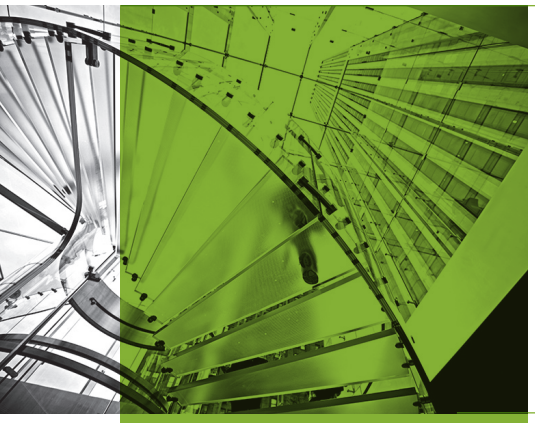

EVPN-VXLAN CAMPUS FABRICS

Challenge

Traditional campus networks are proprietary and too rigid to support the needs of endpoints in larger enterprises. These networks must be flexible enough to accommodate IoT devices and provide consistent security at every layer, both within and across campuses.

Solution

Juniper's AI-driven campus fabrics, based on a VXLAN overlay with an EVPN control plane, deliver an efficient and scalable way to build and interconnect enterprise networks with consistency. Cloud managed Juniper Networks® EX Series Ethernet switches are onboarded, deployed, and managed at scale with Juniper Mist™ cloud, assuring great wired and wireless experiences through the automation, insight, and simplicity of AI-driven campus management.

Benefits

- *Control plane-based L2/L3 information exchange*
- *Efficient host mobility*
- *Open, nonproprietary solution*
- *Scalability at all network layers*
- *Faster convergence*
- *Flexible and secure architecture*
- *AI-driven campus management*

Enterprise networks around the world are adopting cloud and cloud-based applications to improve their competitiveness, lower IT costs, and provide users with anytime, anywhere access to resources and data. This trend, driven largely by the widespread use of mobile devices, social media, and collaboration tools, along with the growing number of IoT devices, is having a significant impact on enterprise campus networks. A growing number of network endpoints, coupled with rapidly evolving business needs, is driving demand for highly scalable networks that are not only simple, scalable, and programmable, but also built on a standards-based architecture that is common across both the campus and data center.

The Challenge

Most campus networks are based on conventional Layer 2 Ethernet-based architectures that leverage technologies like Juniper's Virtual Chassis technology, eliminating the need for Spanning Tree Protocol (STP). While these architectures work well in small or medium-sized campuses catering to traditional requirements where services are limited to a single network, they are sometimes too rigid to support the scalability required by larger enterprises.

Cloud-based applications enable new business models, provide greater business agility, and support the adoption of key technologies such as unified communications, video, and other latency-sensitive applications. The increasing use of IoT devices also means that these same networks are expected to scale rapidly without adding complexity. Since many IoT devices have limited networking capabilities, they require L2 adjacency across buildings or campuses. Traditionally, this problem was solved by extending VLANs across buildings and campuses using data plane "flood-and-learn." This approach, however, is inefficient and hard to manage—inefficient due to excess consumption of network bandwidth, and difficult to manage because VLANs need to be extended to new network ports.

Security, which is no longer just a perimeter problem, also poses a unique challenge. Modern enterprises want security to be embedded into their network architectures—not just inside the campus, but through segmentation and policies extended across the entire organization, including data centers.

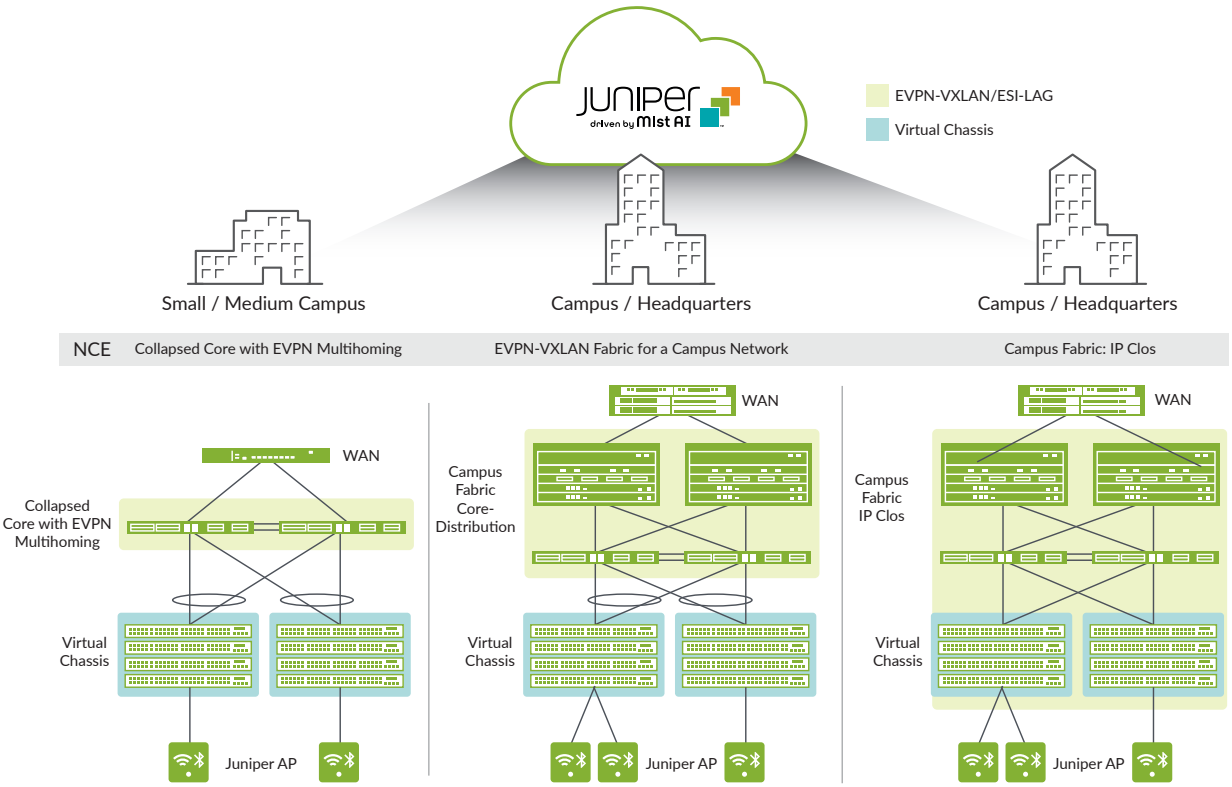

Figure 1: Campus fabrics showing Virtual Chassis and EVPN-VXLAN-based architectures

Juniper Networks Validated Campus Fabric Designs

Juniper Networks campus fabrics provide a single, standards-based Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) solution that can be deployed in any campus, whether a two-tier network with a collapsed core distribution or a campus-wide system that involves multiple buildings with separate distribution and core layers. Figure 1 shows the validated fabric architectures.

- **EVPN multihoming (on collapsed core or distribution):** A collapsed core architecture merges the core and distribution layers into a single switch, turning the traditional three-tier hierarchal network into a two-tiered network. This eliminates the need for STP across campus networks by providing multihoming capabilities from the access to the core layer.
- **Campus fabric core-distribution:** A pair of interconnected core or distribution Juniper Networks® EX Series Ethernet Switches provides L2 EVPN and L3 VXLAN gateway support. The IP Clos network between the distribution and core layers offers two modes: centrally or edge-routed bridging overlay.
- **Campus fabric IP Clos:** This IP Clos architecture pushes VXLAN L2 gateway functionality to the access layer. This model is also referred to as "end-to-end," given that VXLAN tunnels are terminated at the access layer.

Campus fabric architectures let you manage your campus as a single IP fabric, with over-the-top (OTT) policy and control provided by Juniper. Any number of switches can be connected in a Clos network or IP fabric; EVPN-VLAN extends the fabric and connects multiple enterprise buildings, while VXLAN stretches L2 across the network. This flexible, campus-fabric, IP Clos network between the distribution and core layers can exist in two modes: centrally routed bridging overlay or edge-routed bridging overlay.

Features and Benefits

Juniper's standards-based EVPN-VXLAN solution offers the following features and benefits when operating as a campus control plane protocol.

- Greater network efficiency:
  - Supports multipath traffic to active/active dual-homed access layer switches
  - Reduces unknown unicast flooding with control plane media access control (MAC) learning
  - Reduces Address Resolution Protocol (ARP) flooding by enabling MAC-to-IP binding in the control plane
  - Supports multipath traffic over multiple core switches (VXLAN entropy)

- Flexibility:
  - Easily enables L2 stretch
  - Enables micro- and macrosegmentation with group-based policies (GBPs)
  - Enables easy integration with L3 and L2 VPNs
- Fast convergence:
  - Enables rapid reconvergence when links to dual-homed access switches fail (aliasing)
  - Supports faster reconvergence when endpoints move
- Scalability:
  - Offers scalable BGP-based control plane
  - Allows seamless expansion of core, aggregation, and access layers as business needs grow
  - Supports seamless expansion of campuses as business needs grow
- Nonproprietary:
  - Supports multivendor core, aggregation, and access layers with standards-based protocols
- Security:
  - Implements network microsegmenation using GBP that leverages the underlying VXLAN technology
  - Provides location-agnostic endpoint access control with consistent security policies
  - Blocks lateral threats
  - Simplifies network configurations by avoiding the need to configure a large number of firewall filters

Solution Components

The EVPN-VXLAN-based campus architecture decouples the overlay network from the underlay.

VXLAN, an encapsulation/tunneling protocol, does not change the flood-and-learn behavior of the Ethernet protocol. Instead, its control protocol—in this case, EVPN—uses Multiprotocol BGP (MP-BGP) to allow the network to carry both L2 MAC and L3 IP information. By making the combined set of MAC and IP information available for forwarding decisions, EVPN—together with VXLAN—optimizes routing and switching behavior. Meanwhile, the EVPN extension that allows BGP to transport L2 MAC and L3 IP information offers an alternative to the flood-and-learn behavior, which is considered suboptimal in many use cases.

With overlays, endpoints can be placed anywhere in the network and remain connected to the same logical L2 network, enabling a virtual topology to be decoupled from the physical topology. With an EVPN control plane, enterprises can easily add more core, aggregation, and access layer devices as the business grows without having to redesign the network or perform a forklift upgrade.

The EVPN-VXLAN-based architecture lets you deploy a common set of policies and services across campuses with support for L2 and L3 VPNs. Using an L3 IP-based underlay coupled with an EVPN-VXLAN overlay, campus network operators can deploy much larger networks than would otherwise be possible with traditional L2 Ethernet-based architectures.

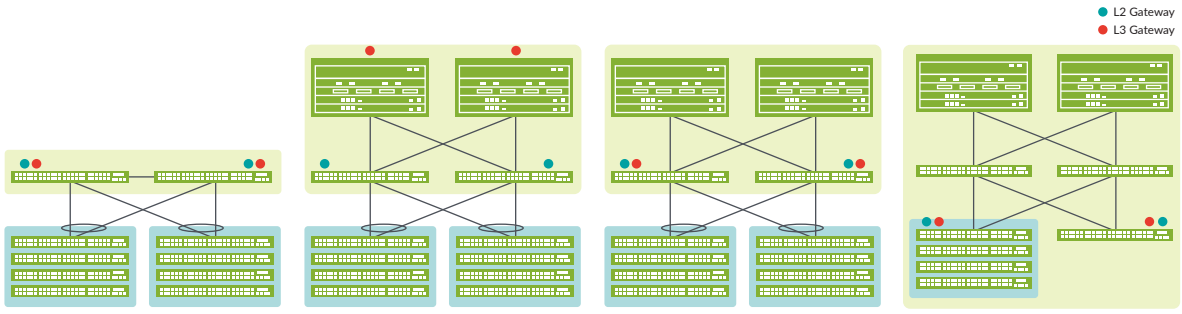

Technology	EVPN Multihoming	Campus Fabric Core-Distribution	Campus Fabric Core-Distribution	Campus Fabric IP Clos
Technology	ESI-LAG	Centrally-routed bridging (CRB)	Edge-routed bridging (ERB)	End-to-End EVPN
Positioning	• Small/medium campus • Collapsed core/distribution • Inter-campus traffic (N/S)	• Small/medium campus • Inter-campus traffic (N/S)	• Medium campus • Intra-campus traffic (E/W)	• Medium/large campus • Intra-campus traffic (E/W) • Recommended when L3 at access
Advantages	• Eliminates STP • Easy migration from Virtual Chassis or MC-LAG	• Simpler as L2/L3 gateway is just on the core • Supports distribution devices that can't perform L3 VXLAN	• Reduced need for gateway re-learning and smaller blast radius • Better multi-vendor interoperability	• Access layer segmentation • Ideal for mobility and IoT devices

Figure 2: Different options for campus fabric

In campus fabric architectures (see Figure 2), the core and aggregation layers form an L3 fabric with an EVPN-VXLAN overlay. Ideally, the underlay would be deployed using the L3 IP Clos model with core and aggregation switches, while the access layer switches would be multihomed to the distribution layer.

The IP Clos model provides an architecture that enables deterministic latency and horizontal scale at the core, aggregation, and access layers. An interior gateway protocol (IGP) like OSPF or EBGP can be used as the underlay routing protocol. Figure 2 shows an IBGP overlay with route reflection where aggregation devices within a given pod or group share endpoint information upstream as EVPN routes to core devices acting as route reflectors. The core devices reflect the routes to downstream aggregation devices using route reflectors to eliminate the need for full-mesh BGP connections, and to simplify the aggregation layer by applying consistent configurations across all aggregation layer switches.

The access layer switches, typically deployed in a Virtual Chassis configuration that allows up to 10 interconnected platforms to operate as a single, logical device, are not part of the EVPN-VXLAN fabric. The access layer, which is L2 only, maps endpoints to VLANs, which are carried in trunk ports to the aggregation layer using the multihomed uplinks from the access layer to the aggregation layer. This vendor-agnostic solution allows enterprises to use their existing access layer infrastructure and upgrade to standards-based access layer switches from Juniper or any other vendor.

VLANs are mapped to VXLANs at the distribution layer, while L3 integrated routing and bridging (IRB) or switch virtual interface (SVI) for the VXLANs are located on the core switch with an anycast gateway address. Flexible and secure configuration options mean that IRBs can be placed in a common routing instance or, if segmentation is required, in separate routing instances. Similar to virtual routing and forwarding (VRF) tables, routing instances enable the network to be segmented for multitenancy and/or security. Based on the enterprise security policy, some routes can be leaked between routing instances for inter-VRF communication, or inter-VRF traffic can be routed through a firewall to enforce advanced security with network segmentation.

Like other Juniper architectures, the campus fabrics do not force customers to invest in new devices; the same devices and technologies can be used in an EVPN-VXLAN campus fabric deployment, as shown below.

- Core layer:
  - Juniper Networks EX9200 Ethernet Switches
  - Juniper Networks QFX10000 Ethernet Switches
  - Juniper Networks QFX5110 and QFX5120 Ethernet Switches
- Distribution layer:
  - Juniper Networks EX4650 Ethernet Switch
  - Juniper Networks QFX5110 and QFX5120 Ethernet Switches
  - Juniper Networks QFX10000 Ethernet Switches
- Access layer:
  - Juniper Networks EX4400/EX4300/EX3400/EX2300 Ethernet Switches
  - Virtual Chassis technology

Connecting Multiple Sites

The benefits of the EVPN-VXLAN-based fabric can be extended across campuses, data centers, and public cloud infrastructure with L2 and L3 VPN support in EVPN (see Figure 3). VXLAN is WAN underlay-agnostic provided the campuses, data centers, and the public cloud infrastructure have IP connectivity. EVPN-VXLAN overlays can be deployed over a variety of WAN technologies, including private MPLS and IPsec over Internet.

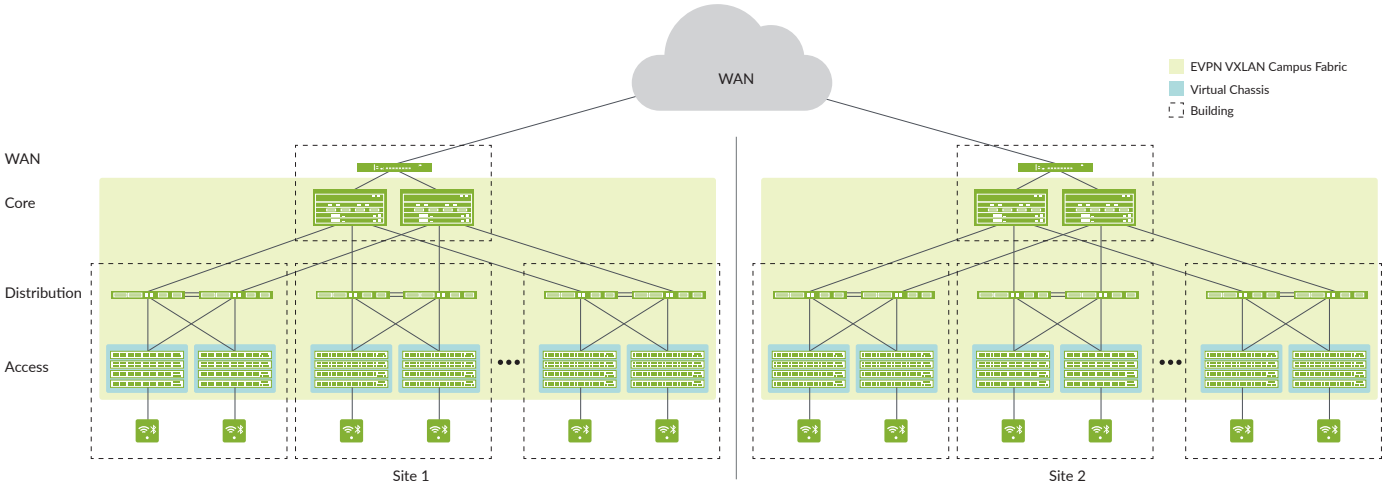

Figure 3: Interconnecting multiple sites with EVPN-VXLAN overlay

Deploying an AI-Driven Campus Fabric

Configuring campus fabrics manually can cause inconsistency and unforced errors in deployments. Juniper solves this operational burden by enabling EVPN-VXLAN campus fabrics to be easily managed via the Juniper Mist cloud. More specifically, administrators can choose a topology (EVPN multihoming, distribution-core, or IP CLOS), and let the software do the rest (see Figure 4). This AI-driven approach unifies management across the LAN, WLAN, and WAN environments in the campus and branch while assuring the wired and wireless campus network delivers great user experiences.

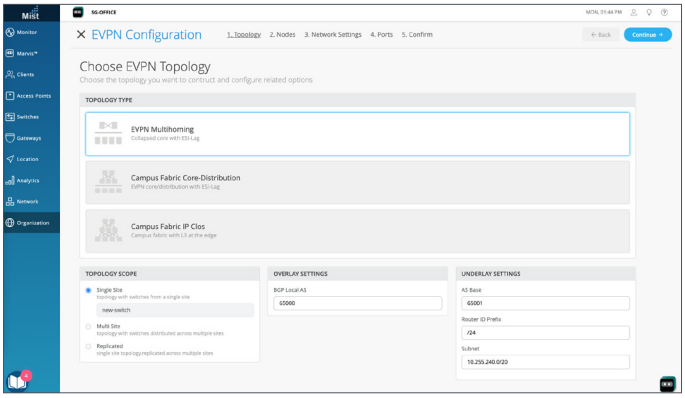

Figure 4: Juniper Mist Wired Assurance campus fabric design

*EVPN-multihoming initially supported, additional architectures supported in future releases.

Operating an AI-Driven Campus Fabric

Juniper Mist™ Wired Assurance claims, configures, manages, and troubleshoots cloud-managed EX Series Ethernet switches. The cloud-based service delivers AI-powered automation and service levels to ensure a better experience for connected devices. Juniper Mist Wired Assurance leverages rich Junos® operating system switch telemetry data to simplify operations, reduce mean time to repair, and improve visibility. Key features for Day 0 through Day 2 operations are:

- **Day 0 operations**—Onboard switches seamlessly by claiming a greenfield switch or adopting a brownfield switch with a single activation code for true plug-and-play simplicity.
- **Day 1 operations**—Implement a template-based configuration model for bulk rollouts of traditional and campus fabric deployments, while retaining the flexibility and control required to apply custom site- or switch-specific attributes. Automate provisioning of ports via Dynamic Port Profiles.
- **Day 2 operations**—Leverage the AI in Juniper Mist Wired Assurance to meet service-level expectations such as throughput, successful connects, and switch health with key pre- and post-connection metrics (see Figure 5). Add the self-driving capabilities in Marvis Actions to detect loops, add missing VLANs, fix misconfigured ports, identify bad cables, isolate flapping ports, and discover persistently failing clients (see Figure 6). And perform software upgrades easily through Juniper Mist cloud.

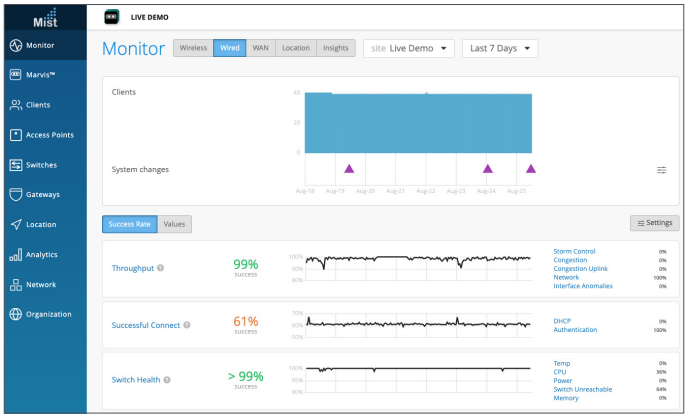

Figure 5: Juniper Mist Wired Assurance service-level expectations

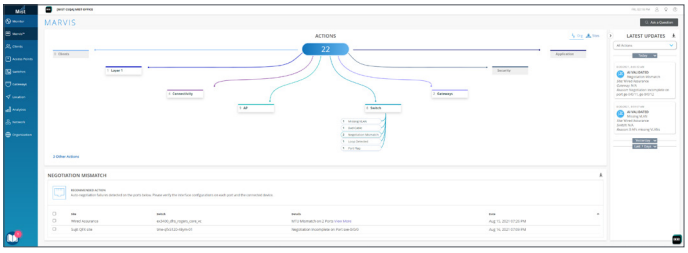

Figure 6: Marvis Actions for wired switches

For more information on Juniper Mist™ Wired Assurance.

Summary—Enterprises Must Embrace EVPN and VXLAN

Cloud-based resources are becoming an increasingly large part of the enterprise's IT strategy, and this requires a network architecture that can accommodate cloud-based services without compromising security or performance. The demands of campus users for anytime, anywhere access and high levels of responsiveness are becoming harder and harder to achieve with traditional network architectures. The increasing prevalence of IoT devices in campus networks demands a network that is not rigid and yet maintains an architecture that is scalable, simple, programmable, open, and supportive of multivendor devices.

Juniper's campus fabrics, based on a VXLAN overlay with an EVPN control plane, is an efficient and scalable way to build campuses and interconnect multiple campuses, data centers, and public clouds. With a robust BGP/EVPN implementation on all QFX Series and EX Series switches, Juniper is uniquely positioned to bring EVPN technology to its full potential by providing optimized, seamless, and standards-compliant L2 or L3 connectivity, both within and across today's evolving campuses and data centers.

Next Steps

For more information, please contact your Juniper representative, or go to www.juniper.net.

About Juniper Networks

At Juniper Networks, we are dedicated to dramatically simplifying network operations and driving superior experiences for end users. Our solutions deliver industry-leading insight, automation, security and AI to drive real business results. We believe that powering connections will bring us closer together while empowering us all to solve the world's greatest challenges of well-being, sustainability and equality.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.207.125.700
Fax: +31.207.125.701

JUNIPER NETWORKS | Driven by Experience™