GET SMART: NETWORK VIRTUALIZATION WITH EVPN-VXLAN

Learn the benefits of EVPN-VXLAN and influence your leadership.

Juniper Networks | Engineering Simplicity

Table of Contents

Juniper Networks | Engineering Simplicity

ABOUT JUNIPER NETWORKS

Juniper Networks simplifies the complexities of networking with products, solutions and services in the cloud era to transform the way we connect, work and live. We remove the traditional constraints of networking to enable our customers and partners to deliver automated, scalable and secure networks that connect the world.

GO BEHIND THE SCENES OF NETWORK VIRTUALIZATION WITH EVPN-VXLAN

EVPN-VXLAN, like other technologies, started out as a proprietary set of vendor specific solutions that evolved into a standard to address several business challenges for customers. The problem space includes how to support both legacy and cloud native workloads, connecting multiple geographically separated sites including data center and campus environments and leveraging the right set of building blocks to address the flexibility necessary in the enterprise space today.

In this e-book, we've pulled together industry experts and analysts to discuss what EVPN-VXLAN can offer to the industry as a whole.

Join us to learn from the innovators who shaped the networking industry to influence your leadership on the business benefits of EVPN-VXLAN.

READY TO UP YOUR DATA CENTER SWITCHING GAME?

The Advanced Data Center Switching (ADCX) course provides in-depth instruction on IP fabric and EVPN-VXLAN data center design and configuration. Additionally, the course covers other data center concepts, including basic and advanced data center design options, Data Center Interconnect (DCI), EVPN multicast enhancements, and an introduction to data center automation concepts. The course ends with a multi-site data center design lab.

Full course details.

GET STARTED NOW AND SAVE

Register for any upcoming ADCX course by June 1st and save 25%. Just enter the promo code EVPN/VXLAN2019 during checkout.

MULTI-VENDOR NETWORK VIRTUALIZATION WITH EVPN

Leadership Discussion Series

JEFF TANTSURA
Head of Networking Strategy
Apstra

Moderator

ANDRE KINDNESS
Principal Analyst
Forrester Research

ALDRIN ISAAC
Senior Director, Portfolio Product
Juniper Networks

Abstract

Network Virtualization (via EVPN-VXLAN), a widely supported open standard, is evolving into the best answer to multi-vendor network virtualization for enterprise networks. Learn why this open standard is integral to customer success, and how to make it relevant to your leadership. EVPN-VXLAN started out as a proprietary set of vendor specific solutions but has since evolved into a standard to address rising business challenges. This session starts with an EVPN overview, then goes into the value and adoption of the technology.

Key Takeaways

- Network virtualization logically partitions the fabric to address different business requirements. EVPN provides robust control plane for L3 discovery, faster convergence, recognizing new IP and MAC addresses, and the ability to scale data centers using BGP.
- An artificial partition between operations of WAN and data center networks creates inefficiencies; there is success in taking the benefits of a single seamless control and data plane across the WAN and Data Center environment.
- “Intelligent network, intelligent operator”—Standards lay out what the technology is, and what it does. Operators who understand the underlying technology can do basic troubleshooting rather than escalating to a vendor’s TAC to resolve problems.
- In early 2000, there were two standards for L2VPN—Kompella, Martini. Solutions did not interoperate with each other causing a lot of pain—EVPN was created to unite solutions and vendors together, moving the industry forward.
- SDN relies on a strong underlying foundation. Protocols, open or proprietary, are built to make things work. EVPN can be the engine for SDN, and SDN makes things better.
- One of the initial goals of EVPN was to connect existing networks together; it’s not all or nothing; operators can start in one place and gradually adopt EVPN.
- EVPN co-exists with an IP network. A portion of the IP network provides network virtualization using EVPN while the rest of it is provides plain IP forwarding.

“When we started EVPN, it was all about universality. Networking has a few set of problems it needs to solve—it really isn’t all that complicated—we’ve just made it really complicated by solving for the same thing with many many different ways.”

Aldrin Isaac, Juniper Networks

“You don’t do EVPN because you like new technology; it helps solves real problems; MC-LAG it’s probably one of the most complicated technologies and completely per-vendor. You cannot put two vendors doing MC-LAG. EVPN abstracts the differences.”

Jeff Tantsura, Apstra

Juniper and EVPN: Virtualization Simplified

Introduction

The process of capturing and determining the value of data is making networks more complex, impacting topologies, hardware, and protocols.

For instance, segmenting data and applications for security and privacy purposes requires building multiple virtual overlay topologies on top of the physical network. Overlays also abstract virtual networks, hiding the complexity of underlay networks that extend from data center fabrics to the cloud. And, because traditional overlay technologies are constrained by a narrow set of requirements, multiple overlay technologies are often needed on a single network, leading to even more complexity.

The Juniper Networks EVPN Solution

To reduce this complexity, Juniper Networks—working with the Internet Engineering Task Force (IETF)—has developed a solution designed to support a wide range of common requirements, including all overlay use cases. This solution is called Ethernet Virtual Private Networks, or EVPNs.

To convey the simplifying power of EVPNs, let's begin with common network underlay technologies, which are primarily built on Ethernet technology. Legacy underlay fabrics typically operate as a single switched broadcast domain, while overlays provide L2 and L3 virtual topologies to connect services to the network.

Modern underlays, by contrast, use a leaf-and-spine topology and support either routed IP or label-switched MPLS forwarding. Multiple overlays must be deployed on both of these fabric types to support L2 virtual topologies (such as OVSDB-VXLAN or VPLS) and L3 (IP-routed) virtual topologies.

EVPN, however, can support both IP-routed and L2 switched virtual topologies on top of Ethernet, IP, and MPLS fabrics, replacing virtual private LAN service (VPLS). EVPN also supports extended overlay topologies, connecting multiple data center, edge computing, and cloud deployments with a single overlay of virtual topologies.

Juniper's EVPN implementation is based on a multiprotocol version of the Border Gateway Protocol, or MP-BGP. This open standard works across vendor implementations, allowing operators to build large-scale multivendor networks. Using MP-BGP also allows EVPN to scale to hundreds of thousands of nodes and millions of reachable destinations—truly an Internet-scale technology. Operators can also use scaling and simplification techniques such as BGP route reflectors to increase scale, manage operational complexity, and improve deployment speed.

Some of the specific overlay support capabilities are shown in Figure 1.

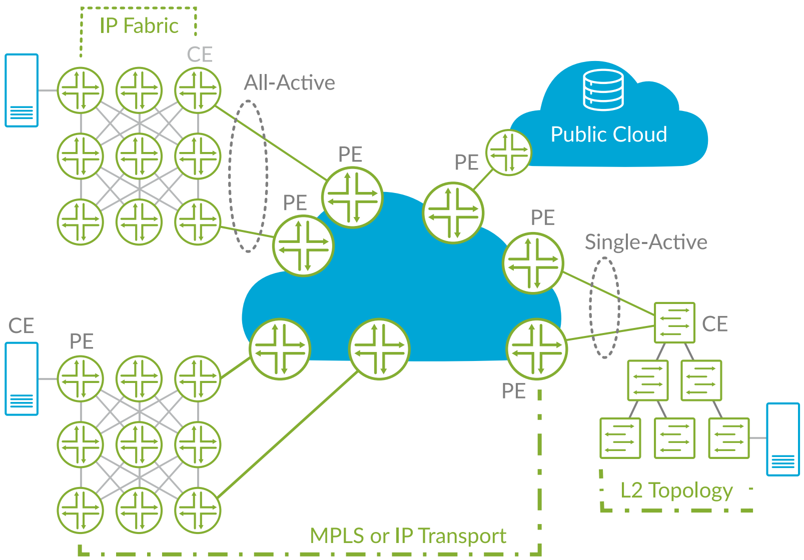

Figure 1: Overlay support technology options being considered for Juniper's EVPN solution

Features of the Juniper EVPN Solution

Juniper's EVPN implementation supports all-active connections between customer edge (CE) and provider edge (PE) devices. In the scenario depicted in Figure 1, an IP leaf-and-spine fabric is connected using an all-active connection to share bandwidth across multiple links as well as resilient connectivity for high-volume loads in the critical business path.

Juniper's EVPN solution also supports single-active connections, which enable fast failover for loads in the critical business path. In Figure 1, a single-active connection is used to connect an L2 data center topology. Virtual devices supporting EVPN connectivity can also be placed in cloud instances; in Figure 1, an EVPN connection connects on-premises loads to a database or data analytics engine located in a public cloud. This approach can also be used to combine multiple edge compute instances to support edge computing.

Mobility is also supported by Juniper's EVPN implementation for both L2 and L3 connected loads. A virtual instance or container can be moved from one topological location in the network to another—say, from an IP fabric to a fabric across which EVPN has been extended—without having to modify the IP address or change the logical L2 connectivity.

Scaling is further enhanced through mass withdraw. In many control planes, when a single link is lost or removed, every destination reachable by that link has to be removed separately (in fact, this is the normal route withdrawal process for MP-BGP). EVPN, however, adds a new convergence mode that proactively removes all destinations reachable through any link that is lost or removed, dramatically improving convergence characteristics.

EVPN supports a distributed anycast gateway, which allows devices connected through an L2 domain to learn a single L2 address in order to reach all devices connected to the network. The distributed anycast gateway feature also forwards traffic along the shortest optimal path through the network, regardless of where the traffic is entering or exiting. Traffic optimization extends this concept to the delivery of all inbound and outbound traffic; rather than forcing flows to follow a trombone path through the network to reach a device with the necessary forwarding information, EVPN distributes forwarding information to network nodes, allowing each to forward traffic along an optimal path.

Technologies with this broad range of capabilities are typically difficult to deploy and manage; the potential for driving up operational costs is a real concern for network operators. The ability to support virtually every kind of overlay using a single technology, however, reduces overall network complexity by replacing many different technologies with a single EVPN deployment. Further, the deeply integrated automation capabilities of Juniper Networks® Junos® operating system, along with the ability to run Junos OS throughout the network, empowers operators to deploy EVPN to solve a wide array of problems quickly and efficiently.

Conclusion

Juniper's EVPN deployment represents a modern, open, standards-based, and automation-native control plane that solves multiple technology problems, helping organizations move and handle data across legacy, on-premises, and cloud-based processing. Juniper's participation in the design, standardization, and implementation of EVPN technology demonstrates we have the skills to solve the world's most complex problems, illustrating Juniper's guiding principle to drive engineering simplicity.

Healthcare and Multicloud

Introduction

Nearly every healthcare organization in today's information-centric world is required to send and receive a wide variety of time-sensitive confidential information. This includes protected health information (PHI), sensitive financial details, images, videos, and other important data which must be accessible at the point of care. Naturally, the sharing of this information must adhere to strict regulatory guidelines and must meet privacy standards to ensure that it is being transmitted safely, securely, and is available at all times.

Cloud adoption is accelerating at a rapid pace throughout the healthcare industry as providers, payers, and life science organizations embrace cloud technology, with its ability to provide easy, secure, and uninterrupted access to clinical and administrative applications and data sets. While most organizations have altered their stance with respect to where data must reside, the conversation is now shifting to how the organization can increase collaboration while ensuring the secure exchange of data across the continuum of care to improve the quality and overall healthcare experience.

Healthcare in the Cloud Era

Managing multicloud environments in healthcare can be extremely challenging. Due to the sensitivity of clinical and financial data, hybrid cloud strategies are very common, as healthcare organizations continue to rely on a mix of public cloud services combined with private clouds.

Healthcare has historically employed an aggressive isolation and containment strategy to deal with complexity. In response, organizations often create domains with hard boundaries around the data center, campus, and backbone, or between clinical and administrative applications. These hard boundaries are precisely why it is so difficult to strike a balance between contexts and limits; they also prevent the distribution of workloads across geographical locations. Hard boundaries also make it difficult to get clear visibility across multiple domains, where a clear line of sight helps with management, orchestration, and control.

In any industry, but especially healthcare, network complexity is an inhibitor to progress. The more complex a network is, the greater the risk of downtime and excessive mean time to repair (MTTR). Healthcare is an industry that doesn't shut down; it needs to be up and running all day, every day. Reducing management complexity is critical.

Juniper Technology to the Rescue

In the use case of a primary data center and a disaster recovery site, Ethernet VPN (EVPN) can be used to connect the two data centers using a Layer 2 bridge, enabling active/active replication and the ability to move virtual workloads as needed. With EVPN, organizations can leverage an open multivendor protocol and have the data center and disaster recovery site in the same address space, which allows for much faster failover should it be needed. While it has long been possible to stretch L2 segments using other technologies, EVPN introduces several capabilities that make doing so much less risky. EVPN is not limited to the data center; it crosses traditional domains to offer a standard that can be also be used in the campus, WAN, and metro area network.

EVPN appeals to the risk-averse nature of healthcare because it is a resilient IP fabric—literally the technology that runs the internet. There is no single point of failure in architectures that implement chassis-style fabrics. Resiliency is achieved at the fabric level, with less reliance on complex per-device reliability.

Another emerging use case involves offering a new cloud infrastructure service model for hosting mission-critical applications like electronic health records (EHR) and imaging, as well as business-essential systems like e-mail, file shares, or any other healthcare administrative workloads for medical centers or hospitals.

Virtual Extensible LAN (VXLAN) enables network segmentation on a greater scale than traditional VLANs and plays a critical role in helping healthcare organizations meet stringent privacy and security requirements. By leveraging VXLAN and a common cloud infrastructure, customers can quickly and safely move workloads between their private data centers and the new cloud service to increase flexibility, security, and overall IT agility.

Summary

The healthcare industry should view refresh and expansion events as opportunities for transformation. Committed organizations will use these opportunities to do two things: make progress toward deploying better technology and avoid making any decisions that unnecessarily limit the number of paths forward. In this way, the set of decisions and changes builds towards a true multicloud infrastructure where teams can manage policies and resources as a whole.

WCU
WEST CHESTER UNIVERSITY

SUMMARY

Company:
West Chester University

Industry:
Education

Business Challenges:
Deliver on student and faculty expectations for connectivity with fast, reliable data center and campus networks

Technology Solution:

- MX480 5G Universal Routing Platform
- QFX10002 and QFX5110 Ethernet Switches
- EX9200 line of Ethernet switches
- Advisory Services and Implementation Services

Business Results:

- Met student and faculty expectations for on-campus digital experience
- Easily scaled the network using EVPN-VLAN in data center and campus networks
- Migrated to new data center network 4X faster than previous upgrade

Whether aspiring poets, doctors, or rocket scientists, students expect an exceptional digital experience. Digital enhances learning, engagement, and collaboration, and it prepares students for the workplace. The increased importance of a comprehensive digital experience is quite clear to West Chester University, a public university in Pennsylvania. The university turned to Juniper Networks routing and switching to build a fast, flexible, at-scale network from data center to campus.

"Students, faculty, and staff require ubiquitous campus-wide network access which is highly available, resilient to unexpected failures, and is easily adaptable to meet their ever-growing coursework, research, and entertainment demands," says Kevin Partridge, assistant director of IT infrastructure services at West Chester University.

Located just outside of Philadelphia in suburban Chester County, West Chester University (WCU) is the largest of Pennsylvania's state universities, with 17,500 students engaged in more than 180 programs of study. WCU is tied for 17th place on the list of Top Public Schools compiled by U.S. News & World Report.[1]

High Expectations

As at most universities, WCU students have grown up online and expect the same levels of technology on campus as in their personal lives. Students and faculty alike need easy access to course materials, video lectures, cloud-based productivity tools, and much more. To deliver on that experience, WCU decided to refresh its network in phases.

WCU first deployed Juniper at the Internet edge, ultimately expanding Juniper networking across its data center and campus networks.

"The partnership is as important to WCU as the infrastructure," says Partridge. "We felt that Juniper would support us in any way we would need, not just at the time of the initial purchase, but all the way through the refresh cycle and in post- sales support."

1 "Top Public Schools, Regional Universities North," Weekly News & World Report, www.usnews.com/best-colleges/rankings/regional- universities-north/top-public

The university deployed the Juniper Networks® MX480 5G Universal Routing Platform for high-performance, reliable, and adaptable connectivity to the Internet.

"We prefer standards-based protocols, and Juniper's implementation of EVPN-VXLAN is very clean."

- Shaun Spence, senior network engineer, West Chester University

Transforming the Data Center

The IT team then turned its attention to the data center. "We wanted a lean, agile architecture for the data center," says Shaun Spence, senior network engineer at WCU.

The server team was planning to migrate its traditional server and SAN-based storage environment to a new hyperconverged infrastructure, with integrated compute, storage, and virtualization technologies. This migration demanded the convergence of both storage and data networks. "We needed a network that could support both storage and data networking and be as robust as possible," Partridge says.

"With the help of Juniper Networks, our agile network architecture continues to support how students learn and faculty teach. This has supported our expansion of online learning in addition to our traditional on- campus learning."

- Kevin Partridge, assistant director of IT infrastructure services, West Chester University

The goal was to simplify and modernize the network, moving from a traditional architecture to a collapsed network for greater scale, improved resilience, and lower latency across devices. "We wanted to improve our dual data center network architecture," Spence says. "With resiliently designed Layer 2 networks, there's generally the introduction of loops, requiring the use of a protocol like spanning tree. This has its drawbacks and is very inefficient. Using EVPN and VXLAN, we were able to provide the required Layer 2 functionality for virtualized server connectivity between our data centers while efficiently utilizing a full mesh of Layer 3 links without the risk of network loops."

WCU chose the Juniper Networks QFX10002, a fixed- configuration 10/25/40/100GbE switch, for its spine-layer switches and the Juniper Networks QFX5110, a high-performance, agile 10/25/40/100GbE switch, for leaf-layer switches.

WCU uses Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN), supported in the Juniper Networks Junos® operating system, as a framework for managing its data center network. "In talking with higher ed colleagues, I was turned on to the idea of EVPN-VXLAN," says Spence. "We prefer standards-based protocols, and Juniper's implementation of EVPN-VXLAN is very clean."

An EVPN-VXLAN architecture supports efficient Layer 2/Layer 3 network connectivity with scale, simplicity, and agility, while also reducing OpEx. EVPN-VXLAN enables WCU to deploy a much larger network than is possible with traditional Layer 2 architectures. And with EVPN-VXLAN, WCU can easily add network capacity without having to redesign the network.

As Partridge explains, "EVPN-VXLAN allowed us to expand our services, such as turning up new Wi-Fi networks to support IoT or BYOD devices, in an efficient, controlled way without impacting users."

The data center network refresh was completed in less than two months—four times faster than the previous upgrade. "The ability to move from one infrastructure to another was greatly enhanced," says Partridge. "We were able to introduce new compute and storage with the underlying Juniper foundation without a massive amount of time and effort."

Building a Scalable Campus Network

With a successful migration to Juniper networking on the Internet edge and data center, WCU deployed Juniper Networks EX9214 Ethernet Switches as a programmable, flexible, and scalable foundation for its campus network core. EVPN-VXLAN plays an important role on the campus, quickly and efficiently backhauling traffic from the wireless controllers to the data centers. "Using EVPN-VXLAN is completely transparent tothe wireless controllers, and we don't need to worry about broadcast storms anymore," Spence says.

Students commonly bring multiple devices—phones, laptops, tablets, gaming consoles, and wearables—to campus, and the number of IoT devices is exploding. The sharp increase in the number of devices made support for IPv6 critical. "With IoT and BYOD, we have more devices connecting to the network, and IPv6 is the answer to that,"

"Juniper's commitment to IPv6 and IPv4 was paramount in our decision. We see higher education as a driving force for IPv6. All of our client networks, including wireless, support IPv6 and we're currently extending IPv6 connectivity into our server networks."

"EVPN-VXLAN allowed us to expand our services, such as turning up new Wi-Fi networks to support IoT or BYOD devices, in an efficient, controlled way without impacting users."

- Kevin Partridge, assistant director of IT infrastructure services, West Chester University

Simplify Network Operations

The WCU team worked with Juniper Advisory Services and Implementation Services on the border routing and data center projects, but quickly became self-sufficient. "It was good to have Juniper's assistance," says Spence. "We had the same system engineers from the border router implementation. They helped us optimize our design."

The network team quickly acclimated to Junos OS, the single operating system that powers Juniper's portfolio of networking and security products. "Junos operating system was very simple to pick up because of my background in programming," says Spence. "Junos is second nature to me."

Ready for the Digital Generation

Technology has fundamentally transformed how college students live and learn, but Partridge is confident.

"The demand for bandwidth continues to expand, and we're able to keep up because of the investments we've made into products which are modular, can be expanded over time, and have software that's continuously improved," he says. "With the help of Juniper Networks, our agile network architecture continues to support how students learn and faculty teach. This has supported our expansion of online learning in addition to our traditional on-campus learning."

GIVERS VS. TAKERS: SEPARATING THE LEADERS FROM THE FOLLOWERS IN NETWORKING

MIKE BUSHONG
VP, Enterprise & Cloud Marketing
Juniper Networks

When the future is certain to be multivendor, how do we collectively go about making technology work?

It's a pretty basic question but it gets to the heart of how technology is developed and propagated throughout the industry. It's not enough to be innovative, developers must do so in a way that facilitates mass adoption by not only customers but competitors too.

Community

The notion of community is straightforward: a unified body of individuals who share a common interest or goal. On the surface, the very nature of business competition seems to run counter to this notion of community. Anyone creating a product or a service naturally wants the target market to develop an affinity or preference for their particular offering. So long as there are multiple suppliers, various communities will emerge each supporting a different solutions. One community splits into many.

There is nothing wrong with this as competition generally benefits the customer.

Except when it doesn't.

Interoperability

When technology solves a particularly difficult problem, it really doesn't matter how competing offerings interact. Merely solving the problem—no matter how unique the solution—is enough to satisfy the market and hopefully win share.

Alternatively, when segment becomes so stable that all solutions begin to look the same, you get a similar dynamic. The building blocks become a commodity and interoperability follows (at a macro level anyway; in networking, there are many issues that occupy the grey middle ground of standards, of course).

But what happens when a technology's reach cannot be contained and the problem space requires change? At these moments, interoperability isn't guaranteed. Worse, the natural competitive nature of building products creates a disincentive to collaborate.

Company or Industry?

This scenario actually runs counter to what industries such as IT require. The days of single-vendor deployments should be a thing of the past. Whether it's the risks inherent in staking the future of your network on a single vendor or the lack of procurement leverage in the absence of other solutions, no enterprise would intentionally choose a locked-in path.

The natural conclusion is that technology needs to move forward in a way that can be proliferated across multiple vendors. The very thing vendors want - differentiation - is at odds with what users need.

So whose interests come first?

The answer varies wildly depending on the supplier. Some have literally built up entire industries developing technology and then contributing back to the communities they serve either through standards body work or open-source projects. Fundamentally putting the needs of the many ahead of the wants of the few.

Others simply don't contribute. They cannot or will not take a leadership position. This strategy can be quite effective if the building blocks are already well known. But it comes at the expense of the very users the company is trying to court.

EVPN, Standards and Moving Forward

That brings us to today.

Data center architectures have been in flux for the past several years. The industry is in the midst of a transition; in a bid to simplify data center operations, EVPN-VXLAN has emerged as the go-to standard. Given its central role in the evolution of enterprise networking, it is critical that EVPN be both capable and ubiquitous. This is why some of the industry's brightest have come together to work through standards bodies to help define EVPN as it forges its way forward.

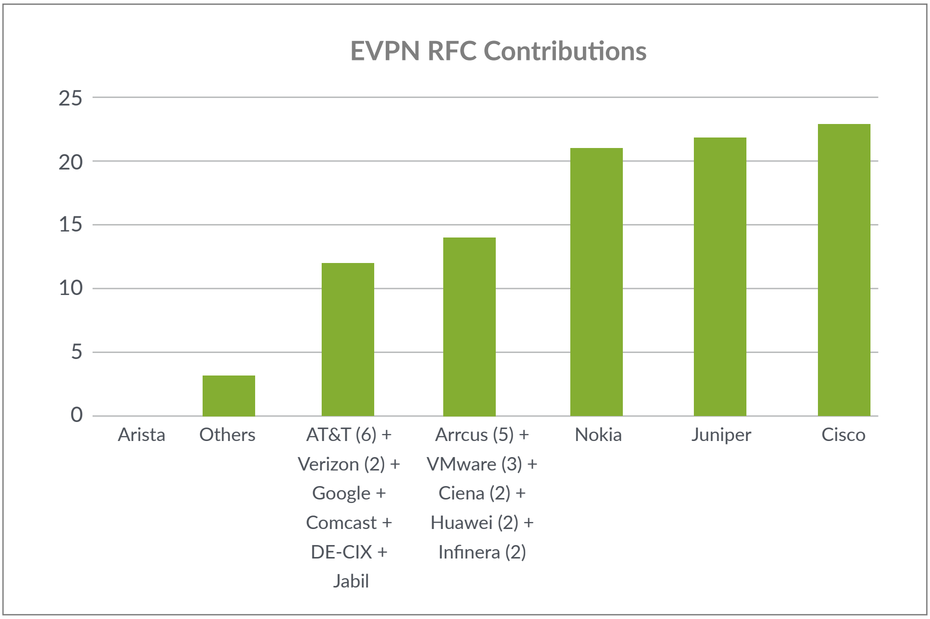

Networking vendors included are leading data center switching vendors as Juniper estimates based on Dell'Oro Group 4Q18 Ethernet Data Center report.

Juniper has always been a strong advocate for opening up key technologies. We have played central roles in virtually every major networking protocol over the past 20 years and it's no different with EVPN.

Our participation is a reflection of both our ability and our contribution. You cannot lead if you do not know how to move forward. A leader without followers is no leader. That Juniper has both contributed and attracted others to EVPN demonstrates our leadership in the space.

As enterprises look to first learn and then deploy EVPN, they are going to need leaders. Today, we kick off a series of content designed not to sell but rather to illuminate. With our GetSmart series on SDxCentral, we will look at why EVPN matters, providing an industry view of what EVPN can do as the de facto architectural foundation for the modern data center. Empower your leadership, learn the benefits of EVPN-VXLAN.

EVPN BUILDING BLOCKS

Leadership Discussion Series

Moderator

RAHUL AGGARWAL
Founder & CEO
Augtera Networks

KEITH TOWNSEND
Co-founder
The CTO Advisor

ALDRIN ISAAC
Senior Director, Portfolio Product
Juniper Networks

Abstract

After covering the foundation of EVPN-VXLAN in the first session, this one will cover EVPN building blocks with an emphasis on the data center. RFC authors themselves will discuss what network virtualization and EVPN is, how it simplifies the management of your infrastructure and how to increase the performance and efficiency of your applications and workloads. You'll also find out what EVPN-VXLAN business outcomes a CTO is looking for and how network virtualization can resolve typical problems, especially in a multi-vendor data center.

Key Takeaways

- EVPN is the defacto standard solution for network virtualization over L3 with broad adoption across the networking industry and major market segments.
- EVPN solves network virtualization using broadly tested technologies—IP networks as its base and BGP for signaling infrastructure.
- EVPN brings the attributes of IP routing and forwarding to L2 by leveraging BGP to distribute MAC addresses and by using IP or MPLS based tunnels (ex. VXLAN) for transport.
- EVPN supports several use cases for L2 only, L3 only and L2-L3 integrated environments at scale.
- EVPN removes the requirement to configure the data center fabric when connecting servers within or between data centers allowing changes without requiring changes to the data center fabric.
- EVPN enhances VXLAN exchanging Data Plane Learning in favor of Control Plane Learning
- EVPN opens networks to N-way multihoming, as opposed to 2-way with MC-LAG based technologies, thus increasing reliability.
- SDN is not an either-or with EVPN—it can be both which includes multi-vendor SDN and the ability to bring in different solutions to operate this multi-vendor technology.
- IP fabrics are the future for data centers. Proprietary solutions solve common data center problems but do not offer the complete solution to network virtualization. EVPN is different because it supports L2, L3 or a mixed set of requirements in an IP Fabric, whereas STP, MC-LAG, and even TRILL do not work well with an IP transport.

"What I'm seeing in the market is given the enormous value IP fabrics bring, if people have an existing L2 infrastructure or L2 fabric, they are deploying an IP fabric in a smaller footprint in one POD, connecting the new workloads to that, gradually migrating existing workloads if they can, to the new fabric and of course the fabric grows over time."

Rahul Aggarwal, Augtera Networks

TELEHOUSE

KDDI

SUMMARY

Company:
Telehouse America

Industry:
Web Services

Business Challenges:
Enhance data center services and simplify operations by adopting network automation, EVPN, and VXLAN, and a highly flexible data center fabric.

Technology Solution:

- QFX5100 line of Ethernet switches

Business Results:

- Enhanced agility of multitenant data center environment
- Reduced network changes from hours to minutes
- Simplified operations withnetwork automation

AT TELEHOUSE AMERICA, A FEW ELITE ENGINEERS OPERATE NEW YORK'S LARGEST DATA CENTER

Telehouse America, which offers carrier-neutral data center services in the U.S., opened its first data center on Staten Island in New York City in 1989. Today, it operates two massive data centers in New York as well as 46 other data center facilities around the world, including London, Paris, Los Angeles, Hong Kong, Seoul, Singapore, and Vietnam. Telehouse also offers international exchange (IX) services and connectivity with more than 750 carriers and ISPs. A subsidiary of KDDI, a Global Fortune 300 Japanese telecommunications company, Telehouse has more than 3000 customers.

Business Challenge

With a history of innovation and growing demand for data center services, Telehouse America wanted to enhance the network connecting its two New York facilities. The company is known for its rock-solid infrastructure and exceptional customer service. The Staten Island facility has provided uninterrupted service for nearly three decades—including during power failures caused by an unprecedented natural disaster. Telehouse America's Chelsea location is a central network infrastructure facility, and is the location of Telehouse NYIIX, the largest peering exchange in New York and one of the top 10 exchanges in the world.

“In the future, the importance of networks will undoubtedly continue to increase,” says Akio Sugano, vice president at Telehouse America. “Establishing powerful connectivity inside and outside of the data center and providing a stable network is an important mission that we have been tasked with.”

There were three firm requirements for the new network: The use of Ethernet VPN (EVPN) and Virtual Extensible LAN (VXLAN) technology to accommodate a multitenant environment; an Ethernet fabric to reduce the operational load; and an API to automate network control.

Reducing the workload of the operations staff while expanding capabilities was critical. A small team ensures highly stable data center operations at Telehouse America, but as Sugano notes, network engineers are scarce and they do not have sufficient manpower in the face of rapidly growing customer demand.

“Junos OS is stable and easy to use, and our engineers can quickly configure the network without worrying about detailed conditions.”

- Gregory Grant, Manager of Operations, Telehouse America

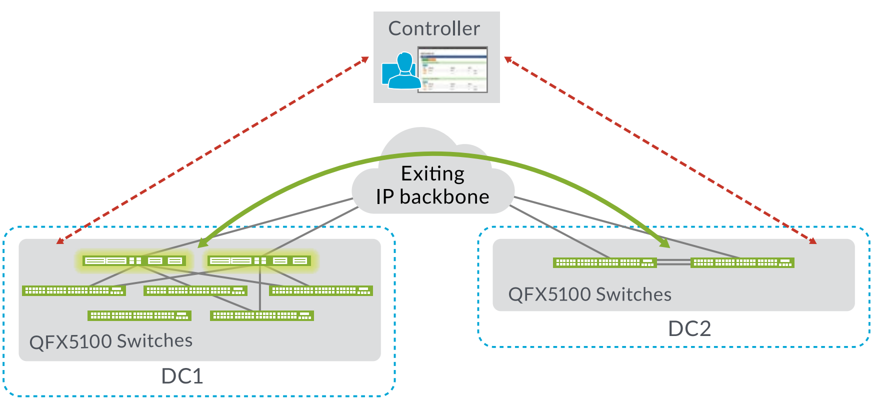

Figure 1: Multi-tenant L2 overlay with EVPN/VXLAN technology

Technology Solution

Sugano determined that automating network configuration, along with an Ethernet fabric foundation, would reduce the operational load within the multitenant environment.

Telehouse America selected the Juniper Networks® QFX5100 line of Ethernet switches as the network solution that satisfied all of its requirements.

"We tested various devices, but only the QFX5100 switch enables the proper level of EVPN/VXLAN functionality we expected," says Gregory Grant, manager of operations at Telehouse America.

Telehouse America leveraged EVPN/VXLAN technologies on the QFX5100 switches to scale out its multitenant network. A Layer 3 IP-based underlay network, coupled with a VXLAN/ EVPN overlay network, enabled the company to deploy a much larger network than with a traditional Layer 2 Ethernet architecture. By decoupling the virtual topology from the physical topology, servers and virtual machines can be placed anywhere in the network and remain connected to the same logical L2 network.

With the QFX5100, Telehouse has a low-latency, highperformance fabric architecture that can be managed as a single device. Multiple QFX5100 switches can be interconnected into a single logical device, vastly simplifying management and operations.

Juniper Networks Junos® operating system was also a factor in the decision. "Junos OS is stable and easy to use, and our engineers can quickly configure the network without worrying about detailed conditions," says Grant. "We also focused on the extensive functions for data center network management and the high performance and functionality that each provides."

Grant appreciates the many ways the Junos operating system simplifies operations: "Junos OS is commitment-based, so the work can be done without stress, and no mistakes or problems occur. I feel that Junos OS is a very important element in providing stable services," he says.

Grant uses the Junos OS REST API to automate network configuration changes and other network operations. Based on open-source technology, the Junos OS API offers a high level of freedom and flexibility.

"We tested various devices, but only the QFX5100 switch enables the proper level of EVPN/VXLAN functionality we expected."

- Gregory Grant, Manager of Operations, Telehouse America

Business Results

Telehouse America built more than a network to connect its New York facilities. It reduced the operational load on its engineering team by leveraging EVPN/VXLAN and the Junos OS API—while enhancing its services.

Its previous network had a traditional three-layer design, and every time a customer made a request, the configuration work had to be performed within the physical environment. Changes to the network configuration could take several hours or even up to half a day—a heavy load for just a few engineers.

Now when the configuration of a single switch is changed via software, that change is propagated across the network. The operational load has been greatly reduced, and the team says that work that had previously taken a half a day can now be completed in just five minutes.

A simpler, more agile network, powered by Junos OS, enables a small engineering team to be more effective and efficient as the business evolves. "A data center must adapt to user needs and environmental changes. For that reason, there are some major difficulties when it comes to securing human resources," says Grant. "Junos OS is an operator-friendly network OS that enables staff who do not possess advanced technical skills to control the network relatively easily."

"The Junos OS API is very convenient, and we've developed our own monitoring tool to help reduce operating costs," says Grant. "I like that it can be flexibly configured according to detailed technical requirements. I hope to acquire new knowledge by converting management information into database format and tying that into advanced business applications such as sales and marketing."

Next Steps

Currently, the QFX5100-based EVPN/VXLAN infrastructure is comprised of the two facilities in New York. Telehouse America has expressed its intent to expand the EVPN/VXLAN to its Los Angeles facility, home to Telehouse LAIIX, which provides peering exchange to the Pacific Rim. It also has set its sights on rolling out the new network to the Telehouse Group in Europe and the Pacific Rim.

That deployment can be accelerated, now that Telehouse has standardized on Juniper. "We hope to increase the quality of our services by actively adopting cutting-edge technologies," says Sugano. "There are many promising technologies, but in each case, standardization takes a long time. Juniper Networks was quick to adopt advanced open-source technology and is proactively engaged in standardization. I hope they will continue to maintain both the historic network OS and advanced opensource technology."

With Juniper, Telehouse and Telehouse America have built the digital foundation of cutting-edge, robust services that enable customers around the world to serve their customers and accelerate their own digital transformation .

Collapsed EVPN: Juniper's Approach

EVPN-VXLAN Overview

Operating a network as an IP fabric not only provides greater scale and resiliency, it also offers Layer 3 connectivity to applications. When Layer 2 is required, Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN)—an industry standard that allows networking teams to build fabrics without relying on proprietary protocols—provides end-to-end Ethernet connectivity on top of the IP fabric. This makes the deployment of multivendor networks possible, enabling enterprises to evolve without unnecessarily stranding existing assets.

EVPN-VXLAN does not dictate whether overlays should be L2- or L3-based. EVPN can blend Layers 1, 2, and 3 in a seamless solution. Operators can leverage EVPN for L3-only overlays without being forced to take L2 Ethernet and L1 multihoming services if they are not needed. Operators can also bind L3 routing services to their L2 overlays in the form of integrated routing and bridging (IRB), which connects bridging tables to routing tables without requiring a physical interface. Logical IRB interfaces allow packets addressed to the default gateway media access control (MAC) address to be routed instead of bridged. Integrated routing can be performed at a "centralized" network node or directly on all of the network nodes to which tenant end systems are connected. The former is referred to as centrally routed bridging, while the latter is referred to as edge-routed bridging (ERB).

Benefits of EVPN-VXLAN

- **Lower operating costs:** Highly scalable support for both L2 and L3 services reduces the number of protocols the enterprise needs to operate the network, making it simpler and more reliable.
- **Higher service availability:** EVPN doesn't need to flood broadcast, unknown unicast, and multicast (BUM) traffic in Ethernet segments and includes Address Resolution Protocol (ARP) suppression, eliminating unnecessary traffic replication, improving scalability, and reducing traffic noise for more efficient monitoring and troubleshooting.

- **Multihoming and better ROI:** First-class support for multihomed attachments at both L2 and L3 with fast failure recovery ensures a more efficient use of resources by eliminating standby links. It also eliminates the need for multichassis LAG protocols, leading to better ROI and higher service availability.
- **No MPLS license required:** EVPN-VXLAN can operate over MPLS or IP networks, using VXLAN encapsulation in the latter to transport Ethernet frames.

Technical Solution: Collapsed EVPN

Generally speaking, this solution involves two or more spine switches (for the purposes of this example, Juniper Networks® QFX10000 Ethernet Switches) configured in an EVPN fabric. Leaf switches, which don't have to be EVPN devices, can be configured either as standalone platforms or in a Virtual Chassis configuration with multichassis link aggregation group (MC-LAG). Routing between the bridged domains is performed on the spine switches, while endpoints connect to leaf switches via either a bridged domain with Spanning Tree configured or LAG VLAN trunks. The EVPN fabric is connected to the external world (northbound) through either IP equal-cost multipath (ECMP) or LAG using direct links between the external router gateway(s) and each spine node. Spine devices can advertise host routes for attached endpoints to northbound neighbors, providing one mechanism for endpoint IP mobility.

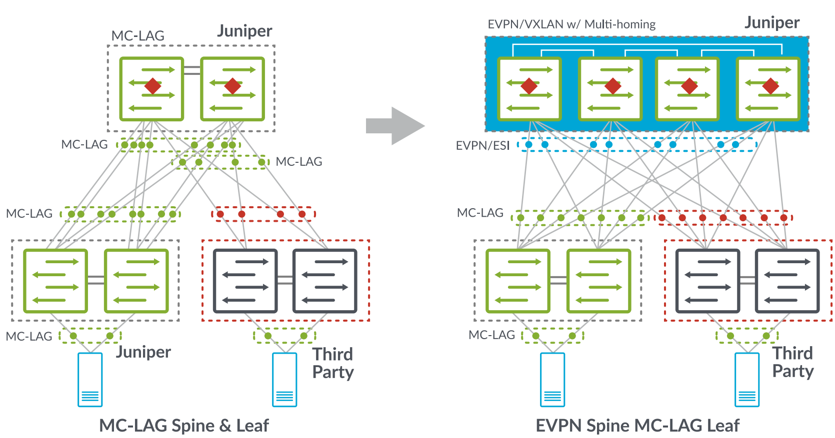

Figure 1: Moving from MC-LAG to collapsed EVPN

This approach is especially useful in brownfield deployments with a mix of leaf node types that may or may not support EVPN.

While larger scale multitenancy is possible in this solution, it is best for small to mid-sized deployments supporting fewer than 100 tenants. Multitenancy allows for administrative segmentation beyond only the subnet or VLAN, creating a more efficient policy enforcement model.

Juniper Networks EVPN-VXLAN

While EVPN technology itself is based on a set of open standards, implementation details will either enhance or dampen its benefits. In Juniper's case, customers benefit from operational simplicity, greater ROI, and future-proofing through:

- Standards-based Ethernet multihoming.
- Consistent EVPN implementations available on all major data center, WAN, campus, branch, and cloud platforms, including Juniper Networks MX Series 5G Universal Routing Platforms, vMX Virtual Router, QFX Series Switches, EX Series Ethernet Switches, and SRX Series Services Gateways, as well as Contrail vRouter.
- Support for high-scale EVPN deployments when needed down the line through availability on the MX Series and QFX10000 switches.
- Hybrid EVPN configurations that combine physical and virtual platforms to extend EVPN down to virtual machines and containers.
- Unified EVPN configurations and operations, enabled by the Juniper Networks Junos® operating system CLI and API available across all platforms.

Conclusion

Ethernet VPNs offer an industry-standard solution that allows networking teams to build fabrics without relying on proprietary protocols. This in turn makes the deployment of multivendor networks possible, enabling enterprises to evolve without stranding existing technologies that run on the network today. Juniper's implementation offers an efficient and scalable way to build and interconnect campuses, data centers, and public clouds, uniquely positioning the company to realize the full potential of EVPN technology by providing optimized, seamless, and standards-compliant L2 or L3 connectivity for the enterprise.

WHEN OPERATIONS DRIVE ARCHITECTURE—A MISSIVE ON EVPN

MIKE BUSHONG
VP, Enterprise & Cloud Marketing
Juniper Networks

Despite all of the recent talk of automation, it's hard to argue that operations is little more than an afterthought for most enterprise networks. It's not that operations aren't important. It's just something that needs to be figured out after the architecture has been determined.

At least, that's the way it used to be. But that's changing, and in a huge way. Why?

Enterprise IT as a Provider of Services

There is far too much baggage with the term "service provider", so for the purposes of this blog, I will use the more difficult term "provider of services." Enterprise IT has been on a steady march from largely an enabler of the business to a provider of services. As technology moves from a supporting function to an integral part of the offering, IT moves from an entity that people complain about to part of the product supply chain. And enterprises treat their supply chains differently.

The supply chain has to be efficient; when it's not, costs soar, margins shrink and profits dwindle. The supply chain also has to be reliable; when it fails, the production line grinds to a halt. And the supply chain has to be repeatable; it's not about building one product for one customer—it's about building a pipeline of products for a market.

Manufacturing operations are a science. And as IT joins the mainstream and becomes a more traditional manufacturing process, it should come as no surprise that operations are following a similar track. And the end result is to make IT a more efficient provider of the very services needed to develop and ship technology products.

Learning from Traditional Service Providers

As enterprises become providers of services, they can learn from service providers who preceded them and have been operating this way for decades.

Service providers elevated operations long ago. Operational expenses, not capital outlay for equipment, dominates service provider economics. When your customer base is expansive and has a myriad of needs, the underlying infrastructure naturally becomes more diverse. The key to service provider economics is managing this diverse infrastructure with a uniform operating model.

Building blocks that abstract the lower layers of the service deliver the common operational platforms on top of which the business can be built. They are critical components of any would-be solution.

Operational Uniformity in the Data Center

In many ways, enterprises have it worse than service providers. At least service providers have control over the product they offer, whereas enterprises must answer to lines of business that have traditionally been the power brokers in enterprise IT. This makes it virtually impossible to remove things, forcing enterprises to drive toward the "new" while simultaneously supporting the old.

How do you create operational uniformity when you are forced to keep decades-old technology alive? For these cases, there needs to be a bridging technology capable of supporting the old while enabling the new. In the language of enterprise data center networking, it's about providing secure connectivity for Layer 2 applications while building an infrastructure that can efficiently handle Layer 3 applications.

EVPN as a Bridge

EVPN is that bridge between the old and the new.

Originally developed for service providers, EVPN was about providing secure connectivity to virtualized compute resources that connect at the edge of the network. The key to efficient operations was that regardless of whether an application required Layer 2 or Layer 3, operators could leverage a single protocol with a common interface.

EVPN abstracts the underlying transport, which is exactly why it is so well suited to provide a bridge from old to new. Even moderately sophisticated enterprises have scores of applications.

As they upgrade, rewrite and replace these applications, enterprise IT is forced to operate in a hybrid environment. Without a common protocol to unify the operating model, there is simply no path to becoming an efficient part of the technology supply chain.

What Does EVPN Do?

EVPN basically gives enterprises the ability to run an overlay across an enterprise-wide L3 network. It provides the control plane required to route between virtual segments while maintaining independence from the physical underlay network. Within that virtual network, applications can be separated, effectively creating a multi-tenant domain managed through the administration of distributed network policy.

If that's what EVPN does, what does it enable?

- Highly scalable support for both L2 and L3 services, allowing it to be used in the data center, data center interconnect (DCI), WAN, campus and metro. EVPN reduces the number of protocols the enterprise needs to operate the network, making it simpler and more reliable, lowering operating costs and improving service availability.
- EVPN removes the need to flood broadcast, unknown unicast and multicast (BUM) traffic through stretched ethernet segments. This suppresses ARP, eliminating unnecessary traffic replication, improving scalability and reducing traffic noise for more efficient monitoring and troubleshooting, leading to higher service availability.
- EVPN provides first-class support for multi-homed attachment at both L2 and L3 with fast-failure recovery, allowing more efficient use of resources by eliminating standby links. It also removes the need for multi-chassis LAG protocols, leading to better ROI and higher service availability.

The original problem statement and early design considerations for EVPN are well documented in the following standards drafts:

- Problem statement: https://tools.ietf.org/html/rfc7364
- NVO3 draft: https://tools.ietf.org/html/draft-ietf-nvo3-evpn-applicability-01

Juniper Networks and EVPN

The rise of EVPN is consistent with Juniper's belief that simplicity will ultimately be on the right side of change. The key to scaling operations for enterprises is converging on a narrower set of protocols that provide a common operating model despite the diversity with which enterprises must grapple.

EVPN has been developed in the public eye with the explicit objective of remaining open, which is also consistent with the principles upon which Juniper was founded and continues to operate. The future of enterprise IT cannot be a fractured mess of vendor-specific solutions to industry-wide problems.

If the key to progress is simplification, it makes perfect sense to leverage a common operating system—Juniper Networks' Junos® software—to deliver EVPN across the entire portfolio. The future of the data center has been under development for years, both within standards bodies and within Juniper's development teams. As enterprise IT shifts from an enabler to a provider of services, it will be built on a foundation that is purpose-built exactly for this type of transition.

Today, Juniper kicks off the second video in a series of content designed not to sell, but rather to illuminate. With our GetSmart series on SDxCentral, Juniper will look at why EVPN matters, providing an industry view of what EVPN can do as the de facto architectural foundation for the modern data center. For additional information please download the E-Book on EVPN-VXLAN.

EVPN EXPANDED USE CASES

Leadership Discussion Series

Moderator

JORGE RABADAN
Senior Product Line Manager
Nokia

RUSS WHITE
Infrastructure Architect
Juniper Networks

ALDRIN ISAAC
Senior Director, Portfolio Product
Juniper Networks

Abstract

In the first two parts of this series we covered EVPN-VXLAN foundations and building blocks with an emphasis on data center. Industry experts discuss some of the expanded use cases like DCI, EVPN and the public cloud, WAN, campus (multi-tenancy across sites), etc. Hear about support for multiple encapsulation and multicast distribution and it relates to business requirements as well as possible future use cases showcasing the extensibility of EVPN. You'll also find out how EVPN-VXLAN addresses the business outcomes a CTO is looking for, especially in a multi-vendor data center.

Key Takeaways

- EVPN began as an open solution for a couple of use cases. It has since been picked up by several vendors to address challenging problems in more efficient ways than current technologies.
- EVPN with DCI offers the ability to scale across multiple domains while offering policy enforcement, policy boundaries, and reduction of tunnels removing typical scale limitations.
- EVPN replaces multiple technologies with one—it offers multicast using the same control plane as a unicast which simplifies the overall use case.
- Multicast with EVPN—generally creates tenant endpoint state in the network core—with overlay multicast, state is not maintained in the core, offering efficiency and scale.
- EVPN decouples the overlay from the underlay (VXLAN, IP, MPLS) and makes networking seamless regardless of the underlay technology used.
- In the instance that L2 stretch is required, EVPN enhances the experience by getting rid of typical L2 issues (e.g. Turning off traffic broadcasts and removing L2 loops).
- EVPN offers an effective underlying control plane for a service mesh with multicast services because the underlying architecture supports agnostic transport.
- EVPN has been tested for interoperability with more than 11 major vendors working together in a single network.

> *"This concept of universal and extensible means that you can replace many different technologies with one. So even if the single technology may be more complex than any of those one technologies...having the single solution simplifies the entire network in a really important way."*
>
> Russ White, Juniper Networks

EVPN and Data Center Interconnect

Introduction

Network operators will often build multiple geographically separated data centers to improve resilience, increase scale, and to run older fabrics alongside new ones during a transition. Running multiple geographically dispersed data centers requires a technology that allows organizations to interconnect these locations across metro, regional, or even national areas. This is a significant challenge.

The Challenge

Such a solution does exist. Called Data Center Interconnect (DCI) technology, it enables businesses to link their widely dispersed data centers on a common network. It also poses a number of challenges for network operators.

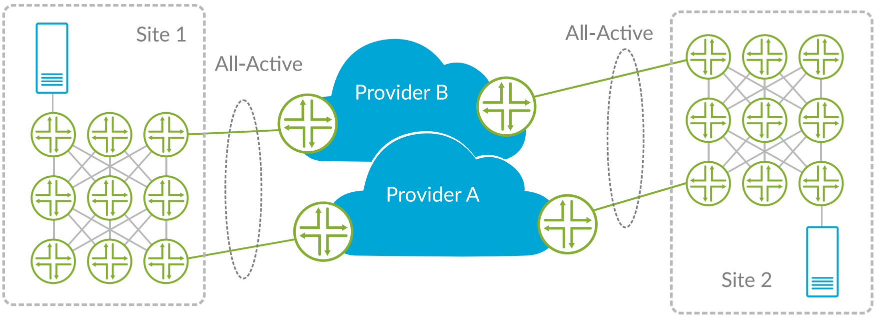

Figure 1: Using DCI to interconnect two geographically dispersed data centers.

For example, services sharing a single subnet on either side of a DCI deployment require some form of a stretched Ethernet segment. Stretching subnets across geographically dispersed data center fabrics with DCI extends failure domains, enabling a problem in one site to cause a failure in another site residing on the other side of the DCI. This can increase mean time to repair (MTTR) and negatively impact application performance and network availability.

Stretching the path across a geographical region between two services, or between virtual machines or containers where services are hosted, causes scalability issues as well. As the number of physical and virtual hosts grows, the amount of reachability information carried in the control plane swells, increasing the load on processors and memory in network devices.

Stretching reachability information, both Ethernet and IP, across geographically dispersed data center fabrics causes scale issues. Data center control planes can grow over time, requiring higher end hardware to support that growth.

Resilience and technology support are also a challenge in DCI deployments. Many different technologies are used to carry traffic between disparate data center fabrics, including Virtual Extensible LAN (VXLAN), MPLS over generic routing encapsulation (GRE), and network virtualization using GRE (NVGRE). Control plane decisions are driven by the available transport layer, rather than by efficiency or simplicity. While resilience requires carrying traffic across multiple providers, this is a complex operation, and if active/standby solutions are used, then organizations must pay for all links reserved, whether or not they are being used.

Finally, many DCI solutions are proprietary, leading to vendor hardware and software lock-in for customers.

The Juniper EVPN DCI Solution

The Juniper Networks EVPN DCI solution overcomes the challenges of connecting geographically dispersed data centers. Ethernet VPN (EVPN) is based on Multiprotocol BGP (MBGP), which advertises both Ethernet and IP reachability information through a single control plane. This means a single control plane protocol can connect data center fabrics within and between subnets without a complex mixture of multiple technologies.

Stretched failure domains at the transport layer are the result of stretched broadcast domains. Combining a network that forwards every broadcast or multicast with hosts that use broadcast for service discovery or basic communication is a recipe for disaster that, should a problem occur, will require a lot of engineering hours to recover from.

EVPN solves this problem by automatically limiting broadcasts and multicasts to topologically isolated segments. For instance, if a host at Site 1 requests the Ethernet address of a host attached to Site 2 using broadcast Address Resolution Protocol (ARP), EVPN will stop the request before it reaches the provider connections, replacing it with an intelligent process for obtaining the same information without broadcasting across the DCI links.

EVPN only sends Ethernet control plane information where it is needed to forward packets. IP reachability in EVPN, like all IP-based control planes, can be aggregated at boundary points. These two mechanisms allow operators with EVPN deployments to easily and intelligently manage their control plane state.

EVPN includes a number of other techniques that reduce control plane load—for instance, single link-failure notification to disable reachability for all destinations through that link, and the use of default Ethernet addresses on topologically isolated portions of a subnet to relieve hosts from having to maintain large reachability tables.

With EVPN, multiple DCI paths—either through diverse links or multiple providers—can be active at the same time, both within and outside a single subnet. This allows traffic to be load-shared through both Provider A and Provider B in Figure 1, enabling both links to add value even when there is no outage to route around. Juniper's EVPN solution supports transport over MPLS, MPLS over GRE, VXLAN, and NVGRE DCI transport, allowing a number of DCI connectivity options.

While some solutions are proprietary, EVPN technology was developed through the Internet Engineering Task Force (IETF) with the cooperation of both providers and vendors. EVPN implementations are not dependent on a single hardware implementation, so this single control plane technology can outlive the hardware and software you deploy today, promoting long-term stability for network management and operations.

Summary

The need to connect multiple, geographically separated data centers is a common one. In the past, available DCI options have been complex and difficult to deploy, they increased the likelihood of domain failures, they strained control plane scalability, and they provided limited support for underlay technologies and resilience.

EVPN , in contrast, is a standards-based, automation-native approach to DCI that solves many of the challenges facing data center operators today. By supporting the design, standardization, and implementation of EVPN technology, Juniper is committed to applying its skills to solving complex problems while demonstrating its intent to deliver engineering simplicity.

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 Evolved Campus solution, based on VXLAN overlay with an EVPN control plane, offers an efficient and scalable way to build and interconnect multiple campuses, data centers, and public clouds.

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

THE EVOLVED CAMPUS

EVPN-VXLAN-Based Enterprise Networks

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 Internet of Things (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 eliminate the need for Spanning Tree protocol. While these architectures work well in small or medium-sized campuses, where services are limited to a single network and cater to traditional campus requirements, they are simply too rigid to support the scalability needs of 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 these same networks are expected to scale rapidly without adding complexity. Since many of these 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.

The Juniper Networks Evolved Campus Solution

The EVPN-VXLAN-based campus architecture decouples the overlay network from the underlay with technologies such as Virtual Extensible LAN (VXLAN) and Ethernet VPN (EVPN). This approach addresses the needs of the modern enterprise network by allowing network administrators to create logical L2 networks across different L3 networks.

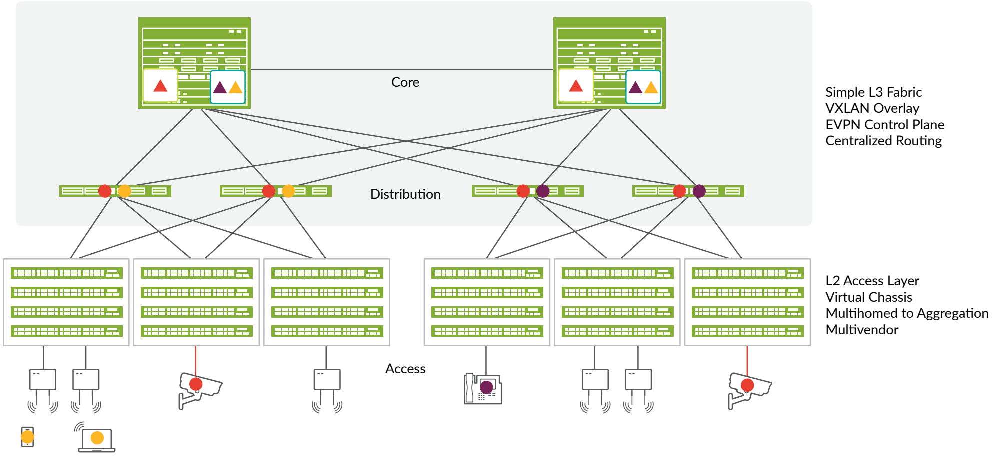

Figure 1: Evolved campus high-level architecture

VXLAN, an encapsulation/tunneling protocol, does not change the flood and learn behavior of the Ethernet protocol. Instead, the VXLAN control protocol—in this case EVPN—uses MP-BGP to allow the network to carry both L2 media access control (MAC) and L3 IP information in the control plane. 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 the flood-and-learn behavior, which is considered suboptimal in several use cases.

The standards-based EVPN solution offers the following benefits when operating as a campus control plane protocol.

- Greater network efficiency
  - Reduces unknown unicast flooding with control plane 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)
  - Supports multipath traffic to active/active dual-homed access layer switches
- 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
- Flexibility
  - Enables easy integration with L3 and L2 VPNs
  - Delivers BGP-based control plane that allows application of fine-grained policy control
- Nonproprietary
  - Supports multivendor core, aggregation, and access layers with standards-based protocols

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 a Layer 3 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.

In an evolved campus architecture (see Figure 1), the core and aggregation layers form a Layer 3 fabric with an EVPN-VXLAN overlay. Ideally, the underlay would be deployed using the L3 Clos model with core and aggregation switches, while the access layer switches would be multihomed to the distribution layer.

The Clos model provides an architecture that enables deterministic latency and horizontal scale at the core, aggregation, and access layers. You can use either an interior gateway protocol (IGP) like OSPF as the underlay or EBGP as the underlay routing protocol; this particular solution uses an IBGP overlay design with route reflection where aggregation devices within a given pod or group, as shown in Figure 1, share endpoint information upstream as EVPN routes to core devices acting as route reflectors. The core devices reflect the routes to downstream aggregationi devices using route reflectors to eliminate the need for full-mesh BGP connections and 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 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 for advanced security enforcement with network segmentation.

Like other Juniper architectures, the evolved campus does not force customers to invest in new devices. The same devices and technologies used in other Juniper architectures can be used in an evolved campus deployment, as shown below.

- Core layer
  - EX9204/EX9208/EX9214 Ethernet Switches
  - EX9251/EX9253 Ethernet Switches
- Aggregation layer
  - EX4600 Ethernet Switches
  - EX4650 Ethernet Switches
- Access layer
  - EX4300/EX3400/EX2300 Ethernet Switches
  - Virtual Chassis technology
  - Non-Juniper or legacy Juniper access layer switches

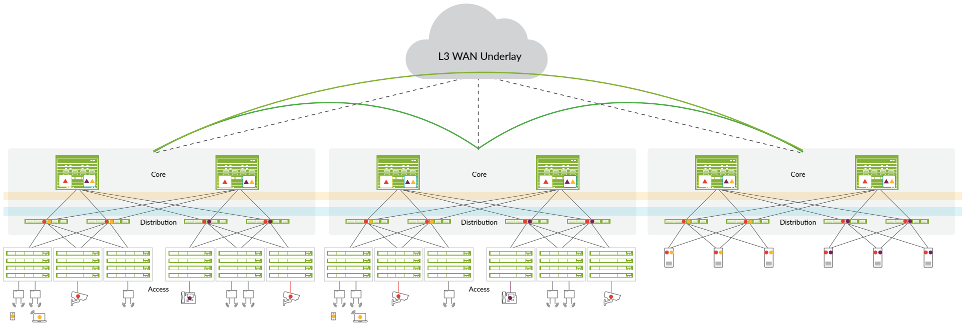

Figure 2: Interconnecting multiple campuses and data centers with EVPN-VXLAN overlay

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. VXLAN is WAN underlay-agnostic provided the campuses, data centers, and the public cloud infrastructure have IP connectivity. EVPN VXLAN overlay can be deployed over a variety of WAN technologies, including private MPLS and IPsec over Internet.

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 Evolved Campus solution, based on 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 platforms—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.

EVPN AND THE FUTURE OF DATA CENTERS

MIKE BUSHONG
VP, Enterprise & Cloud Marketing
Juniper Networks

It is rare when the future of virtually anything is completely divorced from its present. But that is the case with enterprise data centers. Looking at how mainstream enterprise data centers will evolve—both in the near and medium term—it's important to understand that progress will be evolutionary. Despite every technologist's preference, pragmatism dictates that our ties to the past will not be completely severed.

This means that the architectural building blocks responsible for data centers of the future need to be effective at handling both legacy and modern applications, providing a bridge from where enterprises find themselves today to where they need to be tomorrow.

Requirements for the Data Center

First and foremost, it must be acknowledged that most data centers will host legacy applications for the foreseeable future. Despite widespread adoption of cloud-delivered applications, the reality is that there is a long trail of legacy applications that are either not well-suited to the cloud or there is no business justification to modernize. So despite the best of intentions, most mainstream enterprises will find it impossible to completely separate from their past.

Of course, some applications will be rewritten while others will be added. Any application developed today will likely leverage microservices architectures and be built on common application platforms such as Kubernetes and CloudFoundry. For applications hosted in the cloud, the data center architecture is left to the cloud providers, but not all workloads will run in the public cloud. For some applications, infrastructure will be on-premises (as with multi-access edge computing required for distributed IoT workloads, for example). In these instances, the data center is somewhat more nebulously defined, not only for private cloud but also hyperconverged infrastructure running closer to the user. This means the data center of the future must account for diverse deployment models leveraging these application platforms.

Finally, these applications will largely be consumed on-demand. This places the operational burden on the architectures, as user demand will dictate when and where connectivity will be required. As applications become more distributed, data centers will need to support a self-service ability to create connectivity services between endpoints within the data center, which will be a combination of on-premises infrastructure and one or more public clouds.

Operations as the Unifying Problem Space

The unifying problem space in the future of data center networking is operations. Whether it's providing a common operating model for both legacy and modern applications or supporting self-service connectivity services across distributed data centers, the key is converging on a narrow set of operational tools and protocols that serve as an abstraction layer for the underlying infrastructure.

Without a simpler, more efficient and highly-repeatable operating model, enterprises will find that they are incapable of satisfying the disparate needs of the future without compromising on lead times, integrated security and mean time to repair.

EVPN enables this operating model. It allows data center operators to create virtualized L2 and L3 networks that meet the needs of both legacy and modern applications directly on top of the IP underlay, without introducing the additional complexity of MPLS (which was required in the past).

Supporting the Past

For most enterprises, the majority of legacy applications will require L2 connectivity. Where data centers of the past used to feature Spanning Tree, the current best practice for building scalable, highly-available data center underlays is to use L3 protocols. In modern data centers, L2 connectivity services for legacy applications must be virtualized on top of L3 IP connectivity. EVPN can use VXLAN to accomplish this.

EVPN provides a significant step forward, even in L2 environments. There are no forwarding loops in EVPN L2 domains, which greatly reduces the risk of outage due to human error. Additionally, EVPN features a more efficient handling of broadcast and unknown traffic, helping scale VPNs with legacy applications deployed in large L2 domains. EVPN also supports rich multihoming capabilities, allowing applications to maintain multiple connections to their virtual network segments—useful for improving resiliency and distributing application load across the network.

All of this is provided while maintaining strict tenant separation. Each application is connected via its own virtual segments and traffic is passed based on separate edge policy. This creates a separation of dependencies as each application is encapsulated into its own VPN, granting fine-grained control.

So not only does EVPN support legacy applications, it does so in an inherently more scalable and operationally friendly way.

Supporting the Future

For modern applications, the future involves microservices leveraging containerized application components orchestrated via application platforms like Kubernetes. Distributing application components creates a dependency on standardized connectivity services that can extend to containers, VMs and bare-metal servers deployed across both private and public clouds.

In these environments, connectivity services extend beyond merely passing packets. If there is a routed instance that acts as the virtual network edge, that virtual device can also serve as a point of policy enforcement. By building policy and control into these virtual network endpoints, connectivity services can be easily extended to include security (i.e., microsegmentation).

Of course, modern applications are not necessarily static. One benefit of microservices is that they can scale up and down dynamically, effectively right-sizing infrastructure consumption based on demand. This places a real-time operational burden on connectivity. Combining EVPN's control plane with an over-the-top orchestration solution, for example, can provide the dynamism required to service a modern application.

Diverse Environments, Common Operations

If the future of data centers requires maintaining one foot in the past and the other stepping toward the future, technology building blocks will be required that support both. The key is doing so without adding to the already burdensome operational load for most enterprises. EVPN's support for virtualized Layer 2 Ethernet and Layer 3 IP services, with first-class connectivity multihoming, robust multicast features and optimized handling of "network noise," makes it a natural choice for forward-looking data center architectures focusing on simplified networking.

Today, we kick off the third and final video in a series of content designed not to sell but rather to illuminate. With our GetSmart series on SDxCentral, we will look at why EVPN matters, providing an industry view of what EVPN can do as the de facto architectural foundation for the modern data center. For additional information please download our E-Book on EVPN-VXLAN.

Evolve Your Campus with EVPN-VXLAN

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 Internet of Things (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 the campus and data centers. The increasing use of IoT devices also means that these networks are expected to scale rapidly and since many of these IoT devices have limited networking capabilities, they mandate unconventional requirements from the network without adding complexity and resources.

Most traditional campus architectures are based on proprietary vendor lock-in technologies designed to address the needs of traditional campuses with static requirements and single vendor deployments. While these architectures work well in small campuses that are static, they are simply too rigid to support the scalability and ever-changing needs of modern large enterprises.

Juniper's EVPN-VXLAN based campus architecture uses a Layer 3 IP-based underlay with EVPN-VXLAN overlay. This architecture addresses the needs of the modern enterprise by decoupling virtual topology from the physical topology. This enables simple IP based layer 3 network underlay that limits the layer 2 broadcast domain and a flexible overlay with efficient layer 3 or layer 2 connectivity depending on business needs. Endpoints that require layer 2 adjacency like some IoT devices can be placed anywhere in the network and remain connected to the same logical L2 network with consistent network access policies.

Juniper's EVPN-VXLAN based campus architecture supports multiple deployment models. Enterprise can deploy IP fabric with EVPN-VXLAN overlay only at the core/distribution layer and the access layer can be a layer 2 Virtual Chassis. Or enterprises can deploy end-to-end IP fabric underlay with EVPN-VXLAN overlay across core, distribution and access layers.

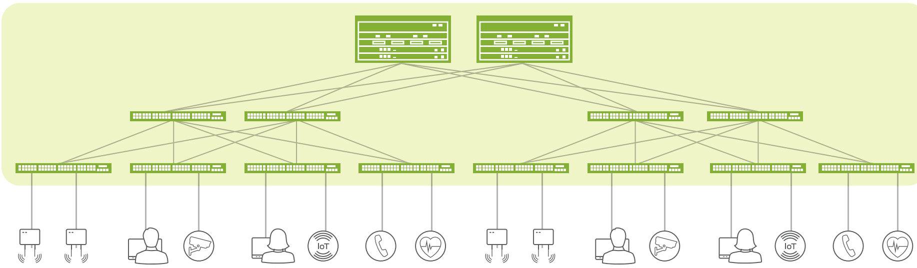

With control plane based L2/L3 learning, this architecture reduces the flood and learn issues associated with data plane learning in traditional architectures. Juniper's EVPN-VXLAN based campus architecture allows enterprises to easily add additional core, distribution, and access layer devices as the business grows without having to redesign the network or perform a forklift upgrade. This is a big shift from the rigid chassis based traditional architectures. This vendor-agnostic solution allows enterprises to utilize their existing access layer infrastructure and gradually migrate to access layer switches from Juniper.

Juniper's EVPN-VXLAN based campus networks provide the following benefits:

- Consistent architecture for any scale
  - Enterprises typically have multiple campuses or sites with different size requirements. These requirements can be met with a common EVPN-VXLAN based campus architecture that is consistent across all sites, irrespective of the size of the site. The architecture also provides the flexibility to scale out or scale in as the requirements of a site change
- Non-proprietary architecture, enabling multi-vendor deployment
  - The architecture is uses standards-based protocols like EVPN and VXLAN enabling Enterprises to deploy campus networks using best-in-class multi-vendor network equipment
- Control plane L2/L3 learning reduces flood and learn
  - Learning MAC addresses in the forwarding plane has an adverse impact on network performance as the number of endpoints grow in large enterprise networks. Juniper's EVPN-VXLAN campus architecture uses an EVPN control plane for exchange and learning of routes, avoiding exchanging newly learned MAC addresses in the forwarding plane.

- Location agnostic connectivity enables L2 mobility and consistent user/endpoint experience
  - Some endpoints require L2 adjacency. For example, legacy building security systems or some IoT devices are L2-based. Juniper's EVPN-VXLAN campus architecture can easily provide layer 2 adjacency inside a campus and across campuses with an L2 VXLAN overlay without any changes to the underlay. With Juniper's standards-based NAC integration, an endpoint can be connected anywhere in the network allowing that endpoint to have the same level of access to the network
- Flexible WAN overlay design because VXLAN is underlay agnostic
  - One of the key benefits of VXLAN is that it is underlay agnostic. You can connect multiple campuses using an L2VPN or L3VPN service from a WAN provider or connect campuses using IPSec over Internet and still provide a VXLAN overlay on top of these WAN connections
- Consistent network segmentation across campuses and datacenters
  - With a consistent EVPN-VXLAN based architecture across campuses and datacenters, Enterprises can now easily provide consistent end-to-end network segmentation for endpoints and applications
- Common design means common skillsets and tools to manage datacenter and campus
  - When enterprises design campuses and datacenters based on a common EVPN-VXLAN design, they can use common tools and common network teams to deploy and manage campus and datacenter networks

Juniper's Evolved Campus solution based on VXLAN overlay with EVPN control plane is an efficient and scalable way to build campuses and interconnect multiple campuses, datacenters and public cloud. With a robust BGP/EVPN implementation on all platforms, 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, data centers and cloud infrastructure.

ADDITIONAL RESOURCES

Architecture Guide: Cloud Data Center
https://www.juniper.net/documentation/en_US/release-independent/solutions/information-products/pathway-pages/sg-005-cloud-data-center.html

Blog: Why Does the Enterprise Data Center Need EVPN-VXLAN
https://forums.juniper.net/t5/Enterprise-Cloud-and/Why-does-the-Enterprise-Data-Center-need-EVPN-and-VXLAN/ba-p/347678

Whitepaper: Multicloud Technical Guide for Network and Cloud Architects
https://www.juniper.net/assets/us/en/local/pdf/whitepapers/2000722-en.pdf

Point of View: The 5-Step Multicloud Migration Framework
https://www.juniper.net/assets/us/en/local/pdf/pov/3200064-en.pdf

Webpage: The Enterprise Multicloud 5-step Journey
https://www.juniper.net/us/en/solutions/data-center/data-center-journey/

Training: Advanced Data Center Switching (ADCX)
https://learningportal.juniper.net/juniper/user_activity_info.aspx?id=9685

Engineering
Simplicity

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: +31.0.207.125.700

Fax: +31.0.207.125.701