Juniper Networks | Engineering Simplicity

Campus Fabric EVPN Multihoming Using Juniper Mist™ Wired Assurance—Juniper Validated Design (JVD)

Published
2025-12-18

Table of Contents

About this Document

Overview

This document covers how to deploy a Campus Fabric EVPN Multihoming architecture to support a campus networking environment using Juniper Mist Wired Assurance. The use case shows how you can deploy a single campus fabric that uses Ethernet VPN (EVPN) in the control plane, Virtual Extensible LAN (VXLAN) tunnels in the overlay network, and BGP in the underlay using integration with Juniper® Series of High-Performance Access Points (APs).

Solution Benefits

IN THIS SECTION

Most traditional campus architectures use single-vendor, chassis-based technologies that work well in small, static campuses with few endpoints. Campus architectures are too rigid to support the scalability and changing needs of modern large enterprises. Multi-chassis link aggregation group (MC-LAG) is a good example of a single-vendor technology that addresses the collapsed core deployment model. In this model, two chassis-based platforms are typically in the core of your network; deployed to handle all Layer 2 and Layer 3 requirements while providing an active/backup resiliency environment. MC-LAG does not interoperate between vendors, creating lock-in, and is limited to two devices.

A Juniper Networks EVPN multihoming solution based on EVPN-VXLAN addresses the collapsed core architecture and is simple, programmable, and built on a standards-based architecture that is common across campuses and data centers. See RFC 8365 for more information on this architecture.

EVPN multihoming uses a Layer 3 IP-based underlay network and an EVPN-VXLAN overlay network between the collapsed core Juniper switches. Broadcast, unknown unicast, and multicast (BUM) traffic, is handled natively by EVPN and eliminates the need for Spanning Tree Protocols (STP/RSTP). A flexible overlay network based on VXLAN tunnels combined with an EVPN control plane, efficiently provides Layer 3 or Layer 2 connectivity. This architecture decouples the virtual topology from the physical topology, which improves network flexibility and simplifies network management. Endpoints that require Layer 2 adjacency, such as Internet of Things (IoT) devices, can be placed anywhere in the network and remain connected to the same logical Layer 2 network.

With an EVPN multihoming deployment, up to four collapsed core devices are supported and all of them use EVPN-VXLAN. This standard is vendor-agnostic, so you can use the existing access layer infrastructure such as Link Aggregation Control Protocol (LACP) without the need to retrofit this layer of your network. Connectivity with legacy switches is accomplished with standards-based ESI-LAG. ESI-LAG uses standards-based LACP to interconnect with legacy switches.

Benefits of Campus Fabric EVPN Multihoming

The traditional Ethernet switching approach is inefficient because it leverages broadcast and multicast technologies to announce Media Access Control (MAC) addresses. It is also difficult to manage because you need to manually configure VLANs to extend them to new network ports. This problem is compounded significantly when considering the explosive growth of mobile and IoT devices.

EVPN multihoming's underlay topology is supported with a routing protocol that ensures loopback interface reachability between nodes. In the case of EVPN multihoming, Juniper Mist Wired Assurance supports eBGP between the core switching platforms. These devices support the EVPN-VXLAN function as VXLAN Tunnel Endpoint (VTEPs) that encapsulate and decapsulate the VXLAN traffic. VTEP represents the construct within the switching platform that originates and terminates VXLAN tunnels. In addition to this, these devices route and bridge packets in and out of VXLAN tunnels as required. EVPN multihoming addresses the collapsed core model traditionally supported by technologies like MC-LAG and Virtual Router Redundancy Protocol (VRRP). In this case, you can retain the investment at the access layer while supporting the fiber or cabling plant that terminates connectivity up to four core devices.

Figure 1: Campus Fabric EVPN Multihoming

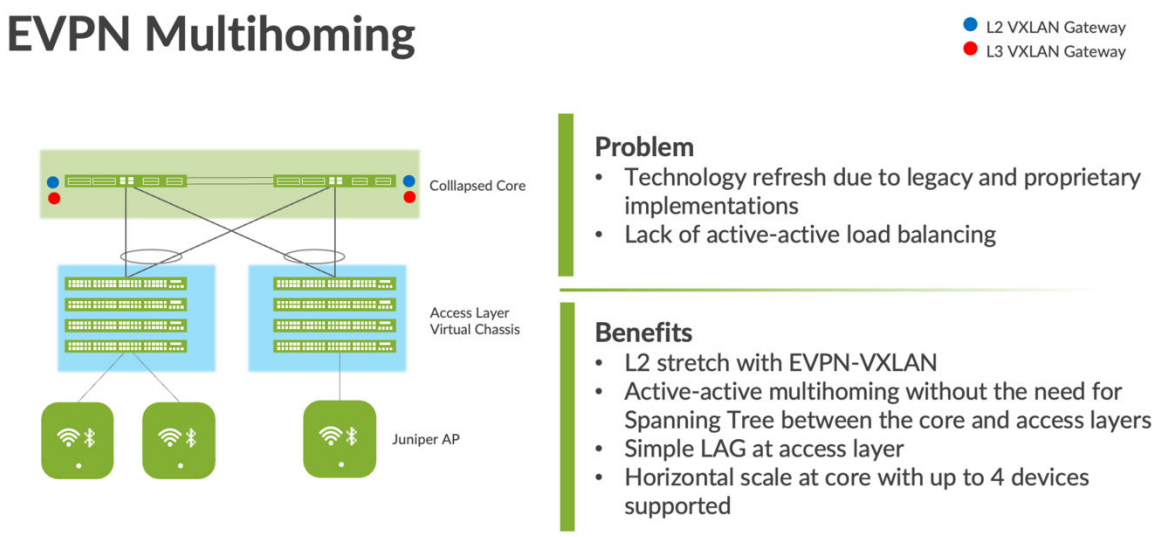

This architecture provides optimized, seamless, and standards-compliant Layer 2 or Layer 3 connectivity. Juniper Networks EVPN-VXLAN campus networks provide the following benefits:

- Consistent, scalable architecture—Enterprises typically have multiple sites with different size requirements. A common EVPN-VXLAN-based campus architecture is consistent across all sites, irrespective of the size. EVPN-VXLAN scales out or scales in as a site evolves.
- Multi-vendor deployment—The EVPN-VXLAN architecture uses standards-based protocols so enterprises can deploy campus networks using multi-vendor network equipment. There is no single vendor lock-in requirement.
- Reduced flooding and learning—Control plane-based Layer 2 and Layer 3 learning reduces the flood and learn issues associated with data plane learning. Learning MAC addresses in the forwarding plane has an adverse impact on network performance as the number of endpoints grows. The EVPN control plane handles the exchange and learning of routes, so newly learned MAC addresses are not exchanged in the forwarding plane.
- Location-agnostic connectivity—The EVPN-VXLAN campus architecture provides a consistent endpoint experience no matter where the endpoint is located. Some endpoints require Layer 2 reachability, such as legacy building security systems or IoT devices. The Layer 2 VXLAN overlay provides Layer 2 reachability across campuses without any changes to the underlay network. With our standards-based network access control integration, an endpoint can be connected anywhere in the network.
- Underlay agnostic—VXLAN as an overlay is underlay agnostic. With a VXLAN overlay, you can connect multiple campuses with a Layer 2 VPN or Layer 3 VPN service from a WAN provider or by using IPsec over Internet.
- Consistent network segmentation—A universal EVPN-VXLAN-based architecture across campuses and data centers means consistent end-to-end network segmentation for endpoints and applications.
- Simplified management—Campuses and data centers based on a common EVPN-VXLAN design can use common tools and network teams to deploy and manage campus and data center networks.

Solution Architecture

IN THIS SECTION

Juniper Mist Wired Assurance Overview

Juniper Mist Wired Assurance is a cloud service that brings automated operations and service levels to the campus fabric for switches, IoT devices, APs, servers, and printers. It is about simplifying every step of the way, starting from Day 0 for seamless onboarding and autoprovisioning through Day 2 and beyond for operations and management. Juniper Networks® EX Series Switches provide rich Junos OS streaming telemetry that enable the insights for switch health metrics and anomaly detection, as well as Mist AI™ capabilities.

Mist's AI engine and Marvis® Virtual Network Assistant further simplify troubleshooting while streamlining helpdesk operations by monitoring events and recommending actions. Marvis is one step towards the Self-Driving Network™, turning insights into actions and fundamentally transforming Information Technology (IT) operations from reactive troubleshooting to proactive remediation.

Juniper Mist™ cloud services are 100% programmable using open Application Programming Interfaces (APIs) for full automation, integration with your operational support systems, or both. Operational support systems include IT applications, ticketing systems, and IP management systems.

Juniper Mist™ delivers unique capabilities for the WAN, LAN, and Wireless networks such as the following:

- User Interface (UI) or API-driven configuration at scale.
- Service-level expectations (SLEs) for key performance metrics such as throughput, capacity, roaming, and uptime.

- Marvis® Virtual Network Assistant—An integrated AI engine that provides rapid troubleshooting of full stack network issues, trending analysis, anomaly detection, and proactive problem remediation.
- Single management system.
- License management.
- Premium Analytics for long term trending and data storage.

For more information about Juniper Mist Wired Assurance, see the following datasheet: https://www.juniper.net/content/dam/www/assets/datasheets/us/en/cloud-services/juniper-mist- wired-assurance-datasheet.pdf

Campus Fabric Core-Distribution High-Level Architecture

EVPN multihoming, with an EVPN-VXLAN architecture, decouples the overlay network from the underlay network. This approach addresses the needs of the modern enterprise network by allowing network administrators to create logical Layer 2 networks across one or more Layer 3 networks. In an EVPN multihoming deployment, the use of EVPN VXLAN supports native traffic isolation using routing instances; commonly called virtual routing and forwarding (VRFs) for macro-segmentation purposes.

The Juniper Mist™ portal workflow makes it easy to create campus fabrics.

Figure 2: High-level Campus Fabric Creation

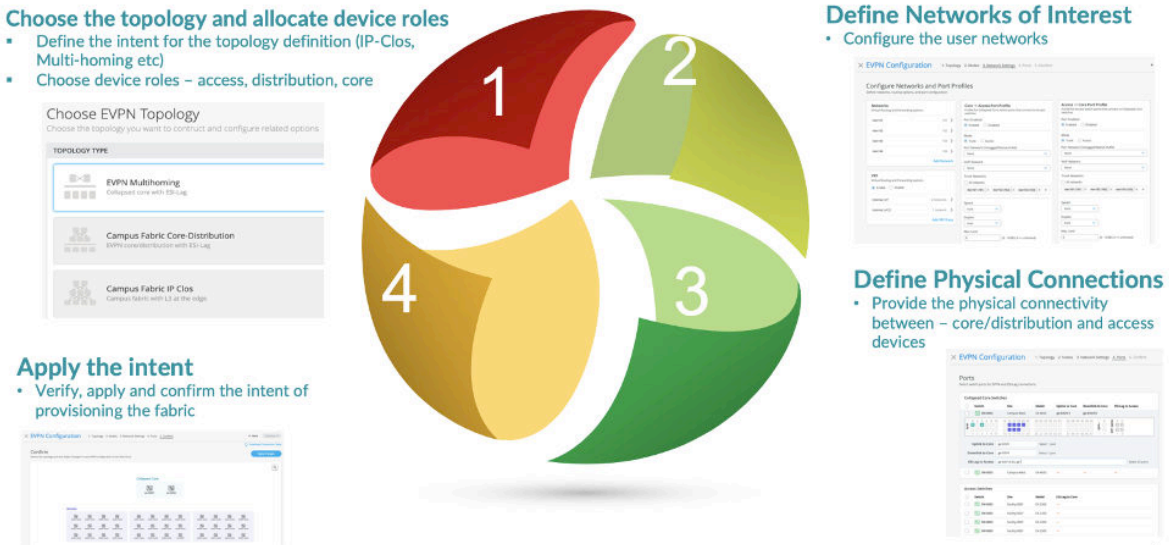

Underlay Network

An EVPN-VXLAN fabric architecture makes the network infrastructure simple and consistent across campuses and data centers. All the collapsed core devices must be connected to each other using a Layer 3 infrastructure.

You can use any Layer 3 routing protocol to exchange loopback addresses between the core and distribution devices. BGP provides benefits such as better prefix filtering, traffic engineering, and route tagging. Mist configures eBGP as the underlay routing protocol in this example. Juniper Mist automatically provisions private autonomous system numbers and all BGP configuration for the underlay and overlay for only the campus fabric. There are options to provide additional BGP speakers to allow you to peer with external BGP peers.

Underlay BGP is used to learn loopback addresses from peers so that the overlay BGP can establish neighbors using the loopback address. The overlay is then used to exchange EVPN routes.

Figure 3: Pt-Pt Links Using /31 Addressing Between Collapsed Core Switches

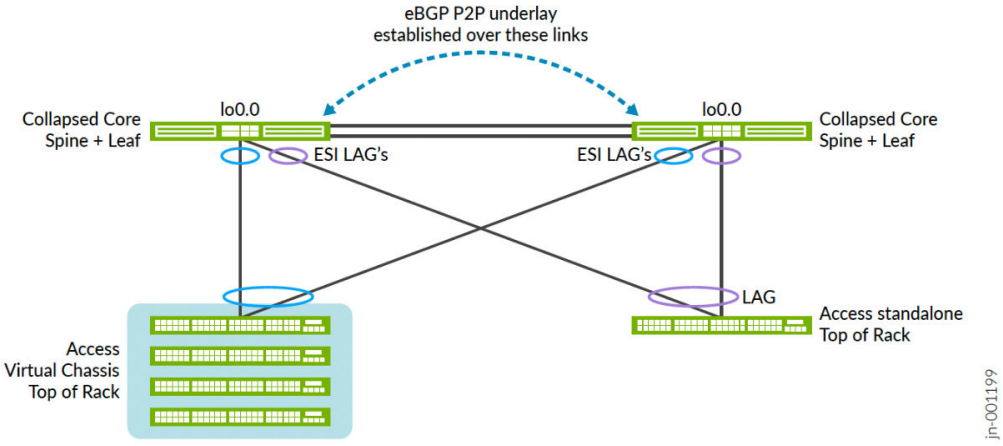

Network overlays enable connectivity and addressing independent of the physical network. Ethernet frames are wrapped in IP UDP datagrams, which are encapsulated into IP for transport over the underlay. VXLAN enables virtual Layer 2 subnets or VLANs to span underlying physical Layer 3 network.

In a VXLAN overlay network, each Layer 2 subnet or segment is uniquely identified by a Virtual Network Identifier (VNI). A VNI segments traffic the same way that a VLAN ID does. This mapping occurs on the core, distribution, and border gateway, which can reside on the core or services block. As is the case with VLANs, endpoints within the same virtual network can communicate directly with each other.

Endpoints in different virtual networks require a device that supports inter-VXLAN routing, which is typically a router, or a high-end switch known as a Layer 3 gateway. The entity that performs VXLAN encapsulation and decapsulation is called a VTEP. Each VTEP is known as the Layer 2 gateway and

typically assigned with the device's loopback address. This is also where VXLAN (commonly known as VNI) to VLAN mapping exists.

Figure 4: VXLAN VTEP Tunnels

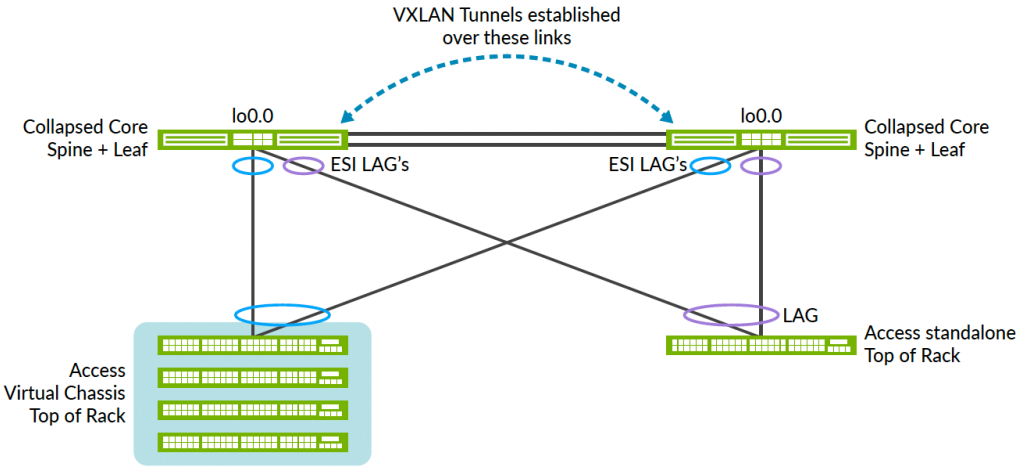

VXLAN can be deployed as a tunnelling protocol across a Layer 3 IP campus fabric without a control plane protocol. However, the use of VXLAN tunnels alone does not change the flood and learn behavior of the Ethernet protocol.

The two primary methods for using VXLAN without a control plane protocol are static unicast VXLAN tunnels and VXLAN tunnels. These methods are signaled with a multicast underlay and do not solve the inherent flood and learn problem and are difficult to scale in large multitenant environments. These methods are not in the scope of this documentation.

Understanding EVPN

Ethernet VPN is a BGP extension to distribute endpoint reachability information such as MAC and IP addresses to other BGP peers. This control plane technology uses Multiprotocol BGP (MP-BGP) for MAC and IP address endpoint distribution, where MAC addresses are treated as type 2 EVPN routes. EVPN enables devices acting as VTEPs to exchange reachability information with each other about their endpoints.

Juniper supported EVPN standards: https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn.html

What is EVPN-VXLAN: https://www.juniper.net/us/en/research-topics/what-is-evpn-vxlan.html

The benefits of using EVPNs include:

- MAC address mobility
- Multitenancy
- Load balancing across multiple links
- Fast convergence
- High availability
- Scale
- Standards-based interoperability

EVPN provides multipath forwarding and redundancy through an all-active model. The collapsed core layer can have up to four devices in a ring or mesh topology. If one core device fails, traffic flows use the remaining active links.

The technical capabilities of EVPN include:

- Minimal flooding—EVPN creates a control plane that shares end host MAC addresses between VTEPs.
- Multihoming—EVPN supports multihoming for client devices. A control protocol like EVPN that enables synchronization of endpoint addresses between the distribution switches is needed to support multihoming, because traffic traveling across the topology needs to be intelligently moved across multiple paths.
- Aliasing—EVPN leverages all-active multihoming when connecting devices to the distribution layer of a campus fabric. The connection from the multihomed distribution layer switches is called an ESI-LAG, while the access layer devices connect to each distribution switch using standard LACP.
- Split horizon—Split horizon prevents the looping of broadcast, unknown unicast, and multicast (BUM) traffic in a network. With split horizon, a packet is never sent back over the same interface it was received on, which prevents loops.

Overlay Network (Data Plane)

VXLAN is the overlay data plane encapsulation protocol that tunnels Ethernet frames between network endpoints over the underlay network. Devices that perform VXLAN encapsulation and decapsulation for the network are referred to as a VTEP. Before a VTEP sends a frame into a VXLAN tunnel, it wraps the original frame in a VXLAN header that includes a VNI. The VNI maps the packet to the original VLAN at the ingress switch. After applying a VXLAN header, the frame is encapsulated into a UDP/IP packet for transmission to the remote VTEP over the IP fabric, where the VXLAN header is removed and the VNI to VLAN translation happens at the egress switch.

Figure 5: VXLAN Header

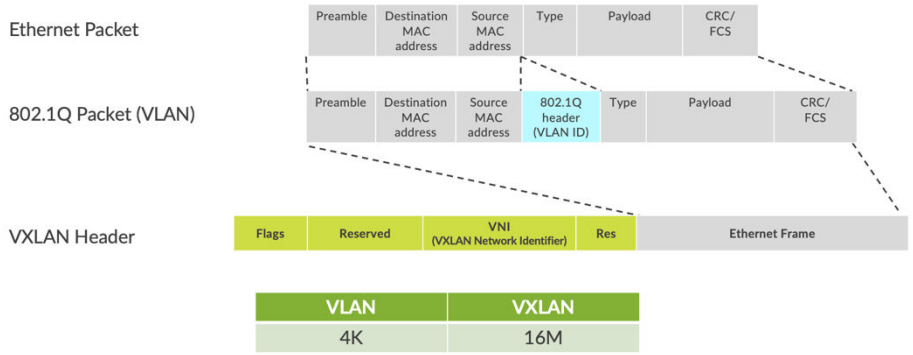

VLAN	VXLAN
4K	16M

VTEPs are software entities tied to a device's loopback address that source and terminate VXLAN tunnels. VXLAN tunnels in an EVPN multihoming fabric are provisioned only at the collapsed core switches.

Overlay Network (Control Plane)

MP-BGP with EVPN signalling acts as the overlay control plane protocol. Adjacent switches peer using their loopback addresses using next hops announced by the underlay BGP sessions. The collapsed core devices establish eBGP sessions between each other. When there is a Layer 2 forwarding table update on any switch participating in campus fabric, it sends a BGP update message with the new MAC route to other devices in the fabric. Those devices then update their local EVPN database and routing tables. In EVPN multihoming fabrics, the control plane exchange happens through interior BGP and each collapsed core switch acts as a route reflector.

Figure 6: EVPN Overlay Network iBGP Synchronization

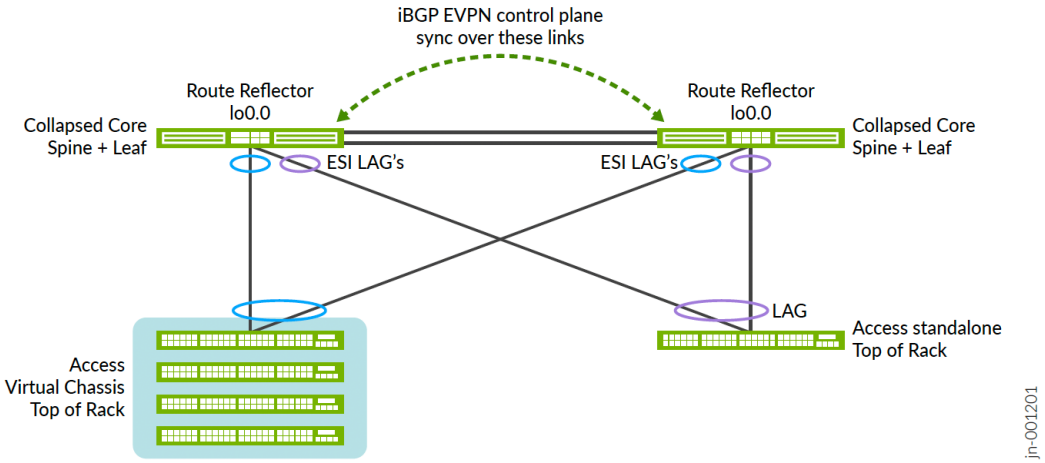

Resiliency and Load Balancing

We support Bidirectional Forwarding Detection (BFD) as part of the BGP protocol implementation. This provides fast convergence in the event of a device or link failure without relying on the routing protocol's timers. Mist configured BFD minimum intervals of 1000ms and 3000ms in the underlay and overlay respectively. Load balancing, per packet by default, is supported across all core-distribution links within the campus fabric using ECMP enabled at the forwarding plane.

Ethernet Segment Identifier (ESI)

When the access layer multihomes to the distribution layer devices in a campus fabric, an ESI-LAG is formed on the distribution layer devices. This ESI is a 10-octet integer that identifies the Ethernet segment amongst the distribution layer switches participating in the ESI. MP-BGP is the control plane protocol used to coordinate this information. ESI-LAG enables link failover in the event of a bad link, supports active-active load balancing, and is automatically assigned by Juniper Mist.

Figure 7: Resiliency and Load Balancing

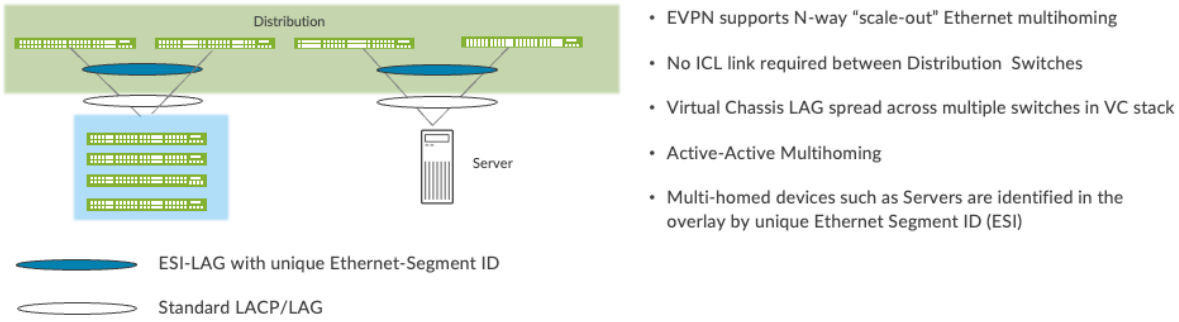

Access Layer

The access layer provides network connectivity to end-user devices, such as personal computers, VoIP phones, printers, and IoT devices as well as connectivity to wireless APs. In this example, we use Juniper APs as the access point devices. Evolving IT departments are looking for a cohesive approach for managing wired and wireless networks. Juniper Networks has a solution that can simplify and automate operations and end-to-end troubleshooting, ultimately evolving into the Self-Driving Network™.

The access switch itself is only demanded to support IEEE 802.3ad Link Aggregation and active LACP on two uplinks towards the collapsed core switches of the EVPN multihoming fabric. The VLANs configured on the ports where the wired client and APs are attached are then multiplexed and tagged on the uplinks.

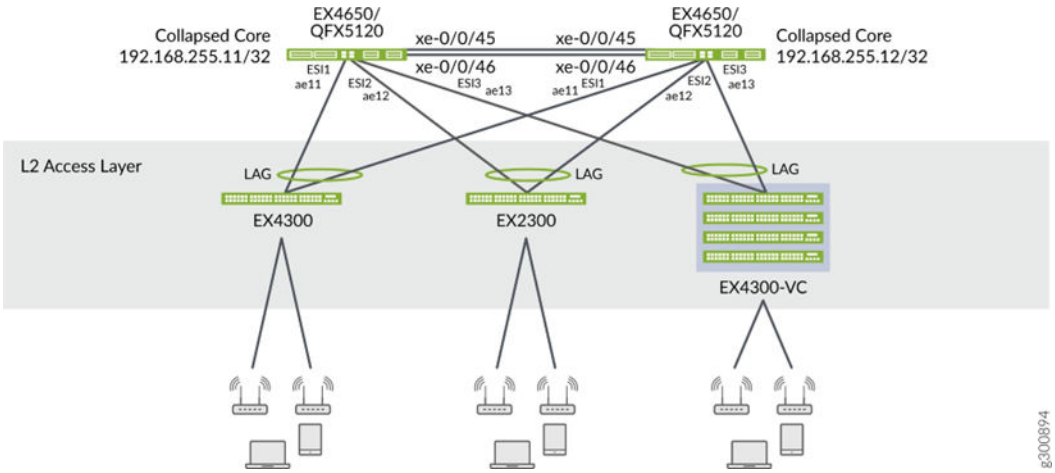

Single or Multi PoD Design

Juniper Mist campus fabrics support deployments with only one PoD (formally called Site-Design) or multiple PoDs. The organizational deployment shown below targets enterprises who need to align with a multi-POD structure:

Figure 8: Multiple PoD Design Example

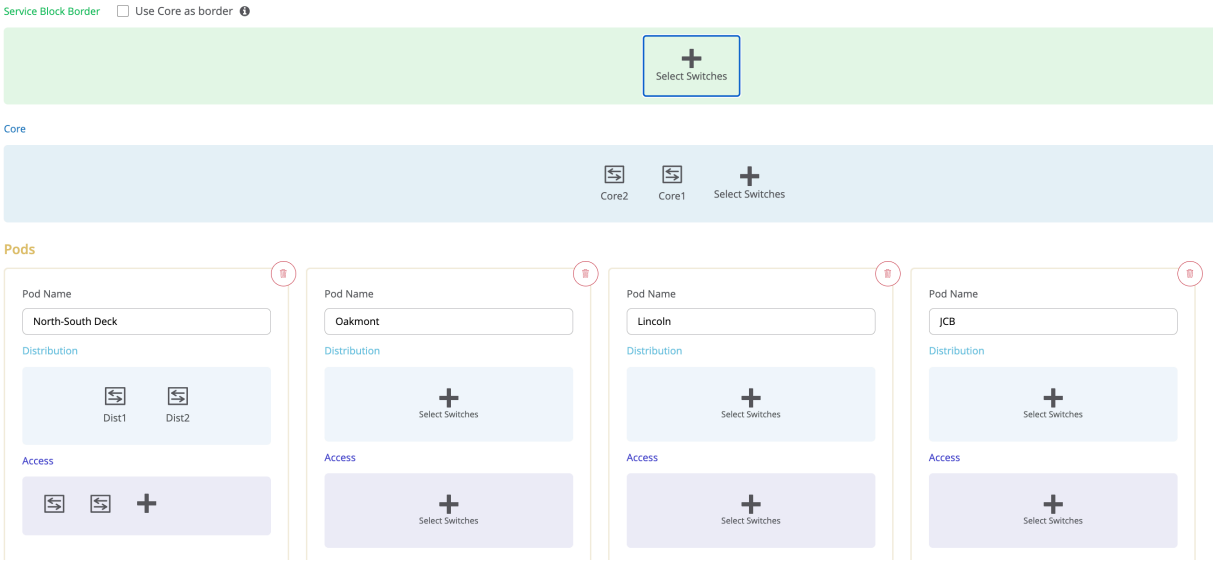

NOTE: This multi-PoD option is not available with EVPN multihoming fabrics. However, you can instead build multiple sites each with a small EVPN multihoming fabric with the limit that you cannot stretch VLANs between those sites.

Juniper Access Points

In our network, we choose Juniper APs as our preferred access point devices. They are designed from the ground up to meet the stringent networking needs of the modern cloud and smart device era. Juniper Mist delivers unique capabilities for both wired and wireless LAN:

- Wired and wireless assurance—Juniper Mist is enabled with Wired and Wireless Assurance. Once configured, service-level expectations (SLEs) for key wired and wireless performance metrics such as throughput, capacity, roaming, and uptime are addressed in the Juniper Mist platform. This JVD uses Juniper Mist Wired Assurance services.

- Marvis—An integrated AI engine that provides rapid wired and wireless troubleshooting, trending analysis, anomaly detection, and proactive problem remediation.

VRF Segmentation

VRF segmentation is used to organize users and devices in groups on a shared network while separating and isolating the different groups. The routing devices on the network create and maintain a separate virtual routing and forwarding (VRF) table for each group. The users and devices in a group are placed in one VRF segment and can communicate with each other, but they cannot communicate with users in another VRF segment. If you want to send and receive traffic from one VRF segment to another VRF segment, you must configure the routing path on the WAN router of the fabric which can also implement stateful firewalls.

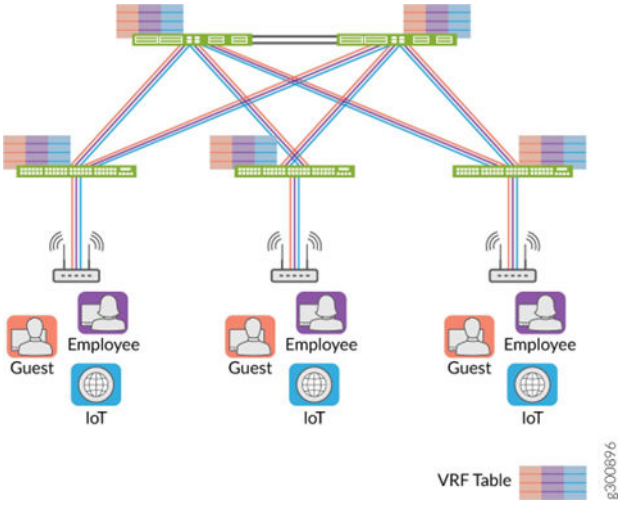

Supported Platforms for Campus Fabric EVPN Multihoming

To review the software versions and platforms on which this JVD was validated by Juniper Networks, see the Validated Platforms and Software section in this document.

Table 1: Supported Platforms for Campus Fabric EVPN Multihoming Deployment

Campus Fabric EVPN Multihoming Deployment	Supported Platforms
Access layer	EX2300 EX3400 EX4300 EX4100 EX4400
Collapsed Core layer	EX4400-24X EX4650 QFX5120 QFX5130 QFX5700 EX92xx

Validation Framework

IN THIS SECTION

Test Bed

The diagram below shows the suggested topology used for the JVD lab evaluating various EVPN multihoming fabric topologies.

Figure 9: JVD Lab Proposal

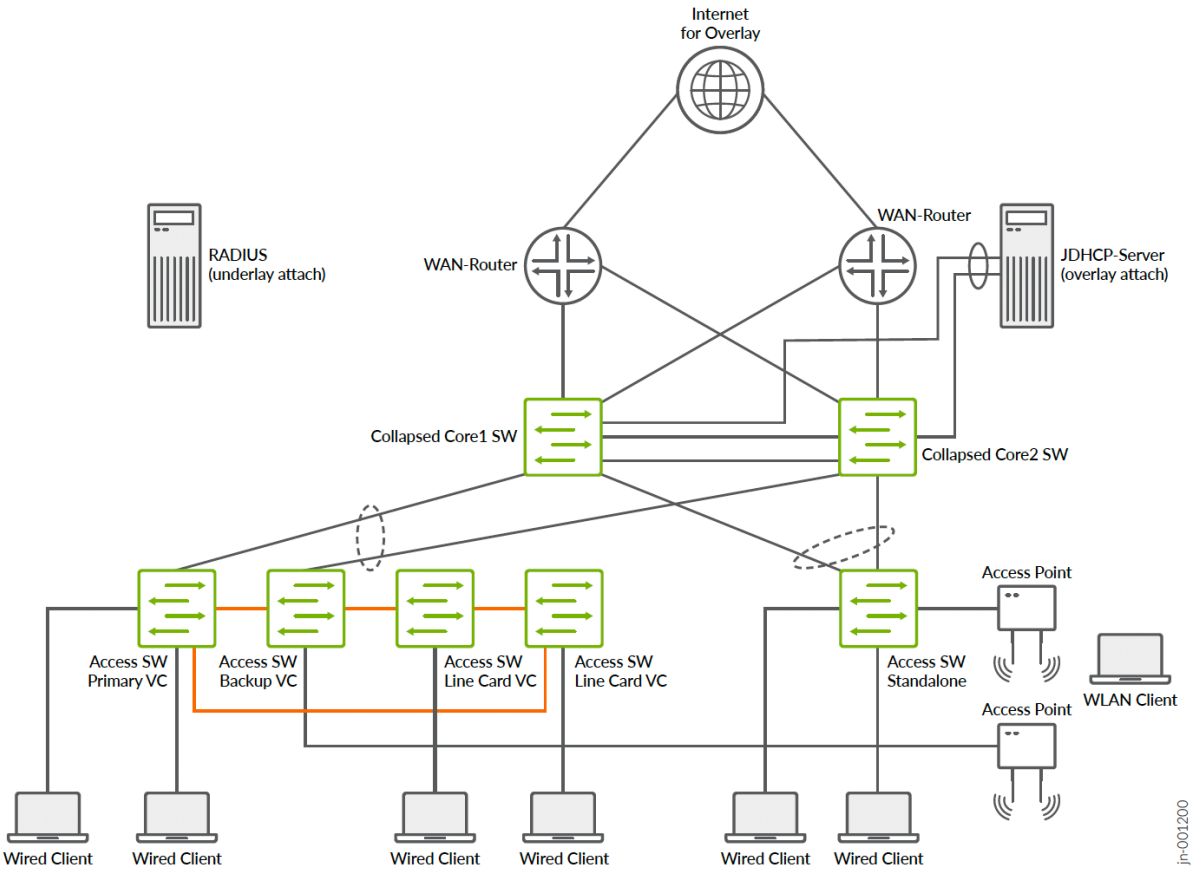

The suggested lab design allows for the evaluation of the following:

- EVPN multihoming fabrics with:
  - Two collapsed core switches acting as spine and leaf.
    - One 4-Member Virtual Chassis access switch acting as ToR.
    - One standalone access switch acting as ToR.
- Service block function through:
  - Integrated into existing collapsed core switches and acting as service leaf and collapsed core at the same time.

    - Attached WAN routers through Layer 2 or Layer 3 exit.
    - Attached servers through ESI-LAG redundant links.
- WAN router integration:
    - Layer 2 fabric exit.
        - ESI-LAG-based trunks.
    - Layer 3 fabric exit.
        - OSPF as routing protocol.
        - eBGP as routing protocol
    - Attached to:
        - Collapsed core switch.
    - Redundant WAN router design:
        - Two Juniper MX routers.
        - Two Juniper SRX firewalls in cluster configuration.
- Wi-Fi access points:
    - Local-attached to the access switches with Power over Ethernet (PoE).
    - Various Wi-Fi clients.
    - Basic Wi-Fi roaming.
- Overlay server attached to a service block functionality:
    - DHCP server.
    - Other services.
- RADIUS server:
    - Server location:
        - Local server attached to underlay network.
        - Remote Juniper Mist Access Assurance through public cloud.
    - Authentication for the following clients:
        - Wired clients attached to access switches.
        - Wi-Fi clients using the access points.
    - Authentication based on clients:

    - MAC address.
    - 802.1X EAP authentication.
  - Dynamic authorization profiles:
    - Single VLAN assigned.
    - Multiple VLANs assigned.
    - Filter-Id ACL assigned.
- Testing fabric features such as:
  - DHCP relay
  - Protect RE-filter
  - DHCP snooping
  - Storm control
  - MAC address limit with aging
  - DNS
  - NTP

The following network topology was tested as part of this JVD.

Figure 10: EVPN Multihoming Fabric with Two Collapsed Cores

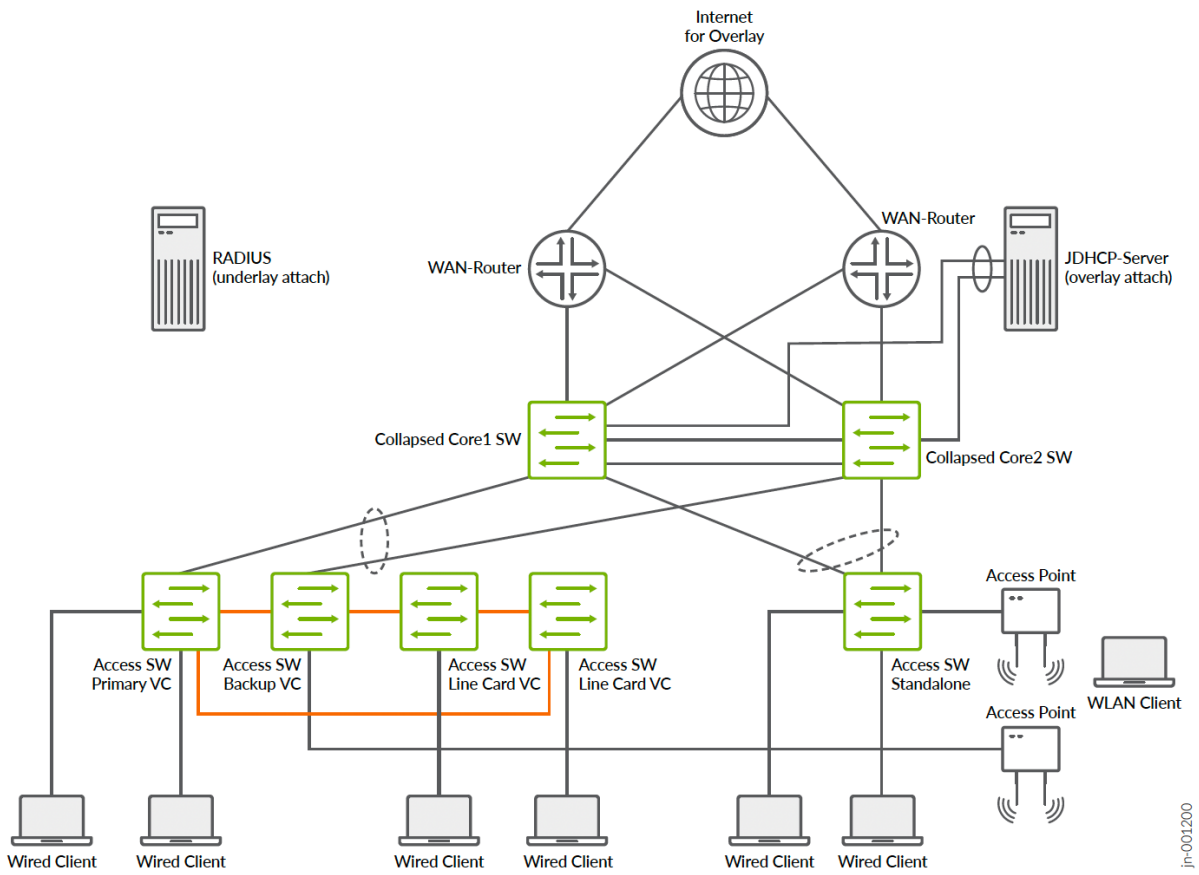

Platforms / Devices Under Test (DUT)

To review the software versions and platforms on which this JVD was validated by Juniper Networks, see the Validated Platforms and Software section in this document.

Test Bed Configuration

In the appendix section of this JVD, we are sharing information on exactly how some of the tests were performed. Contact Juniper or your Juniper account representative to obtain the full archive of the test bed configuration used for this JVD.

Test Objectives

IN THIS SECTION

Test Goals

The testing for this JVD was performed with the following goals in mind. Also, consult the separate Test Report for more information. The testing was executed with a focus on the following:

- Testing with Junos OS version 23.4R2.
- Testing with two collapsed core switches.
- Testing with Virtual Chassis that have two or four members.
- Testing with features that are activated as combinations at the same time.

The scale testing for this design was done with:

- Up to 20 VRFs.
- Up to 500 VLANs (across all VRFs).
- Up to 45K IP and MAC addresses of simulated wired clients.

Test Non-Goals

The testing for this JVD was not performed, for various reasons, on the following items:

- Testing this fabric with redundant WAN routers. This is already described in a separate JVD extension in common for all fabrics.
- Juniper Mist Edge integration for Wi-Fi scaling.
- Testing with IPv6 underlay and overlay was not available in time and was moved to a later phase.

- Testing with four collapsed cores in a ring or mesh topology was not available in time and was moved to a later phase.

Recommendations

The following simple guidelines will help you to successfully implement a campus fabric EVPN Multihoming design into your network.

- Review the JVD extension for WAN router integration.
- All fabric networks should be configured in the following way to avoid inconsistency:
  - First, create them as part of your switch template for a site.
  - Then, import the created networks as part of the campus fabric dialogue and assign to VRFs.
  - Even if the system allows you a local network creation on a switch, do not use this option.
- Do not manually configure VRFs locally on any switch. The fabric usually does this automatically on an as-needed basis.
  - The current exception to this rule is for a Layer 2 WAN router integration through transport VLAN. Review the JVD extension for WAN router integration and follow the example in the appendix.
- When using DHCP relay configuration for the fabric:
  - Review the JVD extension which covers DHCP relay configuration.
  - Only use the fabric dialogue for configuring DHCP relay and no local configuration directly on a switch.
- When designing and using Virtual Chassis:
  - Virtual Chassis can only be used at the access switch layer of a campus fabric environment:
  - When designing a Virtual Chassis, it is not advised to use the maximum number of supported members listed in the Virtual Chassis Overview (Juniper Mist). A good rule of thumb is to use roughly half of the stated maximum. This helps prevent bandwidth oversubscription on the VCPs that form the ring between the chassis members.
  - Create and assign separate templates for Virtual Chassis systems that have the same number of members. Avoid applying identical port configurations to Virtual Chassis setups of different sizes. This approach allows the system to apply configuration changes directly, without repeatedly checking whether the ports defined in the template actually exist on the local Virtual Chassis.

    - All Virtual Chassis configurations should be done through the Juniper Mist cloud and the Modify Virtual Chassis dialogue. Additional CLI or CLI commands should not be used for managing a Virtual Chassis.

- Unassigned access ports should be configured with a quarantine VLAN or disabled ports using a template. Please review the example here.
    - If possible, use a different VRF for the quarantine VLAN to isolate this traffic.
    - Best practice is also enabling “STP Edge” in the quarantine port profile.
- When deciding how to manage port configurations dynamically:
    - Using RADIUS or a NAC system to assign VLANs and filters is the recommended method, particularly for customers using Juniper Mist Access Assurance.
    - Dynamic Port Configuration is considered a less preferred option.
- When using Dynamic Port Configuration:
    - Avoid matching by MAC address if the device supports LLDP.
    - Don’t match by MAC address if ports are enabled with dot1x.
    - The use of a filter-id should be avoided. In most cases, this is unnecessary when ports are 802.1X-enabled and a dynamic VLAN can be assigned through RADIUS.
    - Avoid a high number of port flaps for a DPC-configured port.
    - Refer switch insights to ascertain the individual configuration is applied.
- Traffic towards a third-party RADIUS Server is expected to use inet.0 via the management port, same as the management traffic towards the Juniper Mist cloud, for example, underlay. This allows you to fine-tune the MTU for the UDP Packets send towards such a service in case it is needed.

APPENDIX: Example EVPN Multihoming Fabric Creation

IN THIS SECTION

NOTE: The examples shown in the entire appendix section are made with functional testing in mind. Shortcuts are also made on WAN router integration which is not a production grade design. However, with the design below, you can easily evaluate how a new campus fabric is deployed.

Campus Fabric EVPN Multihoming Components

This configuration example uses the following devices:

- Two QFX5110 switches as distribution devices, software version: Junos OS Release 22.4R3-S2 or later.
- Two access layer EX4100 switches, software version: Junos OS Release 22.4R3-S2 or later.
- One SRX345 WAN router, software version: 21.2R3-S7 or later.
- Juniper APs.
- Two Linux desktops that act as wired clients.

Figure 11: Topology

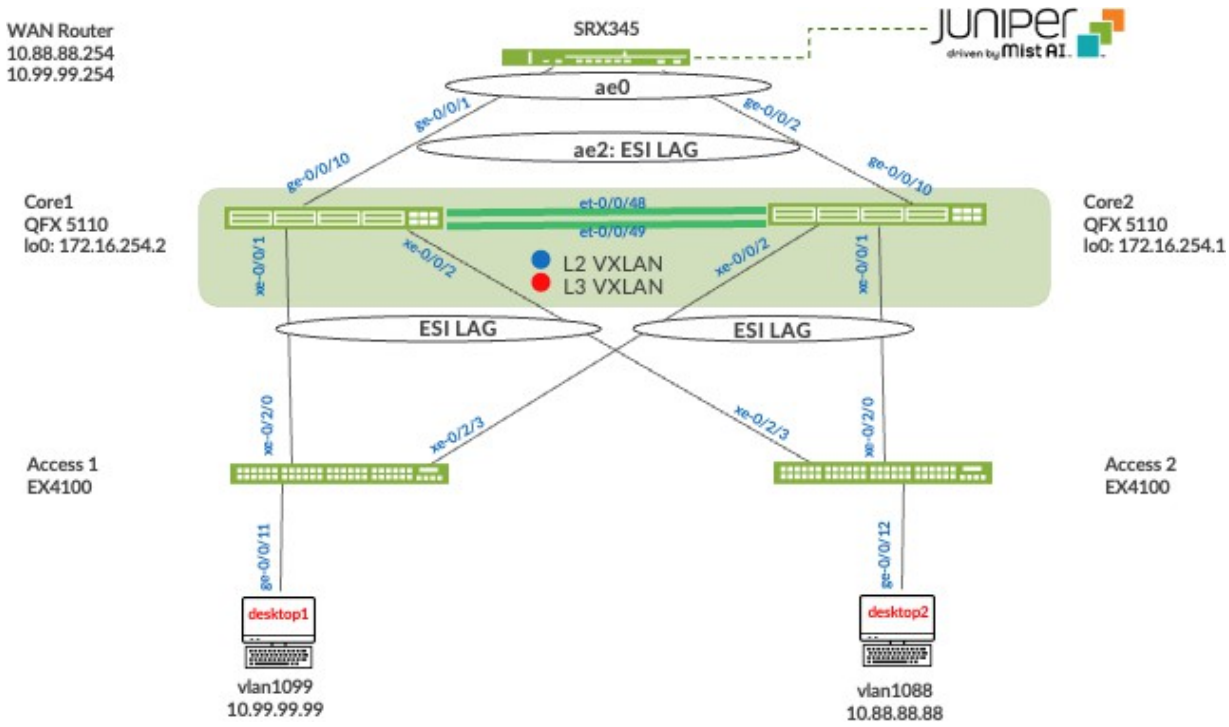

Juniper Mist Wired Assurance

Juniper Mist Wired Assurance, through the Juniper Mist portal, can be used to centrally manage all Juniper switches. Juniper Mist Wired Assurance gives you full visibility into the devices that comprise your network's access layer. The portal provides a user interface to access your architecture through the AI-driven cloud services with your Juniper Mist account. You can monitor, measure, and get alerts on key compliance metrics on the wired network. This includes switch version and PoE compliance, switch AP affinity, and VLAN insights.

The following link describes onboarding Juniper switches to the Juniper Mist cloud: https://www.juniper.net/documentation/us/en/quick-start/hardware/cloud-ready-switches/topics/topic-map/step-1-begin.html

Juniper Mist Wired Assurance, through the portal, is used to build Campus Fabric EVPN Multihoming from the ground up. This includes the following:

- Assignment of point-to-point (P2P) links between the core and distribution layers.
- Assignment of unique BGP AS numbers per device participating in the underlay and overlay.

- The creation of VRF instances allows you to logically segment traffic. This also includes the assignment of new or existing VLANs to each representative VRF.
- IP addressing of each Layer 3 gateway integrated routing and bridging (IRB) interface assigned to the distribution layer.
- IP addressing of each loopback interface.
- Configuration of routing policies for underlay and overlay connectivity.
- Optimized maximum transmission unit (MTU) settings for P2P underlay, Layer 3 IRB, and ESI-LAG bundles.
- Downloadable connection table (CSV format) that can be used by those involved in the physical buildout of the campus fabric.
- Graphical interface depicting all devices with BGP peering and physical link status.

For more information on Juniper Mist Wired Assurance, see: https://www.mist.com/documentation/category/wired-assurance/

Juniper Mist Wired Assurance Switches

You must validate that each device participating in the campus fabric has been adopted or claimed and assigned to a site. The switches are named for respective layers in the fabric to facilitate building and operating the fabric.

Figure 12: Switch Inventory

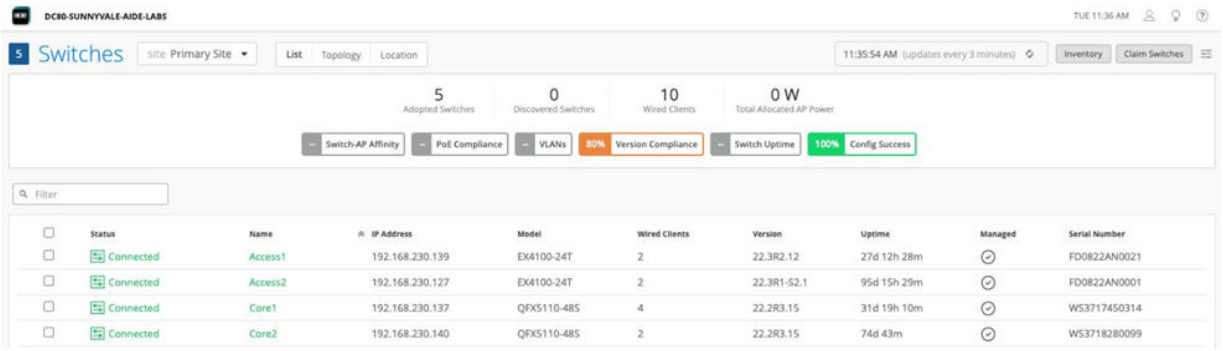

Templates

A key feature of switch management through the Juniper Mist cloud is to use templates and a hierarchical model to group the switches and make bulk updates. Templates provide uniformity and convenience, while the hierarchy (site and switch) provides both scale and granularity.

Templates and the hierarchical model mean that you can create a template configuration and then all the devices in each group inherit the template settings. When a conflict occurs, for example, when there are settings at both the site and organizational levels that apply to the same device, the narrower settings (in this case, the site settings) override the broader settings defined at the organization level.

Individual switches, at the bottom of the hierarchy, can inherit all or part of the configuration defined at the organization level, and again at the site level. Of course, individual switches can also have their own unique configurations.

You can include individual CLI commands at any level of the hierarchy, which are then appended to all the switches in that group on an “AND” basis—that is, individual CLI settings are appended to the existing configuration (existing settings might be replaced or appended).

NOTE: If you run CLI commands for items not native to the portal, this configuration data is applied last; overwriting existing configuration data within the same stanza. You can access the CLI command option from the switch template or individual switch configuration.

Figure 13: Adding Additional CLI

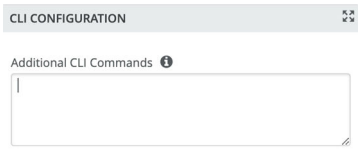

Under **Organization -> Switch Templates**, we use the following template:

Figure 14: Switch Templates

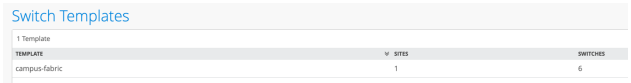

We provide a copy of the following template in JSON format for importing into your own system for verification:

```
{
  "ntp_servers": [],
  "dns_servers": [
    "8.8.8.8",
    "9.9.9.9"
  ],
  "dns_suffix": [],
  "additional_config_cmds": [],
  "networks": {
    "vlan1099": {
      "vlan_id": 1099,
      "subnet": "10.99.99.0/24"
    },
    "vlan1088": {
      "vlan_id": 1088,
      "subnet": "10.88.88.0/24"
    },
    "vlan1033": {
      "vlan_id": 1033,
      "subnet": "10.33.33.0/24"
    }
  },
  "port_usages": {
    "myaccess": {
      "mode": "trunk",
      "disabled": false,
      "port_network": "vlan1033",
      "voip_network": null,
      "stp_edge": false,
      "port_auth": null,
      "all_networks": false,
      "networks": [
        "vlan1033",
        "vlan1088",
        "vlan1099"
      ],
      "speed": "auto",
      "duplex": "auto",
      "mac_limit": 0,
```

```
      "persist_mac": false,
      "poe_disabled": false,
      "enable_qos": false,
      "storm_control": {},
      "mtu": 9018,
      "description": ""
    },
    "myesilag": {
      "mode": "trunk",
      "disabled": false,
      "port_network": null,
      "voip_network": null,
      "stp_edge": false,
      "port_auth": null,
      "all_networks": true,
      "networks": [],
      "speed": "auto",
      "duplex": "auto",
      "mac_limit": 0,
      "persist_mac": false,
      "poe_disabled": false,
      "enable_qos": false,
      "storm_control": {},
      "mtu": 9014,
      "description": ""
    },
    "dynamic": {
      "mode": "dynamic",
      "reset_default_when": "link_down",
      "rules": []
    },
    "vlan1099": {
      "mode": "access",
      "disabled": false,
      "port_network": "vlan1099",
      "voip_network": null,
      "stp_edge": false,
      "all_networks": false,
      "networks": null,
      "port_auth": null,
      "speed": "auto",
      "duplex": "auto",
      "mac_limit": 0,
```

```
      "persist_mac": false,
      "poe_disabled": false,
      "enable_qos": false,
      "storm_control": {},
      "mtu": 9014,
      "description": "Corp-IT",
      "disable_autoneg": false,
      "mac_auth_protocol": null,
      "enable_mac_auth": null,
      "mac_auth_only": null,
      "guest_network": null,
      "bypass_auth_when_server_down": null
    },
    "vlan1088": {
      "mode": "access",
      "disabled": false,
      "port_network": "vlan1088",
      "voip_network": null,
      "stp_edge": false,
      "all_networks": false,
      "networks": null,
      "port_auth": null,
      "speed": "auto",
      "duplex": "auto",
      "mac_limit": 0,
      "persist_mac": false,
      "poe_disabled": false,
      "enable_qos": false,
      "storm_control": {},
      "mtu": 9014,
      "description": "Developers",
      "disable_autoneg": false,
      "mac_auth_protocol": null,
      "enable_mac_auth": null,
      "mac_auth_only": null,
      "guest_network": null,
      "bypass_auth_when_server_down": null
    },
    "vlan1033": {
      "mode": "access",
      "disabled": false,
      "port_network": "vlan1033",
      "voip_network": null,
```

```
      "stp_edge": false,
      "all_networks": false,
      "networks": null,
      "port_auth": null,
      "speed": "auto",
      "duplex": "auto",
      "mac_limit": 0,
      "persist_mac": false,
      "poe_disabled": false,
      "enable_qos": false,
      "storm_control": {},
      "mtu": 9014,
      "description": "Guest-WiFi",
      "disable_autoneg": false,
      "mac_auth_protocol": null,
      "enable_mac_auth": null,
      "mac_auth_only": null,
      "guest_network": null,
      "bypass_auth_when_server_down": null
    }
  },
  "switch_matching": {
    "enable": true,
    "rules": [
      {
        "name": "core",
        "match_model": "EX9204",
        "port_config": {},
        "additional_config_cmds": [
          ""
        ],
        "ip_config": {
          "type": "dhcp",
          "network": "default"
        },
        "oob_ip_config": {
          "type": "dhcp",
          "use_mgmt_vrf": false
        }
      },
      {
        "name": "distribution",
        "port_config": {},
```

```
        "additional_config_cmds": [
          ""
        ],
        "ip_config": {
          "type": "dhcp",
          "network": "default"
        },
        "oob_ip_config": {
          "type": "dhcp",
          "use_mgmt_vrf": false
        },
        "match_model[0:7]": "QFX5120"
      },
      {
        "name": "access",
        "port_config": {
          "ge-0/0/16": {
            "usage": "myaccess",
            "dynamic_usage": null,
            "critical": false,
            "description": "",
            "no_local_overwrite": true
          }
        },
        "additional_config_cmds": [
          ""
        ],
        "ip_config": {
          "type": "dhcp",
          "network": "default"
        },
        "oob_ip_config": {
          "type": "dhcp",
          "use_mgmt_vrf": false
        },
        "match_model[0:6]": "EX4400"
      }
    ]
  },
  "switch_mgmt": {
    "config_revert_timer": 10,
    "root_password": "juniper123",
    "protect_re": {
```

```
      "enabled": false
    },
    "tacacs": {
      "enabled": false
    }
  },
  "radius_config": {
    "auth_servers": [],
    "acct_servers": [],
    "auth_servers_timeout": 5,
    "auth_servers_retries": 3,
    "fast_dot1x_timers": false,
    "acct_interim_interval": 0,
    "auth_server_selection": "ordered",
    "coa_enabled": false,
    "coa_port": ""
  },
  "vrf_config": {
    "enabled": false
  },
  "remote_syslog": {
    "enabled": false
  },
  "snmp_config": {
    "enabled": false
  },
  "dhcp_snooping": {
    "enabled": false
  },
  "acl_policies": [],
  "mist_nac": {
    "enabled": true,
    "network": null
  },
  "name": "campus-fabric"
}
```

Topology

Juniper Mist Wired Assurance provides the template for LAN and loopback IP addressing for each collapsed core device once the device's management IP address is reachable. Each device is provisioned with a /32 loopback address and /31 point-to-point interfaces that interconnect collapsed core devices within the campus fabric. Devices such as the access layer switches connect to the distribution layer using standard LAGs; while the collapsed core uses ESI-LAGs in a multihoming, load balancing manner.

The WAN router can be provisioned through the portal but is separate from the campus fabric workflow. The WAN router has a southbound LAG configured to connect to the ESI-LAG on the core switches. WAN routers can be standalone or built as a high availability cluster. In this document, a single SRX router is used as the WAN router.

NOTE: There is a JVD extension available covering more details on WAN router integration especially for production-grade installations. What is shown here is a quick method that has known limits not feasible for production usage.

Create the Campus Fabric

1. From **Organization** on the left-hand side of the portal, select **Campus Fabric**.

Figure 15: Campus Fabric creation

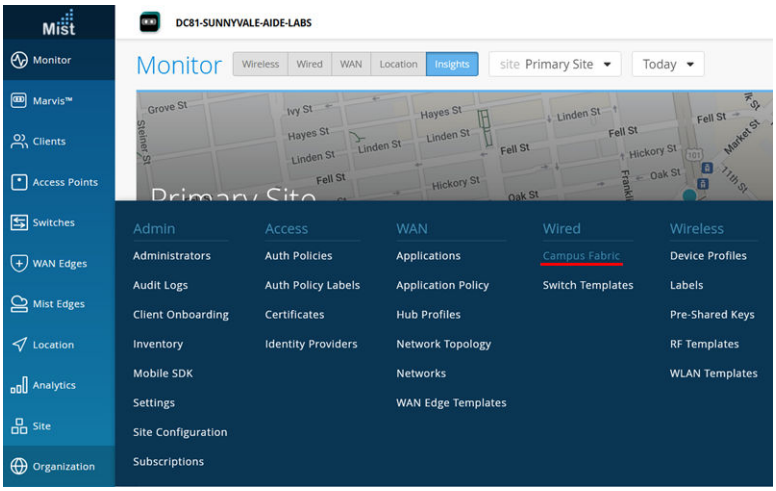

Mist provides the option of deploying a campus fabric at the organizational or site level noted in the upper-left side of the campus fabric menu shown below. Both designs now enable you to build

fabrics with just a single PoD or multiple PoDs based on customer requirements to connect multiple buildings. In EVPN multihoming fabrics, only site-level deployments can be made without PoDs.

In the example shown here, the fabric was built at the site level:

Figure 16: Fabric Site Level Creation

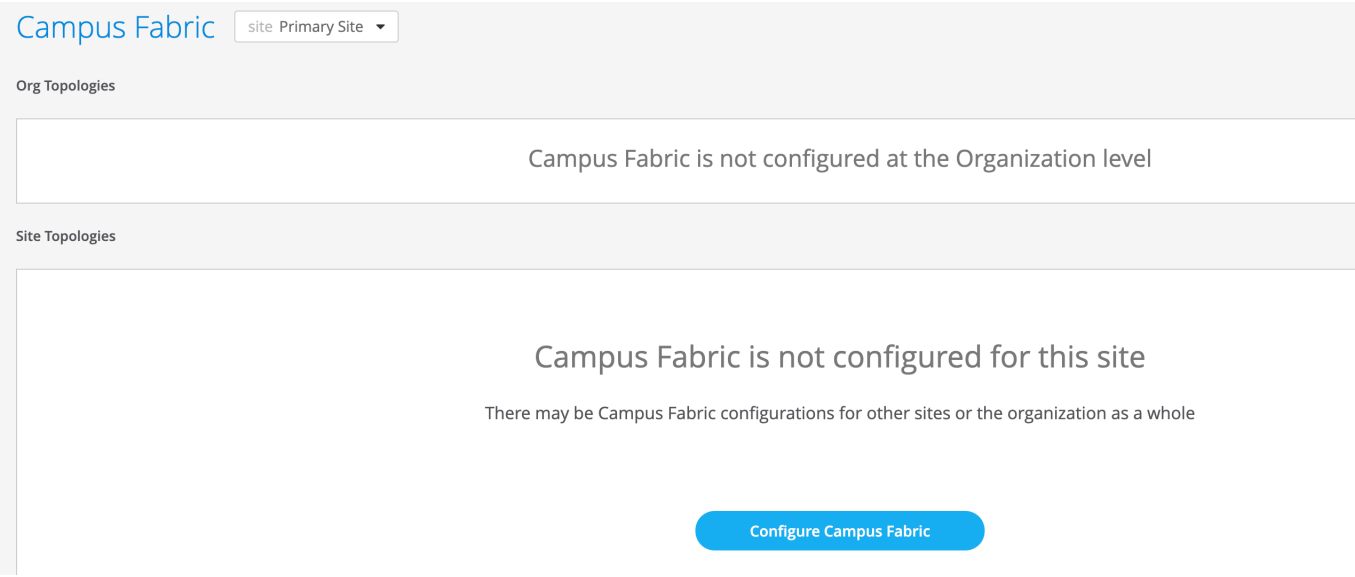

Choose the Campus Fabric Topology

2. Select the campus fabric EVPN Multihoming option below:

Figure 17: EVPN Multihoming Fabric Creation

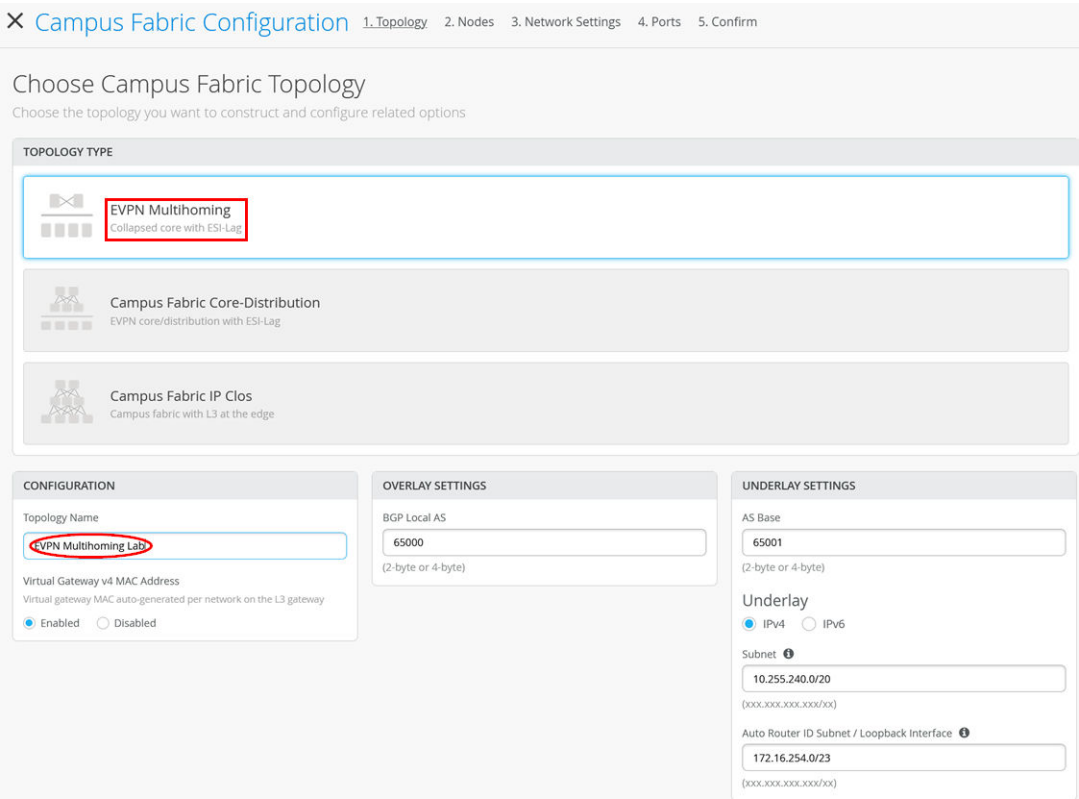

Mist provides a section to name the campus fabric EVPN Multihoming:

- Configuration—Provide a name in accordance with company standards.
- Virtual Gateway v4 MAC Address—Here, you can configure the global MAC address of the gateway used for all VLANs (the assigned default). Or, the MAC address can be unique per VLAN (which is preferred when troubleshooting).

Topology Settings

- BGP Local AS—The BGP AS number used for all control plane interactions.
- AS Base—Represents the starting point of private BGP AS numbers that are automatically allocated per collapsed core device. You can use whatever private BGP AS number range suits your deployment. The routing policy provisioned by Juniper Mist ensures the AS numbers are never advertised outside of the fabric.
- Subnet—Represents the pool of IP addresses used for P2P links between devices. You can use whatever range suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. This number can be modified to suit the specific deployment scale. For example, /24 provides up to 128 P2P /31 subnets.

- Auto Router ID Subnet—Represents the pool of IP addresses associated with each device's loopback address. Each device will automatically get a loopback IP address of /32 assigned from this pool. You can use whatever range suits your deployment. VXLAN tunnelling using a VTEP is associated with this address. The loopback IP addresses assigned here are only visible in the underlay transport network. The definition of these underlay loopback IP addresses is critical for the operation of the EVPN-VXLAN fabric to function at all.

> **NOTE**: We recommend default settings for all options unless it conflicts with other networks attached to the campus fabric. The P2P links between each layer utilize /31 addressing to conserve addresses.

Select Campus Fabric Nodes

3. Select devices to participate in each layer of the Campus Fabric EVPN Multihoming. We recommend that you validate each device's presence in the site's switch inventory prior to the creation of the campus fabric.

   The next step is to assign the switches to the layers. Since the switches were named relative to target layer functionality, they can be quickly assigned to their roles.

Figure 18: Select the Fabric Nodes

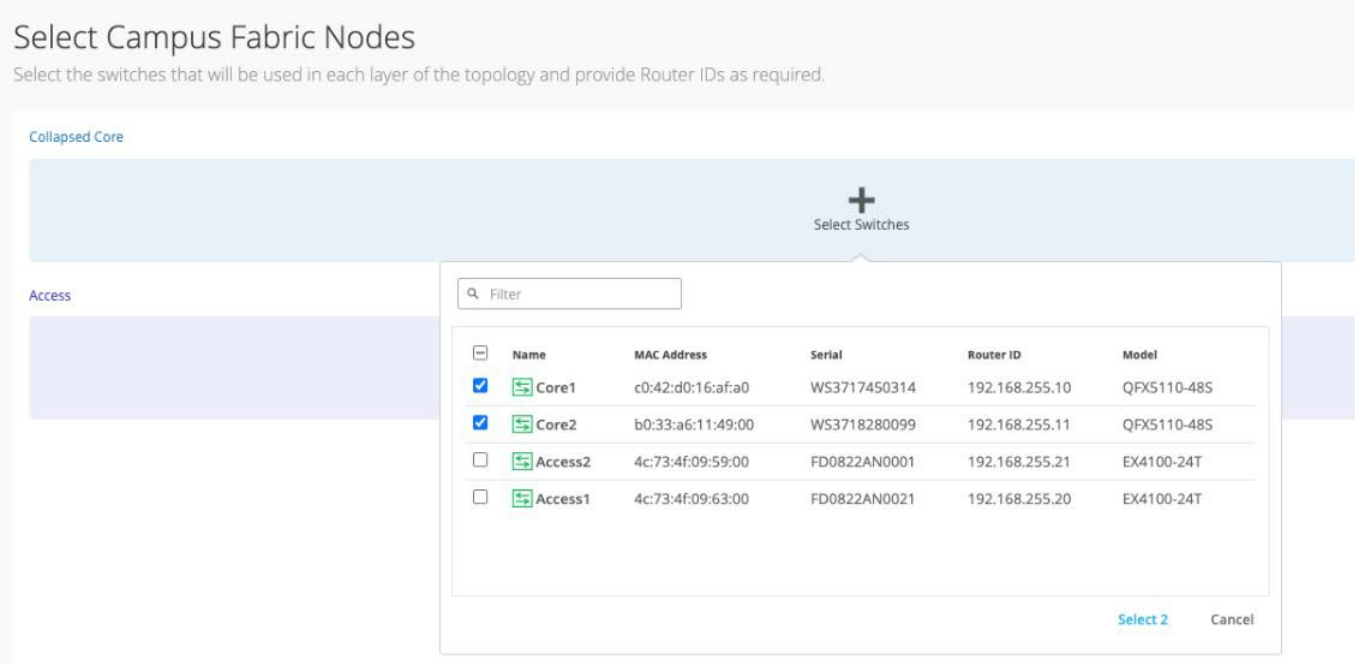

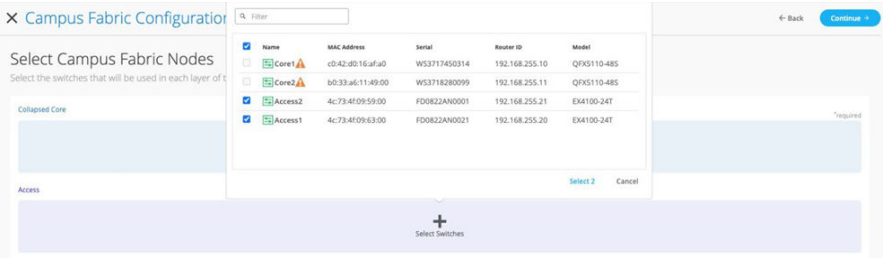

4. Once all devices have been assigned to the appropriate layers, you must provide an underlay loopback IP address for each device (with the exception of the access switches). This loopback interface is associated with a logical construct called a VTEP; used to source the VXLAN tunnel. The campus fabric EVPN multihoming has VTEPs for VXLAN tunnelling on the collapsed core switches.

   When defining an auto router ID subnet prefix, the underlay loopback IP address and router ID assignments happen automatically. There is no need to manually assign them. Make use of this best practice.

Configuring Networks

5. Enter the network information such as VLANs and VRF options. VLANs are mapped to VNIs and can optionally be mapped to VRFs to provide a way to logically separate traffic such as IoT device traffic from Corp IT traffic.

Figure 19: Configure Networks

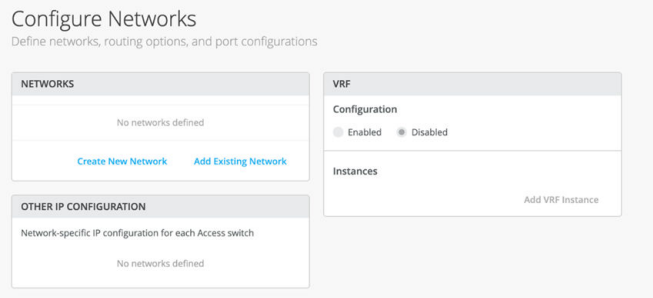

Networks

6. VLANs can be created or imported under this section including the IP subnet and default gateway per each VLAN.

   The **Shared Elements** section of the campus fabric template includes the networks section mentioned above where VLANs are created.

Figure 20: Networks inherited by Switch Template

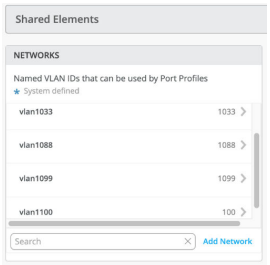

7. Back to the campus fabric build, select the existing template which includes Layer 2 VLAN information. All VLAN and IP information is inherited from the template.

Figure 21: Network import from Template

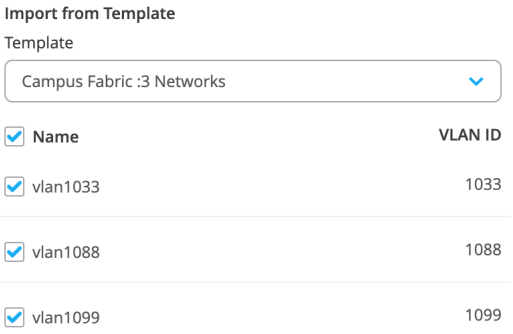

Networks can be edited, newly added, or added from an existing template:

Figure 22: Edit a Network

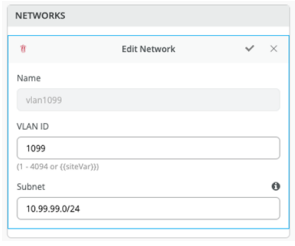

For each network, add the information of the subnet and virtual gateways following the examples below:

Figure 23: Network 1099 and VGA

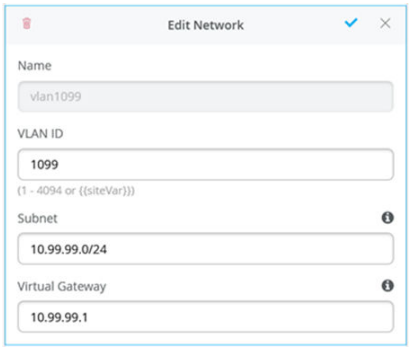

Figure 24: Network 1088 and VGA

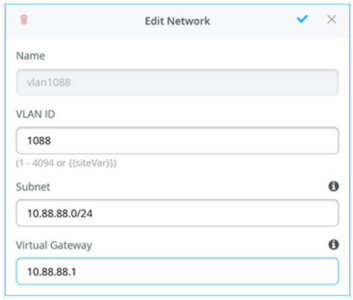

Figure 25: Network 1033 and VGA

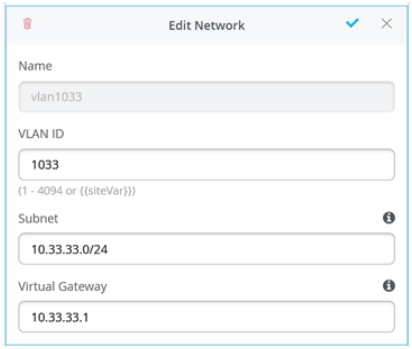

Other IP Configuration

Juniper Mist Wired Assurance provides automatic IP addressing for IRB interfaces for each of the VLANs. Port profiles and port configurations then associate the VLAN with specified ports. In this case, we selected campus fabric EVPN multihoming at the onset of the campus fabric build. This option uses virtual gateway addressing for all devices participating in the Layer 3 subnet. The Core1 and Core2 switches are configured with a shared IP address for each Layer 3 subnet. This address is

shared amongst both core switches and acts as the default gateway for all devices within the VLAN. Each core device also receives a unique IP address chosen by Juniper Mist. All addresses can be managed per customer requirements. Juniper Mist assigns IP addresses for Core1 and 2 starting at the beginning of each subnet and the end user can modify these IP addresses accordingly. For example, this deployment uses x.x.x.1 as a default gateway for each VLAN and x.x.x.254 as the gateway of last resort (an MX router in this case) for all traffic leaving the VLAN. Therefore, we modify the IP addresses assigned to Core1 from x.x.x.1 to x.x.x.3 allowing the virtual gateway to use x.x.x.1 for all VLANs.

Figure 26: Core1 Static-IP of Overlay VLAN Used

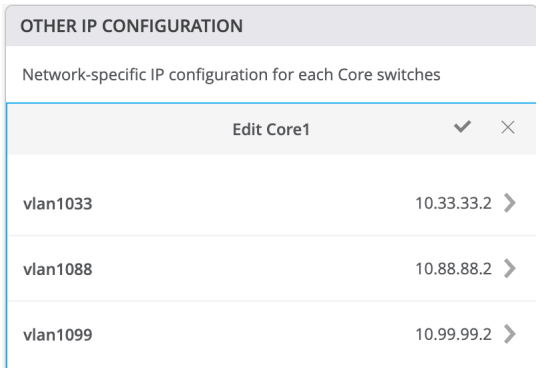

Figure 27: Core2 Static-IP of Overlay VLAN Used

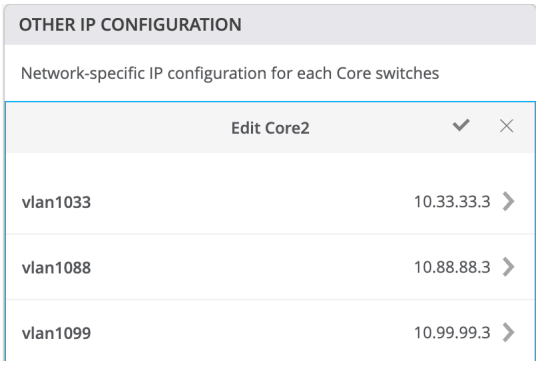

By default, all VLANs are placed in the default VRF. The VRF option allows you to group common VLANs into the same VRF or separate VRFs depending on traffic isolation requirements. This example includes three VRFs or routing instances: corp-it, developers, and guest-wifi.

8. Here, you build the first corp-it VRF and select the pre-defined vlan 1099.

Figure 28: Enable VRF

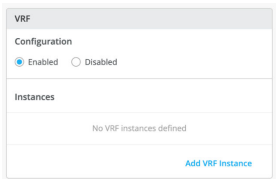

Figure 29: Assign Network to VRF

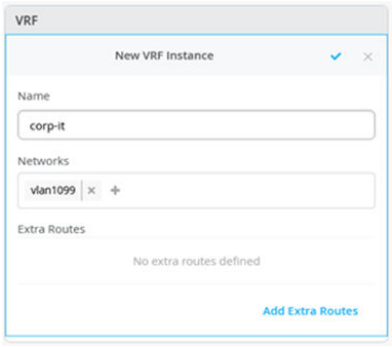

By default, inter-VRF communications are not supported within the campus fabric. If inter-VRF communications is required, each VRF can include extra routes such as a default route that instructs the campus fabric to use an external router or firewall for further security inspection or routing capabilities. In this example, all traffic is trunked over the ESI-LAG and the SRX firewall handles inter-VRF routing. See Figure 11: Topology on page 23.

Notice the SRX firewall participates in the VLANs defined within the campus fabric and is the gateway of last resort for all traffic leaving the subnet.

9. Select the **Add Extra Routes** option to inform Juniper Mist to forward all traffic leaving 10.99.99.0/24 to use the next hop of the MX router: 10.99.99.254.

Figure 30: Add default route

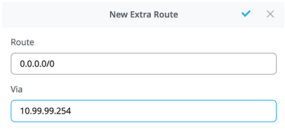

10. Create two additional VRFs:

    a. The developers VRF using vlan 1088 with 0.0.0.0/0 utilizing 10.88.88.254

    b. The guest-wifi VRF using vlan 1033 with 0.0.0.0/0 utilizing 10.33.33.254

Figure 31: Entire Network and VRF Configuration

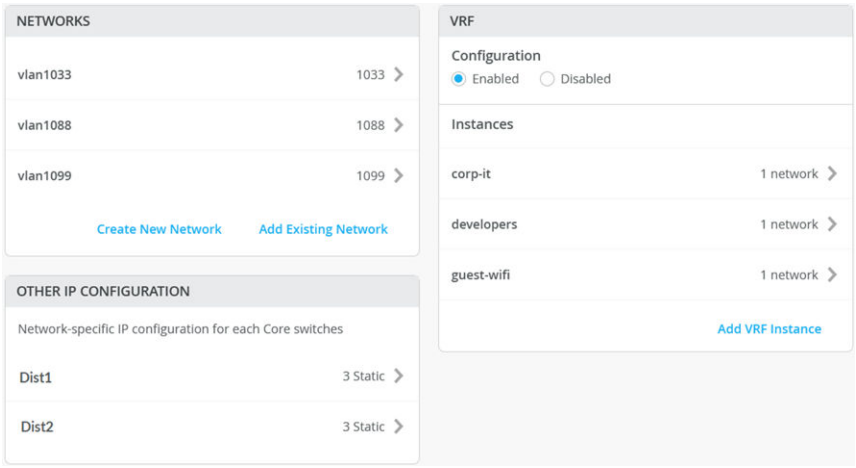

11. As a next step, you need to provide a name such as “fabric-lag” that the fabric will use to establish the redundant LAG interfaces between all access and collapsed core switches. All created VLANs should be automatically added already as future trunk networks.

Figure 32: Fabric LAG Configuration

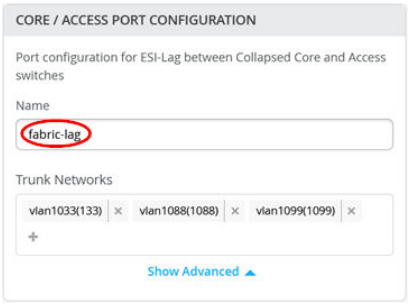

12. The section configures the active-active ESI-LAG trunks between distribution and access switches. Here, we name the port configuration and include VLANs associated with this configuration. The advanced tab provides additional configuration options:

Figure 33: Fabric LAG

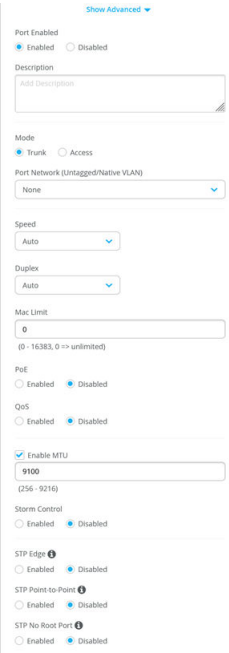

NOTE: We recommend default settings unless specific requirements are needed.

13. Now that all VLANs are configured and assigned to each VRF, and the distribution and access ESI-LAGs have been built, click the **Continue** button in the upper-right corner of the portal to move to the next step.

Configure Campus Fabric Ports

The final step is the selection of physical ports among core and access switches.

Figure 34: Port Overview

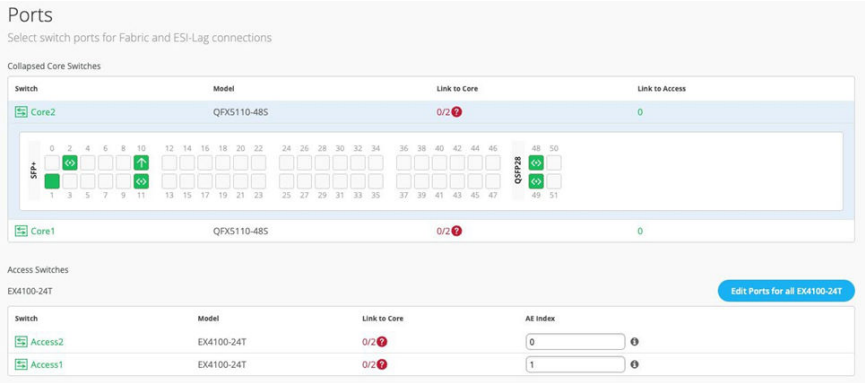

NOTE: To ensure accuracy, we recommend that you run the CLI command “show lldp neighbors” on both collapsed core switches prior to this step in the deployment process.

Collapsed Core Switches

Core1:

14. We are now ready to select the ports that interconnect the collapsed core switches. You must select **et-0/0/48** as a Collapsed core link then choose **Link1.**

Figure 35: First Link core1

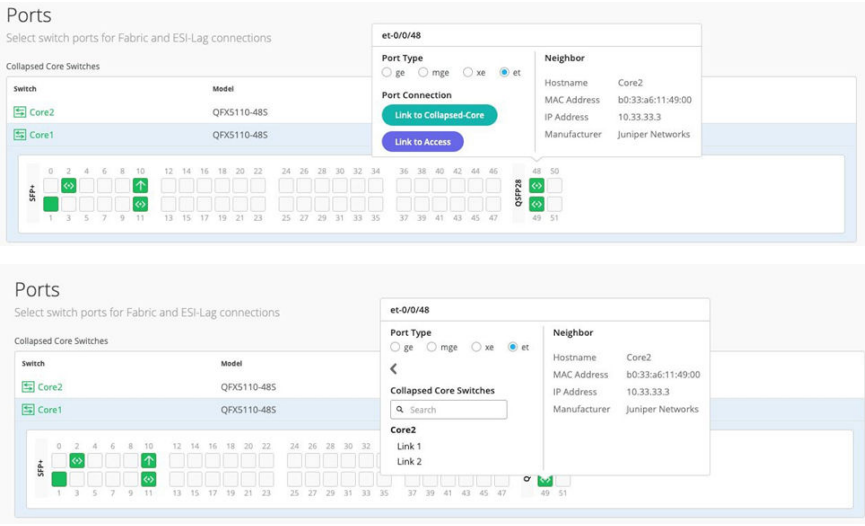

15. Core1 second Link. You must select **et-0/0/49** as a Collapsed Core link then choose **Link2**.

Figure 36: Second Link core1

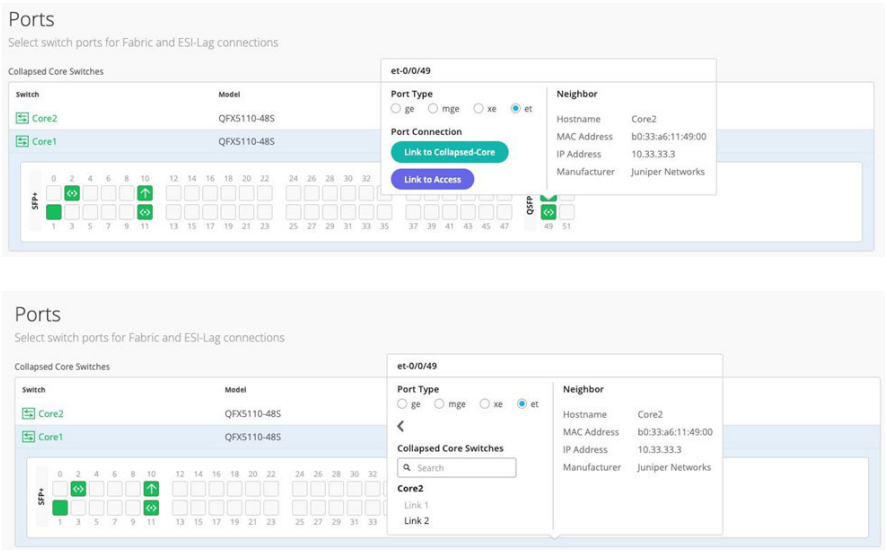

16. Core1 first link to Access1. You can select **xe-0/0/1** as a link to Access1.

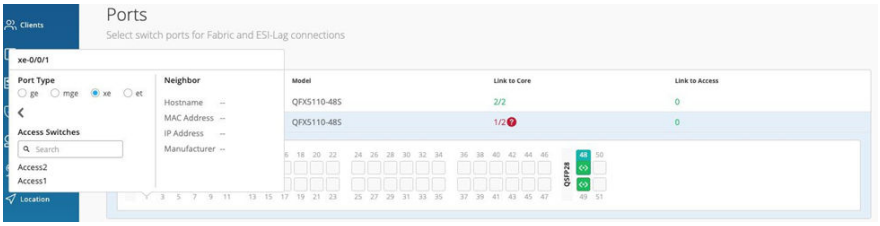

17. This is Core1's second link to Access2. You can select **xe-0/0/2** as a link to Access2.

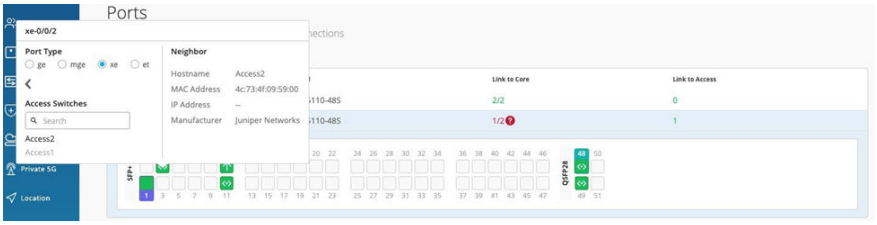

Core2:

18. This is Core2's first link. You must select **et-0/0/48** as a collapsed core link and then choose **Link1**.

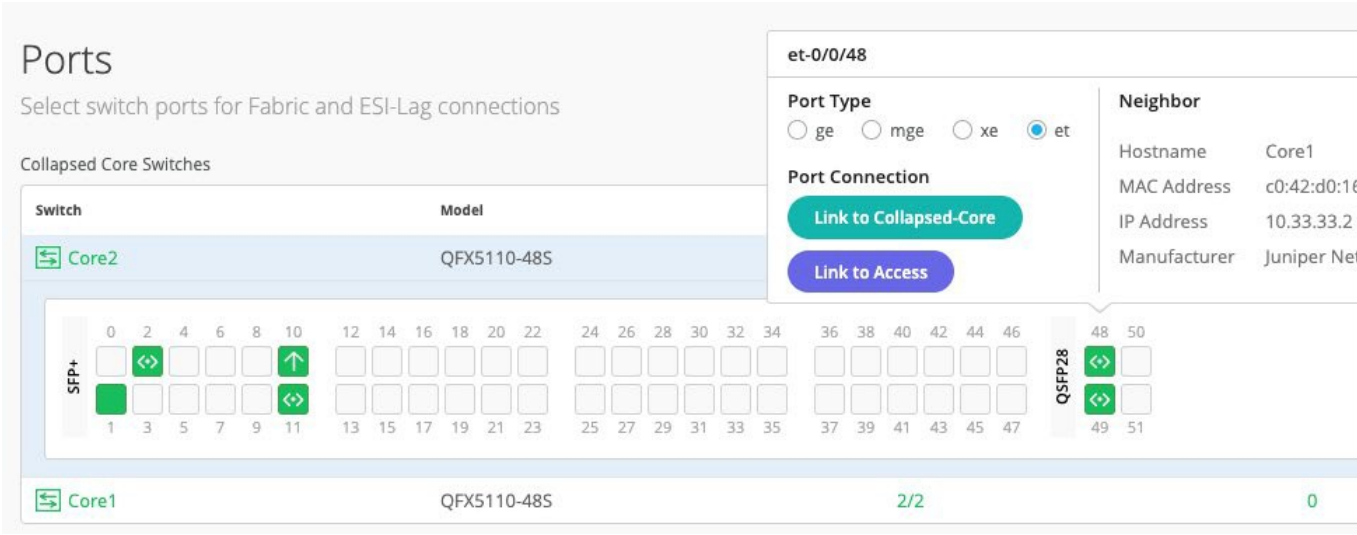

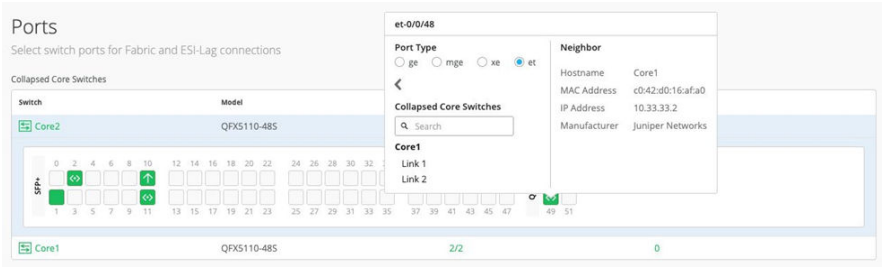

19. This is Core2's second link. You must select **et-0/0/49** as a collapsed core link and then choose **Link2**.

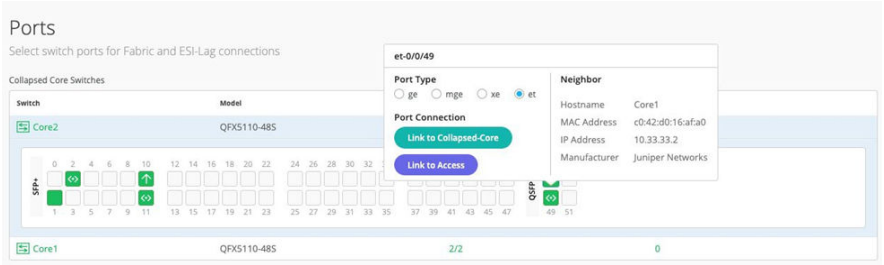

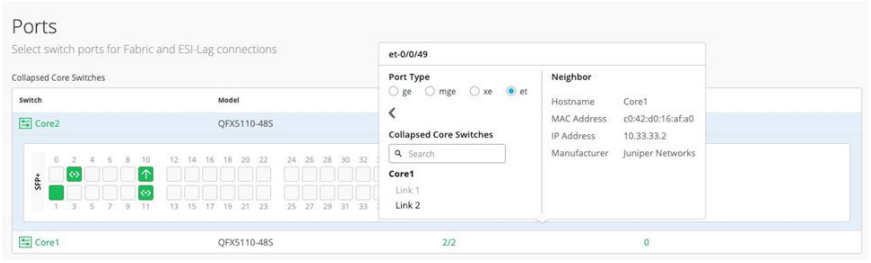

20. This is Core2's first link to Access1. You can select **xe-0/0/1** as a link to Access1.

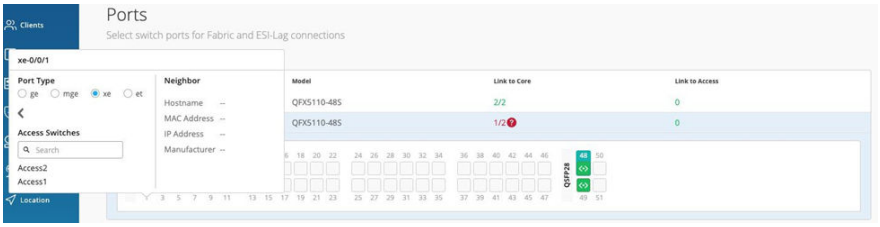

21. This is Core2's first link to Access2. You can select **xe-0/0/2** as a link to Access2.

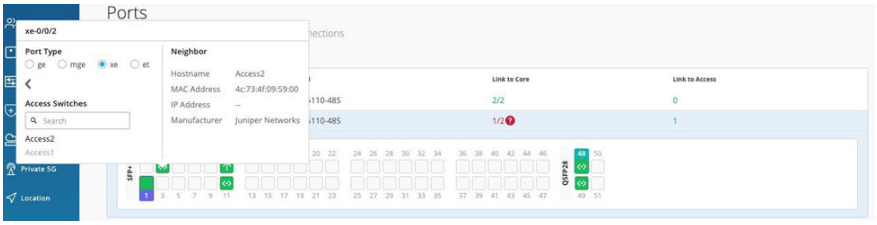

Access Switches

You can now select the ports that interconnect with the collapsed core switches.

Select both uplinks and interface speed, while allowing Juniper Mist to define each AE index. In this case, uplinks xe-0/2/0, xe-0/2/3 are selected as links to the core on both access switches and AE Index 0/1 (system default numbering) on Access2 and 1 respectively.

Access1:

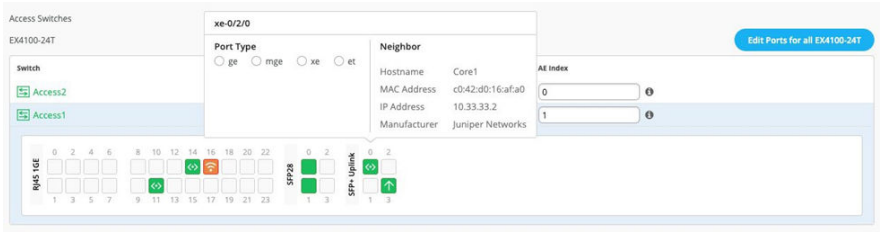

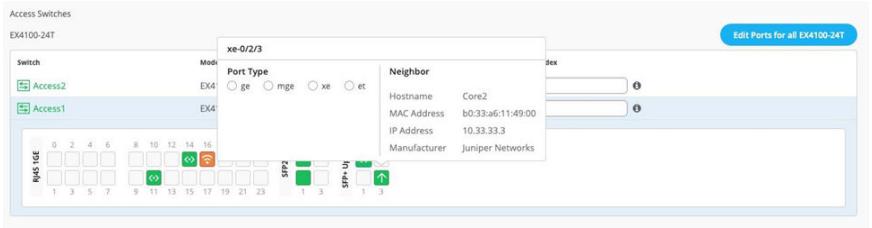

Access2:

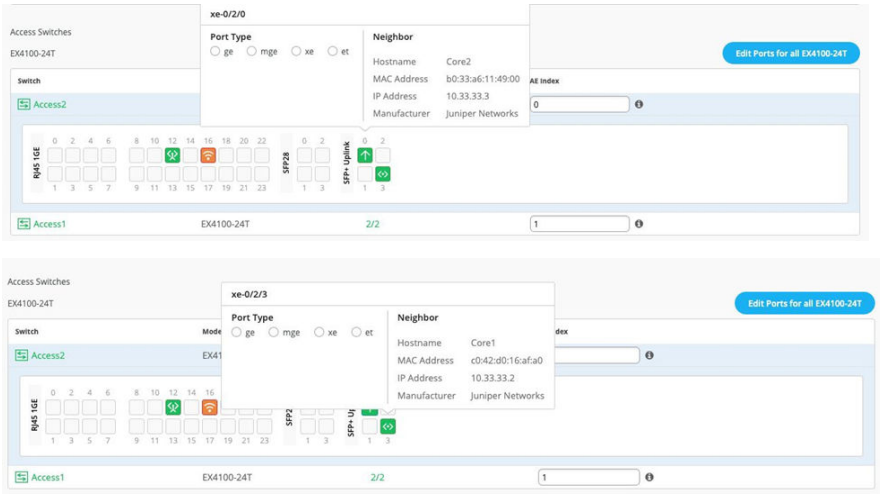

After you select all requisite port combinations, click the **Continue** button in the upper-right corner of the portal.

Campus Fabric Configuration Confirmation

This last section provides the ability to confirm each device's configuration as shown below:

Figure 37: Fabric Confirmation View

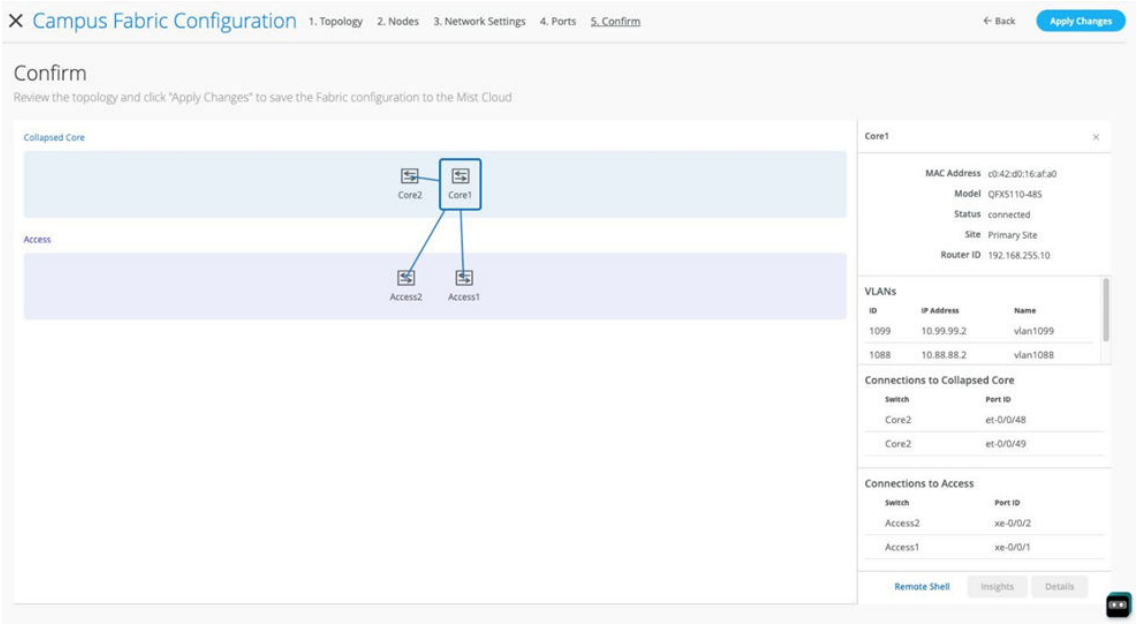

NOTE: As we have configured the usage of auto router ID subnet, the underlay loopback IP addresses may still not be assigned on this page, and warnings may appear like the ones

shown above. Ignore this for now as the assignments happen when you apply the configuration for the first time.

22. Once you have completed verification, select the **Apply Changes** option in the upper-right corner of the portal.

Figure 38: Apply Changes to Fabric

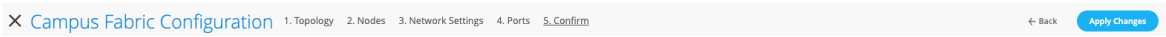

You must complete the second stage confirmation to create the fabric.

Juniper Mist displays the following banner including the estimated time for the campus fabric to be built. The process includes the following:

- Juniper Mist builds the P2P interfaces between distribution and core devices with IP addresses chosen from the range presented at the start of the build.
- Configuring each device with a loopback address from the range presented at the start of the build.
- eBGP is provisioned on each device with unique BGP autonomous system numbers. The primary goal of the underlay is to leverage ECMP for load balancing traffic on a per-packet level for device loopback reachability. The primary goal of the eBGP overlay is for the transport of customer traffic using EVPN-VXLAN.
- Applying IP addresses on each Layer 3 gateway IRB located on Dist1 and Dist2.
- Applying IP addresses on each loopback interface, which is done automatically in this case.
- Configuring routing policies for underlay and overlay connectivity.
- Optimizing MTU settings for P2P underlay, Layer 3 IRB, and ESI-LAG bundles.
- VXLAN-to-VLAN mapping using VNI addresses that are automatically assigned.
- Creating VRFs for corp-it, developers, and guest-wifi and the VLANs associated with each VRF.
- Creating VXLAN tunnels between distribution devices and distribution-core devices (in support of the northbound MX router that is configured in subsequent steps).
- Creating a downloadable connection table (CSV format) that can be used by those involved in the physical buildout of the campus fabric.
- Displaying a graphical interface depicting all devices with BGP peering and the status of physical links.

Figure 39: Applying Changes

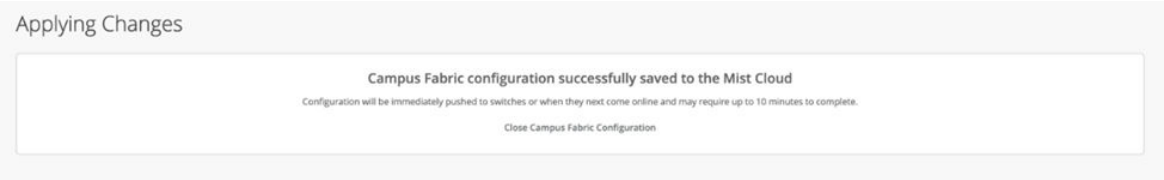

23. Once you click **Close Campus Fabric Configuration**, you can view a summary of the newly created campus fabric EVPN multihoming.

Figure 40: Created EVPN Multihoming Fabric View

With Juniper Mist Wired Assurance, you can download a connection table (CSV format) representing the physical layout of the campus fabric. This can be used to validate all switch interconnects for those participating in the physical campus fabric build. Once the campus fabric is built or in the process of being built, you can download the connection table.

Figure 41: Download Connection Table CSV

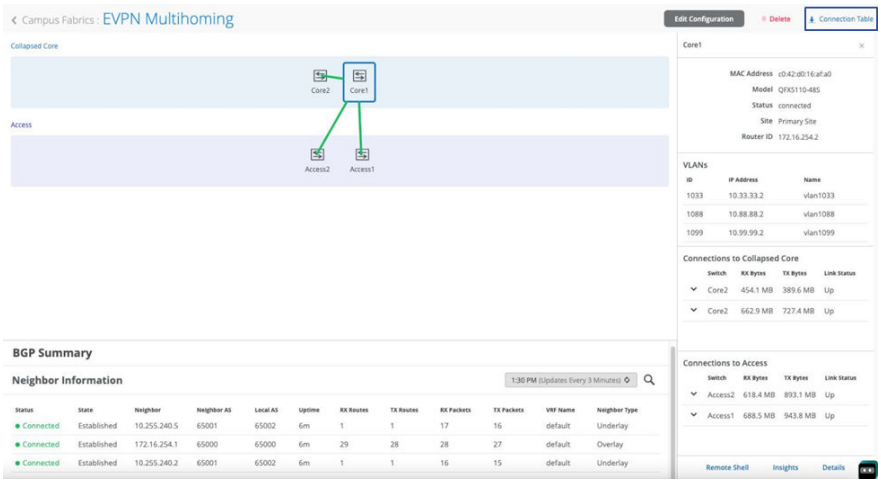

Connection table spreadsheet:

Figure 42: Downloaded Connection Table

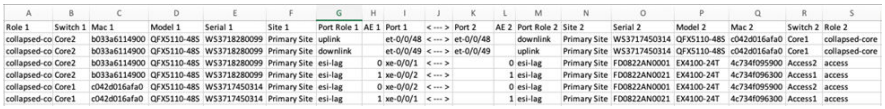

Role 1	Switch 1	Mac 1	Model 1	Serial 1	Site 1	Port Role 1	AE 1	Port 1	< --- >	Port 2	AE 2	Port Role 2	Site 2	Serial 2	Model 2	Mac 2	Switch 2	Role 2
collapsed-co	Core2	b033a6114900	QFX5110-48S	WS3718280099	Primary Site	uplink		et-0/0/48	< --- >	et-0/0/48		downlink	Primary Site	WS3717450314	QFX5110-48S	c042d016afa0	Core1	collapsed-core
collapsed-co	Core2	b033a6114900	QFX5110-48S	WS3718280099	Primary Site	downlink		et-0/0/49	< --- >	et-0/0/49		uplink	Primary Site	WS3717450314	QFX5110-48S	c042d016afa0	Core1	collapsed-core
collapsed-co	Core2	b033a6114900	QFX5110-48S	WS3718280099	Primary Site	esi-lag	0	xe-0/0/1	< --- >		0	esi-lag	Primary Site	FD0822AN0001	EX4100-24T	4c734f095900	Access2	access
collapsed-co	Core2	b033a6114900	QFX5110-48S	WS3718280099	Primary Site	esi-lag	1	xe-0/0/2	< --- >		1	esi-lag	Primary Site	FD0822AN0001	EX4100-24T	4c734f096300	Access1	access
collapsed-co	Core1	c042d016afa0	QFX5110-48S	WS3717450314	Primary Site	esi-lag	0	xe-0/0/2	< --- >		0	esi-lag	Primary Site	FD0822AN0001	EX4100-24T	4c734f095900	Access2	access
collapsed-co	Core1	c042d016afa0	QFX5110-48S	WS3717450314	Primary Site	esi-lag	1	xe-0/0/1	< --- >		1	esi-lag	Primary Site	FD0822AN0001	EX4100-24T	4c734f096300	Access1	access

Apply VLANs to Access Ports

As previously discussed, Juniper Mist provides the ability to templatize well known services such as RADIUS, NTP, DNS, and so on that can be used across all devices within a site. These templates can also include VLANs and port profiles that can be targeted at each device within a site. The last step before verification is to associate VLANs with the requisite ports on each access switch.

In this case, Desktop1 and 2 are associated with different ports on each access switch which requires the configuration to be applied to Access1 and 2, respectively. See Figure 11 on page 23.

It is also noteworthy that Juniper APs connect to the same port on Access1 and 2, allowing the switch template to be customized with this configuration. For example, the following found under the switch template option is customized to associate each switch with its role: core, distribution, and access. Furthermore, all access switches (defined by the Juniper Networks® EX4400 Switch, as an example) associated the AP port profile named “myaccess” with ge-0/0/16 without needing to configure each switch independently.

Figure 43: Port Configuration Through Switch Template

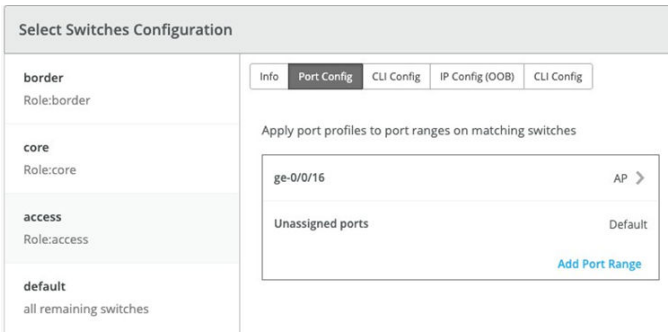

Using Access1 as an example, we apply vlan1099 to port ge-0/0/11 under the **Port Config** section on Access1. In this example, vlan1099 (corp-it), vlan1088 (developers), and vlan1033 (guest-wifi) are defined in the switch template. Here, vlan1099 is selected under the configuration profile:

Figure 44: Assign Port Profile to a Port

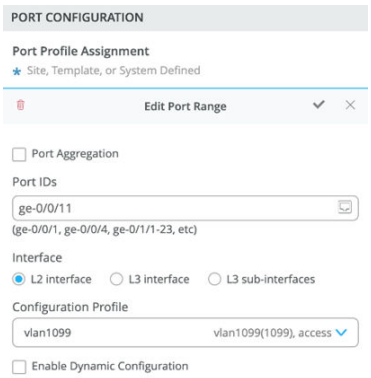

The switch template definition for vlan1099 is shown below, representing attributes associated with VLANs such as dot1x authentication, QoS, and PoE. Vlan1088 and vlan1033 need to be configured in a similar fashion.

Figure 45: Port Profile Example

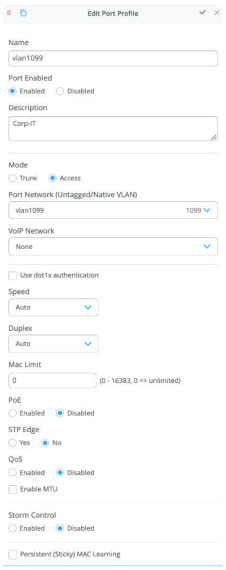

APPENDIX: Fabric Verification (Optional)

IN THIS SECTION

NOTE: You may skip this optional chapter if you want. This information is presented to show more of the internal details on how the fabric is working.

In the following steps, we cover the verification of the campus fabric EVPN multihoming deployment. See Figure 11 on page 23 . Currently, there are two desktops to verify the fabric. Let's take a quick look to see if Desktop1 can connect internally and externally.

Figure 46: Wired Client Connectivity Issue

```
root@desktop1:~# ifconfig vlan1099
vlan1099: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.99.99.99  netmask 255.255.255.0  broadcast 10.99.99.255
        inet6 fe80::5054:ff:fe74:a06f  prefixlen 64  scopeid 0x20<link>
        ether 52:54:00:74:a0:6f  txqueuelen 1000  (Ethernet)
        RX packets 28044  bytes 17108274 (17.1 MB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 26564  bytes 2271495 (2.2 MB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

root@desktop1:~# ip r
default via 10.99.99.1 dev vlan1099
10.99.99.0/24 dev vlan1099 proto kernel scope link src 10.99.99.99
192.168.10.0/24 dev ens3 proto kernel scope link src 192.168.10.61
root@desktop1:~# ping 10.99.99.1 -c 2
PING 10.99.99.1 (10.99.99.1) 56(84) bytes of data.
64 bytes from 10.99.99.1: icmp_seq=1 ttl=64 time=6.45 ms
64 bytes from 10.99.99.1: icmp_seq=2 ttl=64 time=8.86 ms

--- 10.99.99.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 6.452/7.653/8.855/1.201 ms
root@desktop1:~# ping 10.99.99.254 -c 2
PING 10.99.99.254 (10.99.99.254) 56(84) bytes of data.
From 10.99.99.99 icmp_seq=1 Destination Host Unreachable
From 10.99.99.99 icmp_seq=2 Destination Host Unreachable

--- 10.99.99.254 ping statistics ---
2 packets transmitted, 0 received, +2 errors, 100% packet loss, time 1016ms
```

Verification steps:

- Confirmed local IP address, VLAN, and default gateway were configured on Desktop1.
- Can ping default gateway – indicates that we can reach the distribution switch.

- Ping to WAN router failed (10.99.99.254) – we need to troubleshoot.

Start by verifying the campus fabric in the portal by selecting the **Campus Fabric** option under the **Organization** tab on the left side of the portal.

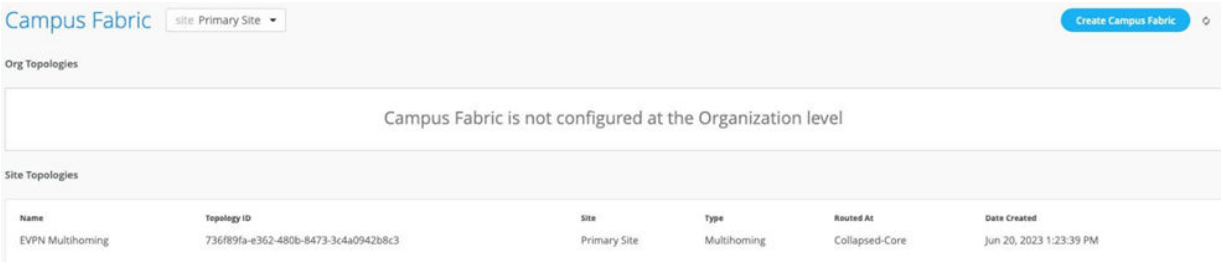

Accessing each device within the campus fabric through remote shell is supported here as well as a visual representation of the following capabilities:

- BGP peering establishment.
- Transmit and receive traffic on a link-by-link basis.
- Telemetry, such as LLDP, from each device that verifies the physical build.

Figure 47: Fabric Health

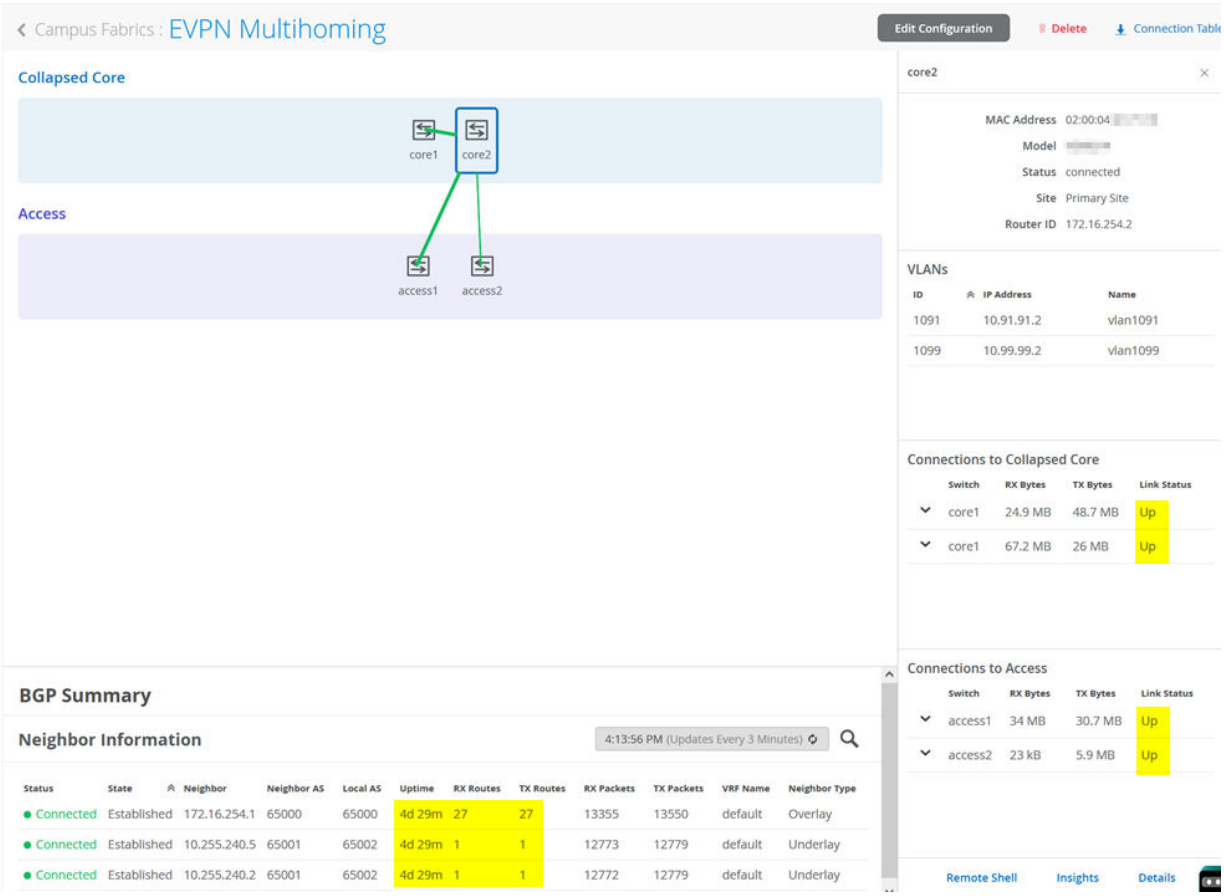

BGP Underlay

Purpose

Verifying the state of eBGP between the collapsed core layers is essential for EVPN-VXLAN to operate as expected. This network of P2P links between each layer supports:

- Load balancing using ECMP for greater resiliency and bandwidth efficiencies.
- BFD to decrease convergence times during failures.
- Loopback reachability to support VXLAN tunnelling.

Due to the automated assignment of loopback IP addresses, for this fabric, we have the following configuration to remember:

Switch Type	Switch Name	Auto assigned Loopback IP
Collapsed Core	Core1	172.16.254.2
Collapsed Core	Core2	172.16.254.1
Access	Access1	N/A
Access	Access2	N/A

Action

Verify that BGP sessions are established between core devices and distribution devices to ensure loopback reachability, BFD session status, and load-balancing using ECMP.

NOTE: Operational data can be gathered through the campus fabric section of the portal using remote shell or using an external application such as SecureCRT or Putty.

Verification of BGP Peering

Core1:

Access the remote shell through the lower-right of the campus fabric, from the switch view, or through Secure Shell (SSH).

Figure 48: show bgp summary on core1

```
mist@Core1> show bgp summary
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       2          2          0          0          0          0
bgp.evpn.0
                      28         28          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/Accepted/Damped.
..
10.255.240.2          65001         38         37       0       0       15:53 Establ
  inet.0: 1/1/1/0
10.255.240.5          65001         39         38       0       0       15:53 Establ
  inet.0: 1/1/1/0
172.16.254.1          65000         52         49       0       0       15:45 Establ
  bgp.evpn.0: 28/28/28/0
  default-switch.evpn.0: 26/26/26/0
  __default_evpn__.evpn.0: 2/2/2/0

{master:0}
mist@Core1>
```

From the BGP summary we can see that the underlay (10.255.240.x) peer relationships are established. This means that the underlay links are attached to the correct devices and the links are up.

It also shows the overlay (172.16.254.x) relationship is established with Core2 and that it is peering at the correct loopback addresses. This demonstrates loopback reachability.

We can also see routes received and time shown when the sessions were established are roughly equal which looks good so far.

If BGP is not established, you can validate the underlay links and addressing, and that the loopback addresses are correct. Loopback addresses should be pingable from other loopback addresses.

The primary goal of eBGP in the underlay is to provide loopback reachability between core switches in an EVPN multihoming deployment. This loopback is used to terminate VXLAN tunnels between devices. The following shows loopback reachability from Core1 to Core2 in the fabric:

Figure 49: Testing Underlay Loopback IP Reachability

```
mist@Core1> ping 172.16.254.1
PING 172.16.254.1 (172.16.254.1): 56 data bytes
64 bytes from 172.16.254.1: icmp_seq=0 ttl=64 time=9.518 ms
64 bytes from 172.16.254.1: icmp_seq=1 ttl=64 time=10.470 ms
64 bytes from 172.16.254.1: icmp_seq=2 ttl=64 time=9.668 ms
^C
--- 172.16.254.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 9.518/9.885/10.470/0.418 ms

{master:0}
mist@Core1>
```

NOTE: eBGP sessions are established between core-distribution layers in the campus fabric. Loopback reachability has also been verified between collapsed core devices.

Let's verify that the routes are established to the collapsed core across multiple paths.

Core1: ECMP Loopback reachability with Core2

Figure 50: Loopback Reachability to Core2

```
mist@Core1> show route forwarding-table destination 172.16.254.1
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.16.254.1/32    user     1                    ulst   131070     4
                            10.255.240.5         ucst     1740     4 et-0/0/48.0
                            10.255.240.2         ucst     1741     4 et-0/0/49.0
```

Core2: ECMP Loopback reachability with Core1

Figure 51: Loopback Reachability to Core1

```
mist@Core2> show route forwarding-table destination 172.16.254.2
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.16.254.2/32    user     1                    ulst   131070     4
                              10.255.240.4       ucst     1737     4 et-0/0/48.0
                              10.255.240.3       ucst     1738     4 et-0/0/49.0
```

Finally, we validate BFD for fast convergence in the case of a link or device failure:

Figure 52: BFD Testing

```
mist@Core2> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.255.240.3             Up        et-0/0/49.0    1.050     0.350        3
10.255.240.4             Up        et-0/0/48.0    1.050     0.350        3
172.16.254.2             Up                       3.000     1.000        3

3 sessions, 3 clients
Cumulative transmit rate 6.7 pps, cumulative receive rate 6.7 pps

{master:0}
mist@Core2>
```

Conclusion: At this point, the BGP underlay and overlay are operational through the verification of eBGP between corresponding layers of the campus fabric and loopback routes are established between collapsed core devices.

EVPN-VXLAN Verification Between Collapsed Core Switches

Since the desktop can ping its default gateway, we can assume the Ethernet switching tables are correctly populated, and VLAN and interface modes are correct. If pinging the default gateway failed, then try troubleshooting the underlay connectivity.

Verification of the EVPN Database on Both Core Switches

Core1:

Figure 53: EVPN DB core1

```
mist@Core1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:24:07
     10001      c0:42:d0:16:af:a0  irb.0                          Jun 20 17:23:45
     11033      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:19:00  Jun 20 17:24:07  10.33.33.1
     11033      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:24:07  10.33.33.3
     11033      c0:42:d0:16:af:a0  irb.1033                       Jun 20 17:23:55  10.33.33.2
     11088      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:50:00  Jun 20 17:24:07  10.88.88.1
     11088      52:54:00:91:ed:5c  00:11:00:00:00:01:00:01:02:00  Jun 20 17:49:04  10.88.88.88
     11088      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:24:07  10.88.88.3
     11088      c0:42:d0:16:af:a0  irb.1088                       Jun 20 17:23:55  10.88.88.2
     11099      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:5b:00  Jun 20 17:24:07  10.99.99.1
     11099      52:54:00:a4:c5:73  00:11:00:00:00:01:00:01:02:01  Jun 20 17:49:14  10.99.99.99
     11099      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:24:07  10.99.99.3
     11099      c0:42:d0:16:af:a0  irb.1099                       Jun 20 17:23:55  10.99.99.2

{master:0}
mist@Core1>
```

Core2:

Figure 54: EVPN DB core2

```
mist@Core2> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      b0:33:a6:11:49:00  irb.0                          Jun 20 17:23:45
     10001      c0:42:d0:16:af:a0  172.16.254.2                   Jun 20 17:24:07
     11033      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:19:00  Jun 20 17:24:07  10.33.33.1
     11033      b0:33:a6:11:49:00  irb.1033                       Jun 20 17:23:55  10.33.33.3
     11033      c0:42:d0:16:af:a0  172.16.254.2                   Jun 20 17:24:07  10.33.33.2
     11088      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:50:00  Jun 20 17:24:07  10.88.88.1
     11088      52:54:00:91:ed:5c  00:11:00:00:00:01:00:01:02:00  Jun 20 17:49:04  10.88.88.88
     11088      b0:33:a6:11:49:00  irb.1088                       Jun 20 17:23:55  10.88.88.3
     11088      c0:42:d0:16:af:a0  172.16.254.2                   Jun 20 17:24:07  10.88.88.2
     11099      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:5b:00  Jun 20 17:24:07  10.99.99.1
     11099      52:54:00:a4:c5:73  00:11:00:00:00:01:00:01:02:01  Jun 20 17:49:13  10.99.99.99
     11099      b0:33:a6:11:49:00  irb.1099                       Jun 20 17:23:55  10.99.99.3
     11099      c0:42:d0:16:af:a0  172.16.254.2                   Jun 20 17:24:07  10.99.99.2

{master:0}
mist@Core2>
```

Both core switches have identical EVPN databases, which is expected. Note that the entries for desktop1 (10.99.99.99) and desktop2 (10.88.88.88) are present on each core switch. These entries are learned through the campus fabric from the ESI-LAGs to each access switch. For example, Desktop1 (10.99.99.99) is associated with shared ESI 10-digit segment between Core1 and 2 facing Access1 and is associated with a VNI of 11099. The fact that we see both Desktop ARP and associated ESI 10-digit segment entries leans towards an issue between the core and the Juniper Networks® SRX Series Firewall. Remember, the SRX Series Firewall is responsible for routing traffic between routing-instances; in this case, between corp-it, developers, and guest-wifi.

Verification of VXLAN Tunnelling Between Collapsed Core Switches

Core1:

Figure 55: vtep remote on core1

```
mist@Core1> show ethernet-switching vxlan-tunnel-end-point remote summary
Logical System Name       Id  SVTEP-IP          IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   172.16.254.2      lo0.0    0
 RVTEP-IP         L2-RTT                IFL-Idx   Interface    NH-Id   RVTEP-Mode  ELP-IP        Flags
 172.16.254.1     default-switch        828       vtep.32769   1760    RNVE

{master:0}
mist@Core1>
```

Core2:

Figure 56: vtep remote on core2

```
mist@Core2> show ethernet-switching vxlan-tunnel-end-point remote summary
Logical System Name       Id  SVTEP-IP          IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   172.16.254.1      lo0.0    0
 RVTEP-IP         L2-RTT                IFL-Idx   Interface    NH-Id   RVTEP-Mode  ELP-IP        Flags
 172.16.254.2     default-switch        821       vtep.32769   1748    RNVE

{master:0}
mist@Core2>
```

NOTE: The EVPN database is confirmed on both core devices and VXLAN tunnels are established between core switches. We have also verified that Desktop1 and Desktop2 are present in both core switches' EVPN databases.

Core1: Ethernet Switching and ARP Tables

```
mist@Core1> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 11 entries, 11 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical              SVLBNH/      Active
   name                 address             flags    interface            VENH Index   source
   default              00:cc:34:f2:ec:80   D        ge-0/0/11.0
   default              00:cc:34:f2:ec:84   D        ge-0/0/11.0
   default              b0:33:a6:11:49:00   DRP      vtep.32769                        172.16.254.1
   vlan1033             00:00:5e:00:01:01   DR       esi.1850             1760         05:00:00:fd:e8:00:00:2b:19:00
   vlan1033             b0:33:a6:11:49:00   DRP      vtep.32769                        172.16.254.1
   vlan1088             00:00:5e:00:01:01   DRP      esi.1848             1760         05:00:00:fd:e8:00:00:2b:50:00
   vlan1088             52:54:00:91:ed:5c   DLR      ae0.0
   vlan1088             b0:33:a6:11:49:00   DRP      vtep.32769                        172.16.254.1
   vlan1099             00:00:5e:00:01:01   DRP      esi.1849             1760         05:00:00:fd:e8:00:00:2b:5b:00
   vlan1099             52:54:00:a4:c5:73   DLR      ae1.0
   vlan1099             b0:33:a6:11:49:00   DRP      vtep.32769                        172.16.254.1

{master:0}
mist@Core1> show arp
MAC Address       Address         Name                      Interface               Flags
b0:33:a6:11:49:00 10.33.33.3      10.33.33.3                irb.1033 [vtep.32769]   permanent remote
b0:33:a6:11:49:00 10.88.88.3      10.88.88.3                irb.1088 [vtep.32769]   permanent remote
52:54:00:91:ed:5c 10.88.88.88     10.88.88.88               irb.1088 [ae0.0]        permanent remote
b0:33:a6:11:49:00 10.99.99.3      10.99.99.3                irb.1099 [vtep.32769]   permanent remote
52:54:00:a4:c5:73 10.99.99.99     10.99.99.99               irb.1099 [ae1.0]        permanent remote
b0:33:a6:11:49:36 10.255.240.2    10.255.240.2              et-0/0/49.0             none
b0:33:a6:11:49:35 10.255.240.5    10.255.240.5              et-0/0/48.0             none
fe:00:00:00:00:80 128.0.0.16      fpc0                      bme0.0                  permanent
c0:42:d0:16:af:a3 192.168.1.1     192.168.1.1               em2.32768               none
72:92:c6:eb:1e:6c 192.168.1.16    192.168.1.16              em2.32768               none
cc:e1:94:ba:39:e0 192.168.230.1   192.168.230.1             vme.0                   none
Total entries: 11

{master:0}
mist@Core1>
```

Core2: Ethernet Switching and ARP Tables

```
mist@Core2> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 9 entries, 9 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical              SVLBNH/      Active
   name                 address             flags    interface            VENH Index   source
   default              c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1033             00:00:5e:00:01:01   DRP      esi.1753             1748         05:00:00:fd:e8:00:00:2b:19:00
   vlan1033             c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1088             00:00:5e:00:01:01   DRP      esi.1752             1748         05:00:00:fd:e8:00:00:2b:50:00
   vlan1088             52:54:00:91:ed:5c   DLR      ae0.0
   vlan1088             c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1099             00:00:5e:00:01:01   DRP      esi.1751             1748         05:00:00:fd:e8:00:00:2b:5b:00
   vlan1099             52:54:00:a4:c5:73   DLR      ae1.0
   vlan1099             c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2

{master:0}
mist@Core2> show arp
MAC Address       Address         Name                      Interface               Flags
c0:42:d0:16:af:a0 10.33.33.2      10.33.33.2                irb.1033 [vtep.32769]   permanent remote
c0:42:d0:16:af:a0 10.88.88.2      10.88.88.2                irb.1088 [vtep.32769]   permanent remote
52:54:00:91:ed:5c 10.88.88.88     10.88.88.88               irb.1088 [ae0.0]        permanent remote
c0:42:d0:16:af:a0 10.99.99.2      10.99.99.2                irb.1099 [vtep.32769]   permanent remote
52:54:00:a4:c5:73 10.99.99.99     10.99.99.99               irb.1099 [ae1.0]        permanent remote
c0:42:d0:16:af:d6 10.255.240.3    10.255.240.3              et-0/0/49.0             none
c0:42:d0:16:af:d5 10.255.240.4    10.255.240.4              et-0/0/48.0             none
fe:00:00:00:00:80 128.0.0.16      fpc0                      bme0.0                  permanent
b0:33:a6:11:49:03 192.168.1.1     192.168.1.1               em2.32768               none
be:be:16:a8:6d:dd 192.168.1.16    192.168.1.16              em2.32768               none
cc:e1:94:ba:39:e0 192.168.230.1   192.168.230.1             vme.0                   none
Total entries: 11

{master:0}
mist@Core2>
```

Result of Our Fabric Checks

Connectivity between the collapsed core switches looks correct since MAC and ARPs are being learned across the fabric on both cores. Let's look at the connection between core and WAN router next.

We need to configure the attachment of the WAN router to complete the entire design. Without the WAN router configuration, the fabric only allows the following communications:

- The same VLAN/VNI on the same access switch but different ports.
- The same VLAN/VNI on different access switches.
- Different VLAN/VNI attached to the same VRF on the same access switch, but different ports.
- Different VLAN/VNI attached to the same VRF on different access switches.

All traffic between VRFs is always isolated inside the fabric. For security reasons, there is no possible configuration to perform route leaking between VRFs. This means that traffic between them is handled directly inside the fabric without the need to traverse through the WAN router as a possible enforcement point.

APPENDIX: WAN Router Integration into the Fabric

In general, there are several possible ways to attach a WAN router to a campus fabric.

- Using a Layer 2 forwarding method:
  - The fabric uplinks are configured as ESI-LAGs and contain one or more tagged VLANs (one for each VRF) to communicate with the WAN router.
  - It is also necessary that you configure the IP address of the WAN router interface manually as the next-hop IP address for default-forwarding on each fabric VRF as already shown above.
  - The WAN router itself needs to understand standard IEEE 802.3ad LAG with active LACP.
  - If you have more than one WAN router attached for redundancy, it is advised to provide failover mechanisms between them for the interface IP addresses towards the fabric. VRRP is recommended.
  - Routes between fabric and WAN router are only statically configured.
- Using a Layer 3 forwarding method:
  - The fabric uplinks are configured as Layer 3 peer-to-peer IP links.
  - Per fabric VRF, a peer-to-peer link needs to be established with the WAN router.

- Usually, there are multiple peer-to-peer links on a single physical uplink. Those are further segmented through tagged VLANs to provide isolation on the uplinks.
- There is no need to manually configure next hops for each VRF inside the fabric as it is assumed that the propagation of the default gateways will be obtained from the WAN router through a routing protocol.
- Between the fabric and the WAN router, a routing protocol must be established to exchange routes.
- The campus fabric supports exterior BGP and OSPF as routing protocols towards the WAN router.

NOTE: The details of such integration are explained in a JVD extension for all fabric types.

For simplicity, in this JVD we have chosen to utilize the Layer 2 exit through the ESI-LAG as the stretched VLAN, which is not intended to be used in production.

Remember that you chose to deploy the border gateway capability on the Juniper Networks® QFX5110 Switches during the Campus Fabric Core-Distribution deployment, represented below:

Figure 57: WAN-Router Integration through ESI-LAG

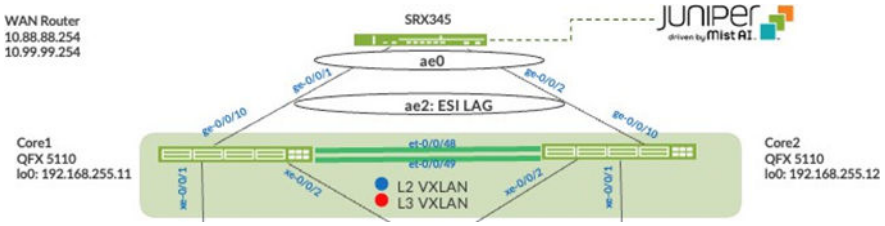

Juniper Mist enables the Juniper Networks® QFX5110 Switch to translate between VXLAN traffic within the campus fabric and standard Ethernet switching for external connectivity. In this case, it is an SRX Series Firewall. Let's verify the ESI status on the core switches:

```
root@Core1> show lacp statistics interfaces
warning: lacp subsystem not running - not needed by configuration.
```

We must configure the ESI-LAG for Layer 2 connectivity between the collapsed core switches and the WAN Router as Juniper Mist does not configure this automatically. You can associate a pre-defined port profile with the requisite ports on each core switch.

The following represents an existing port profile applied to each SRX Series Firewall facing the QFX5110 Switch port:

Figure 58: Port Configuration with ESI-LAG

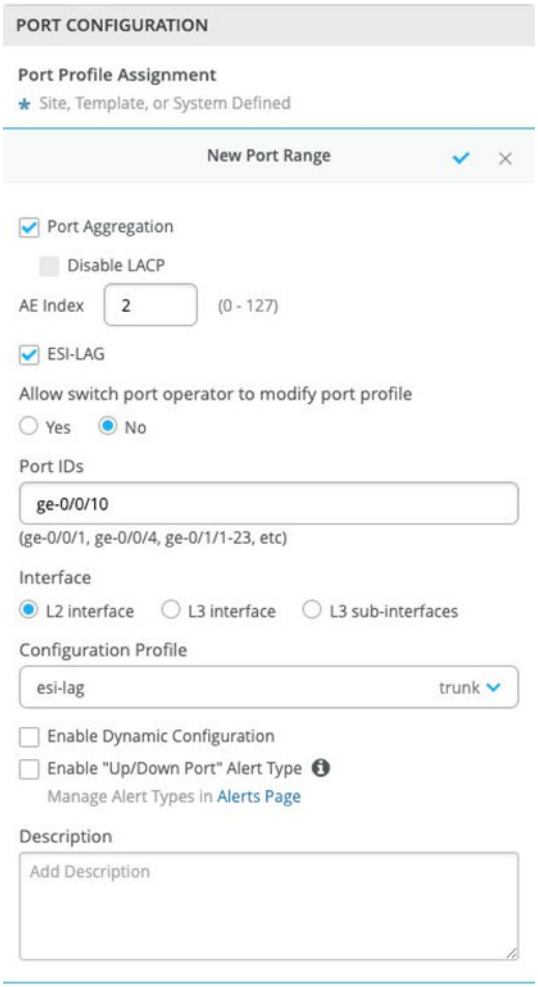

Save the configuration and then verify the changes on the core switch.

Core1: The active status of LACP to the WAN router produces new entries in the switch's ARP tables:

```
mist@Core1> show lacp statistics interfaces ae2
Aggregated interface: ae2
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-0/0/10                358         358            0            0

{master:0}
mist@Core1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:24:07
     10001      c0:42:d0:16:af:a0  irb.0                          Jun 20 17:23:45
     11033      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:19:00  Jun 20 17:55:42  10.33.33.1
     11033      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:55:42  10.33.33.3
     11033      c0:42:d0:16:af:a0  irb.1033                       Jun 20 17:55:42  10.33.33.2
     11033      ee:38:73:9a:b6:a6  00:11:00:00:00:01:00:01:02:02  Jun 20 18:15:50  10.33.33.254
     11088      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:50:00  Jun 20 17:55:42  10.88.88.1
     11088      52:54:00:91:ed:5c  00:11:00:00:00:01:00:01:02:00  Jun 20 18:14:06  10.88.88.88
     11088      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:55:42  10.88.88.3
     11088      c0:42:d0:16:af:a0  irb.1088                       Jun 20 17:55:42  10.88.88.2
     11088      ee:38:73:9a:b6:a6  00:11:00:00:00:01:00:01:02:02  Jun 20 18:15:50  10.88.88.254
     11099      00:00:5e:00:01:01  05:00:00:fd:e8:00:00:2b:5b:00  Jun 20 17:55:42  10.99.99.1
     11099      52:54:00:a4:c5:73  00:11:00:00:00:01:00:01:02:01  Jun 20 18:19:23  10.99.99.99
     11099      b0:33:a6:11:49:00  172.16.254.1                   Jun 20 17:55:42  10.99.99.3
     11099      c0:42:d0:16:af:a0  irb.1099                       Jun 20 17:55:42  10.99.99.2
     11099      ee:38:73:9a:b6:a6  00:11:00:00:00:01:00:01:02:02  Jun 20 18:15:50  10.99.99.254

{master:0}
mist@Core1> show arp
MAC Address       Address         Name                    Interface               Flags
b0:33:a6:11:49:00 10.33.33.3      10.33.33.3              irb.1033 [vtep.32769]   permanent remote
ee:38:73:9a:b6:a6 10.33.33.254    10.33.33.254            irb.1033 [ae2.0]        permanent remote
b0:33:a6:11:49:00 10.88.88.3      10.88.88.3              irb.1088 [vtep.32769]   permanent remote
52:54:00:91:ed:5c 10.88.88.88     10.88.88.88             irb.1088 [ae0.0]        permanent remote
ee:38:73:9a:b6:a6 10.88.88.254    10.88.88.254            irb.1088 [ae2.0]        permanent remote
b0:33:a6:11:49:00 10.99.99.3      10.99.99.3              irb.1099 [vtep.32769]   permanent remote
52:54:00:a4:c5:73 10.99.99.99     10.99.99.99             irb.1099 [ae1.0]        permanent remote
ee:38:73:9a:b6:a6 10.99.99.254    10.99.99.254            irb.1099 [ae2.0]        permanent remote
b0:33:a6:11:49:36 10.255.240.2    10.255.240.2            et-0/0/49.0             none
b0:33:a6:11:49:35 10.255.240.5    10.255.240.5            et-0/0/48.0             none
fe:00:00:00:00:80 128.0.0.16      fpc0                    bme0.0                  permanent
c0:42:d0:16:af:a3 192.168.1.1     192.168.1.1             em2.32768               none
72:92:c6:eb:1e:6c 192.168.1.16    192.168.1.16            em2.32768               none
cc:e1:94:ba:39:e0 192.168.230.1   192.168.230.1           vme.0                   none
Total entries: 14

{master:0}
mist@Core1>
```

Core2: The active status of LACP to the WAN router produces new entries in the switch's ARP tables:

```
mist@Core2> show lacp statistics interfaces ae2
Aggregated interface: ae2
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-0/0/10                601         600            0            0

{master:0}
mist@Core2> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 13 entries, 13 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical              SVLBNH/      Active
   name                address             flags    interface            VENH Index   source
   default             c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1033            00:00:5e:00:01:01   DRP      esi.1753             1748         05:00:00:fd:e8:00:00:2b:19:00
   vlan1033            c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1033            ee:38:73:9a:b6:a6   DLR      ae2.0
   vlan1088            00:00:5e:00:01:01   DRP      esi.1752             1748         05:00:00:fd:e8:00:00:2b:50:00
   vlan1088            52:54:00:91:ed:5c   DLR      ae0.0
   vlan1088            c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1088            ee:38:73:9a:b6:a6   DLR      ae2.0
   vlan1099            00:00:5e:00:01:01   DRP      esi.1751             1748         05:00:00:fd:e8:00:00:2b:5b:00
   vlan1099            52:54:00:a4:c5:73   DLR      ae1.0
   vlan1099            a0:36:9f:bd:0e:a0   DL       ae1.0
   vlan1099            c0:42:d0:16:af:a0   DRP      vtep.32769                        172.16.254.2
   vlan1099            ee:38:73:9a:b6:a6   DLR      ae2.0

{master:0}
mist@Core2> show arp
MAC Address       Address         Name                    Interface               Flags
c0:42:d0:16:af:a0 10.33.33.2      10.33.33.2              irb.1033 [vtep.32769]   permanent remote
ee:38:73:9a:b6:a6 10.33.33.254    10.33.33.254            irb.1033 [ae2.0]        permanent remote
c0:42:d0:16:af:a0 10.88.88.2      10.88.88.2              irb.1088 [vtep.32769]   permanent remote
52:54:00:91:ed:5c 10.88.88.88     10.88.88.88             irb.1088 [ae0.0]        permanent remote
ee:38:73:9a:b6:a6 10.88.88.254    10.88.88.254            irb.1088 [ae2.0]        permanent remote
c0:42:d0:16:af:a0 10.99.99.2      10.99.99.2              irb.1099 [vtep.32769]   permanent remote
52:54:00:a4:c5:73 10.99.99.99     10.99.99.99             irb.1099 [ae1.0]        permanent remote
ee:38:73:9a:b6:a6 10.99.99.254    10.99.99.254            irb.1099 [ae2.0]        permanent remote
c0:42:d0:16:af:d6 10.255.240.3    10.255.240.3            et-0/0/49.0             none
c0:42:d0:16:af:d5 10.255.240.4    10.255.240.4            et-0/0/48.0             none
fe:00:00:00:00:80 128.0.0.16      fpc0                    bme0.0                  permanent
b0:33:a6:11:49:03 192.168.1.1     192.168.1.1             em2.32768               none
be:be:16:a8:6d:dd 192.168.1.16    192.168.1.16            em2.32768               none
cc:e1:94:ba:39:e0 192.168.230.1   192.168.230.1           vme.0                   none
Total entries: 14

{master:0}
mist@Core2>
```

NOTE: The IP address entries with 254 in the last octet now found in Core1 and Core2 are the WAN router's default gateway addresses.

We go back to Desktop1 to see if it can traverse the fabric:

```
root@desktop1:~# ping 1.1 -c 2
PING 1.1 (1.0.0.1) 56(84) bytes of data.
64 bytes from 1.0.0.1: icmp_seq=1 ttl=52 time=2.41 ms
64 bytes from 1.0.0.1: icmp_seq=2 ttl=52 time=2.39 ms

--- 1.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1002ms
rtt min/avg/max/mdev = 2.389/2.399/2.409/0.010 ms
root@desktop1:~# traceroute 1.1
traceroute to 1.1 (1.0.0.1), 30 hops max, 60 byte packets
 1  10.99.99.2 (10.99.99.2)  1.342 ms 10.99.99.3 (10.99.99.3)  4.080 ms  4.026 ms
 2  10.99.99.254 (10.99.99.254)  0.867 ms  0.804 ms  0.769 ms
 3  192.168.230.1 (192.168.230.1)  21.728 ms  21.704 ms  21.672 ms
 4  192.168.70.1 (192.168.70.1)  1.310 ms  1.268 ms  1.204 ms
 5  172.16.80.1 (172.16.80.1)  1.323 ms  1.301 ms  1.269 ms
 6  172.16.254.2 (172.16.254.2)  1.207 ms  0.677 ms  0.684 ms
 7  172.21.0.8 (172.21.0.8)  0.721 ms  0.955 ms  1.054 ms
 8  66.129.246.2 (66.129.246.2)  1.549 ms  1.650 ms  1.578 ms
 9  xe-0-0-54-1.a02.snjsca04.us.bb.gin.ntt.net (157.238.64.89)  27.174 ms  27.147 ms  27.117 ms
10  ae-9.r25.snjsca04.us.bb.gin.ntt.net (129.250.3.102)  3.016 ms ae-9.r24.snjsca04.us.bb.gin.ntt.net (129.250.2.2)  2.141 ms  2.120 ms
11  ae-40.r02.snjsca04.us.bb.gin.ntt.net (129.250.3.121)  2.436 ms ae-19.r01.snjsca04.us.bb.gin.ntt.net (129.250.3.27)  2.402 ms  2.408 ms
12  ae-0.cloudflare.snjsca04.us.bb.gin.ntt.net (128.241.10.23)  7.479 ms ae-1.cloudflare.snjsca04.us.bb.gin.ntt.net (131.103.117.82)  21.064 ms  21.03
7 ms
13  162.158.164.2 (162.158.164.2)  2.213 ms 172.68.188.22 (172.68.188.22)  2.165 ms 162.158.164.2 (162.158.164.2)  2.141 ms
14  one.one.one.one (1.0.0.1)  2.114 ms  2.140 ms  2.070 ms
root@desktop1:~#
```

Next, verify that Desktop1 can ping Desktop2:

```
root@desktop1:~# ping 10.88.88.88 -c 2
PING 10.88.88.88 (10.88.88.88) 56(84) bytes of data.
64 bytes from 10.88.88.88: icmp_seq=1 ttl=62 time=0.945 ms
64 bytes from 10.88.88.88: icmp_seq=2 ttl=62 time=0.844 ms

--- 10.88.88.88 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.844/0.894/0.945/0.050 ms
root@desktop1:~#
```

As the last step, verify Desktop1 can ping Desktop2:

Figure 59: Verify VRF to VRF Traffic

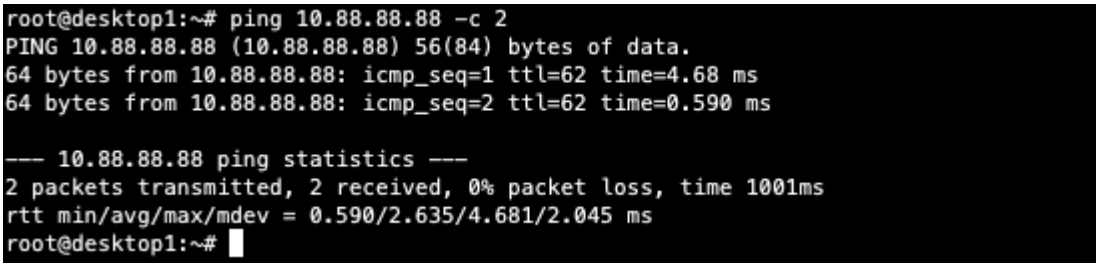

Figure 60: Topology Repeat

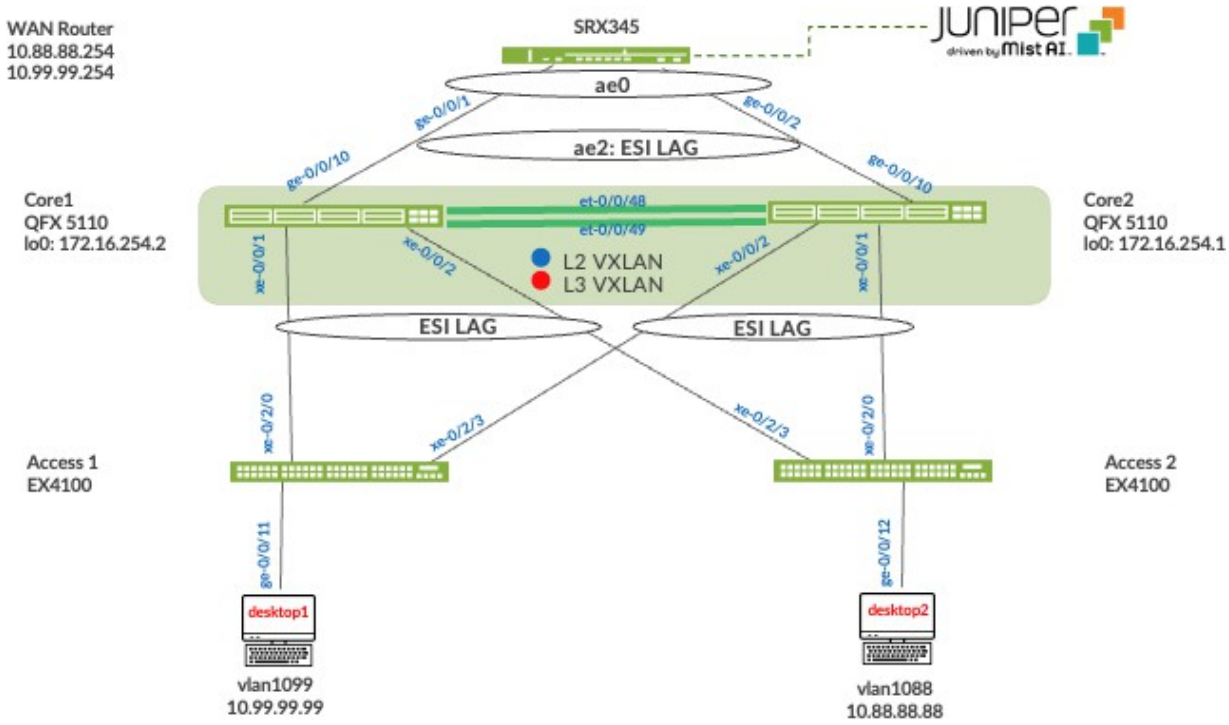

Conclusion: Connectivity within and outside of the campus fabric is verified. Desktops communicate with each other through the fabric, each in an isolated VRF, then forwarded to the SRX Series Firewall through the ESI-LAG on both core devices when accessing services outside of the campus fabric. The campus fabric performs total isolation between the VRFs by default while using the SRX Series Firewall to accept or discard inter-VRF communications.

APPENDIX: EVPN Insights

Juniper Mist Wired Assurance provides real-time status related to the health of the campus fabric EVPN multihoming deployment using telemetry such as BGP neighbor status and TX and RX port statistics. The following screenshots are taken from the campus fabric EVPN multihoming build by accessing the campus fabric option under the **Organization > Wired > Campus Fabric** path of the portal:

Figure 61: core1 Insights

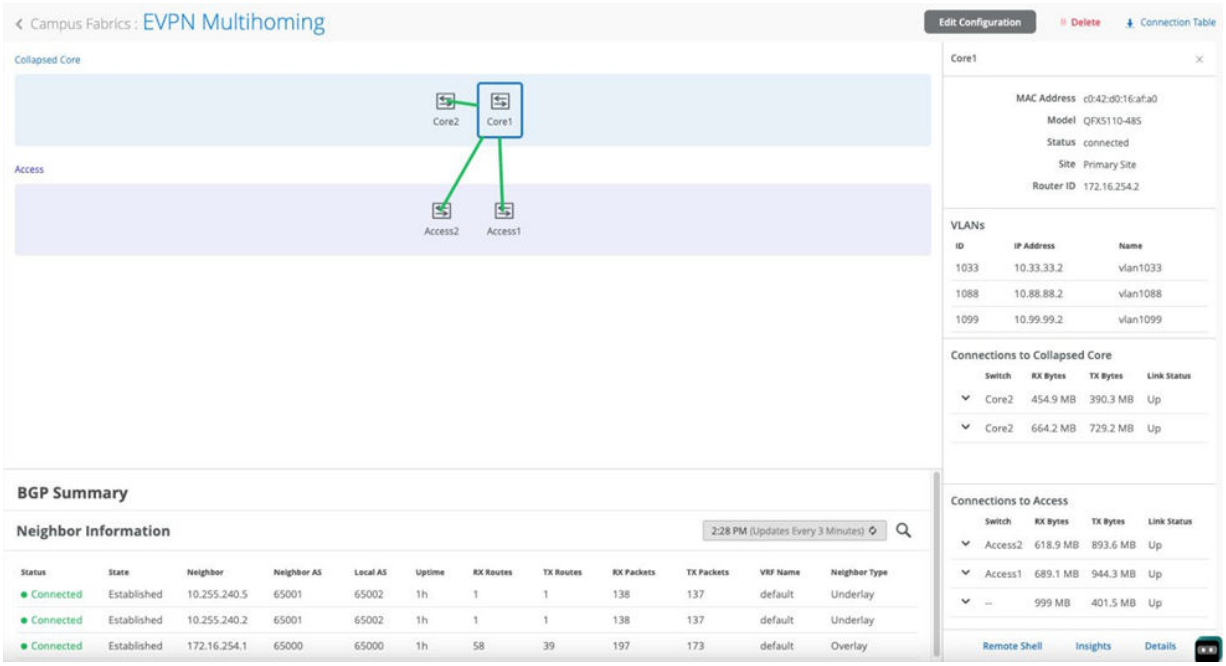

Figure 62: core2 Insights

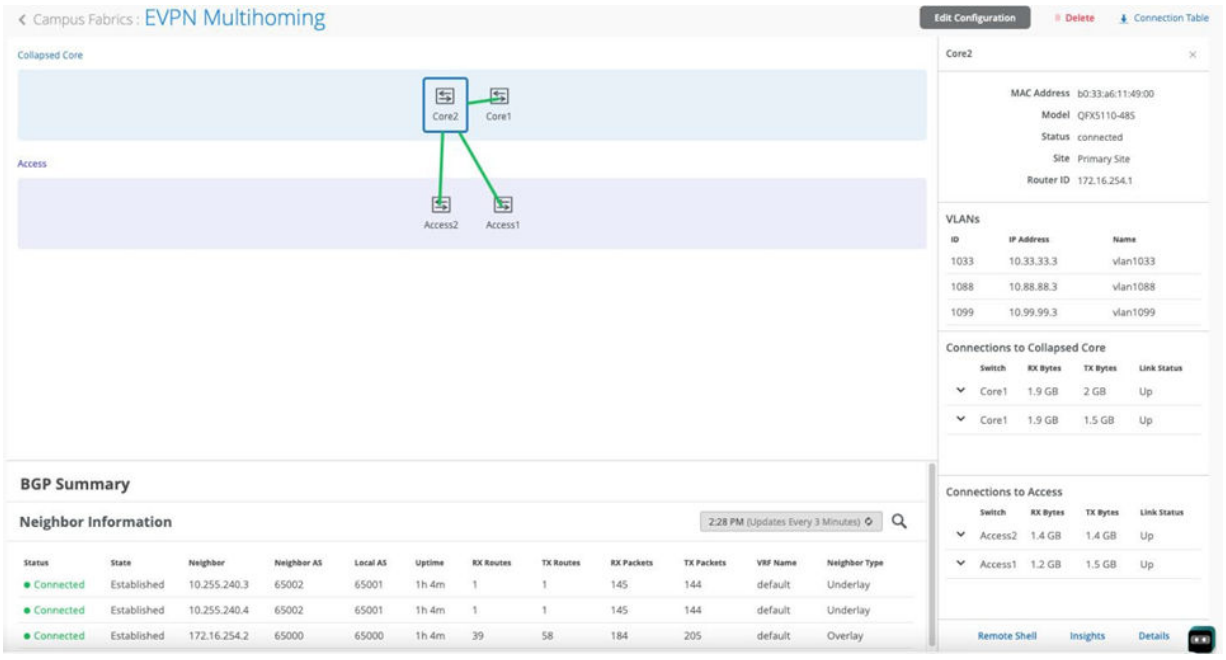

From this view, Juniper Mist also provides remote accessibility into each device's console through the remote shell option as well as rich telemetry through the **Switch Insights** option. Remote shell has been

demonstrated throughout this document when displaying the real-time operational status of each device during the verification stage.

Switch insights of Core1 displays historical telemetry including BGP peering status critical to the health of the campus fabric:

Figure 63: Single Switch Insights

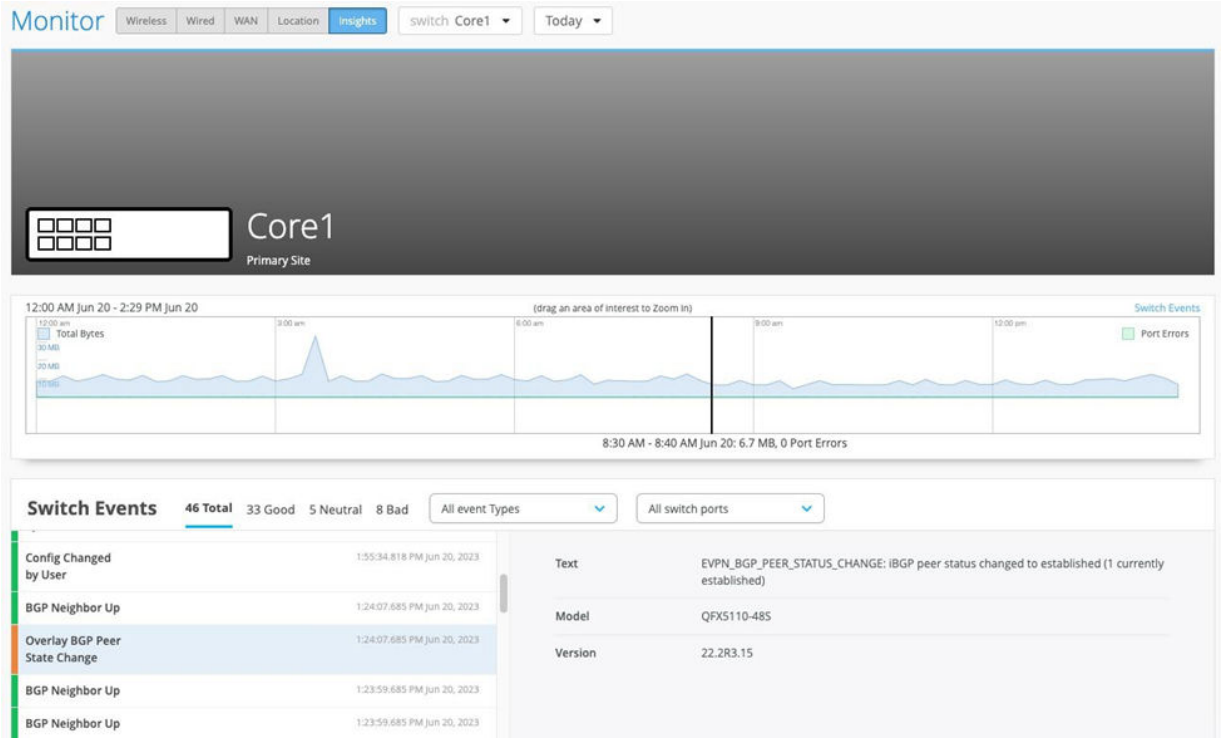

Summary: Juniper Mist Campus Fabric provides an easy method to build an EVPN multihoming deployment to enable EVPN-VXLAN overlay networks. This can be done solely in the Juniper Mist portal. Steps are added in this document to help you understand the troubleshooting steps if deployment is not working correctly.

APPENDIX: Junos OS Configuration from This Fabric

IN THIS SECTION

Campus Fabric EVPN Multihoming Configurations

This section displays the configuration output from the Juniper Mist cloud for the IP Fabric underlay on the core and distribution switches using eBGP.

Juniper Mist provides the following options (default in parentheses):

- BGP Local AS (65000)
- AS Base (65001)
- Loopback pool (172.16.254.0/23)
- Subnet (10.255.240.0/20) – point-to-point interfaces between adjacent layers

Throughout the campus fabric between core and distribution layers, Juniper Mist enables per-packet (Junos OS defines this as per-flow) load balancing using ECMP and fast convergence of BGP in the event of a link or node failure using BFD.

Core1 Configuration:

1. Interconnects with core2:

```
set interfaces et-0/0/48 description evpn_downlink-to-b033a6114900
set interfaces et-0/0/48 unit 0 family inet address 10.255.240.4/31
```

```
set interfaces et-0/0/49 description evpn_uplink-to-b033a6114900
set interfaces et-0/0/49 unit 0 family inet address 10.255.240.3/31
```

2. Loopback interface and router ID:

```
set groups top interfaces lo0 unit 0 family inet address 172.16.254.2/32
set groups top routing-options router-id 172.16.254.2
```

3. Per-packet load-balancing:

```
set groups top policy-options policy-statement ecmp_policy then load-balance per-packet
set groups top policy-options policy-statement ecmp_policy then accept
set groups top routing-options forwarding-table export ecmp_policy
```

4. BGP underlay network between the two distribution switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 65002
set protocols bgp group evpn_underlay multipath multiple-as
set protocols bgp group evpn_underlay bfd-liveness-detection minimum- interval 350
set protocols bgp group evpn_underlay bfd-liveness-detection multiplier 3
set protocols bgp group evpn_underlay neighbor 10.255.240.2 peer-as 65001
set protocols bgp group evpn_underlay neighbor 10.255.240.5 peer-as 65001
```

Core2 Configuration:

1. Interconnects between the two distribution switches:

```
set interfaces et-0/0/48 description evpn_uplink-to-c042d016afa0
set interfaces et-0/0/48 unit 0 family inet address 10.255.240.5/31
set interfaces et-0/0/49 description evpn_downlink-to-c042d016afa0
set interfaces et-0/0/49 unit 0 family inet address 10.255.240.2/31
```

2. Loopback interface and router ID:

```
set groups top interfaces lo0 unit 0 family inet address 172.16.254.1/32
set groups top routing-options router-id 172.16.254.1
```

3. Per-packet load balancing:

```
set groups top policy-options policy-statement ecmp_policy then load-balance per-packet
set groups top policy-options policy-statement ecmp_policy then accept
set groups top routing-options forwarding-table export ecmp_policy
```

4. BGP underlay network between the two distribution switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 65001
set protocols bgp group evpn_underlay multipath multiple-as
set protocols bgp group evpn_underlay bfd-liveness-detection minimum- interval 350
set protocols bgp group evpn_underlay bfd-liveness-detection multiplier 3
set protocols bgp group evpn_underlay neighbor 10.255.240.4 peer-as 65002
set protocols bgp group evpn_underlay neighbor 10.255.240.3 peer-as 65002
```

Configuration of the EVPN-VXLAN Overlay and Virtual Networks

This section displays the Juniper Mist cloud configuration output for the EVPN-VXLAN overlay on the core and distribution switches using eBGP.

Juniper Mist enables load balancing across the overlay network and fast convergence of BGP in the event of a link or node failure using BFD between the core and distribution layers.

Juniper Mist provisions Layer 3 IRB interfaces on the distribution layer.

Juniper Mist enables VXLAN tunnelling, VLAN to VXLAN mapping, and MP-BGP configuration snippets such as vrf-targets on the distribution and core switches.

The VRFs for traffic isolation are provisioned on the distribution switches.

Core1 Configuration:

1. BGP overlay peering between the two distribution switches:

```
set protocols bgp group evpn_overlay type internal
set protocols bgp group evpn_overlay local-address 172.16.254.2
set protocols bgp group evpn_overlay log-updown
set protocols bgp group evpn_overlay family evpn signaling
set protocols bgp group evpn_overlay authentication-key "xyz"
set protocols bgp group evpn_overlay cluster 1.0.0.1
set protocols bgp group evpn_overlay local-as 65000
set protocols bgp group evpn_overlay multipath
set protocols bgp group evpn_overlay bfd-liveness-detection minimum- interval 1000
set protocols bgp group evpn_overlay bfd-liveness-detection multiplier 3
set protocols bgp group evpn_overlay bfd-liveness-detection session-mode automatic
set protocols bgp group evpn_overlay neighbor 172.16.254.1
```

2. Switch options that define vrf-targets and the source loopback interface used for VXLAN:

```
set groups top switch-options vtep-source-interface lo0.0
set groups top switch-options route-distinguisher 172.16.254.2:1
set groups top switch-options vrf-target target:65000:1
set groups top switch-options vrf-target auto
```

3. VXLAN encapsulation:

```
set groups top protocols evpn no-core-isolation
set groups top protocols evpn encapsulation vxlan
set groups top protocols evpn default-gateway no-gateway-community
set groups top protocols evpn extended-vni-list all
```

4. VRFs used for traffic isolation:

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi routing-options static route 0.0.0.0/0 next-hop
```

```
10.33.33.254
set groups top routing-instances guest-wifi routing-options auto-export
set groups top routing-instances guest-wifi interface irb.1033
set groups top routing-instances guest-wifi route-distinguisher 172.16.254.2:103
set groups top routing-instances guest-wifi vrf-target target:65000:103
set groups top routing-instances guest-wifi vrf-table-label
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers routing-options static route 0.0.0.0/0 next-hop
10.88.88.254
set groups top routing-instances developers routing-options auto-export
set groups top routing-instances developers interface irb.1088
set groups top routing-instances developers route-distinguisher 172.16.254.2:102
set groups top routing-instances developers vrf-target target:65000:102
set groups top routing-instances developers vrf-table-label
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it routing-options static route 0.0.0.0/0 next-hop
10.99.99.254
set groups top routing-instances corp-it routing-options auto-export
set groups top routing-instances corp-it interface irb.1099
set groups top routing-instances corp-it route-distinguisher 172.16.254.2:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
```

5. VLAN to VXLAN mapping:

```
set vlans vlan1033 vlan-id 1033
set vlans vlan1033 l3-interface irb.1033
set vlans vlan1033 vxlan vni 11033
set vlans vlan1088 vlan-id 1088
set vlans vlan1088 l3-interface irb.1088
set vlans vlan1088 vxlan vni 11088
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 l3-interface irb.1099
set vlans vlan1099 vxlan vni 11099
```

6. Layer 3 IRB interface enablement with virtual gateway addressing:

```
set interfaces irb unit 1033 virtual-gateway-accept-data
set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet address 10.33.33.2/24 virtual-gateway-address
10.33.33.1
```

```
set interfaces irb unit 1088 virtual-gateway-accept-data
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet address 10.88.88.2/24 virtual-gateway-address
10.88.88.1
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet address 10.99.99.2/24 virtual-gateway-address
10.99.99.1
```

Core2 Configuration:

1. BGP overlay peering between the two distribution switches:

```
set protocols bgp group evpn_overlay type internal
set protocols bgp group evpn_overlay local-address 172.16.254.1
set protocols bgp group evpn_overlay log-updown
set protocols bgp group evpn_overlay family evpn signaling
set protocols bgp group evpn_overlay authentication-key "xyz"
set protocols bgp group evpn_overlay cluster 1.0.0.1
set protocols bgp group evpn_overlay local-as 65000
set protocols bgp group evpn_overlay multipath
set protocols bgp group evpn_overlay bfd-liveness-detection minimum- interval 1000
set protocols bgp group evpn_overlay bfd-liveness-detection multiplier 3
set protocols bgp group evpn_overlay bfd-liveness-detection session-mode automatic
set protocols bgp group evpn_overlay neighbor 172.16.254.2
```

2. Switch options that define vrf-targets and the source loopback interface used for VXLAN:

```
set groups top switch-options vtep-source-interface lo0.0
set groups top switch-options route-distinguisher 172.16.254.1:1
set groups top switch-options vrf-target target:65000:1
set groups top switch-options vrf-target auto
```

3. VXLAN encapsulation:

```
set groups top protocols evpn no-core-isolation
set groups top protocols evpn encapsulation vxlan
set groups top protocols evpn default-gateway no-gateway-community
set groups top protocols evpn extended-vni-list all
```

4. VRFs used for traffic isolation:

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi routing-options static route 0.0.0.0/0 next-hop
10.33.33.254
set groups top routing-instances guest-wifi routing-options auto-export
set groups top routing-instances guest-wifi interface irb.1033
set groups top routing-instances guest-wifi route-distinguisher 172.16.254.1:103
set groups top routing-instances guest-wifi vrf-target target:65000:103
set groups top routing-instances guest-wifi vrf-table-label
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers routing-options static route 0.0.0.0/0 next-hop
10.88.88.254
set groups top routing-instances developers routing-options auto-export
set groups top routing-instances developers interface irb.1088
set groups top routing-instances developers route-distinguisher 172.16.254.1:102
set groups top routing-instances developers vrf-target target:65000:102
set groups top routing-instances developers vrf-table-label
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it routing-options static route 0.0.0.0/0 next-hop
10.99.99.254
set groups top routing-instances corp-it routing-options auto-export
set groups top routing-instances corp-it interface irb.1099
set groups top routing-instances corp-it route-distinguisher 172.16.254.1:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
```

5. VLAN to VXLAN mapping:

```
set vlans vlan1033 vlan-id 1033
set vlans vlan1033 l3-interface irb.1033
set vlans vlan1033 vxlan vni 11033
set vlans vlan1088 vlan-id 1088
set vlans vlan1088 l3-interface irb.1088
set vlans vlan1088 vxlan vni 11088
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 l3-interface irb.1099
set vlans vlan1099 vxlan vni 11099
```

6. Layer 3 IRB interface enablement with virtual gateway addressing.

```
set interfaces irb unit 1033 virtual-gateway-accept-data
set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet address 10.33.33.3/24 virtual-gateway-address
10.33.33.1
set interfaces irb unit 1088 virtual-gateway-accept-data
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet address 10.88.88.3/24 virtual-gateway-address
10.88.88.1
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet address 10.99.99.3/24 virtual-gateway-address
10.99.99.1
```

Configuration of the Layer 2 ESI-LAG Between the Distribution Switches and the Access Switches

This section displays the configuration output from the Juniper Mist cloud for the enablement of the Layer 2 ESI-LAGs between the distribution switches and access switches. This Juniper Mist profile enables all VLANs on the Ethernet bundle with requisite ESI and LACP configuration options. From the perspective of the access switches, the Ethernet bundle that is configured on the access layer views the ESI-LAG as a single MAC address with the same LACP system-ID. This enables load hashing between distribution and access layers without requiring Layer 2 loop-free detection protocols such as RSTP.

Figure 64: Access Switch Attach to Collapsed Core Switches

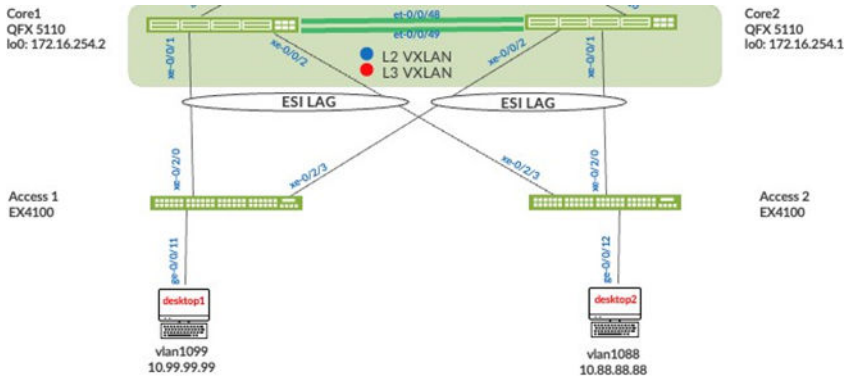

Core1 Configuration:

1. Interface association with the newly created Ethernet bundle that includes ESI and LACP configuration:

```
set interfaces ae0 apply-groups esi-lag
set interfaces ae0 esi 00:11:00:00:00:01:00:01:02:00 set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast set interfaces ae0 aggregated-
ether-options lacp system-id 00:00:00:31:57:00
set interfaces ae0 aggregated-ether-options lacp admin-key 0
set interfaces ae1 apply-groups esi-lag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:02:01 set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast set interfaces ae1 aggregated-
ether-options lacp system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set groups esi-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set interfaces xe-0/0/1 description esilag-to-4c734f095900 set interfaces xe-0/0/1 hold-time
up 120000
set interfaces xe-0/0/1 hold-time down 1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/1 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/0/1 unit 0
set interfaces xe-0/0/2 description esilag-to-4c734f095900 set interfaces xe-0/0/2 hold-time
up 120000
set interfaces xe-0/0/2 hold-time down 1
set interfaces xe-0/0/2 ether-options 802.3ad ae0
```

Core2 Configuration:

2. Interface association with the newly created Ethernet bundle that includes ESI and LACP configuration:

```
set interfaces ae0 apply-groups esi-lag
set interfaces ae0 esi 00:11:00:00:00:01:00:01:02:00 set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:31:57:00
set interfaces ae0 aggregated-ether-options lacp admin-key 0
set interfaces ae1 apply-groups esi-lag
```

```
set interfaces ae1 esi 00:11:00:00:00:01:00:01:02:01 set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set groups esi-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set interfaces xe-0/0/1 description esilag-to-4c734f095900
set interfaces xe-0/0/1 hold-time up 120000
set interfaces xe-0/0/1 hold-time down 1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/1 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/0/1 unit 0
set interfaces xe-0/0/2 description esilag-to-4c734f095900
set interfaces xe-0/0/2 hold-time up 120000
set interfaces xe-0/0/2 hold-time down 1
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/2 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/0/2 unit 0
```

Access1 Configuration:

3. VLANs associated with the new LACP Ethernet bundle:

```
set groups esi-lag interfaces <*> mtu 9200
set groups esi-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set interfaces ae1 apply-groups esi-lag
set interfaces ae1 aggregated-ether-options lacp active
set interfaces xe-0/2/0 ether-options 802.3ad ae1
set interfaces xe-0/2/0 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/2/0 unit 0
set interfaces xe-0/2/3 ether-options 802.3ad ae1
set interfaces xe-0/2/3 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/2/3 unit 0
```

Access2 Configuration:

4. VLANs associated with the new LACP Ethernet bundle:

```
set groups esi-lag interfaces <*> mtu 9200
set groups esi-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups esi-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set interfaces ae0 apply-groups esi-lag
set interfaces ae0 aggregated-ether-options lacp active
set interfaces xe-0/2/0 ether-options 802.3ad ae0
set interfaces xe-0/2/0 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/2/0 unit 0
set interfaces xe-0/2/3 ether-options 802.3ad ae0
set interfaces xe-0/2/3 unit 0 family ethernet-switching storm-control default
deactivate interfaces xe-0/2/3 unit 0
```

Configuration of the Layer 2 ESI-LAG Between the Core Switches and the MX Router

This section displays the configuration output from the Juniper Mist cloud for the enablement of the Layer 2 ESI LAG between the core switches and SRX Series Firewall. This Mist profile enables all VLANs on the Ethernet bundle with requisite ESI and LACP configuration options. From the perspective of the SRX Series Firewall, the Ethernet bundle that is configured on the SRX Series Firewall views the ESI-LAG as a single MAC address with the same LACP system- ID. This enables load hashing between the core and SRX Series Firewall without requiring Layer 2 loop-free detection protocols such as RSTP.

Figure 65: Layer 2 ESI-LAG Supporting Active-Active Load Balancing

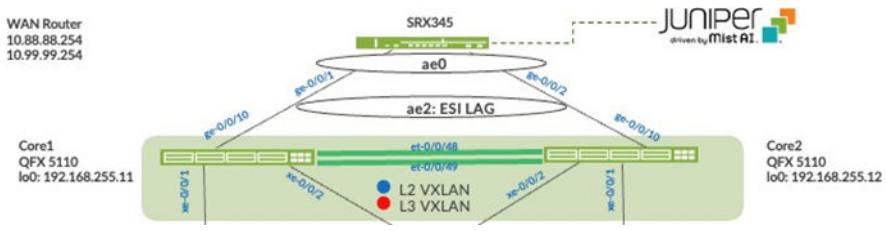

Core 1 Configuration:

1. Interface association with the newly created Ethernet bundle that includes ESI and LACP configuration:

```
set interfaces ge-0/0/10 description esilag-to-4c734f095900
set interfaces ge-0/0/10 hold-time up 120000
set interfaces ge-0/0/10 hold-time down 1
set interfaces ge-0/0/10 ether-options 802.3ad ae2
set interfaces ge-0/0/10 unit 0 family ethernet-switching storm-control default
deactivate interfaces ge-0/0/10 unit 0
set interfaces ae2 apply-groups esi-lag
set interfaces ae2 esi 00:11:00:00:00:01:00:01:02:02
set interfaces ae2 esi all-active
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp periodic fast
set interfaces ae2 aggregated-ether-options lacp system-id 00:00:00:31:57:02
set interfaces ae2 aggregated-ether-options lacp admin-key 2
```

Core 2 Configuration:

2. Interface association with the newly created Ethernet bundle that includes ESI and LACP configuration:

```
set interfaces ge-0/0/10 description esilag-to-4c734f095900
set interfaces ge-0/0/10 hold-time up 120000
set interfaces ge-0/0/10 hold-time down 1
set interfaces ge-0/0/10 ether-options 802.3ad ae2
set interfaces ge-0/0/10 unit 0 family ethernet-switching storm-control default
deactivate interfaces ge-0/0/10 unit 0
set interfaces ae2 apply-groups esi-lag
set interfaces ae2 esi 00:11:00:00:00:01:00:01:02:02 set interfaces ae2 esi all-active
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp periodic fast
set interfaces ae2 aggregated-ether-options lacp system-id 00:00:00:31:57:02
set interfaces ae2 aggregated-ether-options lacp admin-key 2
```

SRX Series Firewall Configuration:

3. Interface association with newly created Ethernet bundle and LACP configuration:

```
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 aggregated-ether-options lacp active
```

```
set interfaces ae0 unit 1033 vlan-id 1033
set interfaces ae0 unit 1033 family inet address 10.33.33.254/24
set interfaces ae0 unit 1088 vlan-id 1088
set interfaces ae0 unit 1088 family inet address 10.88.88.254/24
set interfaces ae0 unit 1099 vlan-id 1099
set interfaces ae0 unit 1099 family inet address 10.99.99.254/24
```

Revision History

Table 2: Revision History

Date	Version	Description
February 2025	JVD-ENTWIRED-EVPNMH-01-01	Initial publish