Campus Fabric Core Distribution CRB Wired Assurance

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
www.juniper.net

Issue Date: 15 December 2022

Contents

Introduction to Campus Fabric Core Distribution Wired Assurance

About This Network Configuration Example
This network configuration example (NCE) describes how to deploy a Campus Fabric Core Distribution ERB 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.

Use Case Overview

Enterprise networks are undergoing massive transitions to accommodate the growing demand for cloud-ready, scalable, and efficient networks, and the plethora of 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 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.

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/rfc7348) 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. The simple IP-based Layer 3 network underlay limits the Layer 2 broadcast domain and eliminates the need for Spanning 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 having to redesign the network. EVPN-VXLAN is vendor-agnostic, so you can use the existing access layer infrastructure and gradually migrate to access layer switches that support EVPN-VXLAN capabilities once the Core and Distribution part of the network is deployed. Connectivity with legacy switches that do not support EVPN VXLAN is accomplished with standards-based ESI-LAG.

Benefits of Campus Fabric Core Distribution
With increasing number of devices connecting to the network, you will need to scale your campus network rapidly without adding complexity. Many IoT devices have limited networking capabilities and require Layer 2 adjacency across buildings and campuses. Traditionally, this problem was solved by extending VLANs between endpoints using data plane-based flood and learn mechanisms inherent with ethernet switching technologies. The traditional ethernet switching approach is inefficient because it leverages inefficient broadcast and multicast technologies to manage MAC addresses. It is also difficult to manage because you need to configure and manually manage VLANs to extend them to new network ports. This problem increases multi-fold when you take into consideration the explosive growth of IoT and mobility.

A campus fabric based on EVPN-VXLAN is a modern and scalable network that uses BGP as the underlay for the core and distribution layer switches. The distribution and core layer switches function as VTEPs that encapsulate and decapsulate the VXLAN traffic. In addition, these devices route and bridge packets in and out of VXLAN tunnels.

The Campus Fabric Core Distribution extends the EVPN fabric to connect VLANs across multiple buildings by stretching the Layer 2 VXLAN network with routing occurring in the Core (CRB) or Distribution (ERB) layers. This network architecture the core and distribution layers of the toplology with integration to access switching via standard LACP

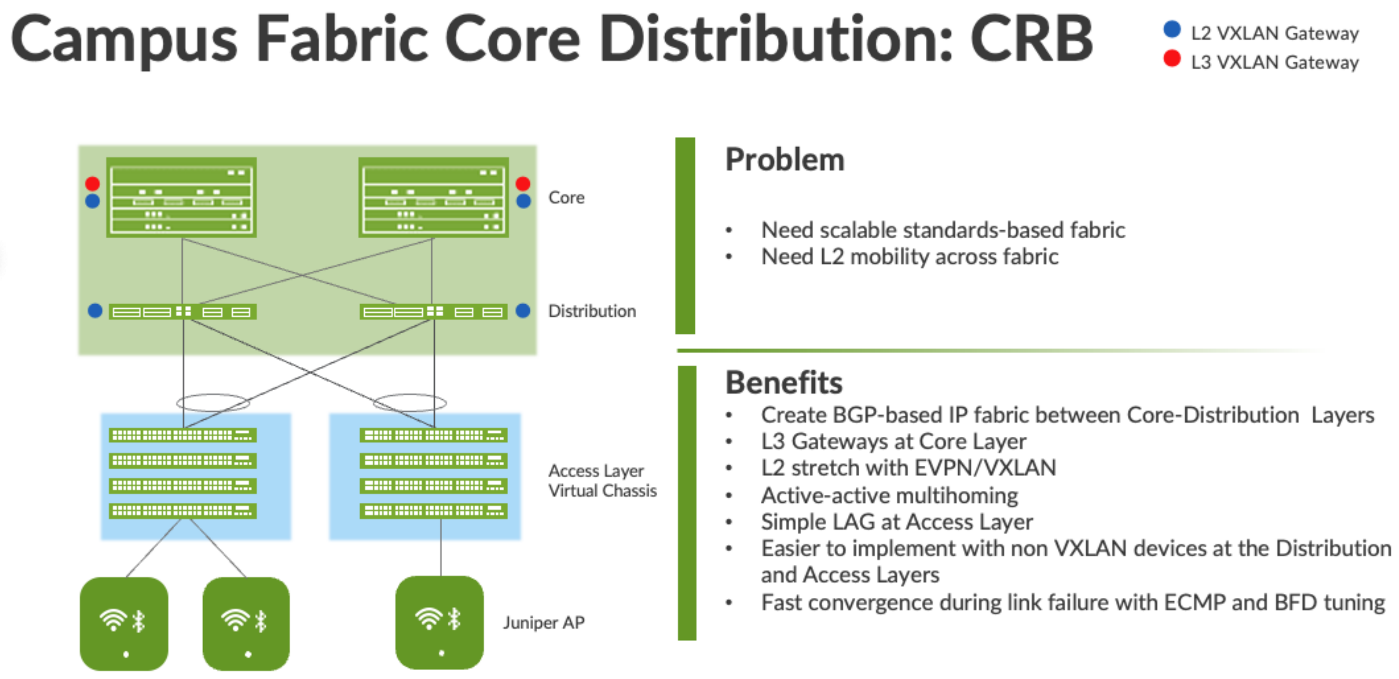

Campus Fabric Core Distribution: CRB

● L2 VXLAN Gateway
● L3 VXLAN Gateway

Problem

- Need scalable standards-based fabric
- Need L2 mobility across fabric

Benefits

- Create BGP-based IP fabric between Core-Distribution Layers
- L3 Gateways at Core Layer
- L2 stretch with EVPN/VXLAN
- Active-active multihoming
- Simple LAG at Access Layer
- Easier to implement with non VXLAN devices at the Distribution and Access Layers
- Fast convergence during link failure with ECMP and BFD tuning

Figure 1 Campus fabric Core Distribution CRB

An EVPN-VXLAN fabric solves these issues and provides the following benefits:

- Reduced flooding and learning—Control plane-based Layer 2/Layer 3 learning reduces the flood and learn issues associated with data plane learning. Learning MAC addresses in the forwarding plane has an adverse impact on network performance as the number of endpoints grows. This is because more management traffic consumes the bandwidth which leaving 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/Layer 3 learning allows the EVPN-VXLAN network to scale up to support hundreds of thousands of endpoints.
- Consistency—A universal EVPN-VXLAN-based architecture across campuses and data-centers enables seamless end-to-end network for endpoints and applications.
- Group Based Policies - With GBT you can enable micro segmentation and macro segmentation with EVPN-VXLAN to minimize Layer 2 flooding, provide traffic isolation, and simplify the network.
- 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.

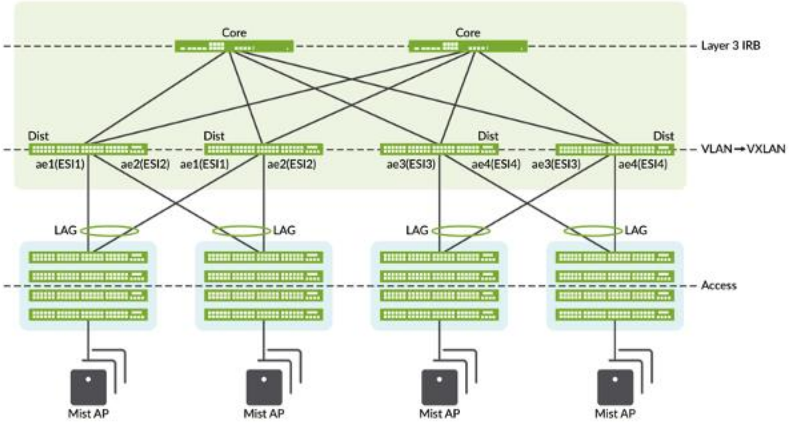

Figure 2 Sample topology

Technical Overview

Network overlays enable connectivity and addressing independent of the physical network. Ethernet frames are wrapped in IP UDP datagrams that are themselves 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. 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 VXLAN Tunnel Endpoint (VTEP) is known as the Layer 2 Gateway and typically assigned with the device's Loopback address.

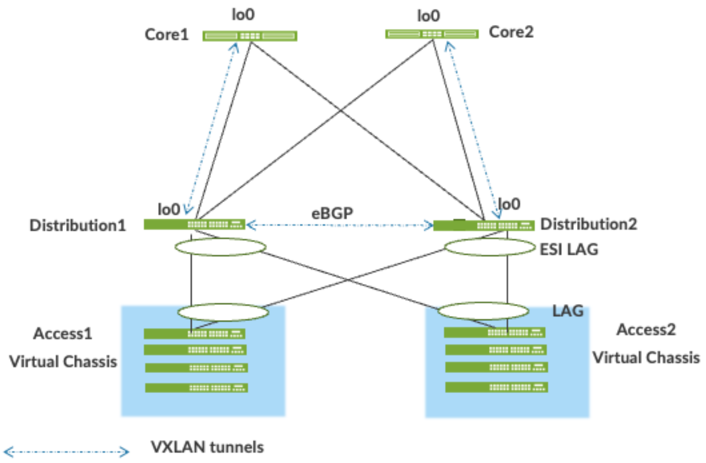

Figure 3. VXLAN VTEP tunnels

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 behaviour of the Ethernet protocol.

The two primary methods for using VXLAN without a control plane protocol—static unicast VXLAN tunnels and VXLAN tunnels that are signalled with a multicast underlay—do not solve the inherent flood and learn problem and are difficult to scale in large multitenant environments. https://www.rfc-editor.org/rfc/rfc7348

Understanding EVPN

Ethernet VPN (EVPN) is an 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 routes. EVPN enables devices acting as VTEPs to exchange reachability information with each other about their endpoints.

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 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.

Underlay Network

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

You could use any Layer 3 routing protocol to exchange loopback addresses between the access, core, and distribution devices. BGP provides benefits like 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 provision additional BGP speakers to allow customers 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 addres. The overlay is then ued to exchange EVPN routes.

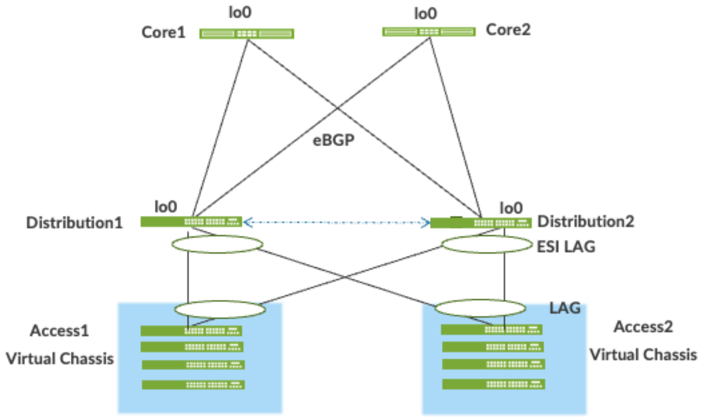

Figure 4. Pt-Pt /31 links Core-Distribution layers running eBGP

Overlay Network (Data Plane)

VXLAN is the overlay data plane encapsulation protocol that tunnels Ethernet frames between network endpoints over the Layer 3 IP network. Devices that perform VXLAN encapsulation and decapsulation for the network are referred to as a VXLAN tunnel endpoint (VTEP). Before a VTEP sends a frame into a VXLAN tunnel, it wraps the original frame in a VXLAN header that includes a 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.

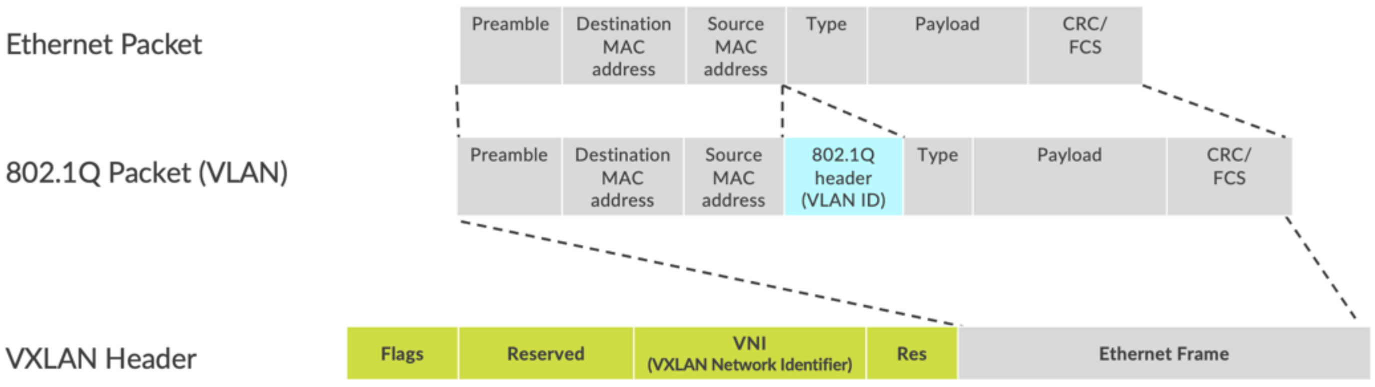

VLAN	VXLAN
4K	16M

Figure 5 VXlan Header

VTEPs are software entities tied to the devices' loopback address that source and terminate VXLAN tunnels. VXLAN tunnels are provisioned on the Core and Distribution Switches.

Overlay Network (Control Plane)

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

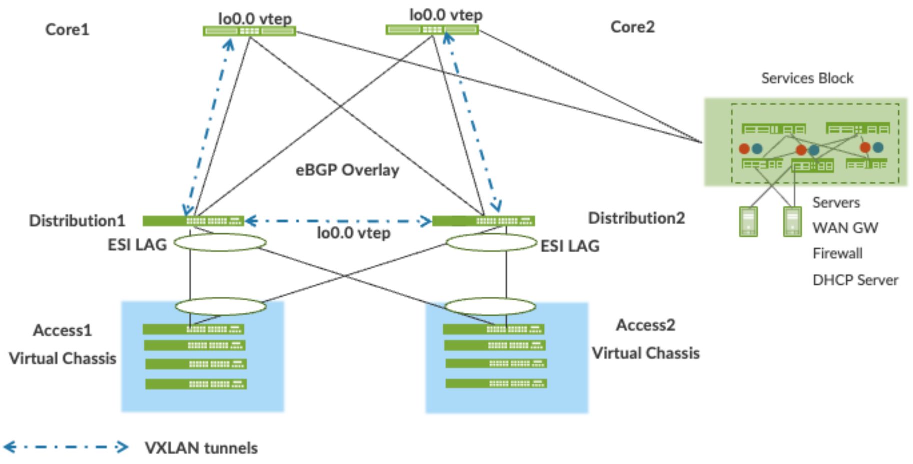

Figure 6

Resiliency and Load Balancing

Juniper supports BFD, Bi-Directional Forwarding, in conjunction with the underlay and overlay

networks. This provides fast convergence in the event of a device or link failure without relying on the routing protocol's timers. Mist configured BFD underlay and over minimum interval of 1000ms and 3000 ms in the underlay and overlay respectively. Load Balancing is supported across all links within the Campus Fabric using ECMP or Equal Cost Multi Pathing enabled at the forwarding plane.

Ethernet Segment Identifier (ESI)

When switches have aggregated ethernet interfaces or LAG and multihomed to two or more switches an ESI is a 10-octet integer that identifies that segment. This enables link failover in the event of a bad link and is automatically assigned by Mist.

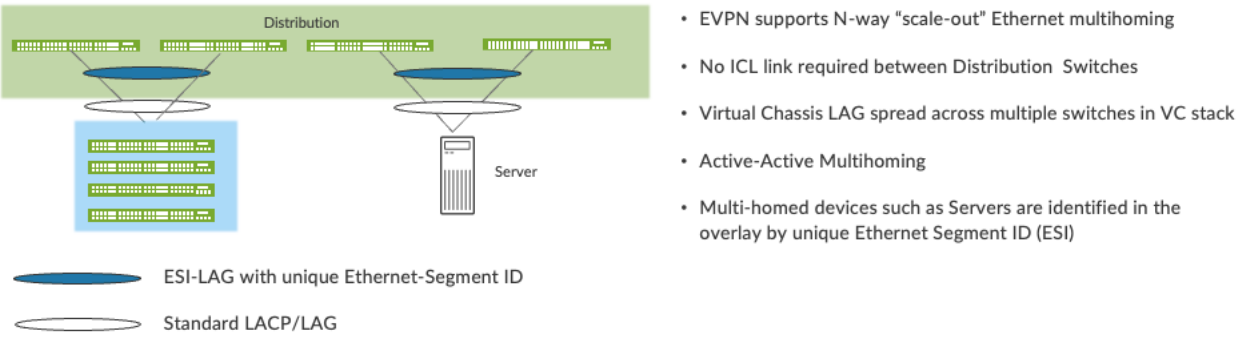

- EVPN supports N-way “scale-out” Ethernet multihoming
- No ICL link required between Distribution Switches
- Virtual Chassis LAG spread across multiple switches in VC stack
- Active-Active Multihoming
- Multi-homed devices such as Servers are identified in the overlay by unique Ethernet Segment ID (ESI)

ESI-LAG with unique Ethernet-Segment ID

Standard LACP/LAG

Figure 7. resiliency and load balancing

Services Block

Customers may wish to position critical infrastructure services off a dedicated Access Pair of Juniper switches which could include WAN and Firewall connectivity. This Services Block Border is supported directly off the Core Layer platforms or as a dedicated pair of switches.

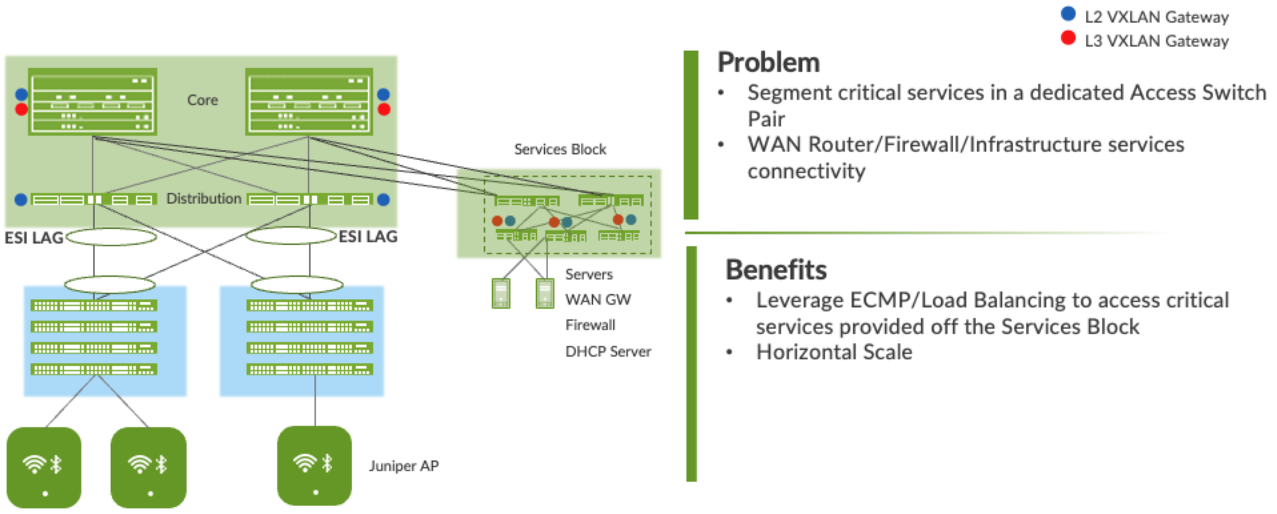

Problem

- 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 8. Services Block

Access Layer

The access layer provides network connectivity to end-user devices, such as personal computers, VoIP phones, printers, IoT devices, as well as connectivity to wireless access points. In this IP Clos campus design, the EVPN-VXLAN network extends all the way to the access layer switches.

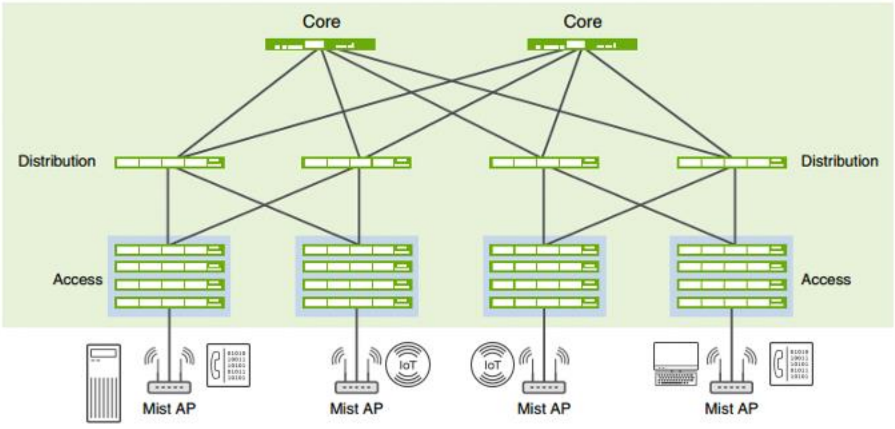

Figure 9. End point access

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

Juniper Access points

For this example, we choose Juniper Access points as our preferred access point devices. They are designed from the ground up to meet the stringent networking needs of the modern cloud and smart-device era.

Juniper Mist Wired Assurance

Mist Wired Assurance is a cloud service that brings automated operations and service levels to the Campus Fabric for switches, IoT devices, access points, servers, printers, etc. It's about simplification 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 the rich Junos streaming telemetry that enable the insights for switch health metrics and anomaly detection.

Mist's AI engine and virtual network assistant, Marvis, further simplifies troubleshooting while streamlining helpdesk operations by monitoring events and recommending actions. Marvis is one step towards the Self-Driving Network™, turning insights into actions and fundamentally transforming IT operations from reactive troubleshooting to proactive remediation.

Mist Cloud services are 100% programmable using open APIs for full automation and/or integration with your IT applications, such as Ticketing Systems, IP Management Systems, etc.

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

- UI 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 wired and wireless troubleshooting, trending analysis, anomaly detection, and proactive problem remediation.

To learn more about Juniper Mist Wired Assurance please access the following datasheet: https://www.juniper.net/content/dam/www/assets/datasheets/us/en/cloud-services/juniper-mist-wired-assurance-datasheet.pdf

Campus Fabric Core Distribituion High Level Architecture

The campus fabric, with an EVPN-VXLAN architecture, decouples the overlay network from the underlay network. This approach addresses the needs of the modern enterprise network by allowing

network administrators to create logical Layer 2 networks across one or more Layer 3 networks. By configuring different routing instances, you can enforce the separation of virtual networks because each routing instance has its own separate routing and switching table.

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

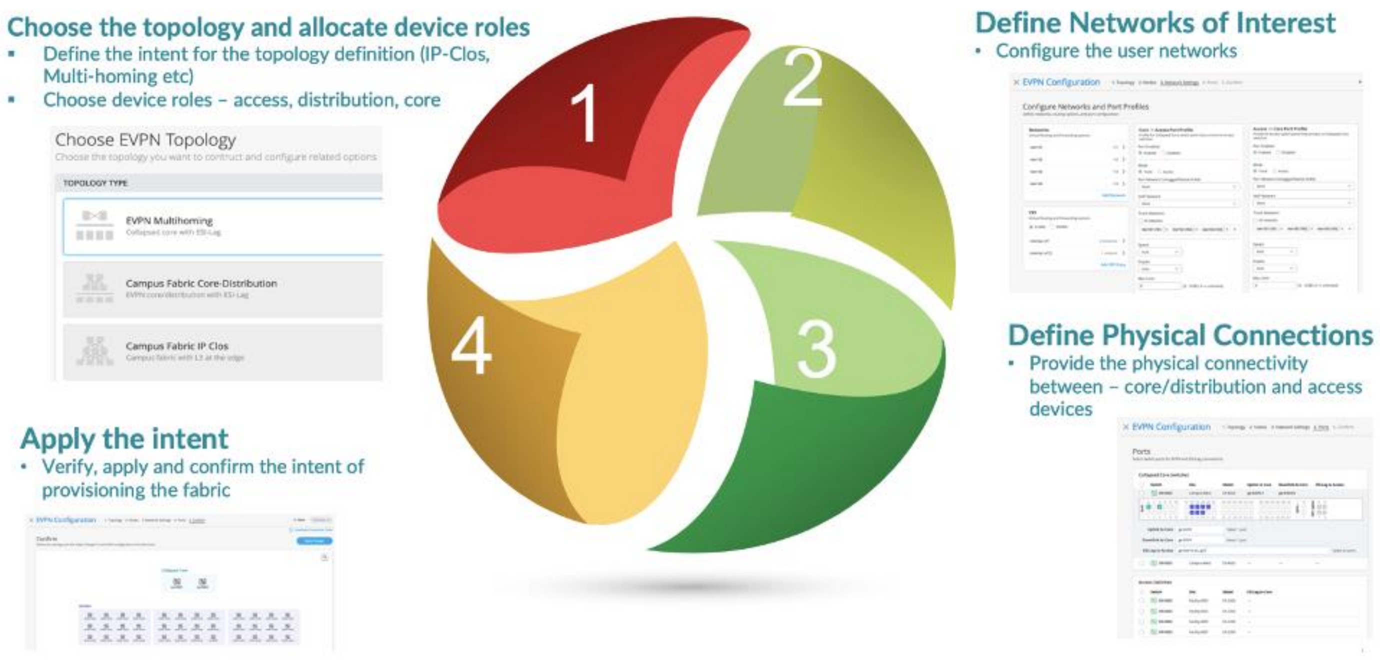

Choose the topology and allocate device roles

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

Define Networks of Interest

- Configure the user networks

Define Physical Connections

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

Apply the intent

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

1
2
3
4
Choose EVPN Topology
Choose the topology you want to contruct and configure related options
TOPOLOGY TYPE
EVPN Multihoming
Collapsed core with ESI-Lag
Campus Fabric Core-Distribution
EVPN core/distribution with ESI-Lag
Campus Fabric IP Clos
Campus fabric with L3 at the edge
EVPN Configuration
Configure Networks and Port Profiles
Ports
Confirm

Campus Fabric Core Distribution CRB Components

This configuration example uses the following devices:

- Two EX9204 switches as core devices, Software version: Junos OS Release 21.4R1.12 or later
- Two QFX5120 switches as distribution devices, Software version: Junos OS Release 21.4R1.12 or later
- Two Access Layer EX4400 switches, Software version: Junos OS Release 22.1R1.10 or later
- One SRX345 wan router, Software version: 20.2R3-S2.5 or later
- Juniper Access Points
- 2 Linux desktops that act as a wired client

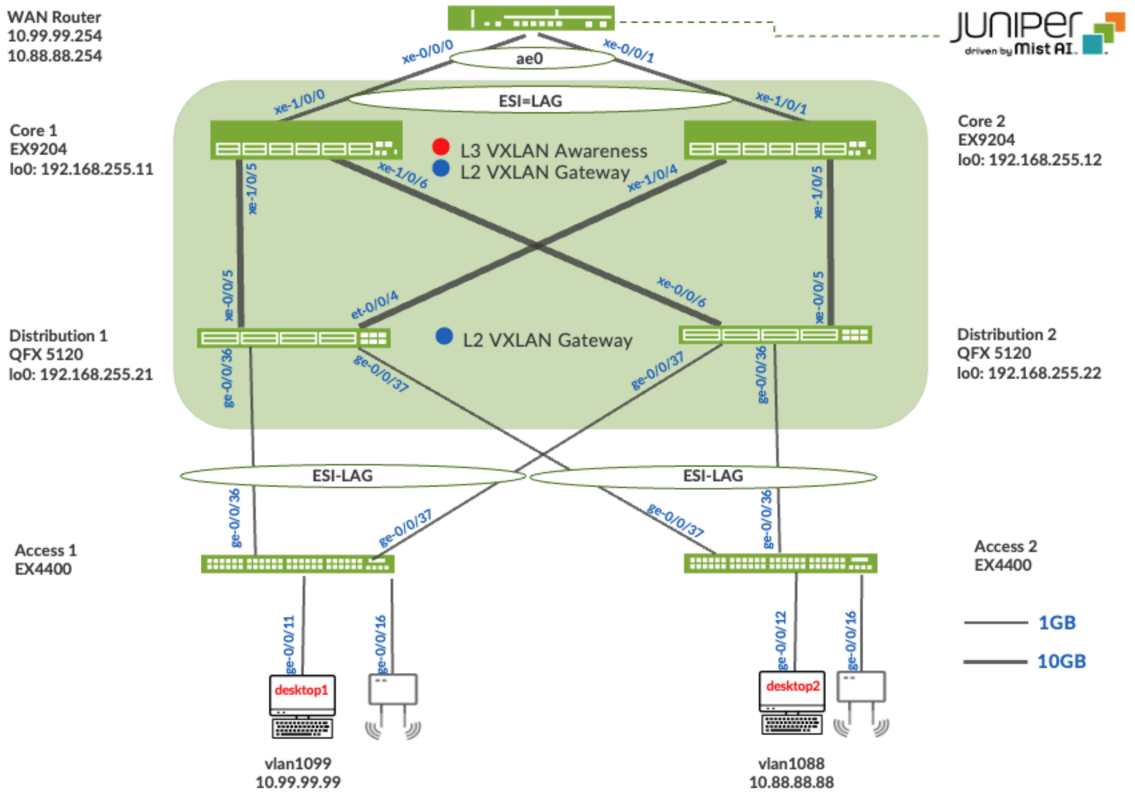

Figure 10. Topology

Juniper Mist Wired Assurance

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

- Assignment of p2p links between all layers of the Campus Fabric
- Assignment of unique BGP AS numbers per device participating in the underlay and overlay.
- Creation of VRF instances to allow the user the ability to logically segment traffic. This also includes the assignment of new or existing VLANs to each representative VRF
- IP addressing of each L3 gateway IRB
- 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
- 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, please leverage the following link:
https://www.mist.com/documentation/category/wired-assurance/

Juniper Mist Wired Assurance Switches Section

The user should validate that each device participating in the Campus Fabric has been adopted or claimed, and assigned to a site. The switches were named for respective layer in the fabric to facilitate building and operating the fabric.

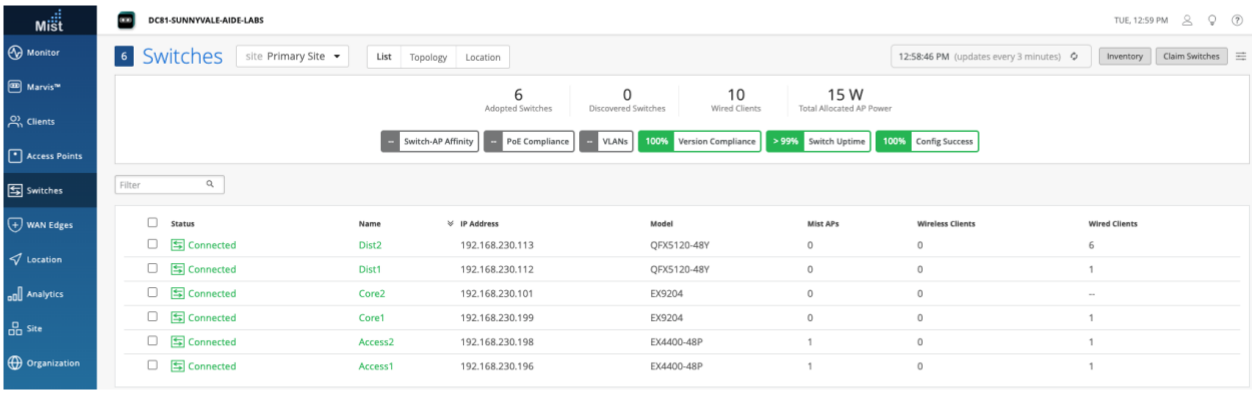

Status	Name	IP Address	Model	Mist APs	Wireless Clients	Wired Clients
Connected	Dist2	192.168.230.113	QFX5120-48Y	0	0	6
Connected	Dist1	192.168.230.112	QFX5120-48Y	0	0	1
Connected	Core2	192.168.230.101	EX9204	0	0	--
Connected	Core1	192.168.230.199	EX9204	0	0	1
Connected	Access2	192.168.230.198	EX4400-48P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

Overview

Use this network configuration example 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 as the underlay routing protocol to exchange loopback routes, and eBGP between the core and distribution 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 the ability to use configuration templates and a hierarchical model to group the switches and make bulk updates. Templates provide uniformity and convenience, while the hierarchy (Site and Switch) provides both scale and granularity.

What templates, and the hierarchical model, means in practice is 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 Network and Organizational levels that apply to the same device, the narrower settings (in this case, Network) 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 Network level. Of course, individual switches can also have their own unique configurations.

You can include individual CLI commands at any level of the hierarchy, which are then appended to all the switches in that group on an "AND" basis– that is, individual CLI settings are appended to the existing configuration (existing setting may deleted or appended).

Note: If a user utilizes CLI commands for items not native to the Mist UI, this configuration data will be applied last; overwriting existing configuration data within the same stanza.

Under Organization and Switch Templates, we utilize the following template

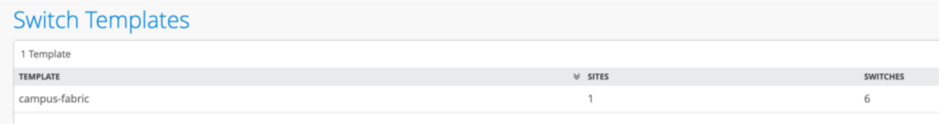

Switch Templates

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 Core and Distribution devices within the Campus Fabric Core Distribution. The devices such as the Access Layer of switches connect to the Distribution layer using standard LAG; while the Distribution utilizes ESI-LAG in a multihoming, load balancing manner.

The WAN router can be provisioned via Mist UI but is separate from the campus fabric workflow. The WAN router has a southbound lag configured to connect to the ESI-LAG on the core switches. WAN router can standalone or built as an HA cluster.

Create the Campus Fabric

From the Organization option on the leaf hand section of the Mist UI, select Wired Campus Fabric

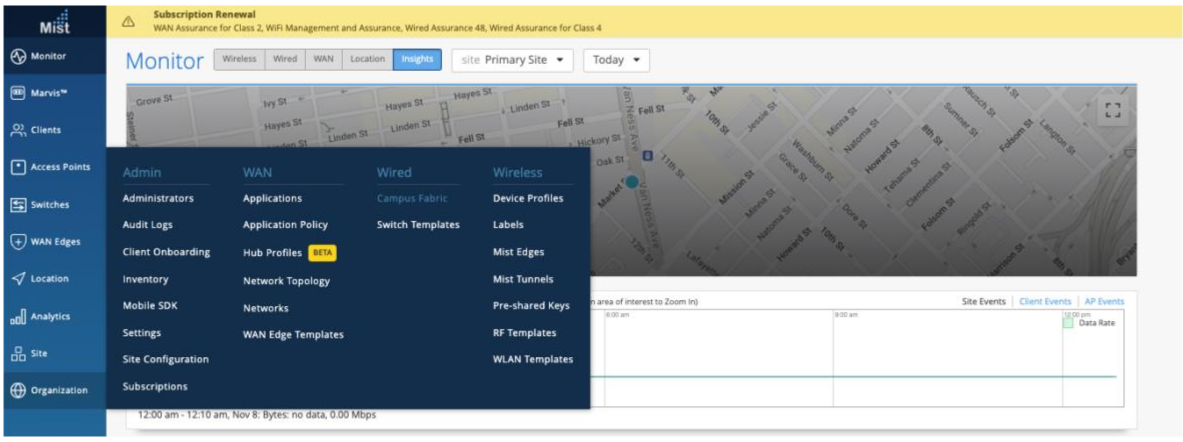

Subscription Renewal
WAN Assurance for Class 2, WiFi Management and Assurance, Wired Assurance 48, Wired Assurance for Class 4

12:00 am - 12:10 am, Nov 8: Bytes: no data, 0.00 Mbps

Mist provides the option of deploying a Campus Fabric at the Org or Site level noted on the upper left hand Campus Fabric pull down menu shown below. For example, those who are building a Campus wide architecture with multiple buildings, each building housing distribution and access switches, could consider building an Org level Campus Fabric that 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 would suffice.

Campus Fabric Org Build

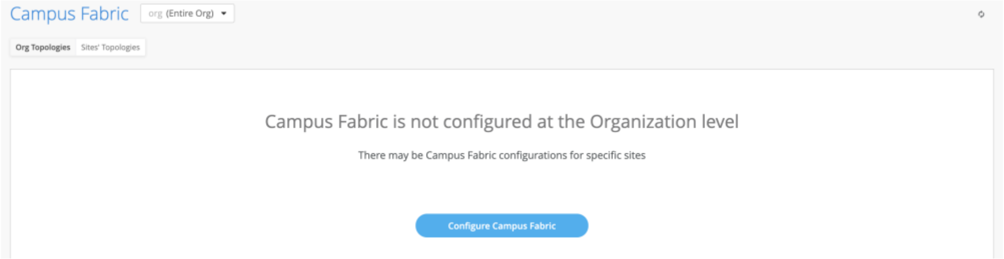

Campus Fabric

org (Entire Org)

Org Topologies | Sites' Topologies

Campus Fabric is not configured at the Organization level

There may be Campus Fabric configurations for specific sites

Configure Campus Fabric

Campus Fabric Site Build

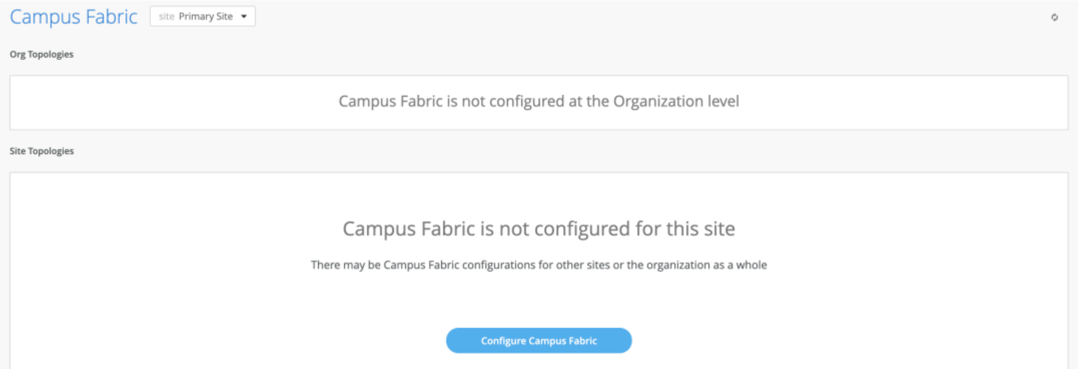

Campus Fabric

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:

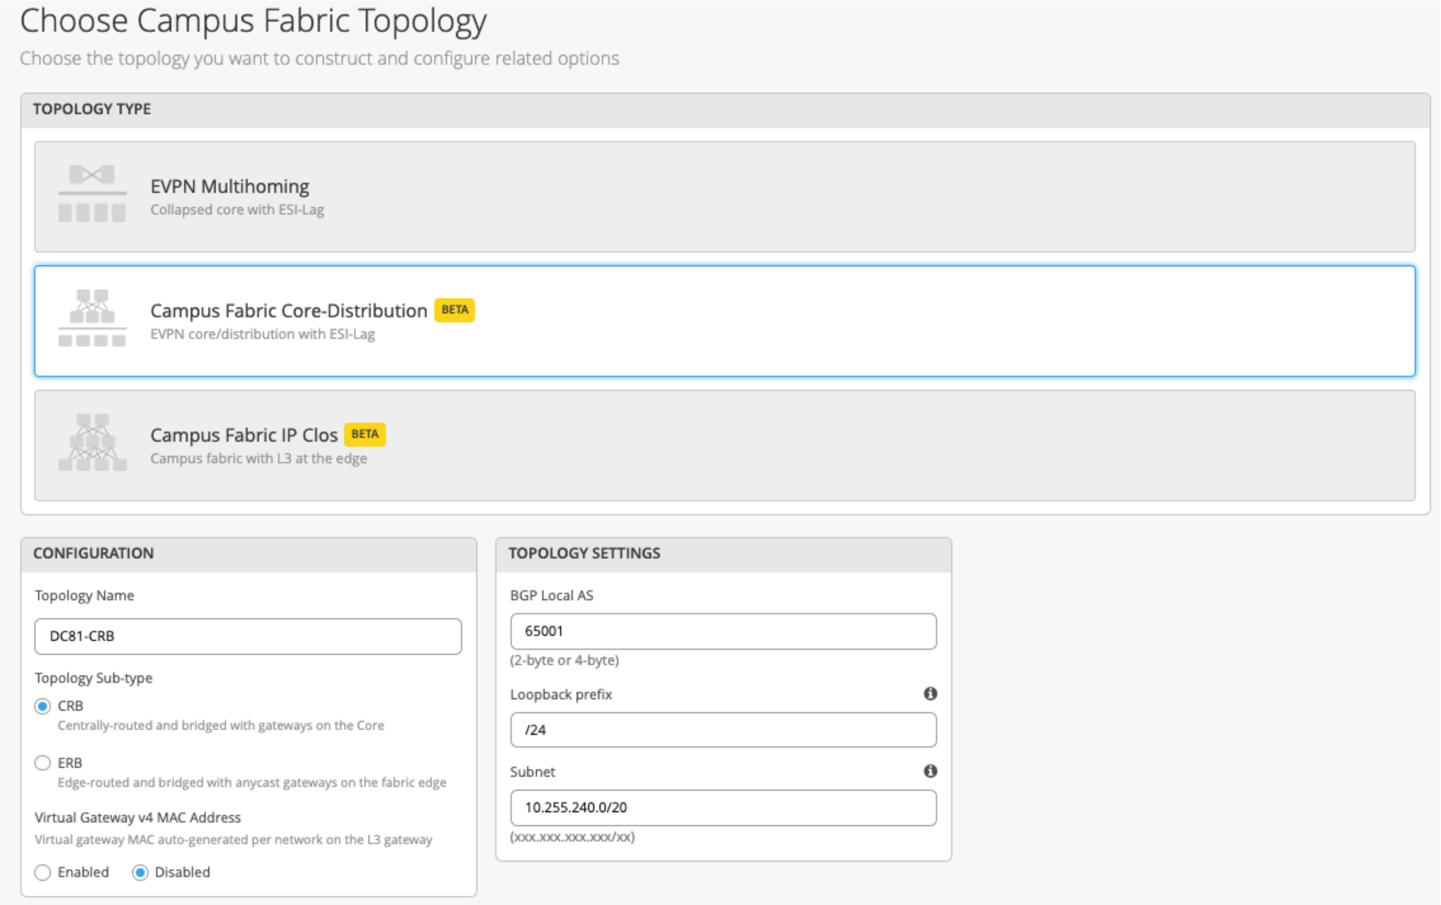

Choose Campus Fabric Topology

Choose the topology you want to construct and configure related options

TOPOLOGY TYPE

EVPN Multihoming
Collapsed core with ESI-Lag

Campus Fabric Core-Distribution BETA
EVPN core/distribution with ESI-Lag

Campus Fabric IP Clos BETA
Campus fabric with L3 at the edge

CONFIGURATION

Topology Name: DC81-CRB

Topology Sub-type

● CRB
Centrally-routed and bridged with gateways on the Core

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

Virtual Gateway v4 MAC Address
Virtual gateway MAC auto-generated per network on the L3 gateway

○ Enabled ● Disabled

TOPOLOGY SETTINGS

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

Loopback prefix: /24

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

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

Configuration

- Provide a name in accordance with company standards

Topology Sub-type

- CRB
- ERB

Virtual Gateway v4 MAC Address

- Only applicable to CRB
- Mist provides a unique MAC address for each L3 IRB

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

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

Topology Settings

- BGP Local AS: represents the starting point of private BGP AS numbers that will automatically be allocated per device. The user can use whatever private BGP AS number range suits their deployment, routing policy will be 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. The user can use whatever range suits their deployment. VXAN tunnelling through the use of a VTEP is associated with this address.
- Subnet: represents the range of IP addresses utilized for point-to-point links between devices. L The user can use whatever range suits their deployment. Mist breaks this subnet into /31 subnet addressing per link. This number can be modified to suit the specific deploymet scale. For example, /24 would provide up to 128 p2p /31 subnets.

Note: Juniper recommends default settings for all options unless it conflicts with other networks attached to the campus fabric. The point-point links between Core and Distribution layers utilize /31 addressing to conserve addresses.

Select campus fabric nodes

The user selects devices to participate at each Layer of the Campus Fabric Core Distribution CRB. Juniper recommends the user validate each device's presence in the site switch inventory prior to the creation of the Campus Fabric.

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

Services Block Router is where the Campus Fabric would interconnect external devices such as firewalls, routers, or critical devices such as DHCP and Radius servers (as an example). Devices to which external services connect to the Campus Fabric are known as Border Leafs. If the user wishes to connect these services/devices to the Campus Fabric Core Distribution CRB in a separate device or pair of devices, the Use Core as border option should be unchecked and the devices chosen by choosing the Select Switches option.

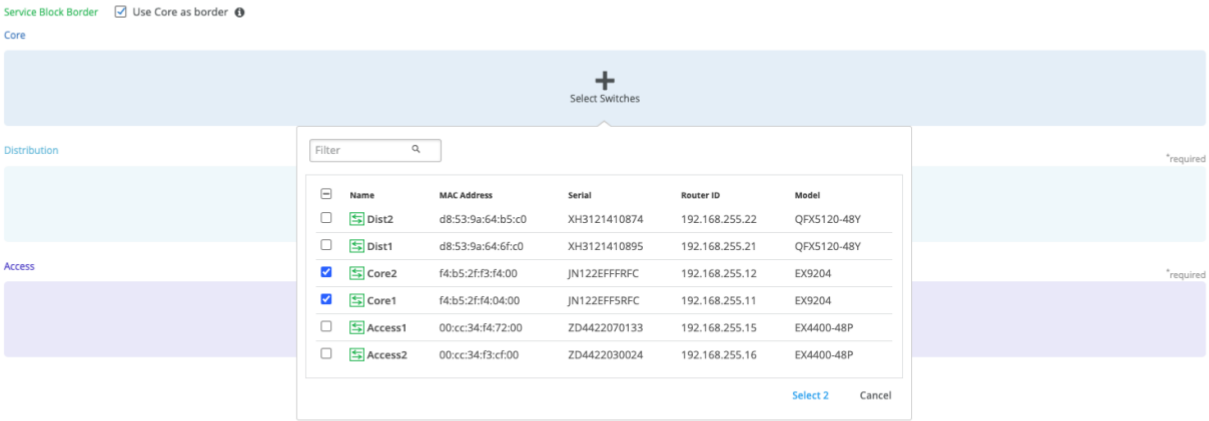
Service Block Border ☑ Use Core as border

Filter

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

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. Users who wish to combine this capability within the Core devices should select the User Core as border option.

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

The loopback addresses and router-ids should be the in same address space. The host-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 nethops. The loopback is also used as the router-id and will be used for overlay eBGP peering and VXLAN tunnel termination.

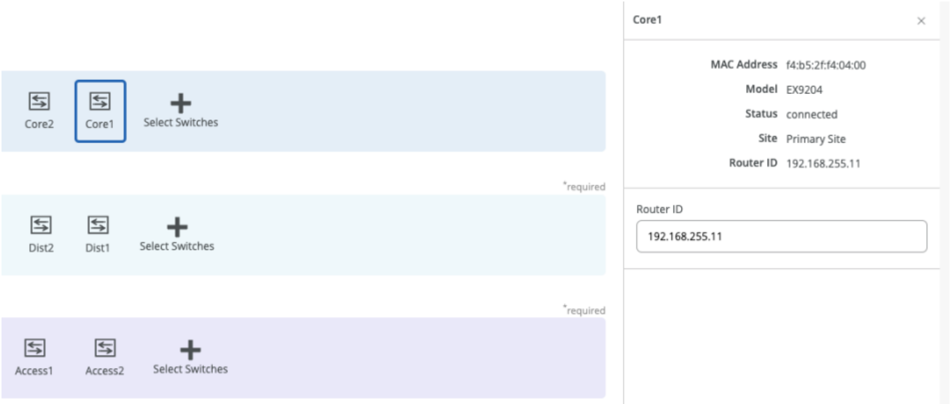

Core1

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

Router ID: 192.168.255.11

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

The loopack prefix is used for import /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 presents the user with default settings, which can be adapted as required.

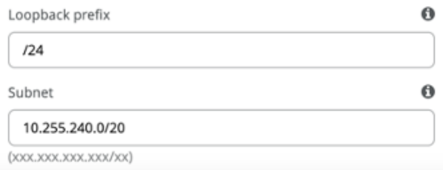
Loopback prefix: /24

Subnet: 10.255.240.0/20

(xxx.xxx.xxx.xxx/xx)

Configure Networks

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

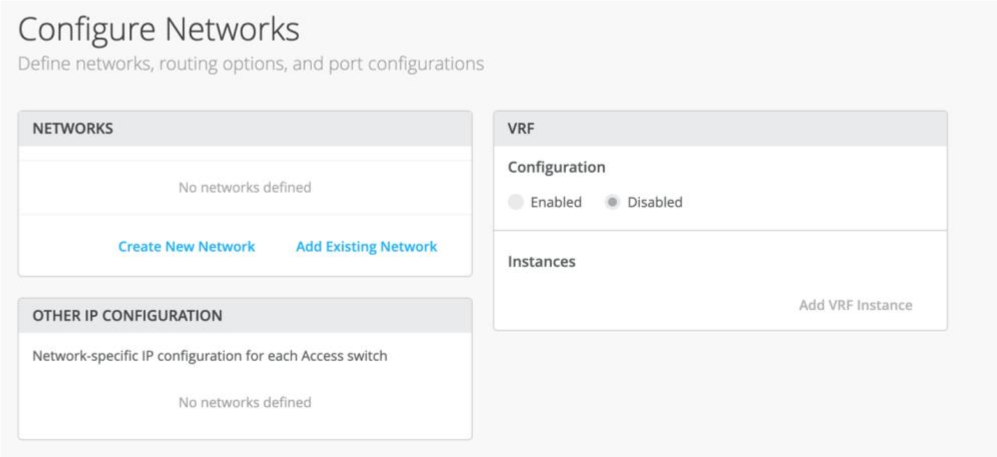

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

No networks defined

Create New Network   Add Existing Network

OTHER IP CONFIGURATION

Network-specific IP configuration for each Access switch

No networks defined

VRF

Configuration

○ Enabled ● Disabled

Instances

Add VRF Instance

Networks

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

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

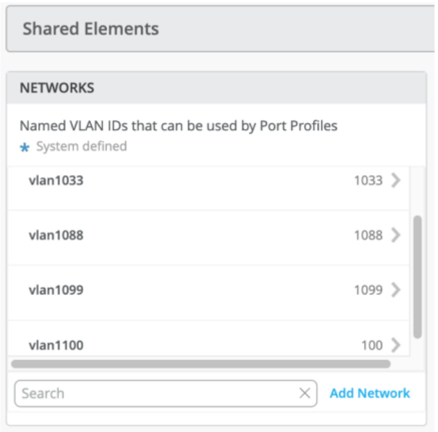

Shared Elements

NETWORKS

Named VLAN IDs that can be used by Port Profiles

★ System defined

vlan1033	1033
vlan1088	1088
vlan1099	1099
vlan1100	100

Search   Add Network

Back to the Campus Fabric build, the user selects the existing template that includes L2 VLAN information. All VLAN and IP information will be inherited from the template

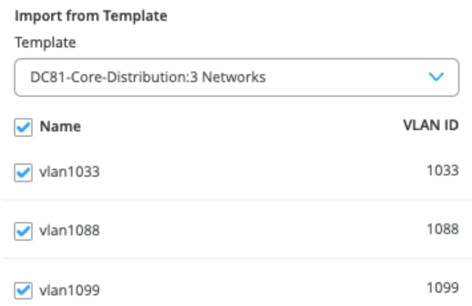

Import from Template

Template

DC81-Core-Distribution:3 Networks

☑ Name	VLAN ID
☑ vlan1033	1033
☑ vlan1088	1088
☑ vlan1099	1099

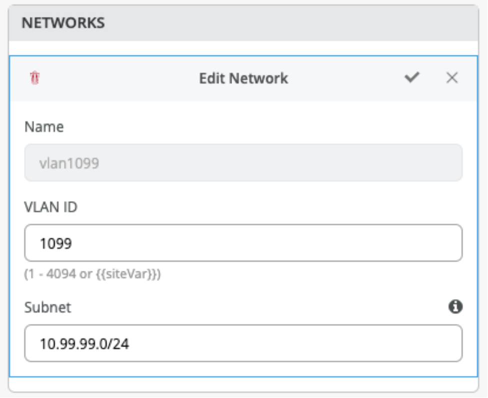

NETWORKS

Edit Network

Name

vlan1099

VLAN ID

1099

(1 - 4094 or {{siteVar}})

Subnet

10.99.99.0/24

Other IP Configuration

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

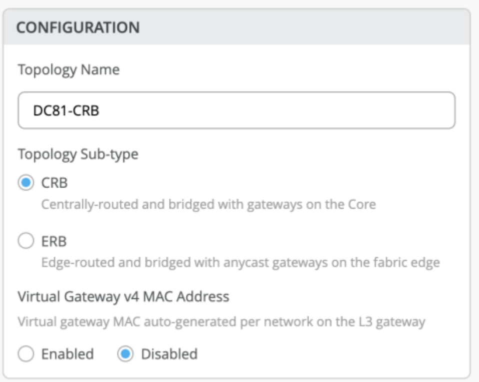

CONFIGURATION

Topology Name

DC81-CRB

Topology Sub-type

● CRB
Centrally-routed and bridged with gateways on the Core

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

Virtual Gateway v4 MAC Address
Virtual gateway MAC auto-generated per network on the L3 gateway

○ Enabled ● Disabled

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

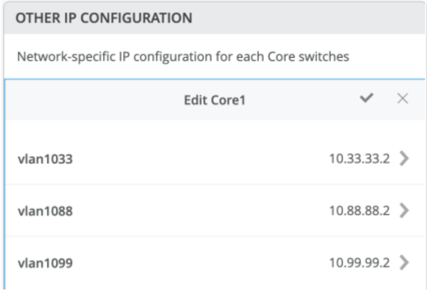

OTHER IP CONFIGURATION

Network-specific IP configuration for each Core switches

Edit Core1

vlan1033	10.33.33.2
vlan1088	10.88.88.2
vlan1099	10.99.99.2

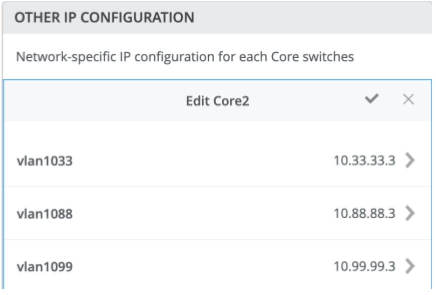

OTHER IP CONFIGURATION

Network-specific IP configuration for each Core switches

Edit Core2

vlan1033	10.33.33.3
vlan1088	10.88.88.3
vlan1099	10.99.99.3

By default, all VLANs are placed in the default VRF. The VRF option allows the user to group common VLANs into the same VRF or separate VRFs depending on traffic isolation requirements. This example includes 3 VRFs or routing instances: corp-it | developers | guest-wifi. Here, the user builds the first corp-it VRF and selects the pre-defined vlan 1099.

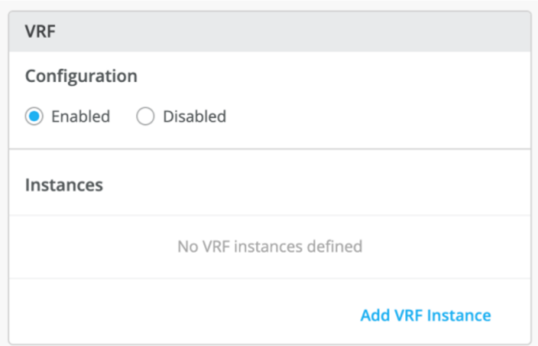

VRF

Configuration

● Enabled ○ Disabled

Instances

No VRF instances defined

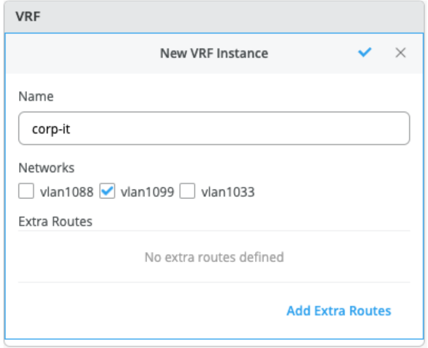

VRF

New VRF Instance

Name

corp-it

Networks

□ vlan1088 ☑ vlan1099 □ vlan1033

Extra Routes

No extra routes defined

By default, inter-VRF communications is not supported within the Campus Fabric. If inter-VRF communications is required, each VRF can include extra routes such as a Default Route that will instruct 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 handles inter-VRF routing. Figure 10. Topology

Notice the SRX participates in the VLANs defined within the Campus Fabric and is the gateway of last resort for all traffic leaving the subnet. The user selects the “Add Extra Routes” option to inform Mist to forward all traffic leaving 10.99.99.0/24 to utilize the next hop of the Juniper SRX firewall: 10.99.99.254

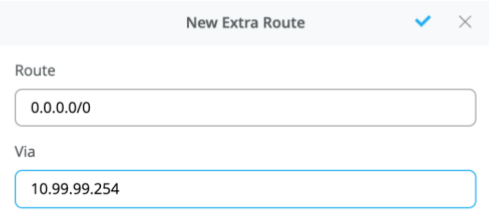

New Extra Route

Route: 0.0.0.0/0

Via: 10.99.99.254

The user creates 2 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

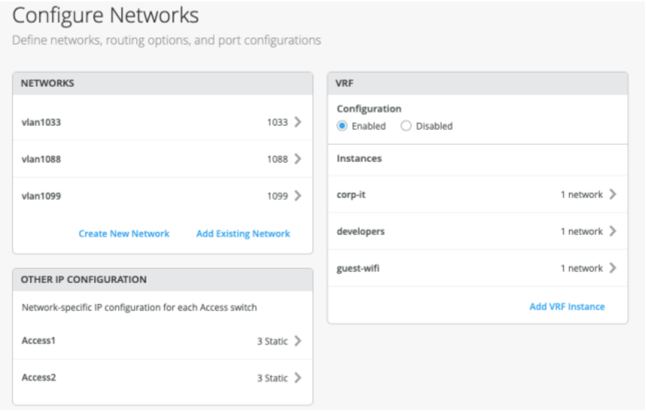

Configure Networks

Define networks, routing options, and port configurations

NETWORKS

vlan1033	1033
vlan1088	1088
vlan1099	1099

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

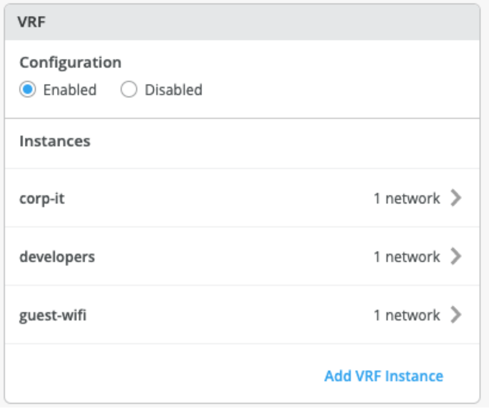

VRF

Configuration

● Enabled ○ Disabled

Instances

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

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

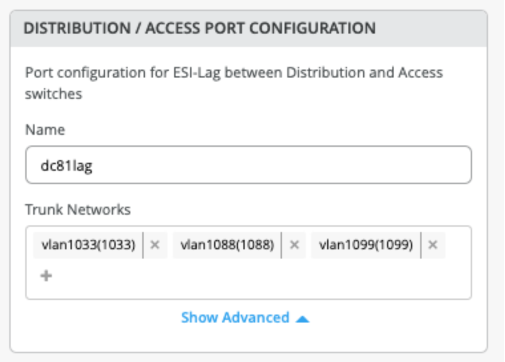

DISTRIBUTION / ACCESS PORT CONFIGURATION

Port configuration for ESI-Lag between Distribution and Access switches

Name: dc81lag

Trunk Networks: vlan1033(1033) | vlan1088(1088) | vlan1099(1099)

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

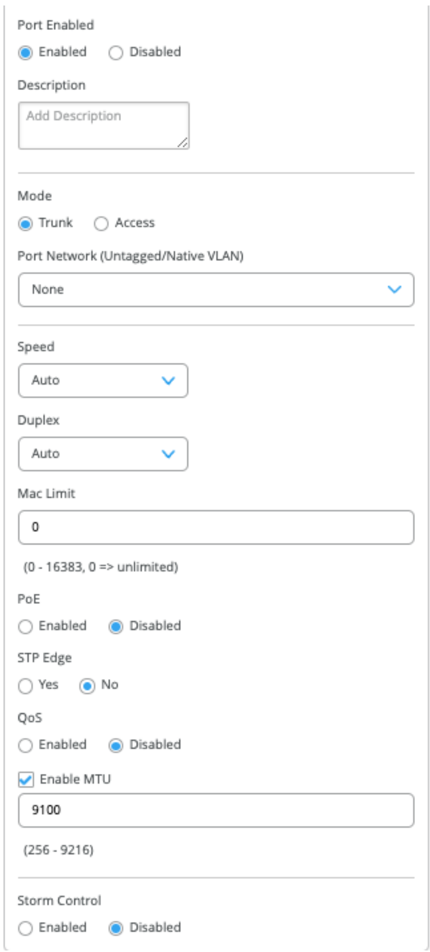

Port Enabled

● Enabled ○ Disabled

Description

Add Description

Mode

● Trunk ○ Access

Port Network (Untagged/Native VLAN)

None

Speed

Auto

Duplex

Auto

Mac Limit

0

(0 - 16383, 0 => unlimited)

PoE

○ Enabled ● Disabled

STP Edge

○ Yes ● No

QoS

○ Enabled ● Disabled

☑ Enable MTU

9100

(256 - 9216)

Storm Control

○ Enabled ● Disabled

Note: Juniper recommends default settings unless particular requirements are needed

Now that all VLANs are configured and assigned to each VRF, and the Distrbution/Access ESI-LAGs have been built, the user can move to the next step by clicking the Continue button at the upper right section of the Mist UI.

Configure campus fabric ports

The final step is the selection of physical ports between Core, Distribution and Access Switches.

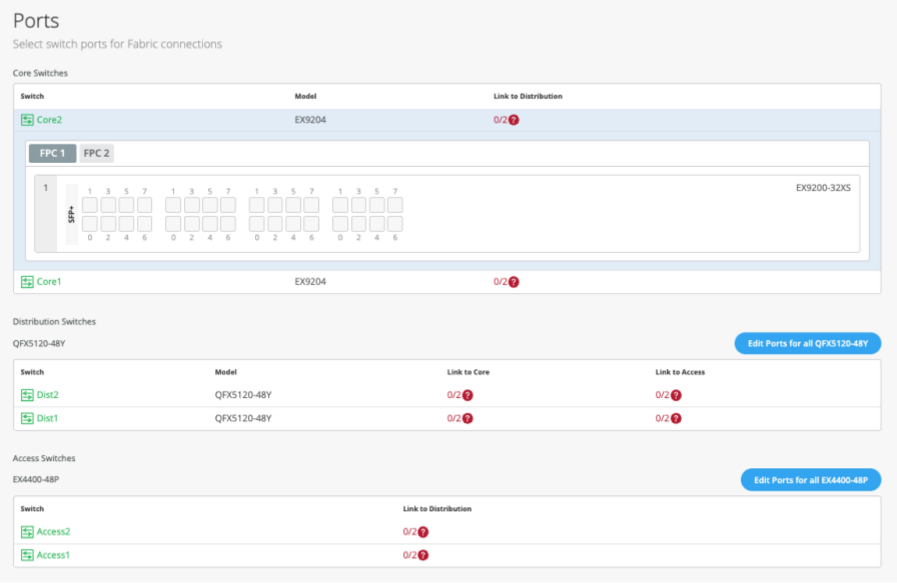

Ports

Select switch ports for Fabric connections

Core Switches

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

Distribution Switches

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

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

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

Core Switches

Core1:
Starting with Core1, the user selects xe-1/0/5 and xe-1/0/6 terminating on Distribution Switches 1 and 2 respectively.

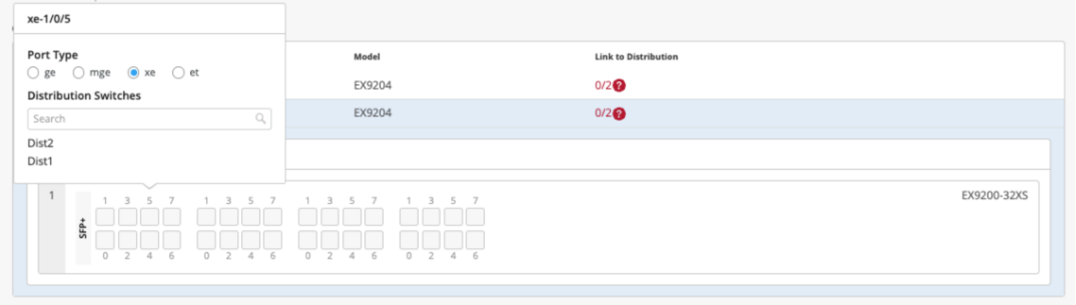

xe-1/0/5

Port Type ○ ge ○ mge ● xe ○ et

Distribution Switches

Search

Dist2

Dist1

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

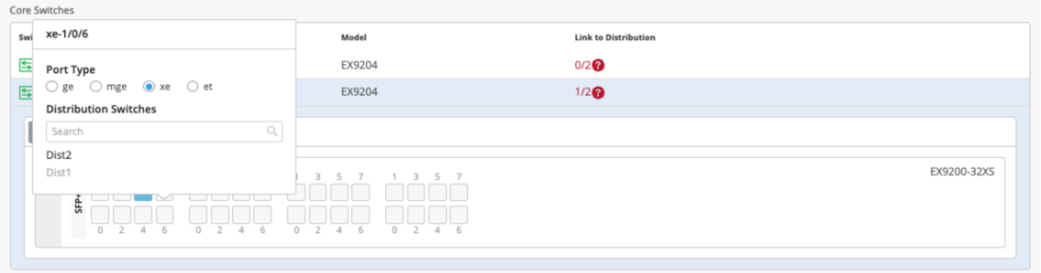

Core Switches

xe-1/0/6

Port Type ○ ge ○ mge ● xe ○ et

Distribution Switches

Search

Dist2

Dist1

Model	Link to Distribution
EX9204	0/2
EX9204	1/2

Core2:
On Core2, the user selects xe-1/0/4 and xe-1/0/5 terminating on Distribution Switches 1 and 2 respectively:

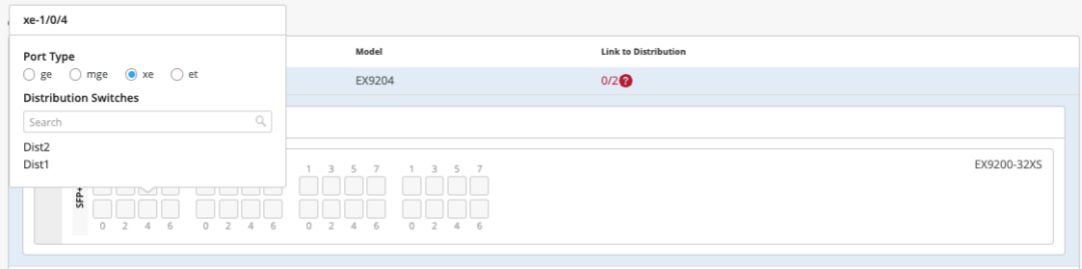

xe-1/0/4

Port Type ○ ge ○ mge ● xe ○ et

Distribution Switches

Search

Dist2

Dist1

Model	Link to Distribution
EX9204	0/2

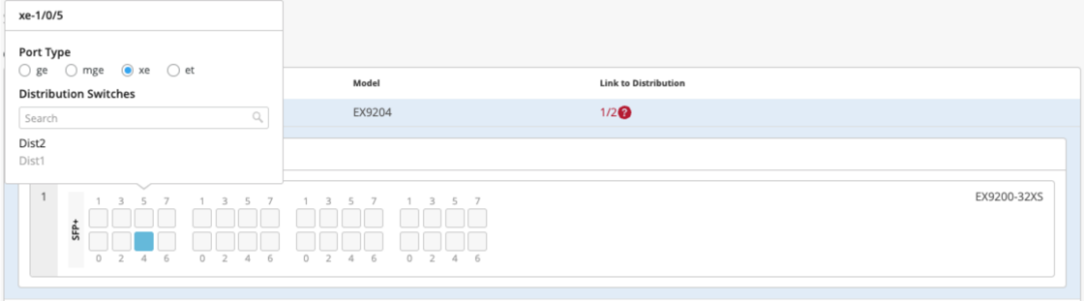

xe-1/0/5

Port Type ○ ge ○ mge ● xe ○ et

Distribution Switches

Search

Dist2

Dist1

Model	Link to Distribution
EX9204	1/2

Distribution Switches

Now moving on to the Distribution Switches, you will notice 2 interconnect options exist

- Link to Core
- Link to Access

Dist1:

The user selects Link to Core and choose xe-0/0/5 and xe-0/0/4 terminating on Core Switches 1 and 2 respectively.

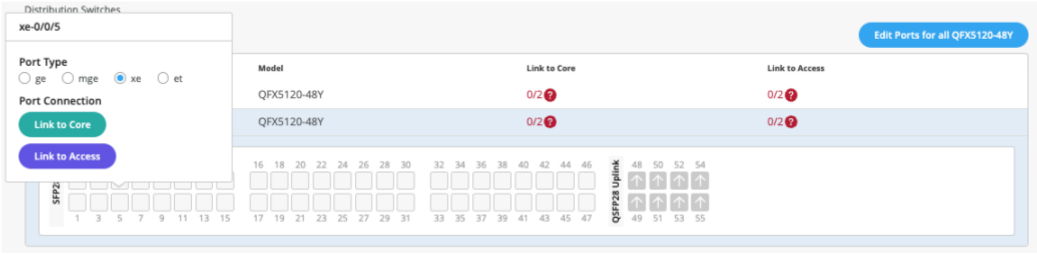

Distribution Switches

xe-0/0/5

Port Type

○ ge ○ mge ● xe ○ et

Port Connection

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

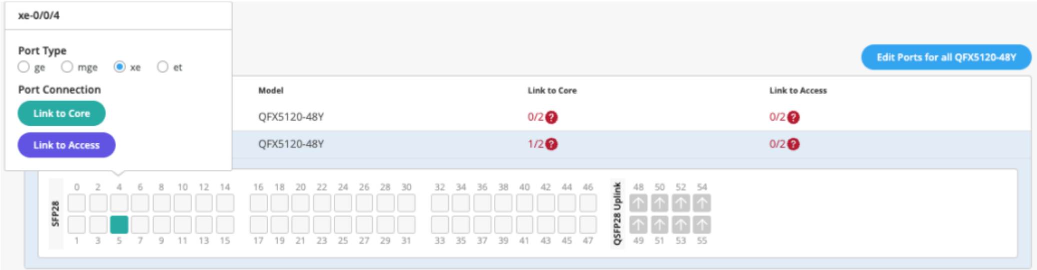

xe-0/0/4

Port Type

○ ge ○ mge ● xe ○ et

Port Connection

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

The user selects Link to Access and choose ge-0/0/36 and ge-0/0/37 terminating on Access Switches 1 and 2 respectively:

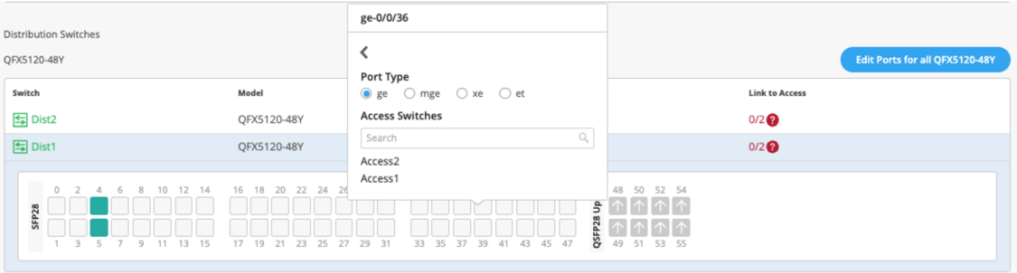

Distribution Switches

QFX5120-48Y

ge-0/0/36

Port Type

● ge ○ mge ○ xe ○ et

Access Switches

Access2

Access1

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

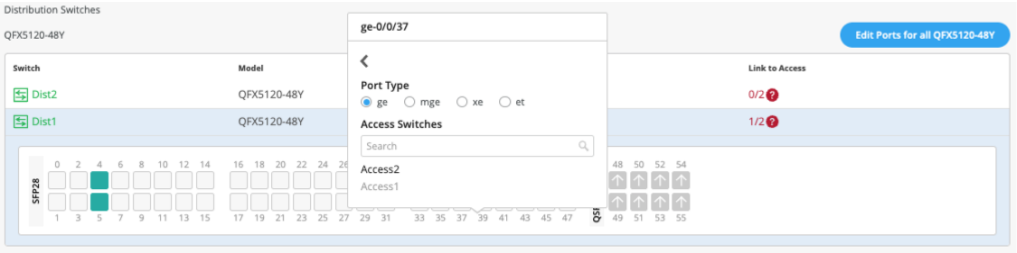

Distribution Switches

QFX5120-48Y

ge-0/0/37

Port Type

● ge ○ mge ○ xe ○ et

Access Switches

Access2

Access1

Switch	Model	Link to Access
Dist2	QFX5120-48Y	0/2
Dist1	QFX5120-48Y	1/2

Next, the user selects the following interconnects off **Dist2**:

- Link to Core
  - xe-0/0/6 – Core1
  - xe-0/0/5 – Core2

- Link to Access
  - ge-0/0/36 – Access2
  - ge-0/0/37 – Access1

Access Switches

The user only needs to know which interfaces will be used to interconnect with the Distribution switch but does not need to know the specific mapping. The system bundles all interfaces into a single ethernet bundle through the AE Index option. This greatly simplifies the physical port build for each access switch

Access1/2:

The user selects both uplinks and interface speed, while allowing Mist to define each AE index. In this case, uplinks ge-0/0/36/37 are selected as Links to Distribution on both Access switches and AE Index 0/1 (system default numbering) on Access1/2 respectively.

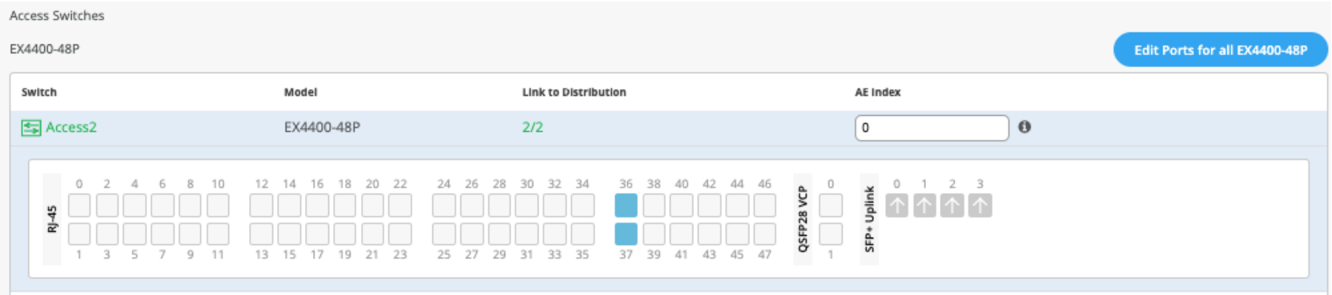

Access Switches

EX4400-48P

Edit Ports for all EX4400-48P

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

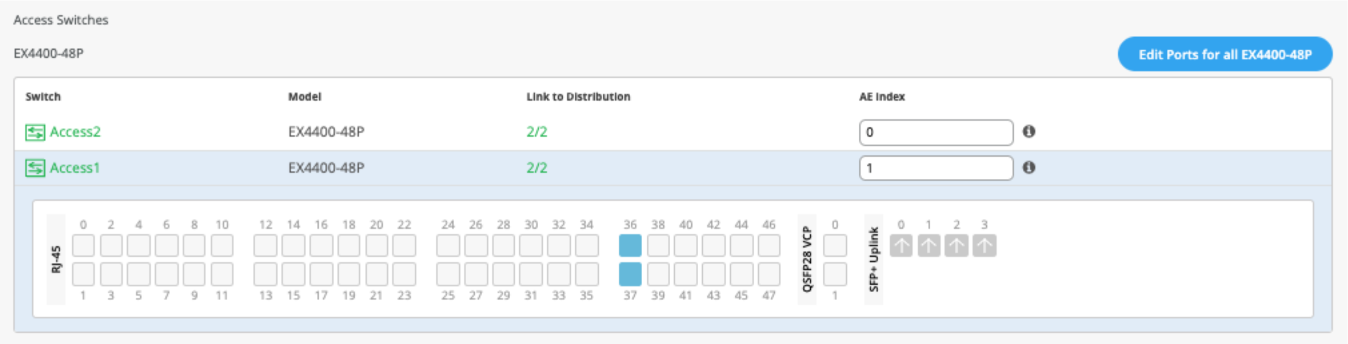

Access Switches

EX4400-48P

Edit Ports for all EX4400-48P

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

Once the user has completed selecting all requisite port combinations, they will select the Continue button at the upper right-hand corner of the Mist UI.

Campus Fabric Configuration Confirmation

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

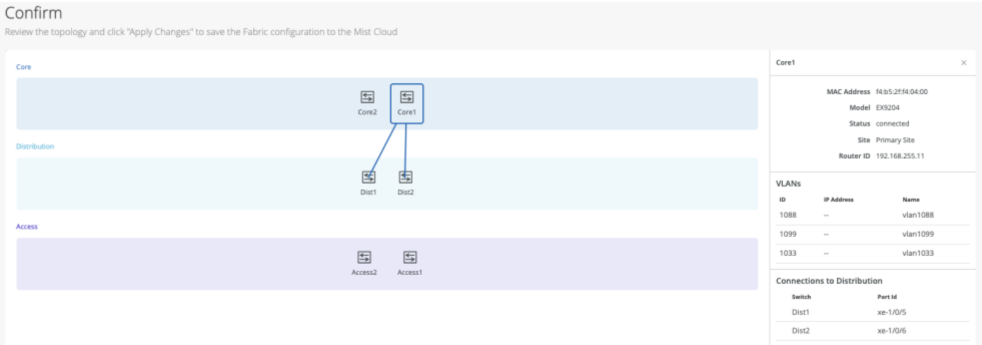

Confirm

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

Core1

MAC Address 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	Port Id
Dist1	xe-1/0/5
Dist2	xe-1/0/6

Once the user has completed verification, they will select the Apply Changes option at the upper right-hand corner of the Mist UI

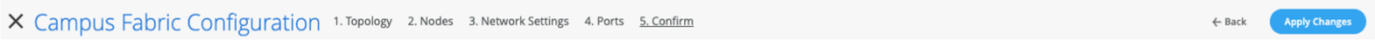

The user is presented a second stage confirmation, confirm to create the fabric.

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

- Mist builds the point-to-point interfaces between Distribution and Core devices with IP addresses chosen from the range presented at the onset of the build.
- Each device is configured with a loopback address from the range presented at the onset of the build.
- eBGP is provisioned at each device with unique BGP autonomous system numbers. The primary goal of the underlay is to leverage ECMP for load balancing traffic on a per packet level for device loopback reachability. The primary goal of the eBGP overlay is support of customer traffic using EVPN-VXLAN.

- IP addressing of each L3 gateway IRB located on Core1 and Core2

- IP addressing of each lo0.0 loopback
- Configuration of routing policies for underlay and overlay connectivity
- Optimized MTU settings for p2p underlay, L3 IRB, and ESI-LAG bundles
- VXLAN to VLAN mapping using VNI addresses that are automatically assigned

- VRF creation of corp-it, developers, and guest-wifi and VLAN associated with each VRF
- VXLAN tunnelling creation between Distribution devices and Distribution-Core devices (in support of the northbound SRX firewall that will be 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

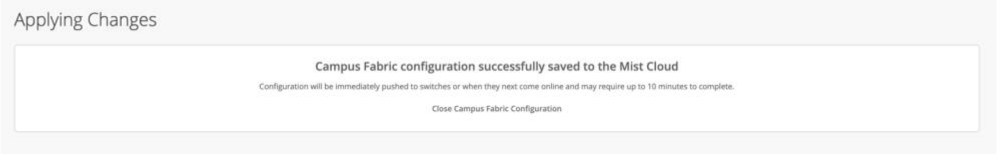

Applying Changes

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

Closing this section provides the user with a summary of the newly created Campus Fabric Core Distribution CRB

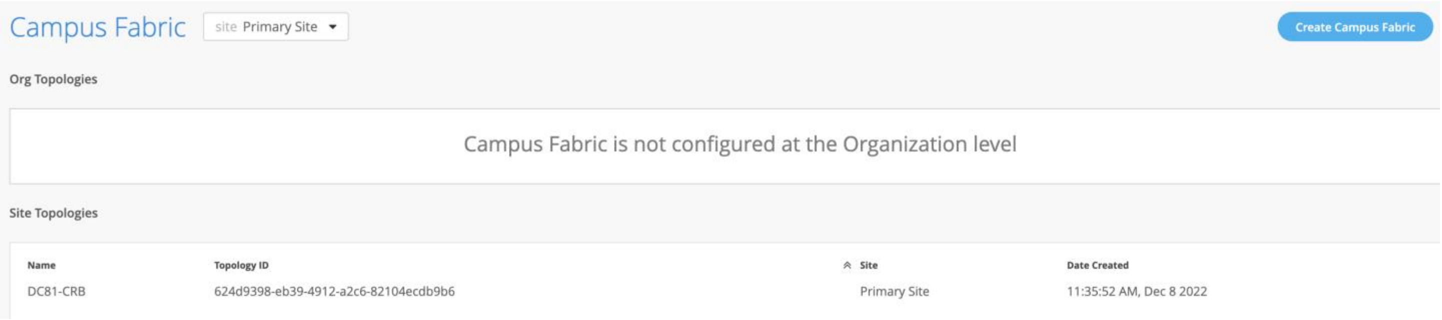

Campus Fabric

site Primary Site

Create Campus Fabric

Org Topologies

Campus Fabric is not configured at the Organization level

Site Topologies

Name	Topology ID	Site	Date Created
DC81-CRB	624d9398-eb39-4912-a2c6-82104ecdb9b6	Primary Site	11:35:52 AM, Dec 8 2022

Juniper Mist Wired Assurance provides the user with the ability to 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, the user can download the connection table:

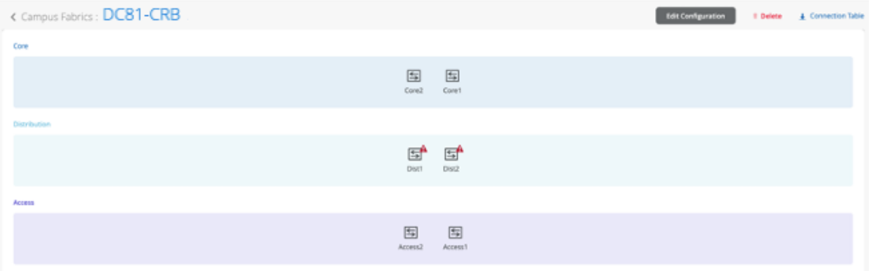

Disclaimer: The RED designations depicted on the Distiribution Switches will be fixed with a future software update.

Connection Table spreadsheet:

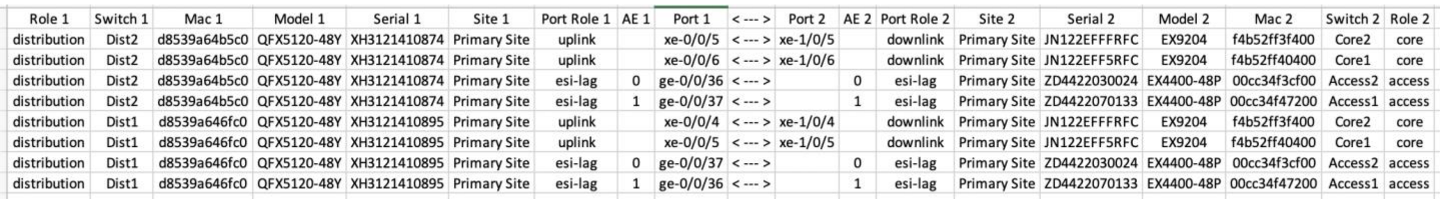

Role 1	Switch 1	Mac 1	Model 1	Serial 1	Site 1	Port Role 1	AE 1	Port 1	< --- >	Port 2	AE 2	Port Role 2	Site 2	Serial 2	Model 2	Mac 2	Switch 2	Role 2
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/5	< --- >	xe-1/0/5		downlink	Primary Site	JN122EFFFRFC	EX9204	f4b52ff3f400	Core2	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	uplink		xe-0/0/6	< --- >	xe-1/0/6		downlink	Primary Site	JN122EFF5RFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag	0	ge-0/0/36	< --- >		0	esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist2	d8539a64b5c0	QFX5120-48Y	XH3121410874	Primary Site	esi-lag	1	ge-0/0/37	< --- >		1	esi-lag	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	uplink		xe-0/0/4	< --- >	xe-1/0/4		downlink	Primary Site	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	JN122EFF5RFC	EX9204	f4b52ff40400	Core1	core
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag	0	ge-0/0/37	< --- >		0	esi-lag	Primary Site	ZD4422030024	EX4400-48P	00cc34f3cf00	Access2	access
distribution	Dist1	d8539a646fc0	QFX5120-48Y	XH3121410895	Primary Site	esi-lag	1	ge-0/0/36	< --- >		1	esi-lag	Primary Site	ZD4422070133	EX4400-48P	00cc34f47200	Access1	access

Apply VLANs to Access ports

As previously discussed, Mist provides the ability to templatize well known services such as Radius, NTP, DNS, etc 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. Figure 10. Topology

It is also noteworthy that Mist Access Points connect to the same port on Access1/2 allowing the Switch Template to be customized with this configuration. For example; the following found under the Switch Template option is customized to associate each switch with its role: Core, Distribution, and Access. Further, all Access switches (defined by Model EX4400 as an example) associated the AP port profile with ge-0/0/16 without needing to configure each independent switch.

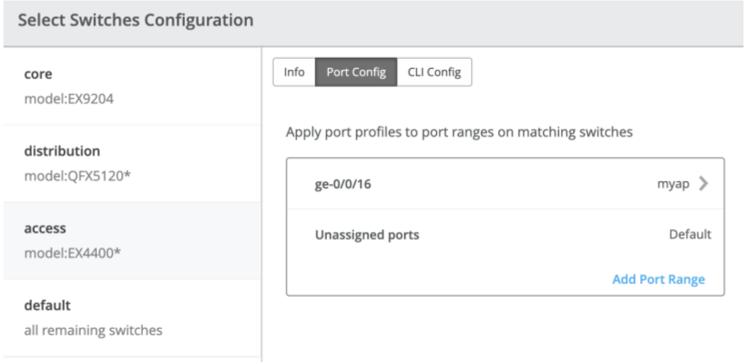

Select Switches Configuration

core
model:EX9204

distribution
model:QFX5120*

access
model:EX4400*

default
all remaining switches

Apply port profiles to port ranges on matching switches

ge-0/0/16	myap
Unassigned ports	Default

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

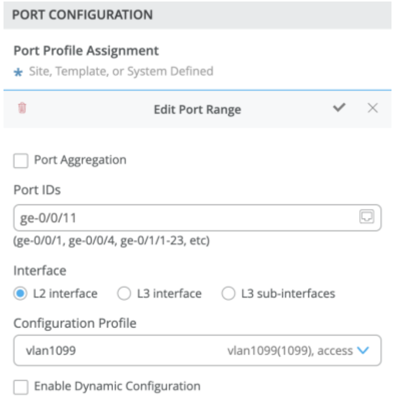

PORT CONFIGURATION

Port Profile Assignment
★ Site, Template, or System Defined

Edit Port Range

□ Port Aggregation

Port IDs
ge-0/0/11
(ge-0/0/1, ge-0/0/4, ge-0/1/1-23, etc)

Interface
● L2 interface ○ L3 interface ○ L3 sub-interfaces

Configuration Profile
vlan1099 vlan1099(1099), access

□ Enable Dynamic Configuration

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

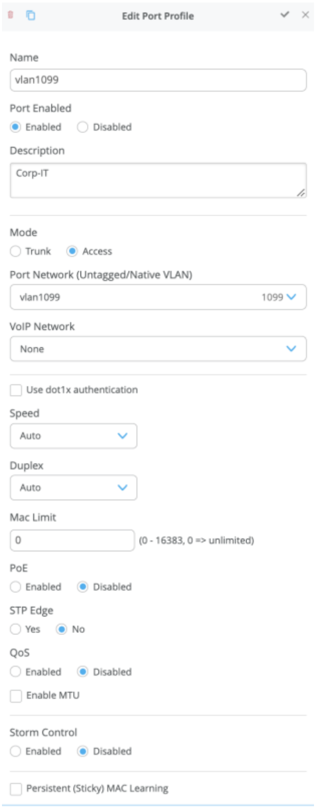

Edit Port Profile

Name

vlan1099

Port Enabled

● Enabled ○ Disabled

Description

Corp-IT

Mode

○ Trunk ● Access

Port Network (Untagged/Native VLAN)

vlan1099 1099

VoIP Network

None

□ Use dot1x authentication

Speed

Auto

Duplex

Auto

Mac Limit

0 (0 - 16383, 0 => unlimited)

PoE

○ Enabled ● Disabled

STP Edge

○ Yes ● No

QoS

○ Enabled ● Disabled

□ Enable MTU

Storm Control

○ Enabled ● Disabled

□ Persistent (Sticky) MAC Learning

VERIFICATION

Verification of the Campus Fabric Core Distribution CRB deployment.Figure 10. Topology Currently there are two desktops that can be used to validate the fabric. Let's take a quick look to see if Desktop1 can connect internally and externally.

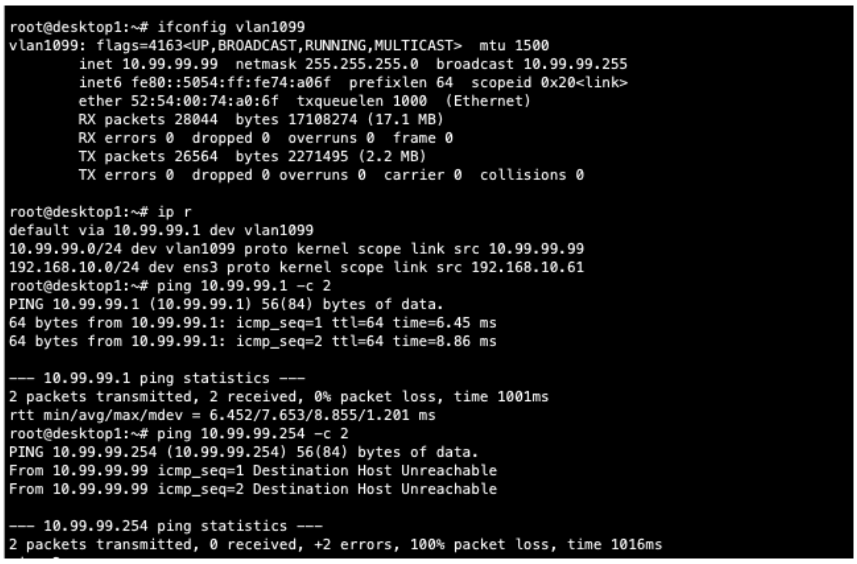

```
root@desktop1:~# ifconfig vlan1099
vlan1099: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.99.99.99  netmask 255.255.255.0  broadcast 10.99.99.255
        inet6 fe80::5054:ff:fe74:a06f  prefixlen 64  scopeid 0x20<link>
        ether 52:54:00:74:a0:6f  txqueuelen 1000  (Ethernet)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

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

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

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

Validation steps

- Confirmed local IP address, vlan and default gateway were configured on Desktop1
- Can ping default gateway – that tells us 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.

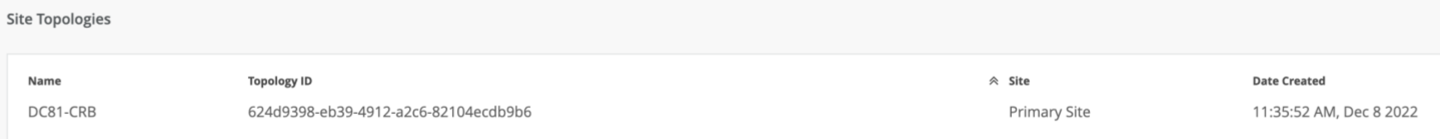

Site Topologies

Name	Topology ID	Site	Date Created
DC81-CRB	624d9398-eb39-4912-a2c6-82104ecdb9b6	Primary Site	11:35:52 AM, Dec 8 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/Receive traffic on a link-by-link basis
- Telemetry, such as lldp, from each device that verifies the physical build

BGP Underlay

Purpose
Verifying the state of eBGP between Core and Disribution layers is essential for EVPN VXLAN to operate as expected. This network of point-to-point links between each layer supports:

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

Without requiring verification at each layer, the focus can be on Core1/2 and their eBGP relationships with Dist1/2. If both Core switches have "established" eBGP peering sessions with both Dist swiches, the user can move to the next phase of verification.

Action

Verify that BGP sessions are established from Core devices with Dist devices to insure 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 Secure CRT or Putty.

Verification of BGP peering

Core1:
Remote Shell can be accessed via the bottom right of the Campus Fabric, from the switch view or via SSH.

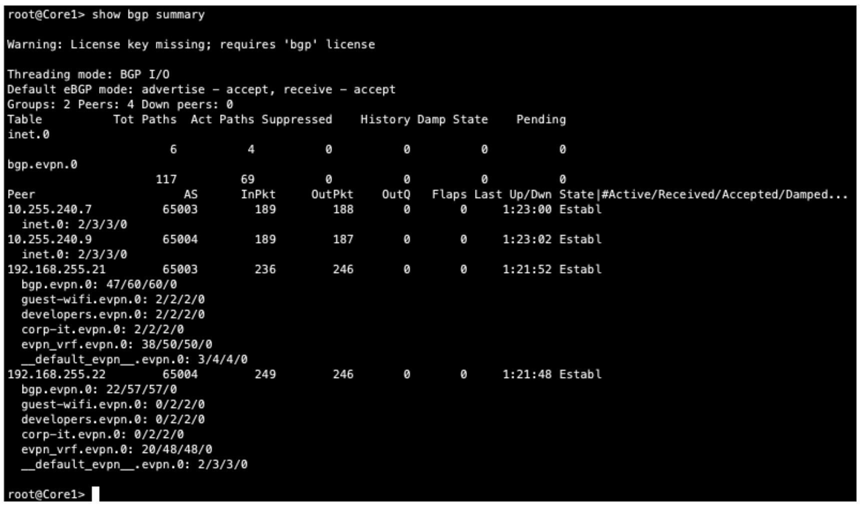

```
root@Core1> show bgp summary

Warning: License key missing; requires 'bgp' license

Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       6          4          0          0          0          0
bgp.evpn.0
                     117         69          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/Accepted/Damped...
10.255.240.7          65003        189        188       0       0     1:23:00 Establ
  inet.0: 2/3/3/0
10.255.240.9          65004        189        187       0       0     1:23:02 Establ
  inet.0: 2/3/3/0
192.168.255.21        65003        236        246       0       0     1:21:52 Establ
  bgp.evpn.0: 47/60/60/0
  guest-wifi.evpn.0: 2/2/2/0
  developers.evpn.0: 2/2/2/0
  corp-it.evpn.0: 2/2/2/0
  evpn_vrf.evpn.0: 38/50/50/0
  __default_evpn__.evpn.0: 3/4/4/0
192.168.255.22        65004        249        246       0       0     1:21:48 Establ
  bgp.evpn.0: 22/57/57/0
  guest-wifi.evpn.0: 0/2/2/0
  developers.evpn.0: 0/2/2/0
  corp-it.evpn.0: 0/2/2/0
  evpn_vrf.evpn.0: 20/48/48/0
  __default_evpn__.evpn.0: 2/3/3/0

root@Core1>
```

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

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

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

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

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

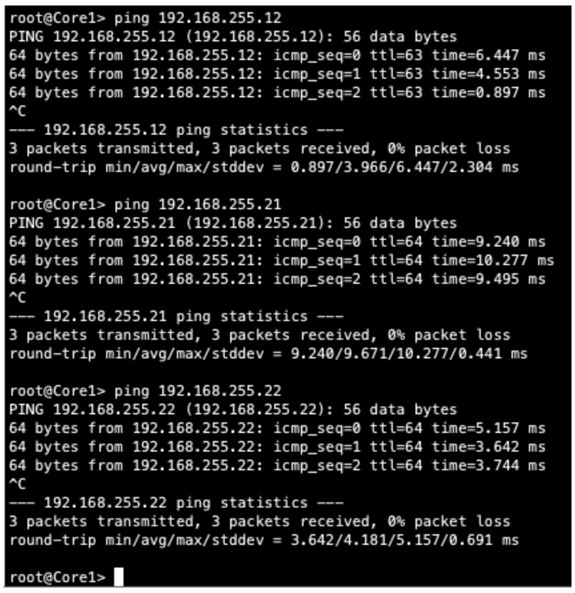

```
root@Core1> ping 192.168.255.12
PING 192.168.255.12 (192.168.255.12): 56 data bytes
64 bytes from 192.168.255.12: icmp_seq=0 ttl=63 time=6.447 ms
64 bytes from 192.168.255.12: icmp_seq=1 ttl=63 time=4.553 ms
64 bytes from 192.168.255.12: icmp_seq=2 ttl=63 time=0.897 ms
^C
--- 192.168.255.12 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.897/3.966/6.447/2.304 ms

root@Core1> ping 192.168.255.21
PING 192.168.255.21 (192.168.255.21): 56 data bytes
64 bytes from 192.168.255.21: icmp_seq=0 ttl=64 time=9.240 ms
64 bytes from 192.168.255.21: icmp_seq=1 ttl=64 time=10.277 ms
64 bytes from 192.168.255.21: icmp_seq=2 ttl=64 time=9.495 ms
^C
--- 192.168.255.21 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 9.240/9.671/10.277/0.441 ms

root@Core1> ping 192.168.255.22
PING 192.168.255.22 (192.168.255.22): 56 data bytes
64 bytes from 192.168.255.22: icmp_seq=0 ttl=64 time=5.157 ms
64 bytes from 192.168.255.22: icmp_seq=1 ttl=64 time=3.642 ms
64 bytes from 192.168.255.22: icmp_seq=2 ttl=64 time=3.744 ms
^C
--- 192.168.255.22 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.642/4.181/5.157/0.691 ms

root@Core1>
```

Note: eBGP sessions are established between Core-Distribution layers in the Campus Fabric

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

Dist1: ECMP Loopback reachability to Dist2 through Dist1/2

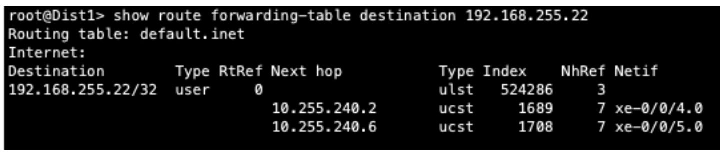

```
root@Dist1> show route forwarding-table destination 192.168.255.22
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.255.22/32  user     0                    ulst   524286     3
                              10.255.240.2       ucst     1689     7 xe-0/0/4.0
                              10.255.240.6       ucst     1708     7 xe-0/0/5.0
```

Dist2: ECMP Loopback reachability with Dist1 through Core1/2

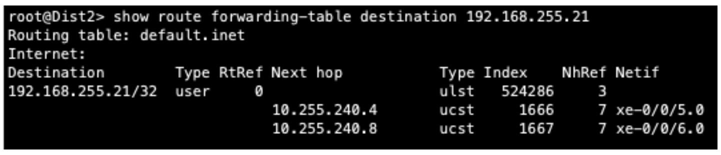

```
root@Dist2> show route forwarding-table destination 192.168.255.21
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.255.21/32  user     0                    ulst   524286     3
                              10.255.240.4       ucst     1666     7 xe-0/0/5.0
                              10.255.240.8       ucst     1667     7 xe-0/0/6.0
```

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

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

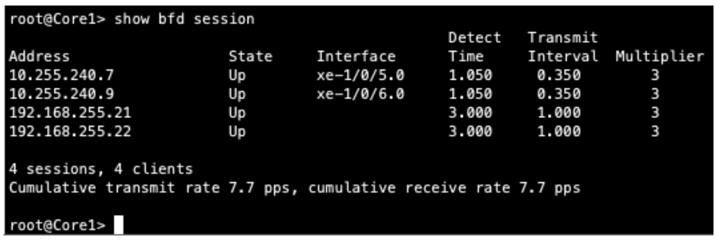

```
root@Core1> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.255.240.7             Up        xe-1/0/5.0     1.050     0.350        3
10.255.240.9             Up        xe-1/0/6.0     1.050     0.350        3
192.168.255.21           Up                       3.000     1.000        3
192.168.255.22           Up                       3.000     1.000        3

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

root@Core1> 
```

Meaning: At this point BGP Underlay and Overlay is operational through the verification of eBGP between corresponding layers of the Campus Fabric and that loopback routes are established between Core and Distribution layers

EVPN VXLAN verification between Core and Distribution switches

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

Verification of the EVPN Database on both Core switches

Core1:

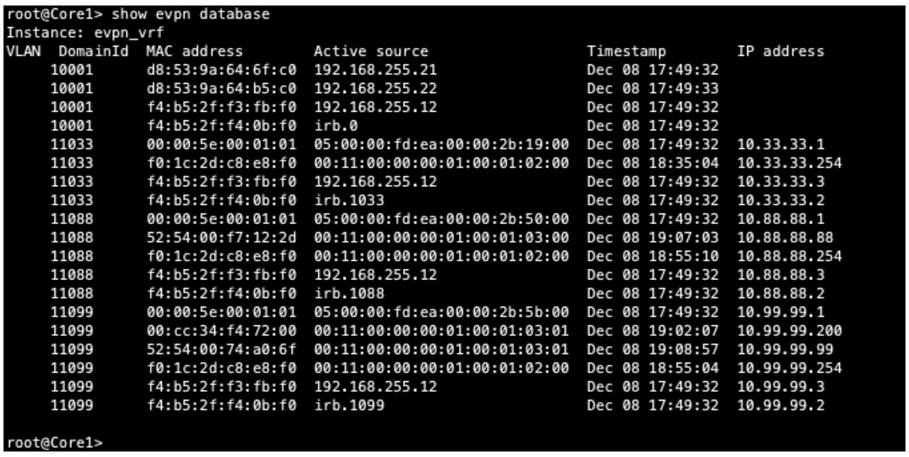

```
root@Core1> show evpn database
Instance: evpn_vrf
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      d8:53:9a:64:6f:c0  192.168.255.21                 Dec 08 17:49:32
     10001      d8:53:9a:64:b5:c0  192.168.255.22                 Dec 08 17:49:33
     10001      f4:b5:2f:f3:fb:f0  192.168.255.12                 Dec 08 17:49:32
     10001      f4:b5:2f:f4:0b:f0  irb.0                          Dec 08 17:49:32
     11033      00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:19:00  Dec 08 17:49:32  10.33.33.1
     11033      f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00  Dec 08 18:35:04  10.33.33.254
     11033      f4:b5:2f:f3:fb:f0  192.168.255.12                 Dec 08 17:49:32  10.33.33.3
     11033      f4:b5:2f:f4:0b:f0  irb.1033                       Dec 08 17:49:32  10.33.33.2
     11088      00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:50:00  Dec 08 17:49:32  10.88.88.1
     11088      52:54:00:f7:12:2d  00:11:00:00:00:01:00:01:03:00  Dec 08 19:07:03  10.88.88.88
     11088      f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00  Dec 08 18:55:10  10.88.88.254
     11088      f4:b5:2f:f3:fb:f0  192.168.255.12                 Dec 08 17:49:32  10.88.88.3
     11088      f4:b5:2f:f4:0b:f0  irb.1088                       Dec 08 17:49:32  10.88.88.2
     11099      00:00:5e:00:01:01  05:00:00:fd:ea:00:00:2b:5b:00  Dec 08 17:49:32  10.99.99.1
     11099      00:cc:34:f4:72:00  00:11:00:00:00:01:00:01:03:01  Dec 08 19:02:07  10.99.99.200
     11099      52:54:00:74:a0:6f  00:11:00:00:00:01:00:01:03:01  Dec 08 19:08:57  10.99.99.99
     11099      f0:1c:2d:c8:e8:f0  00:11:00:00:00:01:00:01:02:00  Dec 08 18:55:04  10.99.99.254
     11099      f4:b5:2f:f3:fb:f0  192.168.255.12                 Dec 08 17:49:32  10.99.99.3
     11099      f4:b5:2f:f4:0b:f0  irb.1099                       Dec 08 17:49:32  10.99.99.2

root@Core1>
```

Core2:

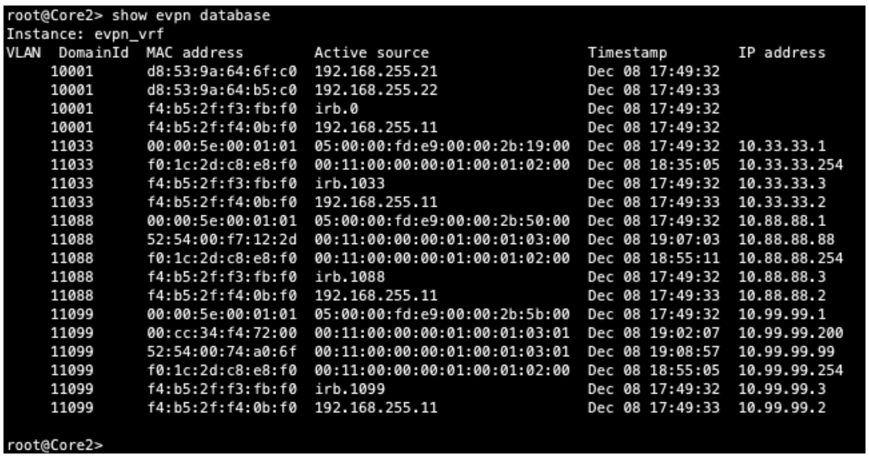

```
root@Core2> show evpn database
Instance: evpn_vrf
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      d8:53:9a:64:6f:c0  192.168.255.21                 Dec 08 17:49:32
     10001      d8:53:9a:64:b5:c0  192.168.255.22                 Dec 08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

root@Core2>
```

Both Core switches have identical EVPN databases which is expected. Notice the entries for desktop1 (10.99.99.99) and desktop2 (10.88.88.88) present in each Core switch. These entries are learned through the Campus Fabric from the ESI LAGs off DIst1/2.

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

Dist

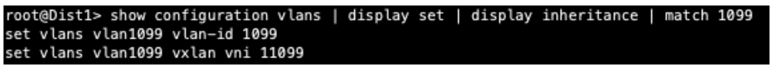

```
root@Dist1> show configuration vlans | display set | display inheritance | match 1099
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 vxlan vni 11099
```

Core

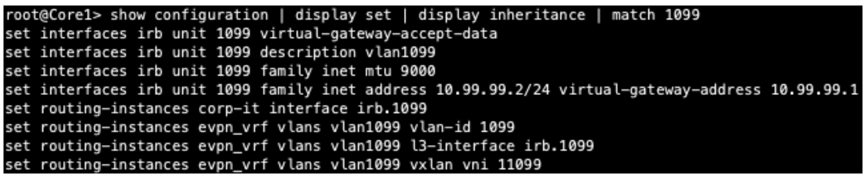

```
root@Core1> show configuration | display set | display inheritance | match 1099
set interfaces irb unit 1099 virtual-gateway-accept-data
set interfaces irb unit 1099 description vlan1099
set interfaces irb unit 1099 family inet mtu 9000
set interfaces irb unit 1099 family inet address 10.99.99.2/24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

Verification of VXLAN tunnelling between Dist and Core switches

Dist1:

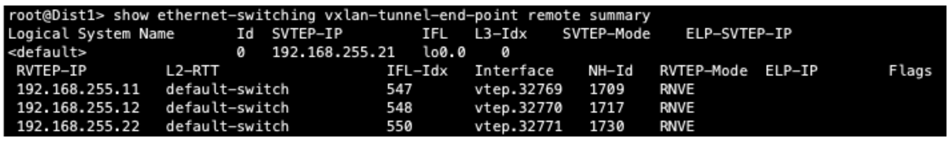

```
root@Dist1> show ethernet-switching vxlan-tunnel-end-point remote summary
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   192.168.255.21   lo0.0    0
 RVTEP-IP         L2-RTT                   IFL-Idx   Interface    NH-Id   RVTEP-Mode  ELP-IP        Flags
 192.168.255.11   default-switch           547       vtep.32769   1709    RNVE
 192.168.255.12   default-switch           548       vtep.32770   1717    RNVE
 192.168.255.22   default-switch           550       vtep.32771   1730    RNVE
```

Core1:

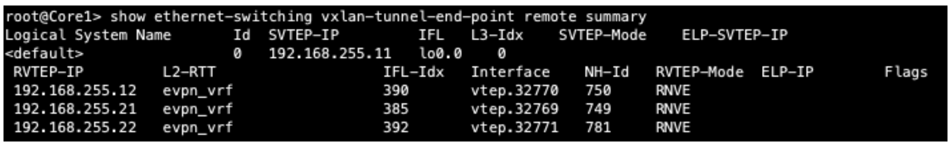

```
root@Core1> show ethernet-switching vxlan-tunnel-end-point remote summary
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
<default>                 0   192.168.255.11   lo0.0    0
 RVTEP-IP         L2-RTT                   IFL-Idx   Interface    NH-Id   RVTEP-Mode  ELP-IP        Flags
 192.168.255.12   evpn_vrf                 390       vtep.32770   750     RNVE
 192.168.255.21   evpn_vrf                 385       vtep.32769   749     RNVE
 192.168.255.22   evpn_vrf                 392       vtep.32771   781     RNVE
```

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

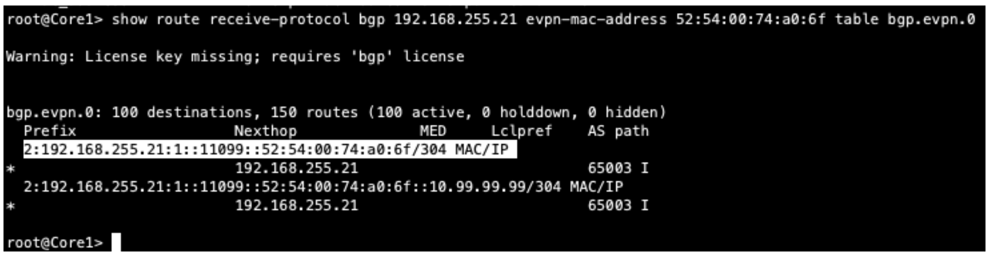

```
root@Core1> show route receive-protocol bgp 192.168.255.21 evpn-mac-address 52:54:00:74:a0:6f table bgp.evpn.0

Warning: License key missing; requires 'bgp' license


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

root@Core1>
```

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

We next verify the MAC address mapped to the correct VTEP interface on Core1:

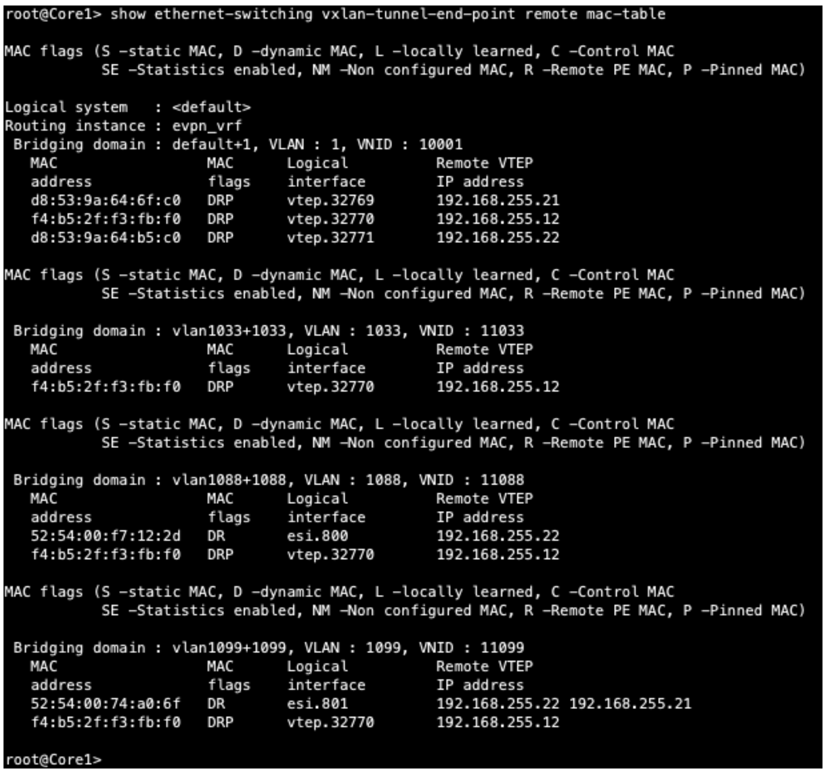

```
root@Core1> show ethernet-switching vxlan-tunnel-end-point remote mac-table

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

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

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

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

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

 Bridging domain : vlan1088+1088, VLAN : 1088, VNID : 11088
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
   52:54:00:f7:12:2d   DR       esi.800          192.168.255.22
   f4:b5:2f:f3:fb:f0   DRP      vtep.32770       192.168.255.12

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

 Bridging domain : vlan1099+1099, VLAN : 1099, VNID : 11099
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
   52:54:00:74:a0:6f   DR       esi.801          192.168.255.22 192.168.255.21
   f4:b5:2f:f3:fb:f0   DRP      vtep.32770       192.168.255.12

root@Core1>
```

Notice Desktop1's MAC address (52:54:00:74:a0:6f) is learned through both Dist1/2 switches

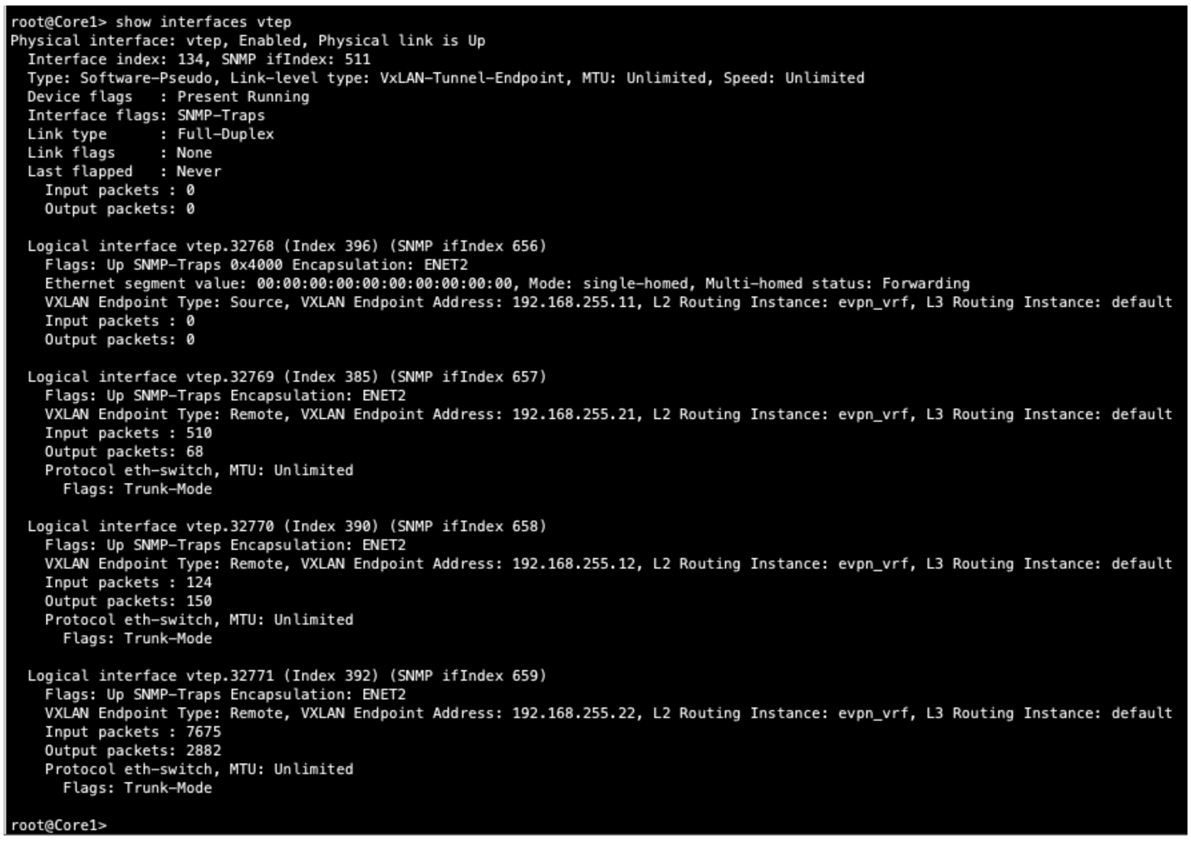

```
root@Core1> show interfaces vtep
Physical interface: vtep, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 511
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 396) (SNMP ifIndex 656)
    Flags: Up SNMP-Traps 0x4000 Encapsulation: ENET2
    Ethernet segment value: 00:00:00:00:00:00:00:00:00:00, Mode: single-homed, Multi-homed status: Forwarding
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.255.11, L2 Routing Instance: evpn_vrf, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32769 (Index 385) (SNMP ifIndex 657)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.255.21, L2 Routing Instance: evpn_vrf, L3 Routing Instance: default
    Input packets : 510
    Output packets: 68
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode

  Logical interface vtep.32770 (Index 390) (SNMP ifIndex 658)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.255.12, L2 Routing Instance: evpn_vrf, L3 Routing Instance: default
    Input packets : 124
    Output packets: 150
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode

  Logical interface vtep.32771 (Index 392) (SNMP ifIndex 659)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.255.22, L2 Routing Instance: evpn_vrf, L3 Routing Instance: default
    Input packets : 7675
    Output packets: 2882
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode

root@Core1>
```

Now, we verify if the Core has Desktop1 and Desktop 2’s MAC and ARP entries:

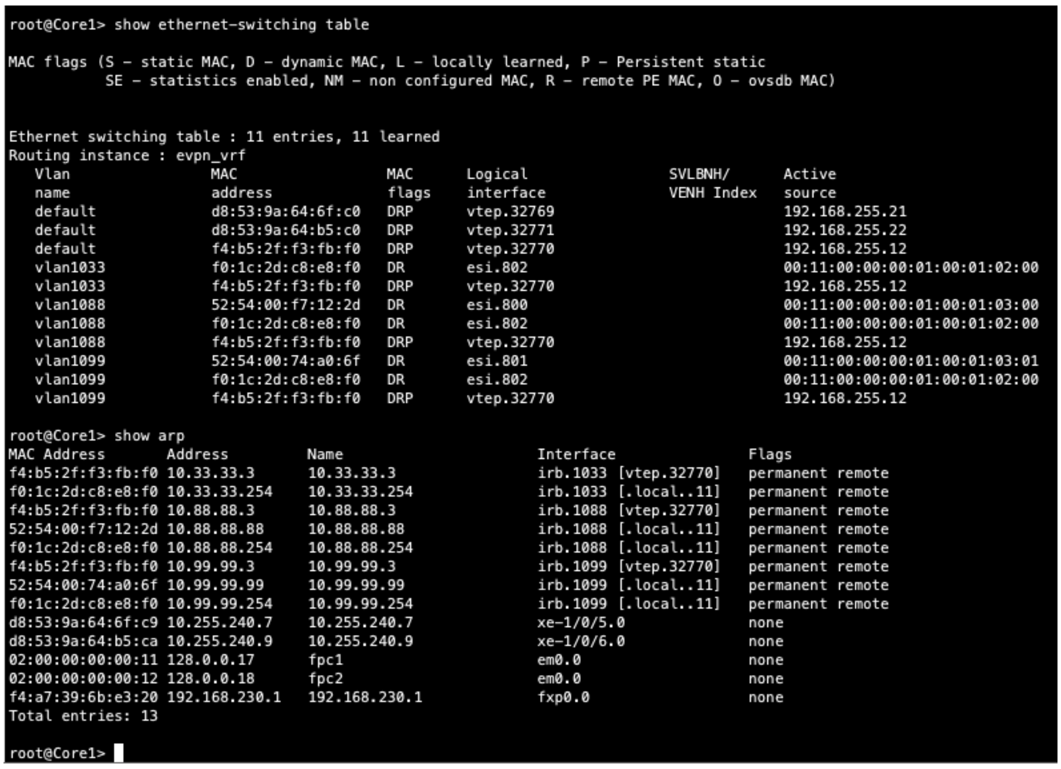

```
root@Core1> show ethernet-switching table

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


Ethernet switching table : 11 entries, 11 learned
Routing instance : evpn_vrf
   Vlan                MAC                 MAC      Logical                SVLBNH/      Active
   name                address             flags    interface              VENH Index   source
   default             d8:53:9a:64:6f:c0   DRP      vtep.32769                          192.168.255.21
   default             d8:53:9a:64:b5:c0   DRP      vtep.32771                          192.168.255.22
   default             f4:b5:2f:f3:fb:f0   DRP      vtep.32770                          192.168.255.12
   vlan1033            f0:1c:2d:c8:e8:f0   DR       esi.802                             00:11:00:00:00:01:00:01:02:00
   vlan1033            f4:b5:2f:f3:fb:f0   DRP      vtep.32770                          192.168.255.12
   vlan1088            52:54:00:f7:12:2d   DR       esi.800                             00:11:00:00:00:01:00:01:03:00
   vlan1088            f0:1c:2d:c8:e8:f0   DR       esi.802                             00:11:00:00:00:01:00:01:02:00
   vlan1088            f4:b5:2f:f3:fb:f0   DRP      vtep.32770                          192.168.255.12
   vlan1099            52:54:00:74:a0:6f   DR       esi.801                             00:11:00:00:00:01:00:01:03:01
   vlan1099            f0:1c:2d:c8:e8:f0   DR       esi.802                             00:11:00:00:00:01:00:01:02:00
   vlan1099            f4:b5:2f:f3:fb:f0   DRP      vtep.32770                          192.168.255.12

root@Core1> show arp
MAC Address       Address         Name                      Interface               Flags
f4:b5:2f:f3:fb:f0 10.33.33.3      10.33.33.3                irb.1033 [vtep.32770]   permanent remote
f0:1c:2d:c8:e8:f0 10.33.33.254    10.33.33.254              irb.1033 [.local..11]   permanent remote
f4:b5:2f:f3:fb:f0 10.88.88.3      10.88.88.3                irb.1088 [vtep.32770]   permanent remote
52:54:00:f7:12:2d 10.88.88.88     10.88.88.88               irb.1088 [.local..11]   permanent remote
f0:1c:2d:c8:e8:f0 10.88.88.254    10.88.88.254              irb.1088 [.local..11]   permanent remote
f4:b5:2f:f3:fb:f0 10.99.99.3      10.99.99.3                irb.1099 [vtep.32770]   permanent remote
52:54:00:74:a0:6f 10.99.99.99     10.99.99.99               irb.1099 [.local..11]   permanent remote
f0:1c:2d:c8:e8:f0 10.99.99.254    10.99.99.254              irb.1099 [.local..11]   permanent remote
d8:53:9a:64:6f:c9 10.255.240.7    10.255.240.7              xe-1/0/5.0              none
d8:53:9a:64:b5:ca 10.255.240.9    10.255.240.9              xe-1/0/6.0              none
02:00:00:00:00:11 128.0.0.17      fpc1                      em0.0                   none
02:00:00:00:00:12 128.0.0.18      fpc2                      em0.0                   none
f4:a7:39:6b:e3:20 192.168.230.1   192.168.230.1             fxp0.0                  none
Total entries: 13

root@Core1> █
```

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

External Campus Fabric connectivity through the Border GW Core EX9204 switches

Remember that the user chose to deploy the Border GW capability on the EX9204 switches during the IP Clos deployment, represented below:

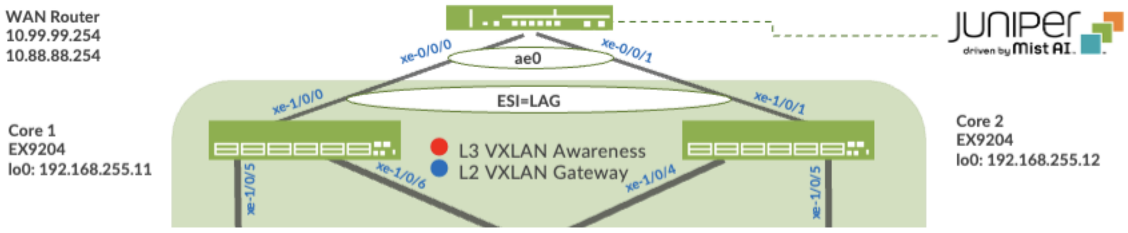

Mist enables the EX9204 to translate between VXLAN traffic within the Campus Fabric and standard ethernet switching for external connectivity, in this case a SRX firewall. Lets verify the ESI status on the Core switches.

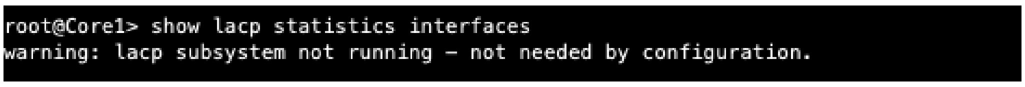

```
root@Core1> show lacp statistics interfaces
warning: lacp subsystem not running - not needed by configuration.
```

We forgot to configure the ESI-LAG: 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 facing EX9204 port.

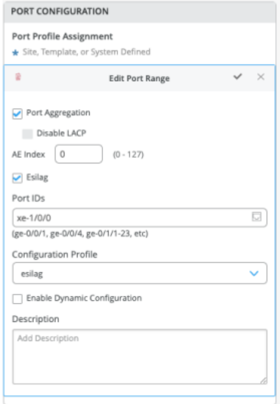
PORT CONFIGURATION

Port Profile Assignment

★ Site, Template, or System Defined

Edit Port Range

☑ Port Aggregation

□ Disable LACP

AE Index: 0 (0 - 127)

☑ Esilag

Port IDs: xe-1/0/0

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

Configuration Profile: esilag

□ Enable Dynamic Configuration

Description: Add Description

Save the config and then verify the changes on the Core switch.

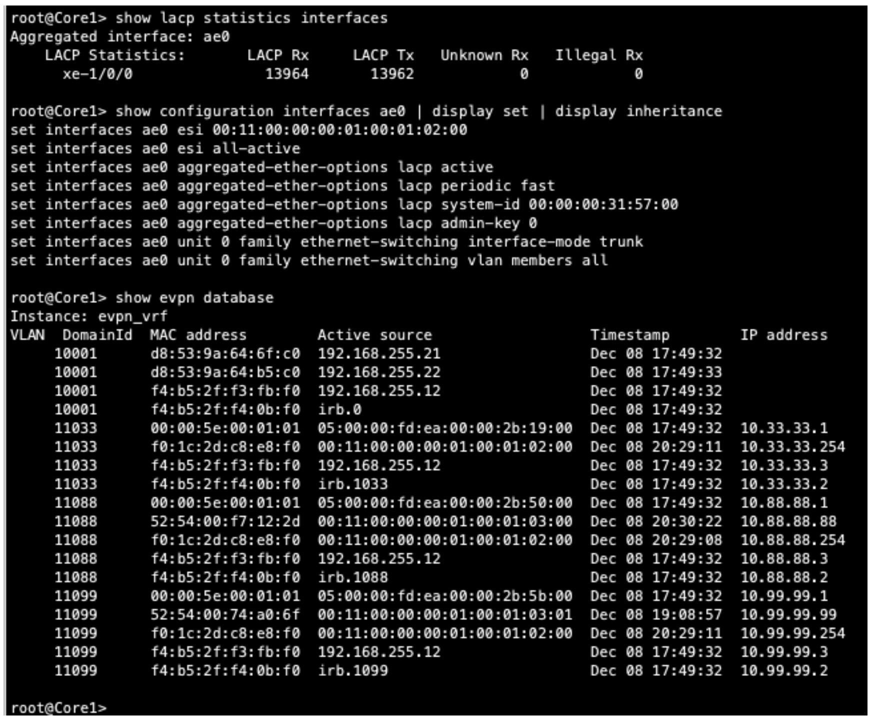
```
root@Core1> show lacp statistics interfaces
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-1/0/0               13964       13962            0            0

root@Core1> show configuration interfaces ae0 | display set | display inheritance
set interfaces ae0 esi 00:11:00:00:00:01:00:01:02:00
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:31:57:00
set interfaces ae0 aggregated-ether-options lacp admin-key 0
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members all

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

root@Core1>
```

Note that LACP is up (this infers there is an existing configuration on the SRX firewall.

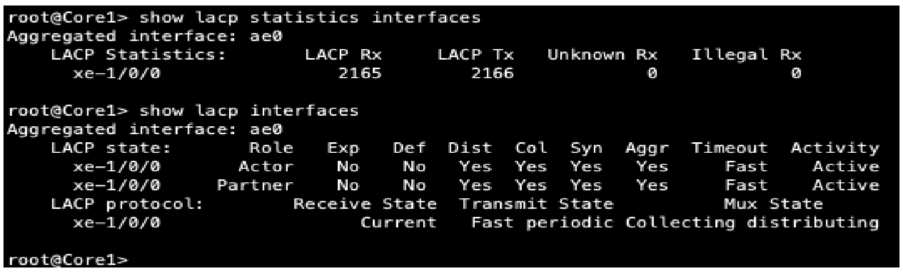
```
root@Core1> show lacp statistics interfaces
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-1/0/0                2165        2166            0            0

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 data base now has the ESI entry. Back to Desktop1 to see if it can cross the fabric.

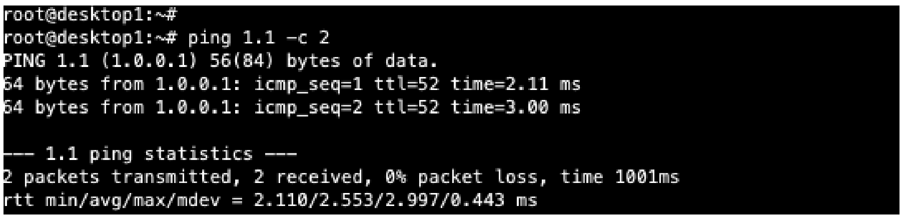
```
root@desktop1:~#
root@desktop1:~# ping 1.1 -c 2
PING 1.1 (1.0.0.1) 56(84) bytes of data.
64 bytes from 1.0.0.1: icmp_seq=1 ttl=52 time=2.11 ms
64 bytes from 1.0.0.1: icmp_seq=2 ttl=52 time=3.00 ms

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

Last step is to verify Desktop1 can ping desktop2

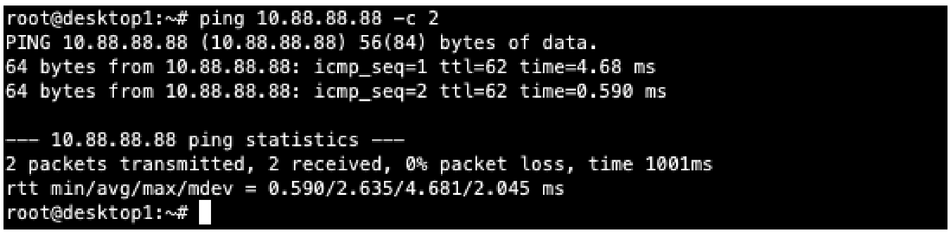
```
root@desktop1:~# ping 10.88.88.88 -c 2
PING 10.88.88.88 (10.88.88.88) 56(84) bytes of data.
64 bytes from 10.88.88.88: icmp_seq=1 ttl=62 time=4.68 ms
64 bytes from 10.88.88.88: icmp_seq=2 ttl=62 time=0.590 ms

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

Meaning: Connectivity within the Campus Fabric and externally have been verified. Desktops communicate with each other through the Campus Fabric, each in an isolated VRF, then forwarded to the SRX firewall 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 the user with realtime status related to the health of the Campus Fabric Core Distribution CRB deployment using telemetry such as BGP neighbor status and TX/RX port statistics. The following screenshots are taken from the Campus Fabric Core Distribution CRB build by accessing the Campus Fabric option under the Organization/Wired of the Mist Portal:

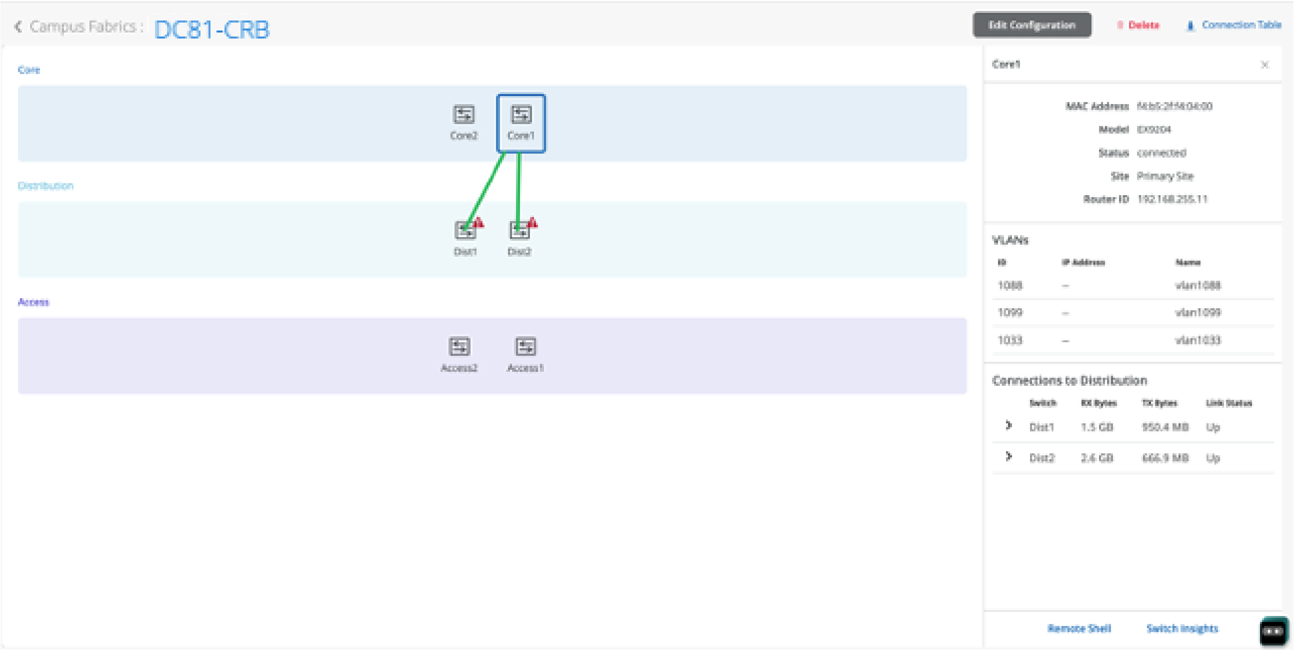

Core1

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	1.5 GB	950.4 MB	Up
Dist2	2.6 GB	666.9 MB	Up

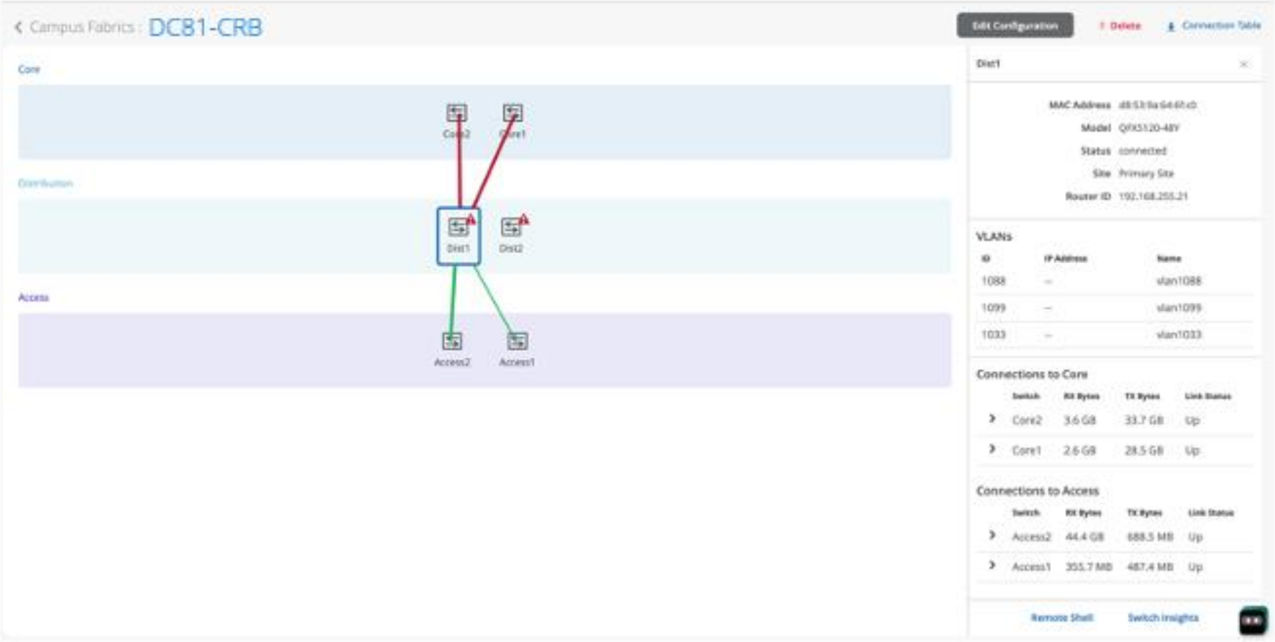

Dist1

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	3.6 GB	33.7 GB	Up
Core1	2.6 GB	28.5 GB	Up

Connections to Access

Switch	RX Bytes	TX Bytes	Link Status
Access2	44.4 GB	688.5 MB	Up
Access1	355.7 MB	487.4 MB	Up

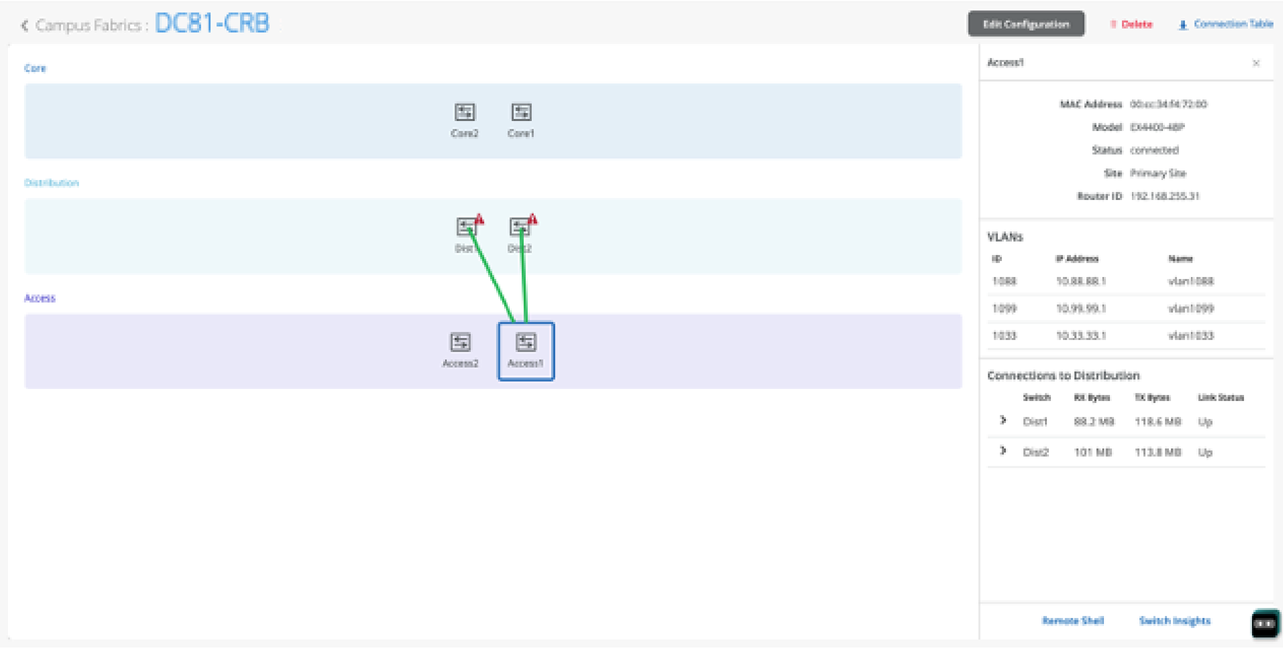

Access1

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 Status
Dist1	88.2 MB	118.6 MB	Up
Dist2	101 MB	113.8 MB	Up

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

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

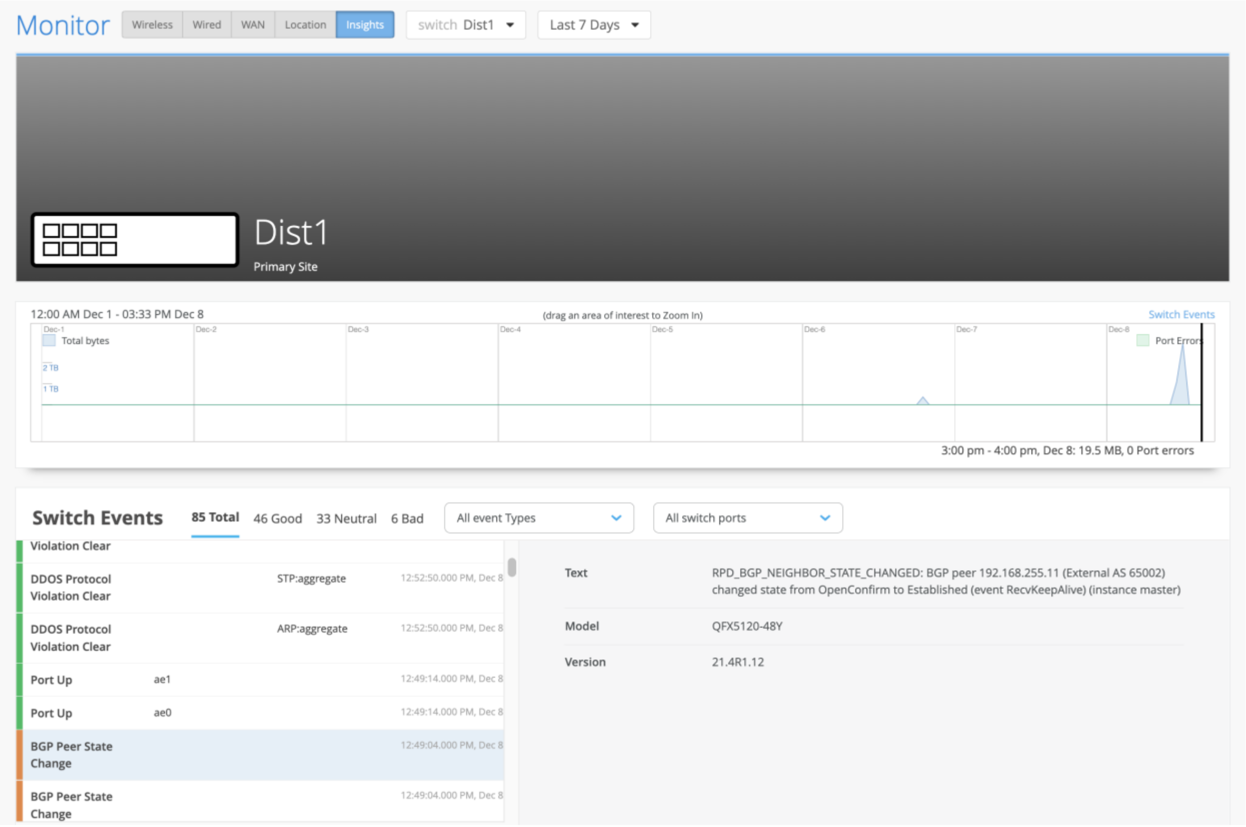

Violation Clear			
DDOS Protocol Violation Clear		STP:aggregate	12:52:50.000 PM, Dec 8
DDOS Protocol Violation Clear		ARP:aggregate	12:52:50.000 PM, Dec 8
Port Up	ae1		12:49:14.000 PM, Dec 8
Port Up	ae0		12:49:14.000 PM, Dec 8
BGP Peer State Change			12:49:04.000 PM, Dec 8
BGP Peer State Change			12:49:04.000 PM, Dec 8

Text	RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 192.168.255.11 (External AS 65002) changed state from OpenConfirm to Established (event RecvKeepAlive) (instance master)
Model	QFX5120-48Y
Version	21.4R1.12

Summary

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

Appendix

Campus Fabric Core Distribution CRB Configurations

Note: Those items in RED are specific to the device and require uniqueness

eBGP Underlay

Dist1

```
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 "$9$wLYZUji.Qz624QF/Cu0-VbsJGmfTz69YgJDk.5TlKv8-
V2gJZjH4o9AtuEh-
Vb2gJqmfzn/PfnCp0hcoJZUqm69A0ORCABEcyW8aZGimf5QF9Cun6evMWx7ikq.PQ/CtOIEn6vWxNbwgoJGUHqmf
TQnmPRhSyW8k.mPz3p0B1hSu0IcSr8LGDjHfT"
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
```

Core1

```
set protocols bgp group evpn_underlay type external
set protocols bgp group evpn_underlay log-updown
set protocols bgp group evpn_underlay import evpn_underlay_import
set protocols bgp group evpn_underlay family inet unicast
set protocols bgp group evpn_underlay authentication-key "$9$deboJGUHf5FwYfT36AtxN-
V4ak.P5Fnbs4ZjHmPSrlvxNws4oGDY2n/9A1IxN-
ws4ik.5z3q.z6CtIR24oJikFn/tpB6/u1RhKvgoaUk.mfTn6AzFyleK8LUjiHqf369pO1zFlK8X-
ds24aJDik.PfzkqBIEhKvjHkq5QCtu0IEAtOREcvMaZGD.P"
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
```

eBGP Overlay in support of EVPN-VXLAN

Dist1

```
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
set protocols bgp group evpn_overlay authentication-key "$9$wLYZUji.Qz624QF/Cu0-VbsJGmfTz69YgJDk.5TlKv8-
V2gJZjH4o9AtuEh-
Vb2gJqmfzn/PfnCp0hcoJZUqm69A0ORCABEcyW8aZGimf5QF9Cun6evMWx7ikq.PQ/CtOIEn6vWxNbwgoJGUHqmf
TQnmPRhSyW8k.mPz3p0B1hSu0IcSr8LGDjHfT"
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
```

Core1

```
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
set protocols bgp group evpn_overlay authentication-key "$9$deboJGUHf5FwYfT36AtxN-
V4ak.P5Fnbs4ZjHmPSrlvxNws4oGDY2n/9A1IxN-
ws4ik.5z3q.z6CtIR24oJikFn/tpB6/u1RhKvgoaUk.mfTn6AzFyleK8LUjiHqf369pO1zFlK8X-
ds24aJDik.PfzkqBIEhKvjHkq5QCtu0IEAtOREcvMaZGD.P"
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
```

VXLAN and EVPN Enablement

Dist1

```
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

set groups top switch-options vtep-source-interface lo0.0
set groups top switch-options route-distinguisher 192.168.255.21:1
set groups top switch-options vrf-target target:65000:1
```

Core1

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

set groups top routing-instances evpn_vrf vtep-source-interface lo0.0
```

Note: Access Switches are L2 switches with no EVPN-VXLAN requirements

VLAN-VXLAN (VNI) mapping

Dist1

```
set vlans default vlan-id 1
set vlans default vxlan vni 10001
set vlans vlan1033 vlan-id 1033
set vlans vlan1033 vxlan vni 11033
set vlans vlan1088 vlan-id 1088
set vlans vlan1088 vxlan vni 11088
set vlans vlan1099 vlan-id 1099
set vlans vlan1099 vxlan vni 11099
```

Core1

```
set groups top routing-instances evpn_vrf vlans default vlan-id 1
set groups top routing-instances evpn_vrf vlans default vxlan vni 10001
set groups top routing-instances evpn_vrf vlans vlan1033 vlan-id 1033
set groups top routing-instances evpn_vrf vlans vlan1033 vxlan vni 11033
set groups top routing-instances evpn_vrf vlans vlan1088 vlan-id 1088
set groups top routing-instances evpn_vrf vlans vlan1088 vxlan vni 11088
set groups top routing-instances evpn_vrf vlans vlan1099 vlan-id 1099
set groups top routing-instances evpn_vrf vlans vlan1099 vxlan vni 11099
```

Routing Policy:

Dist and Core Devices:

```
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 policy-options policy-statement evpn_underlay_export term 01-loopback from route-filter
192.168.255.0/24 orlonger
set groups top policy-options policy-statement evpn_underlay_export term 01-loopback then accept
set groups top policy-options policy-statement evpn_underlay_export term 02-default then reject
set groups top policy-options policy-statement evpn_underlay_import term 01-loopback from route-filter
192.168.255.0/24 orlonger
set groups top policy-options policy-statement evpn_underlay_import term 01-loopback then accept
set groups top policy-options policy-statement evpn_underlay_import term 02-default then reject
```

L3 Interfaces (IRB)

Core1

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

Routing Instances

Core1

```
set groups top routing-instances guest-wifi instance-type vrf
set groups top routing-instances guest-wifi routing-options static route 0.0.0.0/0 next-hop 10.33.33.254
set groups top routing-instances guest-wifi routing-options auto-export
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes advertise direct-nexthop
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes encapsulation vxlan
set groups top routing-instances guest-wifi protocols evpn ip-prefix-routes vni 15560868
set groups top routing-instances guest-wifi interface irb.1033
set groups top routing-instances guest-wifi route-distinguisher 192.168.255.11:103
set groups top routing-instances guest-wifi vrf-target target:65000:103
set groups top routing-instances guest-wifi vrf-table-label
set groups top routing-instances developers instance-type vrf
set groups top routing-instances developers routing-options static route 0.0.0.0/0 next-hop 10.88.88.254
set groups top routing-instances developers routing-options auto-export
set groups top routing-instances developers protocols evpn ip-prefix-routes advertise direct-nexthop
set groups top routing-instances developers protocols evpn ip-prefix-routes encapsulation vxlan
set groups top routing-instances developers protocols evpn ip-prefix-routes vni 15600414
set groups top routing-instances developers interface irb.1088
set groups top routing-instances developers route-distinguisher 192.168.255.11:102
set groups top routing-instances developers vrf-target target:65000:102
set groups top routing-instances developers vrf-table-label
set groups top routing-instances corp-it instance-type vrf
set groups top routing-instances corp-it routing-options static route 0.0.0.0/0 next-hop 10.99.99.254
set groups top routing-instances corp-it routing-options auto-export
set groups top routing-instances corp-it protocols evpn ip-prefix-routes advertise direct-nexthop
set groups top routing-instances corp-it protocols evpn ip-prefix-routes encapsulation vxlan
set groups top routing-instances corp-it protocols evpn ip-prefix-routes vni 11284517
set groups top routing-instances corp-it interface irb.1099
set groups top routing-instances corp-it route-distinguisher 192.168.255.11:101
set groups top routing-instances corp-it vrf-target target:65000:101
set groups top routing-instances corp-it vrf-table-label
set groups top routing-instances evpn_vrf instance-type virtual-switch
set groups top routing-instances evpn_vrf protocols evpn encapsulation vxlan
set groups top routing-instances evpn_vrf protocols evpn default-gateway no-gateway-community
set groups top routing-instances evpn_vrf protocols evpn extended-vni-list all
set groups top routing-instances evpn_vrf protocols rstp interface ae0 disable
set groups top routing-instances evpn_vrf vtep-source-interface lo0.0
set groups top routing-instances evpn_vrf interface ae0.0
set groups top routing-instances evpn_vrf route-distinguisher 192.168.255.11:1
set groups top routing-instances evpn_vrf vrf-target target:65000:1
set groups top routing-instances evpn_vrf vlans default vlan-id 1
set groups top routing-instances evpn_vrf vlans default l3-interface irb.0
set groups top routing-instances evpn_vrf vlans default vxlan vni 10001
set groups top routing-instances evpn_vrf vlans vlan1033 vlan-id 1033
set groups top routing-instances evpn_vrf vlans vlan1033 l3-interface irb.1033
set groups top routing-instances evpn_vrf vlans vlan1033 vxlan vni 11033
set groups top routing-instances evpn_vrf vlans vlan1088 vlan-id 1088
set groups top routing-instances evpn_vrf vlans vlan1088 l3-interface irb.1088
set groups top routing-instances evpn_vrf vlans vlan1088 vxlan vni 11088
set groups top routing-instances evpn_vrf vlans vlan1099 vlan-id 1099
set groups top routing-instances evpn_vrf vlans vlan1099 l3-interface irb.1099
set groups top routing-instances evpn_vrf vlans vlan1099 vxlan vni 11099
```

ESI-LAG between Dist and Access Switches:

Dist1

```
set groups dc81lag interfaces <*> mtu 9100
set groups dc81lag interfaces <*> unit 0 family ethernet-switching interface-mode trunk
set groups dc81lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1088
set groups dc81lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1099
set groups dc81lag interfaces <*> unit 0 family ethernet-switching vlan members vlan1033

set interfaces ge-0/0/36 description esilag-to-00cc34f3cf00
set interfaces ge-0/0/37 description esilag-to-00cc34f3cf00
set interfaces ae0 apply-groups dc81lag
set interfaces ae1 apply-groups dc81lag

set interfaces ae0 apply-groups dc81lag
set interfaces ae0 esi 00:11:00:00:00:01:00:01:03:00
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:31:57:00
set interfaces ae0 aggregated-ether-options lacp admin-key 0

set interfaces ae1 apply-groups dc81lag
set interfaces ae1 esi 00:11:00:00:00:01:00:01:03:01
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp periodic fast
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:31:57:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
```

ESI-LAG between Core1/2 and SRX Firewall

Core1

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

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

SRX

```
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 1033
set interfaces ae0 unit 0 family bridge vlan-id-list 1088
set interfaces ae0 unit 0 family bridge vlan-id-list 1099
```