequisites	15
Role-Based Policy	16
VXLAN and BGP-EVPN Overview	19
Global Client Roles	20
Configuring the Global Client Roles	21
Editing a Role	23
Deleting a Role	24
Performing Selective Enablement of Subnets	24
Deleting a Subnet	25
Performing Selective Enablement of Groups	25
Deleting a Group	26
Network Wizard Overview	26
Setting Up a Campus (3 tier L3 access) Underlay Network	27
Setting Up a of Data Center (Spine and Leaf) Underlay Network	31
Modifying an Underlay Network	34
Deleting an Underlay Network	35
Fabric Wizard Overview	36
Overlay Fabric Overview	37
Overlay Segment Overview	45
Muti-Fabric EVPN Overview	53
Setting up the Muti-Fabric EVPN	53
Modifying a Multi-Fabric EVPN	57
Re-syncing the Multi-Fabric EVPN	57
Static VXLAN Tunnels on AOS-10 Gateways	58
Configuring a Static VXLAN Tunnel on Gateways	58
Mapping VLAN and VNI to a VXLAN Tunnel	60
Viewing the Static VXLAN tunnels on Gateway Clusters	61
Example Use Case	63
Distributed Campus-wide Fabric	63
Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch	65

Centralized Multi-site Fabric with Third-Party SD-WAN	66
---	----

This document describes the NetConductor and provides detailed instructions for setting up, configuring, and managing all supported deployments.

Intended Audience

This guide is intended for network administrators who manage and monitor networks.

Related Documents

In addition to this document, see the following documents for more details on NetConductor:

- *HPE Aruba Networking Central Online Help* at <https://www.arubanetworks.com/techdocs/central/latest/content/home.htm>

Conventions

[Table 1](#) lists the typographical conventions used throughout this guide to emphasize important concepts:

Table 1: *Typographical Conventions*

Type Style	Description
<i>Italics</i>	This style is used to emphasize important terms and to mark the titles of books.
System items	This fixed-width font depicts the following: <ul style="list-style-type: none">▪ Sample screen output▪ System prompts
Bold	<ul style="list-style-type: none">▪ Keys that are pressed▪ Text typed into a GUI element▪ GUI elements that are clicked or selected

The following informational icons are used throughout this guide:



Indicates helpful suggestions, pertinent information, and important things to remember.



Indicates a risk of damage to your hardware or loss of data.



Indicates a risk of personal injury or death.

Contacting Support

Table 2: *Contact Information*

Main Site	arubanetworks.com
Support Site	https://networkingsupport.hpe.com/home
Airheads Social Forums and Knowledge Base	community.arubanetworks.com
North American Telephone	1-800-943-4526 (Toll Free) 1-408-754-1200
International Telephone	arubanetworks.com/support-services/contact-support/
Software Licensing Site	lms.arubanetworks.com
End-of-life Information	arubanetworks.com/support-services/end-of-life/
Security Incident Response Team	Site: arubanetworks.com/support-services/security-bulletins/ Email: aruba-sirt@hpe.com

Chapter 2

NetConductor Overview

With an ever-growing focus on security and scale, the enterprise network is becoming more and more complex in terms of design, deployment, and operations. There is an increasing reliance on BYOD (Bring Your Own Device) and IoT (Internet of Things) for business efficiency and digital transformation initiatives. This increases the risk of security threats to the enterprise due to a sharp increase in the unknown or rogue clients and an ever-expanding threat front. Defining policy manually for these clients using the complex policy constructs available today can prove to be a challenging task for security and network administrators. Furthermore, intent-based networking has become an increasingly popular paradigm that many customers are looking to adopt and implement. The goal of intent-based networking is not only to abstract the underlying complexities of network but instead allow users to design, implement, and operate their networks based on their business intents. Automated network provisioning and orchestration has been identified to achieve this level of abstraction by many network vendors. Thus, the focus has shifted to the security, scalability, and simplification of these networks.

NetConductor is an edge-to-cloud network and security framework designed to tackle these problems for the modern enterprise network. It is tied directly to the HPE Aruba Networking ESP (Edge Services Platform) vision of an edge-to-cloud network. Intelligent overlays are built on highly available underlays and are tied to a full policy-based micro-segmentation model, based on global roles, across the entire network infrastructure of the customer. Role-based policies abstract policy from the underlying network and enable flexible and simplified policy definition and enforcement. This is enhanced by the full automation of the underlay, orchestration of the overlay, a single pane of glass for management and monitoring, and a rich array of complementary services. The NetConductor framework has evolved to enhance the policy and orchestration components to deliver true intent-based network evolution and optimization.

The following are the main pillars of NetConductor:

- **Role-based Segmentation**—NetConductor provides the ability to deploy a zero-trust enforcement model using role-based segmentation. Traditional policies use location specific entities like IP addresses or subnets to define security policies. Role-based policies abstract policy from the underlying network by assigning roles to endpoints or users and using roles to enforce policies. Role-based policies can be enforced in a distributed manner at different parts of the network. NetConductor also provides the ability to automate and simplify policy definition for IoT devices with behavior-based profiling using AI or ML based classification. This greatly simplifies policy definition and ensures consistent policy enforcement across wired and wireless campus networks, the datacenter, and across the WAN.
- **Intelligent Overlays**—Overlay networks provide the ability to deploy flexible services based on ever-changing demands of the endpoints and applications. Decoupling of overlay network from the physical topology enables on-demand deployment of layer 2 and layer 3 services irrespective of underlay physical topology. Overlay networks also enable the ability to carry endpoint or user role information across the network without requiring all devices in the path to understand or manage the roles. NetConductor provides customers the flexibility to choose between centralized overlays or distributed overlays to address their unique requirements. The centralized overlay provides simplified operations and advanced security features for distributed enterprise and smaller campus deployments. For large enterprise campus deployments, NetConductor provides the ability to use

distributed overlays for wired and wireless endpoints. This enables large enterprises to deploy a standards-based and scalable overlay network. Both overlay models support the Colorless Ports feature, which enables automated client on-boarding and access control for ease of operations.

- **Automation and AI Ops**—One of the primary requirements for enterprise campus network is simplicity of deployment, maintenance, and troubleshooting. With cloud-based management provided by HPE Aruba Networking Central, enterprise devices can be on-boarded and managed in a matter of minutes. Intent-based workflows enable architects to deploy and provision the network without the need for technical expertise in the networking protocols and the command line interface. HPE Aruba Networking Central also provides unified policy orchestration for the global network across wired, wireless, and SD-WAN. NetConductor also enhances end user experience and reduces help desk calls with real time problem identification and AI or ML driven actionable insights.

Benefits of NetConductor

The following are some of the key benefits of the NetConductor solution. The objective of this guide is to highlight these capabilities to the customer.

- **Simplified and Consistent Security Policies**
 - Simplified policy definition based on customer identity
 - Security policies agnostic of location, network, and devices
 - Policy follows the endpoint, user, or application across wired and wireless networks
 - Consistent policies across Campus, Branch, and Datacenter
 - Increase scale by eliminating the need for enforcement nodes to maintain endpoint to role mappings to enforce policies
- **Flexible Overlays Agnostic of Underlay Architecture**
 - Flexible choice of centralized or distributed NetConductor fabrics on any underlay physical network architecture
 - Automated stitch-up and tear-down of layer 2 and layer 3 services based on customer on-boarding
 - Address requirements of small, distributed enterprise to a large campus network
- **Simplified Network Deployments and Operations with Intent-Driven Workflows**
 - Abstract complexity of the underlying protocols from network architects or operators
 - Enables global orchestration of roles and role-based policies from HPE Aruba Networking Central
 - Unified monitoring and troubleshooting across all device types and network locations
 - Actionable insights enable ease of troubleshooting for network issues

Features for NetConductor

The following features are available in NetConductor:

1. [Global Client Roles](#)
2. [Network Wizard Overview](#)
3. [Fabric Wizard Overview](#)
4. [Static VXLAN Tunnels on AOS-10 Gateways](#)

NetConductor Vocabulary

The following table provides a brief description of the technical terms used in this guide.

Table 3: *List of Technical Terms used in this Guide*

Term	Description
Border	Border device persona connects the fabric to external networks. For example, connect fabric to WAN or Internet or firewalls.
Border Gateway Protocol (BGP)	BGP is a standardized routing method that enables the internet to exchange routing information between autonomous systems (AS).
Ethernet VPN (EVPN)	EVPN is an extension of the BGP protocol for layer 2 (bridging) and layer 3 (routing) VPNs.
Edge	Edge is a device persona that connects endpoints to the fabric.
External BGP (eBGP)	Refers to BGP connection between external peers.
Fabric	Fabric is a group of AOS-CX Switches that are part of the BGP-EVPN VXLAN overlay. The overlay fabric is created by configuring VXLAN tunnels between stub and edge Switches.
Group-based Policy (GBP)	GBP is used to segment user traffic in a network by grouping the users into roles based on user authentication at the source or VTEP. Source-based roles will remain effective even if a device authenticates at a different location, or if the device is assigned a different IP address.
Internal BGP (iBGP)	Refers to BGP connection between internal peers.
Inter-Switch Link (ISL)	ISL is a layer 2 interface between two VSX peer Switches. Each VSX Switch must be configured with an ISL link connected to its peer VSX Switch.
Open Shortest Path First (OSPF)	OSPF refers to an Interior Gateway Protocol (IGP). OSPF distributes routing information between routers belonging to a single Autonomous System (AS).
Multi-Fabric EVPN	Refers to the Multi-Fabric, where by 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. Multiple data center fabrics and locations can be combined into a single overlay topology.
Policy Identifier	Policy Identifier is a unique identification number mapped to a client role.
Route Reflector	Route Reflector refers to a concept that is specific to BGP which is used to optimize route propagation.
Stub	Stub is a device persona that supports both static VXLAN tunnels and EVPN VXLAN tunnels.
Switch Virtual Interface (SVI)	An SVI (also known as VLAN interface) refers to a logical layer 3 interface on a Switch.
Virtual Extensible LAN (VXLAN)	VXLAN is an Overlay Technology which address the scalability problems associated with large cloud computing deployments.

Table 3: *List of Technical Terms used in this Guide*

Term	Description
Virtual Routing and Forwarding (VRF)	VRF is a technology that allows multiple instances of a routing table to co-exist within the same router simultaneously in an IP-based computer network.
VXLAN Network Identifier (VNI)	Refers to VXLAN network identifier or VXLAN segment ID.
VXLAN Tunnel End Point (VTEP)	An entity that originates and/or terminates VXLAN tunnels.

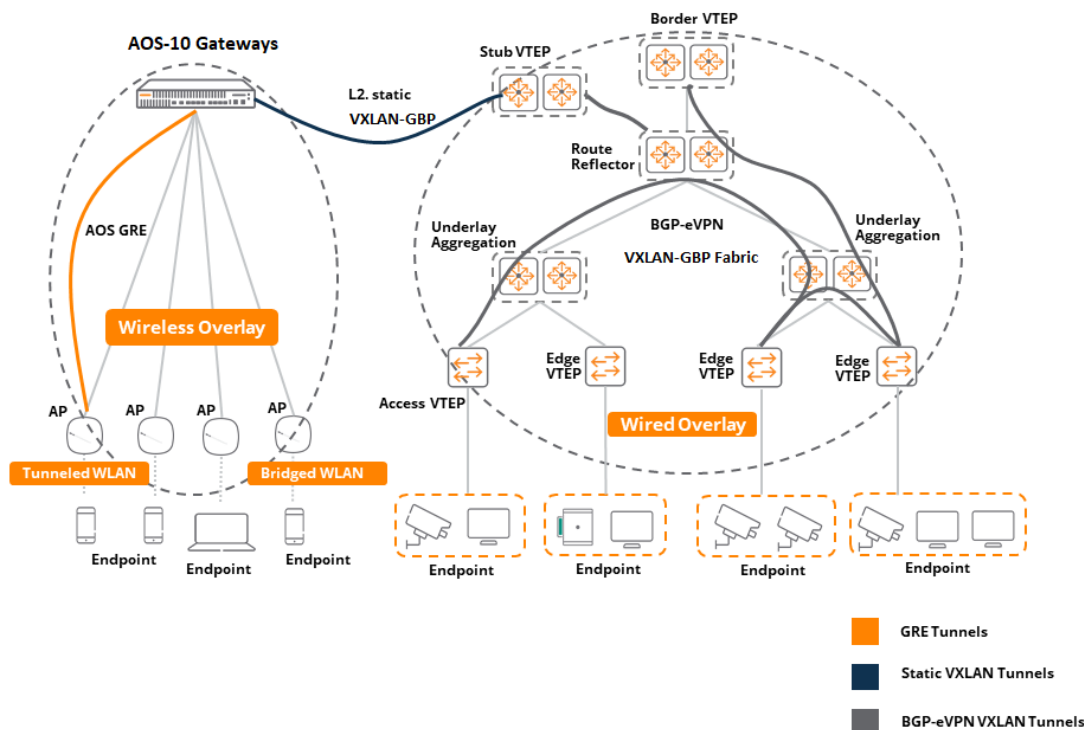
The following section explains the deployment scenarios of NetConductor:

- [Distributed Campus-wide Fabric](#)
- [Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch](#)
- [Centralized Multi-site Fabric with Third-Party Networks](#)

Distributed Campus-wide Fabric

The Distributed Campus-Wide Fabric enables large enterprises to deploy a multi-vendor and scalable overlay across the wired, wireless and wide area network and enables role-based policy enforcement at the edge of the fabric. The fabric consists of a standards-based EVPN-VXLAN fabric on the AOS-CX Switches and is extended to the AOS-10 Gateways using a static VXLAN tunnel. This fabric is based on BGP-EVPN and VXLAN protocols.

Figure 1 *Distributed Campus-wide Fabric Deployment*

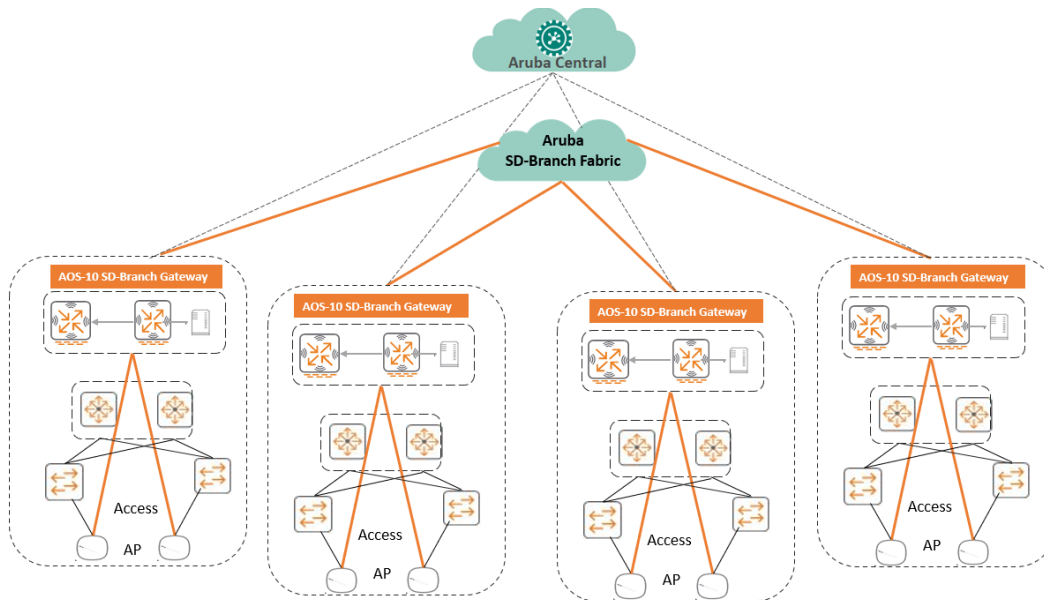


The Distributed Campus-wide Fabric deployment has an IBGP EVPN-VXLAN overlay on the AOS-CX switches, over an OSPF or IBGP underlay network. The underlay routing protocol is enabled to exchange the loopback addresses of all the devices in the fabric. The IBGP fabric is then established by peering all the edge devices in the fabric to the route-reflector, creating a full-mesh VXLAN-EVPN fabric. This fabric is added to the AOS-10 wireless overlay using a static VXLAN tunnel which enables role propagation from the wireless network to the wired network, and vice-versa.

Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch

In the Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch deployment, the AOS-10 HPE Aruba Networking Gateways act as WLAN and user-based tunnel gateways, and enables connectivity and role propagation over the SD-Branch WAN network. The following example explains the Centralized Multi-site Fabric deployment using a AOS-10 HPE Aruba Networking Gateways.

Figure 2 Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch Deployment

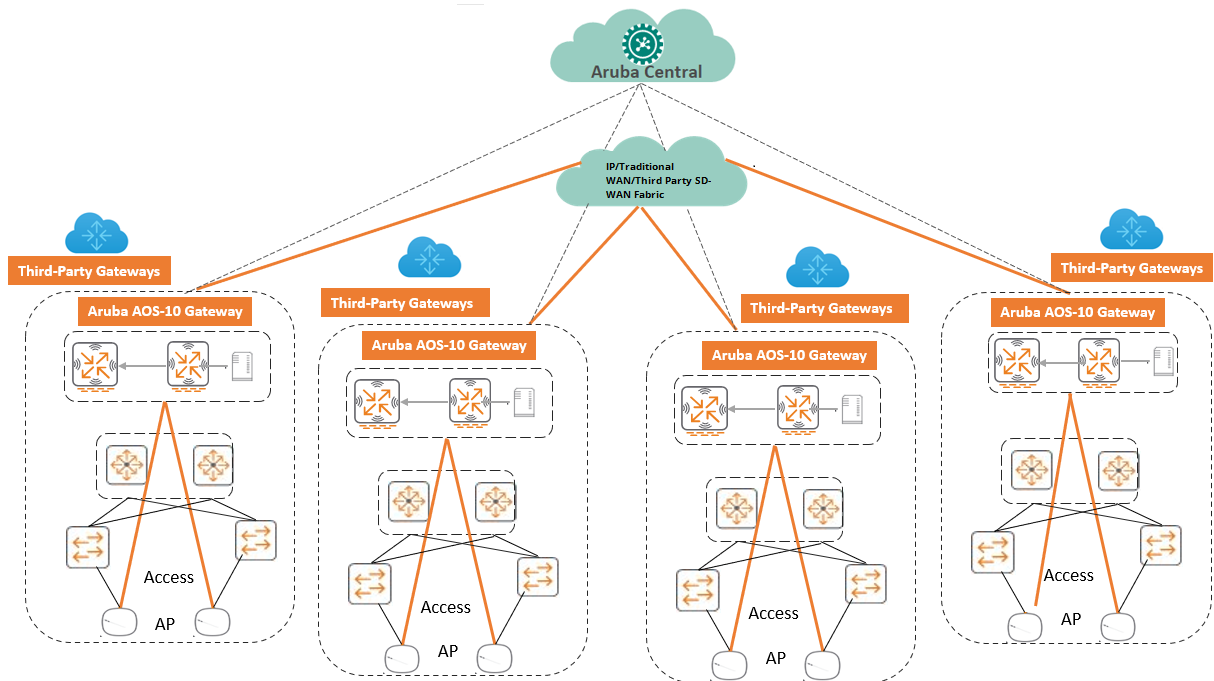


In the Centralized Multi-site Fabric deployment, customers can propagate role information and enforce role-based policies for client traffic across multiple sites connected by an HPE Aruba Networking SD-Branch fabric. The AOS-10 Gateways act as the WLAN and user-based tunnel Gateways for wired and wireless clients within each site, and act as the policy enforcement point for the clients within the site. To enforce role-based policies destined to clients across the fabric. The AOS-10 Gateways encapsulates the client traffic information with VxLAN-BGP and IPSEC, which contain role information in the GPID field in the VXLAN header. The AOS-10 Gateway in the destination site will then enforce the role-based policies for client traffic. Role propagation can also be selectively enabled on a per-group basis for SD-Branch deployments.

Centralized Multi-site Fabric with Third-Party Networks

In the Centralized Multi-site Fabric deployment with Third-Party SD-WAN, the AOS-10 HPE Aruba Networking Gateway act as WLAN and user-based tunnel gateways and enables role propagation over a 3rd-party SD-WAN network. The following example explains Centralized Multi-site Fabric deployment using a third-party SD-WAN Gateway.

Figure 3 Centralized Multi-site with Third-Party SD-WAN Fabric Deployment



In the Centralized Multi-site Fabric deployment, customers can propagate role information and enforce role-based policies for client traffic across multiple sites connected by an IP/Traditional WAN/Third Party SD-WAN Fabric . The AOS-10 Gateways act as the WLAN and UBT Gateways for wired and wireless clients within each site and act as the policy enforcement point for the clients within the site. To enforce role-based policies destined to sites across the IP/Traditional WAN/Third Party SD-WAN Fabric, the AOS-10 Gateways encapsulate the client traffic with VXLAN-GBP and IPSEC which contains the source role information in the GPID field in the VXLAN header. The AOS-10 Gateway in the destination site will then enforce the role-to-role policies for the client traffic. Role propagation can also be selectively enabled on a per subnet basis for third-party WAN deployments.

Lists all the supported Access Points, AOS-CX Switches, and gateways, and their supported software versions supported by NetConductor.

In this section, we cover the supported:

- [Aruba Gateways](#)
- [AOS-CX Switches](#)
- [Aruba Gateways](#)

Aruba Access Points

HPE Aruba Networking APs do not participate in role-based enforcement. The AOS-10 Gateways enforce role-based policies for all the wireless clients.

The following table shows the list of AP models and supported software required for deploying NetConductor.

Table 4: Supported APs

AP Family	AP Model	Supported Software Versions
300 Series	AP-303, AP-303H, AP-303P, AP-304, AP-305	AOS-10.4.0.0 or later
310 Series	AP-314, AP-315, AP-318	
320 Series	AP-324, AP-325	
	NOTE: The 320 Series AP models with 256 MB of SDRAM, manufactured between August 2015 and January 2016, are not supported with AOS-10. These 320 Series AP models have a serial number that begins with DD (for example, DD0003824).	
330 Series	AP-334, AP-335	
340 Series	AP-344, AP-345	
360 Series	AP-365, AP-367	
370 Series	AP-374, AP-375, AP-375EX, AP-375ATEX, AP-377, AP-377EX, AP-387	
500 Series	AP-503H, AP-504, AP-505, AP-505H	
510 Series	AP-514, AP-515, AP-518	
530 Series	AP-534, AP-535	
550 Series	AP-555	
560 Series	AP-565, AP-565EX, AP-567, AP-567EX	
570 Series	AP-574, AP-575, AP-575EX, AP-577, AP-577EX	
580 Series	AP-584, AP-585, AP-585EX, AP-587, AP-587EX	
630 Series	AP-635	
650 Series	AP-655	



As of ArubaOS 10.3 release the 6XX series APs are supported in NetConductor.

AOS-CX Switches

The following table lists the AOS-CX switches, supported software versions, and the supported personas in HPE Aruba Networking Central.

Table 5: Supported AOS-CX Switch Series and Software Versions

Switch Platform	Supported Software Versions	Supported Persona
AOS-CX 6200 Switch Series	10.12.0002 or later	Extended Edge
AOS-CX 6300 Switch Series	10.12.0001 or later	Edge Extended Edge
AOS-CX 6400 Switch Series	10.12.0001 or later	Edge Extended Edge
AOS-CX8100 Switch Series	10.12.0001 or later	Edge Stub Extended Edge
AOS-CX 8325 Switch Series	10.12.0001 or later	Route Reflector Border
AOS-CX 8360 Switch Series	10.12.0001 or later	Edge Stub Border Route Reflector
AOS-CX 8400 Switch Series	10.12.0001 or later	Route Reflector
AOS-CX 9000 Switch Series	10.12.0001 or later	Route Reflector
AOS-CX 10000 Switch Series	10.12.0001 or later	Route Reflector Border

Data sheets and technical specifications for the supported switch platforms are available at:
<https://www.arubanetworks.com/products/networking/switches/>.

Aruba Gateways

The following table shows a list of HPE Aruba Networking Gateway models and the supported software versions, required for deploying the NetConductor.

Table 6: Supported HPE Aruba Networking Gateways

Gateway Family	Gateway Model	Supported Software Versions
7000 Series	7005, 7008, 7010, 7024, 7030	AOS-10.4.0.0 or later
7200 Series	7205, 7210, 72207240XM, 7280	
9000 Series	9004, 9004-LTE, 9012	
9200 Series	9240	

Before you get started with your onboarding and provisioning operations, browse through the list of [Supported Devices](#) in HPE Aruba Networking Central.

This section includes the following topics:

- [Prerequisites](#)
- [Role-Based Policy](#)
- [VXLAN and BGP-EVPN Overview](#)

Prerequisites

The following prerequisites must be fulfilled to add an overlay to the network.

- All Switches must be onboarded on HPE Aruba Networking Central and added to the same UI group.
- All the required HPE Aruba Networking devices are on-boarded and connected to HPE Aruba Networking Central.

NetConductor Licensing

HPE Aruba Networking Central licensing is applicable to NetConductor. Advance licensing is required for Switches, APs, and Gateways.

NetConductor Caveats

NetConductor has the following caveats:

- Roles created globally in HPE Aruba Networking Central will not be applied to APs.
- Existing roles in APs, Gateways, and AOS-CX groups are not populated automatically in Global Clients Page.



Customer can recreate existing roles on Global Clients Page. Check [Global Client Roles](#).

- Multicast traffic is supported on the Switch fabric. The configuration is not supported on HPE Aruba Networking Central.
- Only role-to-role policies are allowed. Policies including protocol and port are not supported.
- Limitations of Policy Configuration for EdgeConnect 9.4.1:
 - NFR-1160—Policies are tied to the VRF this adds to duplication of efforts during configuration.
 - Zero-trust policy is cumbersome to apply in the current version of the the EC releases. This does not impact the end-to-end functionality.
 - NFR-1163—In Edge-Connect UI, configuring the route-target for BGP-EVPN is not a user friendly

when applying the "Route target" for BGP-EVPN on each vrf while configuring the VNI under the VxLAN Tab.

NetConductorFabric VSX Caveats

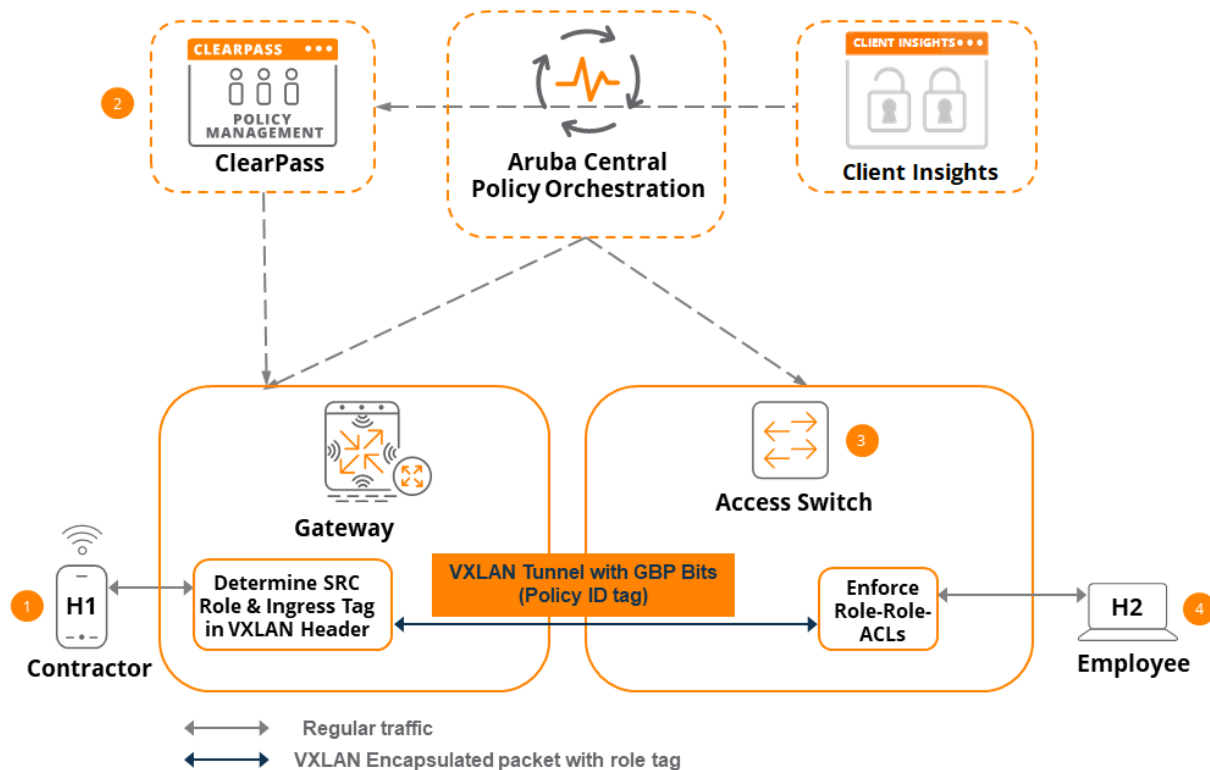
NetConductor Fabric VSX has the following caveats:

- The Fabric only supports pre-formed VSX device pairs.
- You cannot convert existing Fabric devices into VSX pairs.
- Before adding VSX Pairs to Fabric, the VSX Pairs should be fully formed. This can be viewed under **Monitoring Switches and Switch Stacks > Switch > VSX** and verified by the Config Sync Status column displays IN-SYNC.
- Decomposing Fabric VSX members into standalone devices is not supported unless both members are removed from Fabric first.

Role-Based Policy

The following section describes how role based policies are distributed and enforced with examples based on authentication and profiling.

Figure 4 Role-Based Policy—Wired Devices



Role-Based Policy for Wired Devices Using Identity Store Mapping

The following example describes role based policy enforcement using an identity store to assign roles to wired clients.

1. The wired client H2 (laptop) is connected to an access switch in the network. The access switch sends an 802.1x authentication request to ClearPass.
2. ClearPass has knowledge of the active directory group to which this H2 belongs and the respective role mapping group. In this case, the active directory group is Employee, and the role is Employee. If H2 is a Contractor, then the ClearPass assigns the Contractor role to H2. ClearPass communicates the role for the wired client to the access switch using the "HPE-User-Role" RADIUS attribute as part of the authentication response.
3. When the client H1 (Phone) sends a packet to the client H2 (Laptop), the gateway is now aware that H1 is a Contractor and based on this, it encapsulates the data packet in a VXLAN header and tags the group policy ID of the source role in the GBP field of the VXLAN header.
4. When the packet reaches the Switch where H2 is connected to, the access Switch is aware that the destination role of H2 is an Employee and determines whether the Contractor H1 is allowed to communicate to an Employee H2. If the defined policy between Contractor to Employee is "Deny", then the access Switch blocks the packet. If the defined policy between Contractor to Employee is "Allow", then the access Switch sends the packet to H2.

For more information about Role-Based Policy support on AOS-CX, refer to [Group-based policy](#).

Role-Based Policy Enforcement for Wired Clients Using Client Profiling

The following example describes role based policy enforcement using profiling to assign roles to wired clients.

For IoT clients such as IP-based light bulbs and cameras, there is usually no active directory group or MAC address directory for these devices. Hence, roles can be assigned to these devices by profiling these devices based static and dynamic attributes. When the light bulb is connected to the access switch, the device is initially authenticated using MAC-Auth and telemetry is sent to the Client Insights service on HPE Aruba Networking Central. Client Insights then profiles the device using attributes such as MAC OUI, DHCP Options and flow data. It uses this data to profile the client with the "IOT" tag and sends this tag to ClearPass. ClearPass now has the MAC address of the client and IOT tag mapping. When the light bulb authenticates, ClearPass assigns the IOT role to the client.

When the IoT client communicates to the laptop with the Employee role, the access Switch which the IOT client is connected to detects that the role for that MAC address is IoT. It encapsulates the data packet in a VXLAN header and tags the group policy ID of the source role in the GBP field of the VXLAN header. The packet then reaches the destination access Switch and determines that the source role is IoT and the destination role is an Employee. If the defined policy between IoT to Employee roles is deny, then the access Switch blocks the packet. If the defined policy between IoT to Employee roles is allow, then the access Switch sends the packet to the Employee laptop.

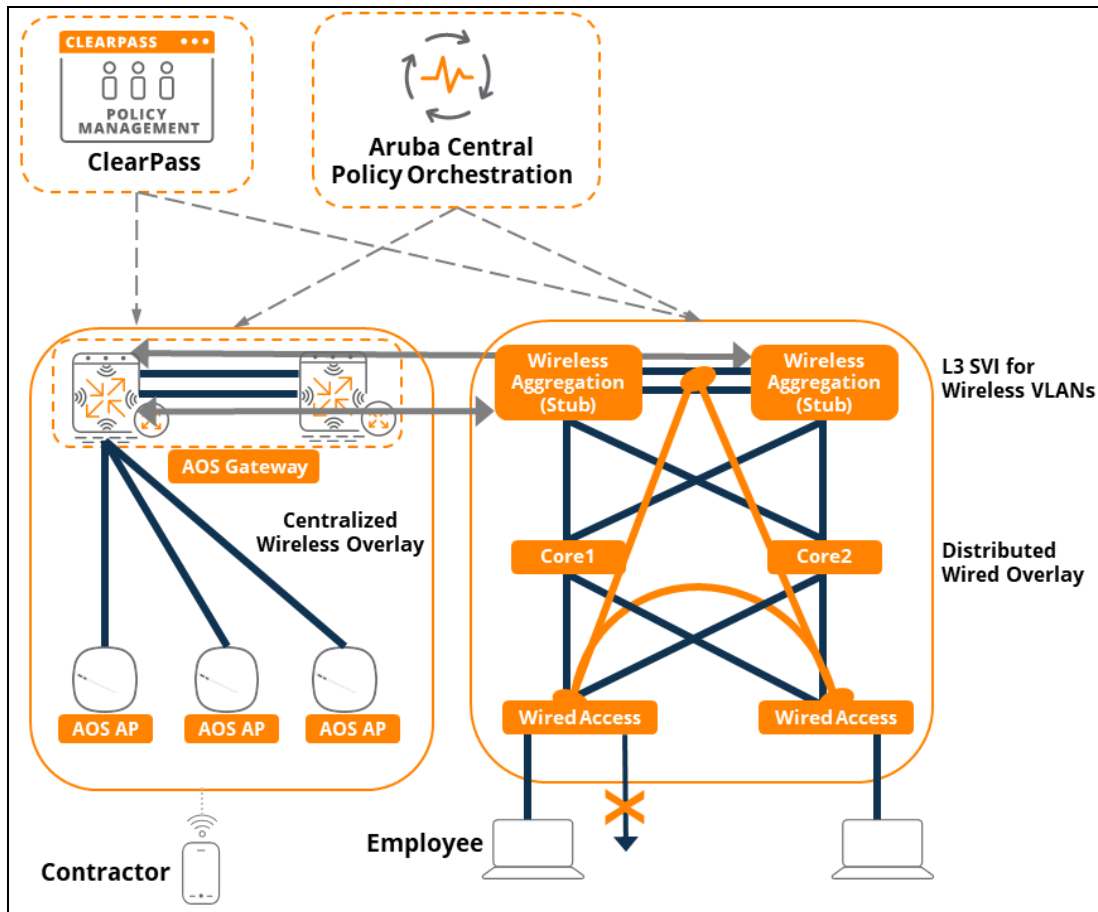
In an IoT device such as a light bulb which has an IP address, there is no active directory group. There is a device profiling which takes place in the network. When the light bulb is connected to the Switch, the telemetry data is sent to CPDI. CPDI then profiles the device, for example the device's make or device series and tags it with an IoT tag and sends it to ClearPass. ClearPass now has the MAC address of the endpoint and the IoT tag on it. When the endpoint tries to authenticate, ClearPass determines if the tag is IoT, then the role assigned is IoT. The access Switches now have the roles for each of the devices.

When the IoT device communicates to a camera at the endpoint, the access Switch detects that the role for that MAC address is IoT and the corresponding group policy ID. The access Switch then adds them in the source header of the VXLAN packet. The packet then reaches the access Switch and determines that the source is IoT and the destination is a camera. If the defined policy IoT to camera is deny, then the access Switch blocks the packet. If the defined policy IoT to camera is allow, then the access Switch delivers the packet.

Role-Based Policy for Wireless Devices

The following example describes role based policy enforcement for wireless clients.

Figure 5 Role-Based Policy—Wireless Devices



Wireless clients that connects to an AOS-10 Tunneled WLAN SSID are authenticated by the AOS-10 Gateway via ClearPass, based on either an identity store mapping or client profiling (see above section). The client traffic is tunneled from the AOS-10 APs to the AOS-10 Gateway through a tunneled WLAN. The AOS-10 Gateway authenticates the client with ClearPass and gets the role of the client from ClearPass using the "HPE-User-Role" RADIUS attribute. If the gateway does not have knowledge of both the source and destination role of the traffic, the Gateway encapsulates the data traffic in a VXLAN header and tags the group policy ID of the source role in the GBP field of the VXLAN header. In this example, ClearPass assigns the Contractor role to Phone. The AOS-10 Gateway sends the data traffic on the static VXLAN tunnel to the Wireless Aggregation switch with the GPID of the role Contractor. The destination endpoint in this example is a wired endpoint with the role Employee. The traffic is allowed or denied by the access switch based on the defined Role-Based Policy.

VXLAN and BGP-EVPN Overview

This section introduces you to VXLAN and EVPN, and explains how they can be implemented using Switches.

VXLAN Overview

Virtual eXtensible LAN (VXLAN) is a MAC-in-UDP technology that provides layer 2 connectivity between networks across an IP network. VXLAN is typically used to extend layer 2 segments across an entire data center or between remote data centers. It is also used to provide multi-tenancy services.

VXLAN addresses the requirements of the Layer 2 and Layer 3 data center network infrastructure in the presence of VMs in a multi-tenant environment. It runs over the existing networking infrastructure and provides a means to stretch a Layer 2 network. In short, VXLAN is a Layer 2 overlay scheme on a Layer 3 network. Each overlay is termed a VXLAN segment. Only VMs within the same VXLAN segment can communicate with each other. Each VXLAN segment is identified through a 24-bit segment ID, termed the VXLAN Network Identifier (VNI). This allows up to 16 M VXLAN segments to coexist within the same administrative domain.

The VNI identifies the scope of the inner MAC frame originated by the individual VM. Thus, you could have overlapping MAC addresses across segments but never have traffic cross over since the traffic is isolated using the VNI. The VNI is in an outer header that encapsulates the inner MAC frame originated by the VM. In the following sections, the term VXLAN segment is used interchangeably with the term VXLAN overlay network.

VXLAN Packet Format

Role-to-role policy is defined at the global level in . Role definitions and policy definitions are global to a customer and are distributed by the underlying device configuration system with all the devices as required. These roles and policies are embedded in the VXLAN header.

For more information about different components in the VXLAN Group Based Policy Extension (VXLAN-GBP) header, refer to [VXLAN Group Policy Option](#).

VXLAN Deployment Models

AOS-CX supports the following VXLAN implementations:

- Centralized Routing
- Distributed Routing

Centralized routing and distributed routing defines where the inter-VLAN routing happens. Whether inter-VLAN routing is happening only on the core/border Switches or inter-VLAN routing is performed on the edge devices such as access Switches.



Both centralized routing and distributed routing models use EVPN control plane. solution utilizes the distributed routing model.

EVPN VXLAN Overview

EVPN-VXLAN refers to a network fabric that extends layer 2 connectivity as a network overlay over an existing physical network. It is an open standards technology that creates more agile, secure, and scalable networks in campuses and data centers. EVPN-VXLAN consists of:

- Ethernet VPN (EVPN) which is used as the overlay control plane and provides virtual connectivity between different layer 2/3 domains over an IP.
- Virtual extensible LANs (VXLAN), a common network virtualization overlay protocol that expands the layer 2 network address space from 4,000 to 16 million.

How EVPN and VXLAN Work Together

EVPN-VXLAN enables businesses to connect geographically dispersed locations using layer 2 virtual bridging. EVPN-VXLAN provides the scale required by cloud service providers and is often the preferred technology for data center interconnections.

EVPN, as an overlay, supports multi-tenancy and is highly extensible, often using resources from different data centers to deliver a single service. It can provide layer 2 connectivity over physical infrastructure for devices in a virtual network or enable layer 3 routing.

Because it serves as a MAC address learning control plane for overlay networks, EVPN can support different data plane encapsulation technologies.

VXLAN encapsulates layer 2 ethernet frames in layer 3 UDP packets, meaning virtual layer 2 subnets can span underlying layer 3 networks. A VXLAN network identifier (VNI) is used to segment each layer 2 subnet similarly to traditional VLAN IDs.

A VXLAN tunnel endpoint (VTEP) is a VXLAN-capable device that encapsulates and de-encapsulates packets. In the physical network, a Switch typically functions as a layer 2 or layer 3 VXLAN Gateway and is considered a hardware VTEP. The virtual network equivalents are known as software VTEPs, which are hosted in hypervisors such as VMware ESXi or vSphere.



For the BGP session for underlay and overlay, currently supports iBGP only.

Global Client Roles



Global Client Roles support is selectively available. Contact your Aruba Account Manager to enable it in your HPE Aruba Networking Central account.

HPE Aruba Networking Central allows you to configure client roles for role-to-role policy enforcement. To make client roles more powerful and consistent within the dynamic-segmentation security framework, you can now manage the client roles globally from HPE Aruba Networking Central. For example, if you have five roles in your network, a role-to-role policy is created by defining permissions from one role to another.

Prerequisites

Ensure that all the required HPE Aruba Networking devices are on-boarded and connected to HPE Aruba Networking Central, prior to configuring global client roles and performing selective enablement of subnets.

For more information, see the topics listed below:

- [Configuring the Global Client Roles](#)
- [Editing a Role](#)
- [Deleting a Role](#)

- [Performing Selective Enablement of Subnets](#)
- [Deleting a Subnet](#)
- [Performing Selective Enablement of Groups](#)
- [Deleting a Group](#)

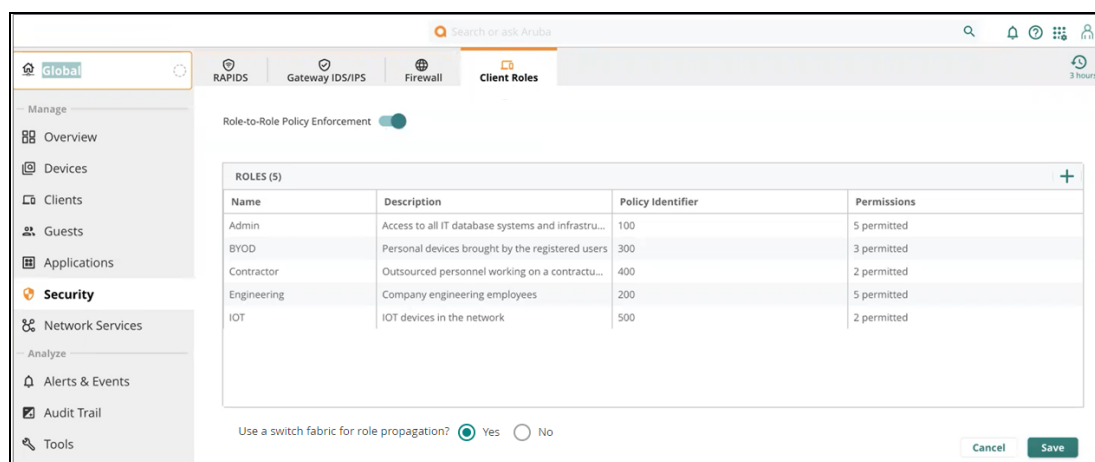
Configuring the Global Client Roles

Global client roles configuration affects all the access Switches and wireless Gateways in the network, as part of the global policy enforcement.

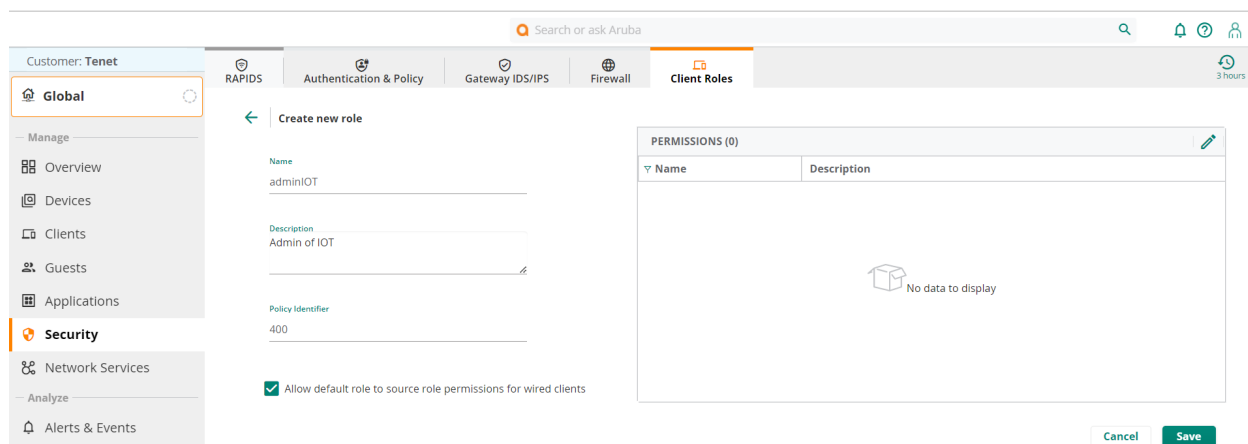
To configure client roles for role-to-role policy enforcement, complete the following procedure:


1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. Toggle the **Role-to-Role Policy Enforcement** switch to the on position, to apply the roles and permissions that are defined.

Figure 6 Client Roles Page



3. In the **Roles** table, click the **+** icon. The **Create new role** page is displayed.



4. In the **Create new role** page, configure the following parameters:
 - **Name**—Enter the name of the role that you want to create. For example, Admin.
 - **Description**—Enter the description of the role. For example, all employees in the organization.
 - **Policy Identifier**—Policy Identifier is a unique auto-generated identifier for the role. You can modify identifier, if required. For example, 100.
5. Enable the **Allow default role to source role permissions for wired clients** checkbox if you want to allow unauthenticated client traffic and external traffic (traffic with Policy Identifier 0) for the wired clients.
6. In the **Permissions** table, click the  edit icon to create bi-directional policies between roles. The **Allow Source to Destination** column is used to enable permissions from the source role to the selected destination role. The **Allow Destination to Source** column is used to enable permissions from the selected destination role to the source role. For example, if you open the **Assign Permissions** window for the Admin role, you must select the check box corresponding to the BYOD role in the **Allow Source to Destination** column to allow network traffic from Admin role to BYOD role. You must click **IOT** under the **Allow Destination to Source** column to allow network traffic from IOT role to Admin role.



Selecting either of the check boxes for self, selects both the check boxes by default. Selecting self enables bi-directional network traffic between devices of the same role.

Figure 7 *Assign Permissions Page*

Assign Permissions

Assign permissions for source role **Admin**

DESTINATION ROLES (5)		
Name	Allow Source to Destination	Allow Destination to Source
Admin (self)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
BYOD	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Contractor	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Engineering	<input type="checkbox"/>	<input checked="" type="checkbox"/>
IOT	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Cancel
Assign

7. In the **Allow Source to Destination** and the **Allow Destination to Source** columns, select the appropriate check box.
8. Click **Assign**.
9. In the **Propagate roles between Aruba Gateways?** options, select **Yes** if you have Branch Gateway or Mobility Gateway for role propagation else Select **No**.
When the **Branch** mode is enabled, the gateways can propagate role information and enforce role-based policies for client traffic across multiple sites across the SD-WAN traffic. When the **Mobility** mode is enabled, gateways can propagate role information and enforce role-based policies for client traffic across multiple gateways, irrespective of the network infrastructure between them.
10. In the **Create new role** page, click **Save**.
The role is created, and the assigned permissions are saved.



Configuring custom policy rules on switches is possible using the Multi-Editor with sequence numbers above 9999.

Editing a Role

To edit a role, complete the following procedure:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.

ROLES (97) +			
▼ Name	Description	Policy Identifier	Permissions ⌵
Admins		3200	0 permitted
COA		4000	0 permitted
CP-PROFILE-100		8100	0 permitted ✎ 🗑
Camera		2400	0 permitted
Campus-IOT		2600	0 permitted
Campus-WPA3-PSK		5900	0 permitted
Campus-employee		3500	0 permitted
Campus-ent-device		2700	0 permitted
Campus-ent-dot1x		5200	0 permitted

2. In the **Roles** table, hover over the row, and click the ✎ edit icon.
3. In the **Edit Role** window, modify the required permissions.




The **Allow default role to source role permissions for wired clients** checkbox cannot be modified while editing a role.

4. Click **Save**.

Deleting a Role

To delete a role, complete the following procedure:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. In the **Roles** table, select the role and click the  delete icon. The **Confirm Action** window is displayed.
3. Click **Yes**.

Performing Selective Enablement of Subnets



Perform selective enablement of subnets if you have not configured distributed overlay fabric for role propagation. For more information about distributed overlay fabric, see [Fabric Wizard Overview](#).

Selective enablement of subnets is used to select the client subnets for which you want to enable role-based policy enforcement on the third-party Gateways. Typically, these client subnets are mapped to different sites in the network.


To selectively enable the subnets that map to the clients in the respective sites using HPE Aruba Networking Central, complete the following procedure:


1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. In the **Use a switch fabric for a role propagation?** options, select **No**.




Delete all configured subnets before enabling a switch fabric for role propagation.

3. In the **What is the role of the Aruba gateways?** options, select **Mobility**. Delete all the configured subnets to change the role to Branch Gateways.

Use a switch fabric for role propagation? Yes No  Please delete all configured subnets and groups before enabling a switch fabric for role propagation. VxLAN tunnels will have to be configured on the switch fabric and gateways for role propagation


What is the role of the Aruba gateways? Branch Mobility  Please delete all configured subnets to change the role to WLAN with branch gateways

ENFORCE ROLES BETWEEN THESE SUBNETS (3)	
Subnet	Description
192.168.201.0/24	Site 1
192.168.202.0/24	Site 2 

4. Click the + icon and then enter the subnet details and the description.
5. Click **Save**.

Deleting a Subnet

To delete a subnet, complete the following procedure:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. In the **ENFORCE ROLES BETWEEN THESE SUBNETS** table, select the subnet and click the  delete icon. The **Confirm Action** window is displayed.
3. Click **Yes**.

Performing Selective Enablement of Groups



Perform selective enablement of groups if you have not configured distributed overlay fabric for role propagation. For more information about distributed overlay fabric, see [Fabric Wizard Overview](#).

Selective enablement of groups is used to select the HPE Aruba Networking Central group for which you want to enable role propagation on AOS-10 EdgeConnect SD-Branch gateway. Typically, these client groups are mapped to different SD-WAN branch sites within the HPE Aruba Networking Central network.

To selectively enable the groups that map to the clients in the respective sites using HPE Aruba Networking Central, complete the following procedure:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. In the **Use a switch fabric for a role propagation?** options, select **No**.



Delete all configured groups before enabling a switch fabric for role propagation.

3. In the **What is the role of the Aruba gateways?** options, select **Branch**. Delete all configured groups to change the role to WLAN gateways.
4. Click the + icon to add the groups.

Use a switch fabric for role propagation? Yes No ⓘ Please delete all configured subnets and groups before enabling a switch fabric for role propagation. VxLAN tunnels will have to be configured on the switch fabric and gateways for role propagation

What is the role of the Aruba gateways? Branch Mobility

SELECT GROUPS TO ENFORCE ROLES	
<input type="checkbox"/>	Group Name
<input type="checkbox"/>	Abhishek-Home
<input type="checkbox"/>	AOS10 AP Lab Malleshwaram
<input checked="" type="checkbox"/>	AOS10 Bangalore HQ IAP 🗑
<input checked="" type="checkbox"/>	AOS10xCampBrGrp 🗑
<input checked="" type="checkbox"/>	AOS10xCampMobGrp 🗑
<input type="checkbox"/>	AOS10xMicroBrGrp
<input type="checkbox"/>	Auto
<input type="checkbox"/>	Bangalore HQ
<input type="checkbox"/>	clarity

5. Select the groups you want to assign.
6. Click **Assign**.
7. Click **Save**.

Deleting a Group

To delete a role, complete the following procedure:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Security > Client Roles**. The **Client Roles** page is displayed.
2. In the **SELECT GROUPS TO ENFORCE ROLES** table, select the group and click the 🗑 delete icon. The **Confirm Action** window is displayed.
3. Click **Yes**.

Network Wizard Overview

Network Wizard is the workflow to set up an underlay network. To establish an underlay network, HPE Aruba Networking Central automates the L3 and OSPF switching connectivity by auto-assigning the loopback IP address, ROP links, and creating the HPE Aruba Networking Virtual switching Extension.



NOTE

Network Wizard workflow requires HPE Aruba Networking Central Advanced licenses. A notification message is displayed indicating a valid advanced license is required for this feature. You can acknowledge the message by clicking the close button. The notification message is user-specific. Once the message is closed, it will not be displayed for the same user, but for other users, the message will be displayed. For more information, refer to the *AOS-CX switches Feature Pack Ordering Guide* and EULA.

Prerequisites

To complete the provisioning of underlay network using HPE Aruba Networking Central, complete the following

- All HPE Aruba Networking CX switches must be running AOS-CX 10.10.1020 or later.
- Move all devices and switches that are on-boarded in HPE Aruba Networking Central into the same UI group.
- All the physical cabling between the switches needs to be setup and running.
- LLDP must be enabled on all devices and ports.
- Ensure all the previously existing underlay network configurations are removed from the switches.

For more information, see the listed topics below:

- [Setting Up a Campus \(3 tier L3 access\) Underlay Network](#)
- [Setting Up a of Data Center \(Spine and Leaf\) Underlay Network](#)
- [Modifying an Underlay Network](#)
- [Deleting an Underlay Network](#)

Setting Up a Campus (3 tier L3 access) Underlay Network

To set an underlay network of Campus or (3 tier L3 access) for the AOS-CX switches, complete the following workflow.

- [Step 1: Configuring the Network Type](#)
- [Step 2: Setting the Structure of the Underlay Network](#)
- [Step 3: Assigning Devices to the Network](#)
- [Step 4: Configuring the Underlay Network](#)
- [Step 5: Viewing the summary Page](#)



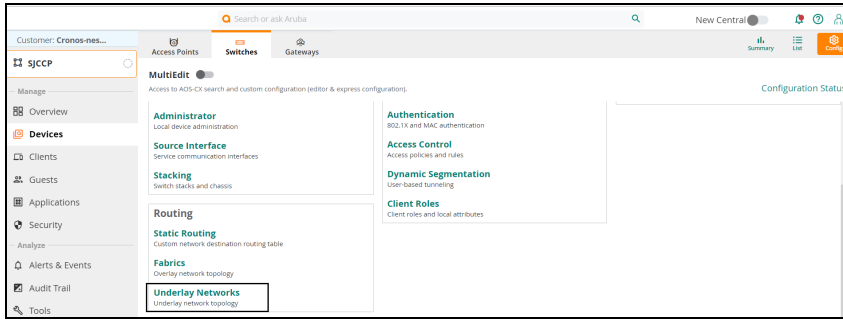
Underlay network requires HPE Aruba Networking Central Advanced licenses. A notification message is displayed indicating a valid advanced license is required for this feature. You can acknowledge the message by clicking the close button. The notification message is user-specific. Once the message is closed, it will not be displayed for the same user, but for other users, the message will be displayed. For more information, refer to the *AOS-CX switches Feature Pack Ordering Guide* and EULA

Step 1: Configuring the Network Type

To configure an underlay network for an AOS-CX switch, complete the following steps:

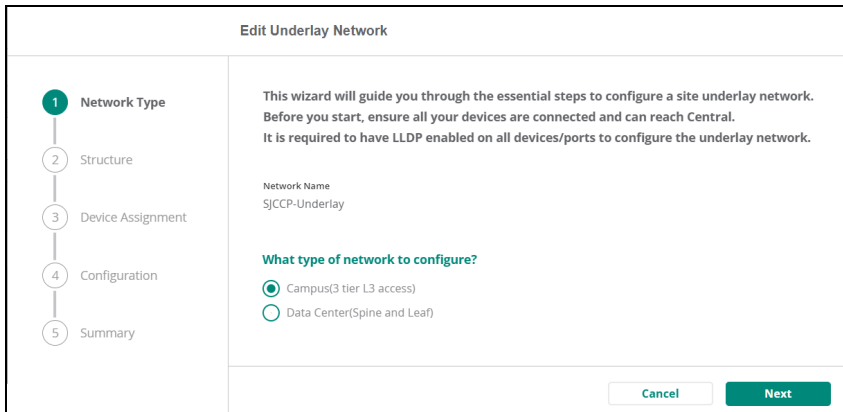
1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > switches**.
 - c. To view the AOS-CX switch configuration dashboard, click the **Config** icon.

2. Click **Routing > Underlay Networks**.



3. In the **Networks** table, click the **+** icon.

The **Guided Setup for Underlay Network** page is displayed.



4. Enter the **Network Name**.

5. Select the **Campus(3 tier L3 access)** checkbox under **What type of Network to configure**.

6. To display the **Structure** section, click **Next**.

Step 2: Setting the Structure of the Underlay Network

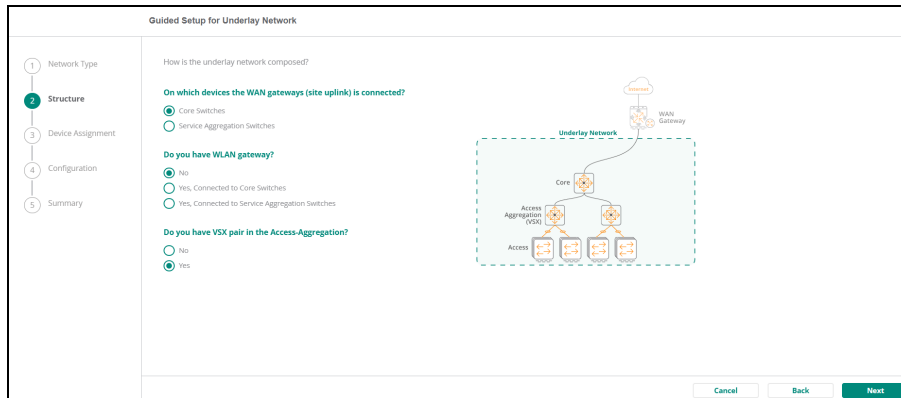
To assign WAN and WLAN gateways to the core switches, complete the following step.

1. In the **Structure** section, select the type of switch to be connected to the WAN gateway.
 - **Core switches**—Connects the WAN gateway directly to the core switches.
 - **Service Aggregation switches**—Connects the WAN gateway to the core switches through a service aggregation switch.
2. Select one of the checkboxes if you are connected to the WLAN gateway.
 - **No**—Select if there is no WLAN gateway connected directly to one or more core switches.
 - **Yes, Connected to Core switches**—Select when there is a WLAN gateway connected to the core switches.
 - **Yes, Connected to Service Aggregation switches**—Select when there is a WLAN gateway connected to the core switches through a service aggregation switch.



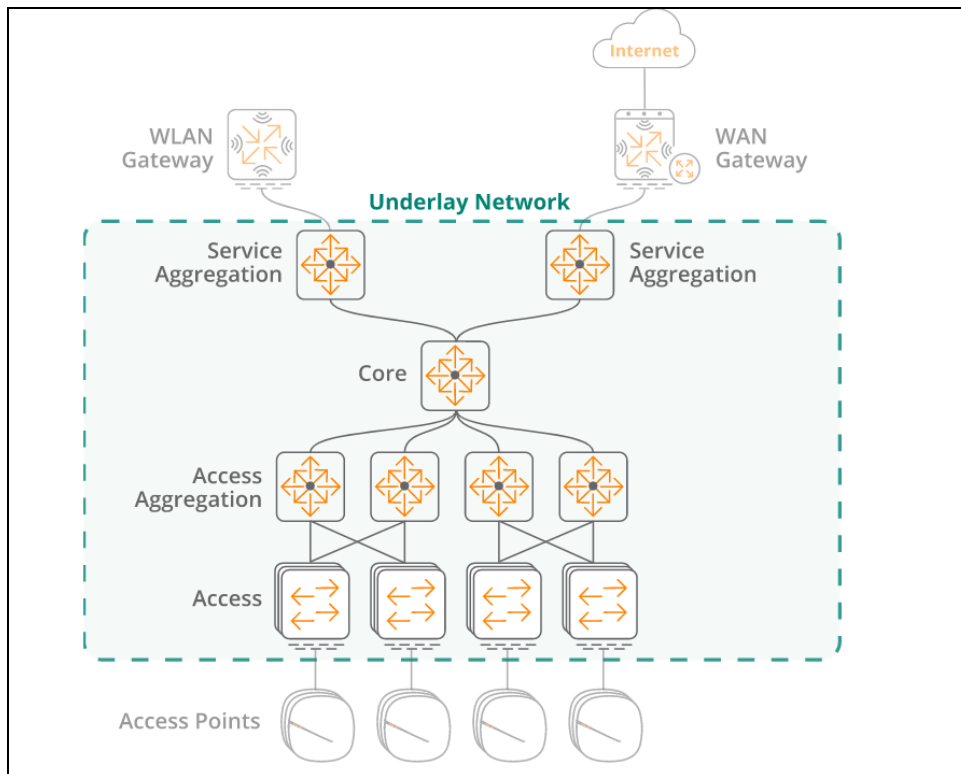
The **Do you have VSX pair in Access Aggregation** checkbox is an allowlisted option. Contact your administrator for access.

3. Select the checkbox for VSX pair in the Access-Aggregation.
 - **No**—Select if you want to deploy routed access design for the switches.
 - **Yes**—Select if you want to deploy scaled access design for the switches.
4. To display the **Device Assignment** section, click Next.



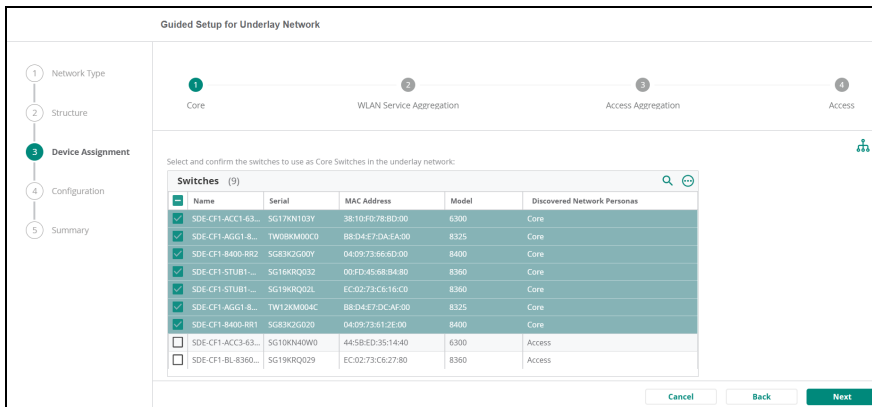
Step 3: Assigning Devices to the Network

This section displays the device persona suggestions for the core switches present in the network.



The **Device assignment** section consists of subsections that auto populate suggestions depending on the selections that you made in the **Structure** section. The **switches** table displays the **Name**, **Serial**, **MAC Address**, **Model**, and **Discovered Network Personas** across all subsections of the **Device Assignment** page.

- **Core**—Select new switches or edit the selection based on your requirements and click **Next**.
- **Access Aggregation**—Select and confirm the switches to use as access aggregation switches in the underlay network and click **Next**.
- **WAN Service Aggregation**—This section is displayed when WAN gateways (site uplink) are connected to the service aggregation switches. Select new switches or edit the selection based on your requirements and click **Next**.
- **WLAN Service Aggregation**—This section is displayed when WLAN gateways are connected to the service aggregation switches. Select new switches or edit the selection based on your requirements and click **Next**.
- **Access**—Select new switches or edit the selection based on your requirements and Click **Next** to display **Configuring your underlay network** page.





Step 4: Configuring the Underlay Network

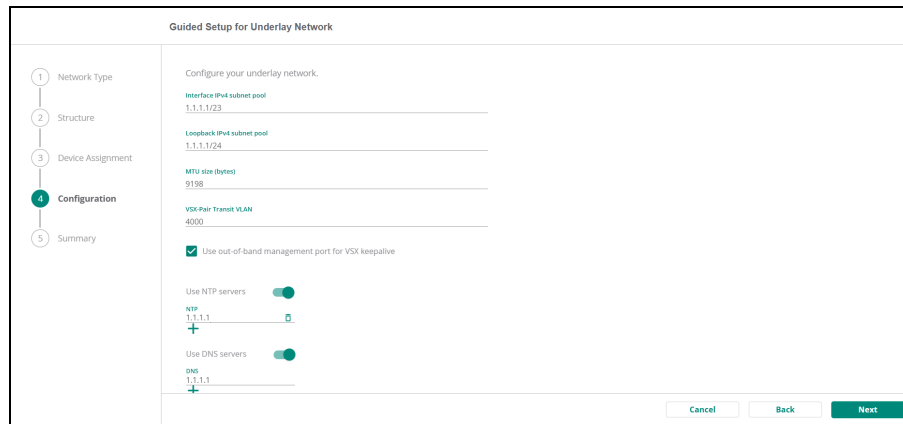
The **Configuring your underlay network** page displays various parameters to configure the underlay network. After providing the required information for the parameters. Click **Next** to open the Summary Page.

The following table describes the parameters in the **Configuring your underlay network** table.

Table 7: Configuring your Underlay Network Parameters.

Column	Function
Interface IPv4 subnet pool	Enter the ROP interface IPv4 subnet pool address for the underlay network.
Loopback IPv4 subnet pool	Enter the loopback IPv4 subnet pool address for the underlay network.
MTU size (bytes)	Enter the MTU size within the range of 1–9168 bytes.
VSX-Pair Transit VLAN	Enter the VSX-pair transit VLAN within the range of 1–4098.
Use out-of-band management port for VSX keepalive	The toggle remains disabled by default to use dedicated ports for the VSX keepalive. Enable the toggle to use out-of-band ports for VSX keepalive

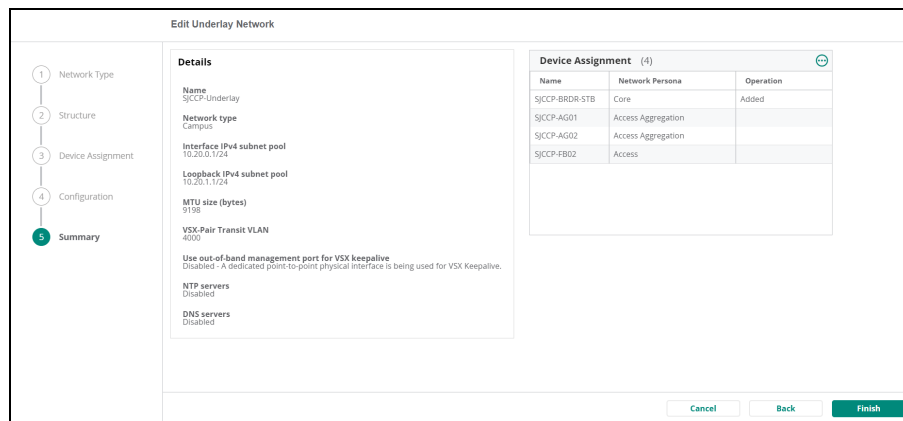
Column	Function
Use NTP servers	Enable the toggle to enter the NTP server address for the underlay network. You can add more than one address for the server by clicking the add (+) icon. Clicking the delete  icon deletes the NTP server address.
Use DNS servers	Enable the toggle to enter the DNS server address for the underlay network. You can add more than one address for the server by clicking the add (+) icon. Clicking the delete  icon deletes the DNS server address.



Step 5: Viewing the summary Page

This section displays all the parameter that you have configured under the **Details** and **Device Management** sections.

Click **Finish** to complete the underlay network configuration.



Setting Up a of Data Center (Spine and Leaf) Underlay Network

To set an underlay network of Data Center (Spine and Leaf) for the AOS-CX switches, complete the following workflow.

- [Step 1: Configuring the Network Type](#)
- [Step 2: Setting the Structure of the Underlay Network](#)
- [Step 3: Assigning Devices to the Network](#)
- [Step 4: Configuring the Underlay Network](#)
- [Step 5: Viewing the Summary Page](#)

Step 1: Configuring the Network Type

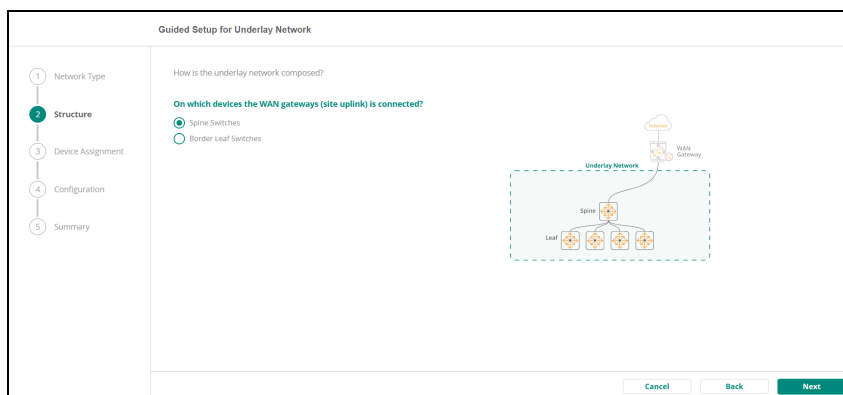
To configure an underlay network for the AOS-CX switch, complete the following steps:

1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > switches**.
 - c. To view the AOS-CX switch configuration dashboard, click the **Config** icon.
2. Click **Routing > Underlay Networks**.
3. In the **Networks** table, click the + icon.
The **Guided Setup for Underlay Network** page is displayed.
4. Enter the **Network Name**.
5. Select the **Data Center(Spine and Leaf)** checkbox under **What type of Network to configure**.
6. Click Next to display the **Structure** section.

Step 2: Setting the Structure of the Underlay Network

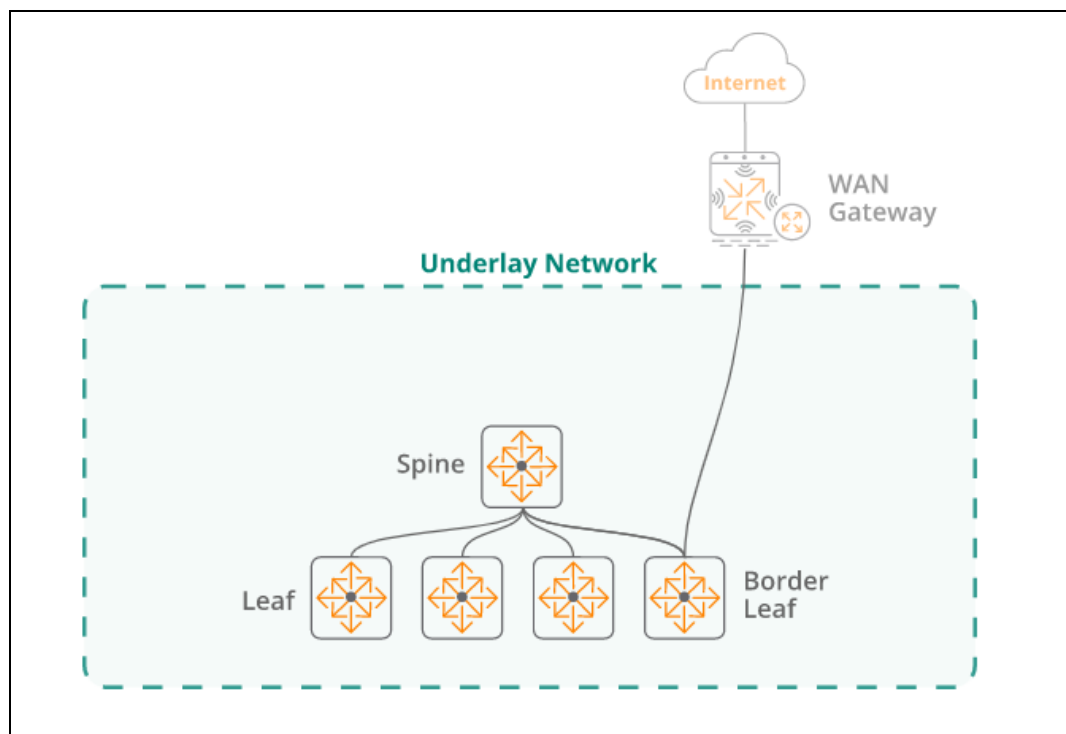
To assign WAN gateways to the spine or border switches, complete the following steps.

1. In the **Structure** section, select the type of switch to be connected to the WAN gateway.
 - a. **Spine switches**—Connects the WAN gateway to the spine switches.
 - b. **Border switches**—Connects the WAN gateway to the border switches.
2. To display the **Device Assignment** section, click **Next**.



Step 3: Assigning Devices to the Network

This section displays the device persona suggestions for the switches present in the network.



The **Device assignment** section consists of three subsections of **Spine**, **Border**, and **Leaf**. The **switches** table displays the **Name**, **Serial**, **MAC Address**, **Model**, and **Discovered Network Personas** across all subsections of the Device assignment page.

- **Spine**—Select the switches to use as spine switches in the underlay network and click **Next**.
- **Border**—Select the switches to use as border switches (uplink) in the underlay network and click **Next**.
- **Leaf**—Select the switches to use as leaf switches in the underlay network and click **Next** to display **Configuring your underlay network** page.

Step 4: Configuring the Underlay Network



The **Configuring your underlay network** page displays various parameters to configure underlay network. After providing the required information for the parameters.

Click **Next** to open the **Summary** Page.

The following table describes the parameters in the **Configuring your underlay network** table.

Table 8: *Configuring your Underlay Network Parameters.*


Column	Function
Interface IPv4 subnet pool	Enter the ROP interface IPv4 subnet pool address for the underlay network.
Loopback IPv4 subnet pool	Enter the loopback IPv4 subnet pool address for the underlay network.
MTU size (bytes)	Enter the MTU size within the range of 1-9198 bytes.

Column	Function
VSX-Pair Transit VLAN	Enter the VSX-pair transit VLAN within the range of 1–4094.
Use out-of-band management port for VSX keepalive	The toggle remains disabled by default to use dedicated ports for the VSX keepalive. Enable the toggle to use out-of-band ports for VSX keepalive
Use NTP servers	Enable the toggle to enter the NTP server address for the underlay network. You can add more than one address for the server by clicking the add (+) icon. Clicking the delete  icon deletes the NTP server address.
Use DNS servers	Enable the toggle to enter the DNS server address for the underlay network. You can add more than one address for the server by clicking the add (+) icon. Clicking the delete  icon deletes the DNS server address.

Step 5: Viewing the Summary Page

This section displays all the parameter that you have configured under the **Details** and **Device Management** sections.

Click **Finish** to complete the underlay network configuration.

To delete an underlay network, click the  delete icon in the tunnel row.

Modifying an Underlay Network




The **Network Type** and **Structure** sections are not editable.

You can add devices to the underlay network incrementally by editing the underlay network. To edit or modify an underlay network, complete the following steps:

- In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - Set the filter to a group.

The dashboard context for the group is displayed.

- b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Underlay Networks**.
3. Hover over the tunnel row in the **Networks** table and click the  edit icon.
4. Edit the details for **Device Assignment** and complete this section if you have added new devices to your underlay network.
5. Edit the details for **Configuration** and complete this section if you have added new devices to your underlay network.
6. Click Next to display the **Summary** section
7. The **Summary** section highlights the new devices that were added in the network.
8. Click **Finish** in the **Summary** section to save the new underlay network.


Deleting an Underlay Network

To delete an underlay network, complete the following steps:



Deleting an underlay network only removes the network configuration from HPE Aruba Networking Central but not the actual switches. You have to manually remove the switch configurations by using the multi-editor mode.

1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.

The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Underlay Networks**.
3. Hover over the tunnel row in the **Networks** table and click the  delete icon.
4. Click **Yes** to delete.

The underlay network is deleted.

HPE Aruba Networking Central allows you to create an overlay fabric by configuring VXLAN tunnels between Stub and Edge Switches, and overlay segments for the existing overlay fabrics.

Prerequisites

Before provisioning the overlay fabric using HPE Aruba Networking Central, complete the following prerequisites:

- All HPE Aruba Networking CX switches should be on firmware version 10.10.1020 and above.
- All Switches and the required HPE Aruba Networking devices must be on-boarded on HPE Aruba Networking Central and added to the same UI group.
- The underlay network must be set up before attempting to add any overlay to the network.
- At least two loopback (loopback 0 and loopback 1) interfaces are configured on all devices of the overlay fabric.
- The IP addresses of the loopback interfaces are unique on each device of the overlay fabric except on the VSX pair (VSX pair has at least one loopback interface whose IP address is the same on both primary and secondary).
- At least one instance of OSPF or BGP should be configured on all devices of the overlay fabric. OSPF or BGP router ID should be the IP address of an existing loopback interface.
- Modifying configurations of the Fabric Wizard using UI config or Multi-Edit is not supported. This causes the devices going out of sync and / or traffic loss.
- Fabric configuration referenced outside the Fabric context can cause undesirable outcomes such as devices going out of sync when the Fabric configurations are modified by the Fabric Wizard.
- OSPF or BGP must be configured on at least two loopback interfaces (one is used for OSPF and BGP, and the other for the VXLAN source IP).
- If the device has VSX configuration, the following validations must be performed:
 - The VSX pair should have minimum one loopback interface with the same IP address (used as VXLAN source IP).
 - Keep alive is configured.
 - ISL link—LAG interface with member ports is configured.
- In the case fabric was created with firmware version lower than 10.12.000 version on HPE Aruba Networking CX 8325, 10000, it is necessary to configure 'gbp enable' on these devices using the Multi-editor when Role-to-Role Policy Enforcement is enabled.

For more information on overlay fabric and overlay segment see the following topics:

- [Overlay Fabric Overview](#)
- [Overlay Segment Overview](#)

Overlay Fabric Overview

The distributed overlay fabric is a group of AOS-CX Switches that are part of the BGP-EVPN VXLAN overlay. The overlay fabric is created by configuring VXLAN tunnels between Stub and Edge Switches. The AOS-10 Gateways also participate in the overlay fabric, using static VXLAN tunnels to the Stub VTEP. HPE Aruba Networking Central allows you to configure the overlay fabric on top of a virtual network using AOS-CX Switches. You can assign a list of fabric person as such as Border VTEP, Route Reflectors, Edge VTEP, or Stub VTEP to the overlay fabric.

Fabrics requires HPE Aruba Networking Central Advanced Switch Licenses for all switches. A notification message is displayed indicating that a valid advanced license is required for this feature. You can acknowledge the message by clicking the close button. The notification message is user-specific. Once the message is closed, it will not be displayed for the same user, but for other users, the message will be displayed. The notification message will re-appear after 90 days and the user will have to acknowledge it every 90 days. For more information, refer to the *AOS-CX switches Feature Pack Ordering Guide* and EULA.



Fabric Design

NetConductor provides two types of Fabric Design options **Routed-Access** and **Scaled-Access Design**.

- The **Routed-Access** Fabric Design offers:
 - Layer 3 Routed Access network.
 - Eliminates L2 Spanning Tree complexity at the edge.
 - Small to Medium deployment sizes.
 - Authentication and Group-based Policy at the Access (Edge) layer.
 - Supports all fabric personas supported in NetConductor except Extended Edge.
- The **Scaled-Access Design** Fabric Design offers:
 - Layer 3 Routed Access till Access-Aggregation layer.
 - Routed with L2 Multi-Chassis between Access-Aggregation and Access layers.
 - Orchestrated Static VXLAN between Access-Aggregation and Access layers.
 - Scales to high number of VRFs and access switches (32 VRFs with upto 32 Access-aggregation and 1000 extended-edge nodes).
 - Scales to extremely large deployment sizes.
 - Authentication and Group-based Policy at the Access (extended-edge) layer.
 - Supports all fabric personas supported in NetConductor.

Fabric Personas

This section describes the different fabric personas and their functions in the HPE Aruba Networking Central BGP EVPN overlay fabric workflow.

- **Border VTEP**—Refers to a Gateway or an external facing device from the overlay fabric to the external network. Border VTEP participates in the Layer 3 VNI or VRF but does not have overlay VLANs or SVIs configured on this device.
- **Route Reflector**—Refers to a concept that is specific to iBGP that uses to optimize route propagation. Route reflector reduces the configuration required on all the devices and optimizes the way BGP sessions establishes. In a typical scenario, every edge device in the network peers with

every other edge device. To avoid this, a route reflector is introduced where all the devices in the overlay fabric peer with the route reflector. Optimizing the configuration and route distribution. Core Switches are generally used as route reflectors in an overlay fabric.

- **Edge VTEP**—Refers to a traditional access Layer where the clients are onboarded. It is also the Layer 3 Gateway for all the clients that are associated with it. Edge VTEP typically has the SVI interface, and it is the entry point of all the wired clients onto the BGP EVPN overlay fabric.
- **Stub VTEP**—Supports static VXLAN and EVPN VXLAN. Stub VTEP is used to establish static VXLAN tunnel to the Gateway and EVPN-VXLAN towards the overlay fabric. Stub VTEP helps to carry role information between EVPN-VXLAN overlay fabric and devices like Gateway that only support Static VXLAN.
- **Extended Edge**—Supports VXLAN extended to 6000 series platforms, including the HPE Aruba Networking CX-6200 switches supporting Micro-segmentation and Group-based Policy. Aruba CX 6200 Switch performs authentication, authorization, VLAN, Role assignment and Group-based Policy enforcement. The switch connects to Edge or Stub node via trunk or LAG (MC-LAG) interface for underlay reachability. The Upstream Edge (or) Stub node provisioned with overlay network segment anycast-gateway registers client to BGP EVPN control plane and relays Role information between static VXLAN (from 6200) and dynamic VXLAN tunnels (rest of the Fabric).

For more information, see the topics listed below:

- [Setting Up an Overlay Fabric](#)
- [Modifying an Overlay Network Fabric](#)
- [Re-syncing the Fabric](#)
- [Deleting an Overlay Network Fabric](#)
- [Viewing the Fabric Configuration Status](#)

Setting Up an Overlay Fabric

The distributed overlay fabric is a group of AOS-CX Switches that are part of the BGP-EVPN VXLAN overlay. The overlay fabric is created by configuring VXLAN tunnels between Stub and Edge Switches. The AOS-10 Gateways also participate in the overlay fabric, using static VXLAN tunnels to the Stub VTEP. HPE Aruba Networking Central allows you to configure the overlay fabric on top of a virtual network using AOS-CX Switches. You can assign a list of fabric personas such as Border VTEP, Route Reflectors, Edge VTEP, or Stub VTEP to the overlay fabric.

-
- Configuring the overlay fabric workflow is applicable only to AOS-CX Switches.
 - For more information, refer to the *AOS-CX switches Feature Pack Ordering Guide* and EULA.
 - The overlay fabric workflow requires HPE Aruba Networking Central Advanced licenses. A notification message is displayed indicating a valid advanced license is required for this feature. You can acknowledge the message by clicking the close button. The notification message is user-specific. Once the message is closed, it will not be displayed for the same user, but for other users, the message will be displayed. For more information, refer to the *AOS-CX switches Feature Pack Ordering Guide* and EULA.
-



The following workflow describes the steps to set up the overlay fabric for the AOS-CX Switches.

1. [Step 1: Create the Overlay Fabric Name](#)
2. [Step 2: Assign Devices and Personas](#)
3. [Step 3: Add the Overlay Network](#)
4. [Step 4: Creating Static Tunnels from Stub VTEPs to a Gateway Cluster](#)
5. [Step 5: Creating Static Tunnels from Stub VTEP to an Extended Edge Switch](#)
6. [Step 6: View Fabric Summary](#)

Step 1: Create the Overlay Fabric Name

To configure an overlay fabric for the AOS-CX Switch, complete the following steps:

1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. In the **Fabrics** table, click the + icon.
The **Create a New Fabric** page is displayed.

Figure 8 Create a New Fabric Page

4. In the **Fabric Name** field, enter a name.



The value in the **BGP AS Number** field is auto-generated. BGP AS Number can be modified, if necessary.

5. In the **BGP AS Number** field, enter the autonomous number for the fabric BGP.
6. Enable the **VLAN Client Presence Detect** to disable the VNI mapped VLANs if there are no authenticated clients on the VLAN, or if the VLAN has no statically configured ports and those ports are up. When not enabled, VNI mapped VLANs are always up. This is an optional step.

7. Enable the **FIB Optimization**, to optimize unused host routes from Forwarding Information Base on the device.
 - **Ageout Time (seconds)** -Enter the time between 60 to 3600 seconds. The default time is set at 90 seconds.
 - **Nexthops to Exclude** - Enter the nexthops addresses you want to exclude from FIB optimization.



The **FIB Optimization** is supported only on switch series 6300, 6400, and 8360 when the switch persona is Edge.

8. Click **Next**. The **Devices** table is displayed with a list of all the Switches in the network.
9. In the **Host Name** column, in the **Devices** table, select the devices that you want to assign to the Switches on the overlay network. The Assign selected device to window is displayed with a list of Switch personas.

Figure 9 Create a New Fabric Page

Name	Firmware Version	MAC Address	Serial Number	Persona
ACC1-6300-VSF	10.12.1012	8c85c1-635c40	SG12KMY04S	EDGE
ACC1-EE-6200-1	10.12.1000	9460d5-d562c0	VN2ZLBB058	
ACC1-EE-6200-2	10.12.1000	9460d5-d57000	VN2ZLBC011	
ACC2-6300	10.12.1012	ec6794-670580	SG37KMYDMD	EDGE
ACC2-EE-6200-1	10.12.0006	a40e75-be53c0	SG31KW52JS	
ACC3-6400	10.12.1012	ec0273-d60000	SG1AKM703P	EDGE

Step 2: Assign Devices and Personas

To assign the selected devices to a fabric on the virtual network, complete the following steps:

1. In the **Assign selected device to** window, select the persona to which you want to assign the devices.

Figure 10 Create a New Fabric Page

✓ Name Fabric
 2 Add Devices
 3 Add Overlay Network
 4 Stub Tunnels to Gateway

Devices (14)				
Name	Firmware Version	MAC Address	Serial	Series
ACC1-6300-VSF	10.12.1012	8c85c1-635c40	SG12	
ACC1-EE-6200-1	10.12.1000	9460d5-d562c0	VN2	
ACC1-EE-6200-2	10.12.1000	9460d5-d57000	VN2	
ACC2-6300	10.12.1012	ec6794-670580	SG37	
ACC2-EE-6200-1	10.12.0006	a40e75-be53c0	SG31	
ACC3-6400	10.12.1012	ec0273-d60000	SG1A	

1 item(s) selected

Assign selected device to

Edge

Border

Stub

RR

Extended Edge

Cancel Apply

2. Click **Apply**.
3. Click **Next**. The **Overlay Networks** section is displayed.



Personas under the VSX pair should be the same.

Step 3: Add the Overlay Network



A default overlay network is created automatically. You can create, move or delete an additional overlay network, if required.

To create an additional overlay network, complete the following steps:

1. In the **Overlay Networks** table, click the **+** icon.

Figure 11 *Overlay Networks Table*

Overlay Networks (3)				+ ⋮
Name	VNI	Route Target	Route Distinguisher	
server	20011	65001:20011	20011	
VRF2	10012	65001:10012	10012	
VRF1	10011	65001:10011	10011	

2. In the **Name** field, enter a name for the overlay network.
3. In the **VNI** field, enter a virtual tunnel identifier.



The VNI is auto-generated, but it can be modified if necessary.

4. In the **Route Target** field, enter the route target range.
5. In the **Route Distinguisher** field, enter the number for route distinguisher.



Hover over an overlay network and click on to edit an overlay network. To delete the overlay network click on .

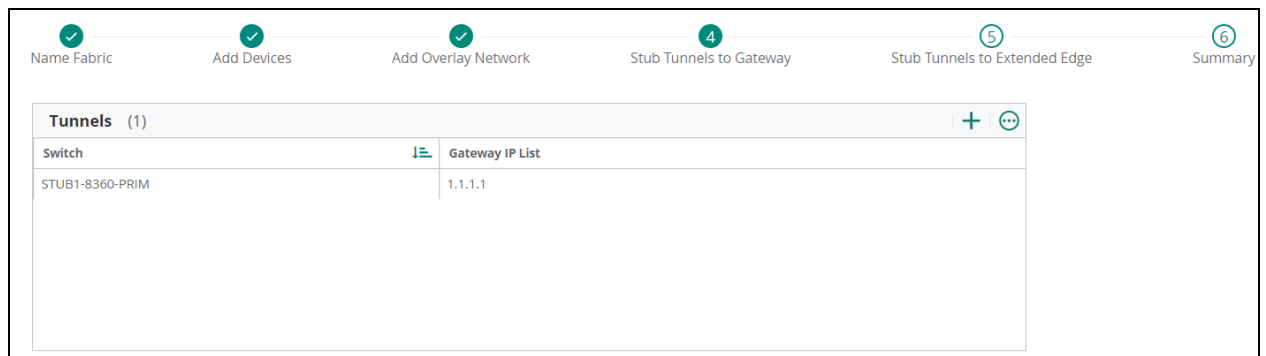
6. Click **Next**. The **Stub Tunnels to Gateway** section is displayed.

Step 4: Creating Static Tunnels from Stub VTEPs to a Gateway Cluster

To create a static tunnel from Stub VTEP to a Gateway cluster, complete the following steps:

1. Click **Skip** if you don't want to add a Stub VTEP to a Gateway cluster.
2. In the **Tunnels** table, click the **+** icon.
The list of Switches and Gateway IP addresses is displayed.

Figure 12 *Stub Tunnels to Gateways Section*



3. Select a Switch from the **Switch** drop-down list.
4. In the **Gateway IP List** field, enter the IP address of the AOS-10 Gateway, which will be the tunnel termination point for the static VXLAN tunnel.



To input multiple values in the **Gateway IP List** parameter, use commas between the IP addresses. For example, 10.1.1.2, 10.1.1.3.

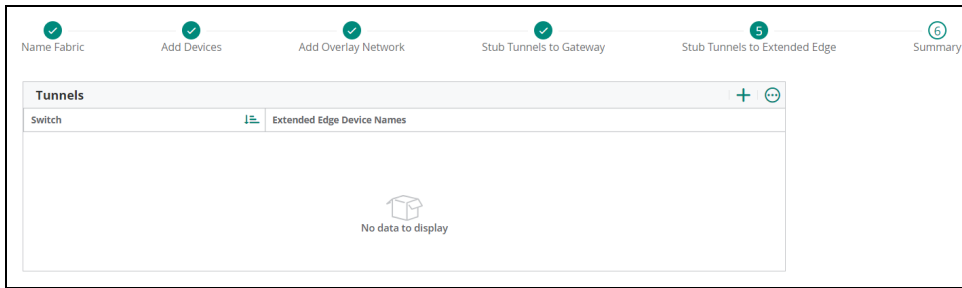
5. Click **Next** to continue.

Step 5: Creating Static Tunnels from Stub VTEP to an Extended Edge Switch

To create a static tunnel from Stub VTEP to an Extended Edge Switch, complete the following steps:

1. Click **Skip** if you don't want to add a Stub VTEP to an Extended Edge Switch
2. In the **Tunnels** table, click the **+** icon.
The list of Switches and Host name of the extended edge switches is displayed.

Figure 13 *Stub Tunnels to Extended Edge*



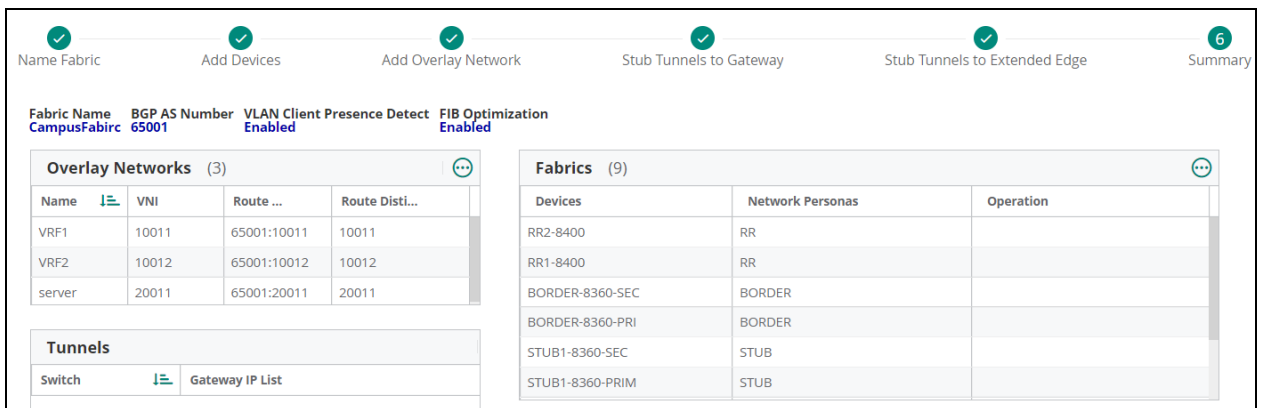
3. Select a Switch from the **Switch** drop-down list.
4. In the **Extended Edge Device Name** field, enter the name of the device which will be the tunnel termination point for the static VXLAN tunnel.
5. Click **Next** to continue.

Step 6: View Fabric Summary

To view the fabric summary, complete the following steps:

1. In the **Stub Tunnels to Gateway** page, click **Next**. The **Summary** section is displayed.


Figure 14 *Create a New Fabric Section*



2. In the **Summary** page, view the fabric summary and check the following:
 - Device personas are assigned correctly.
 - Overlay networks are created properly.
 - Stub tunnels to the Gateways are configured properly,
 - Vlan Client Presence Detect is enabled or not.
 - FIB Optimization is ... enabled or not.
3. Click **Save**.

Modifying an Overlay Network Fabric

To edit or modify a fabric, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. Hover over the tunnel row in the **Fabrics** table and click the  edit icon.
4. Input the tunnel details.
5. Click **Save**.
The overlay network fabric is saved.

Re-syncing the Fabric



The re-sync button is only visible when there are configuration conflicts on the devices.

The re-sync button resolves configuration conflicts within existing overlay fabric devices. You no longer need to run the entire workflow of configuring a fabric to resolve configuration conflicts on devices.


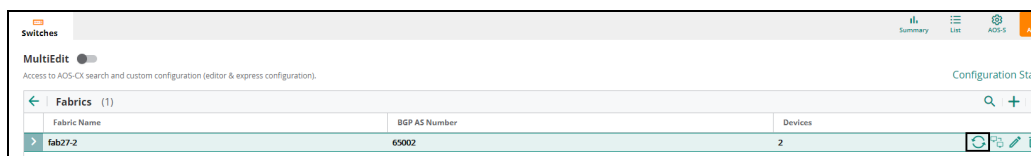



1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. In the **Fabrics** table, hover over the Fabric that you want to re-sync and click the  re-sync icon.

Figure 15 Re-sync Button



Fabric Name	BGP AS Number	Devices	Actions
fab27-2	65002	2	  




To view the status of the fabric, go to the **Device Level Configuration** screen's **Config Status** column.

Deleting an Overlay Network Fabric

To delete a fabric, follow the steps below:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.

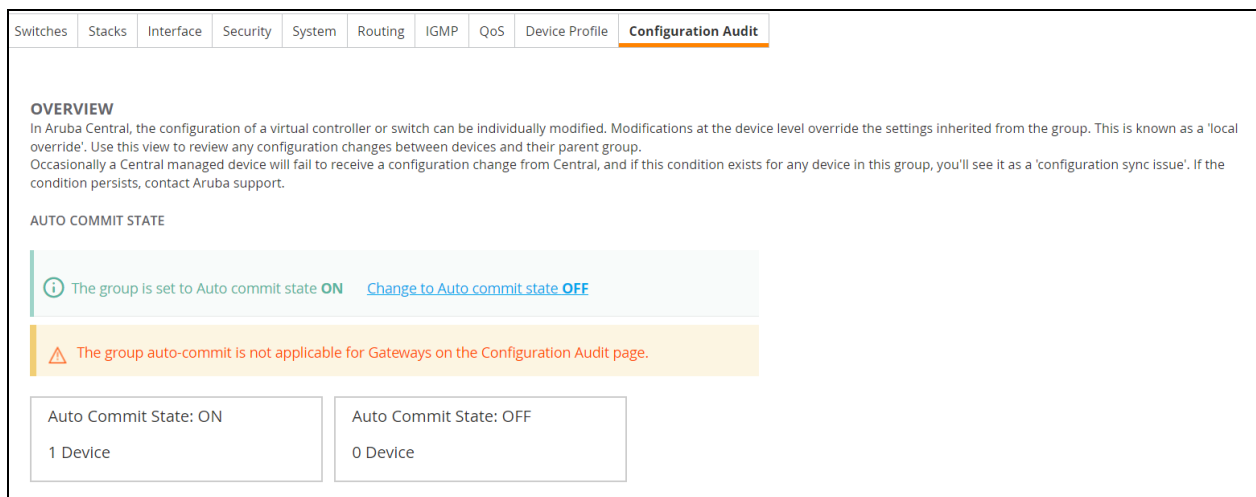
- b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
 3. Hover over the tunnel row in the **Fabrics** table and click the  delete icon.
 4. Click **Yes** to delete the fabric.
The overlay network fabric is deleted.

Viewing the Fabric Configuration Status

To view the fabric configuration status, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**
 - c. Click the **Config** icon.
 - d. Click **Configuration Audit**.
The **Configuration Status** page is displayed.

Figure 16 Configuration Status Page



2. In the **Configuration Status** tab, you can view the configuration status of the overlay network fabric.

Overlay Segment Overview

Overlay segment defines a broadcast domain. Overlay segment can be of two types, Layer 2 segment and Layer 3 segment. The overlay segment maps to the VLAN/VNI by deploying on the overlay network. You can define the VLAN ID, IP Subnet, default Gateway, DHCP relay server, and roles present on this segment. Then define the devices that this segment enables for.

For more information, see the topics listed below:

- [Setting Up a Layer 2 Segment](#)
- [Setting Up a Layer 3 Segment](#)
- [Viewing the Overlay Fabric](#)
- [Modifying the Overlay Segment](#)
- [Deleting the Overlay Segment](#)
- [Viewing the Fabric Configuration Status](#)

Setting Up a Layer 2 Segment

Layer 2 segments supports use cases where the default gateway of the segment is external to the fabric. Example of such use-cases includes, IoT devices connecting to gateway or fabric traffic that terminates at the firewall.


The following workflow describes the steps for configuring Layer 2 segments:


1. [Step 1: Configure the Layer 2 Segment](#)
2. [Step 2: Assign Roles](#)
3. [Step 3: Select Edge and Stub Devices for the Layer 2 Segment](#)
4. [Step 4: View the Layer 2 Segment summary](#)

Step 1: Configure the Layer 2 Segment

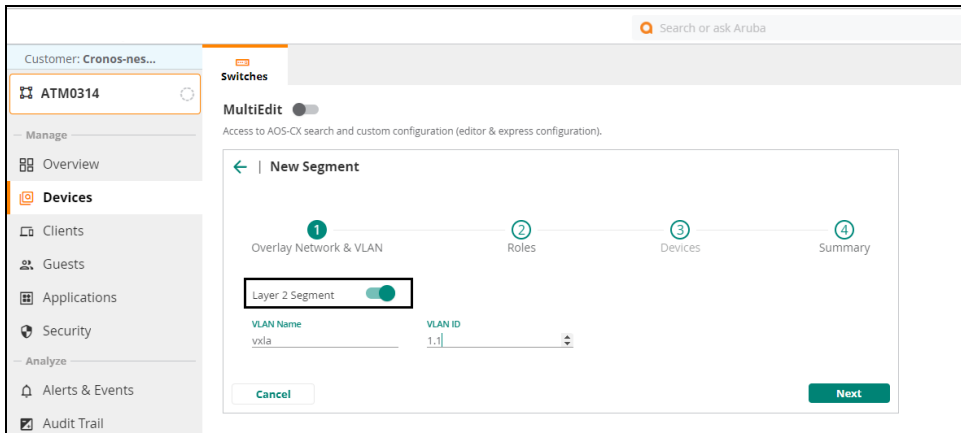
To configure an overlay segment for the AOS-CX Switch, complete the following steps:

1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.

In the **Fabrics** table, select the Fabric on which you want to enable Layer-2 segment and click the  icon.



3. Enable the **Layer 2 Segment** toggle.



4. In the **Overlay Network & VLAN** table, populate the Overlay network and VLAN table with the relevant information.
The following table describes the Layer 2 segment parameters.

Table 9: Overlay Network & VLAN Table Parameter

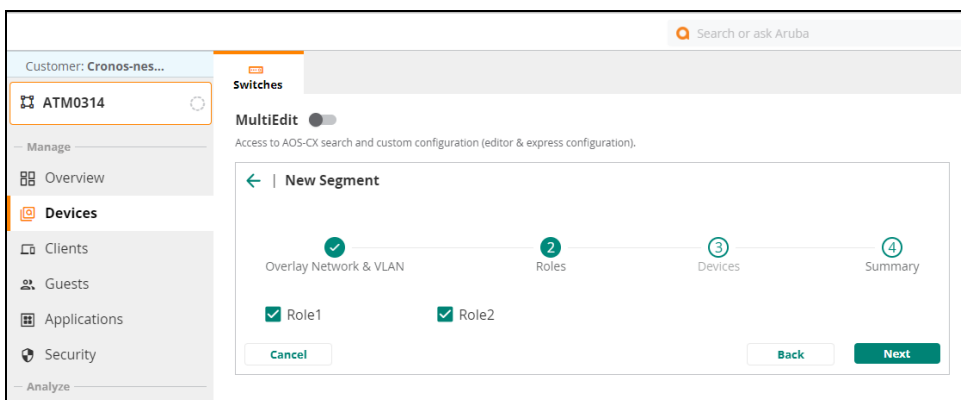
Column	Function
VLAN Name	Option to provide the VLAN name.
VLAN ID	The VLAN ID number.

5. Click **Next**. The **Roles** section is displayed with a list of all the global roles defined on the **Client Roles** page.

Step 2: Assign Roles

This section describes the steps to assign roles to endpoints or users that are part of the Layer 2 segment.

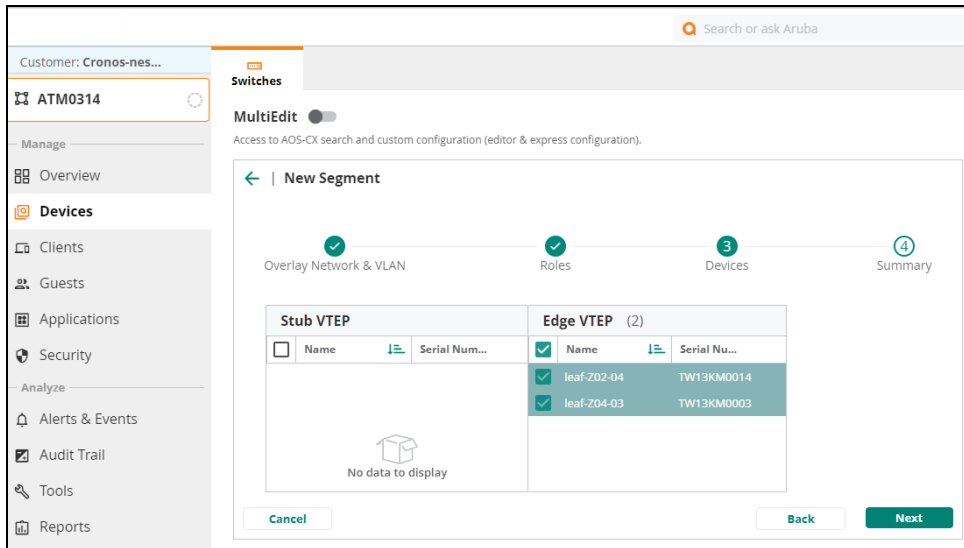
1. Select the roles for assigning the Layer 2 overlay network segment.
2. Click **Next**. The **Devices** section is displayed.



Step 3: Select Edge and Stub Devices for the Layer 2 Segment

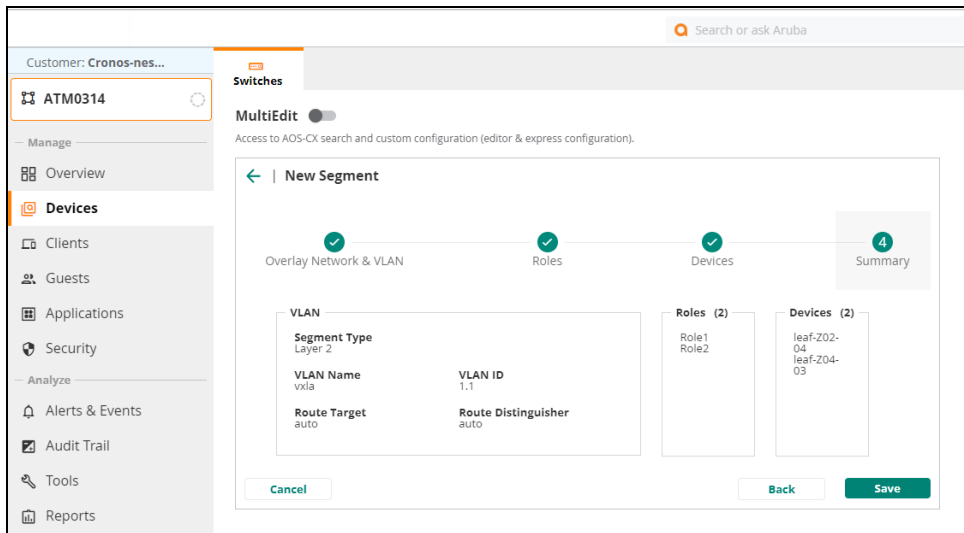
In the **Devices** section, select the devices to apply on the segment.

1. Select individual devices or select all the devices and assign them to the overlay fabric.
2. Click **Next**. The **Summary** section is displayed.



Step 4: View the Layer 2 Segment summary

1. The **Summary** section is displayed where you can view the summary of your configuration.



2. Click **Save** to complete the configuration.

Setting Up a Layer 3 Segment

The following workflow describes the steps for configuring overlay segments for the existing overlay fabrics in HPE Aruba Networking Central.

1. [Step 1: Configure the Overlay Segment](#)
2. [Step 2: Assign Roles](#)
3. [Step 3: Select Edge and Stub Devices for the Overlay Segments](#)
4. [Step 4: View the Overlay Segment Summary](#)

Step 1: Configure the Overlay Segment

To configure an overlay segment for the AOS-CX Switch, complete the following steps:


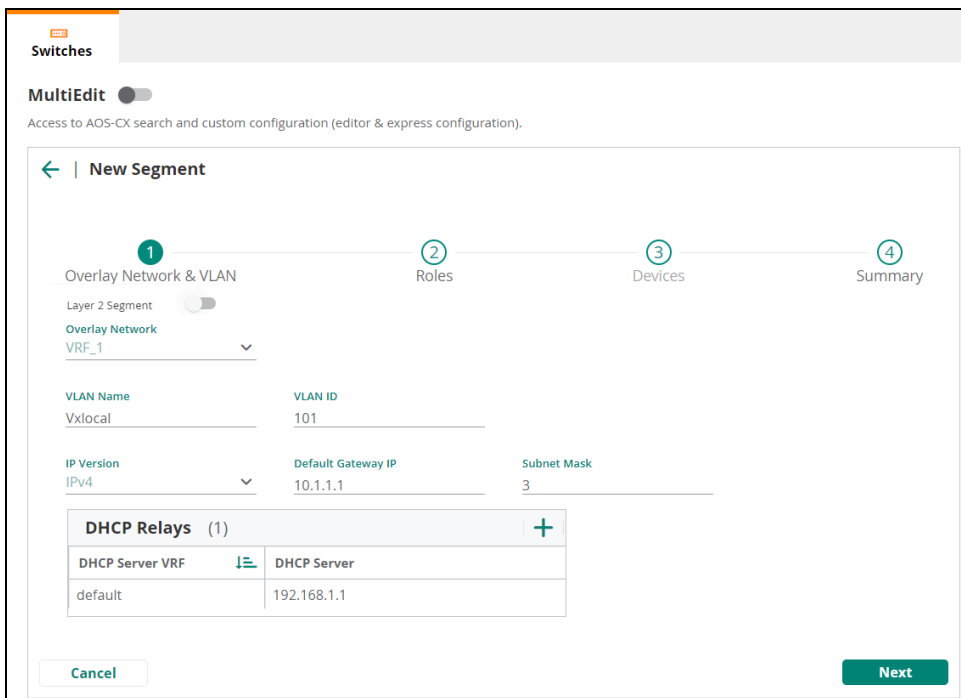
1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. In the **Fabrics** table, to select the row for the fabric, click the  icon.
The **Overlay Network & VLAN** section is displayed.

Figure 17 *Overlay Network & VLAN Page*



The following table describes the columns in the **Devices** table.

Table 10: *Overlay Network & VLAN Table Parameters*

Column	Function
Layer 2 Segment	The toggle is enabled when configuring Layer 2 segment only.
Overlay Network	List of overlay networks. You can select the overlay network that you want to create the service segment for, from the list.
VLAN Name	Option to provide the VLAN name.

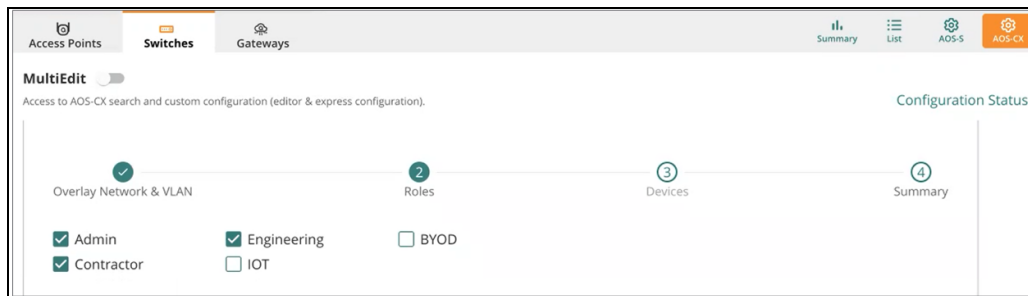
Column	Function
VLAN ID	The VLAN ID number.
IP Version	Option to select the IP version. <ul style="list-style-type: none"> ▪ IPv4 ▪ IPv6 ▪ Dual Stack
Default Gateway IP	Option to define the default Gateway IP for the overlay network.
Subnet Mask	Option to provide the subnet mask for the IP address.
DHCP Server VRF	The default DHCP server VRF value or you can add manually. Note: You can manually add multiple DHCP server relays.
DHCP Server	The IP address of the DHCP relay server.

4. In the **Overlay Network & VLAN** table, provide all the details.
5. Click **Next**. The **Roles** section is displayed with a list of all the global roles defined on the **Client Roles** page.

Step 2: Assign Roles

This section indicates the roles that can be assigned to endpoints or users that are part of this segment. To assign roles that you want to allow to participate in the overlay network segment, complete the following steps:

Figure 18 Roles Page

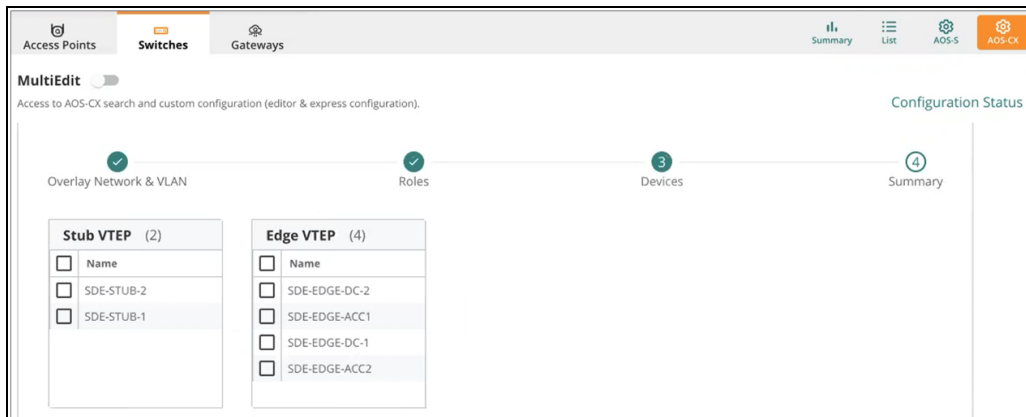


1. Select the roles to participate in the overlay network segment.
2. Click **Next**. The **Devices** section is displayed.

Step 3: Select Edge and Stub Devices for the Overlay Segments

In the **Devices** section, select the devices which you want to apply on the segment. You can either select individual devices or select all the devices and attach to the overlay fabric.

Figure 19 Devices Page

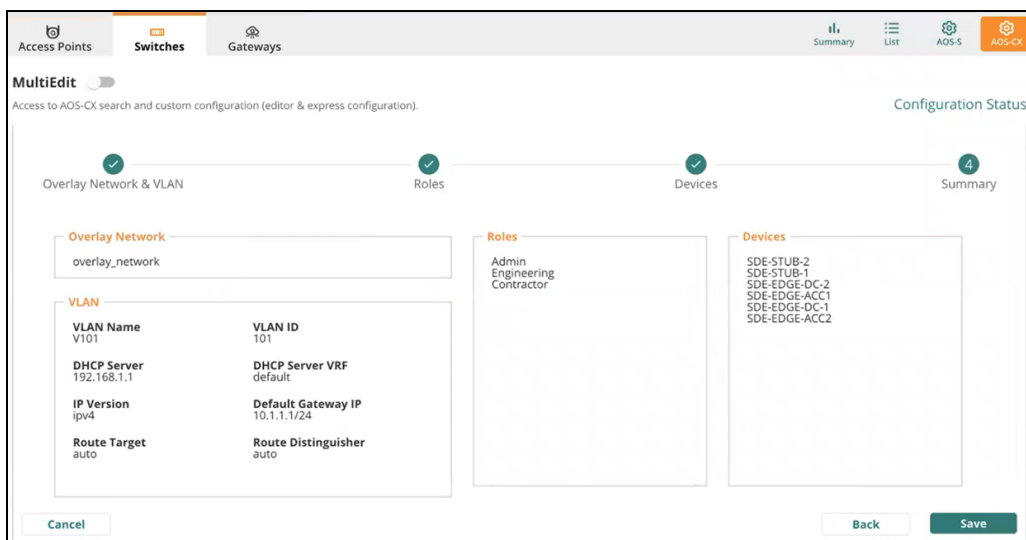


Step 4: View the Overlay Segment Summary

To view the overlay segment summary, complete the following steps:

In the **Devices** section, click **Next**. The **Summary** section is displayed where you can view the summary of your configuration.


Figure 20 Summary Page



Viewing the Overlay Fabric

To view the newly created overlay fabric or view an existing fabric, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices** > **Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing** > **Fabrics**.
3. In the overlay segment row, click the arrow icon to expand the fabric.

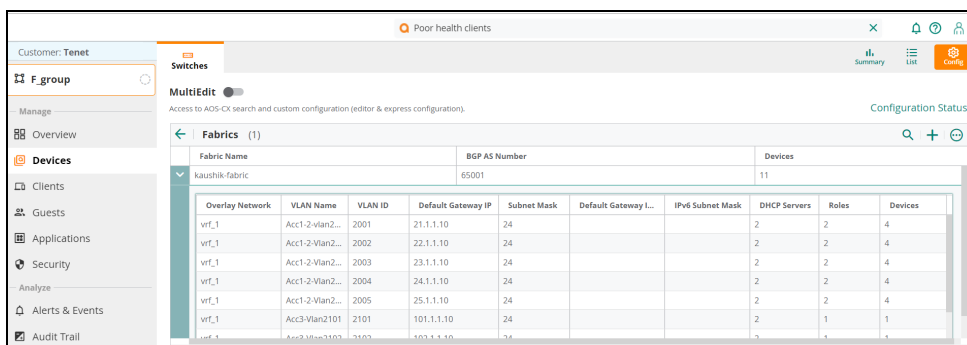


4. View the overlay segment.

The following table describes the columns in the **Fabrics** table.


Table 11: Overlay Segment Table Parameters

Column	Function
Overlay Network	List of overlay networks.
VLAN Name	The VLAN name.
VLAN ID	The VLAN ID number.
IP Version	The IP version. <ul style="list-style-type: none"> IPv4 IPv6 Dual Stack
Default Gateway IP	The default Gateway IP for the overlay network.
Subnet Mask	The subnet mask for the IP address.
DHCP Server VRF	The default DHCP server VRF value or you can add manually. Note: You can manually add multiple DHCP server relays.
DHCP Server	The IP address of the DHCP relay server.
Roles	The number of roles assigned to this network. NOTE: When you delete a role in the Clients Role page. The number of updated role is not reflected in this column, unless you edit the overlay segment.
Devices	The number of devices assigned to this network.




Modifying the Overlay Segment

To edit or modify an overlay fabric, complete the following steps:

1. In the WebUI, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. In the overlay segment row, click the  edit icon and provide the required details.
4. Click **Save**.
The overlay segment is saved.

Deleting the Overlay Segment

To delete an overlay segment, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - To select a group in the filter:
 - a. Set the filter to a group.
The dashboard context for the group is displayed.
 - b. Under **Manage**, click **Devices > Switches**.
 - c. To view the AOS-CX Switch configuration dashboard, click the **AOS-CX** or **Config** icon.
2. Click **Routing > Fabrics**.
3. In the overlay segment row, click the  delete icon.
The overlay segment is deleted.

Muti-Fabric EVPN Overview

The Muti-Fabric EVPN in NetConductor allows you to build independent fabric for administrative construct allowing to scale for large, geo-dispersed deployments. Independent fabric are interconnected together across metro Ethernet or IP Network. Muti-Fabric EVPN allows to extend segmentation (VLANs, VRFs, Roles) with consistent Group-based Policy. Fabric inter-connectivity requires network to support higher MTU (≥ 1550 bytes) to accommodate VXLAN encapsulation between sites. Maximum of 32 Fabrics can be interconnected horizontally in a Multi-LayerFabric Design.

For more information on Muti-Layer Fabric Overview see the following topics:

- [Setting up the Muti-Fabric EVPN](#)
- [Modifying a Multi-Fabric EVPN](#)
- [Re-syncing the Multi-Fabric EVPN](#)

Setting up the Muti-Fabric EVPN

NetConductor supports Multi-Fabric by 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. Multiple data center fabrics and locations can be combined into a single

overlay topology. HPE Aruba Networking Central allows you to configure the Multi-Fabric on top of a virtual network using AOS-CX Switches. You can assign a list of fabric personas such as Border VTEP, Extended-dge, Route Reflectors, Edge VTEP, or Stub VTEP to the Multi-Fabric.

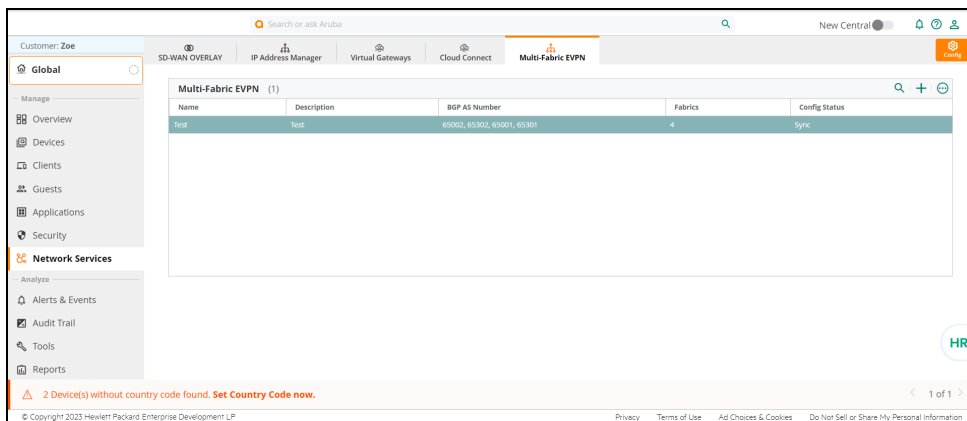
The following workflow describes the steps to set up the Multi-Fabric for the AOS-CX Switches.

- [Navigating to Multi-Fabric EVPN](#)
- [Adding the Fabrics Name to the Multi-Layer Network](#)
- [Selecting Border Leader for Multi-Fabric EVPN](#)
- [Viewing Overlay Networks](#)
- [Selecting the Stretched Segments](#)
- [Viewing Multi-Fabric EVPN Summary](#)

Navigating to Multi-Fabric EVPN

The following steps describes the navigation to the Multi-Fabric EVPN page:

1. In the WebUI, set the filter to **Global**.
2. Go to **Manage > Network Services > Multi-Fabric EVPN**.
3. Click on the plus **+** icon to add a new Multi-Fabric EVPN.
The guided setup page for Multi-Fabric EVPN is displayed.



Adding the Fabrics Name to the Multi-Layer Network

The following steps describes how to add name to the Multi-Fabric EVPN:

1. Enter the **Name** of the Multi-Fabric EVPN.
2. Enter the **Description** of the Multi-Fabric EVPN.

- Click **Next** to display the Fabrics section.

Selecting Border Leader for Multi-Fabric EVPN

Borders can only AOS-CX Switch Series 10000, 8325, 8360, 9300, and 6400.

The minimum firmware version is 10.11.000 or later for Border-leader/Border for Switch Series 6400, 8360 and 9300

The minimum firmware version is 10.12.000 or later for Border-leader/Border for Switch Series 10000 and 8325

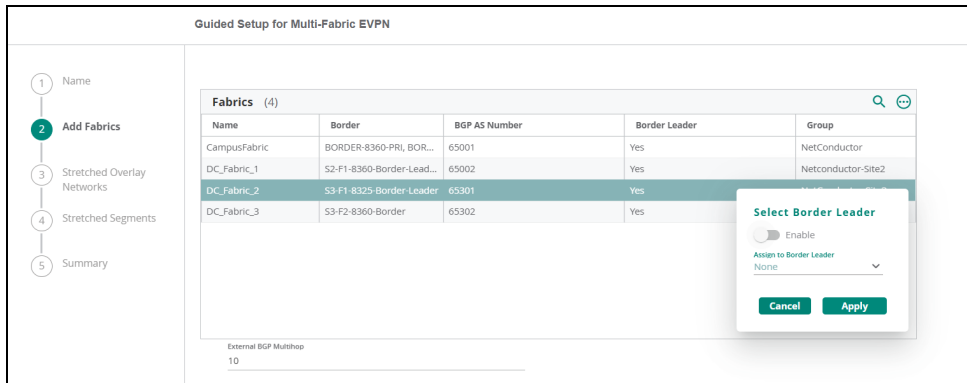


The following steps describes how to add a border leader:

- The Fabrics section lists all the stretched overlay networks information, the Fabrics columns describes:
 - Name**—Displays the name of fabric available in the network.
 - Border**—Displays the name of the border this fabric belongs to.
 - BGP As Number**—Displays the identifying number of the fabric connected to the BGP.
 - Border Leader**—Displays whether this fabric is a border leader or not.
 - Group**—Displays the name of the group this fabric belongs from.
- Click the Fabric of choice to populate the **Select Border Leader** pop-up and click on the **Enable** toggle to make it a border leader. Click **Apply** to add this fabric as border leader.

Name	Border	BGP AS Number	Border Leader	Group
CampusFabric	BORDER-8360-PR1, BOR...	65001	Yes	NetConductor
DC_Fabric_1	S2-F1-8360-Border-Lead...	65002	Yes	Netconductor-Site2
DC_Fabric_2	S3-F1-8325-Border-Leader	65301	Yes	
DC_Fabric_3	S3-F2-8360-Border	65302	Yes	

- Click on the remaining Fabric(s) to disable the **Enable** toggle. From the drop-down select the **Assign to Border Leader**.



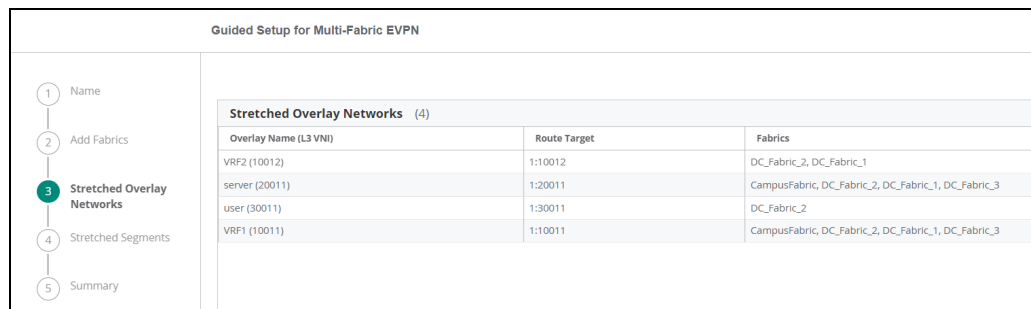
4. Enter the **External BGP Multihop** number or use the up and down arrow icons to select the number of choice.
5. Click **Next** to continue.

Viewing Overlay Networks

It displays the list of available Overlay Networks (L3VNI) that are common across all the selected fabrics.




Clicking on the edit icon enables you to edit the **Route Target** details.



Selecting the Stretched Segments

It displays the list of available Overlay Segments (L2VNI) that are common across all the selected fabrics. This is an optional step, if you want to add stretched segments to your fabric follow the steps below:

1. Select the check-box against the **VLAN Name** you want to add as stretched segments.
2. Click **Next** to continue.



Viewing Multi-Fabric EVPN Summary

1. In the **Summary** section, view the fabric summary and check the following sections:
 - Details
 - Fabrics
 - Configuration
 - Stretched Overlay Networks
 - Stretched Segments
2. Click **Finish** to save the Multi-Fabric EVPN.

The screenshot displays the 'Guided Setup for Multi-Fabric EVPN' interface. On the left, a vertical navigation pane shows five steps: 1. Name, 2. Add Fabrics, 3. Stretched Overlay Networks, 4. Stretched Segments, and 5. Summary (highlighted). The main content area is divided into several sections:

- Details:** A form with 'Name' (Test) and 'Description' (Test).
- Fabrics (4):** A table listing fabric configurations.
- Stretched Overlay Networks (4):** A table listing overlay networks and their associated fabrics.
- Stretched Segments (2):** A table listing VLANs and their associated overlay networks.
- Configuration:** A section for 'External BGP Multihop' with a value of 10.

Name	Borders	BGP AS Number	Border Leader
CampusFabric	BORDER-8360-...	65001	BORDER-8360-PR1, B...
DC_Fabric_1	S2-F1-8360-Bor...	65002	S2-F1-8360-Border-Le...
DC_Fabric_2	S3-F1-8325-Bor...	65301	S3-F1-8325-Border-Le...
DC_Fabric_3	S3-F2-8360-Bor...	65302	S3-F2-8360-Border

Overlay Name (L3 VNI)	Route Target	Fabrics
VRF1 (10011)	1:10011	CampusFabric, DC...
VRF2 (10012)	1:10012	DC_Fabric_2, DC_F...
server (20011)	1:20011	CampusFabric, DC...
user (30011)	1:30011	DC_Fabric_2

VLAN Name	L2 VNI	Route Target	Overlay Name (L3 ...)	Fabrics
vlan81	81	1:81	VRF1 (10011)	DC_Fabric_...
vlan84	84	1:84	VRF2 (10012)	DC_Fabric_2

Modifying a Multi-Fabric EVPN

To edit a fabric, complete the following steps:

1. In the HPE Aruba Networking Central app, set the filter to **Global**.
2. Go to **Manage > Network Services > Multi-Fabric EVPN**.
3. Hover over the tunnel row in the **Fabrics** table and click the edit icon.
4. Change the details as required.
5. Click **Save**.
The multi-fabric EVPN is saved.

To delete a fabric, complete the following steps:

1. Hover over the tunnel row in the **Fabrics** table and click the delete icon.
2. Click **Delete**.

Re-syncing the Multi-Fabric EVPN

The re-sync button is only visible when there are configuration conflicts on the devices.



NOTE

The re-sync button resolves configuration conflicts within existing overlay fabric devices. You no longer need to run the entire workflow of configuring a fabric to resolve configuration conflicts on devices.



1. In the HPE Aruba Networking Central app, set the filter to **Global**.
2. Go to **Manage > Network Services > Multi-Fabric EVPN**.
3. In the **Fabrics** table, hover over the Fabric that you want to re-sync and click the  re-sync icon.

Figure 21 *Re-sync Button*



To view the status of the fabric, go to the **Device Level Configuration** screen's **Config Status** column.



Static VXLAN Tunnels on AOS-10 Gateways

HPE Aruba Networking Central allows you to configure a static VXLAN tunnel on Gateways.

Prerequisites

The following prerequisites must be fulfilled to configure static VXLAN tunnels:

- AOS-10 Gateways are configured as Mobility Gateways.
- Gateway Cluster is enabled on HPE Aruba Networking Central.
- The **Role-to-Role Policy Enforcement** toggle must be defined on the **Client Roles** page to enable the global role policies.
- The **Use a switch fabric for a role propagation?** option must be set to **Yes** on the **Client Roles** page to enforce the global role policies on the Switch fabric.

For more information, see the topics listed below:

- [Configuring a Static VXLAN Tunnel on Gateways](#)
- [Mapping VLAN and VNI to a VXLAN Tunnel](#)
- [Viewing the Static VXLAN tunnels on Gateway Clusters](#)

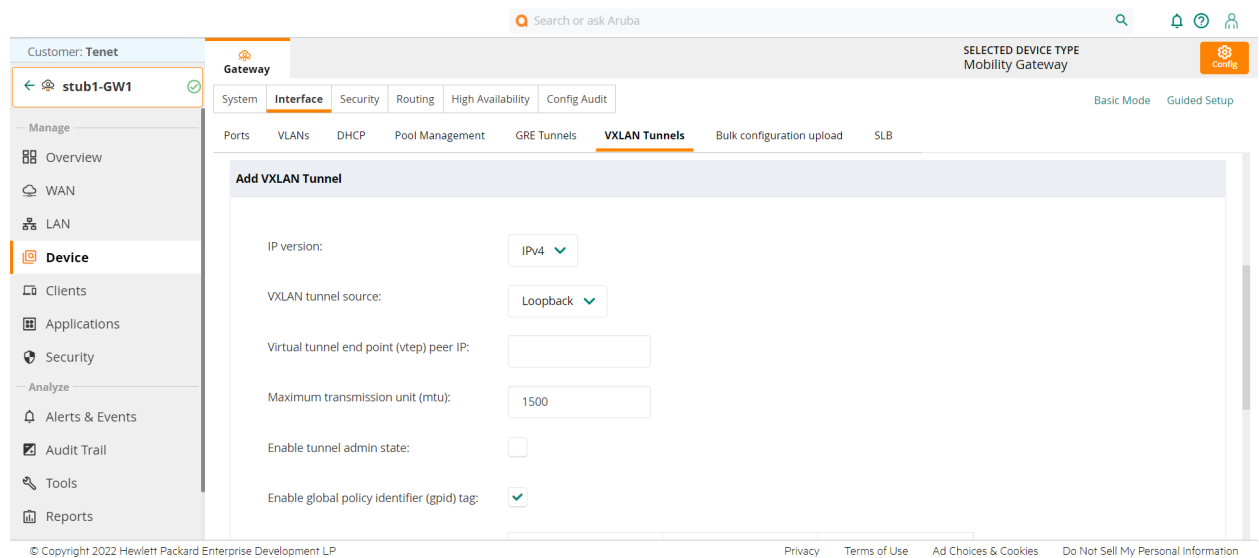
Configuring a Static VXLAN Tunnel on Gateways

To configure a static VXLAN tunnel on a Gateway or VPNC group, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - a. Set the filter to a group containing at least one Branch Gateway.
 - b. Under **Manage**, click **Devices > Gateways**.
 - c. Click **Config**.
2. If you are in the **Basic Mode**, click **Advanced Mode** to access the advanced configuration options.

3. Click **Interface > VXLAN Tunnels**. The **VXLAN Tunnels** page is displayed.
4. In the **VXLAN Tunnels** pane, click the **+** icon.
The **Add VXLAN Tunnel** pane is displayed.

Figure 22 Add VXLAN Tunnel Page



The following table describes the parameters in the **Add VXLAN Tunnel** table:

Table 12: Add VXLAN Tunnel Parameters

Parameter	Description
IP Version	Select one of the IP versions for the HPE Aruba Networking Gateways, to select the system IP address.
VXLAN tunnel source	Select one of the following VXLAN tunnel sources for the HPE Aruba Networking Gateways, to select the system IP address: <ul style="list-style-type: none"> ▪ System IP—Enables the system IP as your tunnel source. ▪ Loopback—Selecting this option automatically fetches the cluster loopback and the cluster system IPs as your tunnel source. ▪ VLAN—Selecting this option allows you to choose an L3 VLAN that is configured on the Gateway as your tunnel source. ▪ Static IP—Selecting this option allows you to input a source IP address.
Source IP address	This field allows you to add the source IP address. This field appears when the Static IP option is selected in the VXLAN Tunnel Source drop-down.
VLAN interface	This field displays the auto-populated list of L3 VLANs.
Virtual Tunnel End Point (VTEP) peer IP	This is the VTEP peer IP address. For a standalone Switch, this can be the loopback 1 IP address of the Switch. If the peer is a VSX pair, use the logical VTEP IP address of the VSX pair.
Maximum transmission unit (MTU)	Provides the MTU setting for the VXLAN. The value should be between 1024 and 9198 to avoid fragmentation within the fabric.
	NOTE: Enable the Jumbo Frames Processing toggle at Devices > Gateways > Security > Firewall and check the Jumbo MTU box at

Parameter	Description
	Devices > Gateways > Interface > Ports.
Enable tunnel admin state	Select this option to enable the admin state of the VXLAN interface.
Enable Global Policy Identifier (GPID) tag	Select this option to enable the GPID tag. This is the GPID that is propagated across the network.

5. In the **Add VXLAN Tunnel** pane, configure the VXLAN parameters as described in the [Add VXLAN Tunnel Parameters](#) table based on your network requirements.

Mapping VLAN and VNI to a VXLAN Tunnel

To map the VLAN and VNI to the VXLAN tunnel, complete the following steps:

1. In the HPE Aruba Networking Central app, complete the following steps:
 - a. Set the filter to a group containing at least one Branch Gateway.
 - b. Under **Manage**, click **Devices > Gateways**.
 - c. Click **Config**.
2. If you are in the **Basic Mode**, click **Advanced Mode** to access the advanced configuration options.
3. Click **Interface > VXLAN Tunnels**. The **VXLAN Tunnels** page is displayed.
4. In the **VLAN/VNI mappings** page, click the + icon.
The **Add VNI** table is displayed.

Figure 23 Add VNI table

The following table describes the parameters in the **Add VNI** table:

Table 13: Add VNI Tab Parameters

Parameter	Description
VLAN ID	Select a VLAN ID from the list, to on-board the VLAN to the VXLAN tunnel. Example: 101, 102, and 103.
Virtual network identifier	VNI is auto-generated. For auto-generation of VNI, 1000 is added to the VLAN ID. Example: If the VLAN ID is 101, the VNI is auto-generated as 10101.

5. In the drop-down list, select the **VLAN ID** to on-board the VLAN to the VXLAN tunnel.
6. Click **OK**.
7. Click **Save Settings**.

Viewing the Static VXLAN tunnels on Gateway Clusters

To view the static VXLAN tunnels on gateways complete the following steps:

1. In the WebUI, complete the following steps:
 - a. Set the filter to **Global**.
 - b. Under **Manage**, click **Devices > Gateways**.
 - c. A list of Gateways is displayed.
2. Click **Clusters**.
 - a. The list of **Gateway Clusters** is displayed.
 - b. Click on a **Gateway Clusters**
3. Under **Overview**, click **Tunnels**
4. A list of **Tunneled Down** and **Tunnel** destination devices are displayed for cluster gateways.

Table 14: *Tunnel Down Parameters*

Parameter	Description
Destination Device	Displays the name of the location of the tunnel down device.
IP Address	Displays the IP address of the destination device.
Last Connected	Displays when this device was connected to the gateway.
Reason	Displays the description for connecting the destination device to the gateway.
Gateway Name	Displays the name of the gateway.

Table 15: *Tunnels Parameters*

Parameter	Description
Destination Device	Displays the name of the location of the tunnel device.
IP Address	Displays the IP address of the destination device.
Encapsulation	Displays whether this is a IPsec or VxLAN encapsulation.
Status	Displays whether this tunnel is Up or Down .
SSID	Displays the name given to a WLAN and is used by the client to access a WLAN network.
VNI	Displays the VNI of the device.

Customer: HPE Yoda

Search or ask Aruba

Summary Gateways **Tunnels**

← auto_gwcluster_s... ✓

Manage


Overview

Analyze

Alerts & Events

Audit Trail

Tunnel Down

Destination Device	IP Address	Last Connected	Reason	Gateway Name
 No data to display right now				

GATEWAYS | 7024-116.15

Tunnels (1)

Destination Device	IP Address	Encapsulation	Status	SSID	VNI (VxLAN)
> e8:26:89:cc:d9:14	10.16.122.117	IPSec	• Up	--	--

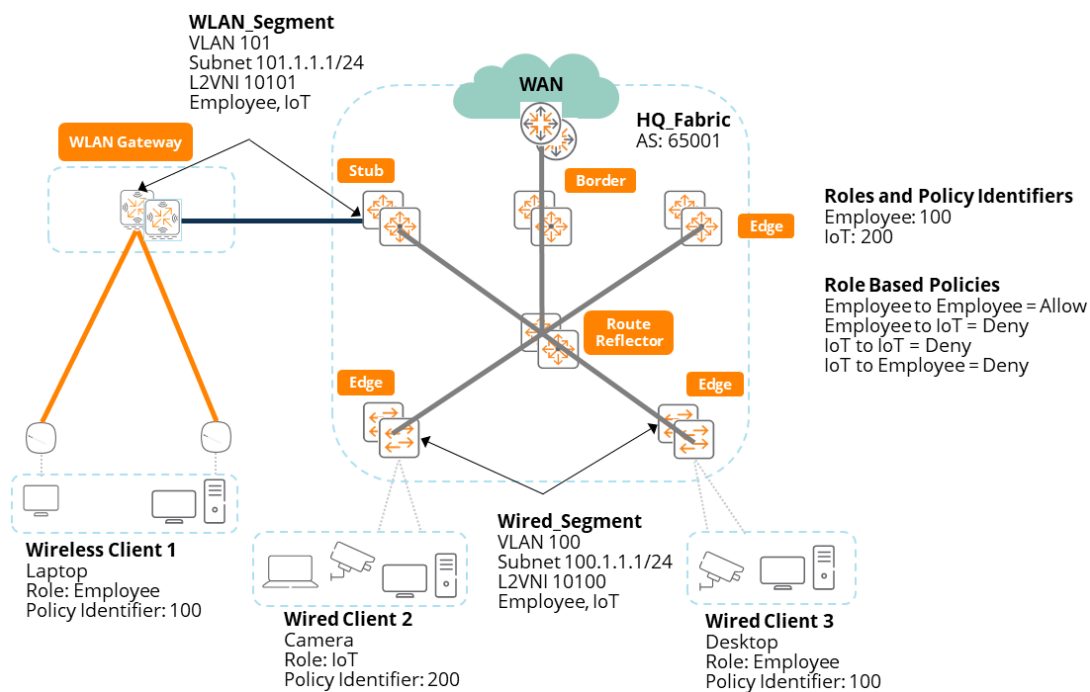
This section describes example use cases for the following deployments:

1. [Distributed Campus-wide Fabric](#)
2. [Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch](#)
3. [Centralized Multi-site Fabric with Third-Party SD-WAN](#)

Distributed Campus-wide Fabric

This section describes an example use case for the Distributed Campus-wide Fabric deployment using HPE Aruba Networking Central.

Figure 24 *Distributed Campus-wide Fabric Use Case*



In the use case, the goal is to enable role-based micro-segmentation across all wired and wireless clients in an enterprise campus. Roles and role-based policies are orchestrated centrally from the **Client Roles** page on HPE Aruba Networking Central. These policies are enforced on the edge VTEPs in a distributed EVPN fabric for wired clients and on the WLAN gateways for wireless clients. The VXLAN-EVPN fabric is provisioned on the AOS-CX switches using the Fabric Wizard on HPE Aruba Networking Central.

As depicted in the above diagram, wired and wireless clients that are assigned the employee role are allowed to communicate with other clients with the Employee role, but not to the clients with the IoT role. Conversely, clients with the IoT role are denied access to clients of both the IoT and employee role.

Step 1: Creating the Role and Role-based Policy

The roles and role-based policies are defined in the **Client Roles** page on HPE Aruba Networking Central. In the **Client Roles** page, the Employee role is created with a policy Identifier 100 and the IoT role is created with a Policy Identifier of 200.

Policies for these roles are defined next. The permission for the Employee role is assigned to allow source to destination. This allows communication between the clients with the Employee role. By default, all other role-based communication is denied.

After the permissions are assigned, this role and policy definition is configured on all gateways and switches on the network.

For more information about how to create role and role-based policy using HPE Aruba Networking Central, refer to [Global Client Roles](#).

Step 2: Creating the Overlay Fabric

The next step is to provision an EVPN-VXLAN fabric on the AOS-CX switches to enable propagation of roles across the network and enforce role-based policies for wired clients. In the **Create a New Fabric** page, a new fabric with the name HQ_Fabric is created with BGP AS 65001, which is auto-populated by HPE Aruba Networking Central. After the fabric is created, the configurations for the EVPN-VXLAN fabric and static VXLAN tunnels are pushed to the AOS-CX devices.

For more information about how to create fabric using HPE Aruba Networking Central, refer to [Fabric Wizard Overview](#).

Step 3: Creating the Wired and Wireless Segment

The Segment Creation wizard is used to create a wired segment and apply the segment to all the Edge VTEPs in a fabric. A segment is created with the name Wired_Segment, VLAN ID 100, and Default Gateway IP of 100.1.1.1/24. The L2VNI of 10100 is auto-assigned to the segment by HPE Aruba Networking Central. The Employee and IoT role is assigned to the segment and applied to all the wired Edge VTEPs in the fabric.

Similarly, a wireless segment is created and applied to the Stub VTEPs that connect to the WLAN gateways in the network. This segment has a name Wireless_Segment, VLAN ID 101 and Default Gateway IP of 101.1.1.1/24. The Employee and IoT role is also assigned to the segment and applied to all the Stub VTEPs in the fabric.

For more information about how to create wired and wireless segment using HPE Aruba Networking Central, refer to [Fabric Wizard Overview](#).

Step 4: Creating Static VXLAN Tunnel on the Gateway

The final step is to create the static VXLAN tunnel on the WLAN gateway. Using the VXLAN Tunnel UI Configuration, a VXLAN tunnel is created using the System IPs of the cluster as the tunnel source and the VTEP IP of the AOS-CX switches as the tunnel peer. The VLAN ID 101 is added to the tunnel to map to the Wireless Segment created on the Stub VTEPs. The tunnel is enabled and GBP role propagation is also enabled on the tunnel. After the mapping is completed, the configurations for the static VXLAN tunnel are pushed to the WLAN gateway.

For more information about how to create static VXLAN tunnel on the WLAN gateway using HPE Aruba Networking Central, refer to [Static VXLAN Tunnels on AOS-10 Gateways](#).

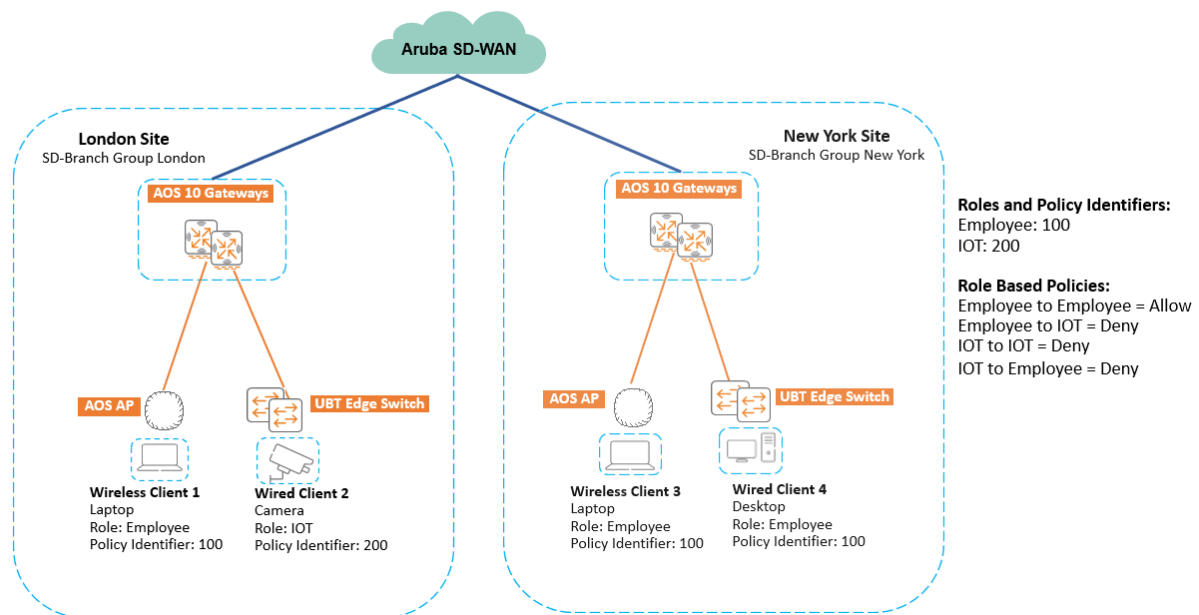
Upon successful configuration, role-based policies are enforced across all clients in the network. In the example above, Wireless Client 1 which is authenticated with the role Employee communicates with the

Wired Client 3 with role Employee but is denied communication with the Wired Client 2 with role IoT. Similarly, the Wired Client 2 cannot communicate with either Wireless Client 1 or Wired Client 3.

Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch

This section describes an example use case for the Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch deployment using HPE Aruba Networking Central.

Figure 25 Centralized Multi-site Fabric with HPE Aruba Networking SD-Branch Use case



The SD-Branch group London maps to the London site and the SD-Branch group New York maps to the New York site. The goal is to enable role-based micro-segmentation across multiple geographic sites connected over an HPE Aruba Networking SD-WAN fabric. Roles and role-based policies are centrally in the **Client Roles** page on HPE Aruba Networking Central. The AOS-10 Gateway is the WLAN gateway for the wireless clients and user-based tunnel Gateway for the wired clients within the site. The role-based policies are enforced on the AOS-10 Gateways for all the clients within the site. Role propagation and role-based policy enforcement is selectively enabled per groups.

As depicted in the above diagram, wired and wireless clients that are assigned the Employee role is allowed to communicate with other clients with the Employee role, but not to the clients with the IoT role. Conversely, clients with the IoT role are denied access to clients of both the IoT and Employee role.

Step 1: Creating Role and Role-based Policy

The roles and role-based policies are defined on the **Client Roles** page on HPE Aruba Networking Central. In the **Client Roles** page, the Employee role is created with a Policy Identifier 100 and the IoT role is created with a Policy Identifier of 200.

Policies for these roles are defined. The permission for the Employee role is assigned to allow source to destination. This allows communication only between the clients with the Employee role. By default, all other role-based communication to the

After the permissions are assigned, the same role and policy definition is configured on all gateways and switches on the network.

For more information about how to create role and role-based policy using HPE Aruba Networking Central, refer to [Configuring the Global Client Roles](#).

Step 2: Selective Enablement of Groups

The next step is to selectively enable the groups that maps to the respective sites. In the **Client Roles** page, select **No** for **Use a switch fabric for a role propagation?** option, and select the **Branch** option. The respective groups are configured to enable role propagation and role-based policy enforcement for those sites.

For more information about how to selectively enable groups for multi-site using HPE Aruba Networking Central, refer to [Performing Selective Enablement of Groups](#).

After successfully applying the configuration, role-based policies are enforced on all clients in the network.

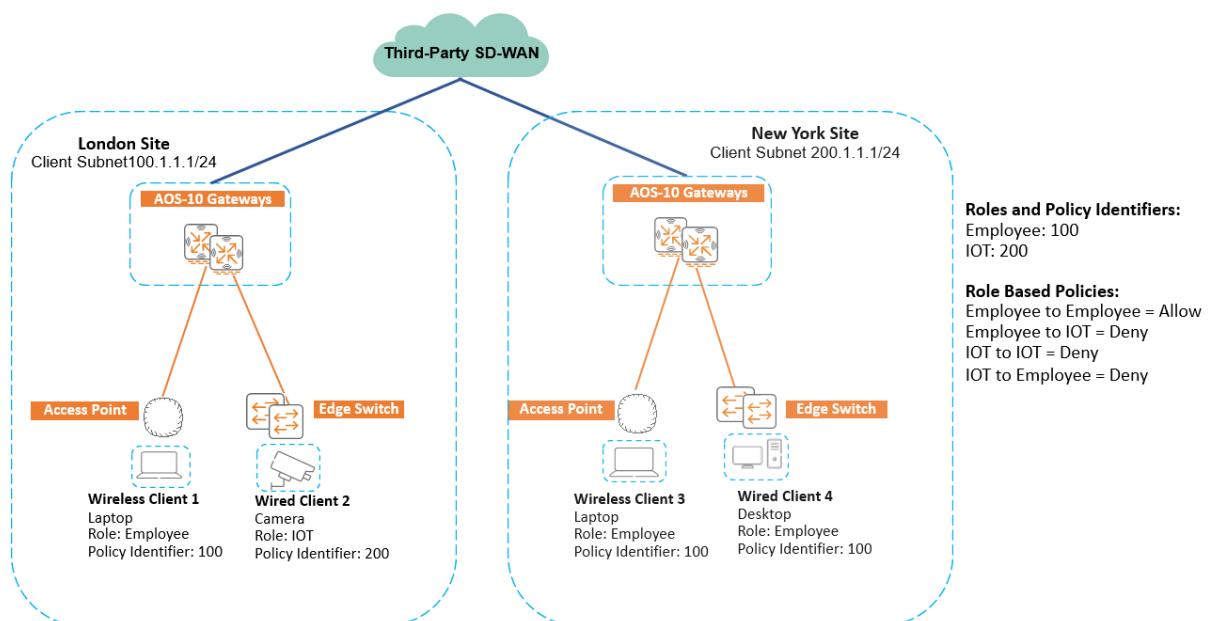
In the above example:

- Wireless Client 1 in the London site, which is authenticated with the role Employee communicates with Wired Client 3 with the role Employee in the New York site, but is denied communication with Wired Client 2 with the role IoT, although it is in the same site and group.
- Similarly, Wired Client 2 cannot communicate with either Wireless Client 1 or Wired Client 3.

Centralized Multi-site Fabric with Third-Party SD-WAN

This section describes an example use case for the Centralized Multi-site Fabric with Third-Party SD-WAN deployment using HPE Aruba Networking Central.

Figure 26 Centralized Multi-site Fabric with Third-Party SD-WAN Use Case



In the use case, the client subnet 100.1.1.1/24 maps to clients in the London site and the client subnet 200.1.1.1/24 maps to clients in the New York site. The goal is to enable role-based micro-segmentation across multiple geographic sites connected over a Third-Party SD-WAN fabric. Roles and role-based policies are centrally in the **Client Roles** page on HPE Aruba Networking Central. The third-party gateway is the WLAN gateway for the wireless clients and user-based tunnel Gateway for the wired clients within the site. The role-based policies are enforced on the Third-Party Gateways for all the clients within the site. Role propagation and role-based policy enforcement is selectively enabled per subnets.

As depicted in the above diagram, wired and wireless clients that are assigned the Employee role is allowed to communicate with other clients with the Employee role, but not to the clients with the IoT role. Conversely, clients with the IoT role are denied access to clients of both the IoT and Employee role.

Step 1: Creating Role and Role-based Policy

The roles and role-based policies are defined on the **Client Roles** page on HPE Aruba Networking Central. In the **Client Roles** page, the Employee role is created with a Policy Identifier 100 and the IoT role is created with a Policy Identifier of 200.

Policies for these roles are defined. The permission for the Employee role is assigned to allow source to destination. This allows communication between the clients with the Employee role. By default, all other role-based communication is denied.

After the permissions are assigned, the same role and policy definition is configured on all gateways and switches on the network.

For more information about how to create role and role-based policy using HPE Aruba Networking Central, refer to [Global Client Roles](#).

Step 2: Selective Enablement of Subnets

The next step is to selectively enable the subnets that maps to the clients in the respective sites. In the **Client Roles** page, select **No** for **Use a switch fabric for a role propagation?** option, and select the **Mobility** option. The respective client subnets are configured to enable role propagation and role-based policy enforcement for those sites.

For more information about how to selectively enable subnets for multi-site using HPE Aruba Networking Central, refer to [Performing Selective Enablement of Subnets](#).

After successfully applying the configuration, role-based policies are enforced on all clients in the network.

In the above example:

- Wireless Client 1 in the London site, which is authenticated with the role Employee communicates with Wired Client 3 with the role Employee in the New York site, but is denied communication with Wired Client 2 with the role IoT, although it is in the same site and subnet.
- Similarly, Wired Client 2 cannot communicate with either Wireless Client 1 or Wired Client 3.