

Campus Fabric Core-Distribution ERB Using Juniper Mist Wired Assurance— Juniper Validated Design (JVD)

Published
2024-07-22

Table of Contents

About this Document	1
Solution Benefits	1
Solution Architecture	4
Validation Framework	16
Test Objectives	20
Recommendations	21
APPENDIX: Example ERB Fabric Creation	22
APPENDIX: ERB Fabric Verification (Optional)	57
APPENDIX: WAN Router Integration into the Fabric	66
APPENDIX: EVPN Insights	71
APPENDIX: Junos Configuration from This Fabric	74

Campus Fabric Core-Distribution ERB Using Juniper Mist Wired Assurance— Juniper Validated Design (JVD)

Juniper Networks Validated Designs provide customers with a comprehensive, end-to-end blueprint for deploying Juniper solutions in their network. These designs are created by Juniper's expert engineers and tested to ensure they meet the customer's requirements. Using a validated design, customers can reduce the risk of costly mistakes, save time and money, and ensure that their network is optimized for maximum performance.

About this Document

Overview

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

Solution Benefits

IN THIS SECTION

- [Benefits of Campus Fabric Core-Distribution | 2](#)

Enterprise networks are undergoing massive transitions to accommodate the growing demand for cloud-ready, scalable, and efficient network. There's also demand for the plethora of Internet of Things (IoT) and mobile devices. As the number of devices grows, so does network complexity with an ever-greater

need for scalability, segmentation, and security. To meet these challenges, you need a network with automation and artificial intelligence (AI) for operational simplification.

Most traditional campus architectures use single-vendor, chassis-based technologies that work well in small, static campuses with few endpoints. However, they are too rigid to support the scalability and changing needs of modern large enterprises.

A Juniper Networks EVPN-VXLAN fabric is a highly scalable architecture that is simple, programmable, and built on a standards-based architecture (<https://www.rfc-editor.org/rfc/rfc8365>) that is common across campuses and data centers.

The Juniper campus architecture uses a Layer 3 IP-based underlay network and an EVPN-VXLAN overlay network. Broadcast, unknown unicast, and multicast (BUM) traffic is handled natively by EVPN and eliminates the need for spanning-tree protocols (STP or 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.

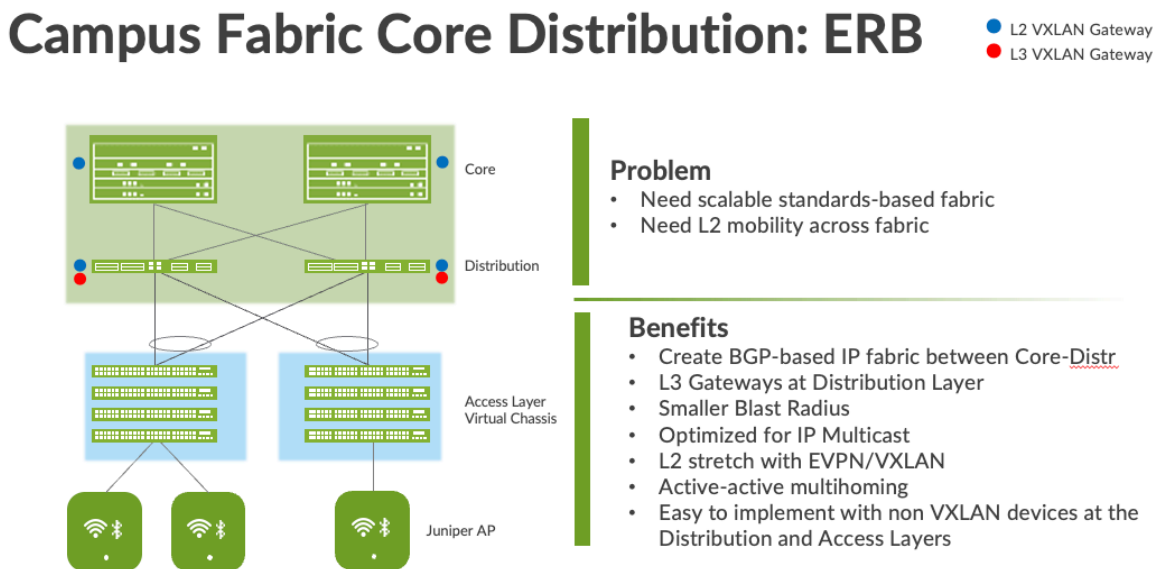
With an EVPN-VXLAN campus architecture, you can easily add core, distribution, and access layer devices as your business grows without a need for redesigning. As EVPN-VXLAN is vendor-agnostic, you can use the existing access layer infrastructure and gradually migrate to access layer switches. This supports EVPN-VXLAN capabilities once the core and distribution part of the network is deployed. Connectivity with legacy switches that do not support EVPN-VXLAN is accomplished with standards-based ESI-LAG.

Benefits of Campus Fabric Core-Distribution

- With the increasing number of devices connecting to the network, you need to scale your campus network rapidly without adding complexity. Many IoT devices have limited networking capabilities and require Layer 2 adjacency across buildings and campuses. Traditionally, this problem was solved by extending VLANs between endpoints using data plane-based flood and learning mechanisms inherent with Ethernet switching technologies. 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 configure and manually manage VLANs to extend them to new network ports. This problem increases multi-fold when you take into consideration the explosive growth of IoT and mobility.
- A campus fabric based on EVPN-VXLAN is a modern and scalable network that uses BGP as the underlay for the core and distribution layer switches. The distribution and core layer switches function as VXLAN Tunnel Endpoints (VTEPs) that encapsulate and decapsulate the VXLAN traffic. In addition, these devices route and bridge packets in and out of VXLAN tunnels.

- The Campus Fabric Core-Distribution extends the EVPN fabric to connect VLANs across multiple buildings. This is done by stretching the Layer 2 VXLAN network with routing occurring in the core (Centrally-Routed Bridging (CRB)) or distribution (Edge Routed Bridging (ERB)) layers. This network architecture supports the core and distribution layers of the topology with integration to access switching via standard Link Aggregation Control Protocol (LACP).

Figure 1: Campus Fabric Core-Distribution ERB



A Campus Fabric Core-Distribution ERB deployment provides the following benefits:

- **Reduced flooding and learning**—Control plane-based Layer 2 and Layer 3 learning reduces the flood and learn issues associated with data plane learning. Learning MAC addresses in the forwarding plane has an adverse impact on network performance as the number of endpoints grows. 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. For example, in a Campus Fabric IP Clos, core switches only learn the access layer switches addresses instead of the device endpoint addresses.
- **Consistency**—A universal EVPN-VXLAN-based architecture across disparate campus and data center deployments enables a seamless end-to-end network for endpoints and applications.
- **Investment protection**—The only requirement to integrate at the access layer is standards-based LACP and LAG. This provides investment protection for the section of the network that has the highest cost and footprint.

- Location-agnostic connectivity—The EVPN-VXLAN campus architecture provides a consistent endpoint experience no matter where the endpoint is located. Some endpoints require Layer 2 reachability, such as legacy building security systems or IoT devices. VXLAN overlay provides Layer 2 extension across campuses without any changes to the underlay network. We use optimal BGP timers between the adjacent layers of the Campus Fabric with Bi-directional Forwarding Detection (BFD) that support fast convergence in the event of a node or link failure and equal-cost multipath (ECMP). For more information, see [Configuring Per-Packet Load Balancing](#).

Solution Architecture

IN THIS SECTION

- [Campus Fabric Core-Distribution High-Level Architecture | 4](#)
- [Underlay Network | 5](#)
- [Understanding EVPN | 7](#)
- [Overlay Network \(Data Plane\) | 8](#)
- [Overlay Network \(Control Plane\) | 9](#)
- [Resiliency and Load Balancing | 10](#)
- [Ethernet Segment Identifier \(ESI\) | 10](#)
- [Services Block | 11](#)
- [Access Layer | 12](#)
- [Single or Multi PoD Design | 12](#)
- [Juniper Access Points | 13](#)
- [Juniper Mist Edge | 14](#)
- [Supported Platforms for Campus Fabric Core-Distribution ERB | 14](#)
- [Juniper Mist Wired Assurance | 15](#)

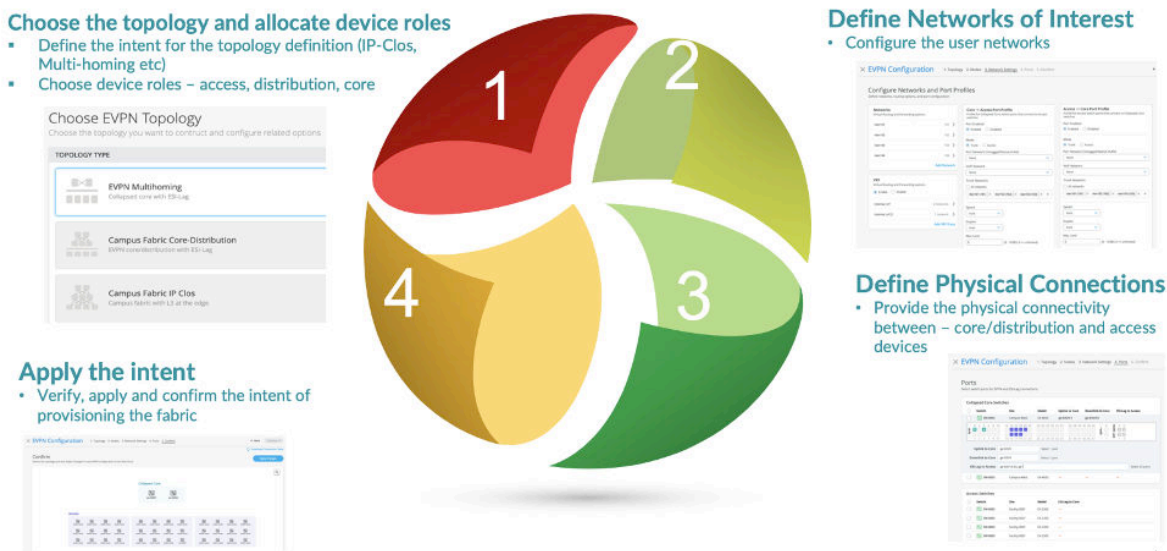
Campus Fabric Core-Distribution High-Level Architecture

The campus fabric, 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. By

configuring different routing instances, you can enforce the separation of virtual networks because each routing instance has its own separate routing and switching table.

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

Figure 2: High-level Campus Fabric Creation




Underlay Network

An EVPN-VXLAN fabric architecture makes the network infrastructure simple and consistent across campuses and data centers. All the core and distribution devices must be connected to each other using a Layer 3 infrastructure. We recommend deploying a Clos-based IP fabric to ensure predictable performance and to enable a consistent, scalable architecture.

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

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

Figure 3: Pt-Pt Links Using /31 Addressing Between Core and Distribution Layers




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

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

Endpoints in different virtual networks require a device that supports inter-VXLAN routing, which is typically a router, or a high-end switch known as a Layer 3 gateway. The entity that performs VXLAN encapsulation and decapsulation is called a VXLAN tunnel endpoint (VTEP). Each VTEP is known as the Layer 2 gateway and typically assigned with the device's loopback address. This is also where VXLAN (commonly known as VNI) to VLAN mapping exists.

Figure 4: VXLAN VTEP Tunnels



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

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

Understanding EVPN

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

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

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

The benefits of using EVPNs include:

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

EVPN provides multipath forwarding and redundancy through an all-active model. The core layer can connect to two or more distribution devices and forward traffic using all the links. If a distribution link or core device fails, traffic flows from the distribution layer toward the core layer using the remaining active links. For traffic in the other direction, remote core devices update their forwarding tables to send traffic to the remaining active distribution devices connected to the multihomed Ethernet segment.

The technical capabilities of EVPN include:

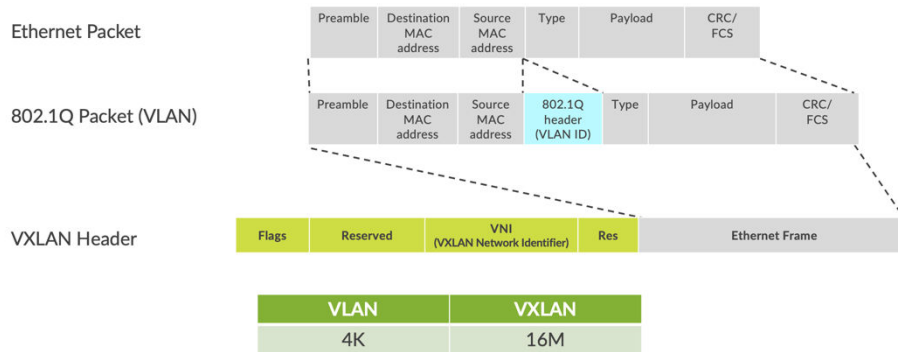
- Minimal flooding—EVPN creates a control plane that shares end host MAC addresses between VTEPs.
- Multihoming—EVPN supports multihoming for client devices. A control protocol like EVPN that enables synchronization of endpoint addresses between the distribution switches is needed to support multihoming,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

Overlay Network (Data Plane)

VXLAN is the overlay data plane encapsulation protocol that tunnels Ethernet frames between network endpoints over the underlay network. Devices that perform VXLAN encapsulation and decapsulation for the network are referred to as a VXLAN tunnel endpoints (VTEPs). Before a VTEP sends a frame into a

VXLAN tunnel, it wraps the original frame in a VXLAN header that includes a VNI. The VNI maps the packet to the original VLAN at the ingress switch. After applying a VXLAN header, the frame is encapsulated into a UDP/IP packet for transmission to the remote VTEP over the IP fabric, where the VXLAN header is removed and the VNI to VLAN translation happens at the egress switch.

Figure 5: VXLAN Header



VTEPs are software entities tied to a device's loopback address that source and terminate VXLAN tunnels. VXLAN tunnels in a core distribution fabric are provisioned on the following:

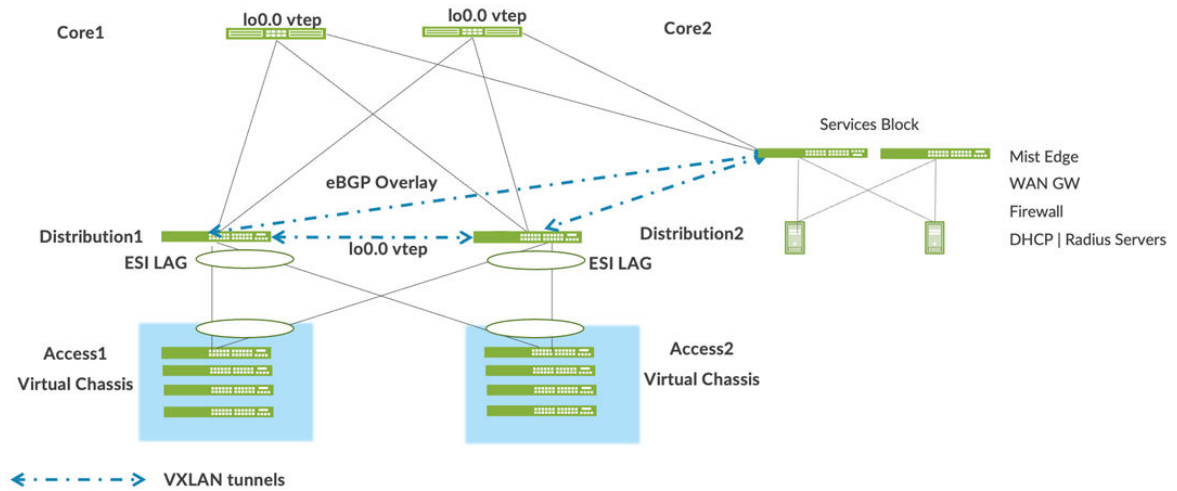
- Distribution switches to extend services across the campus fabric.
- Core switches, when acting as a border router, interconnect the campus fabric with the outside network.
- Services Block devices that interconnect the campus fabric with the outside network.

Overlay Network (Control Plane)

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

Figure 6: EVPN VXLAN Overlay Network with a Services Block

Overlay Control Plane



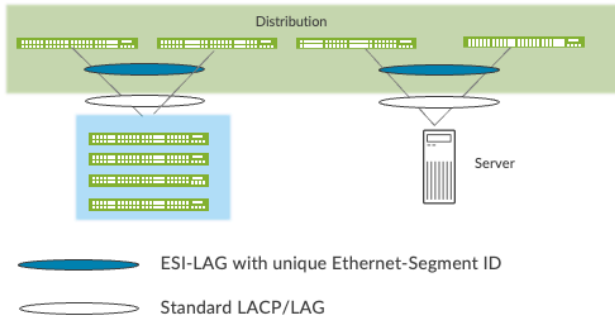
Resiliency and Load Balancing

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

Ethernet Segment Identifier (ESI)

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

Figure 7: Resiliency and Load Balancing

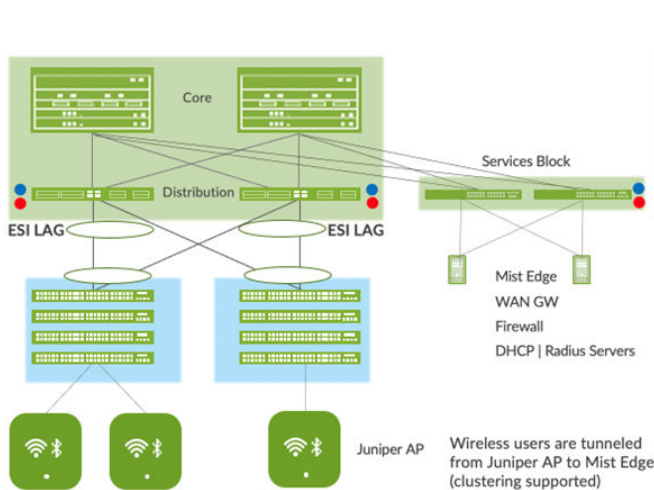


- EVPN supports N-way "scale-out" Ethernet multihoming
- No ICL link required between Distribution Switches
- Virtual Chassis LAG spread across multiple switches in VC stack
- Active-Active Multihoming
- Multi-homed devices such as Servers are identified in the overlay by unique Ethernet Segment ID (ESI)

Services Block

You might wish to position critical infrastructure services off of a dedicated access pair of Juniper switches. This can include WAN and firewall connectivity, RADIUS, and DHCP servers as an example. For those who wish to deploy a lean core, the dedicated services block mitigates the need for the core to support encapsulation and de-encapsulation of VXLAN tunnels as well as additional capabilities such as routing instances and additional L3 routing protocols. The services block border capability is supported directly off of the core layer or as a dedicated pair of switches.

Figure 8: Services Block



Problem

- Segment critical services in a dedicated Access Switch Pair
- WAN Router/Firewall/Infrastructure services connectivity


Benefits

- Leverage ECMP/Load Balancing to access critical services provided off the Services Block
- Horizontal Scale

Access Layer

The access layer provides network connectivity to end-user devices, such as personal computers, VoIP phones, printers, IoT devices, as well as connectivity to wireless access points. In this Campus Fabric Core-Distribution design, the EVPN-VXLAN network extends between the core and distribution layer switches.

Figure 9: End Point Access

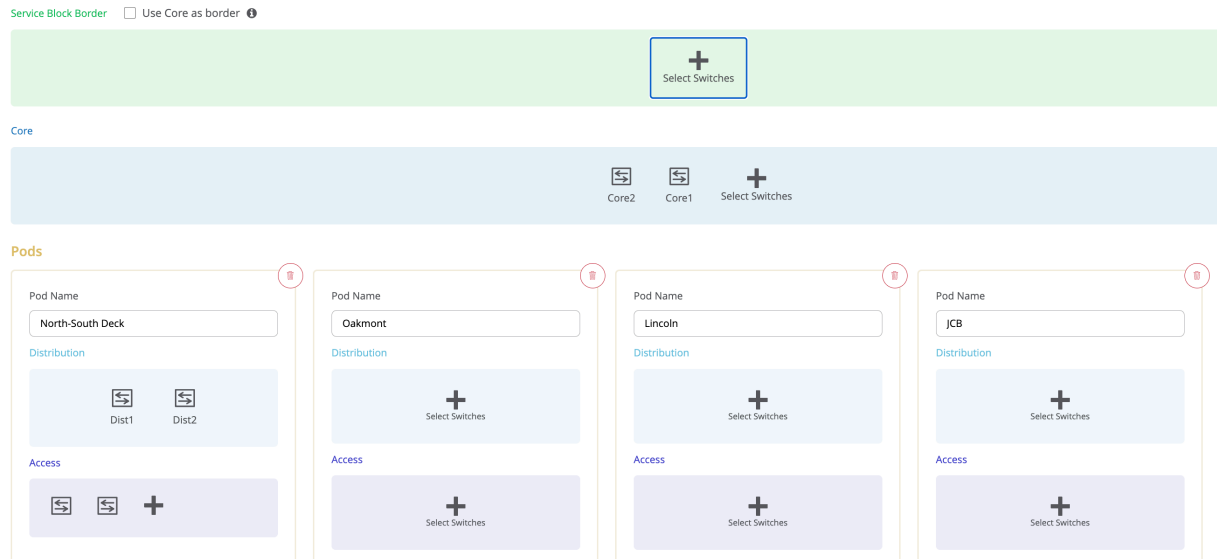


In this example, each access switch or Virtual Chassis is multihomed to two or more distribution switches. Juniper's Virtual Chassis reduces the number of ports required on distribution switches and optimizes availability of fiber throughout the campus. The Virtual Chassis is also managed as a single device and supports up to 10 devices (depending on switch model) within a Virtual Chassis. With EVPN running as the control plane protocol, any distribution switch can enable active-active multihoming to the access layer. EVPN provides a standards-based multihoming solution that scales horizontally across any number of access layer switches.

Single or Multi PoD Design

Juniper Mist Campus Fabric supports deployments with only one PoD (formally called Site-Design) or multiple PoD's. The organizational deployment shown below, targets enterprises who need to align with a multi-POD structure:

Figure 10: Multiple PoD Design Example



NOTE: The example in the appendix section of this document only shows a single PoD for an easy start design.

Juniper Access Points

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

- **Wired and wireless assurance**—Mist is enabled with Wired and Wireless Assurance. Once configured, service-level expectations (SLEs) for key wired and wireless performance metrics such as throughput, capacity, roaming, and uptime are addressed in the Mist platform. This JVD uses Juniper Mist Wired Assurance services.
- **Marvis**—An integrated AI engine that provides rapid wired and wireless troubleshooting, trending analysis, anomaly detection, and proactive problem remediation.

Juniper Mist Edge

For large campus networks, Juniper Mist™ Edge provides seamless roaming through on-premises tunnel termination of traffic to and from the Juniper access points. Juniper Mist Edge extends select microservices to the customer premises while using the Juniper Mist™ cloud and its distributed software architecture for scalable and resilient operations, management, troubleshooting, and analytics. Juniper Mist Edge is deployed as a standalone appliance with multiple variants for different size deployments.

Evolving IT departments look for a cohesive approach to managing wired, wireless, and WAN networks. This full stack approach simplifies and automates operations, provides end-to-end troubleshooting, and ultimately evolves into the Self-Driving Network™. The integration of the Mist platform in this JVD addresses both challenges. For more details on Mist integration and Juniper Networks® EX Series Switches, see [How to Connect Mist Access Points and Juniper EX Series Switches](#).

Supported Platforms for Campus Fabric Core-Distribution ERB

Table 1: Supported Platforms for Campus Fabric Core-Distribution ERB deployment

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

Table 1: Supported Platforms for Campus Fabric Core-Distribution ERB deployment (Continued)

Campus Fabric Core-Distribution ERB Deployment	Supported Platforms
Core layer	EX4650
	QFX5120
	QFX5130
	QFX5700
	QFX10000
	EX92xx
Services block	EX4400-24X
	EX4650
	QFX5120
	QFX5130
	QFX5700
	EX92xx

Juniper Mist Wired Assurance

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 Junos streaming telemetry that enable the insights for switch health metrics and anomaly detection, as well as 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 transforming IT operations from reactive troubleshooting to proactive remediation.

Juniper Mist cloud services are 100% programmable using open APIs for full automation and/or integration with your operational support systems. For example, IT applications such as ticketing systems and IP management systems.

Juniper Mist delivers unique capabilities for the WAN, LAN, and wireless networks:

- UI or API-driven configuration at scale.
- Service-level expectations (SLEs) 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.

To learn more 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>

Validation Framework

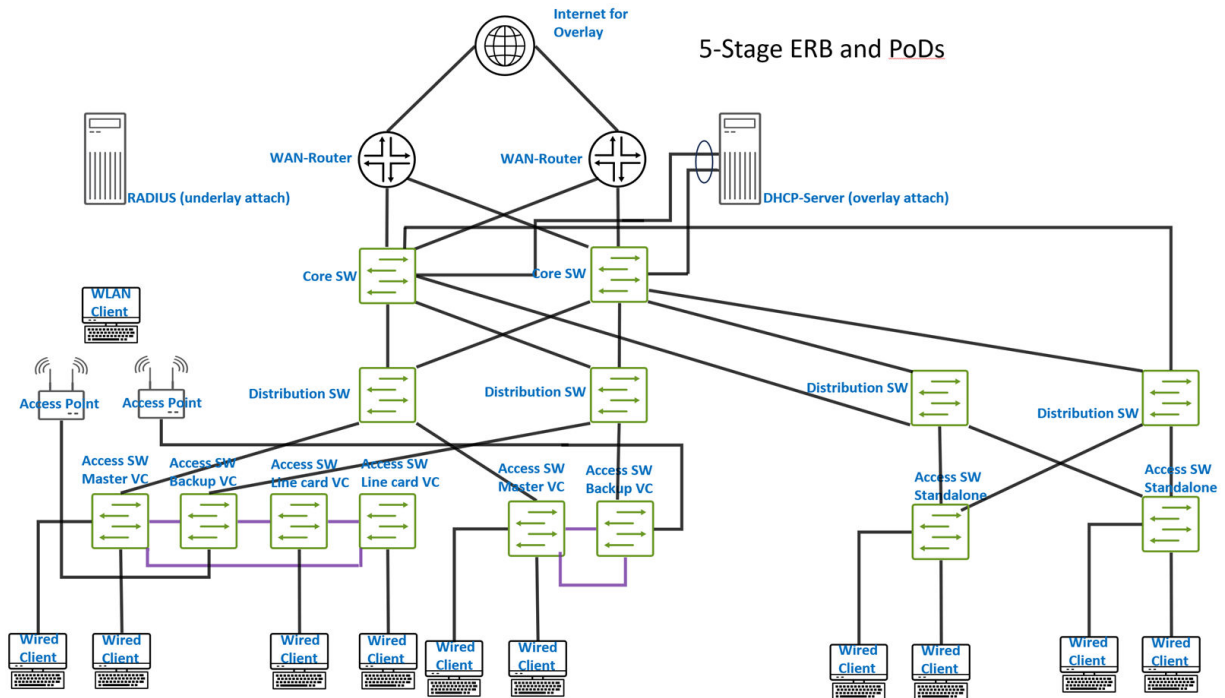
IN THIS SECTION

- Test Bed | 16
- Platforms / Devices Under Test (DUT) | 19
- Test Bed Configuration | 20

Test Bed

In the diagram below, you will see the suggested topology used for the phase 2 lab evaluating an ERB fabric with multiple PoDs design.

Figure 11: JVD Lab Proposal



The suggested lab design provides the ability to evaluate the following:

- Five-stage ERB multi-site fabric with:
 - Two redundant core switches acting as spines.
 - Pod1/Building1:
 - Two redundant distribution switches acting as leaves.
 - One 4 Member Virtual Chassis Access Switch acting as ToR.
 - One 2 Member Virtual Chassis Access Switch acting as ToR.
 - Pod2/Building2:
 - Two redundant distribution switches acting as leaves.
 - Two standalone access switches acting as ToR.
- Service block function via:
 - Integrated to existing core switches (default) acting as service leaf and core at the same time.
 - Attached WAN routers via Layer 2 or Layer 3 exit.
 - Attached servers via ESI-LAG redundant links.

- WAN router integration:
 - Layer 2 fabric exit.
 - ESI-LAG-based trunks.
 - Layer 3 fabric exit.
 - OSPF as routing protocol.
 - eBGP as routing protocol
 - Attached to:
 - Core switch.
 - Redundant WAN router design:
 - Two Juniper MX routers.
 - Two Juniper SRX Firewalls in cluster configuration.
- Wi-Fi Access Points:
 - Local-attached to the access switches with Power over Ethernet (PoE).
 - Various Wi-Fi clients.
 - Basic Wi-Fi roaming.
- Overlay server attached to service block functionality:
 - DHCP server.
 - Other services.
- RADIUS server:
 - Server location:
 - Local server attached to underlay network.
 - Remote Juniper Mist Access Assurance via public cloud.
 - Authentication for the following clients:
 - Wired clients attached to access switches.
 - Wi-Fi clients using the access points.
 - Authentication based on Clients:
 - MAC address.
 - 802.1X EAP authentication.

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

Platforms / Devices Under Test (DUT)

The devices tested and on which Junos version is reported in the table below:

Table 2: Devices Under Test

Devices Under Test		
Fabric function	Device	Junos Release
Core Switches	QFX5120	22.4R3-S2
Distribution Switches Pod1	QFX5120	22.4R3-S2
Distribution Switches Pod2	QFX5120	22.4R3-S2
Access Virtual Chassis Pod1	EX4400	22.4R3-S2
Access Virtual Chassis Pod1	EX4100	22.4R3-S2
Access Switch Pod2	EX3400	22.4R3-S2

Table 2: Devices Under Test *(Continued)*

Devices Under Test		
Fabric function	Device	Junos Release
Access Switch Pod2	EX3400	22.4R3-S2
WAN-Router	SRX1500	

Test Bed Configuration

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

Test Objectives

IN THIS SECTION

- Test Goals | 20
- JVD Non-Goals | 21

Test Goals

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

- Testing with Junos version 22.4R3-S2.
- Testing with multiple PoDs.
- Testing with Virtual Chassis that have 2 or 4 members.

- Testing with features that are activated as combinations at the same time.

The scale testing for this design was done with:

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

JVD Non-Goals

Following are the non-goals for the current JVD qualification:

- Testing this fabric with redundant WAN routers. This is already described in a separate JVD extension in common for all fabrics.
- Five-stage ERB single site testing as this was already performed in phase 1.
- Juniper Mist Edge integration for Wi-Fi scaling.

Recommendations

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

- Review the JVD extension for WAN router integration.
 - For this fabric type, we recommend using the Layer 3 eBGP integration approach.
- All fabric networks should be configured in the following way to avoid inconsistency:
 - First, create them as part of your switch template for a site.
 - Then, import the created networks as part of the campus fabric dialogue and assign to VRFs.
 - Even if the system allows you a local network creation on a switch, do not use this option.
- Do not manually configure VRFs locally on any switch. The fabric usually does this automatically on an as-needed basis.
 - The current exception to this rule is Layer 2 WAN router integration via transport VLAN. Please review the JVD extension for WAN router integration and follow the example in the appendix.

- When using DHCP relay, configuration for the fabric:
 - Configure a “Loopback per-VRF subnet” pool range.
 - Include the pool range as sharing host routes with your WAN router as the loopback IP addresses get assigned as /32 across all of the VRFs shared on the fabric.
 - Only use the fabric dialogue for configuring DHCP relay and no local configuration directly on a switch.
 - There is a planned JVD extension which covers DHCP relay configuration. Please look to see if it is available.
- Consider Juniper Mist Edge integration when you have more than 2,000 wireless clients.
 - Each Juniper Mist Edge should only connect via LAG to a single service block function.
 - Design for stickiness at a single Juniper Mist Edge until failover must happen.
- Unassigned access ports should be configured with a quarantine VLAN via a template.
 - If possible, use a different VRF for the quarantine VLAN to isolate this traffic.
 - Best practice is also enabling “STP Edge” in the quarantine Port Profile.

APPENDIX: Example ERB Fabric Creation

IN THIS SECTION

- [Campus Fabric Core-Distribution ERB Components | 23](#)
- [Juniper Mist Wired Assurance | 24](#)
- [Juniper Mist Wired Assurance Switches | 25](#)
- [Templates | 26](#)
- [Topology | 33](#)
- [Create the Campus Fabric | 33](#)
- [Apply VLANs to Access Ports | 55](#)

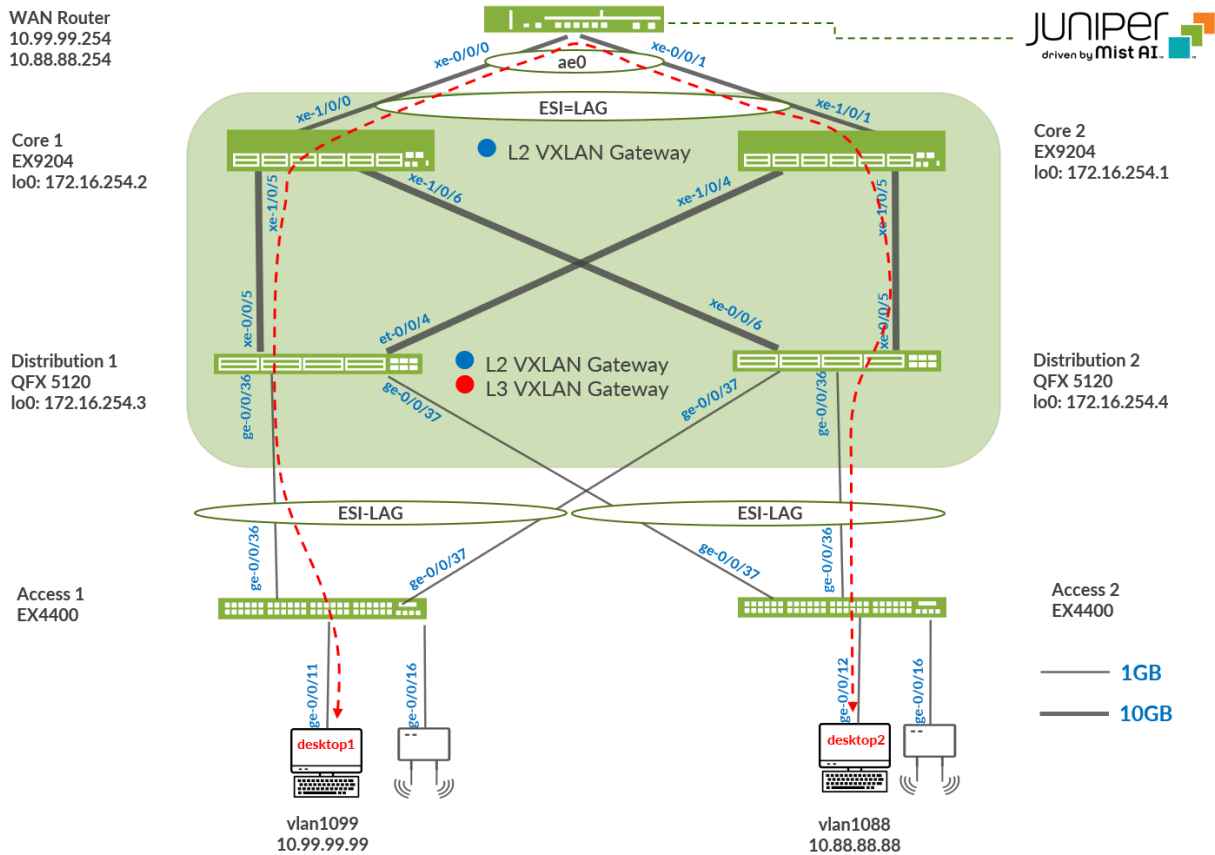
NOTE: The configuration example shown below was made as part of phase 1 of this JVD. Please review the topology of phase 2. We've kept this example to introduce how to build a minimal ERB fabric.

Campus Fabric Core-Distribution ERB Components

This configuration example uses the following devices:

- Two EX9204 switches as core devices, software version: Junos OS Release 22.4R2-Sx or later.
- Two QFX5120 switches as distribution devices, software version: Junos OS Release 22.4R2-Sx or later.
- Two access layer EX4400 switches, software version: Junos OS Release 22.4R2-Sx or later.
- One MX80 WAN router, software version: Junos OS Release 21.2R3-Sx or later.
- Juniper access points.
- Two Linux desktops that act as wired clients.

Figure 12: Figure : Topology



Juniper Mist Wired Assurance

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

Juniper switch onboarding to the Juniper Mist cloud:

<https://www.juniper.net/documentation/us/en/quick-start/hardware/cloud-ready-switches/topics/topic-map/step-1-begin.html>

Juniper Mist Wired Assurance, through the portal, is used to build a Campus Fabric Core-Distribution ERB from the ground up. This includes the following:

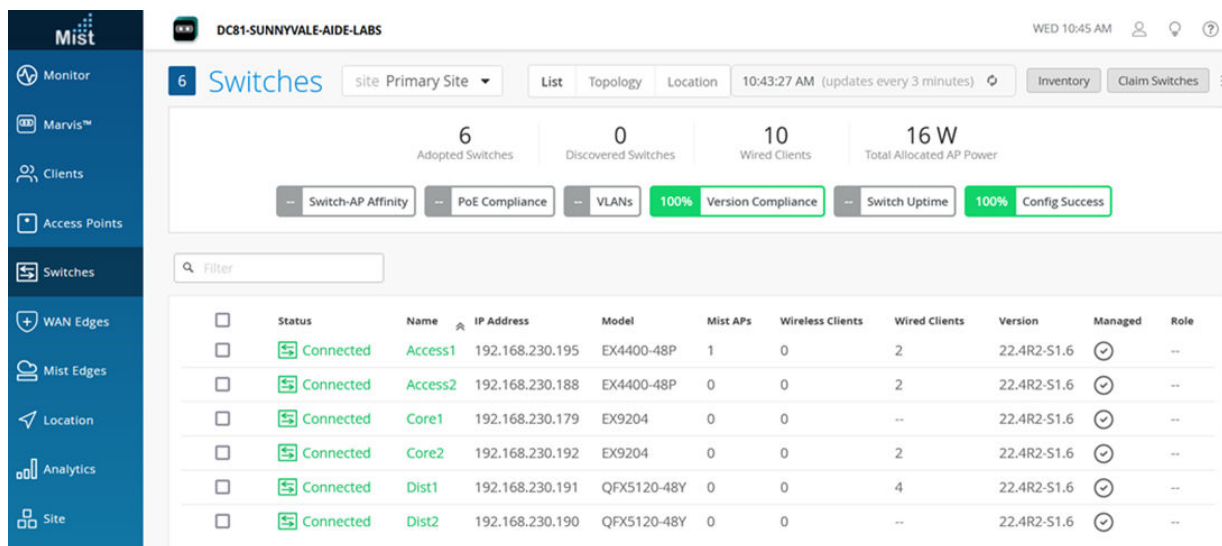
- Assignment of P2P links between the core and distribution layers.
- Assignment of unique BGP AS numbers per device participating in the underlay and overlay.
- The creation of VRF instances allows you to logically segment traffic. This also includes the assignment of new or existing VLANs to each representative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

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

Juniper Mist Wired Assurance Switches

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

Figure 13: Switch Inventory



Templates

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

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

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

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

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

Figure 14: Adding Additional CLI

The screenshot shows a window titled "CLI CONFIGURATION" with a maximize button. Below the title bar, there is a section labeled "Additional CLI Commands" with an information icon. A large, empty text input field is provided for entering commands.

Under organization and switch templates, we use the following template:

Figure 15: Switch Templates

The screenshot shows a table titled "Switch Templates" with the following data:

TEMPLATE	SITES	SWITCHES
campus-fabric	1	6

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

```
{
  "ntp_servers": [],
  "dns_servers": [
    "8.8.8.8",
    "9.9.9.9"
  ],
  "dns_suffix": [],
  "additional_config_cmds": [],
  "networks": {
    "vlan1099": {
      "vlan_id": 1099,
      "subnet": "10.99.99.0/24"
    },
    "vlan1088": {
      "vlan_id": 1088,
      "subnet": "10.88.88.0/24"
    },
    "vlan1033": {
      "vlan_id": 1033,
```

```
    "subnet": "10.33.33.0/24"
  }
},
"port_usages": {
  "myaccess": {
    "mode": "trunk",
    "disabled": false,
    "port_network": "vlan1033",
    "voip_network": null,
    "stp_edge": false,
    "port_auth": null,
    "all_networks": false,
    "networks": [
      "vlan1033",
      "vlan1088",
      "vlan1099"
    ],
    "speed": "auto",
    "duplex": "auto",
    "mac_limit": 0,
    "persist_mac": false,
    "poe_disabled": false,
    "enable_qos": false,
    "storm_control": {},
    "mtu": 9018,
    "description": ""
  },
  "myesilag": {
    "mode": "trunk",
    "disabled": false,
    "port_network": null,
    "voip_network": null,
    "stp_edge": false,
    "port_auth": null,
    "all_networks": true,
    "networks": [],
    "speed": "auto",
    "duplex": "auto",
    "mac_limit": 0,
    "persist_mac": false,
    "poe_disabled": false,
    "enable_qos": false,
    "storm_control": {}
  }
}
```

```
"mtu": 9014,
  "description": ""
},
"dynamic": {
  "mode": "dynamic",
  "reset_default_when": "link_down",
  "rules": []
},
"vlan1099": {
  "mode": "access",
  "disabled": false,
  "port_network": "vlan1099",
  "voip_network": null,
  "stp_edge": false,
  "all_networks": false,
  "networks": null,
  "port_auth": null,
  "speed": "auto",
  "duplex": "auto",
  "mac_limit": 0,
  "persist_mac": false,
  "poe_disabled": false,
  "enable_qos": false,
  "storm_control": {},
  "mtu": 9014,
  "description": "Corp-IT",
  "disable_autoneg": false,
  "mac_auth_protocol": null,
  "enable_mac_auth": null,
  "mac_auth_only": null,
  "guest_network": null,
  "bypass_auth_when_server_down": null
},
"vlan1088": {
  "mode": "access",
  "disabled": false,
  "port_network": "vlan1088",
  "voip_network": null,
  "stp_edge": false,
  "all_networks": false,
  "networks": null,
  "port_auth": null,
  "speed": "auto",
```

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



```

"rules": [
  {
    "name": "core",
    "match_model": "EX9204",
    "port_config": {},
    "additional_config_cmds": [
      ""
    ],
    "ip_config": {
      "type": "dhcp",
      "network": "default"
    },
    "oob_ip_config": {
      "type": "dhcp",
      "use_mgmt_vrf": false
    }
  },
  {
    "name": "distribution",
    "port_config": {},
    "additional_config_cmds": [
      ""
    ],
    "ip_config": {
      "type": "dhcp",
      "network": "default"
    },
    "oob_ip_config": {
      "type": "dhcp",
      "use_mgmt_vrf": false
    },
    "match_model[0:7]": "QFX5120"
  },
  {
    "name": "access",
    "port_config": {
      "ge-0/0/16": {
        "usage": "myaccess",
        "dynamic_usage": null,
        "critical": false,
        "description": "",
        "no_local_overwrite": true
      }
    }
  }
]

```

```
    },
    "additional_config_cmds": [
      ""
    ],
    "ip_config": {
      "type": "dhcp",
      "network": "default"
    },
    "oob_ip_config": {
      "type": "dhcp",
      "use_mgmt_vrf": false
    },
    "match_model[0:6]": "EX4400"
  }
],
},
"switch_mgmt": {
  "config_revert_timer": 10,
  "root_password": "juniper123",
  "protect_re": {
    "enabled": false
  },
  "tacacs": {
    "enabled": false
  }
},
"radius_config": {
  "auth_servers": [],
  "acct_servers": [],
  "auth_servers_timeout": 5,
  "auth_servers_retries": 3,
  "fast_dot1x_timers": false,
  "acct_interim_interval": 0,
  "auth_server_selection": "ordered",
  "coa_enabled": false,
  "coa_port": ""
},
"vrf_config": {
  "enabled": false
},
"remote_syslog": {
  "enabled": false
},
}
```

```

"snmp_config": {
  "enabled": false
},
"dhcp_snooping": {
  "enabled": false
},
"acl_policies": [],
"mist_nac": {
  "enabled": true,
  "network": null
},
"name": "campus-fabric"
}

```

Topology

Juniper Mist Wired Assurance provides the template for LAN and loopback IP addressing for each core and distribution 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 and distribution devices within the Campus Fabric Core-Distribution. Devices such as the access layer switches connect to the distribution layer using standard LAGs; while the distribution uses ESI-LAG in a multihoming, load balancing manner.

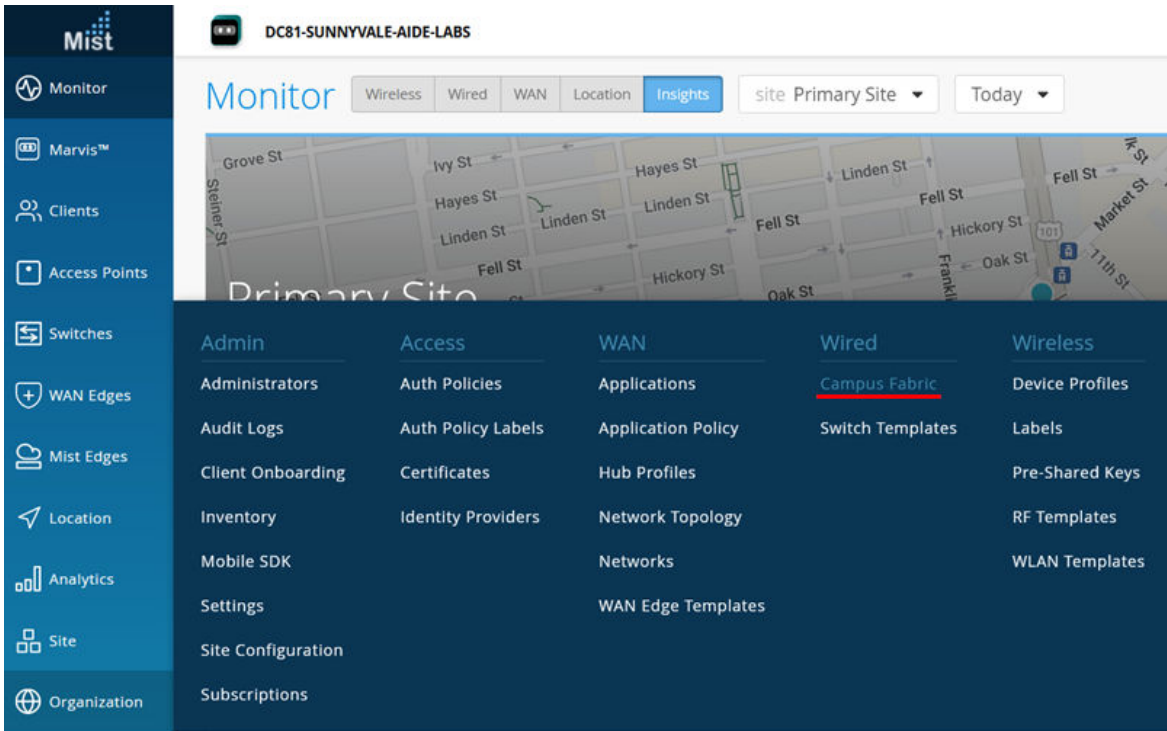
The WAN router can be provisioned through the portal but is separate from the campus fabric workflow. The WAN router has a southbound LAG configured to connect to the ESI-LAG on the core switches. WAN routers can be standalone or built as a high availability cluster. In this document, a single MX router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

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

Create the Campus Fabric

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

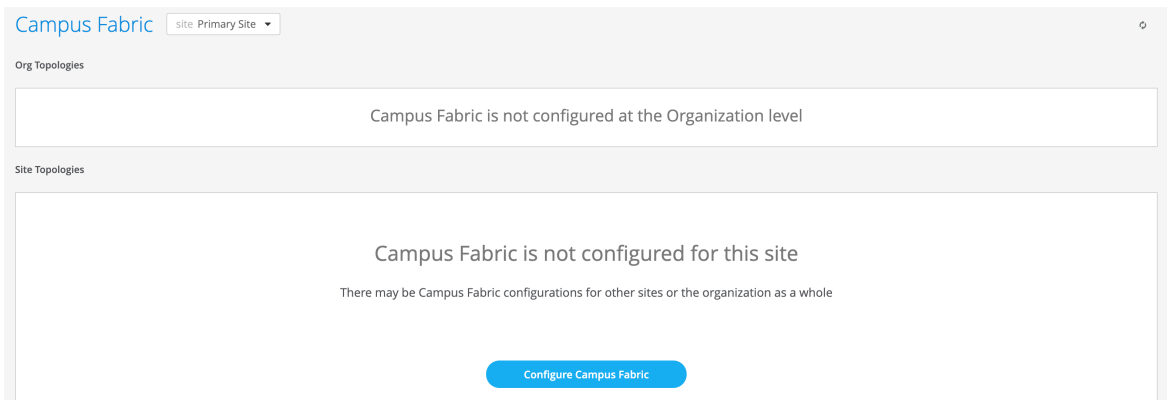
Figure 16: Campus Fabric Creation



Mist provides the option of deploying a campus fabric at the organizational or site level noted in the upper-left side of the campus fabric menu shown below. Both designs now allow you to build fabrics with just a single PoD or multiple PoDs based on customer requirements to connect multiple buildings.

In our example here, the fabric was built on the site level with a single PoD only.

Figure 17: Fabric Site Level Creation



2. Choose the Campus Fabric Topology

Select the Campus Fabric Core-Distribution option below:

Figure 18: ERB Fabric Creation

Campus Fabric Configuration 1. Topology 2. Nodes 3. Network Settings 4. Ports 5. Confirm ← Back Continue

Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options

TOPOLOGY TYPE

- EVPN Multihoming**
Collapsed core with ESI-Lag
- Campus Fabric Core-Distribution**
EVPN core/distribution with ESI-Lag
- Campus Fabric IP Clos**
Campus fabric with L3 at the edge

CONFIGURATION

Topology Name
Campus Fabric ERB

Topology Sub-type

- CRB
Centrally-routed and bridged with gateways on the Core
- ERB
Edge-routed and bridged with anycast gateways on the fabric edge

TOPOLOGY SETTINGS

BGP Local AS
65001
(2-byte or 4-byte)

Subnet
10.255.240.0/20
(xxx.xxx.xxx.xxx/xx)

Auto Router ID Subnet
172.16.254.0/23
(xxx.xxx.xxx.xxx/xx)

Loopback per-VRF subnet
172.16.192.0/19
(xxx.xxx.xxx.xxx/xx)

Mist provides a section to name the Campus Fabric Core-Distribution ERB:

- Configuration—Provide a name in accordance with company standards
- Topology Sub-type:
 - CRB
 - ERB

NOTE: ERB uses anycast addressing which provides shared IP addresses among all distribution layer devices participating in the L3 IRB. Deployments that require a routing protocol on the IRB must use CRB with virtual gateway addressing. You must choose CRB if

most network traffic is north-south while ERB should be selected if mainly east-west traffic patterns exist.

Topology Settings

- **BGP Local AS**—Represents the starting point of private BGP AS numbers that are automatically allocated per device. You can use whatever private BGP AS number range suits your deployment, routing policy is provisioned by Mist to ensure the AS numbers are never advertised outside of the fabric.
- **Subnet**—Represents the pool of IP addresses used for point-to-point links between devices. You can use whatever range suits your deployment. Mist breaks this subnet into /31 subnet addressing per link. This number can be modified to suit the specific deployment scale. For example, /24 provides up to 128 P2P /31 subnets.
- **Auto Router ID Subnet**—Represents the pool of IP addresses associated with each device's loopback address. Each device will automatically get a loopback IP address of /32 assigned from this pool. You can use whatever range suits your deployment. VXLAN tunnelling using a VTEP is associated with this address. The loopback IP addresses assigned here are only visible in the underlay transport network. The definition of these underlay loopback IP addresses is critical for the operation of the EVPN-VXLAN fabric to function at all.
- **Loopback per-VRF-subnet**—Represents a second pool of loopback IP addresses which are each associated with an L3 VRF and switch of the overlay fabric network. It is designed for scale-out services in the overlay network where some services, like DHCP relay, share a single IP address external to the fabric. This is the case for anycast fabrics like ERB and IP Clos. If those L3 VRFs use a dedicated loopback IP address per VRF and switch, it is easy to send return answers to an originating VRF/switch.

NOTE: In previous documentation, you did not have the default configuration fields for auto router ID subnet and loopback per-VRF subnet. Instead, you had a field for loopback prefix definition like shown below and then you had to assign the loopback IPs for each fabric node manually. This has changed in favor of automatic loopback assignments via the configuration of the prefix pool.

Figure 19: Older Fabric Configuration Options

TOPOLOGY SETTINGS

BGP Local AS

(2-byte or 4-byte)

Loopback prefix ?

Subnet ?

(xxx.xxx.xxx.xxx/xx)

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

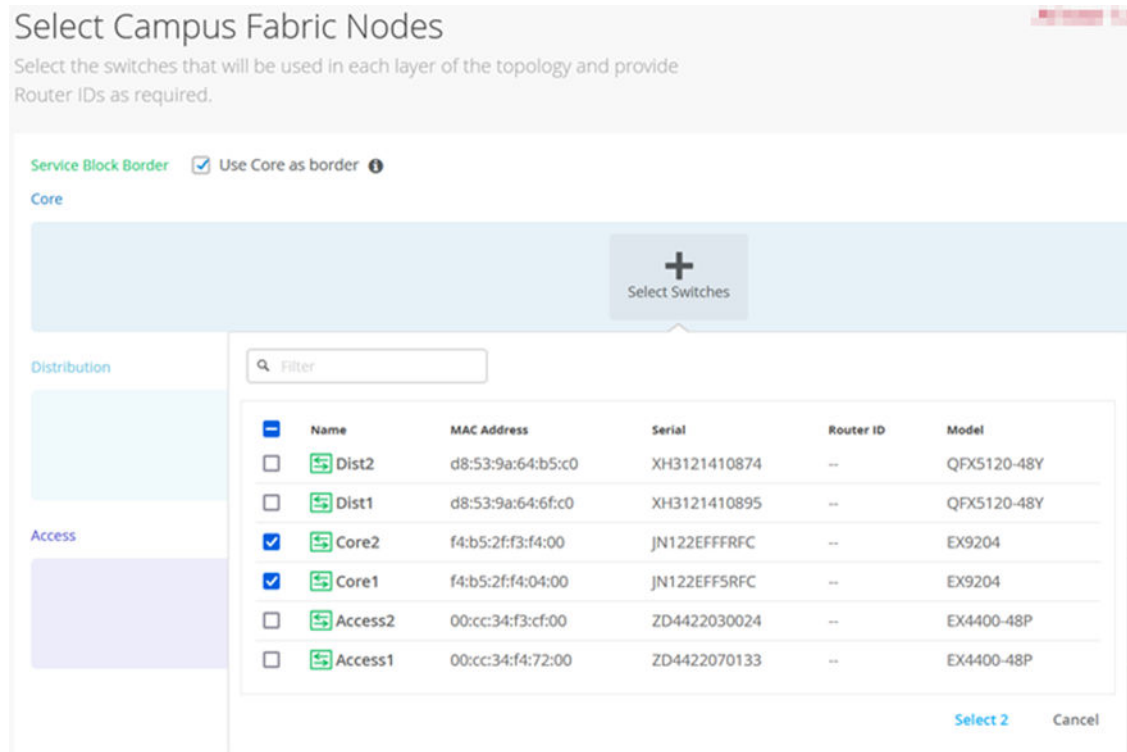
3. Select Campus Fabric Nodes

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

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

The services block router is where the campus fabric interconnects external devices such as firewalls, routers, or critical devices. For example, DHCP and RADIUS servers. Devices to which external services connect to the campus fabric are known as border leafs. If you want to connect these services or devices to the Campus Fabric Core-Distribution ERB in a separate device or pair of devices, clear the **Use Core as border** option and select the **Select Switches** option to choose the devices.

Figure 20: Select the Fabric Nodes



NOTE: Placing the services block router on a dedicated pair of switches (or single switch) alleviates the encapsulation and de-capsulation of VXLAN headers from the core layer. If you want to combine this capability within the core devices, you must select the **Use Core as border** option.

- Once all layers have selected the appropriate devices, you must provide an underlay loopback IP address for each device (with the exception of the access switches). This loopback interface is associated with a logical construct called a VTEP; used to source the VXLAN tunnel. The Campus Fabric Core-Distribution ERB has VTEPs for VXLAN tunnelling on the distribution switches and the core switches when enabling the core border option.

When defining an auto router ID subnet prefix, the underlay loopback IP address and router ID assignments happens automatically. There is no need to manually assign them. You may still see warnings like the one below about a 5 router ID, you can ignore those as the automatic assignment happens at a later phase.


Figure 21: Router ID Not Assigned Yet


Select Campus Fabric Nodes


Select the switches that will be used in each layer of the topology and provide Router IDs as required.

Service Block Border Use Core as border ⓘ


Core *required



Core1



Core2


Select Switches


Distribution *required



Dist1



Dist2


Select Switches

Access


Access1


Access2


Select Switches

Core1

MAC Address f4:b5:2f:f4:04:00

Model EX9204

Status connected

Site Primary Site

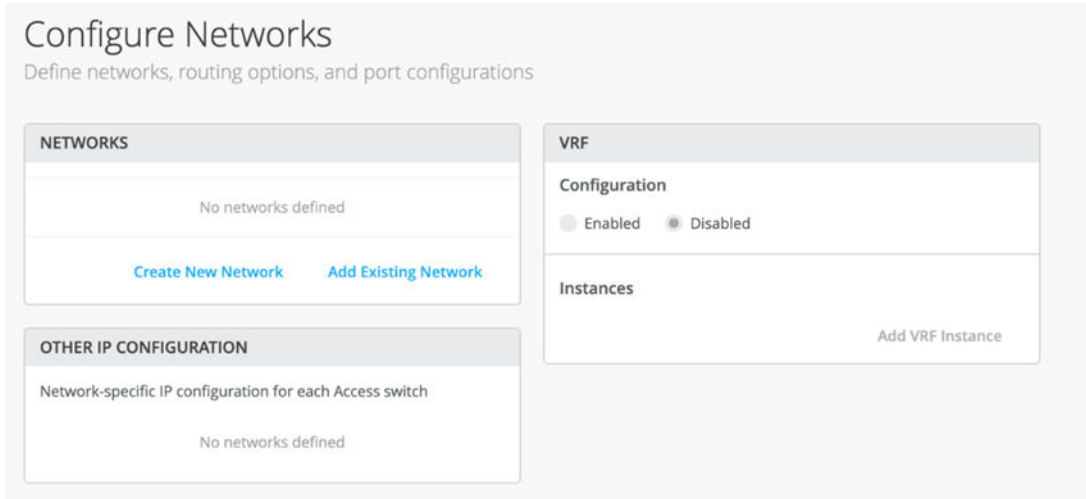
Router ID ---

NOTE: If the auto router ID subnet field is not configured or is empty, you can use the previous mode of operation and manually assign the underlay loopback IP addresses as router IDs to each device needing one. Make sure that all IP addresses are in the same subnet as required by the Mist cloud fabric configuration.

5. Configuring Networks

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

Figure 22: Configure Networks

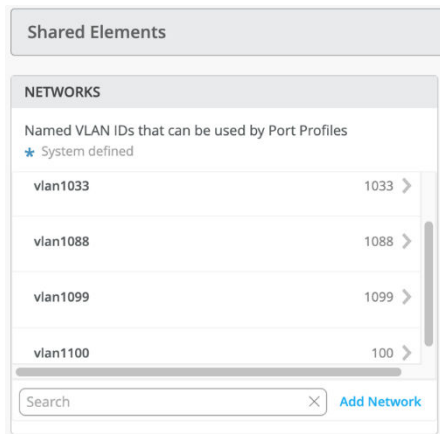


6. Networks

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

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

Figure 23: Networks Inherited by Switch Template



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

Figure 24: Network Import from Template

Import from Template

Template

Campus Fabric :3 Networks

<input checked="" type="checkbox"/>	Name	VLAN ID
<input checked="" type="checkbox"/>	vlan1033	1033
<input checked="" type="checkbox"/>	vlan1088	1088
<input checked="" type="checkbox"/>	vlan1099	1099

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

Figure 25: Edit a Network

NETWORKS

Edit Network

Name

vlan1099

VLAN ID

1099

(1 - 4094 or {{siteVar}})

Subnet

10.99.99.0/24

Other IP Configuration

Juniper Mist Wired Assurance provides automatic IP addressing for IRB interfaces for each of the VLANs. Port profiles and port configurations then associate the VLAN with specified ports. In this case, we selected campus fabric ERB at the onset of the campus fabric build.

Figure 26: ERB Selection

CONFIGURATION

Topology Name

Campus Fabric ERB

Topology Sub-type

CRB
Centrally-routed and bridged with gateways on the Core

ERB
Edge-routed and bridged with anycast gateways on the fabric edge

This option uses anycast addressing for all devices participating in the L3 subnet. In this case, Dist1 and Dist2 switches are configured with the same IP address for each L3 subnet.

More information on anycast gateways can be found here: <https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-mclag-irb-gateway-anycast-address.html>

Figure 27: Anycast GW on Distribution Switch 1

OTHER IP CONFIGURATION	
Network-specific IP configuration for each Distribution switch	
Edit Dist1	✓ ✕
vlan1033	10.33.33.1 >
vlan1088	10.88.88.1 >
vlan1099	10.99.99.1 >

Figure 28: Anycast GW on Distribution Switch 2

OTHER IP CONFIGURATION	
Network-specific IP configuration for each Distribution switch	
Edit Dist2 ✓ ✕	
vlan1033	10.33.33.1 >
vlan1088	10.88.88.1 >
vlan1099	10.99.99.1 >

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

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

Figure 29: Enable VRF

VRF

Configuration

Enabled Disabled

Instances

No VRF instances defined

[Add VRF Instance](#)

Figure 30: Assign Network to VRF

VRF

New VRF Instance ✓ ✕

Name

Networks

Extra Routes

No extra routes defined

[Add Extra Routes](#)

By default, inter-VRF communications are not supported within the campus fabric. If inter-VRF communications is required, each VRF can include extra routes such as a default route that instructs the campus fabric to use an external router or firewall for further security inspection or routing capabilities. In this example, all traffic is trunked over the ESI-LAG and the MX router handles inter-VRF routing. See [Figure 11: Topology on page 24](#).

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

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

Figure 31: Add Default Route

The screenshot shows a 'New Extra Route' dialog box with a title bar containing a checkmark and a close button. Below the title bar, there are two input fields. The first field is labeled 'Route' and contains the text '0.0.0.0/0'. The second field is labeled 'Via' and contains the text '10.99.99.254'.

10. Create two additional VRFs:
 - a. The developers VRF using vlan 1088 with 0.0.0.0/0 utilizing 10.88.88.254
 - b. The guest-wifi VRF using vlan 1033 with 0.0.0.0/0 utilizing 10.33.33.254

Figure 32: Entire Network and VRF Configuration

The screenshot displays two panels of the Mist configuration interface. The left panel, titled 'NETWORKS', lists three VLANs: 'vlan1033' (ID 1033), 'vlan1088' (ID 1088), and 'vlan1099' (ID 1099). Below the list are links for 'Create New Network' and 'Add Existing Network'. The right panel, titled 'VRF', shows the 'Configuration' section with 'Enabled' selected. Under the 'Instances' section, three VRFs are listed: 'corp-it' (1 network), 'developers' (1 network), and 'guest-wifi' (1 network). A link for 'Add VRF Instance' is visible at the bottom of the VRF panel.

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

Figure 33: Fabric LAG Configuration

DISTRIBUTION / ACCESS PORT CONFIGURATION

Port configuration for ESI-Lag between Distribution and Access switches

Name

erb-lag

Trunk Networks

vlan1033(1033) x vlan1088(1088) x vlan1099(1099) x

+

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

Figure 34: Fabric LAG

Port Enabled
 Enabled Disabled

Description
 Add Description

Mode
 Trunk Access

Port Network (Untagged/Native VLAN)
 None

Speed
 Auto

Duplex
 Auto

Mac Limit
 0
 (0 - 16383, 0 => unlimited)

PoE
 Enabled Disabled

STP Edge
 Yes No

QoS
 Enabled Disabled

Enable MTU
 9100
 (256 - 9216)

Storm Control
 Enabled Disabled

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

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

Configure Campus Fabric Ports

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

Figure 35: Port Overview

Ports
Select switch ports for Fabric connections

Core Switches

Switch	Model	Link to Distribution
Core2	EX9204	0/2

FPC 1 FPC 2

1

SFP+	1	3	5	7	1	3	5	7	1	3	5	7	1	3	5	7
	0	2	4	6	0	2	4	6	0	2	4	6	0	2	4	6

EX9200-32XS

Core1	EX9204	0/2
-------	--------	-----

Distribution Switches

QFX5120-48Y Edit Ports for all QFX5120-48Y

Switch	Model	Link to Core	Link to Access
Dist2	QFX5120-48Y	0/2	0/2
Dist1	QFX5120-48Y	0/2	0/2

Access Switches

EX4400-48P Edit Ports for all EX4400-48P

Switch	Link to Distribution
Access2	0/2
Access1	0/2

14. Core Switches

Core1:

Starting with Core1, select xe-1/0/5 and xe-1/0/6 terminating on distribution switches 1 and 2, respectively.

Figure 36: First Port Core1

xe-1/0/5

Port Type
 ge mge xe et

Distribution Switches
 Q

Dist2
Dist1

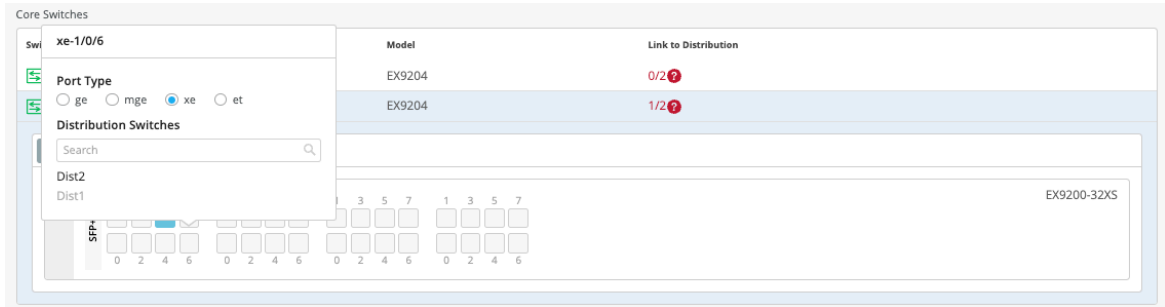
Model	Link to Distribution
EX9204	0/2
EX9204	0/2

1

SFP+	1	3	5	7	1	3	5	7	1	3	5	7	1	3	5	7
	0	2	4	6	0	2	4	6	0	2	4	6	0	2	4	6

EX9200-32XS

Figure 37: Figure :Second port core1



15. Core2:

On Core2, select xe-1/0/4 and xe-1/0/5 terminating on distribution switches 1 and 2, respectively.

Figure 38: Figure :First Port Core2

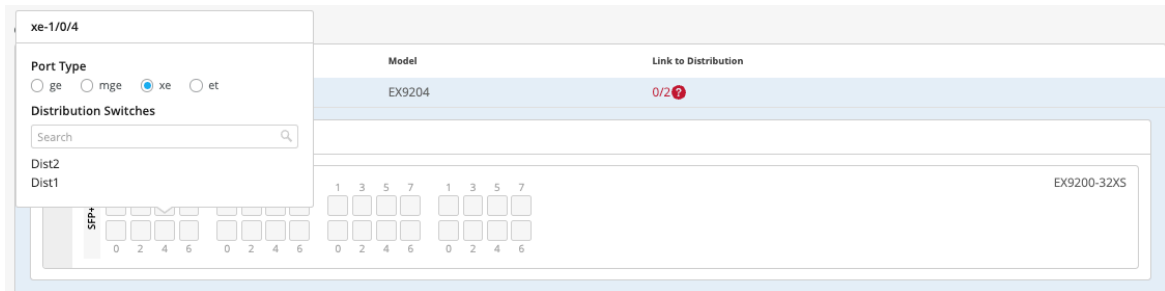
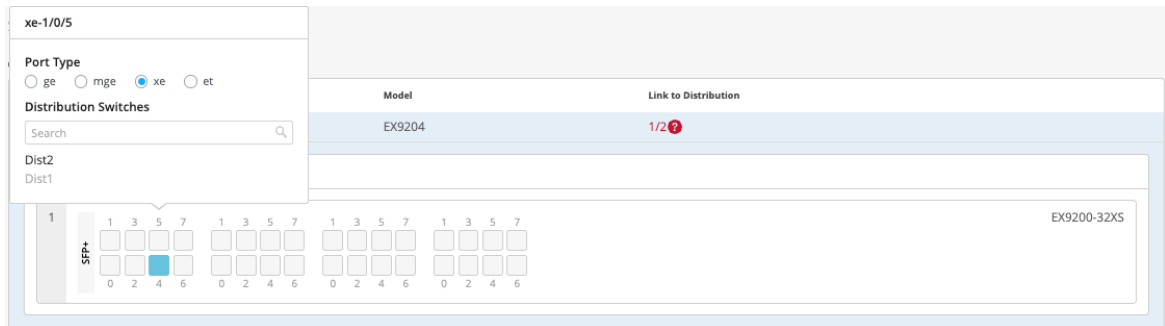


Figure 39: Second Port Core2



16. Distribution Switches

Now moving on to the distribution switches, you notice two interconnect options exist:

- Link to Core

- Link to Access

Dist1:

Select **Link to Core** and choose xe-0/0/5 and xe-0/0/4 terminating on Core Switches 1 and 2, respectively.

Figure 40: First Uplink Port Dist1

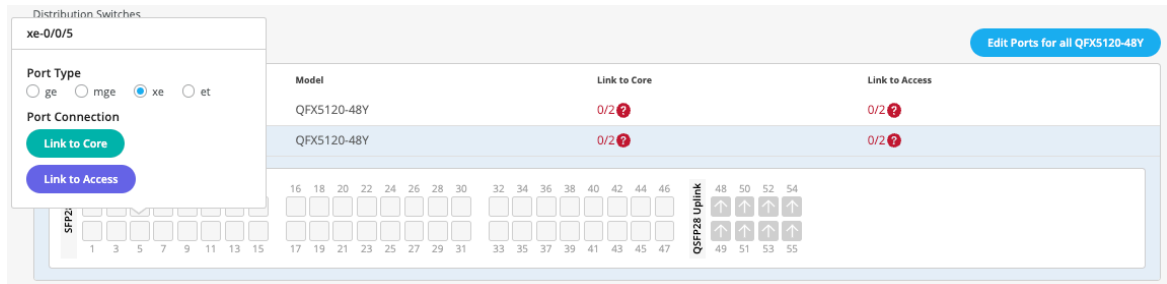
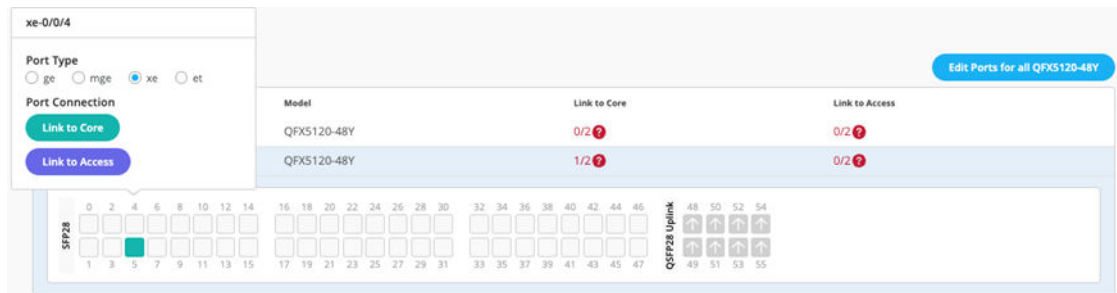


Figure 41: Second Uplink Port Dist1



17. Select **Link to Access** and choose ge-0/0/36 and ge-0/0/37 terminating on access switches 1 and 2, respectively.

Figure 42: First Downlink Port Dist1

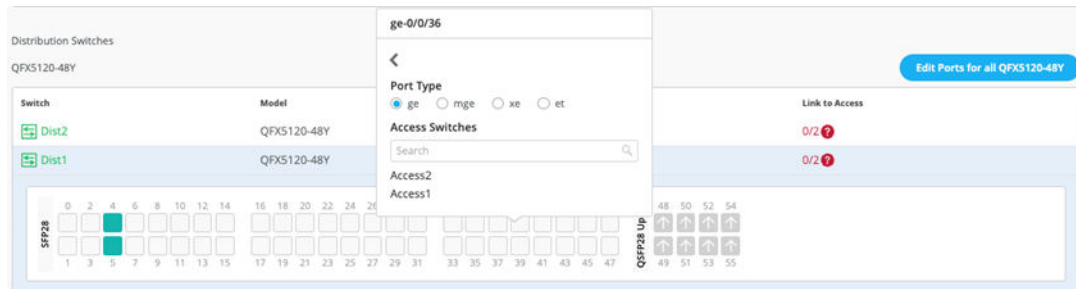


Figure 43: Second Downlink Port Dist1



18. Next, select the following interconnects off **Dist2**:

a. Link to Core

- i. xe-0/0/6 – Core1
- ii. xe-0/0/5 – Core2

b. Link to Access

- i. ge-0/0/36 – Access2
- ii. ge-0/0/37 – Access1

19. Access Switches

You only need to know which interfaces are used to interconnect with the distribution switch but do not need to know the specific mapping. The system bundles all interfaces into a single Ethernet bundle through the AE index option. This greatly simplifies the physical port build for each access switch.

Access1 and 2:

Select both uplinks and interface speed, while allowing Mist to define each AE index. In this case, uplinks ge-0/0/36+37 are selected as Links to Distribution on both access switches and AE Index 11 and 12 on Access1 and 2, respectively.

Figure 44: Uplink and AE# on Access1

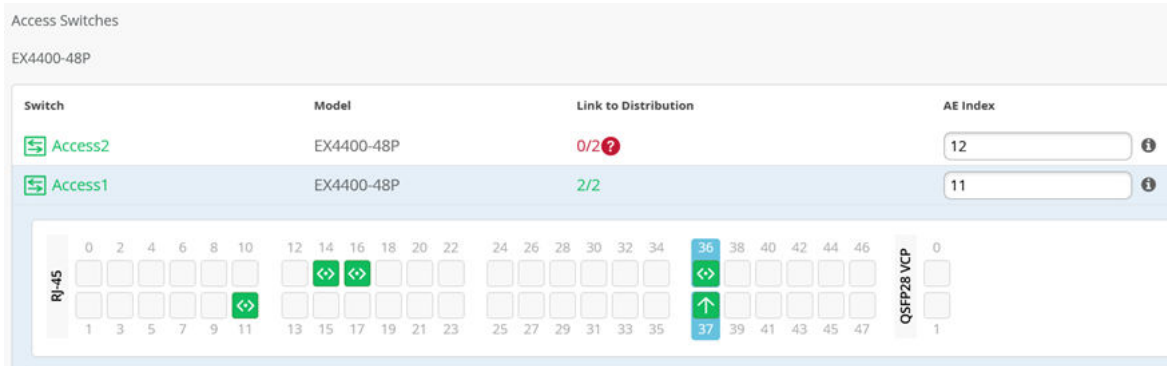
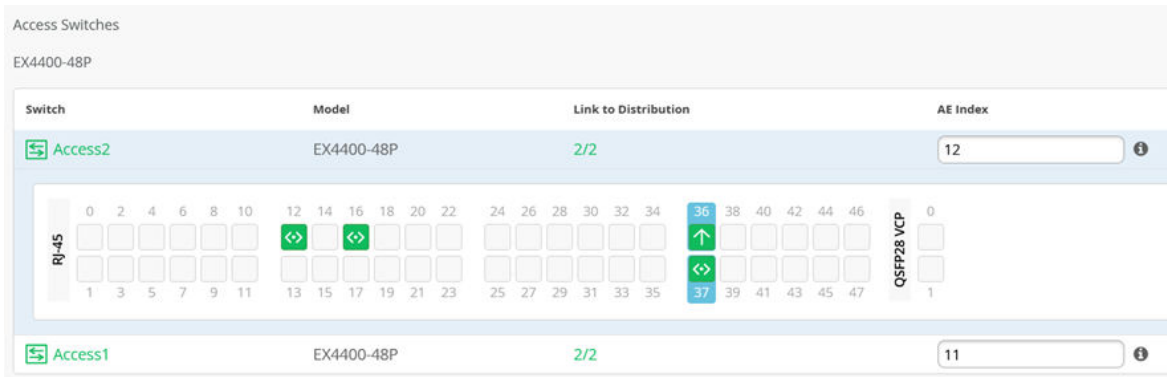


Figure 45: Uplink and AE# on access2



- Once you have completed selecting all requisite port combinations, select the **Continue** button in the upper-right corner of the portal.

Campus Fabric Configuration Confirmation


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

Figure 46: Fabric Confirmation View

Confirm

Review the topology and click "Apply Changes" to save the Fabric configuration to the Mist Cloud

Core



Distribution

Access

Core1

MAC Address f4:b5:2f:f4:04:00

Model EX9204

Status connected

Site Primary Site

Router ID --

VLANs

ID	IP Address	Name
1033	--	vlan1033
1088	--	vlan1088
1099	--	vlan1099

Connections to Distribution

Switch	Port Id
Dist1	xe-1/0/5
Dist2	xe-1/0/6

NOTE: As we have configured the usage of auto router ID subnet, the underlay loopback IP addresses may still not be assigned on this page and warnings may appear like the ones shown above. Please ignore this for now as the assignments happen when you apply the configuration for the first time.

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

Figure 47: Apply Changes to Fabric


X Campus Fabric Configuration 1. Topology 2. Nodes 3. Network Settings 4. Ports 5. Confirm ← Back [Apply Changes](#)

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

Mist displays 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 distribution and 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 on each device with unique BGP autonomous system numbers. The primary goal of the underlay is to leverage ECMP for load balancing traffic on a per packet level for device loopback reachability. The primary goal of the eBGP overlay is support of customer traffic using EVPN-VXLAN.
- IP addressing of each L3 gateway IRB located on Dist1 and Dist2.
- IP addressing of each lo0.0 loopback, which is done automatically in this case.
- 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 MX router that is configured in 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.

Figure 48: Applying Changes



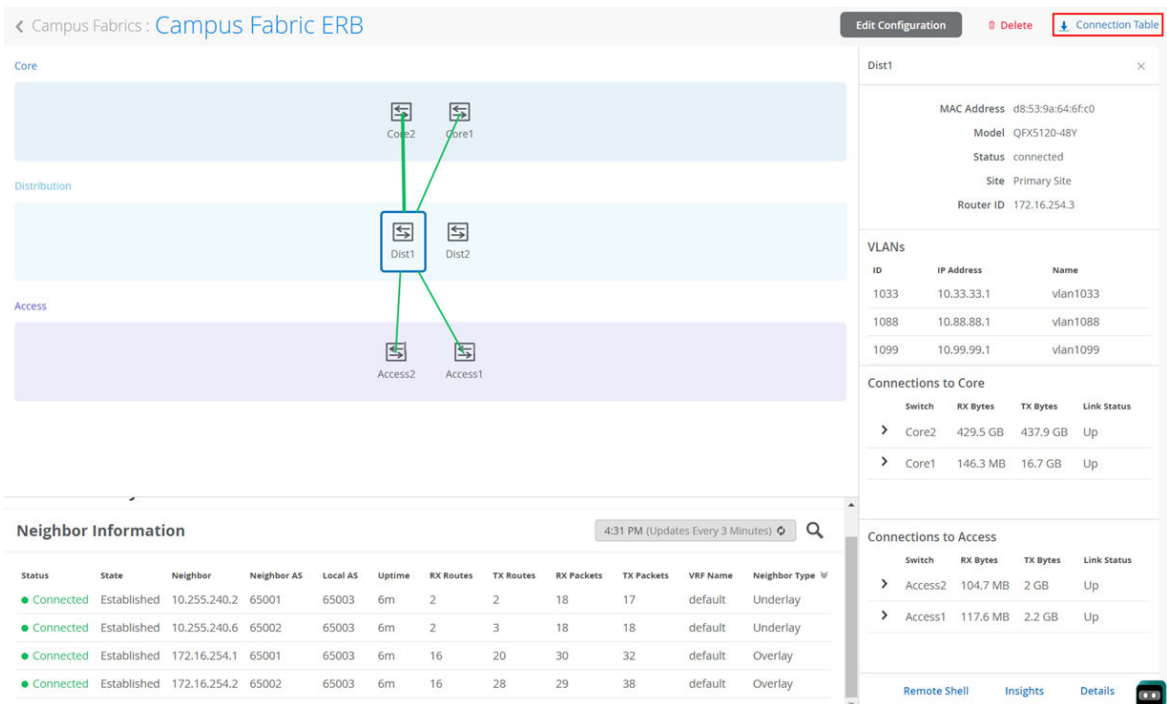
22. Once you click **Close Campus Fabric Configuration**, you can view a summary of the newly created Campus Fabric Core-Distribution ERB.

Figure 49: Created ERB Fabric View



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

Figure 50: Download Connection Table CSV



Connection Table spreadsheet:

Figure 51: Downloaded Connection Table

Role 1	Switch 1	Mac 1	Model 1	Serial 1	Site 1	Port Role	AE 1	Port 1	<-->	Port 2	AE 2	Port Role	Site 2	Serial 2	Model 2	Mac 2	Switch 2	Role 2
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/5	<-->	xe-1/0/5		downlink	Primary Site	JN122EFFFFRFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/6	<-->	xe-1/0/6		downlink	Primary Site	JN122EFFFFRFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag		12 ge-0/0/36	<-->			12 esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag		11 ge-0/0/37	<-->			11 esi-lag	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	uplink		xe-0/0/4	<-->	xe-1/0/4		downlink	Primary Site	JN122EFFFFRFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	uplink		xe-0/0/5	<-->	xe-1/0/5		downlink	Primary Site	JN122EFFFFRFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag		12 ge-0/0/37	<-->			12 esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag		11 ge-0/0/36	<-->			11 esi-lag	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access

Apply VLANs to Access Ports

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

In this case, Desktop1 and 2 are associated with different ports on each access switch which requires the configuration to be applied to Access1 and 2, respectively. See [Figure 12 on page 24](#).

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

Figure 52: Port Configuration Via Switch Template

The screenshot shows a web interface titled "Select Switches Configuration". On the left, there is a list of switch roles: "core" (Model: EX9204), "distribution" (Model: QFX5120*), "access" (Model: EX4400*), and "default" (all remaining switches). The "access" role is currently selected. On the right, there are tabs for "Info", "Port Config", "CLI Config", "IP Config (OOB)", and "CLI Config", with "Port Config" being the active tab. Below the tabs, a message says "Apply port profiles to port ranges on matching switches". A table shows a port range "ge-0/0/16" associated with a profile named "myaccess". Below this, there is a section for "Unassigned ports" with a "Default" profile and a blue "Add Port Range" button.

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.

Figure 53: Assign Port Profile to a Port

The screenshot shows a configuration window titled "PORT CONFIGURATION". Under the "Port Profile Assignment" section, there is a sub-section "Site, Template, or System Defined" with a star icon. Below this is a header "Edit Port Range" with a checkmark and a close button. The "Port IDs" field contains "ge-0/0/11" and a list of examples "(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)". The "Interface" section has three radio buttons: "L2 interface" (selected), "L3 interface", and "L3 sub-interfaces". The "Configuration Profile" field shows "vlan1099" and a dropdown menu with "vlan1099(1099), access" selected. At the bottom, there is a checkbox for "Enable Dynamic Configuration" which is unchecked.

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

Figure 54: Port Profile Example

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

APPENDIX: ERB Fabric Verification (Optional)

IN THIS SECTION

- [BGP Underlay | 59](#)
- [EVPN VXLAN Verification Between Core and Distribution Switches | 63](#)

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

Verification of the Campus Fabric Core-Distribution ERB deployment. See [Figure 12 on page 24](#). Currently, there are two desktops to validate the fabric. Let's take a quick look to see if Desktop1 can connect internally and externally.

Figure 55: Wired Client Connectivity Issue

```

root@desktop1:~# ifconfig vlan1099
vlan1099: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.99.99.99 netmask 255.255.255.0 broadcast 10.99.99.255
    inet6 fe80::5054:ff:fac6:8a58 prefixlen 64 scopeid 0x20<link>
    ether 52:54:00:c6:8a:58 txqueuelen 1000 (Ethernet)
    RX packets 421822 bytes 434750459 (434.7 MB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 106373 bytes 5238868 (5.2 MB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@desktop1:~# ip r
default via 10.99.99.1 dev vln1099 proto static
10.99.99.0/24 dev vln1099 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 -c2
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=5.91 ms
64 bytes from 10.99.99.1: icmp_seq=2 ttl=64 time=5.69 ms
--- 10.99.99.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1002ms
rtt min/avg/max/mdev = 5.690/5.799/5.909/0.109 ms
root@desktop1:~# ping 10.99.99.254 -c2
PING 10.99.99.254 (10.99.99.254) 56(84) bytes of data:
--- 10.99.99.254 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1018ms

```

Validation steps:

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

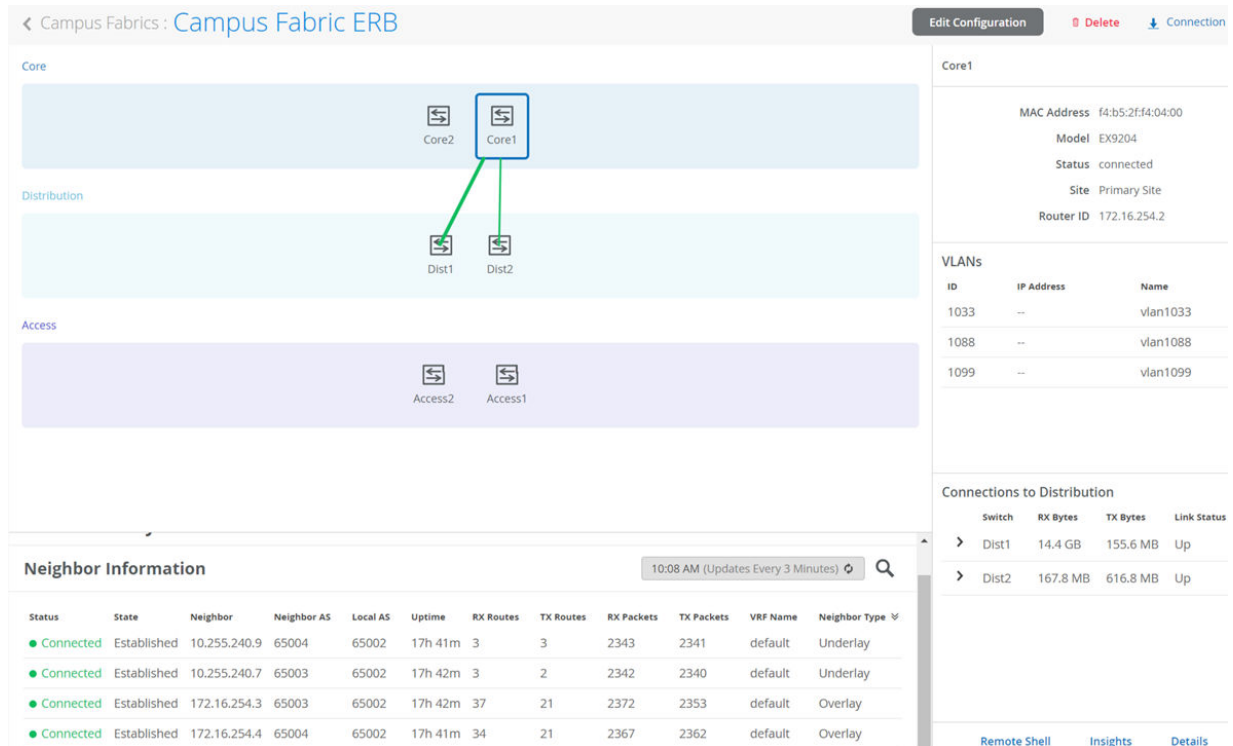
Start by validating the campus fabric in the portal by selecting the **Campus Fabric** option under the **Organization** tab on the left side of the portal.

Site Topologies			
Name	Topology ID	Site	Date Created
DC81-ERB	4904ecde-5266-4a73-99ad-470749e485d6	Primary Site	10:00:33 AM, Dec 12 2022

Remote shell access into each device within the campus fabric is supported here as well as a visual representation of the following capabilities:

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

Figure 56: Fabric Health



BGP Underlay

Purpose

Verifying the state of eBGP between the core and distribution layers 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 to decrease convergence times during failures.
- Loopback reachability to support VXLAN tunnelling.

Without requiring verification at each layer, the focus can be on Dist1 and 2 and their eBGP relationships with Core1 and 2. If both distribution switches have established eBGP peering sessions with both core switches, you can move to the next phase of verification.

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

Switch Type	Switch Name	Auto assigned Loopback IP
Core	Core1	172.16.254.2
Core	Core2	172.16.254.1
Distribution	Dist1	172.16.254.3
Distribution	Dist2	172.16.254.4
Access	Access1	N/A
Access	Access2	N/A

Action

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

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

Verification of BGP Peering

Dist1:

Access the Remote Shell via the lower-right of the campus fabric, from the switch view, or via Secure Shell (SSH).

Figure 57: Show BGP Summary on Dist1

```
mist@Dist1> show bgp summary
-----
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0
bgp.evpn.0
  42         21         0         0         0         0         0
Peer           AS          InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn State|#Active/Received/Accepted/Damped...
10.255.240.2   65001      2359     2358      0        0      17:50:05 Establ
  inet.0: 2/2/2/0
10.255.240.6   65002      2360     2359      0        0      17:50:05 Establ
  inet.0: 2/2/2/0
172.16.254.1   65001      2375     2382      0        0      17:50:03 Establ
  bgp.evpn.0: 16/21/21/0
  default-switch.evpn.0: 9/13/13/0
  default_evpn__evpn.0: 1/2/2/0
  guest-wifi.evpn.0: 2/2/2/0
  developers.evpn.0: 2/2/2/0
  corp-it.evpn.0: 2/2/2/0
172.16.254.2   65002      2374     2390      0        0      17:50:04 Establ
  bgp.evpn.0: 5/21/21/0
  default-switch.evpn.0: 4/13/13/0
  default_evpn__evpn.0: 1/2/2/0
  guest-wifi.evpn.0: 0/2/2/0
  developers.evpn.0: 0/2/2/0
  corp-it.evpn.0: 0/2/2/0
```

From the BGP summary, we can see that the underlay (10.255.240.X) peer relationships are established, which indicates that the underlay links are attached to the correct devices and the links are up.

It also shows the overlay (172.16.254.x) relationships are established and that it is peering at the correct loopback addresses. This demonstrates underlay loopback reachability.

We can also see routes received; time established are roughly equal which looks good so far.

The campus fabric build illustrates per device real-time BGP peering status shown below from Dist1:

Figure 58: BGP Link Status

Neighbor Information											
10:13 AM (Updates Every 3 Minutes)											🔍
Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
● Connected	Established	10.255.240.2	65001	65003	17h 47m	2	2	2352	2351	default	Underlay
● Connected	Established	10.255.240.6	65002	65003	17h 47m	2	3	2352	2351	default	Underlay
● Connected	Established	172.16.254.1	65001	65003	17h 47m	21	26	2368	2375	default	Overlay
● Connected	Established	172.16.254.2	65002	65003	17h 47m	21	37	2367	2383	default	Overlay

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

Verification of BGP connections can be performed on any of the other switches (not shown).

The primary goal of eBGP in the underlay is to provide loopback reachability between core and distribution devices in the campus fabric. This loopback is used to terminate VXLAN tunnels between devices. The following shows loopback reachability from Dist1 to all devices in the campus fabric:

Figure 59: Testing Underlay Loopback IP Reachability

```
mist@Dist1> ping 172.16.254.1 count 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=3.371 ms

--- 172.16.254.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.371/3.371/3.371/0.000 ms

(master:0)
mist@Dist1> ping 172.16.254.2 count 1
PING 172.16.254.2 (172.16.254.2): 56 data bytes
64 bytes from 172.16.254.2: icmp_seq=0 ttl=64 time=7.198 ms

--- 172.16.254.2 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.198/7.198/7.198/0.000 ms

(master:0)
mist@Dist1> ping 172.16.254.4 count 1
PING 172.16.254.4 (172.16.254.4): 56 data bytes
64 bytes from 172.16.254.4: icmp_seq=0 ttl=63 time=11.148 ms

--- 172.16.254.4 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 11.148/11.148/11.148/0.000 ms
```

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

Let's verify that the routes are established to the core and other devices across multiple paths. For example, Dist1 should leverage both paths through Core1 and 2 to reach Dist2 and vice versa.

Dist1: ECMP Loopback reachability to Dist2 through Core1/2

Figure 60: Loopback Reachability to Dist2

```
mist@Dist1> show route forwarding-table destination 172.16.254.4
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
172.16.254.4/32  user   0      10.255.240.2       ucst  524287  3
                  10.255.240.6       ucst  1827   6 xe-0/0/4.0
                  10.255.240.6       ucst  1828   6 xe-0/0/5.0
```

Dist2: ECMP Loopback reachability with Dist1 through Core1/2

Figure 61: Loopback Reachability to Dist1

```
mist@Dist2> show route forwarding-table destination 172.16.254.3
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
172.16.254.3/32  user  0      10.255.240.4      ucst  1827   6   xe-0/0/5.0
                  10.255.240.8      ucst  1828   6   xe-0/0/6.0
```

This can be repeated for Core1 and 2 to verify ECMP load balancing.

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

Figure 62: BFD Testing

```
mist@Dist1> show bfd session

Address          State   Interface      Detect   Transmit
                  Time    Interval      Multiplier
10.255.240.2     Up      xe-0/0/4.0     1.050   0.350   3
10.255.240.6     Up      xe-0/0/5.0     1.050   0.350   3
172.16.254.1     Up      3.000         1.000   3
172.16.254.2     Up      3.000         1.000   3

4 sessions, 4 clients
Cumulative transmit rate 7.7 pps, cumulative receive rate 7.7 pps
```

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

EVPN VXLAN Verification Between Core and Distribution Switches

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

Verification of the EVPN Database on Both Core Switches

Dist1:

Figure 63: EVPN DB Dist1

```
mist@Dist1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address      Active source      Timestamp          IP address
11033  00:00:5e:e4:31:57  irb.1033        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.33.33.1
11033  5c:5b:35:2e:53:61  00:11:00:00:00:01:00:01:03:0c  Sep 07 14:45:42  10.33.33.1
11033  5c:5b:35:af:29:d5  00:11:00:00:00:01:00:01:03:0b  Sep 07 14:45:38  10.33.33.1
11088  00:00:5e:e4:31:57  irb.1088        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.88.88.1
11088  52:54:00:7b:b4:52  00:11:00:00:00:01:00:01:03:0c  Sep 08 08:26:44  10.88.88.88
11099  00:00:5e:e4:31:57  irb.1099        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.99.99.1
11099  52:54:00:c6:8a:58  00:11:00:00:00:01:00:01:03:0b  Sep 07 15:16:07  10.99.99.99
```

Dist2:

Figure 64: EVPN DB Dist2

```
mist@Dist2> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address      Active source      Timestamp          IP address
11033  00:00:5e:e4:31:57  irb.1033        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.33.33.1
11033  5c:5b:35:2e:53:61  00:11:00:00:00:01:00:01:03:0c  Sep 07 14:45:41  10.33.33.1
11033  5c:5b:35:af:29:d5  00:11:00:00:00:01:00:01:03:0b  Sep 07 14:45:38  10.33.33.1
11088  00:00:5e:e4:31:57  irb.1088        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.88.88.1
11088  52:54:00:7b:b4:52  00:11:00:00:00:01:00:01:03:0c  Sep 07 14:57:25  10.88.88.88
11099  00:00:5e:e4:31:57  irb.1099        00:11:00:00:00:01:00:01:03:0c  Sep 07 14:23:33  10.99.99.1
11099  52:54:00:c6:8a:58  00:11:00:00:00:01:00:01:03:0b  Sep 08 08:32:20  10.99.99.99
```

Both core switches have identical EVPN databases which is expected. Notice the entries for desktop1 (10.99.99.99) and desktop2 (10.88.88.88) present on each core switch. These entries are learned through the campus fabric from the ESI-LAGs off Dist1 and 2.

The 10.99.99.99 IP address is associated with irb.1099 and we see a VNI of 11099. Let's just double-check VLAN to VNI mapping on the distribution and core switches and verify the presence of L3 on the distribution switches.

Distribution:

Figure 65: VLAN Configuration on Distribution

```
root@Dist1> show configuration vlans | display set | display inheritance | match 1099
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 l3-interface irb.1099
set vlans vlan1099 vxlan vni 11099
```

Core:

Figure 66: VLAN Configuration on Core

```
root@Core1> show configuration | match 1099 | display set | display inheritance
root@Core1>
```

We now know that there can be an issue with the config or status of the core switches. The VLAN configuration stanza does not show 1099 which points to a lack of configuration on the core devices. We still have control plane output that displays both desktops' IP and MAC addresses. Let's keep troubleshooting.

Verification of VXLAN Tunneling Between Distribution and Core Switches

Dist1:

Figure 67: VTEP Remote on Dist1

```
mist@Dist1> 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   172.16.254.3  lo0.0  0
RVTEP-IP            L2-RTT      IFL-Idx  Interface  NH-Id  RVTEP-Mode  ELP-IP      Flags
172.16.254.4        default-switch  840      vtep.32769  1829   RNVE
```

Core1:

Figure 68: VTEP Remote on Core1

```
mist@Core1> show ethernet-switching vxlan-tunnel-end-point remote summary
mist@Core1> █
```

NOTE: The distribution switches still do not display Core1 and 2 as the remote tunnel destinations. The reason is that we have not yet configured the uplinks to the WAN router at the top of the fabric. Hence, no core switch knows about vlan1099, nor does it have any information about it.

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

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

- Different VLAN/VNI attached to the same VRF on different access switches.

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

APPENDIX: WAN Router Integration into the Fabric

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


- Via a Layer 2 forwarding method:
 - The fabric uplinks are configured as ESI-LAGs and contain one or more tagged VLANs (one for each VRF) to communicate with the WAN router.
 - It is also necessary that you configure the IP address of the WAN router interface manually as the next-hop IP address for default-forwarding on each fabric VRF as already shown above.
 - The WAN router itself needs to understand standard IEEE 802.3ad LAG with active LACP.
 - If you have more than one WAN router attached for redundancy, it is advised to provide failover mechanisms between them for the interface IP addresses towards the fabric. VRRP is recommended.
 - Routes between fabric and WAN router are only statically configured.
- Via a Layer 3 forwarding method:
 - The fabric uplinks are configured as Layer 3 peer-to-peer IP links.
 - Per fabric VRF, a peer-to-peer link needs to be established with the WAN router.
 - Usually, there are multiple peer-to-peer links on a single physical uplink. Those are further segmented via tagged VLANs to provide isolation on the uplinks.
 - There is no need to manually configure next-hops for each VRF inside the fabric as it is assumed that the propagation of the default gateways will be obtained from the WAN router through a routing protocol.
 - Between the fabric and the WAN router, a routing protocol must be established to exchange routes.
 - The campus fabric supports exterior BGP and OSPF as routing protocols towards the WAN router.

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

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

Remember that you chose to deploy the border gateway capability on the Juniper Networks® EX9204 Switches during the Campus Fabric Core-Distribution deployment, represented below:

Figure 69: WAN-Router Integration Via ESI-LAG



Mist enables the EX9204 to translate between VXLAN traffic within the campus fabric and standard Ethernet switching for external connectivity—in this case an MX router. Let's verify the ESI status on the core switches.

Figure 70: LAG with LACP not up

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

We must configure the ESI-LAG since Mist does not configure this automatically. Add a port profile on the core switch interfaces facing the WAN router.

The following represents an existing port profile applied to each MX router facing EX9204 switch ports (on Core2, the switch port is xe-1/0/1).

Figure 71: Port Config with ESI-LAG

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

Figure 72: Check ESI-LAG Configuration and EVPN DB

```
mist@Core1> show configuration interfaces ae0 | display set | display inheritance
set interfaces ae0 mtu 9014
set interfaces ae0 esi 00:11:00:00:01:00:01:02:00
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 ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vLAN members all

mist@Core1> show evpn database
Instance: evpn vs
VLAN DomainId MAC address Active source Timestamp IP address
11033 5c:5b:35:2e:53:61 00:11:00:00:01:00:01:03:0c Sep 08 08:49:27
11033 5c:5b:35:af:29:d5 00:11:00:00:01:00:01:03:0b Sep 08 08:49:27
11033 f0:1e:22:c8:e8:ff 00:11:00:00:01:00:01:02:00 Sep 08 08:49:45 10.33.33.254
11088 52:54:00:7b:b4:52 00:11:00:00:01:00:01:03:0c Sep 08 08:49:27 10.88.88.88
11088 f0:1c:2d:c8:e8:ff 00:11:00:00:01:00:01:02:00 Sep 08 08:49:46 10.88.88.254
11099 52:54:00:c6:8a:58 00:11:00:00:01:00:01:03:0b Sep 08 08:49:27 10.99.99.99
11099 f0:1c:2d:c8:e8:ff 00:11:00:00:01:00:01:02:00 Sep 08 08:49:39 10.99.99.254
```

Note that LACP is up, and this infers that there is an existing configuration on the MX router.

Figure 73: LACP Link to WAN-Router Up

```
root@Core1> show lACP statistics interfaces
Aggregated interface: ae0
LACP Statistics:
xe-1/0/0      LACP Rx  LACP Tx  Unknown Rx  Illegal Rx
                2165      2166      0            0

root@Core1> show lACP interfaces
Aggregated interface: ae0
LACP states:
xe-1/0/0      Actor  No  No  Yes  Yes  Yes  Yes  Fast  Active
xe-1/0/0      Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
LACP protocol:
xe-1/0/0      Receive State Transmit State Mux State
                Current  Fast periodic Collecting distributing

root@Core1>
```

Finally, let us validate that Core1 is receiving Desktop 1's MAC address through MP-BGP via Type 2 EVPN routes.

Figure 74: Check MAC-Address of Client Received Now

```
mist@Core1> show route advertising-protocol bgp 172.16.254.3 evpn-mac-address 52:54:00:c6:8a:58 table bgp.evpn.0
bgp.evpn.0: 57 destinations, 89 routes (57 active, 0 holddown, 0 hidden)
Prefix          Nexthop      MED      Lclpref   AS path
* 2:172.16.254.4:1::11099:52:54:00:c6:8a:58/304 MAC/IP
+                172.16.254.4                65004 I
* 2:172.16.254.4:1::11099:52:54:00:c6:8a:58::10.99.99.99/304 MAC/IP
+                172.16.254.4                65004 I
```

NOTE: The EVPN database is confirmed on both Core1 and 2 and VXLAN tunnels are established between distribution and core switches. We have also verified Desktop1 and 2 are present in Core1 and 2's EVPN databases.

Let's check to see if the core has the Desktop1 MAC address.

Figure 75: Check MAC-Address

```
mist@Core1> show evpn database | match 52:54:00:c6:8a:58
11099      52:54:00:c6:8a:58  00:11:00:00:00:01:00:01:03:0b  Sep 08 08:49:27  10.99.99.99
```

We next verify if Desktop1's MAC address is learned via VTEP interfaces on Core1:

Figure 76: Remote MAC-Table on Core1

```
mist@Core1> show ethernet-switching vxlan-tunnel-end-point remote mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC, P -Pinned MAC)

Logical system : <default>
Routing instance : evpn_vs
Bridging domain : vlan1033+1033, VLAN : 1033, VNID : 11033
MAC      MAC      Logical      Remote VTEP
address  flags  interface    IP address
5c:5b:35:2e:53:61  DR     esi.636      172.16.254.4
5c:5b:35:af:29:d5  DR     esi.637      172.16.254.3

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC, P -Pinned MAC)

Bridging domain : vlan1088+1088, VLAN : 1088, VNID : 11088
MAC      MAC      Logical      Remote VTEP
address  flags  interface    IP address
52:54:00:7b:b4:52  DR     esi.636      172.16.254.4 172.16.254.3

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC, P -Pinned MAC)

Bridging domain : vlan1099+1099, VLAN : 1099, VNID : 11099
MAC      MAC      Logical      Remote VTEP
address  flags  interface    IP address
52:54:00:c6:8a:58  DR     esi.637      172.16.254.4 172.16.254.3
```

Dist1 does show both desktop MAC and ARP entries which once again validates control plane operational status within the campus fabric. Additionally, we have also learned the MAC addresses that are located on the WAN router indicated below in yellow.

Figure 77: Check Other WAN-Router IP's Seen

```

mist@Dist1> 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 : 10 entries, 10 learned
Routing instance : default-switch

```

Vlan name	MAC address	MAC flags	GBP tag	Logical interface	SVLBNH/ VENN Index	Active source
default	5c15b:35:af:29:d5	D		et-0/0/47.0		
default	d8153:9a:64:6f:c0	D		et-0/0/47.0		
default	d8153:9a:64:b5:c0	D		et-0/0/47.0		
vlan1033	5c15b:35:2e:53:61	DR		ae12.0		
vlan1033	5c15b:35:af:29:d5	DL		ae11.0		
vlan1033	f01c:2d:c8:e8:f0	DR		esi.1708	1707	00:11:00:00:01:00:01:02:00
vlan1088	52154:00:7b:b4:52	DLR		ae12.0		
vlan1088	f01c:2d:c8:e8:f0	DR		esi.1708	1707	00:11:00:00:01:00:01:02:00
vlan1099	52154:00:c6:8a:58	DLR		ae11.0		
vlan1099	f01c:2d:c8:e8:f0	DR		esi.1708	1707	00:11:00:00:01:00:01:02:00

```

(master:0)
mist@Dist1> show arp

```

MAC Address	Address	Name	Interface	Flags
f01c:2d:c8:e8:f0	10.33.33.254	10.33.33.254	irb.1033 [.local..8]	permanent remote
52154:00:7b:b4:52	10.88.88.88	10.88.88.88	irb.1088 [ae12.0]	permanent remote
f01c:2d:c8:e8:f0	10.88.88.254	10.88.88.254	irb.1088 [.local..8]	permanent remote
52154:00:c6:8a:58	10.99.99.99	10.99.99.99	irb.1099 [ae11.0]	permanent remote
f01c:2d:c8:e8:f0	10.99.99.254	10.99.99.254	irb.1099 [.local..8]	permanent remote
f41b5:2f:f4:06:9a	10.255.240.2	10.255.240.2	xe-0/0/4.0	none
f41b5:2f:f4:06:9a	10.255.240.6	10.255.240.6	xe-0/0/5.0	none
fe:00:00:00:00:80	128.0.0.16	fpco	bme0.0	permanent
d8153:9a:64:6f:c3	192.168.1.1	hypervisor	em2.32768	none
ee:66:5c:46:24:47	192.168.1.16		em2.32768	none
f4:a7:39:cb:e3:20	192.168.230.1		em0.0	none

```

Total entries: 11

```

Then, confirm the EVPN database now has the ESI entry. The MX router IP address for each VLAN ending in .254 is also present in the EVPN database. Back to Desktop1 to see if it can cross the fabric.

Figure 78: Ping Internet from Client1

```

root@desktop1:~#
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=2.11 ms
64 bytes from 1.0.0.1: icmp_seq=2 ttl=52 time=3.00 ms

--- 1.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 2.110/2.553/2.997/0.443 ms

```

Last step is to verify Desktop1 can ping Desktop2:

Figure 79: Verify VRF to VRF Traffic

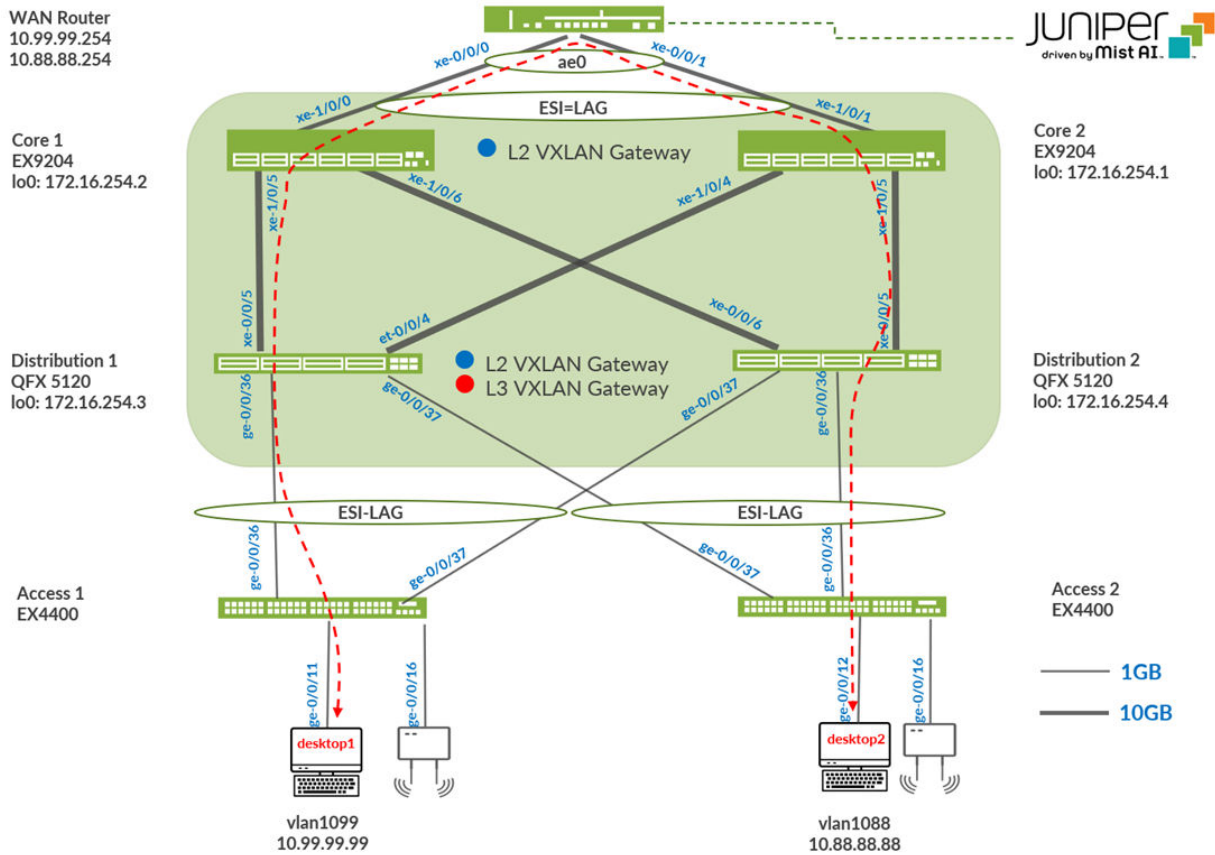
```

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=4.68 ms
64 bytes from 10.88.88.88: icmp_seq=2 ttl=62 time=0.590 ms

--- 10.88.88.88 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.590/2.635/4.681/2.045 ms
root@desktop1:~# █

```


Figure 80: Topology Repeat



Conclusion: The connectivity within the campus fabric and externally have been verified. Desktops can communicate with each other through the campus fabric, each in an isolated VRF, then forwarded to the MX router through the dual-homing ESI-LAG on both Core1 and 2 for routing between VRFs or routing instances. Internet connectivity was also verified from each desktop.

APPENDIX: EVPN Insights

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

Figure 81: Core1 Insights

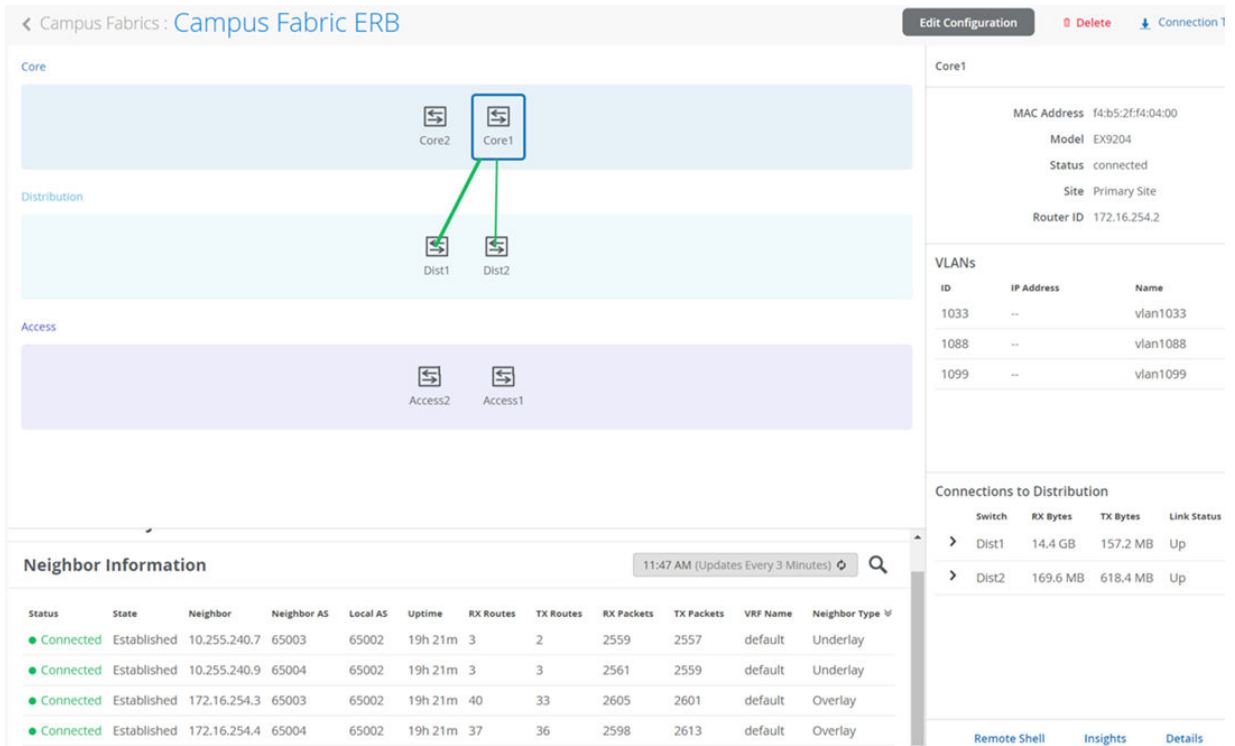


Figure 82: Dist1 Insights

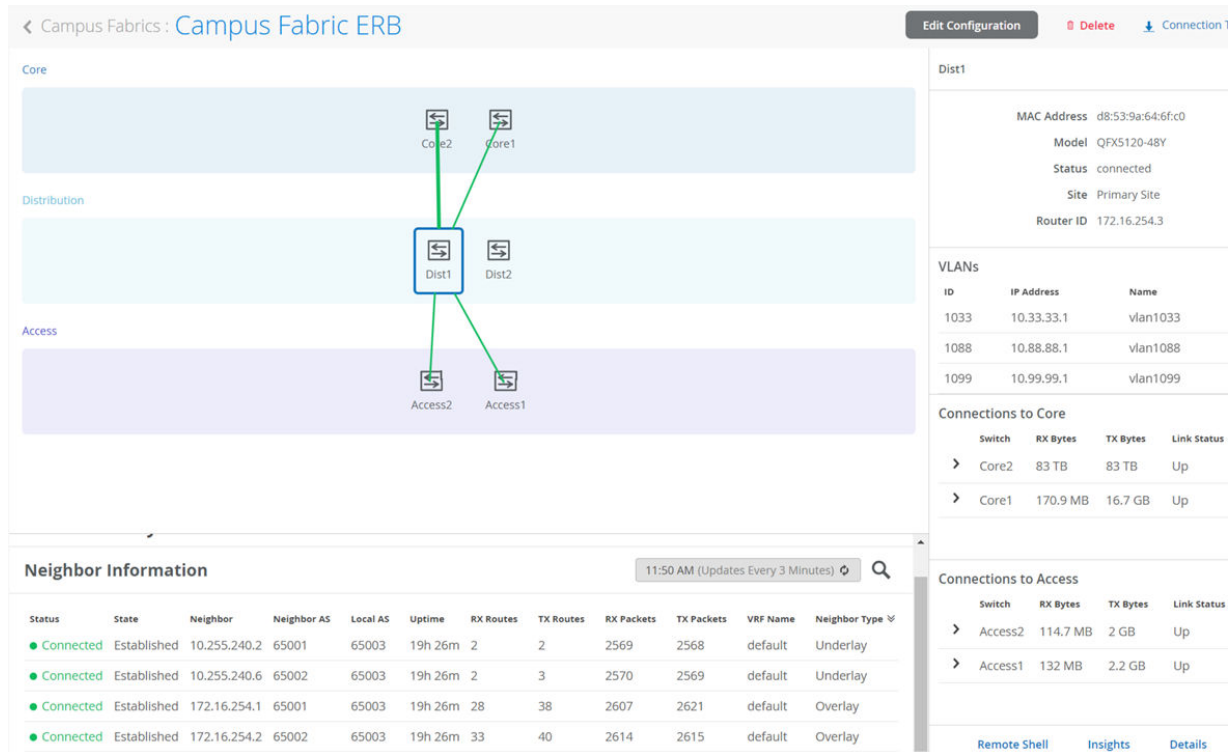
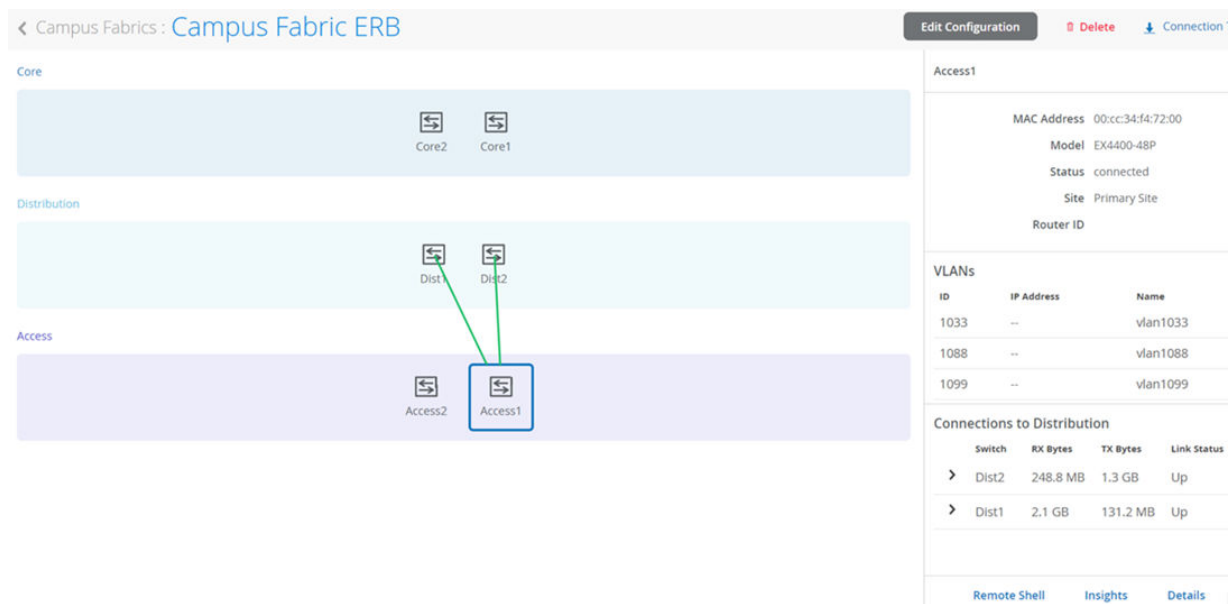


Figure 83: Access1 Insights

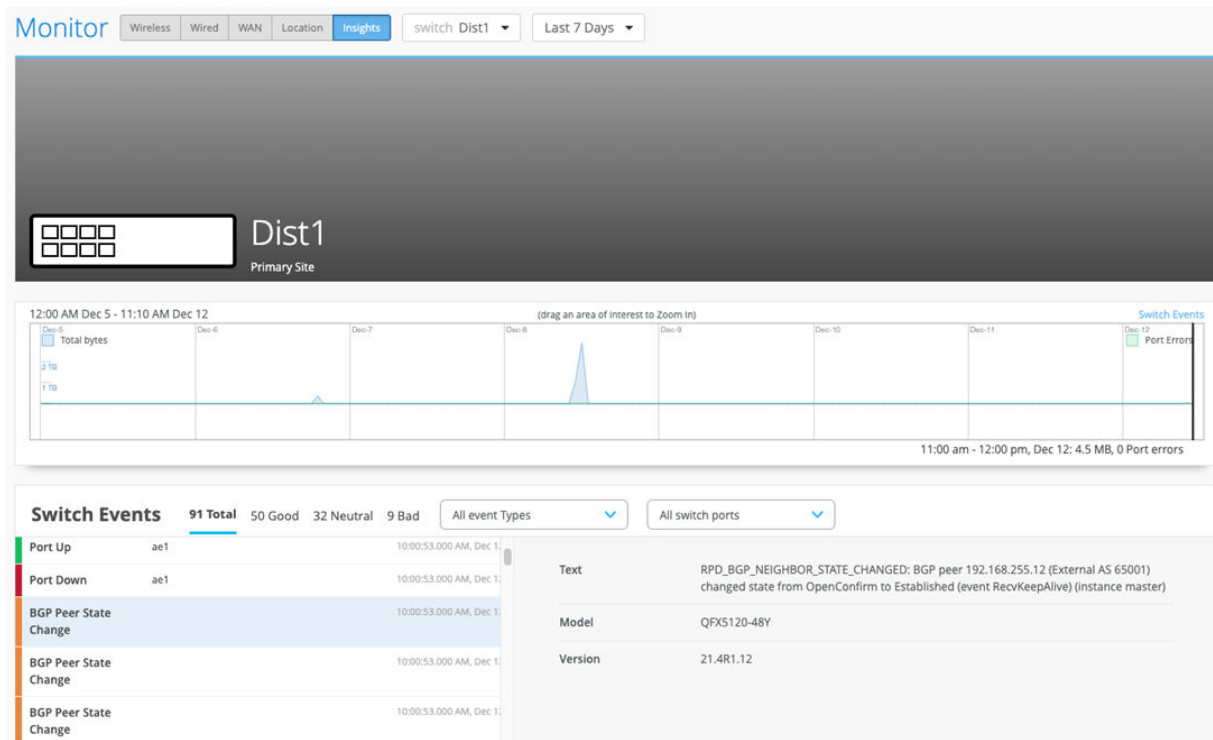


From this view, Mist also provides remote accessibility into each device's console through the Remote Shell option as well as rich telemetry through the **Switch Insights** option. Remote Shell has been

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

Switch Insights of Dist1 displays historical telemetry including BGP peering status critical to the health of the campus fabric:

Figure 84: Single Switch Insights



APPENDIX: Junos Configuration from This Fabric

IN THIS SECTION

- [Campus Fabric Core-Distribution ERB Configurations | 75](#)
- [Configuration of the EVPN VXLAN Overlay and Virtual Networks | 79](#)
- [Configuration of the Layer 2 ESI-LAG Between the Distribution Switches and the Access Switches | 87](#)

- Configuration of the Layer 2 ESI-LAG Between the Core Switches and the MX Router | 89

Campus Fabric Core-Distribution ERB 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 the following options (default in parenthesis):

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

Throughout the Campus Fabric between core and distribution layers, 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 between the two distribution switches:

```
set interfaces xe-1/0/5 description evpn_downlink-to-d8539a646fc0
set interfaces xe-1/0/5 unit 0 family inet address 10.255.240.6/31
set interfaces xe-1/0/6 description evpn_downlink-to-d8539a64b5c0
set interfaces xe-1/0/6 unit 0 family inet address 10.255.240.8/31
```

2. Loopback interface and router ID and AS:

```
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
set groups top routing-options autonomous-system
65002
```

3. Per-packet load-balancing:

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

4. BGP underlay network between the two distribution switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 65002
set protocols bgp group evpn_underlay multipath multiple-as
set protocols bgp group evpn_underlay bfd-liveness-detection minimum-interval 350
set protocols bgp group evpn_underlay bfd-liveness-detection multiplier 3
set protocols bgp group evpn_underlay neighbor 10.255.240.7 peer-as 65003
set protocols bgp group evpn_underlay neighbor 10.255.240.9 peer-as 65004
set protocols bgp graceful-restart
```

Core2 Configuration:

1. Interconnects between the two distribution switches:

```
set interfaces xe-1/0/4 description evpn_downlink-to-d8539a646fc0
set interfaces xe-1/0/4 unit 0 family inet address 10.255.240.2/31
set interfaces xe-1/0/5 description evpn_downlink-to-d8539a64b5c0
set interfaces xe-1/0/5 unit 0 family inet address 10.255.240.4/31
```

2. Loopback interface and router ID and AS:

```
set groups top interfaces lo0 unit 0 family inet address 172.16.254.1/32
set groups top routing-options router-id 172.16.254.1
set groups top routing-options autonomous-system 65001
```

3. Per-packet load-balancing:

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

4. BGP underlay network between the two distribution switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 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 10.255.240.3 peer-as 65003
set protocols bgp group evpn_underlay neighbor 10.255.240.5 peer-as 65004
set protocols bgp graceful-restart
```

Dist1 Configuration:

1. Interconnects between the two core switches:

```
# Core Interfaces:
set interfaces xe-0/0/4 description evpn_uplink-to-f4b52ff3f400
set interfaces xe-0/0/4 unit 0 family inet address 10.255.240.3/31
set interfaces xe-0/0/5 description evpn_uplink-to-f4b52ff40400
set interfaces xe-0/0/5 unit 0 family inet address 10.255.240.7/31
```

2. Loopback interface and router ID:

```
set groups top interfaces lo0 unit 0 family inet address 172.16.254.3/32
set groups top routing-options router-id 172.16.254.3
set groups top routing-options autonomous-system 65003
```

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 to the two core switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 65003
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 10.255.240.6 peer-as 65002
set protocols bgp graceful-restart
```

Dist2 Configuration:

1. Interconnects between the two core switches:

```
# Core Interfaces:
set interfaces xe-0/0/5 description evpn_uplink-to-f4b52ff3f400
set interfaces xe-0/0/5 unit 0 family inet address 10.255.240.5/31
set interfaces xe-0/0/6 description evpn_uplink-to-f4b52ff40400
set interfaces xe-0/0/6 unit 0 family inet address 10.255.240.9/31
```

2. Loopback interface and router ID:

```
set groups top interfaces lo0 unit 0 family inet address 172.16.254.4/32
set groups top routing-options router-id 172.16.254.4
set groups top routing-options autonomous-system 65004
```


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 to the two core switches:

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "xyz"
set protocols bgp group evpn_underlay export evpn_underlay_export
set protocols bgp group evpn_underlay local-as 65004
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 65001
set protocols bgp group evpn_underlay neighbor 10.255.240.8 peer-as 65002
set protocols bgp graceful-restart
```

Configuration of the EVPN VXLAN Overlay and Virtual Networks

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

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 the distribution and core switches.

VRFs for traffic isolation are provisioned on the distribution switches.

Core1 Configuration:

1. BGP overlay peering between the two distribution switches:

```

set protocols bgp group evpn_overlay type external
set protocols bgp group evpn_overlay multihop ttl 1
set protocols bgp group evpn_overlay multihop no-nexthop-change
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 loops 2
set protocols bgp group evpn_overlay authentication-key "xyz"
set protocols bgp group evpn_overlay local-as 65002
set protocols bgp group evpn_overlay multipath multiple-as
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 172.16.254.3 peer-as 65003
set protocols bgp group evpn_overlay neighbor 172.16.254.4 peer-as 65004

```

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

```

set groups top routing-instances evpn_vs vtep-source-interface lo0.0
set groups top routing-instances evpn_vs route-distinguisher 172.16.254.2:1
set groups top routing-instances evpn_vs vrf-target target:65000:1

```

3. VXLAN encapsulation:

```

set groups top routing-instances evpn_vs protocols evpn encapsulation vxlan
set groups top routing-instances evpn_vs protocols evpn default-gateway no-gateway-community
set groups top routing-instances evpn_vs protocols evpn extended-vni-list all

```

4. VLAN to VXLAN mapping:

```

set groups top routing-instances evpn_vs vlans vlan1033 vlan-id 1033
set groups top routing-instances evpn_vs vlans vlan1033 vxlan vni 11033
set groups top routing-instances evpn_vs vlans vlan1088 vlan-id 1088
set groups top routing-instances evpn_vs vlans vlan1088 vxlan vni 11088
set groups top routing-instances evpn_vs vlans vlan1099 vlan-id 1099
set groups top routing-instances evpn_vs vlans vlan1099 vxlan vni 11099

```

Core2 Configuration:

1. BGP overlay peering between the two distribution switches:

```

set protocols bgp group evpn_overlay type external
set protocols bgp group evpn_overlay multihop ttl 1
set protocols bgp group evpn_overlay multihop no-nexthop-change
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 loops 2
set protocols bgp group evpn_overlay authentication-key "xyz"
set protocols bgp group evpn_overlay local-as 65001
set protocols bgp group evpn_overlay multipath multiple-as
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 172.16.254.3 peer-as 65003
set protocols bgp group evpn_overlay neighbor 172.16.254.4 peer-as 65004

```

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

```

set groups top routing-instances evpn_vs vtep-source-interface lo0.0
set groups top routing-instances evpn_vs route-distinguisher 172.16.254.1:1
set groups top routing-instances evpn_vs vrf-target target:65000:1

```

3. VXLAN encapsulation:

```

set groups top routing-instances evpn_vs protocols evpn encapsulation vxlan
set groups top routing-instances evpn_vs protocols evpn default-gateway no-gateway-community
set groups top routing-instances evpn_vs protocols evpn extended-vni-list all

```

4. VLAN to VXLAN mapping:

```

set groups top routing-instances evpn_vs vlans vlan1033 vlan-id 1033
set groups top routing-instances evpn_vs vlans vlan1033 vxlan vni 11033
set groups top routing-instances evpn_vs vlans vlan1088 vlan-id 1088
set groups top routing-instances evpn_vs vlans vlan1088 vxlan vni 11088
set groups top routing-instances evpn_vs vlans vlan1099 vlan-id 1099
set groups top routing-instances evpn_vs vlans vlan1099 vxlan vni 11099

```

Dist1 Configuration:

1. BGP overlay peering between the two core switches:

```

set protocols bgp group evpn_overlay type external
set protocols bgp group evpn_overlay multihop ttl 1
set protocols bgp group evpn_overlay multihop no-nexthop-change
set protocols bgp group evpn_overlay local-address 172.16.254.3
set protocols bgp group evpn_overlay log-updown
set protocols bgp group evpn_overlay family evpn signaling loops 2
set protocols bgp group evpn_overlay authentication-key "xyz"
set protocols bgp group evpn_overlay local-as 65003
set protocols bgp group evpn_overlay multipath multiple-as
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 peer-as 65001
set protocols bgp group evpn_overlay neighbor 172.16.254.2 peer-as 65002

```

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 172.16.254.3:1
set groups top switch-options vrf-target target:65000:1

```

3. VXLAN encapsulation:

```

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 that are used for traffic isolation:

```

set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi interface irb.1033
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 routing-options multipath
set groups top routing-instances guest-wifi route-distinguisher 172.16.254.3:103
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 guest-wifi interface lo0.3
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes advertise direct-
nexthop
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes encapsulation
vxlan
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes vni 15560868
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers interface irb.1088
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 routing-options multipath
set groups top routing-instances developers route-distinguisher 172.16.254.3:102
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 developers interface lo0.2
set groups top routing-instances developers protocols evpn ip-prefix-routes advertise direct-
nexthop
set groups top routing-instances developers protocols evpn ip-prefix-routes encapsulation
vxlan
set groups top routing-instances developers protocols evpn ip-prefix-routes vni 15600414
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it interface irb.1099
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 routing-options multipath
set groups top routing-instances corp-it route-distinguisher 172.16.254.3:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
set groups top routing-instances corp-it interface lo0.1
set groups top routing-instances corp-it protocols evpn ip-prefix-routes advertise direct-
nexthop
set groups top routing-instances corp-it protocols evpn ip-prefix-routes encapsulation vxlan
set groups top routing-instances corp-it protocols evpn ip-prefix-routes vni 11284517

```

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 anycast addressing:

```

set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet mtu 9000
set interfaces irb unit 1033 family inet address 10.33.33.1/24
set interfaces irb unit 1033 mac 00:00:5e:e4:31:57
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet mtu 9000
set interfaces irb unit 1088 family inet address 10.88.88.1/24
set interfaces irb unit 1088 mac 00:00:5e:e4:31:57
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.1/24
set interfaces irb unit 1099 mac 00:00:5e:e4:31:57

```

Dist2 Configuration:

1. BGP overlay peering between the two core switches:

```

set protocols bgp group evpn_overlay type external
set protocols bgp group evpn_overlay local-address 172.16.254.4
set protocols bgp group evpn_overlay local-as 65004
set protocols bgp group evpn_overlay multipath multiple-as
set protocols bgp group evpn_overlay authentication-key "Xyz"
set protocols bgp group evpn_overlay family evpn signaling loops 2
set protocols bgp group evpn_overlay log-updown
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 multihop ttl 1
set protocols bgp group evpn_overlay multihop no-next-hop-change
set protocols bgp group evpn_overlay neighbor 172.16.254.1 peer-as 65001
set protocols bgp group evpn_overlay neighbor 172.16.254.2 peer-as 65002

```

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 172.16.254.4:1
set groups top switch-options vrf-target target:65000:1
```

3. VXLAN encapsulation:

```
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 that are used for traffic isolation:

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi interface irb.1033
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 routing-options multipath
set groups top routing-instances guest-wifi route-distinguisher 172.16.254.4:103
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 guest-wifi interface lo0.3
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes advertise direct-
next-hop
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes encapsulation
vxlan
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes vni 15560868
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers interface irb.1088
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 routing-options multipath
set groups top routing-instances developers route-distinguisher 172.16.254.4:102
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 developers interface lo0.2
set groups top routing-instances developers protocols evpn ip-prefix-routes advertise direct-
```

```

nexthop
set groups top routing-instances developers protocols evpn ip-prefix-routes encapsulation
vxlan
set groups top routing-instances developers protocols evpn ip-prefix-routes vni 15600414
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it interface irb.1099
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 routing-options multipath
set groups top routing-instances corp-it route-distinguisher 172.16.254.4:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
set groups top routing-instances corp-it interface lo0.1
set groups top routing-instances corp-it protocols evpn ip-prefix-routes advertise direct-
nexthop
set groups top routing-instances corp-it protocols evpn ip-prefix-routes encapsulation vxlan
set groups top routing-instances corp-it protocols evpn ip-prefix-routes vni 11284517

```

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 anycast addressing:

```

set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet mtu 9000
set interfaces irb unit 1033 family inet address 10.33.33.1/24
set interfaces irb unit 1033 mac 00:00:5e:e4:31:57
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet mtu 9000
set interfaces irb unit 1088 family inet address 10.88.88.1/24
set interfaces irb unit 1088 mac 00:00:5e:e4:31:57

```



```


set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.1/24
set interfaces irb unit 1099 mac 00:00:5e:e4:31:57

```

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

This section displays the configuration output from the Juniper Mist cloud for the enablement of the Layer 2 ESI-LAGs between the distribution switches and access switches. This 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 RSTP.

Figure 85: Access Switch Attach to Distribution Switches



Dist1 Configuration:

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

```

set interfaces ae11 apply-groups erb-lag
set interfaces ae11 esi all-active
set interfaces ae11 esi 00:11:00:00:00:01:00:01:03:0b
set interfaces ae11 aggregated-ether-options lACP active

```

```

set interfaces ae11 aggregated-ether-options lACP periodic fast
set interfaces ae11 aggregated-ether-options lACP system-id 00:00:00:31:57:0b
set interfaces ae11 aggregated-ether-options lACP admin-key 11
set interfaces ae12 apply-groups erb-lag
set interfaces ae12 esi all-active
set interfaces ae12 esi 00:11:00:00:00:01:00:01:03:0c
set interfaces ae12 aggregated-ether-options lACP active
set interfaces ae12 aggregated-ether-options lACP periodic fast
set interfaces ae12 aggregated-ether-options lACP system-id 00:00:00:31:57:0c
set interfaces ae12 aggregated-ether-options lACP admin-key 12
set groups erb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups erb-lag interfaces <*> unit 0 family ethernet-switching vLAN members [ vLAN1033
vLAN1088 vLAN1099 ]
set groups erb-lag interfaces <*> mtu 9100
set interfaces ge-0/0/36 ether-options 802.3ad ae11
set interfaces ge-0/0/36 description esilag-to-00cc34f3cf00
set interfaces ge-0/0/36 hold-time up 120000 down 1
set interfaces ge-0/0/37 ether-options 802.3ad ae12
set interfaces ge-0/0/37 description esilag-to-00cc34f3cf00
set interfaces ge-0/0/37 hold-time up 120000 down 1

```

Dist2 Configuration:

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

```

set interfaces ae11 apply-groups erb-lag
set interfaces ae11 esi all-active
set interfaces ae11 esi 00:11:00:00:00:01:00:01:03:0b
set interfaces ae11 aggregated-ether-options lACP active
set interfaces ae11 aggregated-ether-options lACP periodic fast
set interfaces ae11 aggregated-ether-options lACP system-id 00:00:00:31:57:0b
set interfaces ae11 aggregated-ether-options lACP admin-key 11
set interfaces ae12 apply-groups erb-lag
set interfaces ae12 esi all-active
set interfaces ae12 esi 00:11:00:00:00:01:00:01:03:0c
set interfaces ae12 aggregated-ether-options lACP active
set interfaces ae12 aggregated-ether-options lACP periodic fast
set interfaces ae12 aggregated-ether-options lACP system-id 00:00:00:31:57:0c
set interfaces ae12 aggregated-ether-options lACP admin-key 12
set groups erb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups erb-lag interfaces <*> unit 0 family ethernet-switching vLAN members [ vLAN1033

```

```

vlan1088 vlan1099 ]
set groups erb-lag interfaces <*> mtu 9100
set interfaces ge-0/0/36 description esilag-to-00cc34f3cf00
set interfaces ge-0/0/36 hold-time up 120000
set interfaces ge-0/0/36 hold-time down 1
set interfaces ge-0/0/36 ether-options 802.3ad ae12
set interfaces ge-0/0/37 description esilag-to-00cc34f3cf00
set interfaces ge-0/0/37 hold-time up 120000
set interfaces ge-0/0/37 hold-time down 1
set interfaces ge-0/0/37 ether-options 802.3ad ae11

```

Access Configuration:

1. VLANs associated with the new LACP Ethernet bundle:

```


set groups erb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups erb-lag interfaces <*> unit 0 family ethernet-switching vlan members [ vlan1033
vlan1088 vlan1099 ]
set groups erb-lag interfaces <*> mtu 9100
set interfaces ae11 apply-groups erb-lag
set interfaces ae11 aggregated-ether-options lacp active
set interfaces ge-0/0/36 ether-options 802.3ad ae11
set interfaces ge-0/0/37 ether-options 802.3ad ae11

```

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

This section displays the configuration output from the Juniper Mist cloud for the enablement of the Layer 2 ESI-LAGs between the core switches and the MX router. This Mist profile enables all VLANs on the Ethernet bundle with requisite ESI and LACP configuration options. From the perspective of the MX router, the Ethernet bundle that is configured on the MX router views the ESI-LAG as a single MAC address with the same LACP system-ID. This enables load hashing between the core and MX router without requiring L2 loop-free detection protocols such as RSTP.

Figure 86: Layer 2 ESI-LAG Supporting Active-Active Load Balancing



Core 1 Configuration:

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

```
set interfaces xe-1/0/0 hold-time up 120000
set interfaces xe-1/0/0 hold-time down 1
set interfaces xe-1/0/0 ether-options 802.3ad ae0
set interfaces xe-1/0/0 unit 0 family ethernet-switching storm-control default
set groups myesilag interfaces <*> mtu 9014
set groups myesilag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups myesilag interfaces <*> unit 0 family ethernet-switching vlan members all
set interfaces ae0 apply-groups myesilag
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
```

Core 2 Configuration:

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

```
set interfaces xe-1/0/1 hold-time up 120000
set interfaces xe-1/0/1 hold-time down 1
set interfaces xe-1/0/1 ether-options 802.3ad ae0
set interfaces xe-1/0/1 unit 0 family ethernet-switching storm-control default
set groups myesilag interfaces <*> mtu 9014
set groups myesilag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups myesilag interfaces <*> unit 0 family ethernet-switching vlan members all
```

```

set interfaces ae0 apply-groups myesilag
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

```

MX Router Configuration:

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

```

set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 mtu 9014
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 1033
set interfaces ae0 unit 0 family bridge vlan-id-list 1088
set interfaces ae0 unit 0 family bridge vlan-id-list 1099
set interfaces irb unit 1033 family inet address 10.33.33.254/24
set interfaces irb unit 1088 family inet address 10.88.88.254/24
set interfaces irb unit 1099 family inet address 10.99.99.254/24
set bridge-domains vlan1033 vlan-id 1033
set bridge-domains vlan1033 routing-interface irb.1033
set bridge-domains vlan1088 vlan-id 1088
set bridge-domains vlan1088 routing-interface irb.1088
set bridge-domains vlan1099 vlan-id 1099
set bridge-domains vlan1099 routing-interface irb.1099

```

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2024 Juniper Networks, Inc. All rights reserved.