JUNIPER NETWORKS | Engineering Simplicity

Network Configuration Example

Campus Fabric EVPN Multihoming Workflow

Published
2023-06-28

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Network Configuration Example Campus Fabric EVPN Multihoming Workflow

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Chapter 1 Campus Fabric EVPN Multihoming Workflow

About this Configuration Example

Scope

Use this Network Configuration Example (NCE) of Wired Assurance EVPN multihoming workflow for building an EVPN multihoming deployment.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. Send your comments to design-center-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Technology Primer: EVPN Multihoming Use Case Overview

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. Multichassis 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 L2/L3 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 https://www.rfc-editor.org/rfc/rfc8365 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 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 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 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 Link Aggregation Control Protocol (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 increases multifold 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 through 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 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.

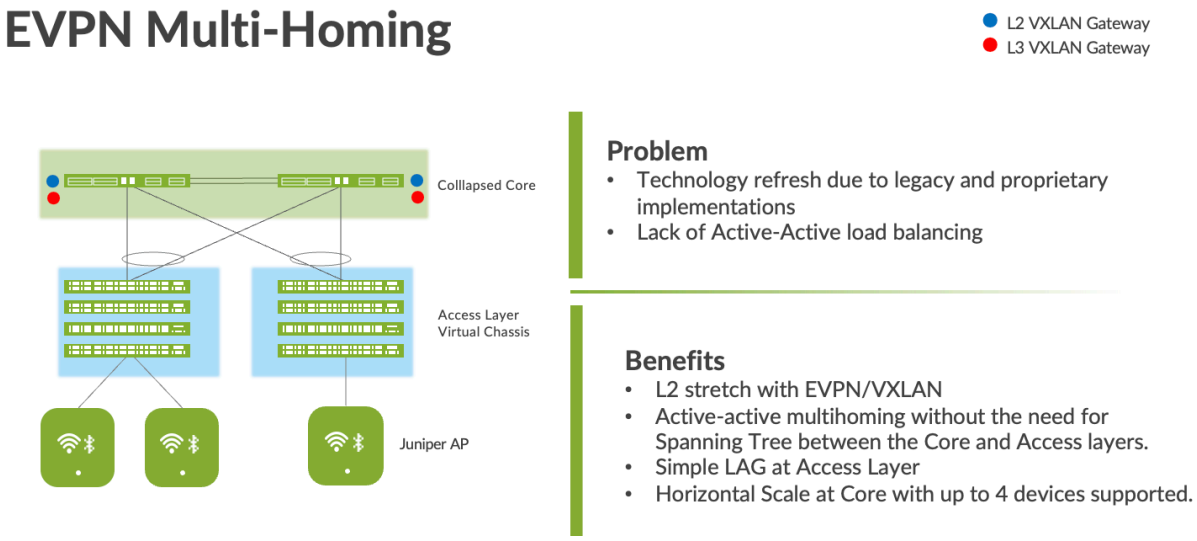

Figure 1: EVPN Multihoming

An EVPN-VXLAN network solves the problems of previous architectures and provides the following benefits:

- Reduced flooding and learning—Control plane-based Layer 2/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. This is because more management traffic consumes the bandwidth which leaves less bandwidth available for production traffic. The EVPN control plane handles the exchange and learning of MAC addresses through eBGP routing, rather than a Layer 2 forwarding plane.

- Scalability—More efficient control-plane based Layer 2 and Layer 3 learning.
- Loop free built in—EVPN VXLAN mitigates the need for Spanning Tree between the access and core layers while supporting active-active load balancing between these layers.

Chapter 2 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, access points, servers, and printers. It is about simplifying every step of the way, starting from Day 0 for seamless onboarding and auto-provisioning through Day 2 and beyond for operations and management. Juniper EX Series Switches provide rich Junos streaming telemetry that enable the insights for switch health metrics and anomaly detection, as well as Juniper Mist AI capabilities.

Mist's AI engine and virtual network assistant, Marvis, further simplifies 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 and/or integration with your Operational Support Systems such as: IT applications, Ticketing Systems, and IP Management Systems.

Mist delivers unique capabilities for the WAN Wide Area Network (WAN), Local Area Network (LAN) and Wireless networks.

- User Interface (UI) or API-driven configuration at scale.
- Service Level Expectations (SLE) for key performance metrics such as throughput, capacity, roaming, and uptime.
- Marvis—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

Chapter 3 EVPN

EVPN Multihoming

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 macrosegmentation purposes.

The Mist UI workflow makes it easy to create campus fabrics.

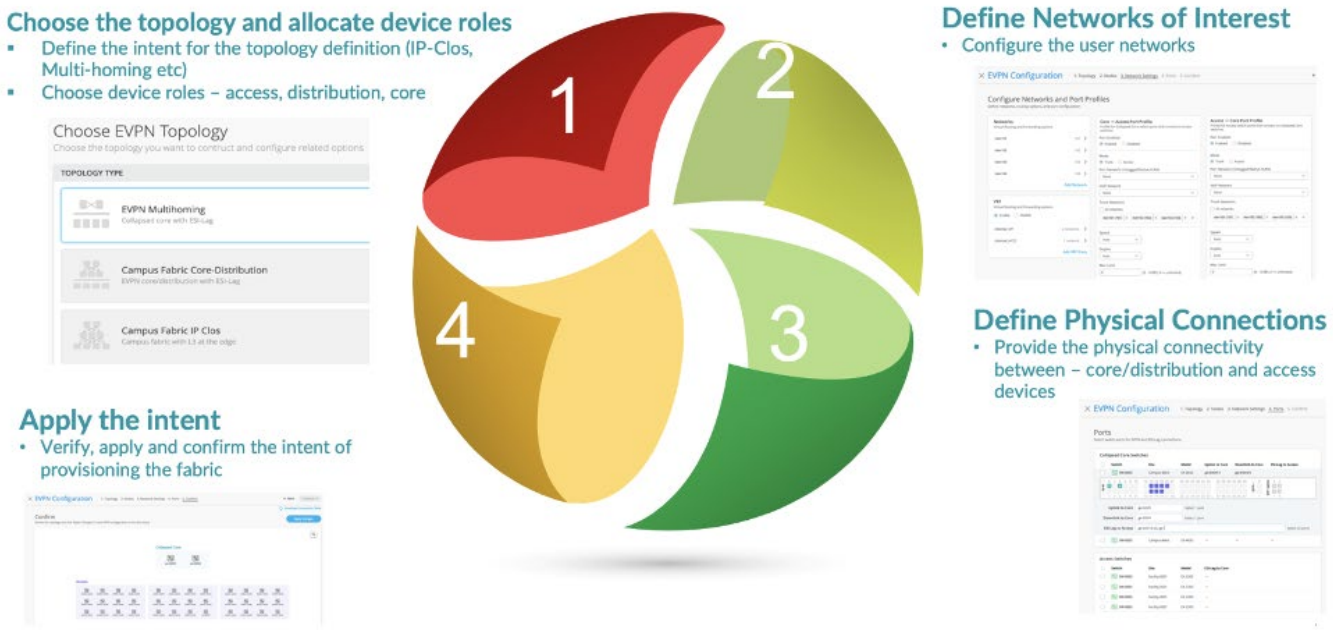

Campus Fabric EVPN Multihoming Components

This configuration example uses the following devices:

- Two QFX5110 switches as distribution devices, software version: Junos OS Release 22.2R3.15 or later.
- Two access layer EX4100 switches, software version: Junos OS Release 22.3R1-S2.1 or later.
- One SRX345 WAN router, software version: 20.2R3-S2.5 or later.
- Juniper Access Points.
- Two Linux desktops that act as wired clients.

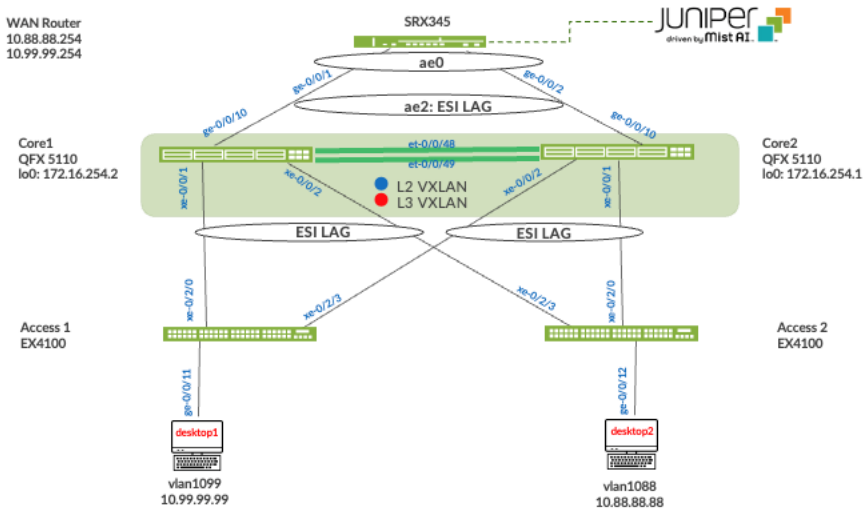

Figure 2: Topology

Juniper Mist Wired Assurance

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

For information on Juniper Switch Onboarding to the Juniper Mist cloud, see:

https://www.juniper.net/documentation/us/en/software/nce/nce-214-midsize-branch-mist-pwp/topics/topic-map/nce-214-midsize-branch-mist-example_part2.html

Wired Assurance, through the Mist UI, is used to build a Campus Fabric EVPN multihoming from ground up and includes the following:

- Assignment of p2p links between the core devices.
- Assignment of unique BGP AS numbers per device participating in the underlay and overlay.
- Creation of VRF instances allow you to logically segment traffic. This also includes the assignment of new or existing VLANs to each representative VRF.
- IP addressing of each L3 gateway Integrated Routing and Bridging (IRB) assigned to the core layer.
- IP addressing of each lo0.0 loopback.
- Configuration of routing policies for underlay and overlay connectivity.
- Optimized Maximum Transmission Unit (MTU) settings for p2p underlay, L3 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 BGP peering between cores 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 were named for the respective layer in the fabric to facilitate building and operating the fabric.

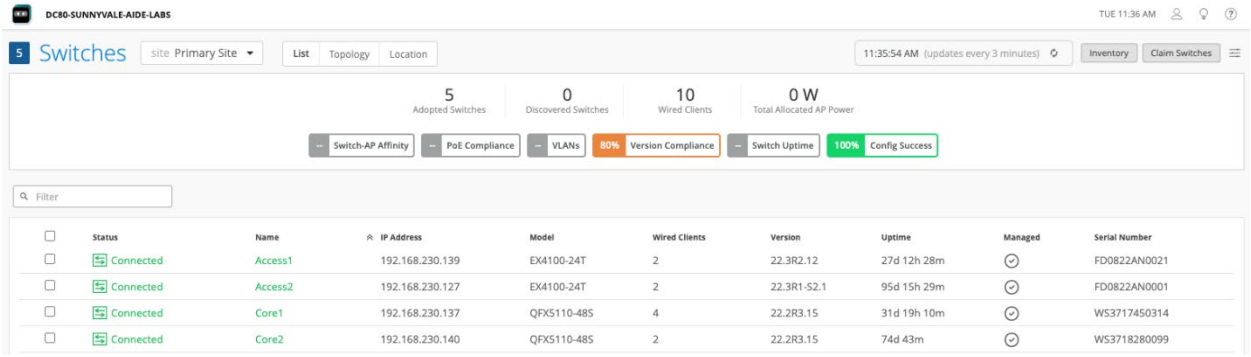

Overview

Templates

A key feature of switch management through the Juniper Mist cloud is the ability to use configuration 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, site) 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. 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. The existing setting may replace or appended.

NOTE: You can include CLI commands for items not native to the Mist UI. 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.

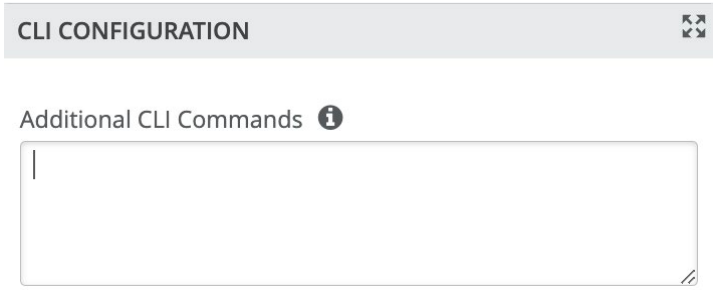

Under Organization and Switch Templates, we use the following template:

Topology

Wired Assurance provides the template for LAN and Loopback IP addressing for each 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 core switches. The devices such as the access layer of switches connect to the access layer using standard Link Aggregation Groups (LAG); while the distribution uses ESI-LAG in a multihoming, load balancing manner.

The WAN router can be provisioned via Mist UI but it 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. The WAN router can be standalone or built as a High Availability cluster. In this document, a single SRX Series Firewall is used as the WAN router.

Chapter 4 Campus Fabric EVPN Multihoming Build Workflow

Create the Campus Fabric

From the Organization option on the left-hand section of the Mist UI, select Wired Campus Fabric.

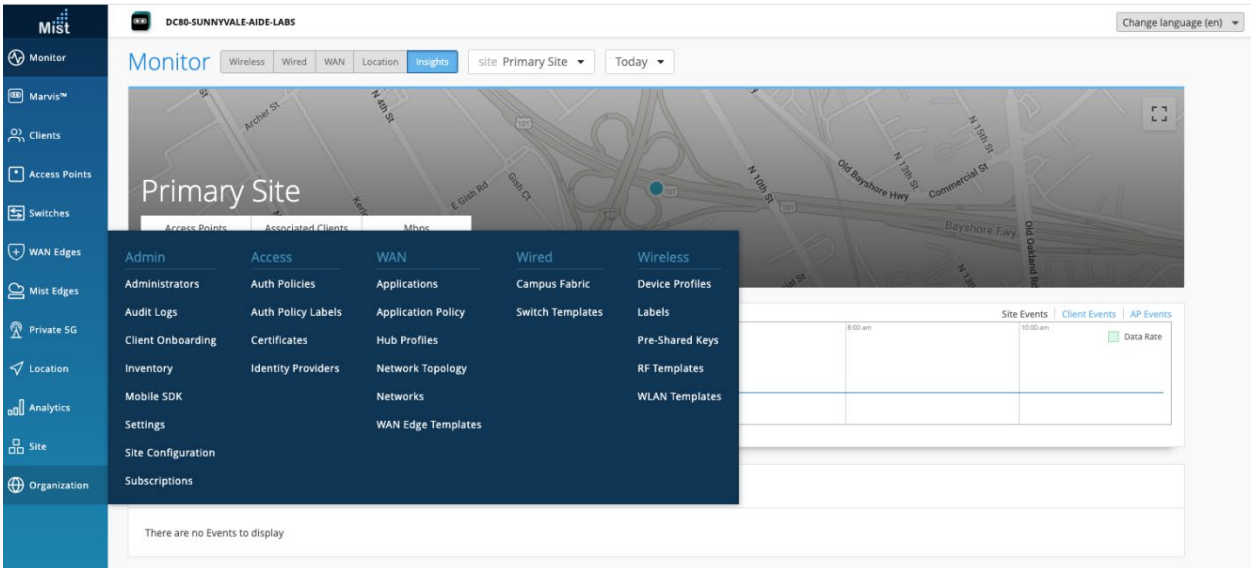

Mist provides the option of deploying a Campus Fabric at the org or site-level noted on the upper left-hand Campus Fabric menu shown below. For example, if you are building a Campus wide architecture with multiple buildings, each building housing distribution and access switches, you can consider building an Org level Campus Fabric. This Campus Fabric ties each of the sites together forming a holistic Campus Fabric. Otherwise, the site build with a single set of core, distribution, and access switches is sufficient.

Campus Fabric Org Build

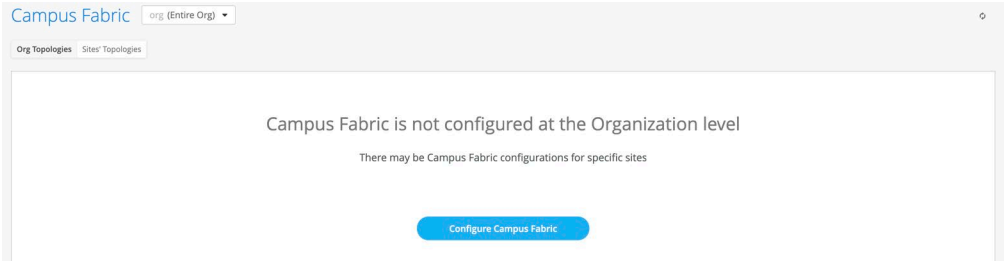

Campus Fabric Site Build

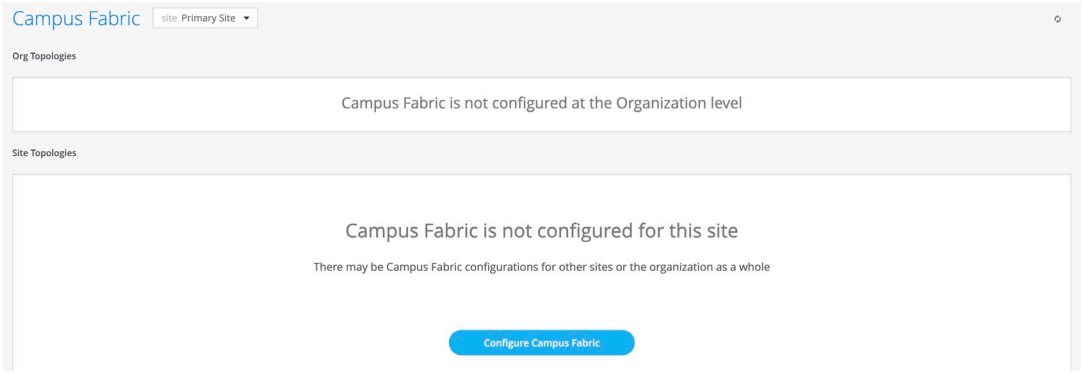

NOTE: EVPN Multihoming is not supported at the Org level.

Choose the Campus Fabric Topology

Select EVPN multihoming and complete the required fields below:

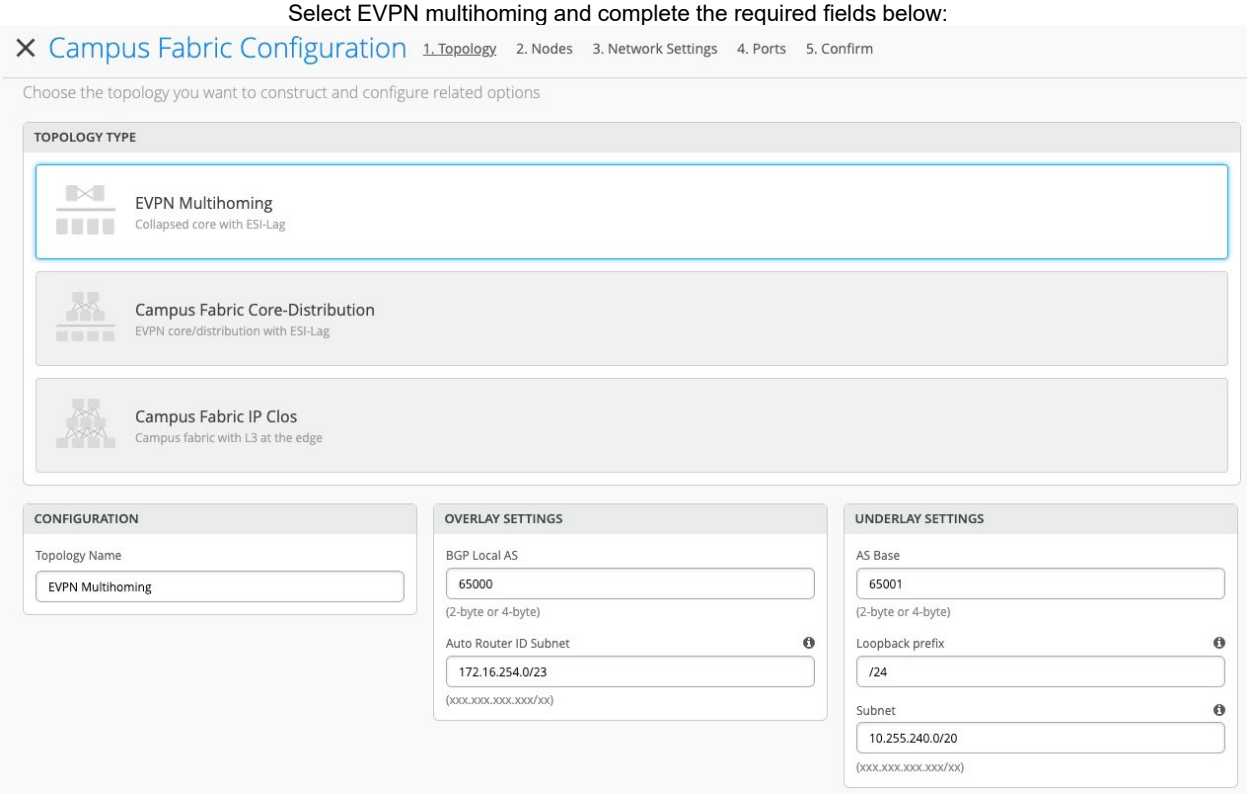

Configuration

Topology Name consistent with the deployment type.

Overlay Settings

- BGP Local AS: represents iBGP AS number that are automatically allocated for every core device. You can use whatever private BGP AS number suits your deployment. The routing policy is provisioned by Mist to ensure that the AS numbers are never advertised outside of the fabric.
- Auto Router ID Subnet: This subnet is used to auto-configure the Router IDs on each device in the Fabric. Router IDs are loopback interfaces (lo0) on each device in the Fabric used for the overlay peering between the core devices.

Underlay Settings

- BGP Local AS: represents the starting point of private BGP AS numbers that are automatically allocated for every device. You can use the private BGP AS number range that suits your deployment. The routing policy is provisioned by Mist to ensure the AS numbers are never advertised outside of the fabric.
- Loopback prefix: represents the range of IP addresses associated with each device's loopback address. You can use the range that suits your deployment. VXLAN tunnelling using a VTEP is associated with this address allocated on all core devices.
- Subnet: represents the range of IP addresses used for point-to-point links between devices. You can use the range that suits your deployment. Mist breaks this subnet into /31 subnet addressing per link.

NOTE: We recommend default settings for all options unless it conflicts with the surrounding environment.

Select Campus Fabric Nodes

You can select devices to participate at each Layer of the EVPN multihoming deployment. We recommend that you validate each device's presence in the site switch inventory before the Campus Fabric is created.

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.

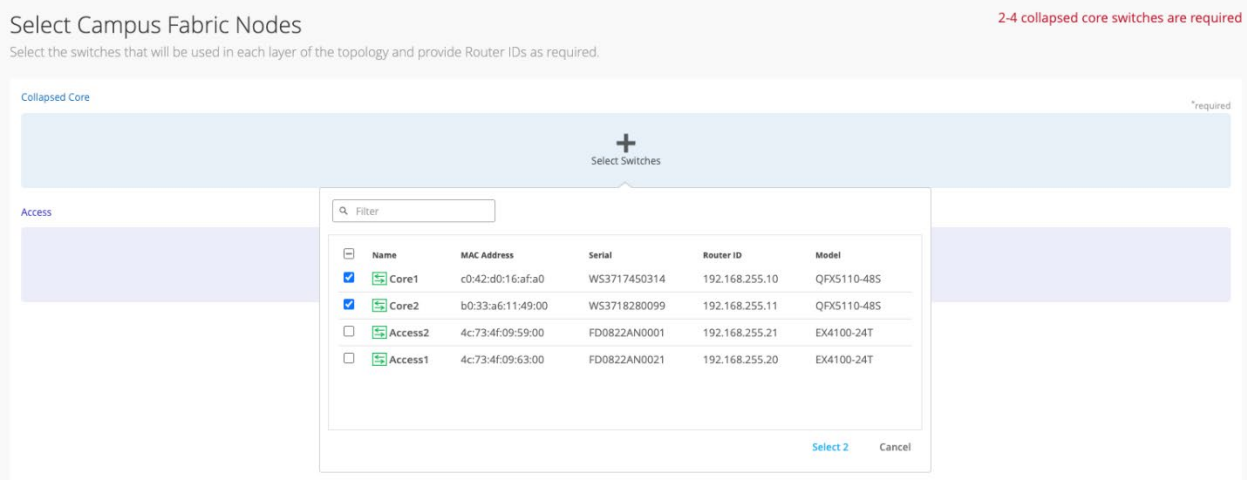

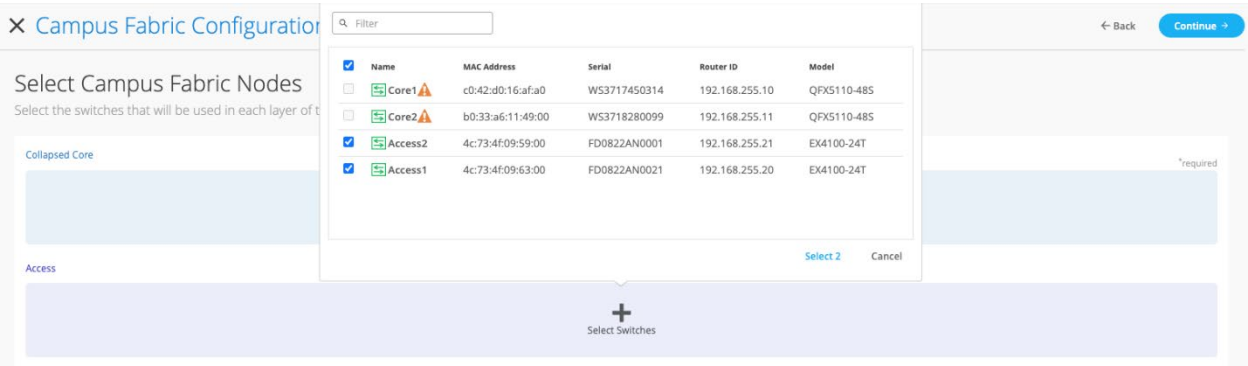

Once all layers have selected the appropriate devices, Mist assigns each device a Router ID associated with a logical loopback address. This loopback is associated with a logical construct called a VTEP to source the VXLAN Tunnel. EVPN multihoming has VTEPs for VXLAN tunnelling on the core switches only.

Mist provided Router-ID/loopback

NOTE: The loopback address and router-ID should be in the same subnet as provided by Mist.

Click the Continue button at the upper-right corner of the page and move to the Network Settings section of the EVPN multihoming deployment process.

Configure Networks

Enter Network information such as VLANs and VRF (routing instances for traffic isolation purposes) options. VLANs are mapped to Virtual Network Identifiers (VNIs) and can optionally be mapped to VRFs to provide customers a way to logically separate traffic patterns such as IoT devices from Corp IT. Port configuration for ESI-LAG between the Collapsed core and access switches are also defined in this section.

VRF

In a Campus Fabric deployment, the use of EVPN VXLAN supports native traffic isolation using routing-instances; commonly called VRFs for macrosegmentation purposes.

For more information on Routing Instance overview, see https://www.juniper.net/documentation/us/en/software/junos/routing-overview/topics/concept/routing-instances-overview.html

VLANs can be placed into a common VRF where all VLANs within each VRF have full connectivity to each other and other external networking resources. A common use case includes most enterprise domains isolating Guest Wi-Fi traffic and saves Internet connectivity. By default, the Campus Fabric provides complete isolation between VRFs forcing inter-VRF communications to traverse a Firewall or security compliance. This aligns with most Enterprise security use cases and compliance and is represented in this document.

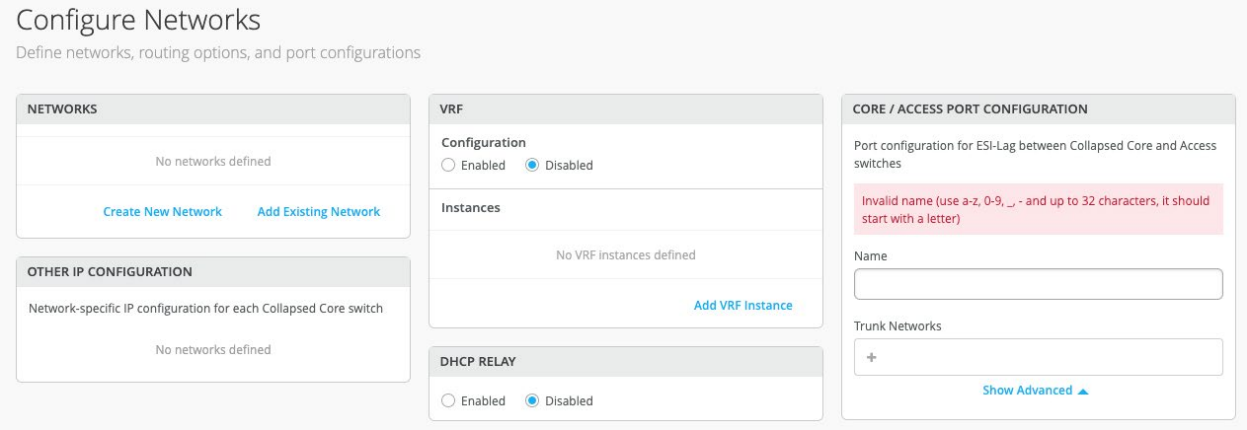

Networks

VLANs can be created or imported under this section including the IP subnet and Default Gateway per each VLANs.

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

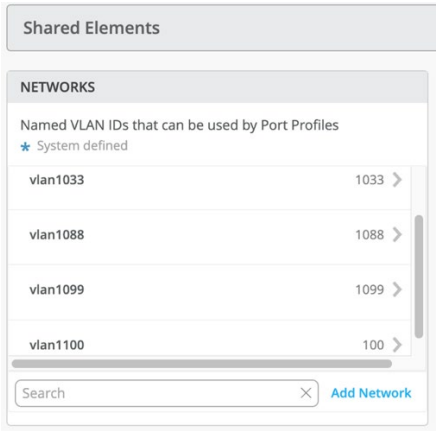

Back to the EVPN multihoming build, you can select the existing template that includes Layer 2 VLAN and IP information. All VLAN and IP information is inherited from the template.

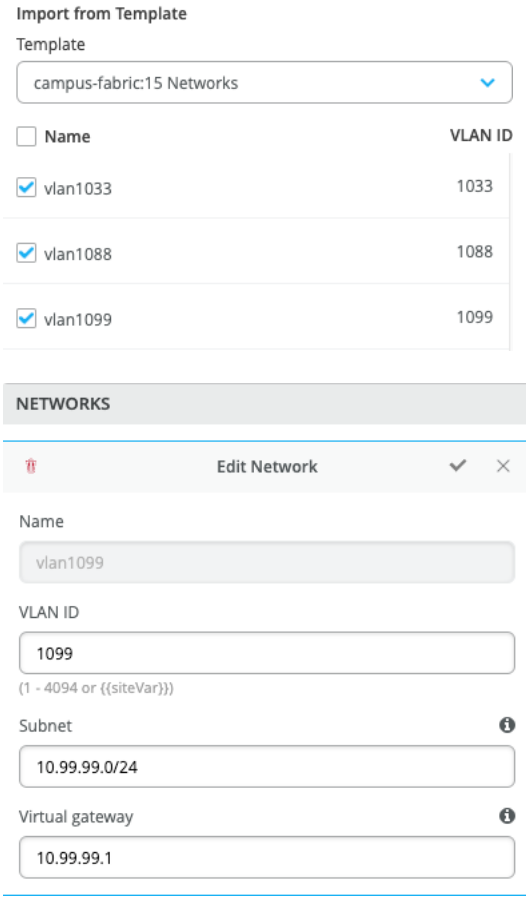

Other IP Configuration

Juniper Mist Wired Assurance provides automatic IP addressing IRBs for each of the VLANs. Then, Port Profiles and Port Configuration associate the VLAN with specified ports.

Once all VLANs are configured, each Collapsed core switch is automatically assigned IP addresses in each VLAN while sharing the Virtual Gateway address, typically used as a Default Gateway address for each VLAN. The following highlights one of the Collapsed Core's IP addressing for each VLAN:

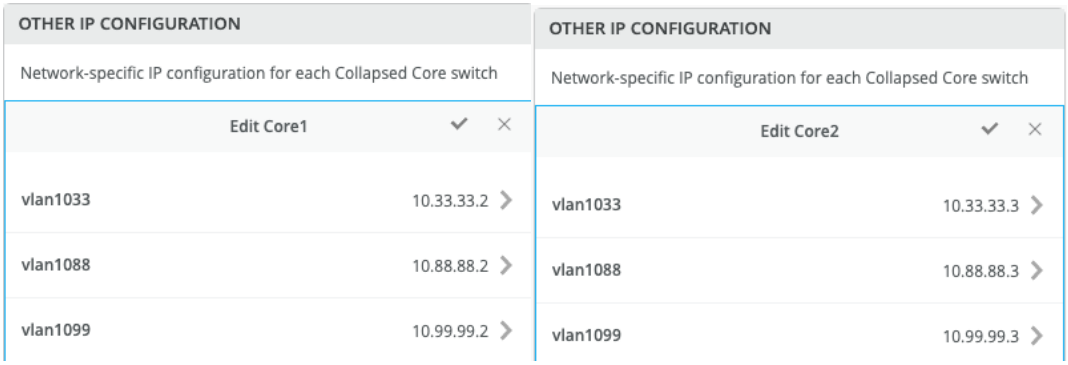

By default, all VLANs are placed in the default VRF. The VRF option allows you to group common VLANs into the same VRF or separate VRFs depending on traffic isolation requirements. This example includes three VRFs or routing instances: corp-it | developers | guest-wifi. Here, you build the first corp-it VRF and select the pre-defined vlan 1099.

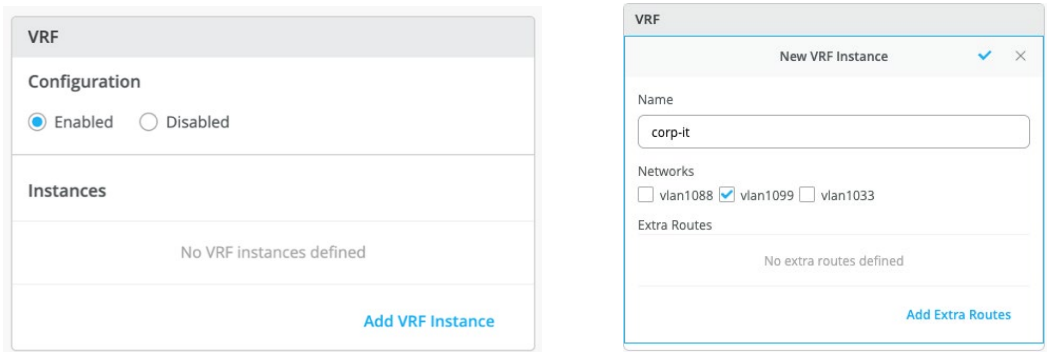

By default, inter-VRF communications is not supported within the Campus Fabric. If inter-VRF communications is required, each VRF can include extra routes such as a Default Route that instructs the Campus Fabric to use an external router or firewall for further security inspection or routing capabilities. In this example, all traffic is trunked over the ESI-LAG and the SRX Series Firewall handles inter-VRF routing. See Figure 2.

SRX Series Firewall participates in the VLANs defined within the Campus Fabric and is the gateway of last resort for all traffic leaving the subnet. You can select the Add Extra Routes option to inform Mist to forward all traffic leaving 10.99.99.0/24 to use the next hop of the SRX Series Firewall: 10.99.99.254

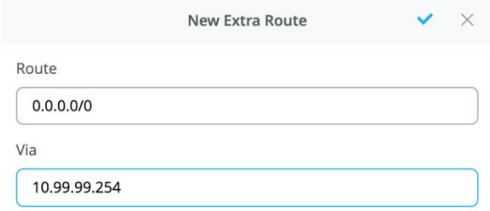

You create two additional VRFs:

- developers using vlan 1088 with 0.0.0.0/0 using 10.88.88.254
- guest-wifi using vlan 1033 with 0.0.0.0/0 using 10.33.33.254

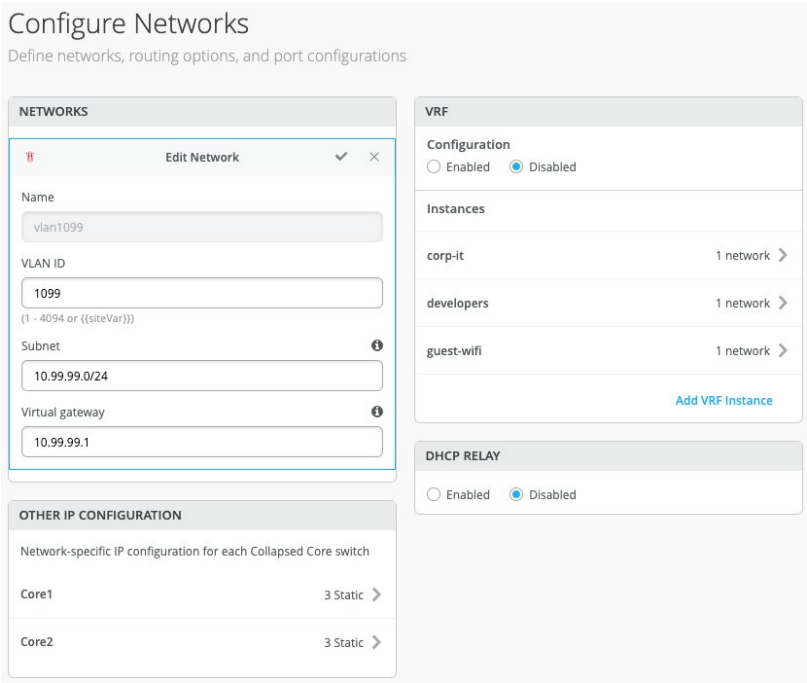

The final step in the Configure Networks section is the Core/Access Port Configuration:

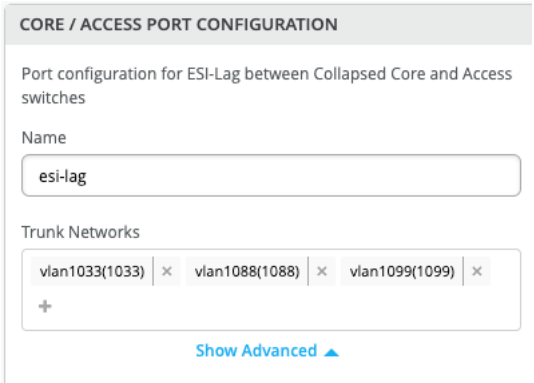

In this section, you configure the active-active ESI-LAG trunks between core and access switches. Here, we name the port configuration and include VLANs associated with this configuration. The Advanced tab provides additional configuration options:

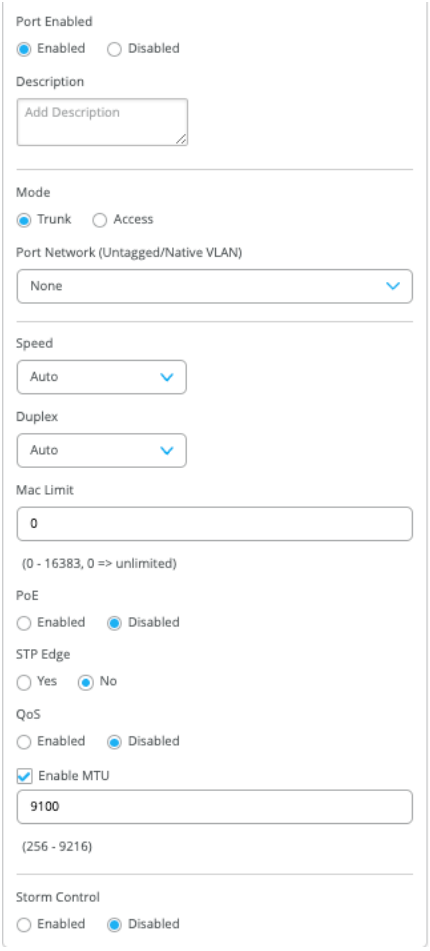

Now that all VLANs are configured and assigned to each VRF, and the core/access ESI-LAGs are built, click the Continue button at the upper-right section of the Mist UI to move to the next step.

Configure Campus Fabric Ports

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

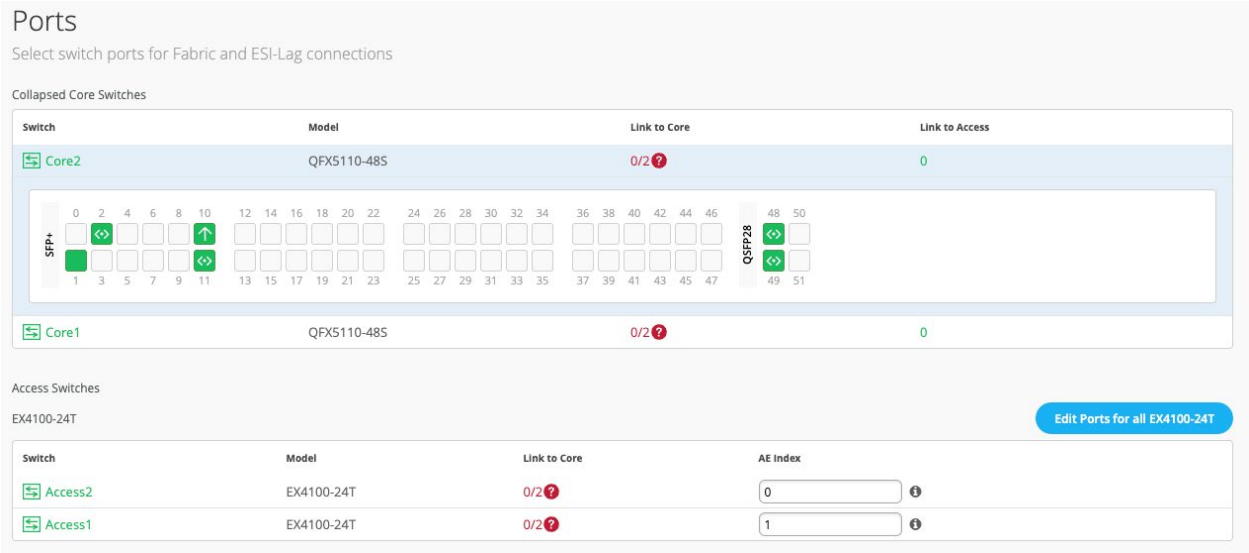

NOTE: To ensure accuracy, we recommend that you run the CLI command “show lldp neighbors” on both core switches prior to this step in the deployment process.

Core Switches

We are now ready to select the ports that interconnect the Collapsed Core Switches.

Core1

You must select et-0/0/48 as a Collapsed Core link.

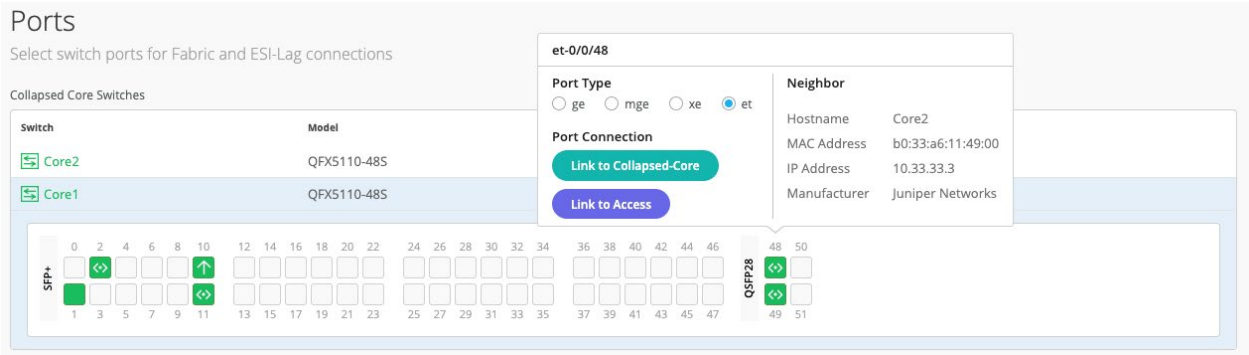

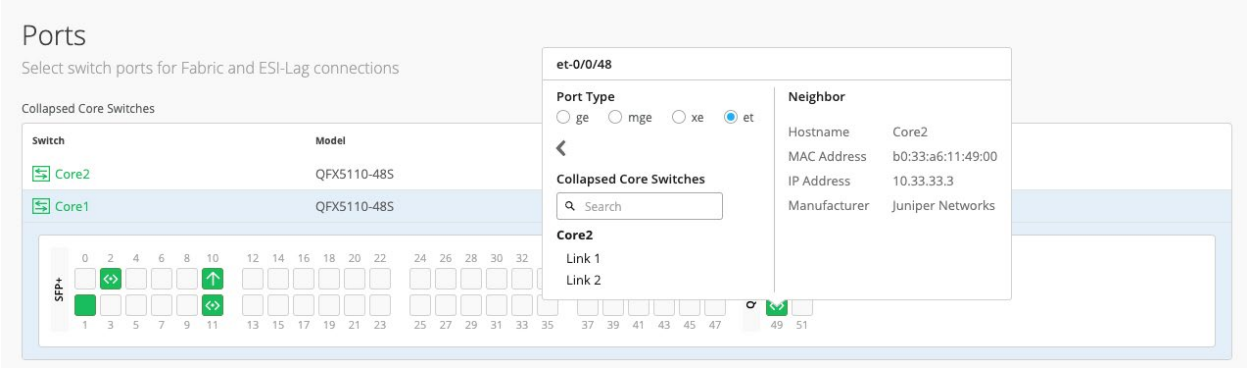

Choose Link1 – Core2.

You must select et-0/0/49 as a Collapsed Core link.

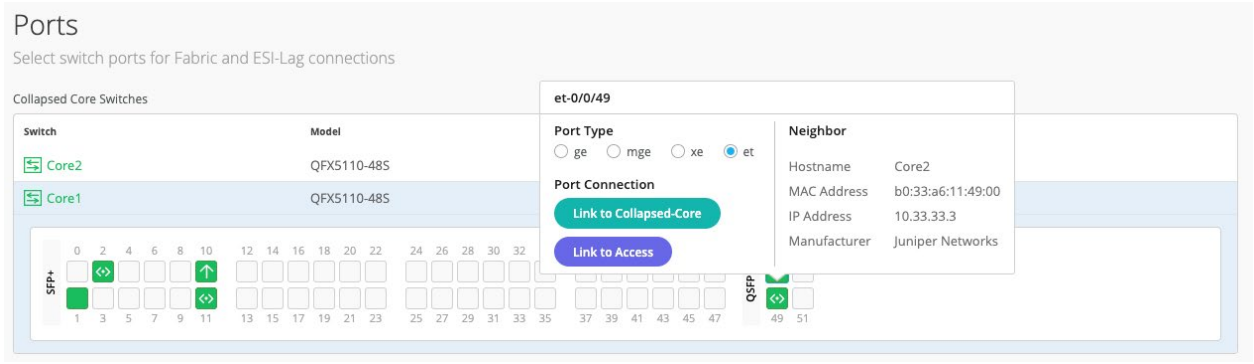

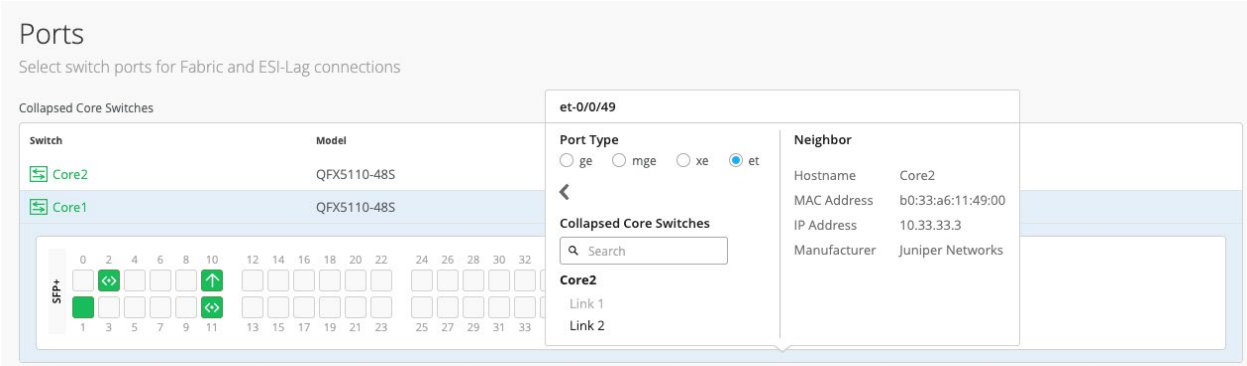

Note that only Link2 is the option left to interconnect with Core2.

Core2

You must select et-0/0/48 as a Collapsed Core link.

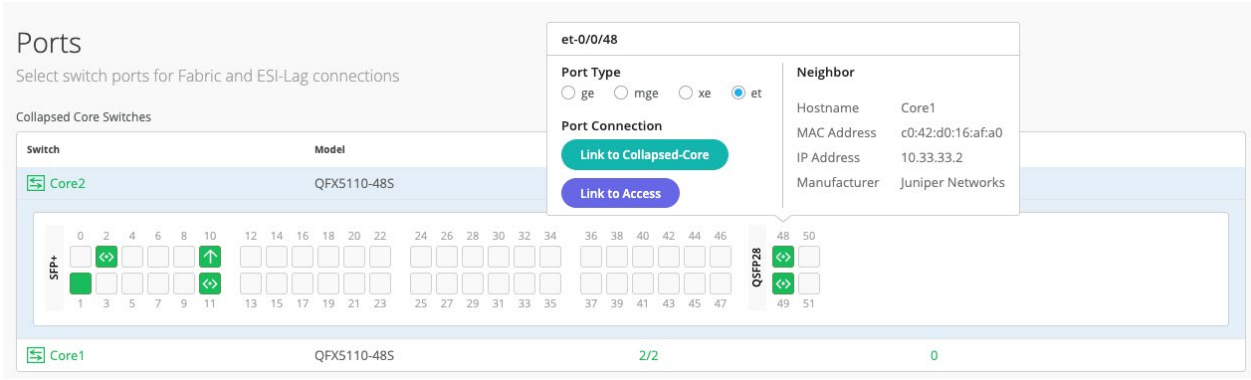

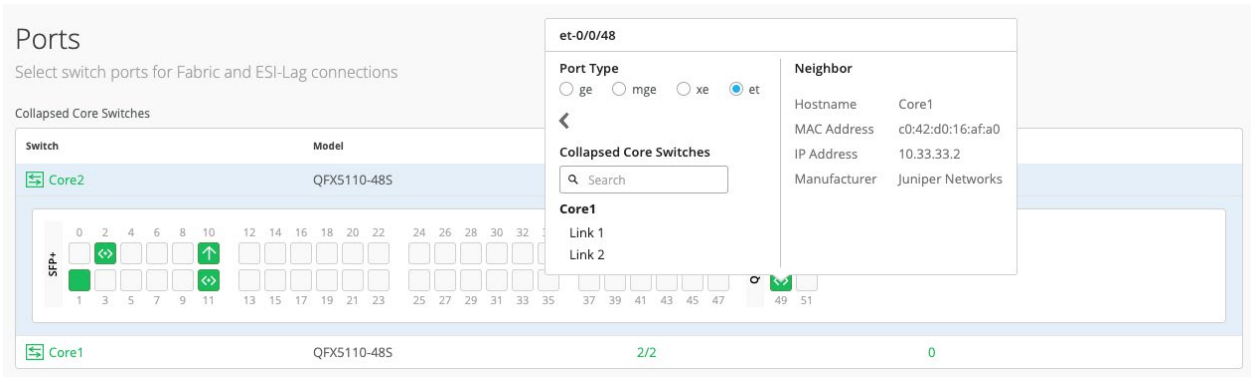

Choose Link1 – Core1.

You must then select et-0/0/49 as a Collapsed Core link.

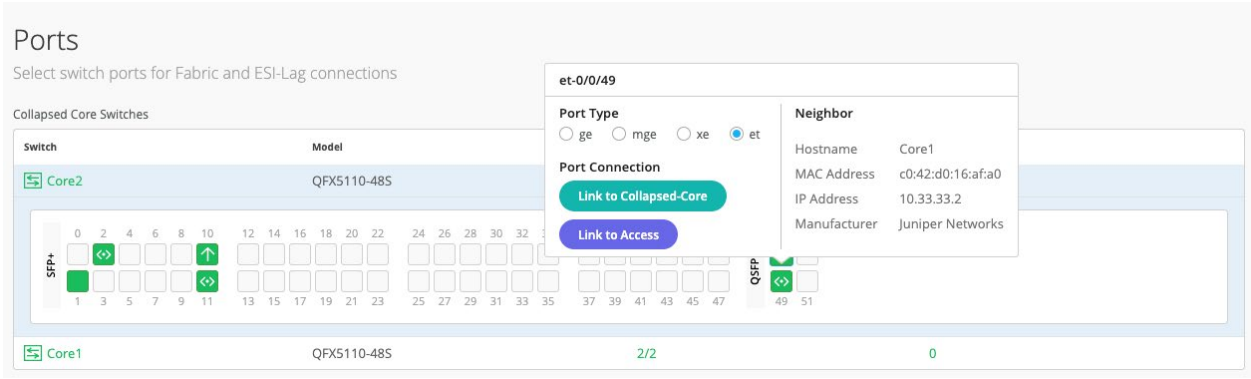

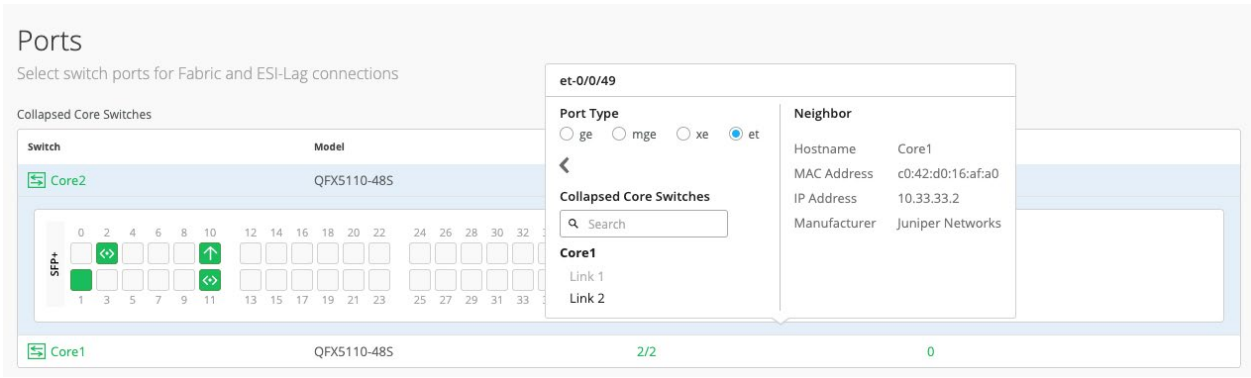

Note that only Link2 is the option left to interconnect with Core1.

You can now select the ports that interconnect the core and access switches.

Core1

You can select xe-0/0/1 as a Link to Access1.

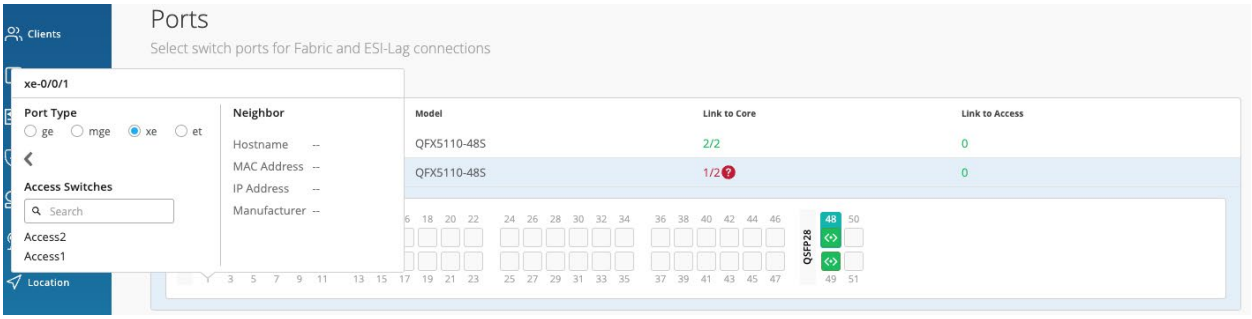

You can select xe-0/0/2 as a Link to Access2.

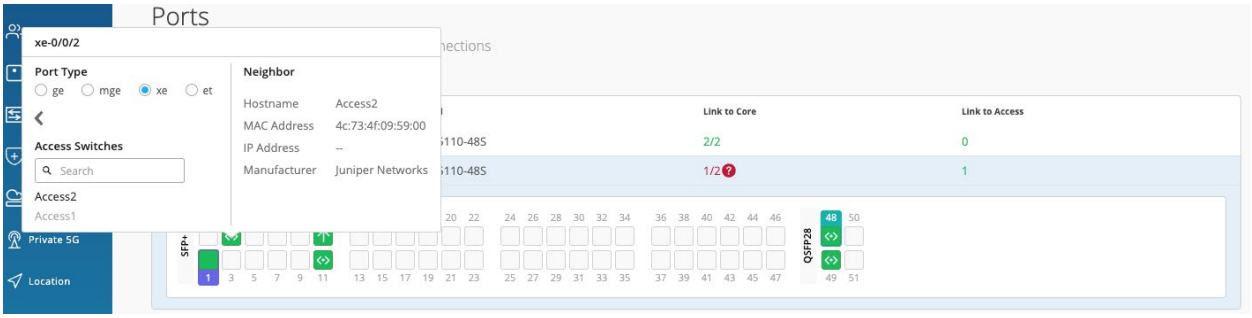

Core2

You can select xe-0/0/1 as a Link to Access2.

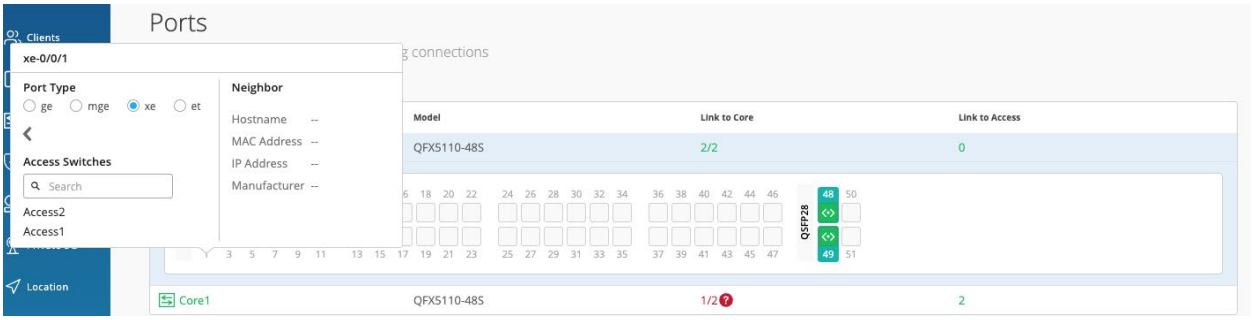

You can select xe-0/0/2 as a Link to Access1.

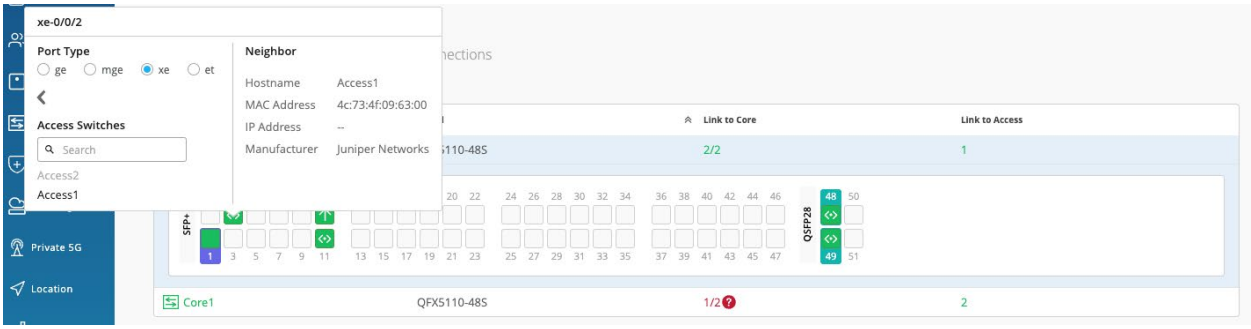

Access Switches

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

You can select both uplinks and interface speed, while allowing Mist to define each AE index. In this case, uplinks xe-0/2/0, xe-0/2/3 are selected as Links to core on both access switches and AE Index 0/1 (system default numbering) on Access 2/1 respectively.

Access1

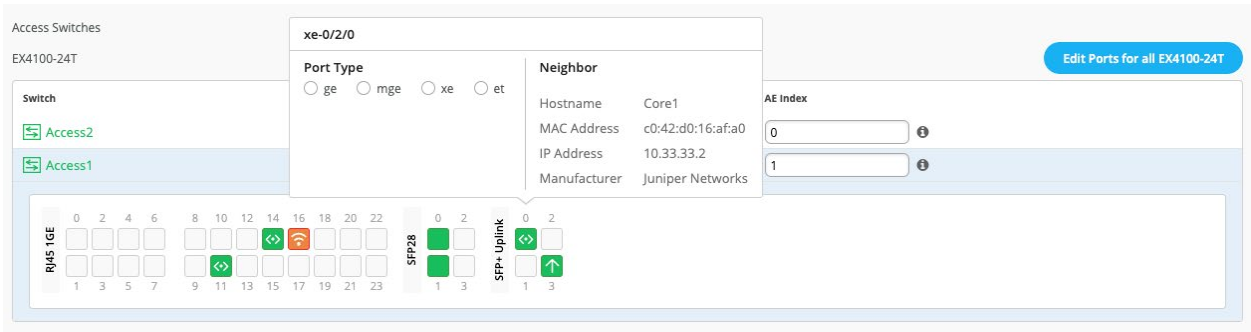

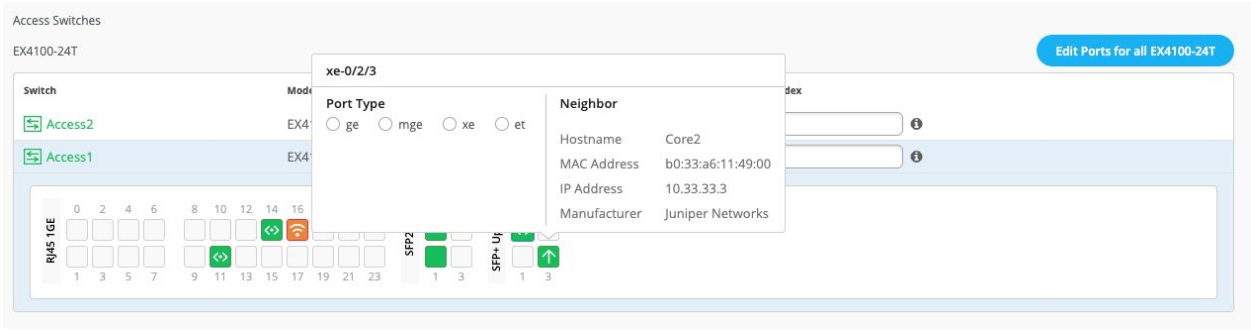

NOTE: LLDP provides you with the core switch name. You can select only the required correct port speed.

Access2

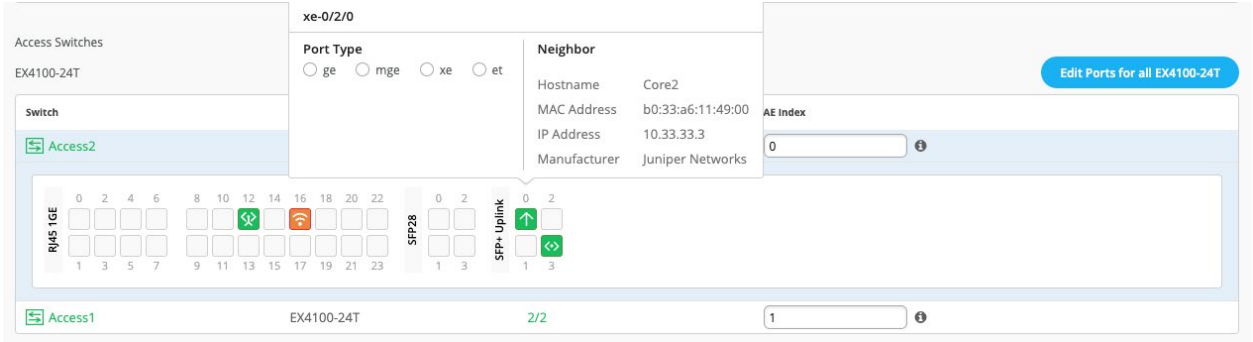

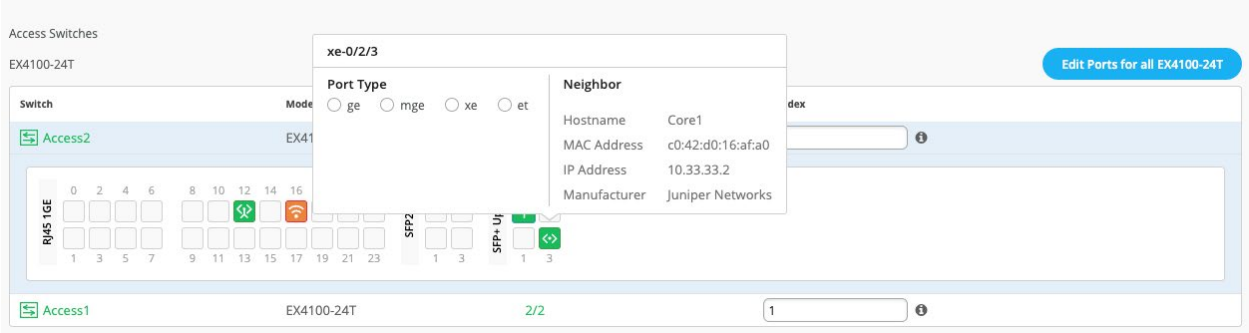

NOTE: LLDP provides you with the core switch name. You can select only the required correct port speed.

After you select all requisite port combinations, you can click the Continue button at the upper-right corner of the Mist UI.

Campus Fabric Configuration Confirmation

On this page, you can confirm each device's configuration as shown below:

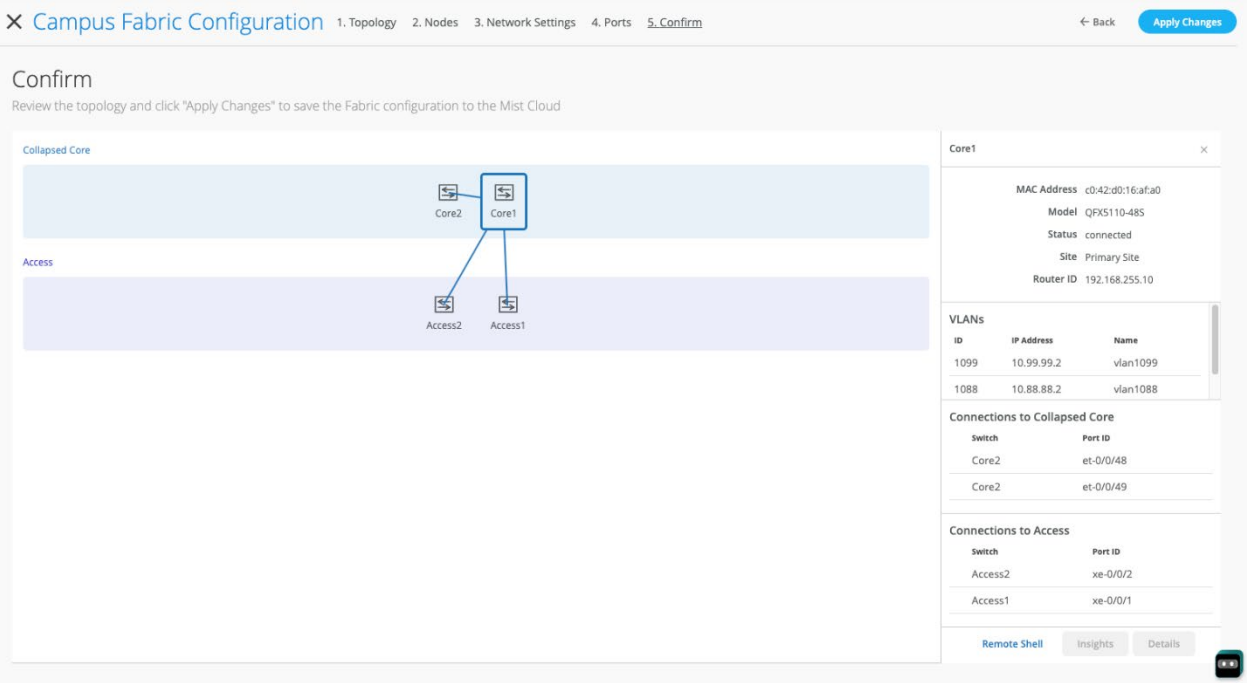

NOTE: The Remote Shell option, available at the lower-right corner, allows you to access each device if needed.

Once you complete the verification, you select the Apply Changes option at the upper-right corner of the Mist UI.

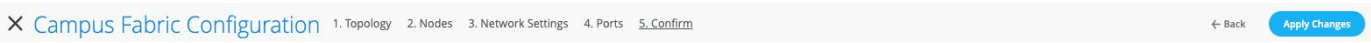

Confirm again to create the fabric.

Mist presents you with the following banner including the estimated time for the Campus Fabric to be built. The process includes the following:

- Mist builds the point-to-point interfaces between Collapsed Core devices with IP addresses chosen from the range presented at the onset of the build.
- Each device is configured with a loopback address from the range presented at the onset of the build.
- eBGP is provisioned at 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 support of customer traffic using EVPN-VXLAN.
- IP addressing of each L3 gateway IRB located on both Collapsed Core switches.
- IP addressing of each lo0.0 loopback.
- Configuration of routing policies for underlay and overlay connectivity.
- Optimized MTU settings for p2p underlay, L3 IRB, and ESI-LAG bundles.
- VXLAN to VLAN mapping using VNI addresses that are automatically assigned.
- VRF creation of corp-it, developers, and guest-wifi and VLAN associated with each VRF.

- VXLAN tunnelling creation between distribution devices and distribution-core devices (in support of the northbound SRX Series Firewall that is configured in the subsequent steps).
- 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.

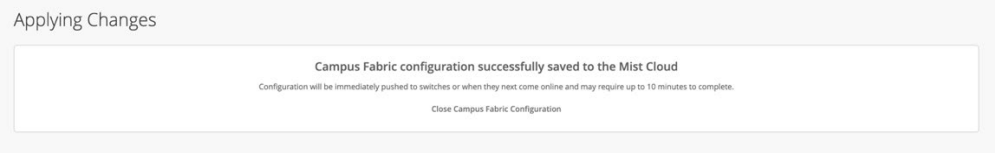

Once you click Close Campus Fabric Configuration, you can view a summary of the newly created Campus Fabric EVPN Multihoming.

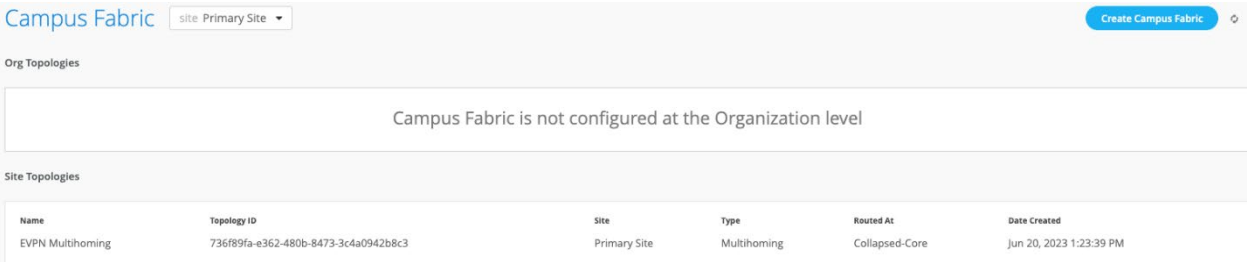

With Juniper Mist Wired Assurance, you download a connection table (.csv format) representing the physical layout of the Campus Fabric. This can be used to validate all switch interconnects for those participating in the physical Campus Fabric build. Once the Campus Fabric is built or is in the process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

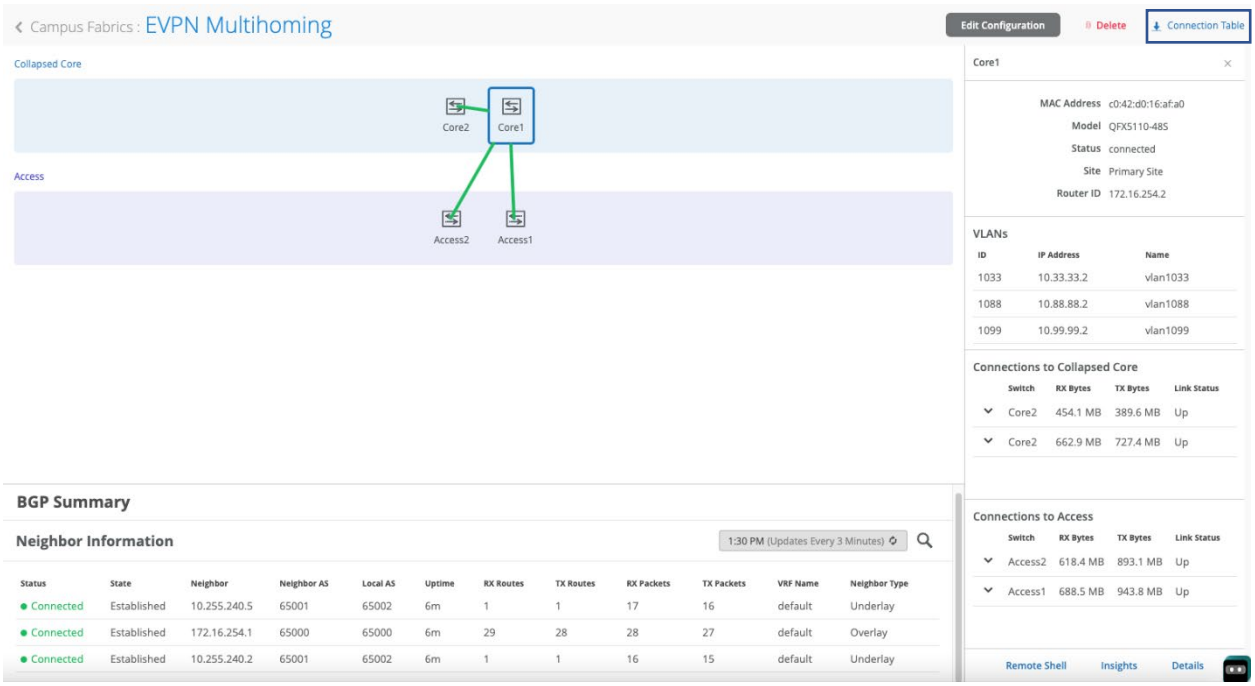

Connection Table spreadsheet:

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
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	FD0822AN0021	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	FD0822AN0021	EX4100-24T	4c734f096300	Access1	access

Apply VLANs to Access Ports

As previously discussed, Mist provides the ability to templatize well known services such as Radius, Network Time Protocol (NTP), and Domain Name System (DNS) 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 Desktop 2 are associated with different ports on each access switch which requires the configuration to be applied to Access1 and Access2 respectively. See Figure 2.

Mist Access Points connect to the same port on Access1/Access2 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. Further, all access switches (defined by Model EX4100 switch as an example) associated the AP port profile with ge-0/0/16 without needing to configure each independent switch.

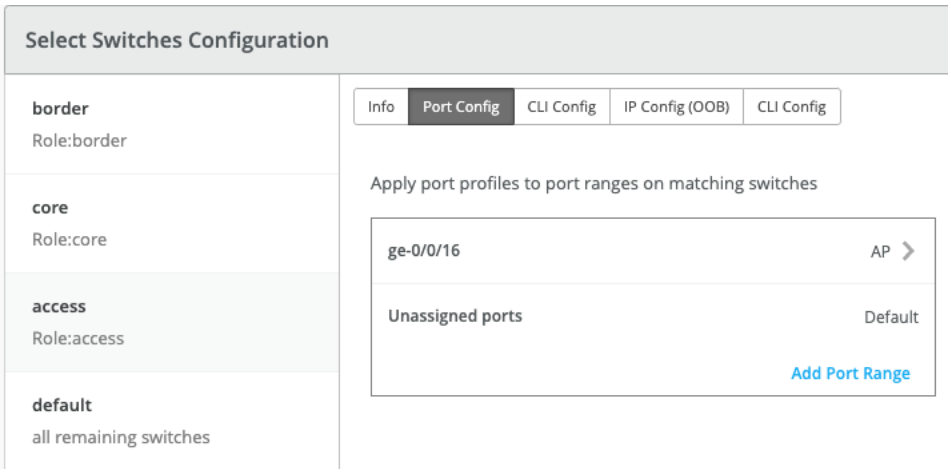

Using Access1 as an example, we apply vlan1099 to port ge-0/0/11 under the Port Configuration section on Access1. In this example, vlan1099 (corp-it), vlan1088 (developers), and vlan1033 (guest-wifi) are defined in the Switch Template. Here, vlan1099 is selected under the configuration profile:

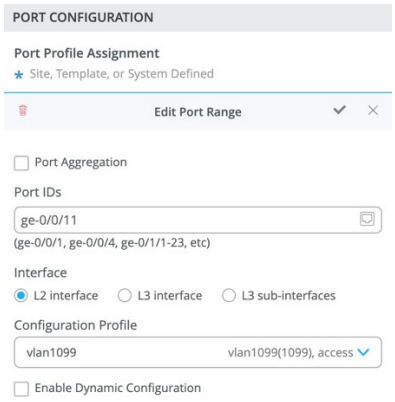

The Switch Template definition for vlan1099 is shown below, representing attributes associated with VLANs such as dot1x authentication, Quality of Service (QoS), and Power over Ethernet. vlan1088 and vlan1033 need to be configured in a similar way.

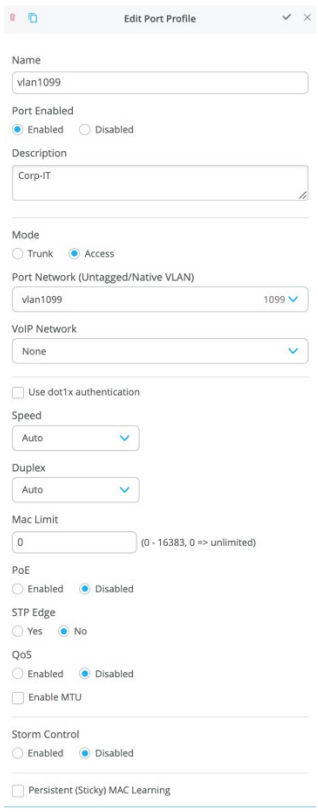
Edit Port Profile
Name
vlan1099
Port Enabled
Enabled
Disabled
Description
Corp-IT
Mode
Trunk
Access
Port Network (Untagged/Native VLAN)
vlan1099
1099
VoIP Network
None
Use dot1x authentication
Speed
Auto
Duplex
Auto
Mac Limit
0
(0 - 16383, 0 => unlimited)
PoE
Enabled
Disabled
STP Edge
Yes
No
QoS
Enabled
Disabled
Enable MTU
Storm Control
Enabled
Disabled
Persistent (Sticky) MAC Learning

Chapter 5 Verification

Verification of the EVPN Multihoming Deployment

For verification of EVPN multihoming deployment, see Figure 2.

There are two desktops to validate the Campus Fabric. Let us take a quick look to see if Desktop1 can connect internally and externally.

```
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
```

Validation Steps

The validation includes the following steps:

- Confirmed local IP address, VLAN and default gateway were configured on Desktop1.
- Can ping default gateway – that indicates that we can reach the access switch.
- Ping to WAN router failed (10.99.99.254) – we need to troubleshoot.

Start by validating Campus Fabric in the Mist UI, by selecting the Campus Fabric option under the Organization tab on the left-hand side of the UI.

Remote shell access into each device within the Campus Fabric is supported here as well as visual representation of the following capabilities:

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

BGP Underlay

Purpose

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

- Load balancing using ECMP for greater resiliency and bandwidth efficiencies.
- bfd, bi-directional forwarding, to decrease convergence times during failures.
- Loopback reachability to support VXLAN tunnelling.

Action

Verify that BGP sessions are established between the core 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 Mist UI or using external applications such as SecureCRT or Putty.

Verification of BGP peering

Core1

Remote Shell can be accessed via the bottom-right of the Campus Fabric, from the switch view or via Secure Shell (SSH).

```
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; time 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 can be pinged 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:

```
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 switches in the Campus Fabric. Loopback reachability is verified between the core devices.

Let us verify that the routes are established between core devices using multiple paths. For example, Dist1 should leverage both paths through Core1/2 to reach Dist2 and vice versa.

Core1: ECMP loopback reachability with 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

```
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 converge in the case of a link or device failure between cores:

```
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>
```

NOTE: At this point, BGP Underlay and Overlay is operational through the verification of BGP between the core devices of the Campus Fabric and that loopback routes are also established between core devices.

EVPN VXLAN Verification Between Core Switches

Since the desktop can ping its default gateway, we can assume the Ethernet-switching tables are correctly populated, VLAN, and interface-mode are correct. If pinging the default gateway fails, then troubleshoot underlay connectivity.

Verification of the EVPN Database on Both Core Switches

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

```
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 in 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 Core1/2 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 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 Core Switches

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

```
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 database.

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>
```

Connectivity between the core and distribution switches looks correct since MAC and ARPs are being learned across the Fabric on both cores. Let us look at the connection between core and WAN router.

External Campus Fabric Connectivity Through the GW QFX5110 Switches

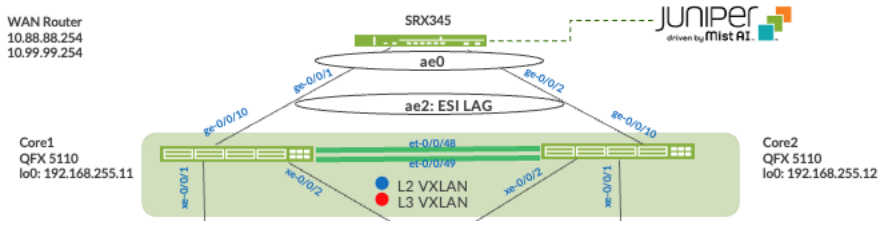

Mist enables the QFX5110 Switch to translate between VXLAN traffic within the Campus Fabric and standard Ethernet switching for external connectivity. In this case, it is a SRX Series Firewall. Let us verify the ESI status on the core switches.

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

We must configure the ESI-LAG for L2 connectivity between the Collapsed Core switches and WAN Router as Mist does not configure this automatically. You can have a pre-defined Port Profile build and associated with the requisite ports on each core switch.

Add a Port profile on core switches interfaces facing the WAN router.

The following represents an existing Port Profile applied to each SRX Series Firewall facing QFX5110 Switch port.

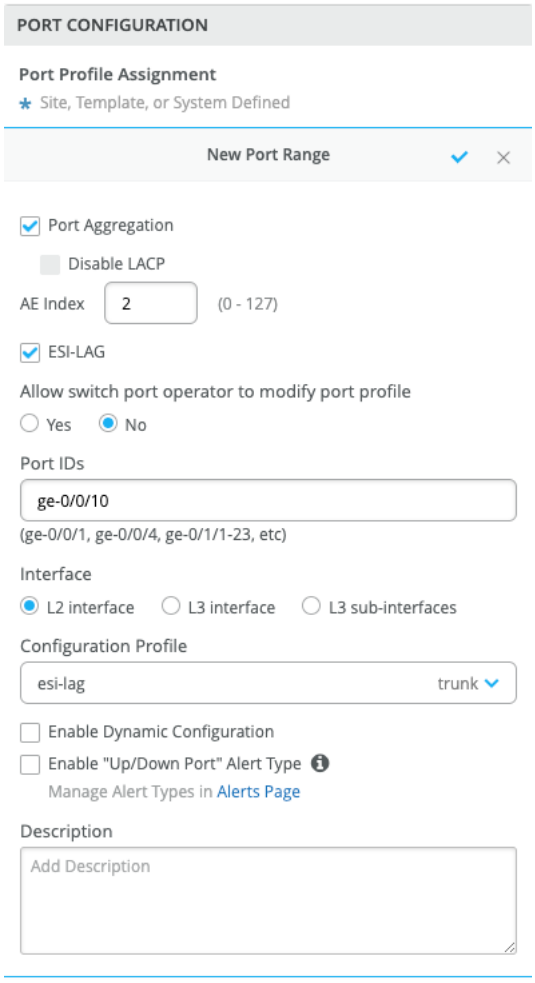

Save the configuration and then verify the changes on the core switch.

Core1: LACP active status to the WAN router produces new entries in the switch/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: LACP active status to the WAN router produces new entries in the switch/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: .254 entries represent the WAN router's Default Gateway IP addresses now found in Core1 and Core2.

We go back to Desktop1 to see if it can cross 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:~#
```

The last step is to 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:~#
```

NOTE: Connectivity within and outside of the Campus fabric is verified. Desktops communicate with each 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 VRF by default while using the SRX Series Firewall to accept or discard inter-VRF communications.

Chapter 6 EVPN Insights

Overview

Juniper Mist Wired Assurance provides you with real-time status related to the health of the EVPN multihoming deployment using telemetry such as BGP neighbor status and TX/RX port statistics. The following screenshots are from the EVPN multihoming build by accessing the Campus Fabric option under the Organization/Wired of the Mist Portal:

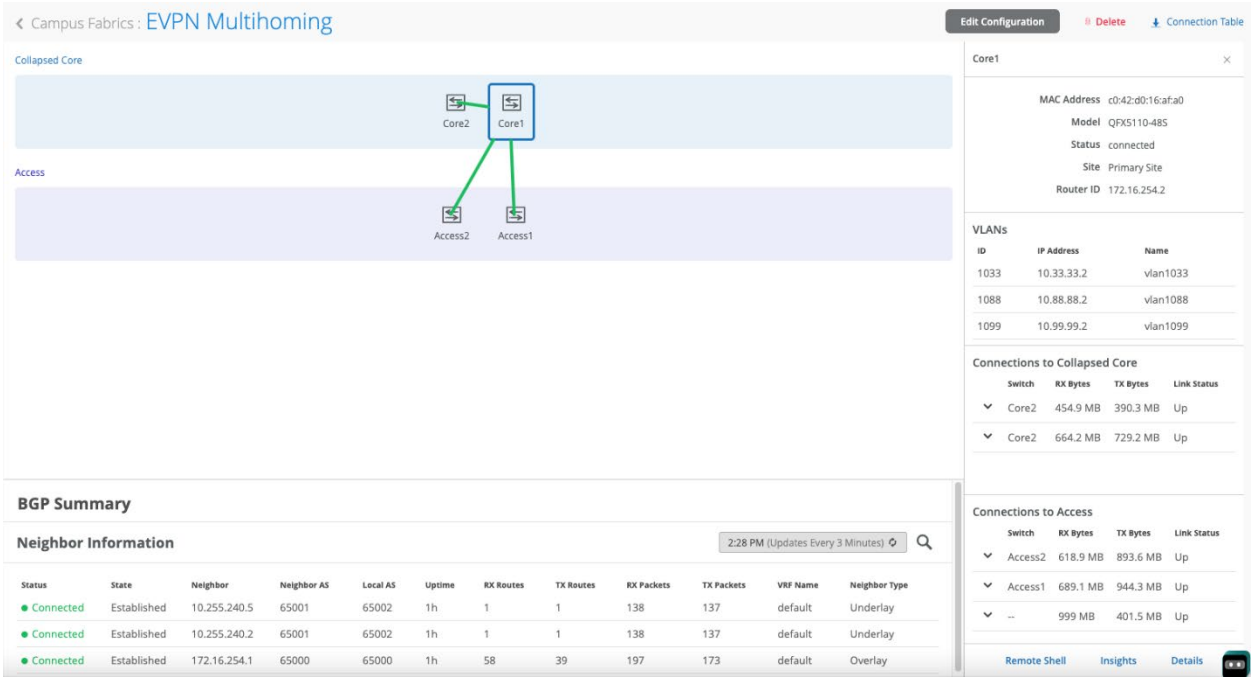

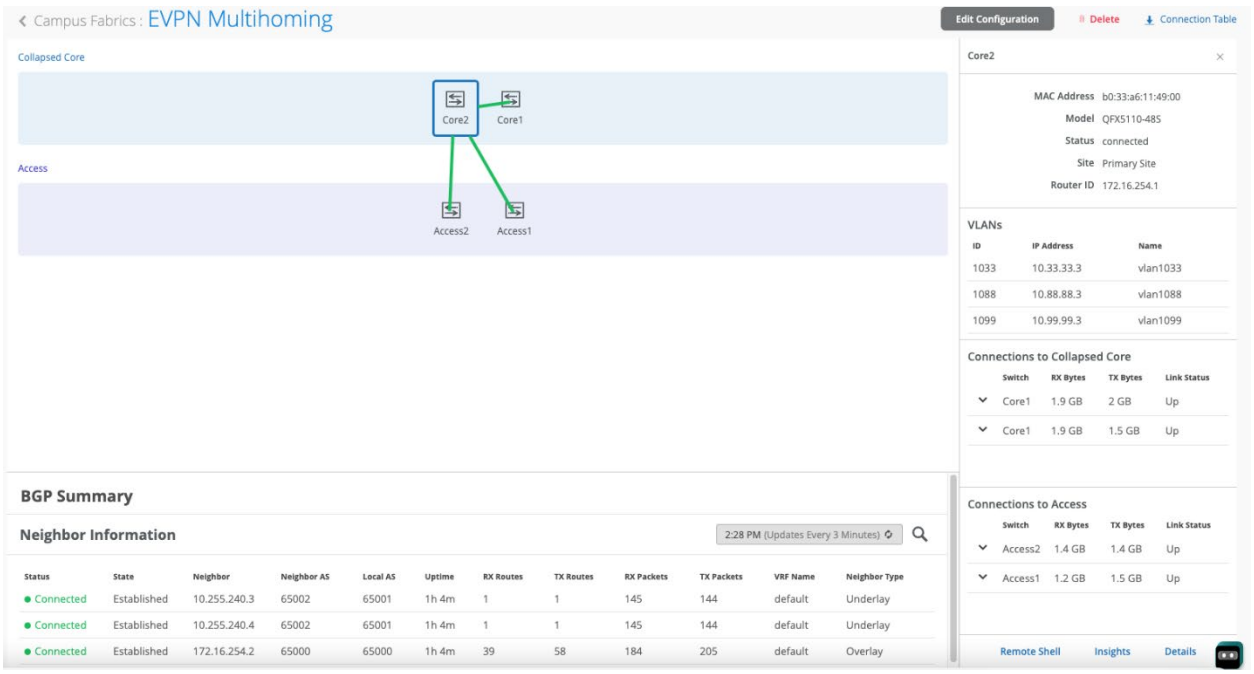

From this view, Mist also provides remote accessibility into each device's console through the Remote Shell option as well as rich telemetry through the Switch Insights option. Remote Shell is demonstrated throughout this document when displaying real-time operational status of each device during the verification stage.

Switch Insights of Core1 displays historical telemetry including BGP peering status critical to the health of the Campus Fabric:

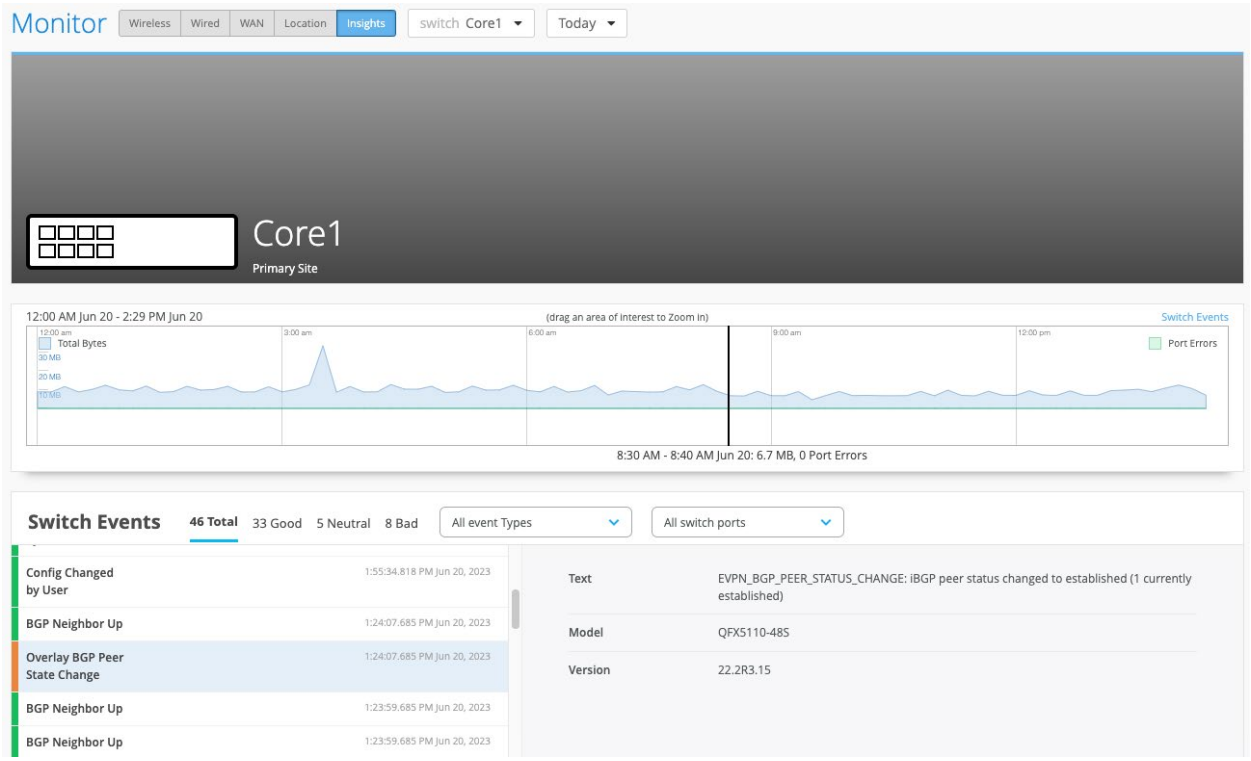

Summary

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 Mist UI. Steps are added in this document to help you understand the troubleshooting steps if deployment is not working correctly.

Chapter 7 Additional Information

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.

Mist provides you with the following options (default in parenthesis):

- BGP Local AS (65001).
- Loopback Prefix (/24).
- Subnet (10.255.240.0/20) – point to point interfaces between core devices.

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 with Core2.

```
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 with Core1.

```
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 172.16.254.2/32
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 with Core1.

```
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 configuration output from the Juniper Mist cloud for the EVPN VXLAN Overlay on the core and distribution switches using eBGP.

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.

Mist provisions L3 IRB interfaces on the distribution layer

Mist enables VXLAN tunnelling, VLAN to VXLAN mapping, and MP BGP configuration snippets such as vrf-targets on core 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. L3 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. L3 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 Core Switches and the Access Switches

This section displays the configuration output from the Juniper Mist cloud for the enablement of the Layer 2 ESI LAG between the distribution switches and access switches. This 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 L2 loop free detection protocols such as Rapid Spanning Tree Protocol (RSTP).

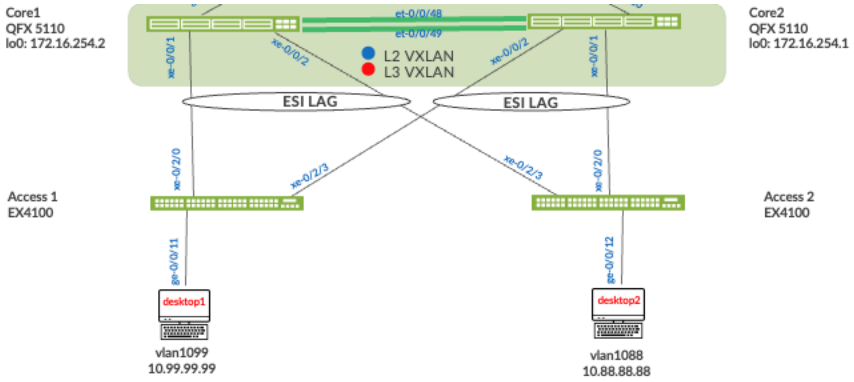

Core1 Configuration

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
```

```
set interfaces xe-0/0/2 unit 0 family ethernet-switching storm-control
default
deactivate interfaces xe-0/0/2 unit 0
```

Core2 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
set interfaces xe-0/0/2 unit 0 family ethernet-switching storm-control
default
deactivate interfaces xe-0/0/2 unit 0
```

Access1 Configuration

1. 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

1. 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 Juniper SRX Series Firewall

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 L2 loop free detection protocols such as RSTP.

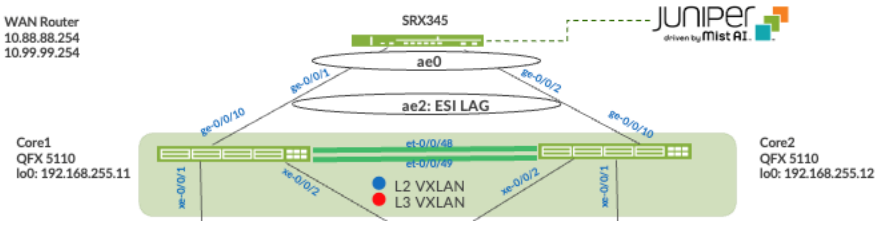

Core 1 Configuration

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

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
```

SRX Series Firewall Configuration

1. 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
```