



Juniper Validated Design (JVD)

Campus Fabric Core Distribution CRB Using Mist Wired Assurance

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 customers 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 Centrally-Routed Bridging (CRB) architecture to support a campus networking environment using Mist Wired Assurance. The use case shows how you can deploy a single campus fabric that uses EVPN in the control plane, VXLAN tunnels in the overlay network, and BGP in the underlay with Juniper Mist Access Points integration.

Documentation Feedback

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

TABLE OF CONTENTS

About this Document	ii
<i>Overview</i>	<i>ii</i>
<i>Documentation Feedback</i>	<i>ii</i>
Overview	1
Benefits of Campus Fabric Core Distribution	2
Technical Overview	4
<i>Underlay Network</i>	<i>4</i>
<i>Understanding EVPN</i>	<i>5</i>
<i>Overlay Network (Data Plane)</i>	<i>6</i>
<i>Overlay Network (Control Plane)</i>	<i>7</i>
<i>Resiliency and Load Balancing</i>	<i>7</i>
<i>Ethernet Segment Identifier (ESI)</i>	<i>8</i>
<i>Services Block</i>	<i>8</i>
<i>Access Layer</i>	<i>9</i>
<i>Juniper Access Points</i>	<i>10</i>
<i>Supported Platforms for Campus Fabric Core Distribution CRB</i>	<i>10</i>
Campus Fabric Core Distribution CRB Unicast Scale	11
Juniper Mist Wired Assurance	11
Campus Fabric Core Distribution High-level Architecture	12
Campus Fabric Core Distribution CRB Components	13
Juniper Mist Wired Assurance	13
Juniper Mist Wired Assurance Switches	14
<i>Overview</i>	<i>15</i>

<i>Templates</i>	15
<i>Topology</i>	15
Create the Campus Fabric	16
<i>Campus Fabric Org Build</i>	16
<i>Campus Fabric Site Build</i>	17
<i>Topology Settings</i>	18
<i>Select Campus Fabric Nodes</i>	18
<i>Configure Networks</i>	20
<i>Networks</i>	21
<i>Other IP Configuration</i>	22
<i>Configure Campus Fabric Ports</i>	26
<i>Core Switches</i>	26
<i>Distribution Switches</i>	27
<i>Access Switches</i>	29
<i>Campus Fabric Configuration Confirmation</i>	30
Verification	34
<i>BGP Underlay</i>	35
<i>Purpose</i>	35
<i>Action</i>	35
<i>Verification of BGP Peering</i>	35
<i>EVPN VXLAN Verification Between Core and Distribution Switches</i>	37
<i>Verification of the EVPN Database on Both Core Switches</i>	37
<i>Verification of VXLAN Tunnelling Between Distribution and Core Switches</i>	38
<i>External Campus Fabric Connectivity Through the Border GW Core EX9204 Switches</i>	41
EVPN Insights	43
Summary	45
Additional Information	45
<i>Campus Fabric Core Distribution CRB Configurations</i>	45
<i>Configuration of the EVPN VXLAN Overlay and Virtual Networks</i>	47
<i>Configuration of the Layer 2 ESI-LAG Between the Distribution Switches and the Access Switches</i>	52
<i>Configuration of the Layer 2 ESI-LAG Between the Core Switches and SRX Series Firewalls</i>	54



Overview

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/Rapid Tree Protocols (STP/RSTP). A flexible overlay network based on a 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 Endpoint (VTEP) 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 (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).

Campus Fabric Core Distribution: CRB

● L2 VXLAN Gateway
● L3 VXLAN Gateway

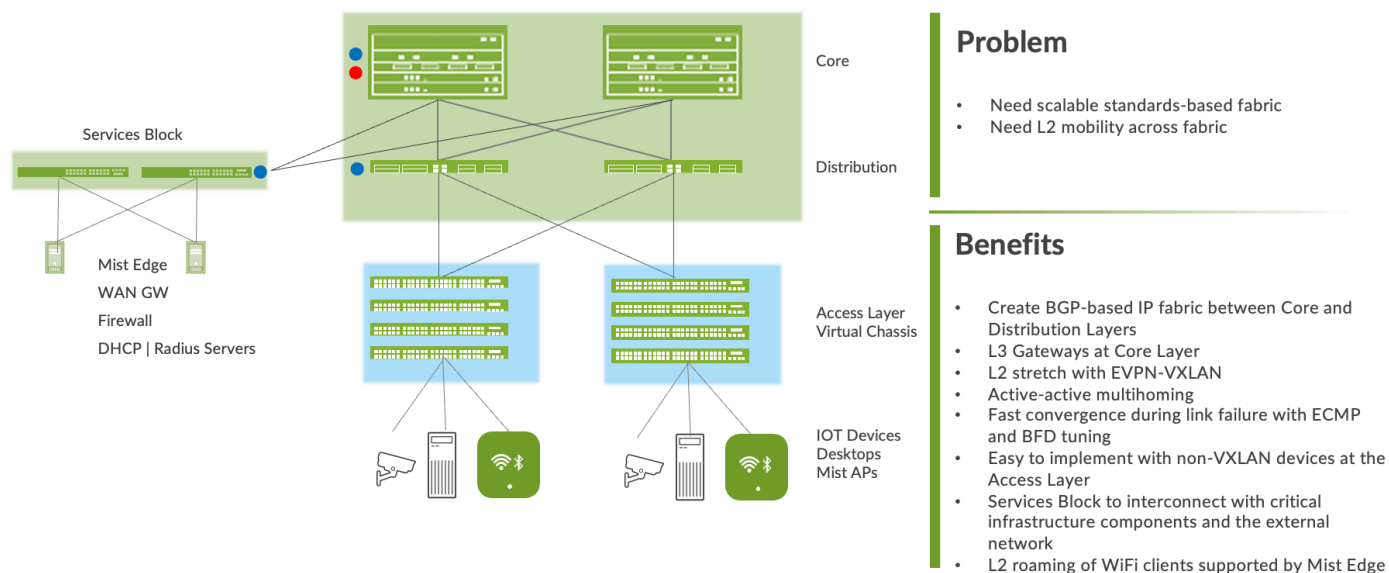


Figure 1: Campus Fabric Core Distribution CRB

A Campus Fabric Core Distribution CRB 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, deployed at the core and distribution layers, 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. L3 Default gateways at the core layer provide higher scale than if placed at the distribution or access layers. This is due to higher performance platforms supported at the layer.
- **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/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. Juniper uses optimal BGP timers between the adjacent layers of the Campus Fabric with Bidirectional Forwarding Detection (BFD) that supports fast convergence in event of a node or link failure and Equal cost multipath (ECMP). For more information, see [Configuring Per-Packet Load Balancing](#).

Technical Overview

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 uses 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 2: 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 3: 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 tunnel. 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 the multihomed Distribution layer switches is called 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 endpoint (VTEP). Before a VTEP sends a frame into a VXLAN tunnel, it wraps the original frame in a VXLAN header that includes a VNI. The VNI maps the packet to the original VLAN at the ingress switch. After applying a VXLAN header, the frame is encapsulated into a UDP/IP packet for transmission to the remote VTEP over the IP fabric, where the VXLAN header is removed and the VNI to VLAN translation happens at the egress switch.

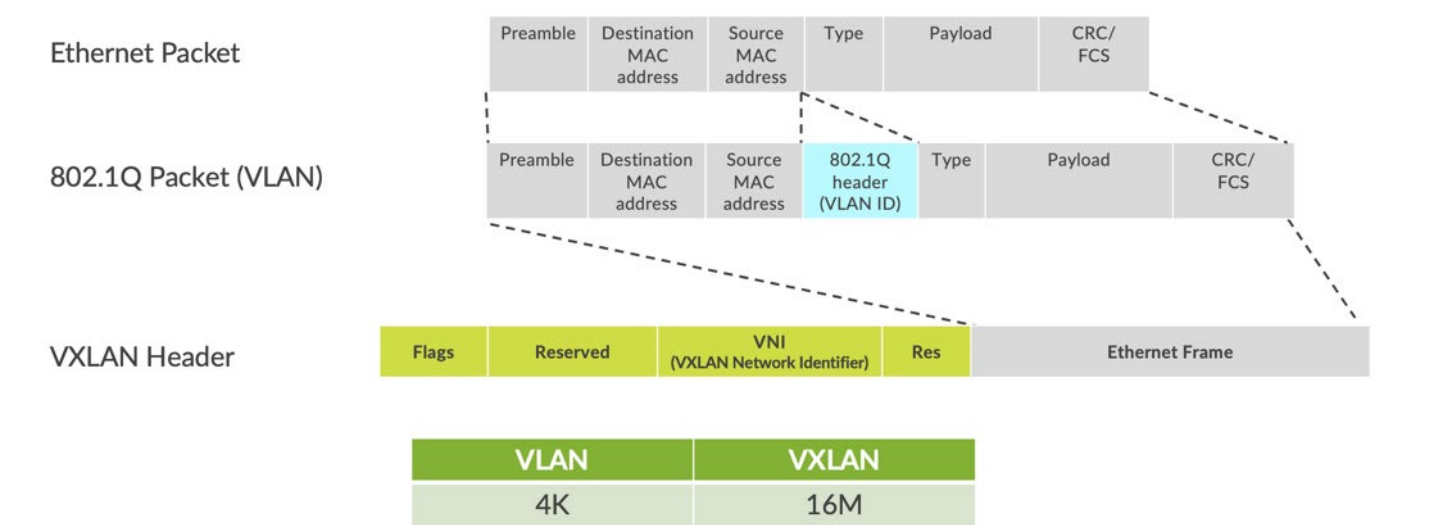


Figure 4: VXLAN Header

VTEPs are software entities tied to the devices' 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 signaling 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.

Overlay Control Plane

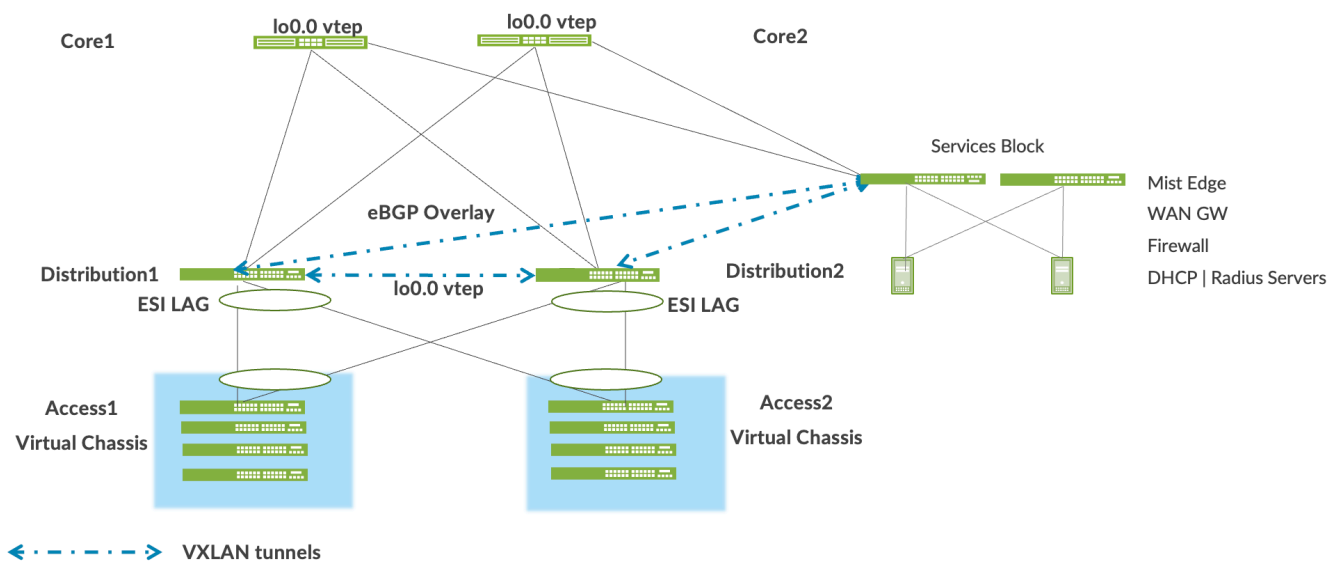


Figure 5: EVPN VXLAN Overlay Network with a Services Block

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.

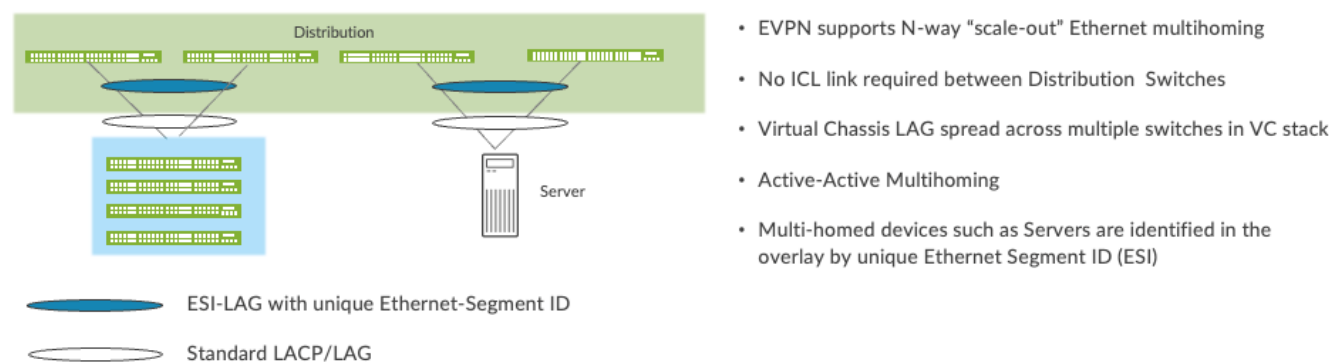


Figure 6: Resiliency and Load Balancing

Services Block

You might need to position critical infrastructure services off a dedicated Access Pair of Juniper switches. This can include WAN and Firewall connectivity, Radius and DHCP Servers as an example. For those customers 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 Instance and additional L3 routing protocols. The Services Block Border capability is supported directly off the core layer or as a dedicated pair of switches.

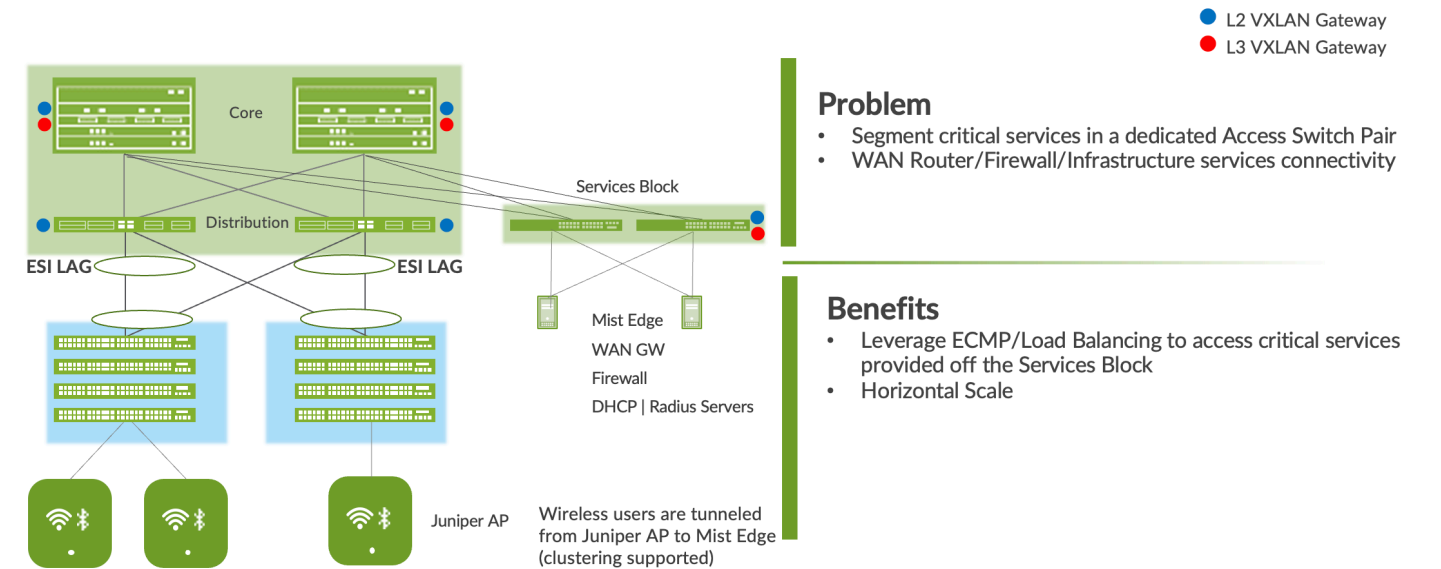
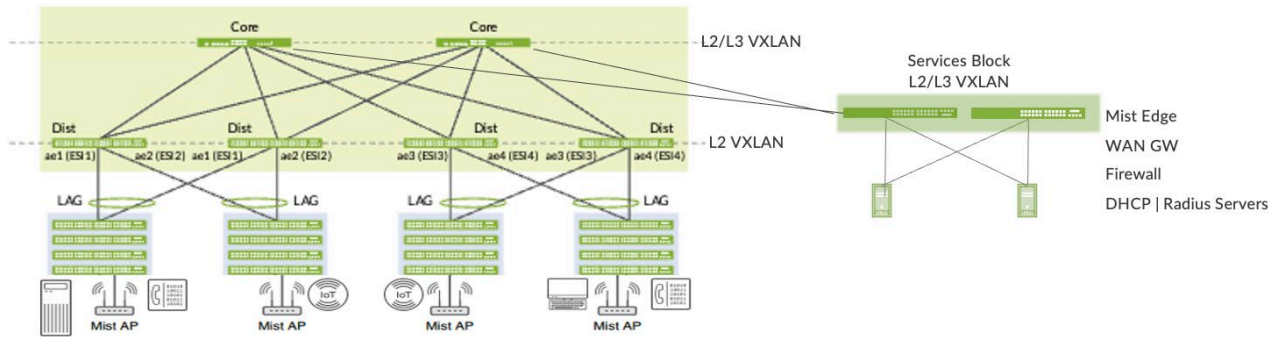


Figure 7: Services Block

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.



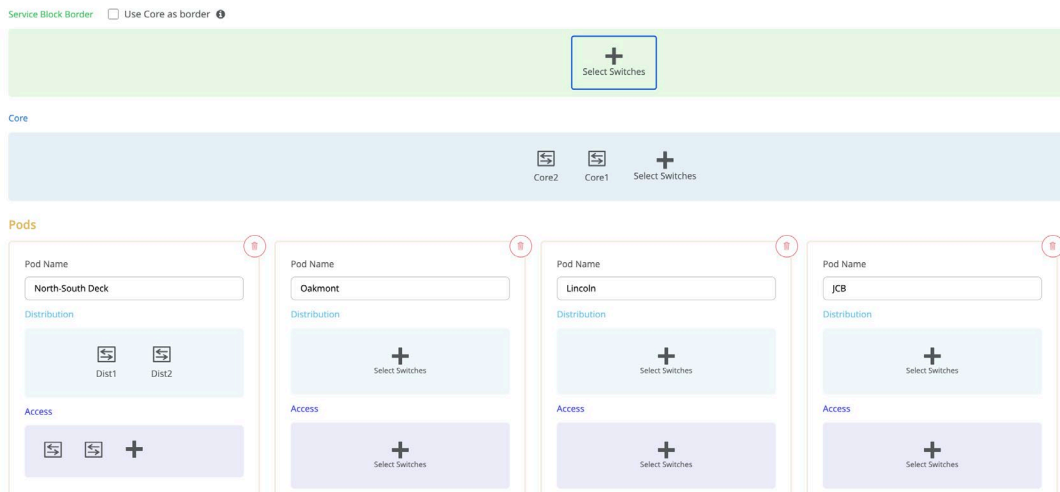
Note: Wireless users are tunneled from Juniper AP to Mist Edge (clustering supported)

Figure 8: 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.

Campus Fabric Organizational Deployment

Mist Campus Fabric supports deployments at the Site and Organizational level. The Organizational deployment shown below, targets Enterprises who need to align with a POD structure:



NOTE: The site level deployment is the focus of this document.

Juniper Access Points

In our network, we choose Mist 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 (SLE) for key wired and wireless performance metrics such as throughput, capacity, roaming, and uptime are addressed in the Mist platform. This JVD uses Mist wired assurance services.
- Marvis—An integrated AI engine that provides rapid wired and wireless troubleshooting, trending analysis, anomaly detection, and proactive problem remediation.

Mist Edge

For large campus networks, 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 for managing wired, wireless, and wan networks. This full stack approach simplifies and automate 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 EX switches, see [How to Connect Mist Access Points and Juniper EX Series Switches](#).

Supported Platforms for Campus Fabric Core Distribution CRB

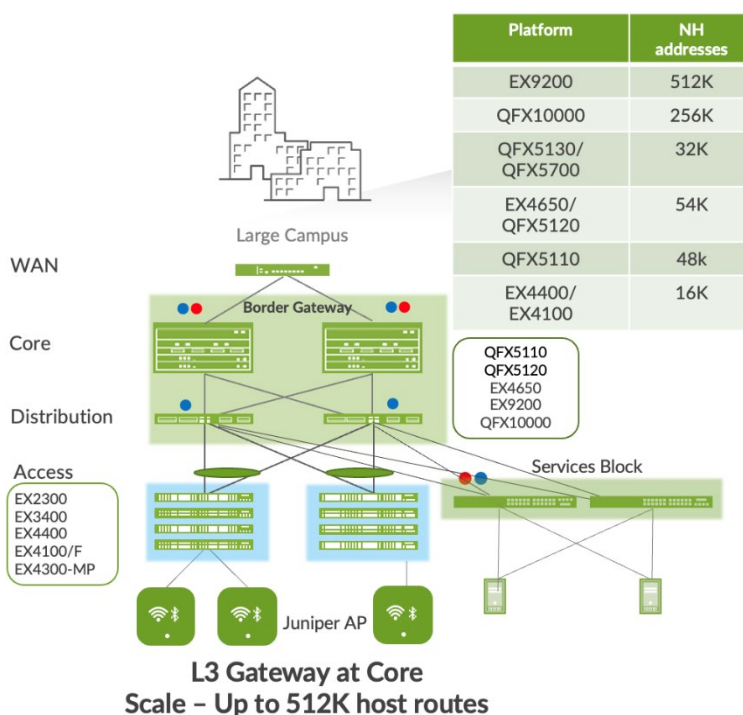
Table 1 lists the supported platforms for Campus Fabric Core Distribution CRB deployment.

Table 1: Supported Platforms for Campus Fabric Core Distribution CRB Deployment

Campus Fabric Core Distribution CRB Deployment	Supported Platforms
Access layer	<ul style="list-style-type: none"> • EX2300 • EX3400 • EX4300 • EX4100 • EX4400
Distribution layer	<ul style="list-style-type: none"> • EX4400-24X • EX4650 • QFX5110 • QFX5120 • QFX5130 • QFX5700
Core layer	<ul style="list-style-type: none"> • EX4650 • EX4400-24X • QFX5110 • QFX5120 • QFX5130 • QFX5700 • QFX10000 • EX92xx

Campus Fabric Core Distribution CRB Deployment	Supported Platforms
Services block	<ul style="list-style-type: none"> EX4400/EX4400-24X EX4650 QFX5110 QFX5120 QFX5130 QFX5700 QFX10000 EX92xx

Campus Fabric Core Distribution CRB Unicast Scale



Juniper Mist Wired Assurance

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 Information Technology (IT) operations from reactive troubleshooting to proactive remediation.

Juniper Mist cloud services are 100% programmable using open Application Programming Interfaces (APIs) for full automation and 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:

- User Interface (UI) or API driven configuration at scale.
- Service Level Expectations (SLE) for key performance metrics such as throughput, capacity, roaming, and uptime.
- Marvis—An integrated AI engine that provides rapid troubleshooting of Full Stack network issues, trending analysis, anomaly detection, and proactive problem remediation.
- Single management system.
- License management.
- Premium analytics for long term trending and data storage.

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>

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.

Choose the topology and allocate device roles

- Define the intent for the topology definition (IP-Clos, Multi-homing etc)
- Choose device roles – access, distribution, core


Choose EVPN Topology

Choose the topology you want to construct and configure related options.




Apply the intent

- Verify, apply and confirm the intent of provisioning the fabric



Define Networks of Interest

- Configure the user networks



Define Physical Connections

- Provide the physical connectivity between – core/distribution and access devices



Campus Fabric Core Distribution CRB Components

This configuration example uses the following devices:

- Two EX9204 switches as core devices, software version: Junos OS Release 21.4R1.12 or later
- Two QFX5120 switches as distribution devices, software version: Junos OS Release 21.4R1.12 or later
- Two access layer EX4400 switches, software version: Junos OS Release 22.1R1.10 or later
- One SRX345 WAN router, software version: Junos OS Release 20.2R3-S2.5 or later
- Juniper Access Points
- Two Linux desktops that act as wired clients

NOTE: Juniper's recommended software version for Campus Fabric IP Core Distributed CRB is available under the EVPN/VXLAN CRB section at: https://supportportal.juniper.net/s/article/Junos-Software-Versions-Suggested-Releases-to-Consider-and-Evaluate?language=en_US

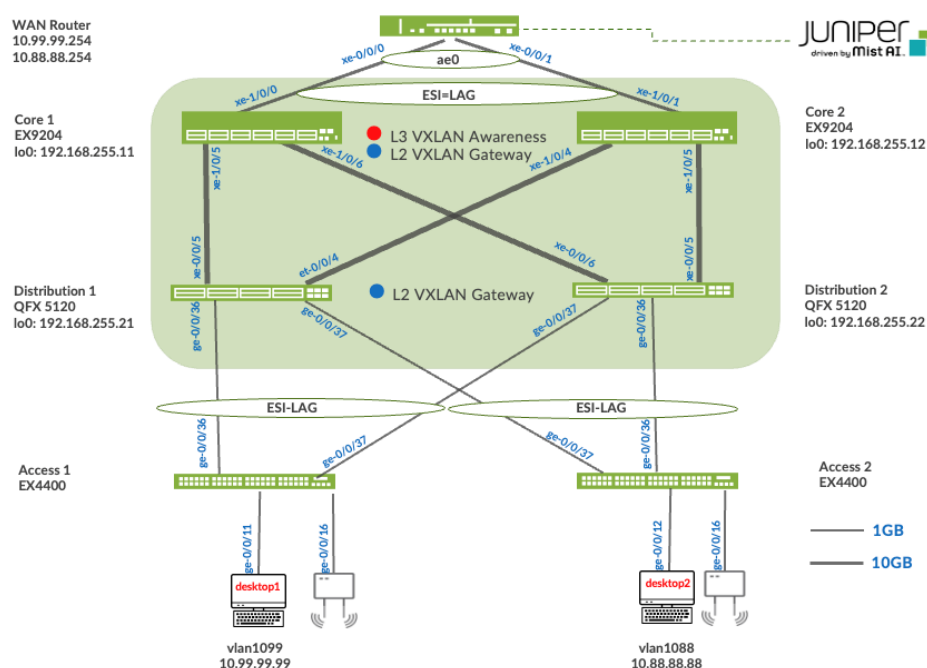


Figure 9: Topology

Juniper Mist Wired Assurance

Wired Assurance, through the Mist UI, can be used to centrally manage all Juniper switches. Juniper Mist Wired Assurance gives you full visibility on the devices that comprise your network's access layer. The Juniper Mist portal provides a user interface to access your architecture through the AI-driven cloud services with your Juniper Mist

account. You can monitor, measure, and get alerts on key compliance metrics on the wired network. This includes switch version and Power Over Ethernet (PoE) compliance, switch-AP affinity, and Virtual LAN (VLAN) insights.

Juniper Switch Onboarding to the Mist Cloud:

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

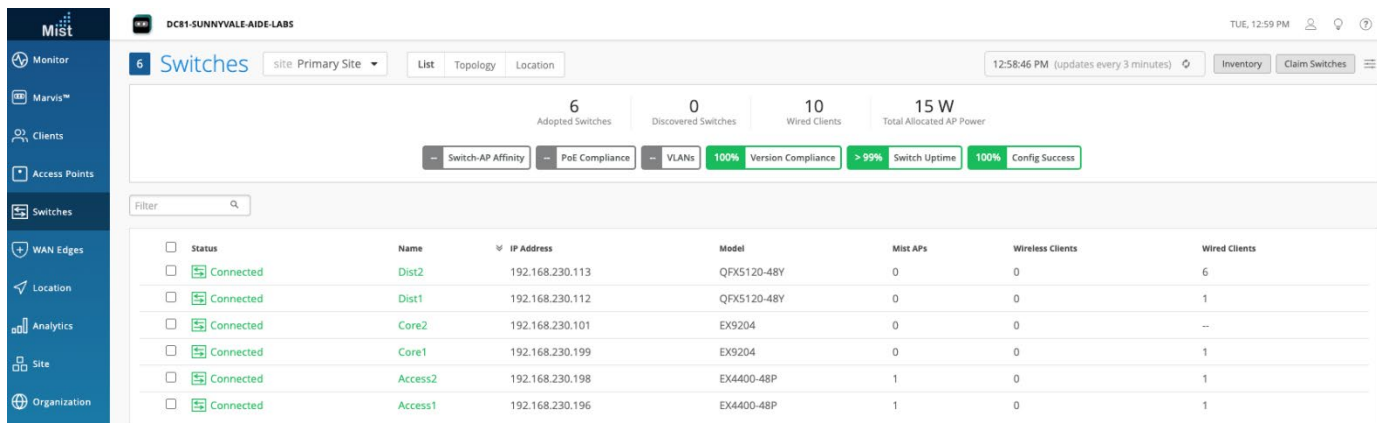
Wired Assurance, through the Mist UI, is used to build a Campus Fabric Core Distribution CRB from ground up. This includes the following:

- Assignment of p2p links between the core and distribution layers.
- Assignment of unique BGP AS numbers per device participating in the underlay and overlay.
- Creation of Virtual Routing and Forwarding (VRF) instances allow you to logically segment traffic. This also includes the assignment of new or existing VLANs to each representative VRF.
- IP addressing of each Layer 3 (L3) gateway Integrated Routing and Bridging (IRB) assigned to the distribution layer.
- IP addressing of each lo0.0 loopback.
- Configuration of routing policies for underlay and overlay connectivity.
- Optimized Maximum Transmission Unit (MTU) settings for p2p underlay, L3 IRB, and ESI-LAG bundles.
- Downloadable connection table (.csv format) that can be used by those involved in the physical buildout of the Campus Fabric.
- Graphical interface depicting 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.



Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients
Connected	Dist2	192.168.230.113	QFX5120-48Y	0	0	6
Connected	Dist1	192.168.230.112	QFX5120-48Y	0	0	1
Connected	Core2	192.168.230.101	EX9204	0	0	—
Connected	Core1	192.168.230.199	EX9204	0	0	1
Connected	Access2	192.168.230.198	EX4400-48P	1	0	1
Connected	Access1	192.168.230.196	EX4400-48P	1	0	1

Figure 10: Switch Inventory

Overview

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 (Organization, 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, Site) override the broader settings defined at the Organization level.

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

You can include individual Command Line Interface (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 setting might be replaced or appended).

NOTE: If you run CLI commands for items not native to the Mist UI, this configuration data is applied last; overwriting existing configuration data within the same stanza. You can access the CLI Command option from the Switch Template or individual Switch configuration.

CLI CONFIGURATION

Additional CLI Commands ⓘ

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

Switch Templates

1 Template

TEMPLATE	SITES	SWITCHES
campus-fabric	1	6

Topology

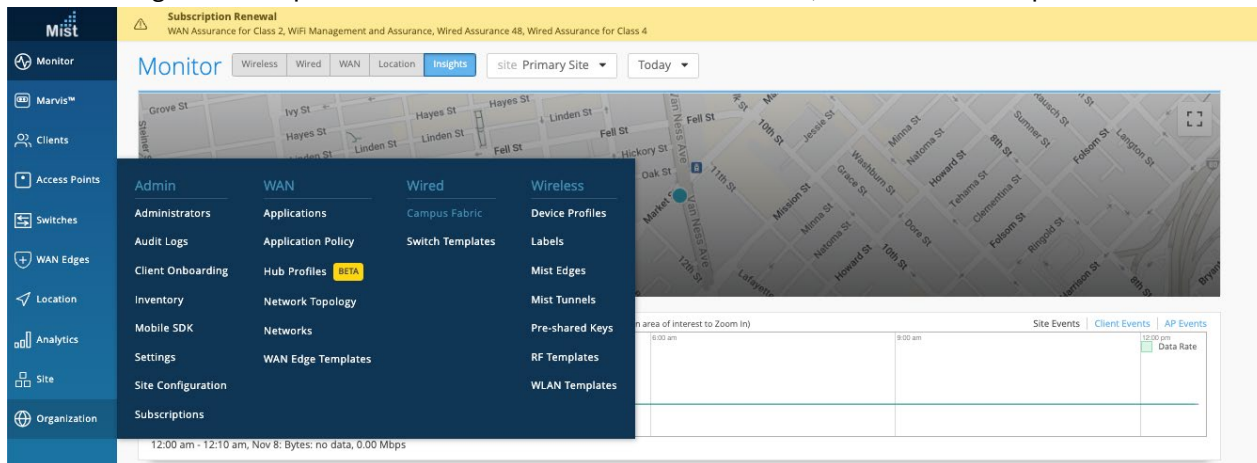
Wired Assurance provides the template for LAN and Loopback IP addressing for each core and distribution device once the device’s management IP address is reachable. Each device is provisioned with a /32 loopback address and /31 point-to-point interfaces that interconnect core and distribution devices within the Campus Fabric Core Distribution.

The devices such as the access layer of switches connect to the distribution layer using standard LAG; while the distribution uses ESI-LAG in a multihoming, load balancing manner.

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

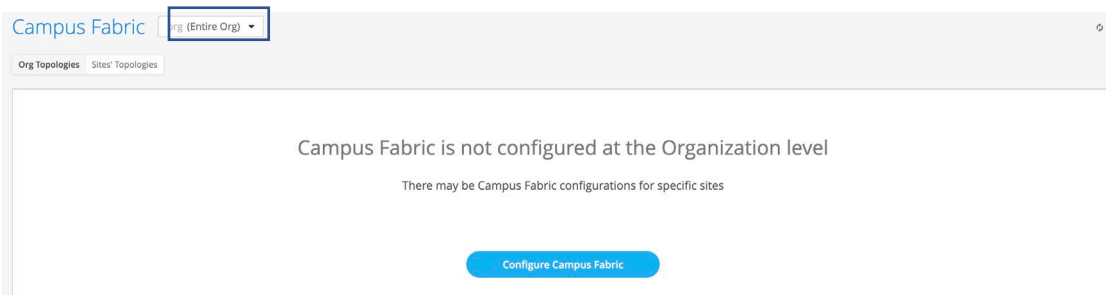
Create the Campus Fabric

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



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

Campus Fabric Org Build



Campus Fabric Site Build

Campus Fabric site: Primary Site

Org Topologies

Campus Fabric is not configured at the Organization level

Site Topologies

Campus Fabric is not configured for this site

There may be Campus Fabric configurations for other sites or the organization as a whole

[Configure Campus Fabric](#)

NOTE: Campus Fabric Site deployment is the focus of this document.


Choose the Campus Fabric Topology


Select the Campus Fabric Core-Distribution option below:


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

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

Virtual Gateway v4 MAC Address
Virtual gateway MAC auto-generated per network on the L3 gateway
☐ Enabled ☒ Disabled

TOPOLOGY SETTINGS

BGP Local AS

(2-byte or 4-byte)

Loopback prefix ?

Subnet ?

(xxx.xxx.xxx.xxx/xx)

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

- Configuration—Provide a name in accordance with company standards
- Topology Sub-type:
 - CRB
 - ERB
- Virtual Gateway v4 MAC Address
 - Only applicable to CRB
 - Mist provides a unique MAC address for each L3 IRB

NOTE: CRB uses virtual-gateway addressing which provides a shared IP addresses among all devices participating in the L3 IRB as well as a unique IP address per device within the IRB/VLAN. Deployments that require a routing protocol on the L3 IRB must use CRB with virtual-gateway addressing.

NOTE: You must choose CRB if most of their traffic patterns are north-south while ERB should be selected if east-west traffic patterns exist as well as IP Multicast.

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.
- Loopback prefix: represents the range of IP addresses associated with each device's loopback address. You can use whatever range suits your deployment. VXAN tunnelling using a VTEP is associated with this address.
- Subnet: represents the range 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.

NOTE: We recommend default settings for all options unless it conflicts with other networks attached to the campus fabric. The point-point links between core and distribution layers use /31 addressing to conserve addresses.

Select Campus Fabric Nodes

Select devices to participate at each layer of the Campus Fabric Core Distribution CRB. 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.

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 CRB 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.

Service Block Border ☒ Use Core as border ⓘ

Core

+
Select Switches

Distribution

Access

Filter Q

<input type="checkbox"/>	Name	MAC Address	Serial	Router ID	Model
<input type="checkbox"/>	Dist2	d8:53:9a:64:b5:c0	XH3121410874	192.168.255.22	QFX5120-48Y
<input type="checkbox"/>	Dist1	d8:53:9a:64:6f:c0	XH3121410895	192.168.255.21	QFX5120-48Y
<input checked="" type="checkbox"/>	Core2	f4:b5:2f:f3:f4:00	JN122EFFFRFC	192.168.255.12	EX9204
<input checked="" type="checkbox"/>	Core1	f4:b5:2f:f4:04:00	JN122EFF5RFC	192.168.255.11	EX9204
<input type="checkbox"/>	Access1	00:cc:34:f4:72:00	ZD4422070133	192.168.255.15	EX4400-48P
<input type="checkbox"/>	Access2	00:cc:34:f3:cf:00	ZD4422030024	192.168.255.16	EX4400-48P

Select 2 Cancel

*required

*required

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 User Core as border option.

Once all layers have selected the appropriate devices, you must provide a loopback IP address for each device. This loopback is associated with a logical construct called a VTEP; used to source the VXLAN Tunnel. Campus Fabric Core Distribution CRB has VTEPs for VXLAN tunnelling on the distribution switches and the core switches when enabling the Core Border option.

The loopback addresses and router-ids should be in the same address space. The router-id of the loopback can be customized to differentiate between core, distribution, and access. This can help identify devices if you are troubleshooting or following next hops. The loopback is also used as the router-id and is used for overlay eBGP peering and VXLAN tunnel termination.

The screenshot shows the Mist configuration interface. On the left, there are three panels for selecting switches: Core (Core2, Core1), Distribution (Dist2, Dist1), and Access (Access1, Access2). The Core1 switch is highlighted with a blue box. On the right, a detailed view of Core1 is shown, including its MAC Address (f4:b5:2f:f4:04:00), Model (EX9204), Status (connected), Site (Primary Site), and Router ID (192.168.255.11).

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

The loopback prefix is used for import and export policies. The subnet addresses are used for point-to-point links throughout the Fabric. Mist automatically creates policies that import, and export loopback addresses used within the Campus Fabric. The selection of fabric type displays with default settings, which can be adapted as required.

The screenshot shows the Mist configuration interface for the Loopback prefix and Subnet fields. The Loopback prefix field is set to /24. The Subnet field is set to 10.255.240.0/20. Below the Subnet field, the format (xxx.xxx.xxx.xxx/xx) is displayed.

Configure Networks

Enter the Network information such as VLANs and VRF (routing instances for traffic isolation purposes) options. VLANs are mapped to Virtual Network Identifier (VNIs) and can optionally be mapped to VRFs to provide customers a way to logically separate traffic patterns such as IoT devices from Corp IT.

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

No networks defined

[Create New Network](#) [Add Existing Network](#)

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

No networks defined

VRF

Configuration

☐ Enabled ☒ Disabled

Instances

[Add VRF Instance](#)

Networks

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

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

Shared Elements

NETWORKS

Named VLAN IDs that can be used by Port Profiles

★ System defined

vlan1033	1033	>
vlan1088	1088	>
vlan1099	1099	>
vlan1100	100	>

Search [Add Network](#)

Back to the Campus Fabric build, select the existing template includes Layer 2 (L2) VLAN information. All VLAN and IP information is inherited from the 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
 Edit Network

Name
 vlan1099

VLAN ID
 1099
(1 - 4094 or {{siteVar}})

Subnet
 10.99.99.0/24

Other IP Configuration

Mist Wired Assurance provides automatic IP addressing Integrated Routing and Bridging (IRBs) for each of the VLANs. Port Profiles and Port Configuration then associate the VLAN with specified ports. In this case, we selected Campus Fabric CRB at the onset of the Campus Fabric build.

CONFIGURATION
 Topology Name
 Campus Fabric CRB

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

Virtual Gateway v4 MAC Address
 Virtual gateway MAC auto-generated per network on the L3 gateway
☐ Enabled ☒ Disabled

This option uses Virtual Gateway addressing for all devices participating in the L3 subnet. Core1 and Core2 switches are configured with shared IP address for each L3 subnet. This address is shared amongst both core switches and acts as the Default Gateway for all devices within the VLAN. Each core device also receives a unique IP address chosen by Mist. All addresses can be managed per customer requirements. Mist assigns IP address for Core1/2 starting at the beginning of each subnet and the end user can modify these IP addresses accordingly. For example, this deployment uses x.x.x.1 as a Default Gateway for each VLAN and x.x.x.254 as the gateway of last resort (MX router in this case) for all traffic leaving the VLAN. Therefore, we modify the IP addresses assigned to Core1 from x.x.x.1 to x.x.x.3 allowing the Virtual Gateway to use x.x.x.1 for all VLANs.

OTHER IP CONFIGURATION			OTHER IP CONFIGURATION		
Network-specific IP configuration for each Core switches			Network-specific IP configuration for each Core switches		
Edit Core1 ✓ ✕			Edit Core2 ✓ ✕		
vlan1033	10.33.33.2	➤	vlan1033	10.33.33.3	➤
vlan1088	10.88.88.2	➤	vlan1088	10.88.88.3	➤
vlan1099	10.99.99.2	➤	vlan1099	10.99.99.3	➤

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

VRF

Configuration

☒ Enabled ☐ Disabled

Instances

No VRF instances defined

[Add VRF Instance](#)

VRF

New VRF Instance ✓ ✕

Name

corp-it

Networks

☐ vlan1088 ☒ vlan1099 ☐ vlan1033

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 Juniper SRX Series Firewalls handles inter-VRF routing. See [Figure 9](#).

Notice the SRX Series Firewalls participates in the VLANs defined within the Campus Fabric and is the gateway of last resort for all traffic leaving the subnet. 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 Juniper SRX Series Firewalls: 10.99.99.254.

New Extra Route

✓ ✕

Route

0.0.0.0/0

Via

10.99.99.254

Create two additional VRFs:

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

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

vlan1033	1033 >
vlan1088	1088 >
vlan1099	1099 >

[Create New Network](#) [Add Existing Network](#)

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

Access1	3 Static >
Access2	3 Static >

VRF

Configuration ☒ Enabled ☐ Disabled

Instances

corp-it	1 network >
developers	1 network >
guest-wifi	1 network >

[Add VRF Instance](#)

The final step in the Configure Networks section is the Distribution/Access Port Configuration.

DISTRIBUTION / ACCESS PORT CONFIGURATION

Port configuration for ESI-Lag between Distribution and Access switches

Name

crb-lag

Trunk Networks

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

[Show Advanced ▲](#)

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:

Port Enabled
☒ Enabled ☐ Disabled

Description

Mode
☒ Trunk ☐ Access

Port Network (Untagged/Native VLAN)

Speed

Duplex

Mac Limit

(0 - 16383, 0 => unlimited)

PoE
☐ Enabled ☒ Disabled

STP Edge
☐ Yes ☒ No

QoS
☐ Enabled ☒ Disabled

☒ Enable MTU

(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/Access ESI-LAGs have been built, click Continue button at the upper-right section of the Mist UI to move to the next step.

Configure Campus Fabric Ports

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

Ports

Select switch ports for Fabric connections

Core Switches

Switch	Model	Link to Distribution
Core2	EX9204	0/2 ?
<div> <div>FPC 1</div> <div>FPC 2</div> </div> <div> <div>1</div> <div>1 3 5 7</div> <div>1 3 5 7</div> <div>1 3 5 7</div> <div>1 3 5 7</div> <div>EX9200-32XS</div> </div>		
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 ?

NOTE: We recommend that you have the output of the show lldp neighbors command from each switch (assuming LLDP is enabled before the switches were selected). This output provides a source of truth for which ports should be selected during at each layer.

Core Switches

Core1:

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

xe-1/0/5

Port Type

☐ ge
 ☐ mge
 ☒ xe
 ☐ et

Distribution Switches

Search

Dist2

Dist1

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

1

SFP+

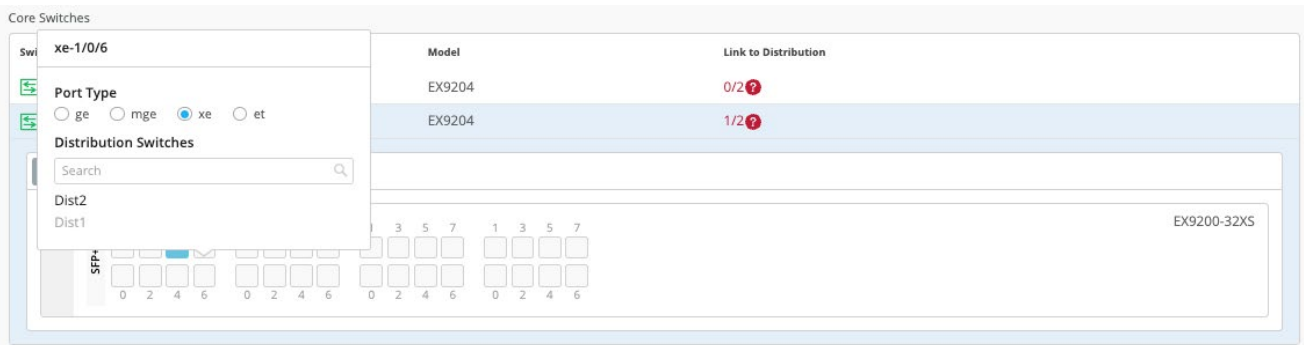
1 3 5 7

1 3 5 7

1 3 5 7

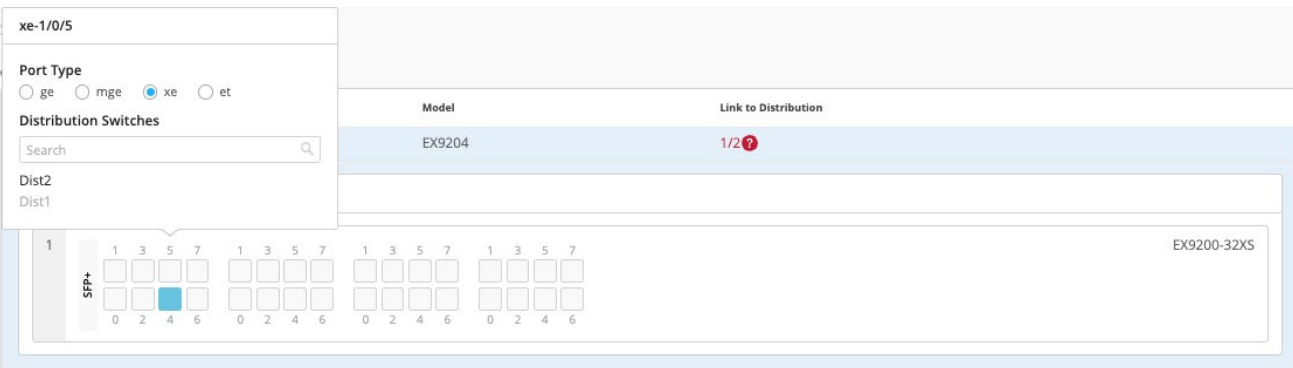
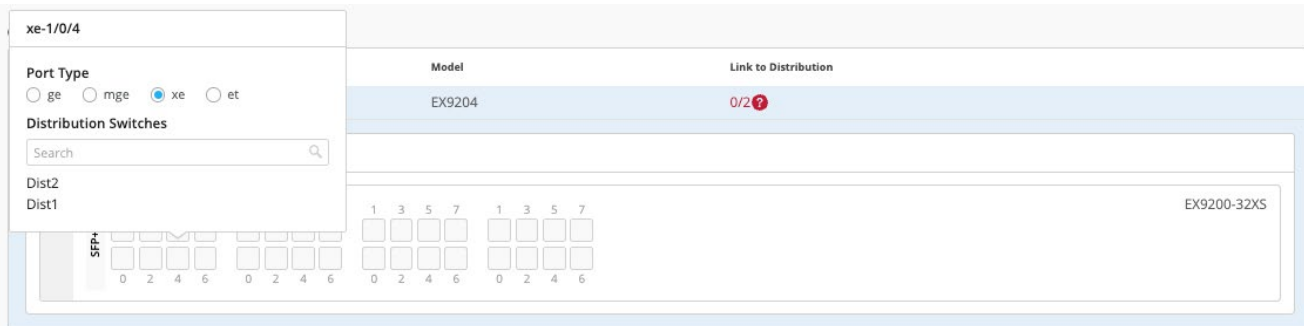
1 3 5 7

EX9200-32XS



Core2:

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



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.

Distribution Switches

xe-0/0/5

Port Type
☐ ge ☐ mge ☒ xe ☐ et

Port Connection
[Link to Core](#)
[Link to Access](#)

Model

Model	Link to Core	Link to Access
QFX5120-48Y	0/2 ?	0/2 ?
QFX5120-48Y	0/2 ?	0/2 ?

SFP28

16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

QSF28 Uplink

48	50	52	54
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Edit Ports for all QFX5120-48Y

xe-0/0/4

Port Type
☐ ge ☐ mge ☒ xe ☐ et

Port Connection
[Link to Core](#)
[Link to Access](#)

Model

Model	Link to Core	Link to Access
QFX5120-48Y	0/2 ?	0/2 ?
QFX5120-48Y	1/2 ?	0/2 ?

SFP28

0	2	4	6	8	10	12	14
<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

QSF28 Uplink

48	50	52	54
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Edit Ports for all QFX5120-48Y

Select Link to Access and choose ge-0/0/36 and ge-0/0/37 terminating on Access Switches 1 and 2 respectively.

Distribution Switches

QFX5120-48Y

Switch

Switch	Model
Dist2	QFX5120-48Y
Dist1	QFX5120-48Y

ge-0/0/36

Port Type
☒ ge ☐ mge ☐ xe ☐ et

Access Switches

Access2
 Access1

Link to Access

Link to Access
0/2 ?
0/2 ?

SFP28

0	2	4	6	8	10	12	14
<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

QSF28 Up

48	50	52	54
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Edit Ports for all QFX5120-48Y

Distribution Switches

QFX5120-48Y

Switch

Switch	Model
Dist2	QFX5120-48Y
Dist1	QFX5120-48Y

ge-0/0/37

Port Type
☒ ge ☐ mge ☐ xe ☐ et

Access Switches

Access2
 Access1

Link to Access

Link to Access
0/2 ?
1/2 ?

SFP28

0	2	4	6	8	10	12	14
<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

QSF28 Up

48	50	52	54
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Edit Ports for all QFX5120-48Y

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

- Link to Core
 - xe-0/0/6 – Core1
 - xe-0/0/5 – Core2
- Link to Access
 - ge-0/0/36 – Access2
 - ge-0/0/37 – Access1

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.

Access1/2:

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 0/1 (system default numbering) on Access1/2 respectively.

Access Switches

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

Switch	Model	Link to Distribution	AE Index
Access2	EX4400-48P	2/2	0

Port Selection Grid:

RJ-45	0	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46
	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47

QSF28 VCP: 0, 1

SFP+ Uplink: 0, 1, 2, 3

Access Switches

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

Switch	Model	Link to Distribution	AE Index
Access2	EX4400-48P	2/2	0
Access1	EX4400-48P	2/2	1

Port Selection Grid:

RJ-45	0	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46
	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47

QSF28 VCP: 0, 1

SFP+ Uplink: 0, 1, 2, 3

Once you have completed selecting all requisite port combinations, select the Continue button at the upper right-hand corner of the Mist UI.

Campus Fabric Configuration Confirmation

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

Confirm

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

The screenshot displays the 'Confirm' page in the Mist UI. The main area shows a network topology with three layers: Core (Core1, Core2), Distribution (Dist1, Dist2), and Access (Access1, Access2). Core1 and Core2 are connected to each other and to the Distribution layer. Dist1 and Dist2 are connected to the Access layer. A sidebar on the right shows details for Core1, including MAC Address, Model, Status, Site, Router ID, and a table of VLANs (1088, 1099, 1033) and Connections to Distribution (Dist1, Dist2).

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

✕ 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 at each device with unique BGP autonomous system numbers. The primary goal of the underlay is to leverage ECMP for load balancing traffic on a per packet level for device loopback reachability. The primary goal of the eBGP overlay is support of customer traffic using EVPN-VXLAN.
- IP addressing of each L3 gateway IRB located on Core1 and Core2.
- IP addressing of each lo0.0 loopback.
- Configuration of routing policies for underlay and overlay connectivity.
- Optimized MTU settings for p2p underlay, L3 IRB, and ESI-LAG bundles.
- VXLAN to VLAN mapping using VNI addresses that are automatically assigned
- VRF creation of corp-it, developers, and guest-wifi and VLAN associated with each VRF.
- VXLAN tunnelling creation between distribution devices and distribution-core devices (in support of the northbound SRX Series Firewalls 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.

Applying Changes

Campus Fabric configuration successfully saved to the Mist Cloud

Configuration will be immediately pushed to switches or when they next come online and may require up to 10 minutes to complete.

[Close Campus Fabric Configuration](#)

Once you click Close Campus Fabric Configuration, you can view a summary of the newly created Campus Fabric Core Distribution CRB.

Campus Fabric site: Primary Site Create Campus Fabric

Org Topologies

Campus Fabric is not configured at the Organization level

Site Topologies

Name	Topology ID	Site	Type	Routed At	Date Created
Campus Fabric CRB	6b2d12a5-3e97-4310-9965-a6e9e3e1891c	Primary Site	Campus Fabric Core-Distribution	Core	12:57:37 PM, Mar 28 2023

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.

< Campus Fabrics : **Campus Fabric CRB** Edit Configuration Delete Download Connection Table

Core

Distribution

Access

BGP Summary

Neighbor Information 2:04 PM (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	1h 3m	2	3	143	143	default	Underlay
Connected	Established	10.255.240.6	65002	65003	1h 3m	2	2	143	142	default	Underlay
Connected	Established	192.168.255.11	65002	65003	1h 3m	54	64	228	236	default	Overlay
Connected	Established	192.168.255.12	65001	65003	1h 3m	57	69	228	219	default	Overlay

Dist1

MAC Address: d8:53:9a:64:6f:c0
 Model: QFX5120-48Y
 Status: connected
 Site: Primary Site
 Router ID: 192.168.255.21

VLANs

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

Connections to Core

Switch	RX Bytes	TX Bytes	Link Status
Core2	3 GB	2.4 GB	Up
Core1	2.7 GB	2.7 GB	Up

Connections to Access

Switch	RX Bytes	TX Bytes	Link Status
Access2	921.5 MB	755.1 MB	Up
Access1	2.9 GB	3.8 GB	Up

[Remote Shell](#) [Switch Insights](#)

Connection Table spreadsheet:

Role 1	Switch 1	Mac 1	Model 1	Serial 1	Site 1	Port Role 1	AE 1	Port 1	< --- >	Port 2	AE 2	Port Role 2	Site 2	Serial 2	Model 2	Mac 2	Switch 2	Role 2
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/5	< --- >	xe-1/0/5		downlink	Primary Site	JN122EFFFFFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/6	< --- >	xe-1/0/6		downlink	Primary Site	JN122EFFFFFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag	0	ge-0/0/36	< --- >		0	esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag	1	ge-0/0/37	< --- >		1	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	JN122EFFFFFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	uplink		xe-0/0/5	< --- >	xe-1/0/5		downlink	Primary Site	JN122EFFFFFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag	0	ge-0/0/37	< --- >		0	esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag	1	ge-0/0/36	< --- >		1	esi-lag	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access

Apply VLANs to Access Ports

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

In this case, Desktop1/2 are associated with different ports on each access switch which requires the configuration to be applied to Access1/2 respectively. See [Figure 9](#).

It is also noteworthy that Mist Access Points connect to the same port on Access1/2 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. Further, all access switches (defined by EX4400 Switch as an example) associated the AP port profile with ge-0/0/16 without needing to configure each independent switch.

Select Switches Configuration

core

model:EX9204

distribution

model:QFX5120*

access

model:EX4400*

default

all remaining switches

Info Port Config CLI Config

Apply port profiles to port ranges on matching switches

ge-0/0/16 myap >

Unassigned ports Default

Add Port Range

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.

PORT CONFIGURATION

Port Profile Assignment

★ Site, Template, or System Defined

Edit Port Range

☐ Port Aggregation

Port IDs

ge-0/0/11

(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Interface

☒ L2 interface
☐ L3 interface
☐ L3 sub-interfaces

Configuration Profile

vlan1099

vlan1099(1099), access

☐ Enable Dynamic Configuration

The Switch Template definition for vlan1099 is shown below, representing attributes associated with VLANs such as dot1x authentication, QoS, and Power over Ethernet (PoE). Vlan1088 and vlan1033 need to be configured in a similar fashion.

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

Verification

Verification of the Campus Fabric Core Distribution CRB deployment. See [Figure 9](#). Currently, there are two desktops to validate the fabric. Let's take a quick look to see if Desktop1 can connect internally and externally.

```

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

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

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

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

```

Validation steps:

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

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

Site Topologies

Name	Topology ID	Site	Type	Routed At	Date Created
Campus Fabric CRB	6b2d12a5-3e97-4310-9965-a6e9e3e1891c	Primary Site	Campus Fabric Core-Distribution	Core	12:57:37 PM, Mar 28 2023

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

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

BGP Underlay

Purpose

Verifying the state of eBGP between core 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, bi-directional forwarding, to decrease convergence times during failures.
- loopback reachability to support VXLAN tunnelling.

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

Action

Verify that BGP sessions are established from core devices with 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 Mist UI or using an external application such as SecureCRT or Putty.

Verification of BGP Peering

Dist1:

Access Remote Shell via the bottom right of the Campus Fabric, from the switch view or via Secure Shell (SSH).

```
root@Dist1> show bgp summary
Warning: License key missing; requires 'bgp' license

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 4 4 0 0 0 0
Peer 85 AS 49 InPkt OutPkt Out0 Flaps Last Up/Dwn State#Active/Received/Accepted/Damped...
10.255.240.2 65001 29 29 0 0 11:33 Establ
inet.0: 2/2/2/0
10.255.240.6 65002 30 29 0 0 11:32 Establ
inet.0: 2/2/2/0
192.168.255.11 65002 67 75 0 0 11:32 Establ
bgp.evpn.0: 16/43/43/0
guest-wifi.evpn.0: 0/2/2/0
developers.evpn.0: 0/2/2/0
corp-it.evpn.0: 0/2/2/0
default-switch.evpn.0: 14/34/34/0
__default_evpn__.evpn.0: 2/3/3/0
192.168.255.12 65001 69 71 0 0 11:33 Establ
bgp.evpn.0: 33/42/42/0
guest-wifi.evpn.0: 2/2/2/0
developers.evpn.0: 2/2/2/0
corp-it.evpn.0: 2/2/2/0
default-switch.evpn.0: 25/33/33/0
__default_evpn__.evpn.0: 2/3/3/0
{master:0}
root@Dist1>
```

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

It also shows the overlay (192.168.255.x) relationships are established and that it is peering at the correct loopback addresses. This demonstrates loopback reachability.

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

If BGP is not established 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:

```

root@Dist1> ping 192.168.255.11
PING 192.168.255.11 (192.168.255.11): 56 data bytes
64 bytes from 192.168.255.11: icmp_seq=0 ttl=64 time=0.832 ms
64 bytes from 192.168.255.11: icmp_seq=1 ttl=64 time=0.858 ms
^C
--- 192.168.255.11 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.832/0.845/0.858/0.013 ms

{master:0}
root@Dist1> ping 192.168.255.12
PING 192.168.255.12 (192.168.255.12): 56 data bytes
64 bytes from 192.168.255.12: icmp_seq=0 ttl=64 time=11.951 ms
64 bytes from 192.168.255.12: icmp_seq=1 ttl=64 time=11.720 ms
^C
--- 192.168.255.12 ping statistics ---
3 packets transmitted, 2 packets received, 33% packet loss
round-trip min/avg/max/stddev = 11.720/11.835/11.951/0.115 ms

{master:0}
root@Dist1> ping 192.168.255.22
PING 192.168.255.22 (192.168.255.22): 56 data bytes
64 bytes from 192.168.255.22: icmp_seq=0 ttl=63 time=0.698 ms
64 bytes from 192.168.255.22: icmp_seq=1 ttl=63 time=0.779 ms
^C
--- 192.168.255.22 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.698/0.738/0.779/0.041 ms

{master:0}
root@Dist1>

```

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

Dist1: ECMP Loopback reachability to Dist2 through Core1/2

```

root@Dist1> show route forwarding-table destination 192.168.255.22
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
192.168.255.22/32 user    0      10.255.240.2      ucst   1689     7 xe-0/0/4.0
                  10.255.240.6      ucst   1708     7 xe-0/0/5.0

```

Dist2: ECMP Loopback reachability with Dist1 through Core1/2

```

root@Dist2> show route forwarding-table destination 192.168.255.21
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
192.168.255.21/32 user    0      10.255.240.4      ucst   1666     7 xe-0/0/5.0
                  10.255.240.8      ucst   1667     7 xe-0/0/6.0

```

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

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

```
root@Dist1> show bfd session
```

Address	State	Interface	Detect Time	Transmit 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
192.168.255.11	Up		3.000	1.000	3
192.168.255.12	Up		3.000	1.000	3

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

Meaning: At this point, BGP Underlay and Overlay are operational through the verification of eBGP between corresponding layers of the Campus Fabric and that 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, vlan and interface-mode are correct. If pinging the default gateway failed, then troubleshoot underlay connectivity.

Verification of the EVPN Database on Both Core Switches

Core1:

```
root@Core1> show evpn database
Instance: evpn_vrf
```

VLAN	DomainId	MAC address	Active source	Timestamp	IP address
10001		d8:53:9a:64:6f:c0	192.168.255.21	Dec 08 17:49:32	
10001		d8:53:9a:64:b5:c0	192.168.255.22	Dec 08 17:49:33	
10001		f4:b5:2f:f3:fb:f0	192.168.255.12	Dec 08 17:49:32	
10001		f4:b5:2f:f4:0b:f0	irb.0	Dec 08 17:49:32	
11033		00:00:5e:00:01:01	05:00:00:fd:ea:00:00:2b:19:00	Dec 08 17:49:32	10.33.33.1
11033		f0:1c:2d:c8:e8:f0	00:11:00:00:00:01:00:01:02:00	Dec 08 18:35:04	10.33.33.254
11033		f4:b5:2f:f3:fb:f0	192.168.255.12	Dec 08 17:49:32	10.33.33.3
11033		f4:b5:2f:f4:0b:f0	irb.1033	Dec 08 17:49:32	10.33.33.2
11088		00:00:5e:00:01:01	05:00:00:fd:ea:00:00:2b:50:00	Dec 08 17:49:32	10.88.88.1
11088		52:54:00:f7:12:2d	00:11:00:00:00:01:00:01:03:00	Dec 08 19:07:03	10.88.88.88
11088		f0:1c:2d:c8:e8:f0	00:11:00:00:00:01:00:01:02:00	Dec 08 18:55:10	10.88.88.254
11088		f4:b5:2f:f3:fb:f0	192.168.255.12	Dec 08 17:49:32	10.88.88.3
11088		f4:b5:2f:f4:0b:f0	irb.1088	Dec 08 17:49:32	10.88.88.2
11099		00:00:5e:00:01:01	05:00:00:fd:ea:00:00:2b:5b:00	Dec 08 17:49:32	10.99.99.1
11099		00:cc:34:f4:72:00	00:11:00:00:00:01:00:01:03:01	Dec 08 19:02:07	10.99.99.200
11099		52:54:00:74:a0:6f	00:11:00:00:00:01:00:01:03:01	Dec 08 19:08:57	10.99.99.99
11099		f0:1c:2d:c8:e8:f0	00:11:00:00:00:01:00:01:02:00	Dec 08 18:55:04	10.99.99.254
11099		f4:b5:2f:f3:fb:f0	192.168.255.12	Dec 08 17:49:32	10.99.99.3
11099		f4:b5:2f:f4:0b:f0	irb.1099	Dec 08 17:49:32	10.99.99.2

```
root@Core1>
```

Core2:

```

root@Core2> show evpn database
Instance: evpn_vrf
VLAN  DomainId  MAC address      Active source      Timestamp          IP address
10001  d8:53:9a:64:6f:c0  192.168.255.21    Dec 08 17:49:32
10001  d8:53:9a:64:b5:c0  192.168.255.22    Dec 08 17:49:33
10001  f4:b5:2f:f3:fb:f0  irb.0             Dec 08 17:49:32
10001  f4:b5:2f:f4:0b:f0  192.168.255.11    Dec 08 17:49:32
11033  00:00:5e:00:01:01  05:00:00:fd:e9:00:00:2b:19:00 Dec 08 17:49:32  10.33.33.1
11033  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 18:35:05  10.33.33.254
11033  f4:b5:2f:f3:fb:f0  irb.1033          Dec 08 17:49:32  10.33.33.3
11033  f4:b5:2f:f4:0b:f0  192.168.255.11    Dec 08 17:49:33  10.33.33.2
11088  00:00:5e:00:01:01  05:00:00:fd:e9:00:00:2b:50:00 Dec 08 17:49:32  10.88.88.1
11088  52:54:00:f7:12:2d  00:11:00:00:00:01:00:01:03:00 Dec 08 19:07:03  10.88.88.88
11088  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 18:55:11  10.88.88.254
11088  f4:b5:2f:f3:fb:f0  irb.1088          Dec 08 17:49:32  10.88.88.3
11088  f4:b5:2f:f4:0b:f0  192.168.255.11    Dec 08 17:49:33  10.88.88.2
11099  00:00:5e:00:01:01  05:00:00:fd:e9:00:00:2b:5b:00 Dec 08 17:49:32  10.99.99.1
11099  00:cc:34:f4:72:00  00:11:00:00:00:01:00:01:03:01 Dec 08 19:02:07  10.99.99.200
11099  52:54:00:74:a0:6f  00:11:00:00:00:01:00:01:03:01 Dec 08 19:08:57  10.99.99.99
11099  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 18:55:05  10.99.99.254
11099  f4:b5:2f:f3:fb:f0  irb.1099          Dec 08 17:49:32  10.99.99.3
11099  f4:b5:2f:f4:0b:f0  192.168.255.11    Dec 08 17:49:33  10.99.99.2
root@Core2>

```

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 in each core switch. These entries are learned through the Campus Fabric from the ESI LAGs off D1st1/2.

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

Dist

```

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

```

Core

```

root@Core1> show configuration | display set | display inheritance | match 1099
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.2/24 virtual-gateway-address 10.99.99.1
set routing-instances corp-it interface irb.1099
set routing-instances evpn_vrf vlans vlan1099 vlan-id 1099
set routing-instances evpn_vrf vlans vlan1099 l3-interface irb.1099
set routing-instances evpn_vrf vlans vlan1099 vxlan vni 11099

```

Verification of VXLAN Tunnelling Between Distribution and Core Switches

Dist1:

```

root@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   192.168.255.21 lo0.0  0
RVTEP-IP      L2-RTT      IFL-Idx  Interface  NH-Id  RVTEP-Mode  ELP-IP      Flags
192.168.255.11 default-switch  547      vtep.32769  1709    RNVE
192.168.255.12 default-switch  548      vtep.32770  1717    RNVE
192.168.255.22 default-switch  550      vtep.32771  1730    RNVE

```

Core1:

```

root@Core1> show ethernet-switching vxlan-tunnel-end-point remote summary
Logical System Name      Id  SVTEP-IP      IFL  L3-Idx  SVTEP-Mode  ELP-SVTEP-IP
<default>
RVTEP-IP                L2-RTT                                IFL-Idx  Interface  NH-Id  RVTEP-Mode  ELP-IP      Flags
192.168.255.12          evpn_vrf                                390      vtep.32770  750    RNVE
192.168.255.21          evpn_vrf                                385      vtep.32769  749    RNVE
192.168.255.22          evpn_vrf                                392      vtep.32771  781    RNVE

```

Finally, let us validate that Core1 is receiving Desktop 1's MAC address through MP-BGP via Type2 EVPN routes.

```

root@Core1> show route receive-protocol bgp 192.168.255.21 evpn-mac-address 52:54:00:74:a0:6f table bgp.evpn.0

Warning: License key missing; requires 'bgp' license

bgp.evpn.0: 100 destinations, 150 routes (100 active, 0 holddown, 0 hidden)
Prefix                Nexthop          MED  Lclpref  AS path
2:192.168.255.21:1::11099::52:54:00:74:a0:6f/304 MAC/IP
*                      192.168.255.21                      65003 I
2:192.168.255.21:1::11099::52:54:00:74:a0:6f::10.99.99.99/304 MAC/IP
*                      192.168.255.21                      65003 I
root@Core1>

```

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

We next verify if Desktop1's MAC address is mapped to the correct VTEP interface on Core1:

```

root@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_vrf
Bridging domain : default+1, VLAN : 1, VNID : 10001
MAC      MAC      Logical      Remote VTEP
address  flags  interface  IP address
d8:53:9a:64:6f:c0  DRP    vtep.32769  192.168.255.21
f4:b5:2f:f3:fb:f0  DRP    vtep.32770  192.168.255.12
d8:53:9a:64:b5:c0  DRP    vtep.32771  192.168.255.22

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 : vlan1033+1033, VLAN : 1033, VNID : 11033
MAC      MAC      Logical      Remote VTEP
address  flags  interface  IP address
f4:b5:2f:f3:fb:f0  DRP    vtep.32770  192.168.255.12

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:f7:12:2d  DRP    esi.800    192.168.255.22
f4:b5:2f:f3:fb:f0  DRP    vtep.32770  192.168.255.12

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
52:54:00:74:a0:6f  DRP    esi.801    192.168.255.22 192.168.255.21
f4:b5:2f:f3:fb:f0  DRP    vtep.32770  192.168.255.12
root@Core1>

```


From an EVPN-VLAN perspective everything is correct. This includes the fact that both Desktop addresses are present once again verifying Campus Fabric connectivity. Maybe we are looking in the wrong place. Let's look at the connection between core and WAN router.

External Campus Fabric Connectivity Through the Border GW Core EX9204 Switches

Remember that you chose to deploy the Border GW capability on the EX9204 switches during the Campus Fabric Core-Distribution deployment, represented below:



Mist enables the EX9204 to translate between VXLAN traffic within the Campus Fabric and standard Ethernet switching for external connectivity, in this case a SRX Series Firewalls. Let's verify the Ethernet Segment Identifier (ESI) status on the core switches.

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

We must configure the ESI-LAG and Mist does not configure this automatically. Add a Port profile on core switches interfaces facing the WAN router.

The following represents an existing Port Profile applied to each SRX Series Firewalls facing EX9204 port:

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

```

root@Core1> show lacp statistics interfaces
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
xe-1/0/0              13964        13962         0                0

root@Core1> show configuration interfaces ae0 | display set | display inheritance
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

root@Core1> show evpn database
Instance: evpn_vrf
VLAN  DomainId  MAC address      Active source      Timestamp      IP address
10001  d8:53:9a:64:6f:c0  192.168.255.21    Dec 08 17:49:32
10001  d8:53:9a:64:b5:c0  192.168.255.22    Dec 08 17:49:33
10001  f4:b5:2f:f3:fb:f0  192.168.255.12    Dec 08 17:49:32
10001  f4:b5:2f:f4:0b:f0  irb.0             Dec 08 17:49:32
11033  00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:19:00 Dec 08 17:49:32  10.33.33.1
11033  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 20:29:11  10.33.33.254
11033  f4:b5:2f:f3:fb:f0  192.168.255.12    Dec 08 17:49:32  10.33.33.3
11033  f4:b5:2f:f4:0b:f0  irb.1033          Dec 08 17:49:32  10.33.33.2
11088  00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:50:00 Dec 08 17:49:32  10.88.88.1
11088  52:54:00:f7:12:2d  00:11:00:00:00:01:00:01:03:00 Dec 08 20:30:22  10.88.88.88
11088  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 20:29:08  10.88.88.254
11088  f4:b5:2f:f3:fb:f0  192.168.255.12    Dec 08 17:49:32  10.88.88.3
11088  f4:b5:2f:f4:0b:f0  irb.1088          Dec 08 17:49:32  10.88.88.2
11099  00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:5b:00 Dec 08 17:49:32  10.99.99.1
11099  52:54:00:74:a0:6f  00:11:00:00:00:01:00:01:03:01 Dec 08 19:08:57  10.99.99.99
11099  f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00 Dec 08 20:29:11  10.99.99.254
11099  f4:b5:2f:f3:fb:f0  192.168.255.12    Dec 08 17:49:32  10.99.99.3
11099  f4:b5:2f:f4:0b:f0  irb.1099          Dec 08 17:49:32  10.99.99.2

root@Core1>

```

Note that LACP is up, and this infers there is an existing configuration on the SRX Series Firewalls.

```

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

root@Core1> show lacp interfaces
Aggregated interface: ae0
LACP state:      Role      Exp      Def      Dist      Col      Syn      Aggr      Timeout      Activity
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:      Receive State      Transmit State      Mux State
xe-1/0/0              Current      Fast periodic      Collecting distributing

root@Core1>

```

Then confirm the EVPN database now has the ESI entry. Back to Desktop1 to see if it can cross the fabric.

```

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.

```
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:~#
```

Meaning: Connectivity within the Campus Fabric and externally have been verified. Desktops communicate with each other through the Campus Fabric, each in an isolated VRF, then forwarded to the SRX Series Firewalls through the dual homing ESI-LAG on both Core1/2 for routing between VRFs or routing instances. Internet connectivity was also verified from each Desktop.

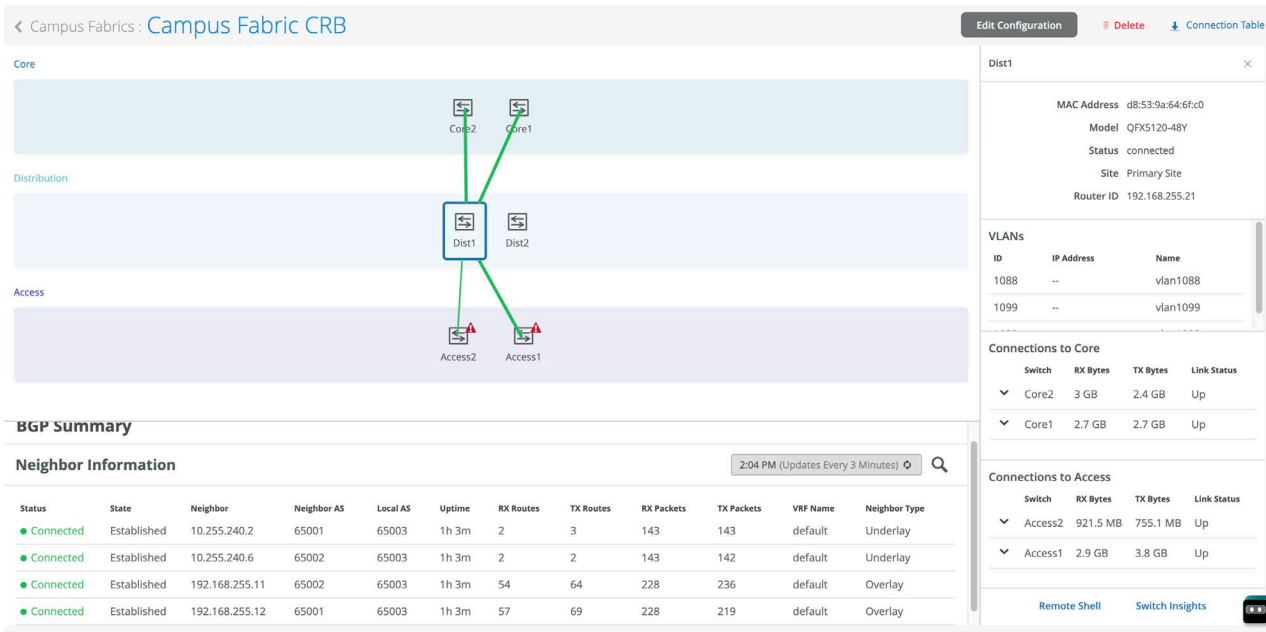
EVPN Insights

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

The screenshot displays the Mist Portal interface for the Campus Fabric Core Distribution CRB build. The topology diagram shows a Core layer with Core1 and Core2, a Distribution layer with Dist1 and Dist2, and an Access layer with Access1 and Access2. Core1 is highlighted with a blue box. The BGP Summary section shows the Neighbor Information table with the following data:

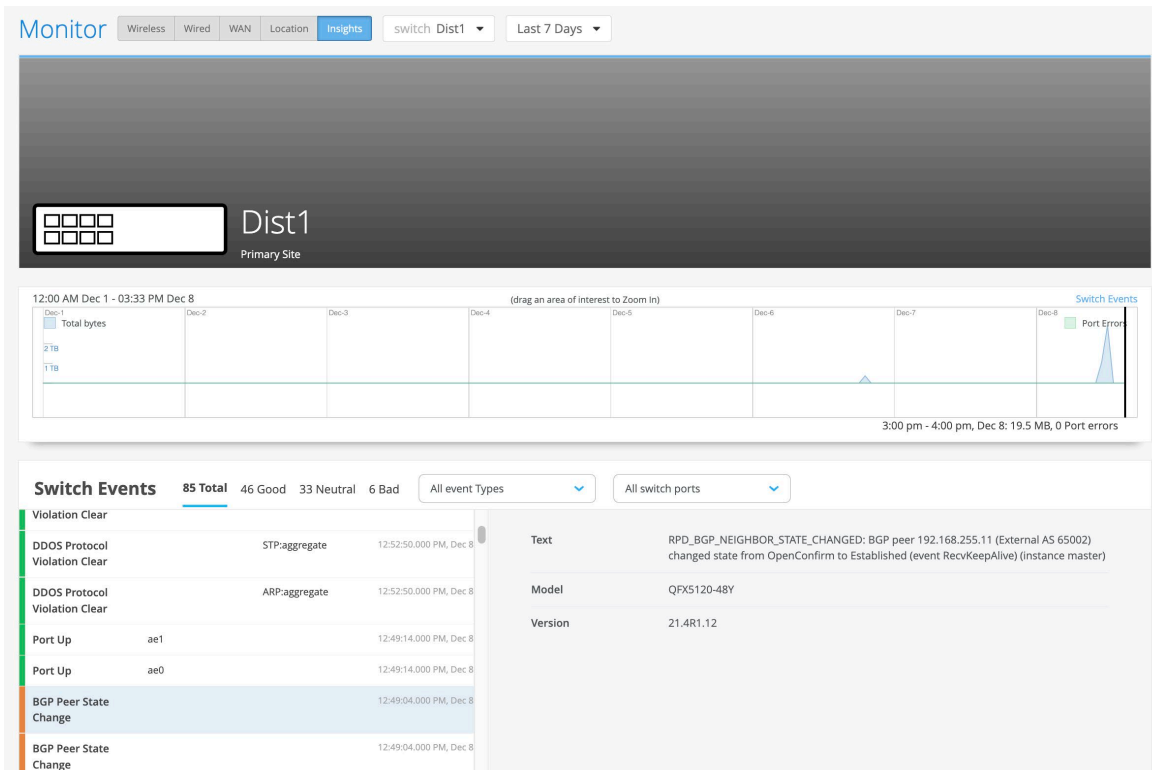
Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
Connected	Established	10.255.240.9	65004	65002	1h 1m	2	3	140	140	default	Underlay
Connected	Established	10.255.240.7	65003	65002	1h 1m	2	2	140	138	default	Underlay
Connected	Established	192.168.255.22	65004	65002	1h 1m	62	93	227	228	default	Overlay
Connected	Established	192.168.255.21	65003	65002	1h 1m	66	56	233	223	default	Overlay

The right sidebar shows details for Core1, including MAC Address, Model, Status, Site, Router ID, and a table of VLANs. The Connections to Distribution section shows the status of connections to Dist1 and Dist2.



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:



Summary

Mist Campus Fabric provides an easy method to build a Core Distribution CRB to enable EVPN-VXLAN overlay networks. This can be done only via Mist UI. Steps are added to this document to help you understand the troubleshooting steps if deployment isn't working correctly.

Additional Information

Campus Fabric Core Distribution CRB 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 Prefix (/24)
- 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.

```
set groups top interfaces lo0 unit 0 family inet address 192.168.255.11/32
set groups top routing-options router-id 192.168.255.11
```

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.

```
set groups top interfaces lo0 unit 0 family inet address 192.168.255.12/32
set groups top routing-options router-id 192.168.255.12
```

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 192.168.255.21/32
set groups top routing-options router-id 192.168.255.21
```

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 core switches and two access 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 192.168.255.22/32
set groups top routing-options router-id 192.168.255.22

```

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 core switches and two access 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 core 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 core 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 192.168.255.11
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 192.168.255.21 peer-as 65003
set protocols bgp group evpn_overlay neighbor 192.168.255.22 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 192.168.255.11:1
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. VRFs that are used for traffic isolation.

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi routing-options static route 0.0.0.0/0 next-hop 10.33.33.254
set groups top routing-instances guest-wifi routing-options multipath
set groups top routing-instances guest-wifi routing-options auto-export
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 guest-wifi interface irb.1033
set groups top routing-instances guest-wifi route-distinguisher 192.168.255.11:103
set groups top routing-instances guest-wifi vrf-target target:65000:103
set groups top routing-instances guest-wifi vrf-table-label
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers routing-options static route 0.0.0.0/0 next-hop 10.88.88.254
set groups top routing-instances developers routing-options multipath
set groups top routing-instances developers routing-options auto-export
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 developers interface irb.1088
set groups top routing-instances developers route-distinguisher 192.168.255.11:102
set groups top routing-instances developers vrf-target target:65000:102
set groups top routing-instances developers vrf-table-label
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it routing-options static route 0.0.0.0/0 next-hop 10.99.99.254
set groups top routing-instances corp-it routing-options multipath
set groups top routing-instances corp-it routing-options auto-export
```



```

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
set groups top routing-instances corp-it interface irb.1099
set groups top routing-instances corp-it route-distinguisher 192.168.255.11:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label

```

5. VLAN to VXLAN mapping.

```

set groups top routing-instances evpn_vs vlans vlan1033 vlan-id 1033
set groups top routing-instances evpn_vs vlans vlan1033 l3-interface irb.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 l3-interface irb.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 l3-interface irb.1099
set groups top routing-instances evpn_vs vlans vlan1099 vxlan vni 11099

```

6. L3 IRB interface enablement with virtual gateway addressing.

```

set interfaces irb unit 1033 proxy-macip-advertisement
set interfaces irb unit 1033 virtual-gateway-accept-data
set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet mtu 9000
set interfaces irb unit 1033 family inet address 10.33.33.2/24 virtual-gateway-address
10.33.33.1
set interfaces irb unit 1088 proxy-macip-advertisement
set interfaces irb unit 1088 virtual-gateway-accept-data
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet mtu 9000
set interfaces irb unit 1088 family inet address 10.88.88.2/24 virtual-gateway-address
10.88.88.1
set interfaces irb unit 1099 proxy-macip-advertisement
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.2/24 virtual-gateway-address
10.99.99.1

```

Core2 Configuration

1. BGP Overlay peering between the two distribution switches.

```

set protocols bgp group evpn_overlay type 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 192.168.255.12
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 192.168.255.21 peer-as 65003
set protocols bgp group evpn_overlay neighbor 192.168.255.22 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 192.168.255.12:1
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. VRFs that are used for traffic isolation.

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi routing-options static route 0.0.0.0/0 next-hop 10.33.33.254
set groups top routing-instances guest-wifi routing-options multipath
set groups top routing-instances guest-wifi routing-options auto-export
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 guest-wifi interface irb.1033
set groups top routing-instances guest-wifi route-distinguisher 192.168.255.12:103
set groups top routing-instances guest-wifi vrf-target target:65000:103
set groups top routing-instances guest-wifi vrf-table-label
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers routing-options static route 0.0.0.0/0 next-hop 10.88.88.254
set groups top routing-instances developers routing-options multipath
set groups top routing-instances developers routing-options auto-export
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 developers interface irb.1088
set groups top routing-instances developers route-distinguisher 192.168.255.12:102
set groups top routing-instances developers vrf-target target:65000:102
set groups top routing-instances developers vrf-table-label
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it routing-options static route 0.0.0.0/0 next-hop 10.99.99.254
set groups top routing-instances corp-it routing-options multipath
set groups top routing-instances corp-it routing-options auto-export
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
set groups top routing-instances corp-it interface irb.1099
set groups top routing-instances corp-it route-distinguisher 192.168.255.12:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
```

5. VLAN to VXLAN mapping.

```
set 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
set groups top routing-instances evpn_vs vlans vlan33 vlan-id 33
set groups top routing-instances evpn_vs vlans vlan33 l3-interface irb.33
set groups top routing-instances evpn_vs vlans vlan33 vxlan vni 10033
set groups top routing-instances evpn_vs vlans vlan99 vlan-id 99
set groups top routing-instances evpn_vs vlans vlan99 l3-interface irb.99
set groups top routing-instances evpn_vs vlans vlan99 vxlan vni 10099
```

6. L3 IRB interface enablement with virtual gateway addressing.

```
set interfaces irb unit 1033 proxy-macip-advertisement
set interfaces irb unit 1033 virtual-gateway-accept-data
set interfaces irb unit 1033 description vlan1033
set interfaces irb unit 1033 family inet mtu 9000
set interfaces irb unit 1033 family inet address 10.33.33.2/24 virtual-gateway-address 10.33.33.1
set interfaces irb unit 1088 proxy-macip-advertisement
set interfaces irb unit 1088 virtual-gateway-accept-data
set interfaces irb unit 1088 description vlan1088
set interfaces irb unit 1088 family inet mtu 9000
set interfaces irb unit 1088 family inet address 10.88.88.2/24 virtual-gateway-address 10.88.88.1
set interfaces irb unit 1099 proxy-macip-advertisement
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.2/24 virtual-gateway-address 10.99.99.1
```

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 192.168.255.21
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 192.168.255.12 peer-as 65001
set protocols bgp group evpn_overlay neighbor 192.168.255.11 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 192.168.255.21:1
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. VLAN to VXLAN mapping.

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

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 multihop ttl 1
set protocols bgp group evpn_overlay multihop no-nexthop-change
set protocols bgp group evpn_overlay local-address 192.168.255.22
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 65004
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 192.168.255.12 peer-as 65001
set protocols bgp group evpn_overlay neighbor 192.168.255.11 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 192.168.255.22:1
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. VLAN to VXLAN mapping.

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

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 LAG between the distribution switches and access switches. This Mist profile enables all VLANs on the Ethernet bundle with requisite ESI and LACP configuration options. From the perspective of the access switches, the Ethernet bundle that is configured on the access layer views the ESI-LAG as a single MAC address with the same LACP system-id. This enables load hashing between distribution and access layers without requiring L2 loop free detection protocols such as RSTP.



Dist1 Configuration

1. BGP Overlay peering between the two core switches.

```
set interfaces ae1 apply-groups crb-lag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:03:01
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1

set groups crb-lag interfaces <*> mtu 9100
set groups crb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033

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 ae1
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 ae0
```

Dist2 Configuration

1. BGP Overlay peering between the two core switches.

```
set interfaces ae1 apply-groups crb-lag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:03:01
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1

set groups crb-lag interfaces <*> mtu 9100
set groups crb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033

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 ae1
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 ae0
```

Access Configuration

1. VLANs associated with the new LACP Ethernet bundle.


```
set groups crb-lag interfaces <*> mtu 9100
set groups crb-lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set groups crb-lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033

set interfaces ae1 apply-groups crb-lag
set interfaces ae1 aggregated-ether-options lacp active
```

```
set interfaces ge-0/0/36 ether-options 802.3ad ae1
set interfaces ge-0/0/37 ether-options 802.3ad ae1
```

Configuration of the Layer 2 ESI-LAG Between the Core Switches and SRX Series Firewalls

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



Core 1 Configuration

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

```
set interfaces 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 ae1
set interfaces xe-1/0/0 unit 0 family ethernet-switching storm-control default

set interfaces ae1 apply-groups esilag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:02:01
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lACP active
set interfaces ae1 aggregated-ether-options lACP periodic fast
set interfaces ae1 aggregated-ether-options lACP system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lACP admin-key 1
```

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 ae1
set interfaces xe-1/0/1 unit 0 family ethernet-switching storm-control default

set interfaces ae1 apply-groups esilag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:02:01
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lACP active
set interfaces ae1 aggregated-ether-options lACP periodic fast
set interfaces ae1 aggregated-ether-options lACP system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lACP admin-key 1
```

SRX Series Firewalls Configuration

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

```
set interfaces ae0 apply-groups lan
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 mtu 9014
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 1033 description vlan1033
set interfaces ae0 unit 1033 vlan-id 1033
set interfaces ae0 unit 1033 family inet address 10.33.33.254/24
set interfaces ae0 unit 1088 description vlan1088
set interfaces ae0 unit 1088 vlan-id 1088
set interfaces ae0 unit 1088 family inet address 10.88.88.254/24
set interfaces ae0 unit 1099 description vlan1099
```

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 © 2023 Juniper Networks, Inc. All rights reserved.