



Juniper Validated Design (JVD)

Campus Fabric IP Clos 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 IP Clos 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 IP Clos	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>7</i>
<i>Services Block</i>	<i>8</i>
<i>Access Layer</i>	<i>9</i>
<i>Juniper Access Points</i>	<i>10</i>
<i>Campus Fabric IP Clos Deployment Types</i>	<i>11</i>
<i>3 Stage Clos</i>	<i>11</i>
<i>5 Stage Clos</i>	<i>11</i>
<i>Supported Platforms for Campus Fabric IP Clos</i>	<i>12</i>
Campus Fabric IP Clos Unicast Scale	13
Juniper Mist Wired Assurance	13
Campus Fabric IP Clos High-Level Architecture	14
Campus Fabric IP Clos Components	14

Juniper Mist Wired Assurance	15
Juniper Mist Wired Assurance Switches	16
<i>Overview</i>	16
<i>Templates</i>	17
<i>Topology</i>	17
Create the Campus Fabric	18
<i>Campus Fabric Org Build</i>	18
<i>Campus Fabric Site Build</i>	19
<i>Choose the Campus Fabric Topology</i>	19
<i>Topology Settings</i>	20
<i>Select Campus Fabric Nodes</i>	20
<i>Configure Networks</i>	22
<i>VRF</i>	22
<i>Networks</i>	23
<i>Other IP Configuration</i>	24
<i>Configure Campus Fabric Ports</i>	27
<i>Core Switches</i>	27
<i>Distribution Switches</i>	28
<i>Access Switches</i>	30
<i>Campus Fabric Configuration Confirmation</i>	30
Verification	35
<i>BGP Underlay</i>	36
<i>Purpose</i>	36
<i>Action</i>	37
<i>Verification of BGP Peering</i>	37
<i>EVPN VXLAN Verification Between Access and Core Switches</i>	39
<i>Verification of the EVPN Database on Both Access Switches</i>	39
<i>Verification of VXLAN Tunnelling Between Access and Core Devices</i>	40
<i>External Campus Fabric Connectivity Through the Border Gateway Core EX9204 Switches</i>	43
EVPN Insights	45
Summary	48
Additional Information	48
<i>Configuration of the Underlay IP Fabric</i>	48
<i>Configuration of the EVPN VXLAN Overlay and Virtual Networks</i>	52
<i>Configuration of the Layer 2 ESI-LAG Between the Core Switches and SRX Series Firewalls</i>	57



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. IP Clos networks provide increased scalability and segmentation using a well-understood standards-based approach (EVPN-VXLAN with GBP).

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. Multi-Chassis Link Aggregation Group (MC-LAG) is a good example of a single-vendor technology that addresses the collapsed core deployment model. In this model, two chassis-based platforms are typically in the core of a customer's network and deployed to handle all Layer 2 and Layer 3 requirements while providing an active and backup resiliency environment. MC-LAG does not interoperate between vendors, creating lock-in, and is limited to two devices.

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 your network. 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. ESI-LAG uses standards-based Link Aggregation Control Protocol (LACP) to interconnect with legacy switches.

Benefits of Campus Fabric IP Clos

- 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 virtual LANs (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 mobile and IoT devices.
- Campus Fabrics have an underlay topology with a routing protocol that ensures loopback interface reachability between nodes. Devices participating in EVPN-VXLAN function as VXLAN tunnel endpoint (VTEP) that encapsulate and decapsulate the VXLAN traffic. VTEP represents construct within the switching platform that originates and terminates VXLAN tunnels. In addition, these devices route and bridge packets in and out of VXLAN tunnels as required.
- The Campus Fabric IP Clos extends the EVPN fabric to connect VLANs across multiple buildings or floors of a single building. This is done by stretching the Layer 2 VXLAN network with routing occurring in the access device instead of the core (Centrally-Routed Bridging (CRB)) or distribution (Edge Routed Bridging (ERB)) layers.

Campus Fabric Core IP Clos

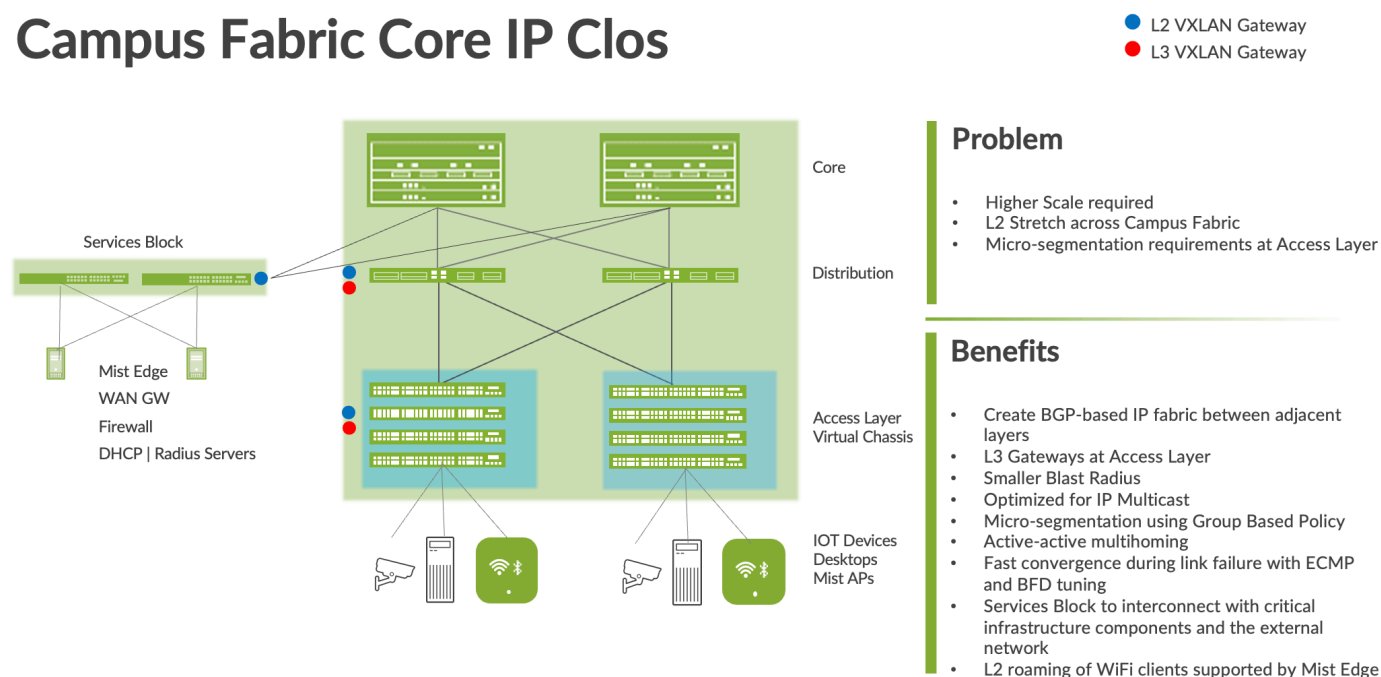


Figure 1: Campus Fabric IP Clos

IP Clos network encompasses the distribution, core, and access layers of your topology.

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

- **Reduced flooding and learning**—Control plane-based Layer 2 and Layer 3 learning reduces the flood and learn issues associated with data plane learning. Learning MAC addresses in the forwarding plane has an adverse impact on network performance as the number of endpoints grows. This is because more management traffic consumes the bandwidth which leaves less bandwidth available for production traffic. The EVPN control plane handles the exchange and learning of MAC addresses through eBGP routing, rather than a Layer-2 forwarding plane.
- **Scalability**—More efficient control-plane based Layer 2 and Layer 3 learning. For example, in a Campus Fabric IP Clos, core switches only learn the access layer switches addresses instead of the device endpoint addresses.
- **Consistency**—A universal EVPN-VXLAN-based architecture across disparate campus and data center deployments enables a seamless end-to-end network for endpoints and applications.
- **Group-based policies**—With group-based policy (GBP), you can enable microsegmentation with EVPN-VXLAN to provide traffic isolation within and between broadcast domains as well as simplify security policies across a Campus Fabric.
- **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, distribution, and access 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 access, core, and distribution devices. BGP provides benefits such as better prefix filtering, traffic engineering, and route tagging. We are using 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/31 Links Between Adjacent Layers Running eBGP

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 access switches 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 tunnels. These methods are signaled with a multicast underlay and do not solve the inherent flood and learn problem and are difficult to scale in large multitenant environments. These methods are not in the scope of this documentation.

Understanding EVPN

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

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

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

The benefits of using EVPNs include:

- MAC address mobility
- Multitenancy
- Load balancing across multiple links
- Fast convergence

- High availability
- Scale
- Standards based interoperability

EVPN provides multipath forwarding and redundancy through an all-active model. The access layer can connect to two or more distribution devices and forward traffic using all the links. If an access link or distribution device fails, traffic flows from the access layer toward the distribution layer using the remaining active links. For traffic in the other direction, remote distribution 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 access 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 access layer of a Campus Fabric. The connection off the multihomed access layer switches is called ESI-LAG, while the access layer devices connect to each access 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 VTEP. Before a VTEP sends a frame into a VXLAN tunnel, it wraps the original frame in a VXLAN header that includes a Virtual Network Identifier (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.

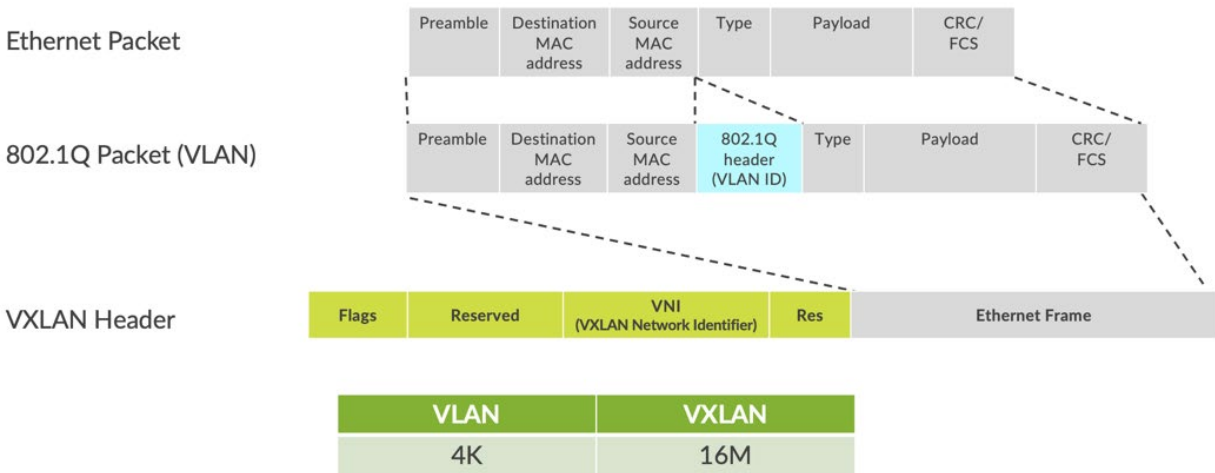


Figure 4: VXLAN Header

VTEPs are software entities tied to the devices' loopback address that source and terminate VXLAN tunnels. VXLAN tunnels in an IP Clos fabric are provisioned on the following:

- Access switches to extend services across the Campus Fabric IP Clos.
- 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 layer switches set up eBGP peers using their loopback addresses using next hops announced by the underlay BGP sessions. For example, core and distribution devices establish eBGP sessions between each other while the access and distribution devices establish eBGP sessions with each other. When there is a Layer 2 forwarding table update on any switch participating in Campus Fabric, it sends a BGP update message with the new MAC route to other devices in the fabric. Those devices then update their local EVPN database and routing tables.

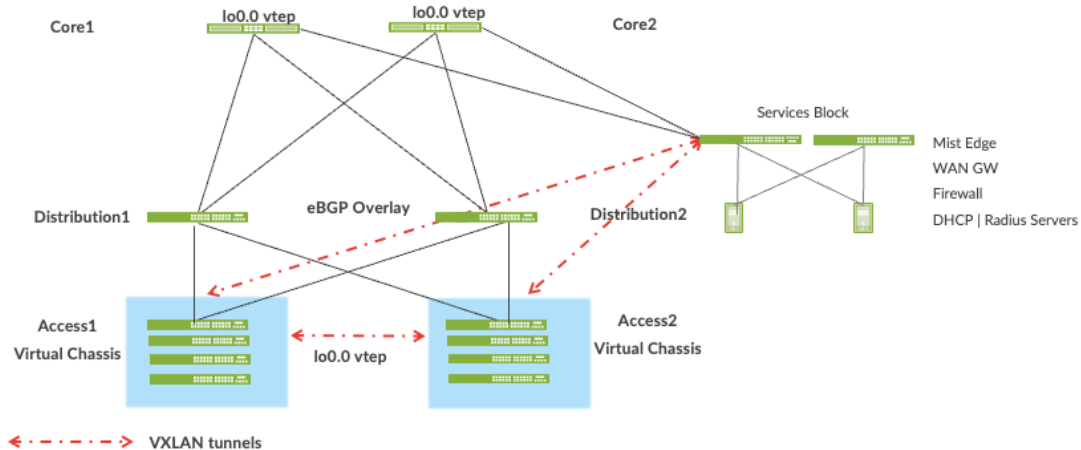


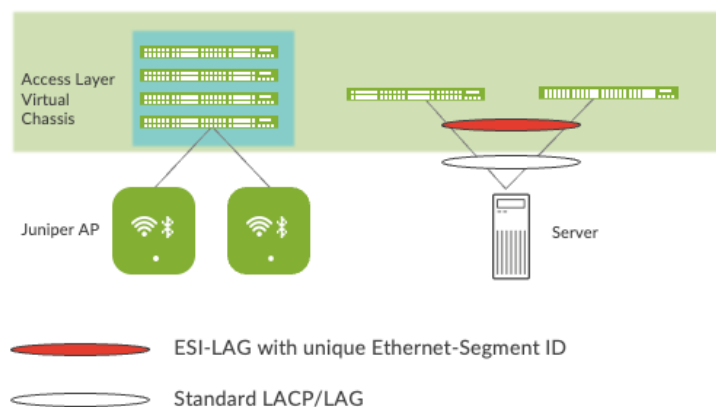
Figure 5: EVPN VXLAN Overlay Network with a Services Block

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 links within the Campus Fabric using Equal Cost Multi Pathing (ECMP) enabled at the forwarding plane.

Ethernet Segment Identifier (ESI)

When devices such as servers and access points are multihomed to two or more switches at the access layer in a Campus Fabric, an ESI-LAG is formed on the access layer devices. This ESI is a 10-octet integer that identifies the Ethernet segment amongst all access 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.

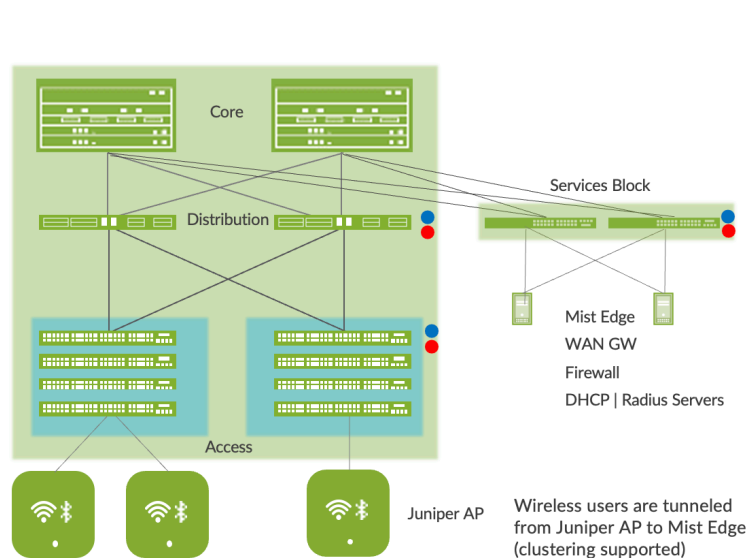


- EVPN supports N-way “scale-out” Ethernet multihoming
- No ICL link required between Access Switches
- Virtual Chassis LAG spread across multiple switches in VC stack
- Active-Active Multihoming
- Multi-homed devices such as Servers are identified in the overlay by unique Ethernet Segment ID (ESI)

Figure 6: Device Resiliency and Load Balancing

Services Block

You 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. If you need 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.



Problem

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

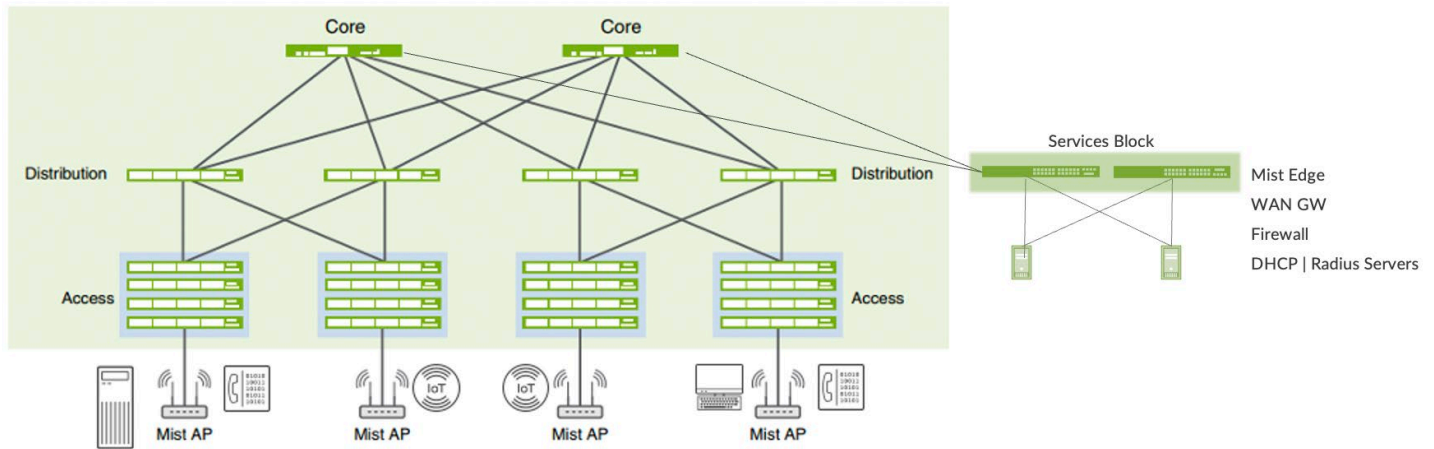
Benefits

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

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. EVPN-VXLAN network extends all the access layer switches.



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

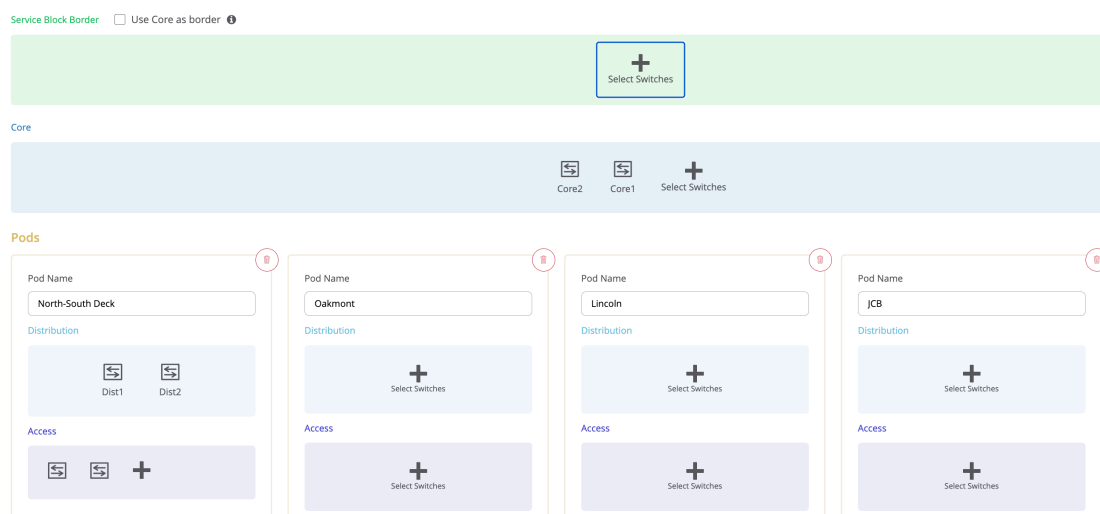
Figure 8: End Point Access

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. See <https://www.juniper.net/documentation/us/en/software/junos/vcf-best-practices-guide/vcf-best-practices-guide.pdf>.

With EVPN running as the control plane protocol, any access switch or Virtual Chassis device can enable active-active multihoming to the distribution layer. EVPN provides a standards-based multihoming solution that scales horizontally across any number of access layer switches. For more information, see [Campus Fabric IP Clos Unicast Scale](#).

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](#).

Campus Fabric IP Clos Deployment Types

Juniper's Wired Assurance supports 3 Stage and 5 Stage IP Clos deployments. The 3 Stage IP Clos is targeted for deployments that do not require a distribution layer and have smaller scale requirements. This also allows for cost effective EX4400, EX4650, QFX5110, and QFX5120 switching platforms to be deployed at the core layer.

3 Stage Clos

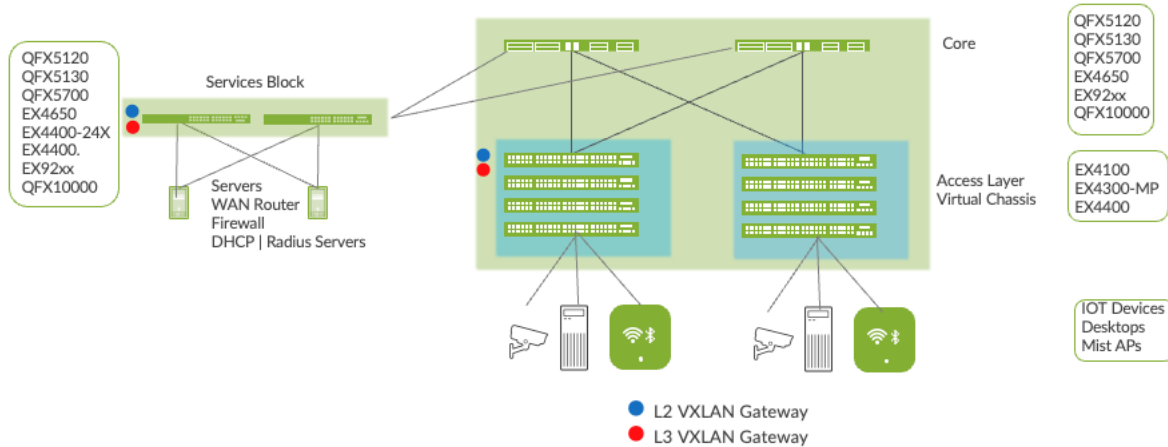


Figure 9: Campus Fabric 3 Stage Clos

5 Stage Clos

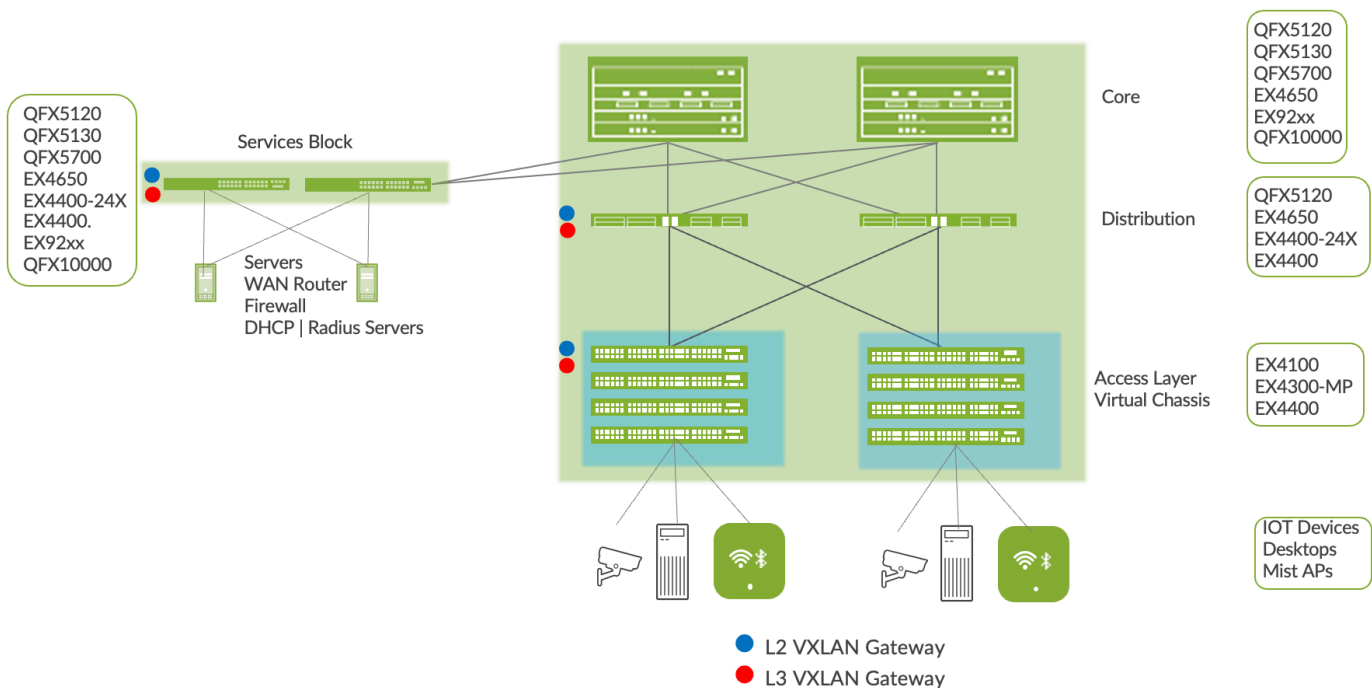


Figure 10: Campus Fabric 5 Stage Clos

NOTE: You can deploy the Services Block in Stage 3 and Stage 5 IP Clos architectures.

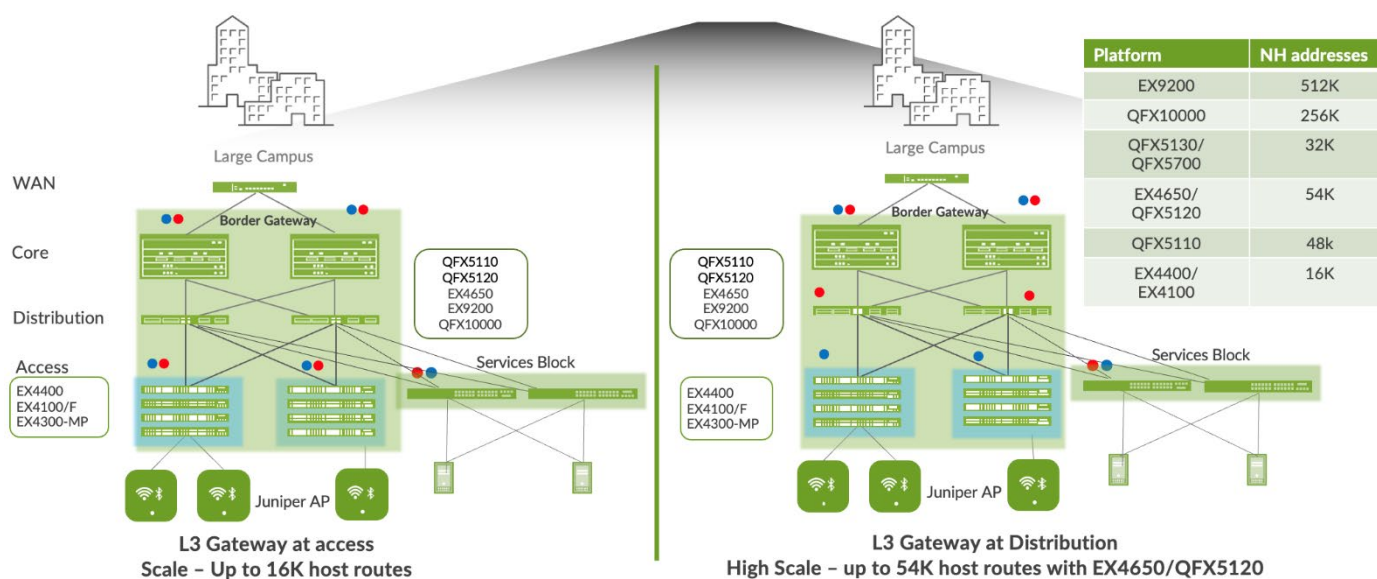
Supported Platforms for Campus Fabric IP Clos

Table 1 lists the supported platforms for Campus Fabric IP Clos deployment.

Table 1: Supported Platforms for Campus Fabric IP Clos Deployment

Campus Fabric IP Clos Deployment	Supported Platforms
Access layer	<ul style="list-style-type: none"> EX4100 EX4300-MP 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
Services block	<ul style="list-style-type: none"> EX4400/EX4400-24X EX4650 QFX5110 QFX5120 QFX5130 QFX5700 QFX10000 EX92xx

Campus Fabric IP Clos Unicast Scale



Juniper Mist Wired Assurance

Juniper Mist Wired Assurance is a cloud service that brings automated operations and service levels to the Campus Fabric for switches, IoT devices, access points, servers, and printers. It is about simplifying every step of the way, starting from Day 0 for seamless onboarding and auto-provisioning through Day 2 and beyond for operations and management. Juniper EX Series Switches provide Junos streaming telemetry that enable the insights for switch health metrics and anomaly detection, as well as Juniper Mist AI capabilities.

Mist's AI engine and virtual network assistant, Marvis, further simplifies troubleshooting while streamlining helpdesk operations by monitoring events and recommending actions. Marvis is one step towards the Self-Driving Network, turning insights into actions and 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 IP Clos 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. In a Campus Fabric deployment, the use of EVPN VXLAN supports native traffic isolation using routing-instances; commonly called Virtual Routing and Forwarding (VRFs) for macrosegmentation purposes.


The Mist UI workflow makes it easy to create Campus Fabrics.

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 IP Clos 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 Clos is available under the EVPN/VXLAN ERB section at: https://supportportal.juniper.net/s/article/Junos-Software-Versions-Suggested-Releases-to-Consider-and-Evaluate?language=en_US

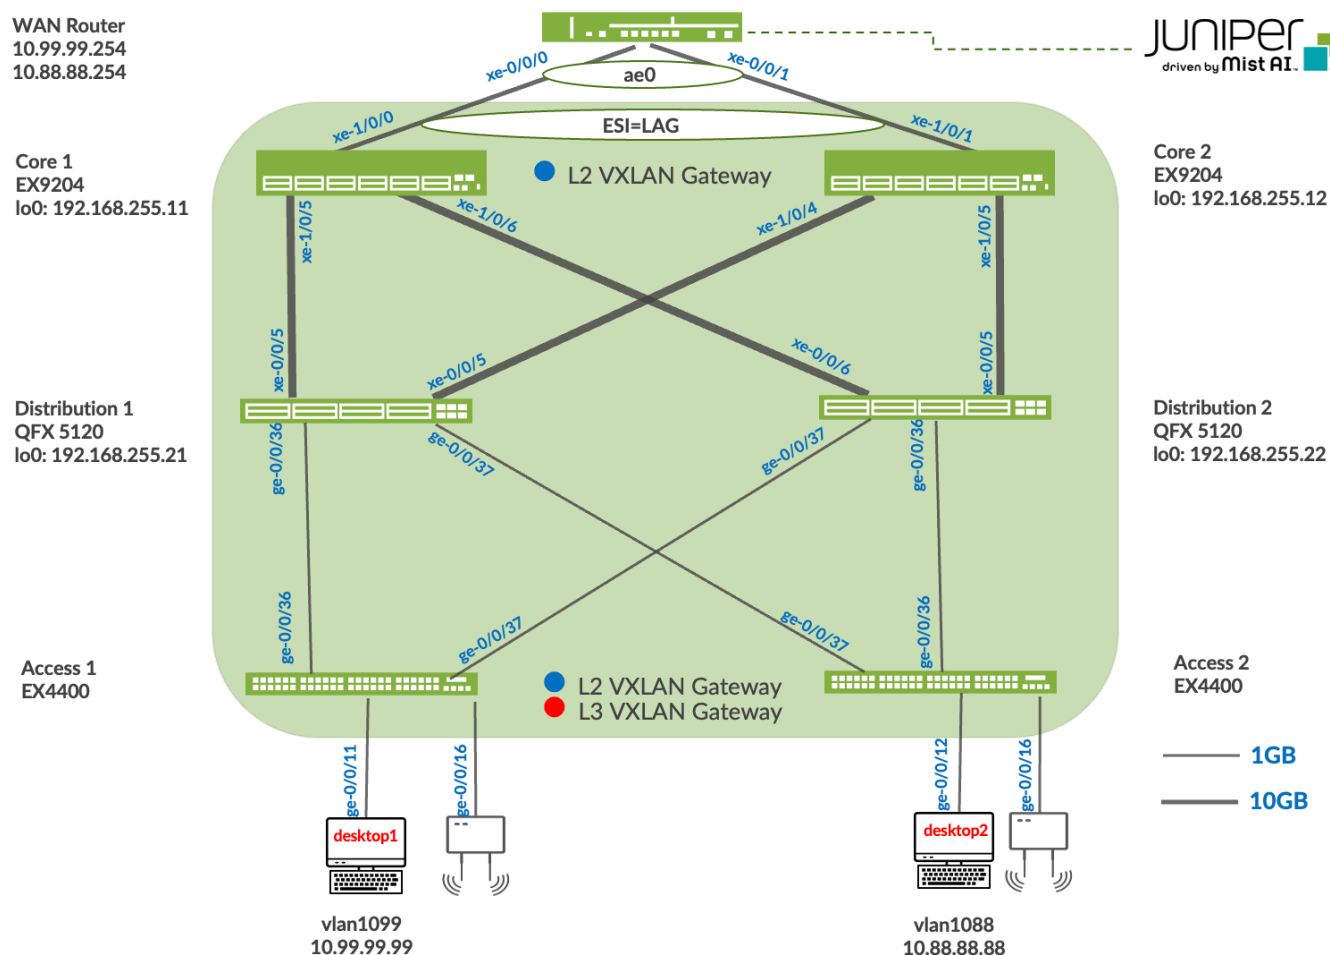


Figure 11: Topology

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

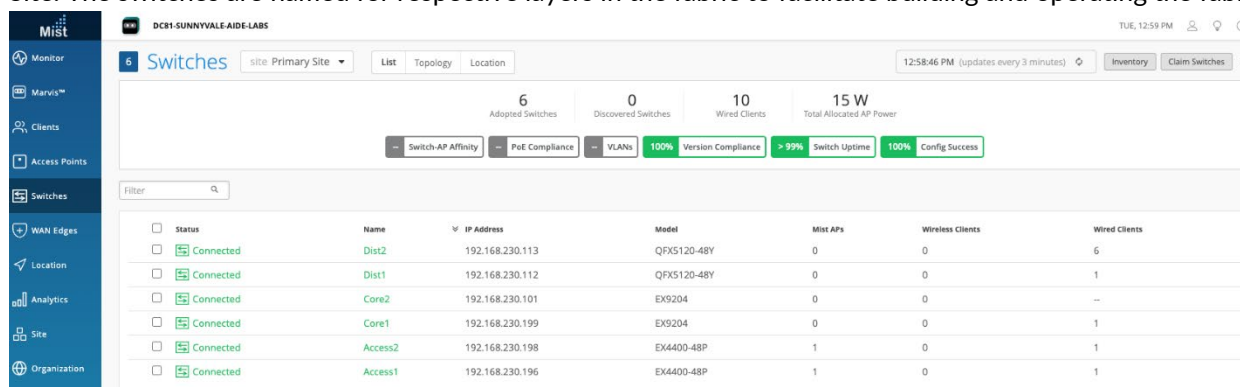
Wired Assurance, through the Mist UI, is used to build a Campus Fabric IP Clos from ground up. This includes the following:

- Assignment of p2p links between all layers of the Campus Fabric.
- 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 gateway Integrated Routing and Bridging (IRB) assigned to the access 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 12: Switch Inventory

Overview

Use this JVD to deploy a single Campus Fabric with a Layer 3 IP-based underlay network that uses EVPN as the control plane protocol and VXLAN as the data plane protocol in the overlay network.

Mist Wired Assurance configures eBGP on the directly connected links to exchange loopback routes, and eBGP between the core and distribution devices and distribution and access devices in the overlay to share reachability information about endpoints in the fabric.

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. 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

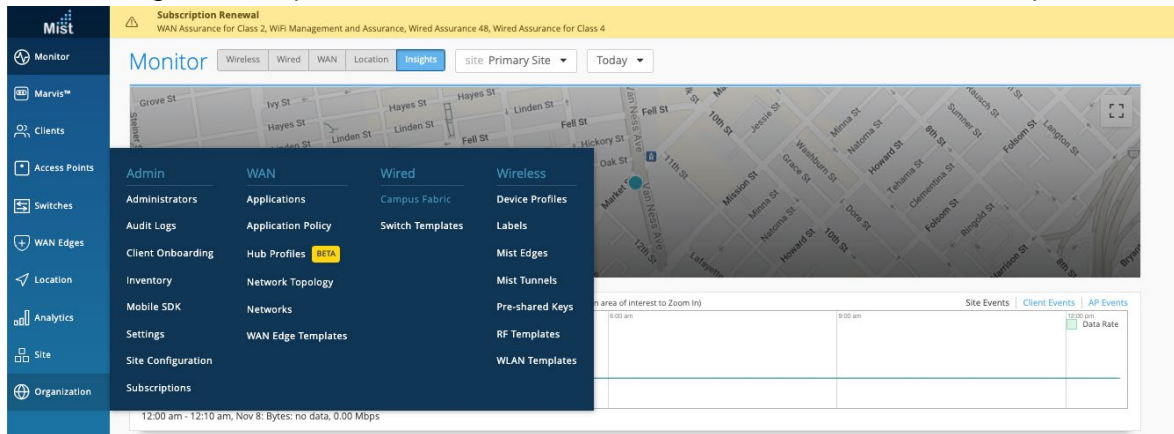
Topology

Wired Assurance provides the template for LAN and Loopback IP addressing for each 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 adjacent devices within the Campus Fabric IP Clos.

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.

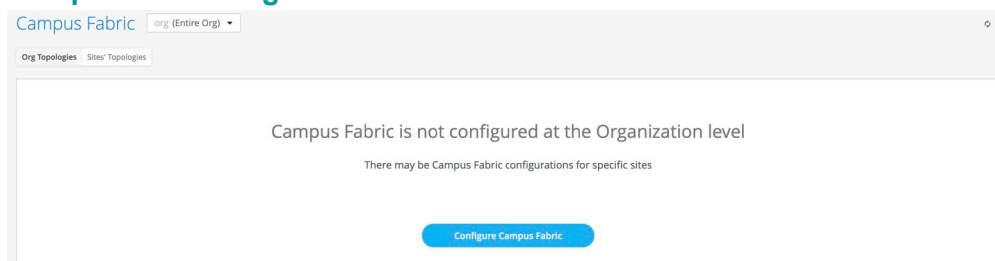
Create the Campus Fabric

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



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

Campus Fabric Org Build



Campus Fabric Site Build

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 IP Clos option below:

× Campus Fabric Configuration [1. Topology](#) [2. Nodes](#) [3. Network Settings](#) [4. Ports](#) [5. Confirm](#)

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


☐ Routed at Distribution
Centrally-routed and bridged with gateways on the Distribution


☒ Routed at Edge
Edge-routed and bridged with anycast gateways on the access

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 IP Clos and where you want to have L3 boundaries (where Default Gateway exists for each VLAN):

Configuration—Provide a name in accordance with company standards.

NOTE: Routed at edge and access layer provides a smaller blast radius for broadcast traffic and is ideal for east-west traffic patterns and IP Multicast environments. Routed at distribution aligns with north-south traffic patterns and configures the access layer as Layer 2 VXLAN gateway only. This deployment is preferred for higher scale deployments.

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. VXLAN 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.

TOPOLOGY SETTINGS

BGP Local AS

 (2-byte or 4-byte)

Loopback prefix ⓘ

Subnet ⓘ

 (xxx.xxx.xxx.xxx/xx)

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

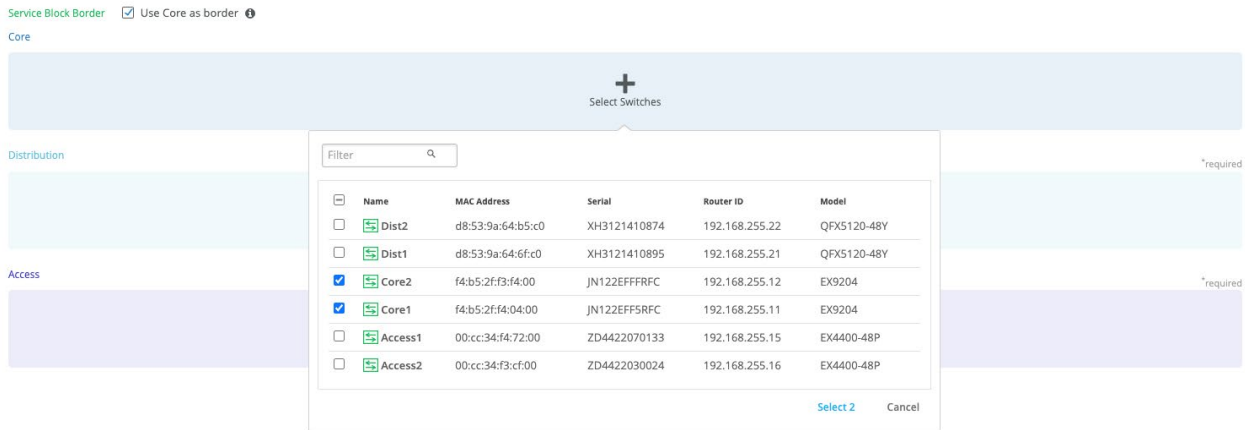
Select Campus Fabric Nodes

Select devices to participate at each layer of the Campus Fabric IP Clos. 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 are 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 IP Clos in a separate

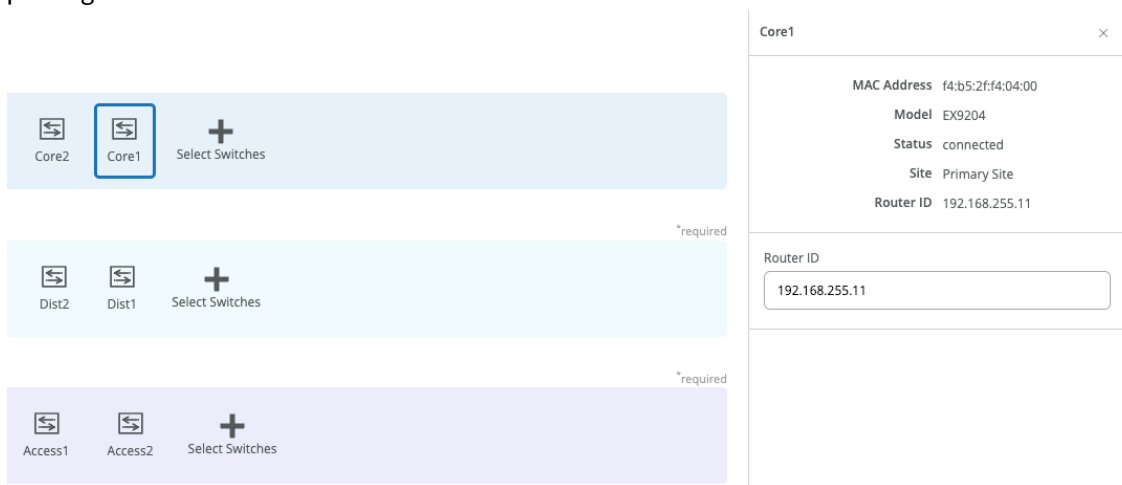
device or pair of devices, clear the Use Core as border option and select the Select Switches option to choose the devices.



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 to source the VXLAN Tunnel. Campus Fabric IP Clos has VTEPs for VXLAN tunnelling on the access 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.



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 a configuration interface with two input fields. The first field is labeled 'Loopback prefix' and contains the value '/24'. The second field is labeled 'Subnet' and contains the value '10.255.240.0/20'. Below the 'Subnet' field, there is a placeholder text '(xxx.xxx.xxx.xxx/xx)'.

Configure Networks

Enter 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.

VRF

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

For more information on Routing Instance Overview, see

<https://www.juniper.net/documentation/us/en/software/junos/routing-overview/topics/concept/routing-instances-overview.html>

VLANs can be placed into a common VRF. Here, all VLANs within each VRF have full connectivity to each other and other external networking resources. A common use case includes most enterprise domains isolating Guest WIFI traffic and save Internet connectivity.

By default, the Campus Fabric provides complete isolation between VRFs forcing inter-VRF communications to traverse a Firewall or security compliance. This aligns with most enterprise security use cases and compliance and is represented in this document.

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 Gateway per each VLANs.

The Shared Elements section of the campus-fabric template includes the Networks section mentioned above where VLANs are created. This can be found under the Organization/Switch Templates section, then choose the appropriate template:

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 Add Existing Network option that includes Layer 2 VLAN information. All VLAN and IP information is inherited from the template.

NETWORKS

Add Existing Network ✓ ✕

Available Networks

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

Import from Template

Template

DC81-IP-Clos: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

Search ✕

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

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 (IRB) for each of the VLANs. Then, Port Profiles and Port Configuration associate the VLAN with specified ports. In this case, we selected Campus Fabric IP Clos routed at Edge at the onset of the Campus Fabric build.

CONFIGURATION

Topology Name

Campus Fabric IPClos

Topology Sub-type

- ☐ Routed at Distribution
 Centrally-routed and bridged with gateways on the Distribution
- ☒ Routed at Edge
 Edge-routed and bridged with anycast gateways on the access

This option uses anycast addressing for all devices participating in the L3 subnet. In this case, Access1 and Access2 switches are configured with shared IP address for each L3 subnet.

For more information on Anycast Gateways, see

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-mclag-irb-gateway-anycast-address.html>

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

Edit Access2

✓

×

vlan1033	10.33.33.1	>
vlan1088	10.88.88.1	>
vlan1099	10.99.99.1	>

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

Edit Access1

✓

×

vlan1033	10.33.33.1	>
vlan1088	10.88.88.1	>
vlan1099	10.99.99.1	>

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

VRF

Configuration

☒ Enabled ☐ Disabled

Instances

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

[Add VRF Instance](#)

Now that all VLANs are configured and assigned to each VRF, 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 ?

FPC 1

FPC 2

1

SFP+

1 3 5 7

0 2 4 6

1 3 5 7

0 2 4 6

1 3 5 7

0 2 4 6

1 3 5 7

0 2 4 6

EX9200-32XS

Core1

EX9204

0/2 ?

Distribution Switches

QFX5120-48Y

Edit Ports for all QFX5120-48Y

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

Access Switches

EX4400-48P

Edit Ports for all EX4400-48P

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

NOTE: We recommend that you have the output of the show lldp neighbors command from each switch. If Juniper enable LLDP out of the box and provides additional LLDP attributes when the switch is added to a Campus Fabric. 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

Switch	Model	Link to Distribution
Core1	EX9204	0/2 ?
Core2	EX9204	0/2 ?

1

SFP+

1 3 5 7

0 2 4 6

1 3 5 7

0 2 4 6

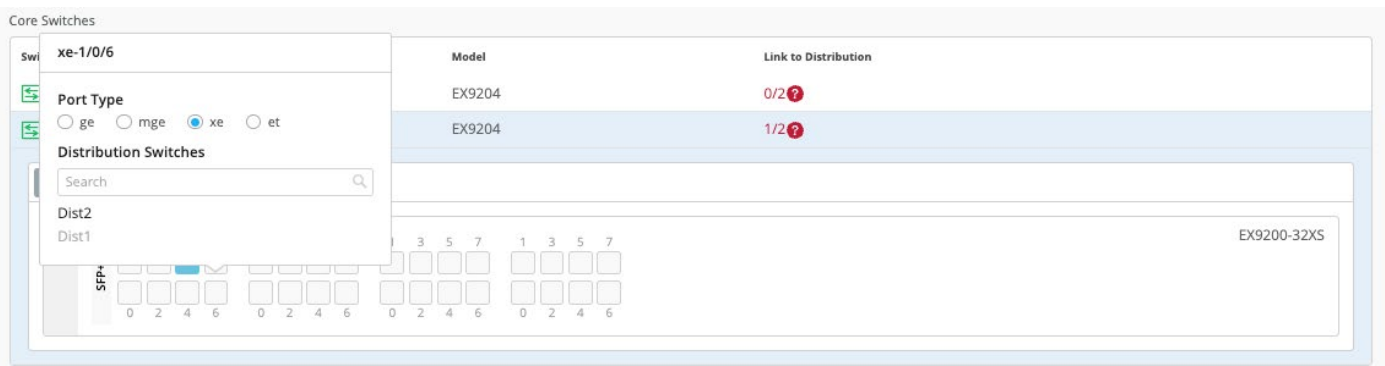
1 3 5 7

0 2 4 6

1 3 5 7

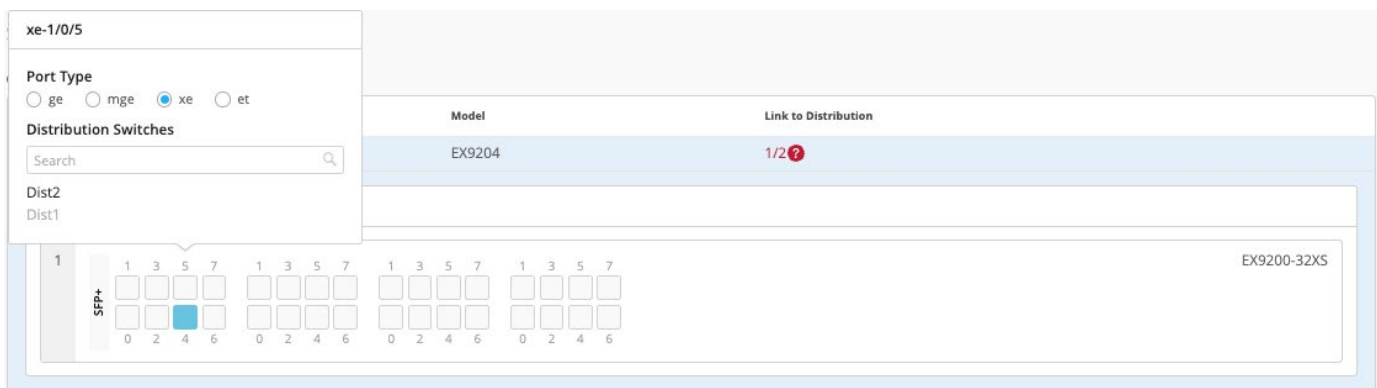
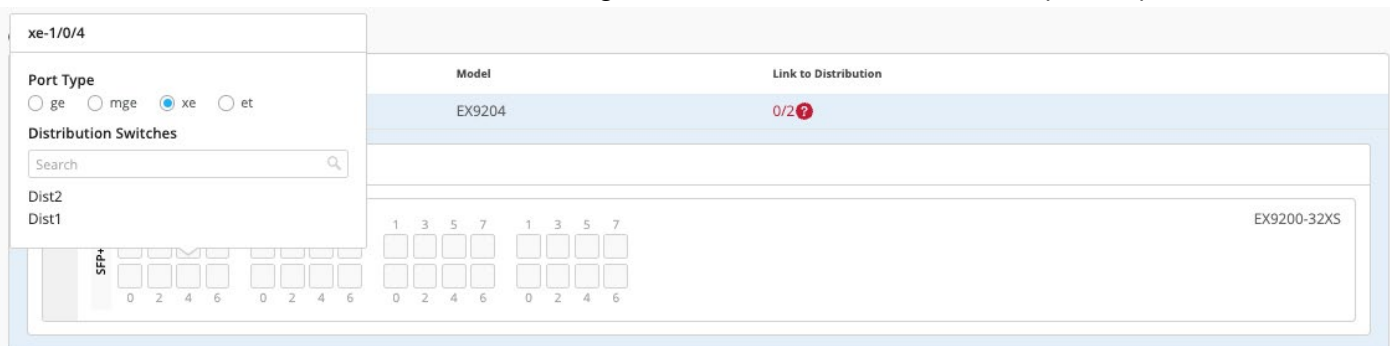
0 2 4 6

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 ?

16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48 50 52 54
 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55

SFP28

QSF28 Uplink

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 ?

0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48 50 52 54
 1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55

SFP28

QSF28 Uplink

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 ?

0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48 50 52 54
 1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55

SFP28

QSF28 Up

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 ?

0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48 50 52 54
 1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55

SFP28

QSF28 Up

Edit Ports for all QFX5120-48Y

Next, select the following interconnects of **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

NOTE: QFX 5120-48Y Switch is an example switch that is targeted for the distribution layer in a Campus Fabric. The device supports blocks of four ports per PHY; Ports0-3, 4-7, and so on. All ports within the same PHY must operate at the same speed.

Access Switches

Finally, select the following interface combinations for Access1 and Access2:

Access1:

- ge-0/0/36 – Distribution Switch – Dist1
- ge-0/0/37 – Distribution Switch – Dist2

Access2:

- ge-0/0/36 – Distribution Switch – Dist1
- ge-0/0/37 – Distribution Switch – Dist2

Once you have completed selecting all requisite port combinations, click Continue 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

Core1

MAC Address: fa:35:2f:f4:04:00
 Model: EX3204
 Status: connected
 Site: Primary Site
 Router ID: 192.168.255.11

VLANs

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

Connections to Distribution

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

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 all 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 assigned to the access layer.
- 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 Virtual Network Identifier (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 access devices and access-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 IP Clos.

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 IPClos	9acf2078-c2cc-40e5-a701-58954d8711b9	Primary Site	Campus Fabric IP Clos	Access	02:36:47 PM, Mar 15 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 IPClos** Edit Configuration Delete Connection Table

BGP Summary

Neighbor Information 2:43 PM (Updates Every 3 Minutes) Q

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

Core1

MAC Address f4b5:2ff4:04:00
Model EX9204
Status connected
Site Primary Site
Router ID 192.168.255.11

VLANs

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

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Status
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 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	JN122EFFFRFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/6	< --- >	xe-1/0/6		downlink	Primary Site	JN122EFFFRFC	EX9204	f4b52ff3f400	Core1	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	downlink		ge-0/0/36	< --- >	ge-0/0/36		uplink	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	downlink		ge-0/0/37	< --- >	ge-0/0/37		uplink	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	JN122EFFFRFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	uplink		xe-0/0/5	< --- >	xe-1/0/5		downlink	Primary Site	JN122EFFFRFC	EX9204	f4b52ff3f400	Core1	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	downlink		ge-0/0/37	< --- >	ge-0/0/37		uplink	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	downlink		ge-0/0/36	< --- >	ge-0/0/36		uplink	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access

Apply VLANs to Access Ports

As previously discussed, Mist provides the ability to template well known services such as Radius, NTP, DNS, and others 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](#).

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 Organizational/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 Access Point (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. These VLANs are defined under the Organization/Switch template section. 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, Quality of Service (QoS), and Power over Ethernet (PoE). Vlan1088 and vlan1033 need to be configured in a similar fashion.

 Edit Port Profile  

Name

Port Enabled


☒ Enabled ☐ Disabled

Description


Mode

☐ Trunk ☒ Access

Port Network (Untagged/Native VLAN)


1099 

VoIP Network




☐ Use dot1x authentication

Speed



Duplex



Mac Limit

(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 IP Clos deployment. See [Figure 9](#).

Currently, there are two desktops to validate the Campus Fabric. Let's take a quick look to see if Desktop1 can connect internally and externally. A third-party tool such as SecureCRT can be used to validate each desktop's configuration with Desktop1 shown below:

```
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	Date Created
DC81-IPClo	1f4467cc-bfaf-4439-b91a-89a66d79d74c	Primary Site	05:46:13 PM, Nov 7 2022

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

< Campus Fabrics: Campus Fabric IPClos Edit Configuration Delete Connection Table

Core1

MAC Address f4b52f14:04:00
Model EX9204
Status connected
Site Primary Site
Router ID 192.168.255.11

VLANs

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

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Status
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

BGP Summary

Neighbor Information 2:43 PM (Updates Every 3 Minutes) Q

Status	State	Neighbor	Neighbor AS	Local AS	Uptime	RX Routes	TX Routes	RX Packets	TX Packets	VRF Name	Neighbor Type
Connected	Established	10.255.240.7	65003	65002	5m	5	2	21	17	default	Underlay
Connected	Established	192.168.255.21	65003	65002	5m	36	4	41	18	default	Overlay
Connected	Established	192.168.255.22	65004	65002	5m	36	40	42	39	default	Overlay

[Remote Shell](#) [Switch Insights](#) 9:48

```

Remote Shell - Core1
integration.mistsys.com/admin/shell.html?siteId=2c65f917-5fa1-4151-bed2-289a219f4c71&deviceId=00000000-00...
Warning: When a device is managed by Mist, the configuration changes made locally via shell
1 will be overwritten with the configuration from the cloud. Please use the UI to make any
config changes.

Last login: Wed Mar 15 17:51:08 2023 from 54.157.92.6
--- JUNOS 22.4R1.10 Kernel 64-bit JNPR-12.1-20221121.c470123_buil
{master}
mist@Core1>

```

BGP Underlay

Purpose

Verifying the state of eBGP between adjacent 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.
- BGP peering as well as loopback VXLAN reachability.

Without requiring verification at each layer, the focus can be on Dist1/2 and their eBGP relationships with Access1/2 and Core1/2. If both distribution switches have “established” eBGP peering sessions with each adjacent layer, you can move to the next phase of verification.

Action

Verify that BGP sessions are established from Dist1/2 with access and core devices to ensure loopback reachability, bfd session status, and load-balancing using ECMP.

NOTE: Operational data can be gathered through the Campus Fabric section of the Mist UI or using an external application such as SecureCRT or Putty.

Verification of BGP Peering

Dist1:

From Switch > Utilities, access Remote Shell via the bottom right of the Campus Fabric, from the switch view or via Secure Shell (SSH).

```
{master:0}
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: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State   Pending
inet.0
bgp.evpn.0
127
59
0
0
0
0
0
0
Peer
10.255.240.2    AS
65001
5709
5658
0
0
0 1d 19:13:28 Establ
inet.0: 2/4/4/0
10.255.240.6    65002
5685
5631
0
1 1d 19:01:29 Establ
inet.0: 2/4/4/0
10.255.240.11   65005
5649
5597
0
0 1d 18:46:54 Establ
inet.0: 2/2/2/0
10.255.240.13   65006
5654
5600
0
0 1d 18:47:02 Establ
inet.0: 2/4/4/0
192.168.255.11  65002
6026
5990
0
1 1d 19:01:20 Establ
bgp.evpn.0: 14/33/33/0
192.168.255.12  65001
6018
6140
0
0 1d 19:13:25 Establ
bgp.evpn.0: 2/28/28/0
192.168.255.31  65006
6029
6019
0
0 1d 18:46:59 Establ
bgp.evpn.0: 21/32/32/0
192.168.255.32  65005
6012
6085
0
0 1d 18:46:52 Establ
bgp.evpn.0: 22/34/34/0
{master:0}
root@Dist1>
```

From the BGP summary, we can see that the underlay (10.255.240.X) peer relationships are established to indicate 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.

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

BGP Summary

Neighbor Information

1:41 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	10.255.240.11	65005	65003	7m	3	4	21	20	default	Underlay
Connected	Established	10.255.240.13	65006	65003	7m	3	5	21	23	default	Underlay
Connected	Established	192.168.255.31	65006	65003	7m	25	22	35	34	default	Overlay
Connected	Established	10.255.240.2	65001	65003	7m	4	5	25	23	default	Underlay
Connected	Established	10.255.240.6	65002	65003	7m	2	5	20	22	default	Underlay
Connected	Established	192.168.255.11	65002	65003	7m	4	37	22	44	default	Overlay
Connected	Established	192.168.255.12	65001	65003	7m	36	37	48	46	default	Overlay
Connected	Established	192.168.255.32	65005	65003	7m	19	27	32	40	default	Overlay

If BGP is not established then go back and validate the underlay links and addressing, and that the loopback addresses are correct. Loopback addresses must be pingable from other loopback addresses. For example, Dist1 can reach Access1 and Core's loopback address once the underlay eBGP peering sessions are established.

```
{master:0}
root@Dist1> ping 192.168.255.12 count 1
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=0.910 ms

--- 192.168.255.12 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.910/0.910/0.910/0.000 ms

{master:0}
root@Dist1> ping 192.168.255.32 count 1
PING 192.168.255.32 (192.168.255.32): 56 data bytes
64 bytes from 192.168.255.32: icmp_seq=0 ttl=64 time=4.299 ms

--- 192.168.255.32 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.299/4.299/4.299/0.000 ms
```

NOTE: eBGP sessions are established between adjacent layers in the Campus Fabric IP Clos.

Let's verify the routes are established to the to the core and other devices across multiple paths. For example, Access1/2 should leverage both paths through Dist1/2 to access Core1/2's loopbacks as well as each other's loopbacks.

Access1: Loopback reachability to Core1 through Dist1/2

```
{master:0}
root@Access1> show route forwarding-table destination 192.168.255.11
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
192.168.255.11/32 user      0
                  10.255.240.12 ucst    1898     6 ge-0/0/36.0
                  10.255.240.16 ucst    1899     6 ge-0/0/37.0
```

Access1: Loopback reachability with Core2 through Dist1/2

```
{master:0}
root@Access1> show route forwarding-table destination 192.168.255.12
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
192.168.255.12/32 user    0          10.255.240.12      ucst   1898    6 ge-0/0/36.0
                  10.255.240.16      ucst   1899    6 ge-0/0/37.0
```

Access1: Loopback reachability with Access2 through Dist1/2

```
{master:0}
root@Access1> show route forwarding-table destination 192.168.255.32
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
192.168.255.32/32 user    0          10.255.240.12      ucst   1898    6 ge-0/0/36.0
                  10.255.240.16      ucst   1899    6 ge-0/0/37.0
```

This can be repeated for Access 2 and so forth to verify ECMP load balancing.

Meaning: At this point, BGP underlay and overlay are operational through the verification of eBGP between corresponding layers of the Campus Fabric and that routes are established to access, core, and distribution.

EVPN VXLAN Verification Between Access and Core Switches

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

Verification of the EVPN Database on Both Access Switches

```
{master:0}
root@Access1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address      Active source      Timestamp      IP address
1      00:cc:34:f3:cf:00  192.168.255.32    Nov 07 23:13:46
1      00:cc:34:f4:72:00  irb.0             Nov 07 23:13:20
1      f4:b5:2f:f3:fb:f0  192.168.255.12    Nov 07 23:13:34
1      f4:b5:2f:f4:0b:f0  192.168.255.11    Nov 07 23:13:34
11099  00:00:5e:e4:31:57  irb.1099          Nov 07 23:13:20  10.99.99.1
11099  52:54:00:74:a0:6f  ge-0/0/11.0       Nov 09 11:11:38  10.99.99.99
21088  00:00:5e:e4:31:57  irb.1088          Nov 08 15:29:02  10.88.88.1
21088  52:54:00:f7:12:2d  192.168.255.32    Nov 07 23:13:46  10.88.88.88
21088  f4:a7:39:6b:e3:20  192.168.255.32    Nov 08 04:21:53  10.88.88.10
31033  00:00:5e:e4:31:57  irb.1033          Nov 07 23:13:20  10.33.33.1
31033  5c:5b:35:2e:53:61  192.168.255.32    Nov 08 15:25:52
31033  5c:5b:35:af:29:d5  ge-0/0/16.0       Nov 08 15:25:52

{master:0}
root@Access1> show evpn database | match 52:54:00:74:a0:6f
11099  52:54:00:74:a0:6f  ge-0/0/11.0       Nov 09 11:11:38  10.99.99.99

{master:0}
root@Access1>
```

You can view the entire database or search by MAC address.

```

root@Access2>
{master:0}
root@Access2> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address      Active source      Timestamp          IP address
1      00:cc:34:f3:cf:00  irb.0            Nov 07 23:13:26
1      00:cc:34:f4:72:00  192.168.255.31   Nov 07 23:13:46
1      f4:b5:2f:f3:fb:f0  192.168.255.12   Nov 07 23:13:46
1      f4:b5:2f:f4:0b:f0  192.168.255.11   Nov 07 23:13:46
11099  00:00:5e:e4:31:57  irb.1099        Nov 08 15:31:24  10.99.99.1
11099  52:54:00:74:a0:6f  192.168.255.31   Nov 07 23:16:31  10.99.99.99
21088  00:00:5e:e4:31:57  irb.1088        Nov 07 23:13:26  10.88.88.1
21088  52:54:00:f7:12:2d  ge-0/0/12.0     Nov 07 23:13:27  10.88.88.88
21088  f4:a7:39:6b:e3:20  ge-0/0/12.0     Nov 09 07:50:17  10.88.88.10
31033  00:00:5e:e4:31:57  irb.1033        Nov 07 23:13:26  10.33.33.1
31033  5c:5b:35:2e:53:61  ge-0/0/16.0     Nov 08 15:25:52
31033  5c:5b:35:af:29:d5  192.168.255.31   Nov 08 15:25:52

{master:0}
root@Access2> show evpn database | match 52:54:00:74:a0:6f
11099  52:54:00:74:a0:6f  192.168.255.31   Nov 07 23:16:31  10.99.99.99

{master:0}
root@Access2>

```

Both access 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 access switch. These entries are learned locally or through the Campus Fabric as represented in the Active Source output.

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 access and core switches.

Access

```

{master:0}
root@Access1> show configuration vlans |display set |display inheritance | match 1099
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 l3-interface irb.1099
set vlans vlan1099 vxlan vni 11099

```

Core

```

root@Core1> show configuration |display s|match 1099
set groups top routing-instances evpn_vrf vlans vlan1099 vxlan vni 11099

root@Core1>

```

Verification of VXLAN Tunnelling Between Access and Core Devices

Access1:

```

{master:0}
root@Access1> 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   192.168.255.31 lo0.0  0
RVTEP-IP            L2-RTT                                IFL-Idx  Interface  NH-Id  RVTEP-Mode  ELP-IP  Flags
192.168.255.11      default-switch        618      vtep.32770  1901    RNVE
192.168.255.12      default-switch        617      vtep.32769  1900    RNVE
192.168.255.32      default-switch        619      vtep.32771  1912    RNVE

```

Access 2:

```
{master:0}
root@Access2> 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   192.168.255.32 lo0.0  0
RVTEP-IP      L2-RTT      IFL-Idx  Interface  NH-Id  RVTEP-Mode  ELP-IP      Flags
192.168.255.11 default-switch  618      vtep.32770  1901    RNVE
192.168.255.12 default-switch  617      vtep.32769  1900    RNVE
192.168.255.31 default-switch  619      vtep.32771  1902    RNVE
```

NOTE: Both access switches display each other in the output as well as Core1/2. The reason VXLAN is supported on Core1/2 is in support of the L2 multihomed connection to the WAN router, in this case a Juniper SRX Series Firewall. The L2 connection was built using Mist Port Profiles and does not require you to leave the Campus Fabric section once built.

Verify if Desktop1's MAC address is advertised via BGP:

```
root@Access1> show route advertising-protocol bgp 192.168.255.21 evpn-mac-address 52:54:00:74:a0:6f table bgp.evpn.0
Warning: License key missing; One or more members of the VC require 'bgp' license

bgp.evpn.0: 46 destinations, 46 routes (46 active, 0 holddown, 0 hidden)
Prefix      Nexthop      MED  Lclpref  AS path
2:192.168.255.31:1::11099::52:54:00:74:a0:6f/304 MAC/IP
*           Self              I
2:192.168.255.31:1::11099::52:54:00:74:a0:6f::10.99.99.99/304 MAC/IP
*           Self              I

{master:0}
root@Access1>
```

Verify if it is received on the core.

```
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: 46 destinations, 53 routes (46 active, 0 holddown, 0 hidden)
Prefix      Nexthop      MED  Lclpref  AS path
2:192.168.255.31:1::11099::52:54:00:74:a0:6f/304 MAC/IP
*           192.168.255.31              65003 65006 I
2:192.168.255.31:1::11099::52:54:00:74:a0:6f::10.99.99.99/304 MAC/IP
*           192.168.255.31              65003 65006 I

{master}
root@Core1>
```

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

```
root@Core1> show evpn database | match 52:54:00:74:a0:6f
11099      52:54:00:74:a0:6f  192.168.255.31      Feb 09 18:59:46  10.99.99.99

{master}
root@Core1>
```

Verify the MAC address mapped to the correct VTEP interface. This is on the core; you can also verify on access switch.

```

root@Core1> show route forwarding-table family ethernet-switching extensive destination 52:54:00:74:a0:6f
Routing table: evpn_vs.evpn-vxlan [Index 8]
Bridging domain: vlan1099.evpn-vxlan [Index 53]
VPLS:
Enabled protocols: Bridging, ACKed by all peers, EVPN VXLAN,
Destination: 52:54:00:74:a0:6f/48
  Learn VLAN: 0
  Route reference: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  IFL generation: 277
  Sequence Number: 0
  L2 Flags: control_dyn
  Flags: sent to PFE
  Nexthop:
  Next-hop type: composite      Index: 670      Reference: 14
{master}
root@Core1>

```

```

root@Core1> show route forwarding-table family ethernet-switching
Routing table: default-switch.bridge
VPLS:
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0
xe-1/1/0.0       intf  0
xe-1/1/1.0       intf  0

Routing table: evpn_vs.evpn-vxlan
VPLS:
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0
vtep.32769        intf  0
vtep.32770        intf  0
vtep.32771        intf  0

```

Finally, the VTEP interface is up and passing traffic.

```

root@Core1> show interfaces vtep.32771
Logical interface vtep.32771 (Index 346) (SNMP ifIndex 578)
Flags: Up SNMP-Traps Encapsulation: ENET2
VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.255.31, L2 Routing Instance: evpn_vs, L3 Routing Instance: default
Input packets : 4138
Output packets: 335
Protocol eth-switch, MTU: Unlimited
Flags: Trunk-Mode
{master}
root@Core1>

```

From an EVPN-VLAN perspective everything is correct. 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 Gateway Core EX9204 Switches

Remember that you chose to deploy the Border Gateway capability on the EX9204 Switches during the Campus Fabric IP Clos deployment, represented below:




Figure 13: Layer 2 ESI-LAG Supporting Active-Active Load Balancing

Mist enables the EX9204 Switch to translate between VXLAN traffic within the Campus Fabric and standard Ethernet switching for external connectivity, in this case a SRX Series Firewall. 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 as 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 Switch port:

PORT CONFIGURATION

Port Profile Assignment
★ Site, Template, or System Defined

Edit Port Range ✓ ✕

☒ Port Aggregation
☐ Disable LACP

AE Index (0 - 127)

☒ Esilag

Port IDs

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

Configuration Profile

☐ Enable Dynamic Configuration

Description

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              101          103           0                0

root@Core1> show configuration interfaces ae0 |display set |display inheritance
set interfaces ae0 hold-time up 120000
set interfaces ae0 hold-time down 1
set interfaces ae0 esi 00:11:11:11:11:11:01:00
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
1      00:cc:34:f3:cf:00 192.168.255.32     Nov 07 23:13:46
1      00:cc:34:f4:72:00 192.168.255.31     Nov 07 23:13:34
1      f4:b5:2f:f3:fb:f0 192.168.255.12     Nov 07 22:59:09
1      f4:b5:2f:f4:0b:f0 irb.0              Nov 07 22:59:10
11099  00:00:5e:e4:31:57 192.168.255.31     Nov 07 23:13:34  10.99.99.1
11099  52:54:00:74:a0:6f 192.168.255.31     Nov 07 23:16:31  10.99.99.99
11099  f0:1c:2d:c8:e8:f0 00:11:11:11:11:11:01:00 Nov 09 17:40:47  10.99.99.254
21008  00:00:5e:e4:31:57 192.168.255.31     Nov 08 15:29:02  10.88.88.1
21008  52:54:00:f7:12:2d 192.168.255.32     Nov 07 23:13:46  10.88.88.88
21008  f0:1c:2d:c8:e8:f0 00:11:11:11:11:11:01:00 Nov 09 17:40:55  10.88.88.254
21008  f4:a7:39:6b:e3:20 192.168.255.32     Nov 08 04:21:53  10.88.88.10
31033  00:00:5e:e4:31:57 192.168.255.31     Nov 07 23:13:34  10.33.33.1
31033  5c:5b:35:2e:53:61 192.168.255.32     Nov 08 15:25:52
31033  5c:5b:35:af:29:d5 192.168.255.31     Nov 08 15:25:52
31033  f0:1c:2d:c8:e8:f0 00:11:11:11:11:11:01:00 Nov 09 17:40:52  10.33.33.254

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

```

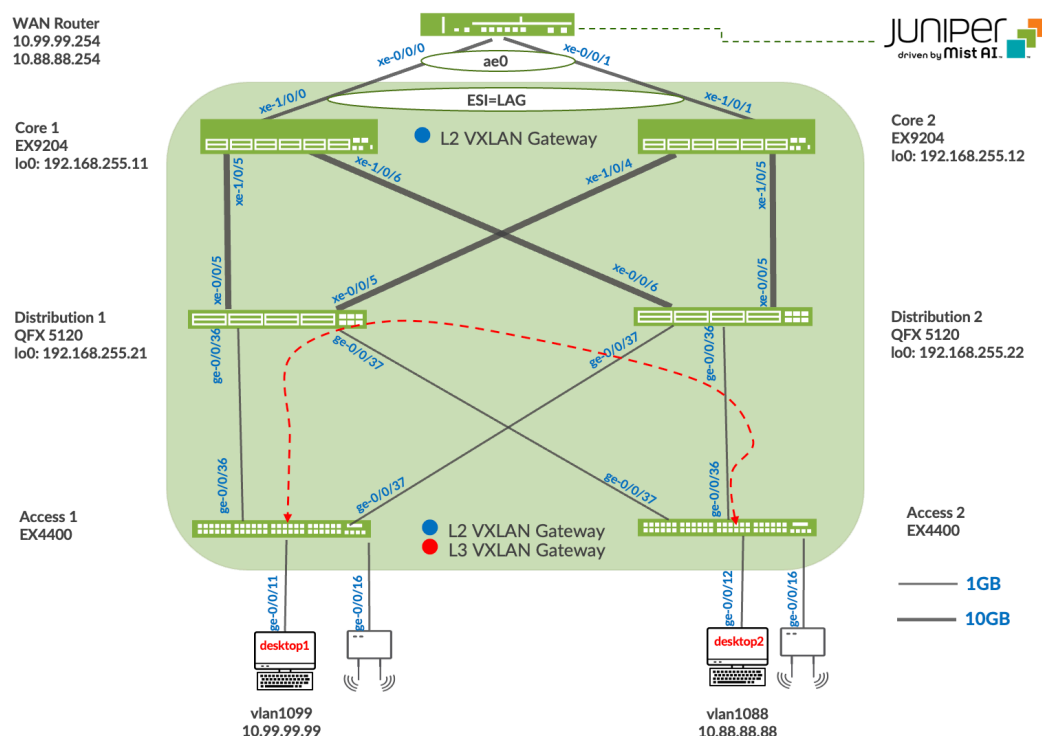
Last step is to verify Desktop1 can ping Desktop2.

```

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 is 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 IP Clos deployment using telemetry such as BGP neighbor status and TX/RX port statistics. The following screens are taken from the Campus Fabric IP Clos build by accessing the Campus Fabric option under the Organization/Wired of the Mist Portal:

< Campus Fabrics : **Campus Fabric IPClos** Edit Configuration Delete Connection Table

Core

Distribution

Access

Neighbor Information 2:50 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	192.168.255.22	65004	65002	11m	37	41	56	54	default	Overlay
Connected	Established	10.255.240.7	65003	65002	11m	5	2	33	30	default	Underlay
Connected	Established	10.255.240.9	65004	65002	11m	5	5	34	34	default	Underlay
Connected	Established	192.168.255.21	65003	65002	11m	37	4	55	31	default	Overlay

Core1

MAC Address f4:b5:2f:f4:04:00
Model EX9204
Status connected
Site Primary Site
Router ID 192.168.255.11

VLANs

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

Connections to Distribution

Switch	RX Bytes	TX Bytes	Link Status
Dist1-	1.2 GB	1.3 GB	Up
Dist2-	1.1 GB	1 GB	Up

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

< Campus Fabrics : **Campus Fabric IPClos** Edit Configuration Delete Connection Table

Core

Distribution

Access

BGP Summary

Neighbor Information 2:52 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	192.168.255.32	65005	65003	13m	19	27	46	54	default	Overlay
Connected	Established	10.255.240.2	65001	65003	13m	4	5	38	36	default	Underlay
Connected	Established	10.255.240.6	65002	65003	13m	2	5	33	35	default	Underlay

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
1033	--	vlan1033

Connections to Core

Switch	RX Bytes	TX Bytes	Link Status
Core2	2.4 GB	2 GB	Up
Core1	2.3 GB	2.3 GB	Up

Connections to Access

Switch	RX Bytes	TX Bytes	Link Status
Access2	388.5 MB	183.7 MB	Up
Access1	2.4 GB	3.2 GB	Up

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

NOTE: The full BGP peering table is not shown.

< Campus Fabrics : **Campus Fabric IPClos** Edit Configuration Delete Connection Table

Core

Core2 Core1

Distribution

Dist1 Dist2

Access

Access2 Access1

Neighbor Information 2:53 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	10.255.240.12	65003	65006	13m	5	3	37	33	default	Underlay
Connected	Established	10.255.240.16	65004	65006	13m	5	4	39	36	default	Underlay
Connected	Established	192.168.255.22	65004	65006	13m	22	35	63	53	default	Overlay
Connected	Established	192.168.255.21	65003	65006	13m	22	25	49	48	default	Overlay

Access1

MAC Address 00:cc:34:f4:72:00
Model EX4400-48P
Status connected
Site Primary Site
Router ID 192.168.255.31

VLANs

ID	IP Address	Name
1088	10.88.88.1	vlan1088
1099	10.99.99.1	vlan1099
1033	10.33.33.1	vlan1033

Connections to Distribution

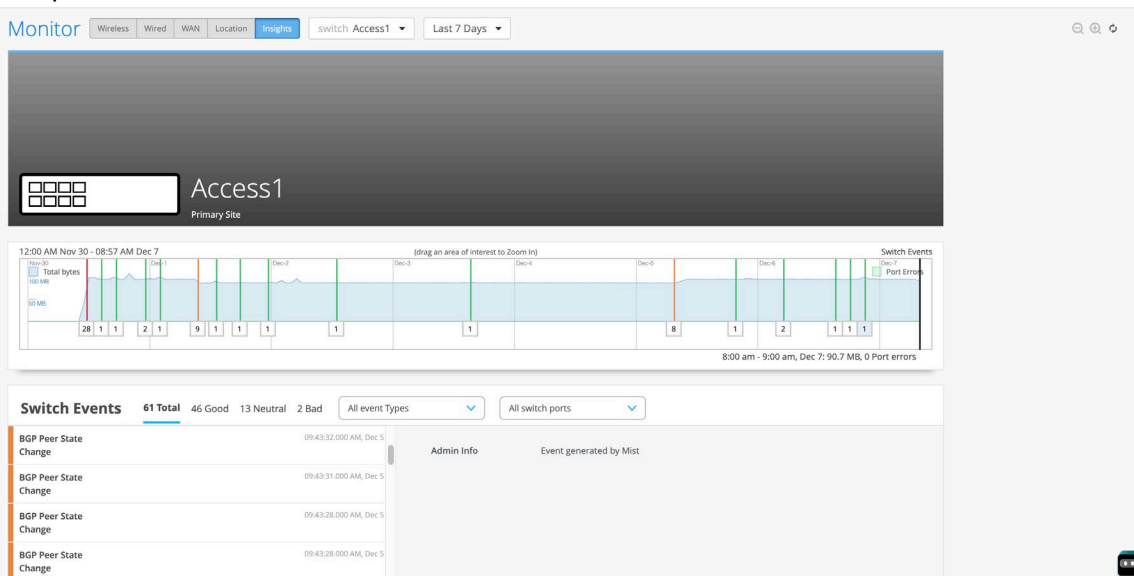
Switch	RX Bytes	TX Bytes	Link Stat
Dist1-	1.4 GB	1.2 GB	Up
Dist2-	22.7 MB	116.1 MB	Up

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

NOTE: The full BGP peering table is not shown.

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

Switch Insights of Access1 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 IP Clos 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

Configuration of the Underlay IP Fabric

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

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

- BGP Local AS (65001)
- Loopback Prefix (/24)
- Subnet (10.255.240.0/20) – point to point interfaces between adjacent layers

Mist enables per-packet (Junos 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 and the two access 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

Access Interfaces:
set interfaces ge-0/0/36 description evpn_downlink-to-00cc34f47200
set interfaces ge-0/0/36 unit 0 family inet address 10.255.240.12/31
set interfaces ge-0/0/37 description evpn_downlink-to-00cc34f3cf00
set interfaces ge-0/0/37 unit 0 family inet address 10.255.240.10/31
```

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 group evpn_underlay neighbor 10.255.240.11 peer-as 65005
set protocols bgp group evpn_underlay neighbor 10.255.240.13 peer-as 65006
set protocols bgp graceful-restart
```

Dist2 Configuration

1. Interconnects between the two core switches and the two access 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

Access Interfaces:
set interfaces ge-0/0/36 description evpn_downlink-to-00cc34f3cf00
set interfaces ge-0/0/36 unit 0 family inet address 10.255.240.14/31
set interfaces ge-0/0/37 description evpn_downlink-to-00cc34f47200
set interfaces ge-0/0/37 unit 0 family inet address 10.255.240.16/31
```

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 group evpn_underlay neighbor 10.255.240.15 peer-as 65005
set protocols bgp group evpn_underlay neighbor 10.255.240.17 peer-as 65006
set protocols bgp graceful-restart
```

Access1 Configuration

1. Interconnects between the two distribution switches.

```
set interfaces ge-0/0/36 description evpn_uplink-to-d8539a646fc0
set interfaces ge-0/0/36 unit 0 family inet address 10.255.240.13/31
set interfaces ge-0/0/37 description evpn_uplink-to-d8539a64b5c0
set interfaces ge-0/0/37 unit 0 family inet address 10.255.240.17/31
```

2. Loopback interface and router ID.

```
set groups top interfaces lo0 unit 0 family inet address 192.168.255.31/32
set groups top routing-options router-id 192.168.255.31
```

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 65006
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 10.255.240.12 peer-as 65003
set protocols bgp group evpn_underlay neighbor 10.255.240.16 peer-as 65004
set protocols bgp graceful-restart
```

Access2 Configuration

1. Interconnects between the two distribution switches.

```
set interfaces ge-0/0/36 description evpn_uplink-to-d8539a64b5c0
set interfaces ge-0/0/36 unit 0 family inet address 10.255.240.15/31
set interfaces ge-0/0/37 description evpn_uplink-to-d8539a646fc0
set interfaces ge-0/0/37 unit 0 family inet address 10.255.240.11/31
```

2. Loopback interface and router ID.

```
set groups top interfaces lo0 unit 0 family inet address 192.168.255.32/32
set groups top routing-options router-id 192.168.255.32
```

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 65005
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 10.255.240.10 peer-as 65003
set protocols bgp group evpn_underlay neighbor 10.255.240.14 peer-as 65004
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, distribution, and access 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 adjacent layers.

Mist provisions L3 IRB interfaces on the access layer if the Routed at Distribution option was chosen during the initial phases of the Campus Fabric build, the L3 IRB interfaces are on the distribution switches.

Mist enables VXLAN tunneling, VLAN to VXLAN mapping, and MP BGP configuration snippets such as vrf-targets on the access layer switches. The core switches have VXLAN tunnelling and VLAN to VXLAN mapping enabled based on the selection of the Core as a Border option.

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 evpn_vs instance-type virtual-switch
set groups top routing-instances evpn_vs protocols evpn encapsulation vxlan
set groups top routing-instances evpn_vs protocols evpn default-gateway do-not-advertise
set groups top routing-instances evpn_vs protocols evpn extended-vni-list all
set groups top routing-instances evpn_vs protocols rstp interface ael disable
set groups top routing-instances evpn_vs protocols rstp bpdu-block-on-edge

```

```

set groups top routing-instances evpn_vs vtep-source-interface lo0.0
set groups top routing-instances evpn_vs interface ael.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

```

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

```

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 evpn_vs instance-type virtual-switch
set groups top routing-instances evpn_vs protocols evpn encapsulation vxlan
set groups top routing-instances evpn_vs protocols evpn default-gateway do-not-advertise
set groups top routing-instances evpn_vs protocols evpn extended-vni-list all
set groups top routing-instances evpn_vs protocols rstp interface ael disable
set groups top routing-instances evpn_vs protocols rstp bpdu-block-on-edge
set groups top routing-instances evpn_vs vtep-source-interface lo0.0
set groups top routing-instances evpn_vs interface ael.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

```

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
```

Dist1 Configuration

1. BGP overlay peering between the two core switches and the two access 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
set protocols bgp group evpn_overlay neighbor 192.168.255.32 peer-as 65005
set protocols bgp group evpn_overlay neighbor 192.168.255.31 peer-as 65006
```

Dist2 Configuration

1. BGP overlay peering between the two core switches and the two access 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
set protocols bgp group evpn_overlay neighbor 192.168.255.32 peer-as 65005
set protocols bgp group evpn_overlay neighbor 192.168.255.31 peer-as 65006
```

Access1 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 192.168.255.31
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 65006
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 switch-options vtep-source-interface lo0.0
set groups top switch-options route-distinguisher 192.168.255.31:1
set groups top switch-options vrf-target target:65000:1
```

3. VXLAN encapsulation.

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

4. VRFs that are used for traffic isolation.

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi 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 192.168.255.31:103
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 192.168.255.31:102
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 192.168.255.31:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
```

5. VLAN to VXLAN mapping.

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

```
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 l3-interface irb.1099
set vlans vlan1099 vxlan vni 11099
```

6. L3 IRB interface enablement with anycast addressing.

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

Access2 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 192.168.255.32
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 65005
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 switch-options vtep-source-interface lo0.0
set groups top switch-options route-distinguisher 192.168.255.32:1
set groups top switch-options vrf-target target:65000:1
```

3. VXLAN encapsulation.

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

4. VRFs that are used for traffic isolation.

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi 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 192.168.255.32:103
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 192.168.255.32:102
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 192.168.255.32:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label

```

5. VLAN to VXLAN mapping.

```

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

```

6. L3 IRB interface enablement with anycast addressing.

```

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

```

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

This section displays the Juniper Mist Cloud configuration output for the enablement of the Layer 2 ESI Link Aggregation Groups (LAG) 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.




Figure 14: Layer 2 ESI-LAG Supporting Active-Active Load Balancing

Core 1 Configuration

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

```
set interfaces xe-1/0/0 hold-time up 120000
set interfaces xe-1/0/0 hold-time down 1
set interfaces xe-1/0/0 ether-options 802.3ad ae1
set interfaces xe-1/0/0 unit 0 family ethernet-switching storm-control default

set groups esilag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esilag interfaces <*> unit 0 family ethernet-switching vlan members all

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 groups esilag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups esilag interfaces <*> unit 0 family ethernet-switching vlan members all

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.