Improve Data Center Interconnect, L2 Services with Juniper's EVPN

The Need for Next-Generation L2 VPN Connectivity

Table of Contents

List of Figures

Executive Summary

Ethernet VPN (EVPN) delivers a wide range of benefits—including greater network efficiency, reliability, scalability, VM mobility, and policy control—that directly impact the bottom line of service providers and enterprises alike.

This white paper provides an overview of EVPN, including its features and benefits. Additionally, this paper explores how Juniper Networks' implementation of this technology extends these benefits, helping enterprise and service providers boost network efficiency and ease deployment while maintaining standards-based compliance and interoperability.

Introduction

Strong demand for Ethernet-based connectivity is driving double-digit growth for Layer 2 VPN (L2VPN) services such as E-LINE/virtual private wire service (VPWS) and E-LAN/virtual private LAN service (VPLS[1]). These technologies have proven attractive for a variety of private line and LAN extension applications across both wide area networks and metro area networks (WANs/MANs).

According to research firm IDC, data center interconnectivity, disaster recovery/business continuity, and data storage replication are the top three applications spurring adoption of VPWS and VPLS[2]. Additionally, dedicated Internet access (DIA), IP VPNs, private clouds, and other applications are also contributing to the rising demand for Ethernet-based connectivity.

Compared to legacy private line and older packet services, VPWS and VPLS are more scalable, flexible, and economical. However, carriers and enterprises alike are facing shortcomings in these L2 technologies, especially with the relentless growth of cloud-originated content, video, real-time traffic, and inter-data center traffic. Scale, performance, and traffic management have become more business critical than ever, driving the need for next-generation L2VPN technology. Network operators need a VPN solution that does the following:

- Scales to the largest data center deployments
- Maximizes bandwidth through active load balancing, address learning, and other mechanisms
- Minimizes latency for optimal user experience
- Speeds service recovery and restoration
- Reduces configuration and operations overhead
- Allows efficient migration of virtual machines (VMs) across data centers
- Provides an integrated L2 and L3 VPN solution that efficiently routes traffic in a mixed L2/L3 environment

Juniper Networks is leading an industry-wide, multivendor initiative to address the shortcomings of current L2VPN solutions. Along with Cisco, Alcatel-Lucent, AT&T, Verizon, Bloomberg, and other participants, Juniper is helping define Ethernet VPN (EVPN), a new standards-based protocol for interconnecting L2/L3 domains over an MPLS-based infrastructure.

EVPN overcomes the shortcoming of current E-LINE and E-LAN offerings by providing integrated L2/L3 connectivity, native support for multihoming, Media Access Control (MAC) address mobility, and network resiliency between edge nodes. It builds on widely deployed VPLS and IP VPN technologies, protecting investments in MPLS infrastructure and knowledge base.

Additionally, Juniper has enhanced its EVPN implementation to further improve scalability, multihoming, and configuration flexibility while preserving standards-based interoperability. With Juniper's EVPN, customers get a combined L2/L3 VPN solution that's more scalable, resilient, and efficient than current technologies.

What Is EVPN and Why Is It a Better Solution?

E-LINE and E-LAN services provide an Ethernet virtual connection between two sites. E-LINE is a point-to-point (P2P) service, whereas E-LAN/VPLS provides a point-to-multipoint (P2MP) service. Defined by the Metro Ethernet Forum, both support class-of-service (CoS) parameters, service multiplexing with virtual LAN (VLAN) tags, services bundling, and security. Both services are typically implemented on an MPLS-based infrastructure.

While these technologies have been a boon to carriers and enterprises alike, each has shortcomings that impede business operations and burden IT staff with configuration and operations tasks. Consequently, Juniper and other IETF members are building on VPLS and IP VPN technologies to define EVPN, the next-generation, standards-based VPN technology for interconnecting L2 domains that also offers integrated L3 gateway functionality.

Similar to IP VPN and VPLS, EVPN provides network virtualization to support multiple customers over a common MPLS infrastructure. Provider edge (PE) devices implement multiple EVPN instances to provide virtual L2 bridged connectivity between customer edge (CE) devices, which can be a host, a router, or a switch.

[1] Source: IDC - http://www.prnewswire.com/news-releases/us-carrier-ethernet-services-2013-2017-forecast-238429701.html
[2] Source: IDC - http://www.idc.com/getdoc.jsp?containerId=prUS24365613

The EVPN specification supports several ways for PE devices to connect, but most implementations use MPLS infrastructure as depicted in Figure 1. By leveraging MPLS, EVPN inherits benefits such as MPLS fast reroute (FRR) and MPLS CoS. Additionally, EVPN PE devices can be connected over an IP infrastructure using MPLS/GRE or VXLAN tunneling. VXLAN, which emerged as a data-plane encapsulation method for data centers with an IP fabric, can also be used in a WAN to carry L2 frames over a pure IP infrastructure. In this case, the native VXLAN encapsulation is used to send packets between PE devices. As shown in Figure 2, EVPN can be used as a control plane for a VXLAN network to communicate virtual network identifiers (VNIs) between endpoints. Additionally, as shown in Figure 3, the ability to offer L2 stretch functionality over IP tunnels signaled by EVPN is also useful for offering L2 reachability for tenants in an IP fabric data center. (Use of EVPN signaling with VXLAN tunnels as a data plane for intra-DC fabric technology is covered in a separate white paper.)

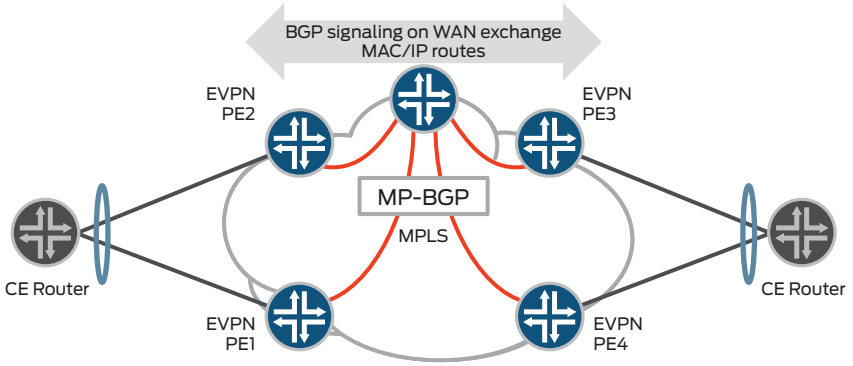

Figure 1: EVPN for next-generation Ethernet services on MPLS

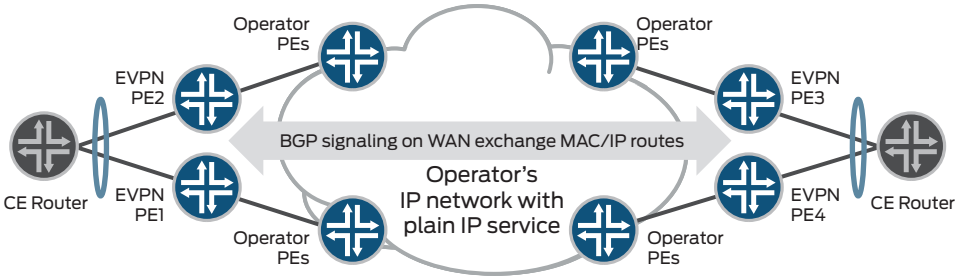

Figure 2: EVPN with VXLAN for L2 VPN over IP WAN

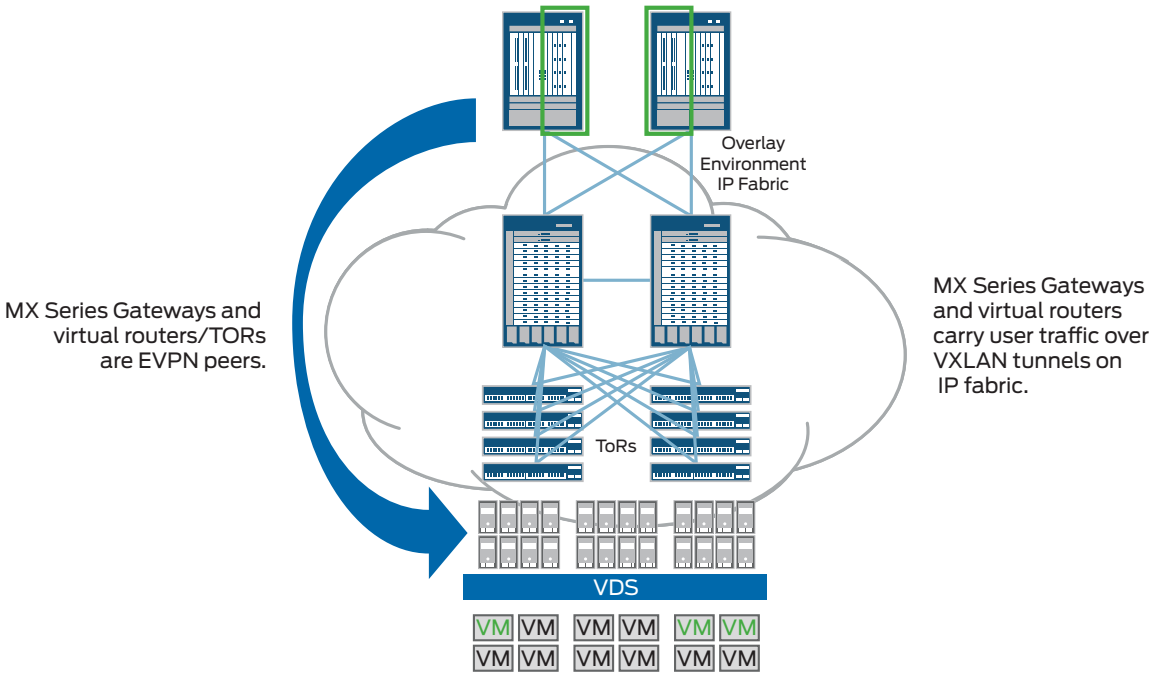

Figure 3: EVPN with VXLAN for data centers with IP overlay networks

EVPN is similar to VPLS in the data plane—for example, both implement MAC learning to establish reachability between various L2 devices. A key difference is that EVPN uses the control plane rather than the data plane to communicate MAC and IP address reachability between PE devices. EVPN uses MP-BGP to distribute MAC routes and their IP bindings. This allows operators to implement very fine-grained control over MAC route distribution and thereby offer more sophisticated L2VPN services than has been possible to date.

More importantly, EVPN addresses the shortcomings of VPLS and other L2 services. These shortcomings, and how EVPN resolves them, are highlighted in the following sections.

Improved Network Efficiency

VPLS and other L2VPN services rely on bridging for their operation and learn MAC addresses in the data plane at every device in a packet's path. VPLS emulates the broadcast domain of a LAN over an MPLS infrastructure. Consequently, unknown unicast and broadcast MAC addresses are always flooded, and all participating ingress PE routers make separate copies of each broadcast or multicast packet to send to all other PE routers that are part of the same extended VPLS-based LAN.

In a large L2VPN, for example, replication overhead can be significant for each ingress router and its attached core-facing links. Especially in data center interconnect applications, as data center traffic grows it's critical that BUM traffic be processed more efficiently. Service providers and enterprises need interconnect technologies that minimize the flooding of multi-destination frames.

EVPN is similar to IP VPNs and uses Multiprotocol BGP (MP-BGP) in the control plane to advertise MAC and IP reachability information. For example, PE devices use MP-BGP to advertise the MAC addresses learned from their connected CE devices, along with an MPLS label, to other PE devices. Essentially, devices aren't required to learn MAC addresses because the EVPN control plane supplies that information. As a result, EVPN greatly reduces flooding the network with unknown unicast traffic, provides greater control over the MAC learning process, and gives operators the ability to apply policies—such as restricting who learns what. Additionally, since MAC routes include MAC addresses and their IP bindings, ARP flooding is minimized PE devices can locally support proxy ARP for remote hosts.

Additionally, with MAC routes-based, IP VPN-like forwarding, the PE devices do not need to maintain point-to-point pseudowires between all PE devices in the network core. Moreover, since EVPN forwards traffic based on per-MAC rules, along with aliasing functionality, it achieves good load balancing in the network core. These properties make EVPN an inherently more scalable solution than VPLS.

Integrated L2/L3 Functionality

Because VPWS and VPLS operate at L2, they require L3 gateway functionality to allow inter-subnet (inter-VLAN) traffic. Even when the traffic is local—for example, when both the subnets (VLANs) are on the same server—traffic must be routed via the L3 gateway. While vendors typically provide L3 gateway solutions, they typically require special configuration—which is an operational burden—and traffic flows aren't always optimal. On the other hand, a pure L3 solution can create issues for intra-subnet traffic, such as duplicate MAC addresses not being detected.

EVPN's integrated routing and bridging (IRB) functionality supports both L2 and L3 connectivity between edge nodes along with built-in L3 gateway functionality. By adding both host and gateway MAC and IP address information in MAC routes, EVPN provides optimum forwarding for both intra-subnets and inter-subnets within and across data centers for unicast as well as multicast traffic. This functionality is especially useful for service providers who offer L2VPN, L3VPN, or direct Internet access services and want to extend all these services to provide cloud services to existing customers, for example.

Multihoming

Data volumes continue to escalate, which makes running the network in active/standby mode increasingly more expensive. It is critical for IT to maximize utilization of all links between data centers. That includes the ability to establish multihomed connections, even between multiple PE devices, and to load-balance across those connections. In addition to better link utilization, multihomed connections also offer greater resiliency and reliability against the potential failure of one connection or node.

A common mechanism for multihoming a CE node to a set of PE nodes is to use multichassis Ethernet link aggregation groups (LAGs) based on IEEE 802.1ax /802.3ad. However, 802.1ax/802.3ad does not define a standard load-balancing algorithm, leading to vendor-specific implementations that behave differently, posing a challenge for IT. Additionally, Ethernet resiliency mechanisms such as Multiple Spanning Tree Protocol (MSTP) and Ethernet ring protection switching (ERPS) pose their own limitations, further complicating multihoming support.

What service providers and enterprises need to maximize bandwidth and ensure reliable, high-performance data center interconnection is multihoming with all-active redundancy and load balancing across all links. EVPN supports both single-active and all-active multihoming for L2 and L3 traffic along with load balancing. With support for both all-active per-service and all-active per-flow multihoming, EVPN enables optimal load balancing across peering PE devices.

High Availability

High availability (HA) is crucial for any network service and is especially critical for data center interconnection where traffic volumes are very large. For availability, VPLS relies on the underlying MPLS capabilities such as Fast Reroute. While these MPLS fast reroute mechanisms aid in boosting network availability, L2 shortcomings pose their own challenges to HA. For example, the lack of all-active multihoming in VPLS makes it difficult to achieve sub-50 ms service restoration in case of an edge node or edge link failure.

The rapid increase in virtualized applications has led to an increase in the volume of MAC addresses that must be handled by the network. For large service providers and enterprises, there can be hundreds of thousands of MAC addresses supported across interconnected sites. If a node or link fails, the need to re-learn a high number of MAC addresses in the broadcast domain can slow network reconvergence, which leads to data loss and negatively affects application performance.

To ensure rapid recovery after a failure, enterprises and service providers need all-active multihoming. Network reconvergence must also be independent of the number of MAC addresses learned by the PE device. EVPN support for all-active multihoming is a key HA enhancement.

In addition, EVPN defines a mechanism to efficiently and quickly signal remote PE devices with the need to update their forwarding tables when a connectivity failure occurs. The withdrawal of the ES route feature allows for the mass withdrawal of MAC addresses whose reachability has been affected due to loss of that particular link.

Data separation and the ability to scale are also key to HA. EVPN provides the ability to preserve the "virtualization" or isolation of groups of hosts, servers, and VMs from each other, which enables network operators to reliably support a significantly higher volume of MAC addresses and VLANs. EVPN also gives network operators greater control over routes and the ability to program remote MAC addresses in the control plane using BGP as the signaling protocol.

In addition to standard EVPN HA techniques, Juniper's EVPN implementation provides gateway redundancy by allowing the same gateway IP and MAC address to be configured on all EVPN PE devices. Such a configuration allows Virtual Router Redundancy Protocol (VRRP)-like default gateway redundancy to protect against PE device failures.

VM Mobility

VMs and workloads are frequently moved to different servers in the same data center, or in remote data center locations, to optimize application workloads or protect against failures. When these migrations occur, it is imperative that VMs maintain the same MAC address, IP address and VLAN ID so that user sessions are not dropped, ensuring a seamless application experience.

In older technologies like VPLS, the only trigger for detecting VM movement is using MAC move detection, which presents a challenge. EVPN makes VM mobility much simpler by allowing PE devices to explicitly signal MAC moves using special MAC mobility community strings in MAC routes. Additionally, EVPN gives operators the mechanisms to control excessive MAC moves, which are typically associated with rogue MAC addresses and not associated with legitimate VM movements. Operators can specify their tolerance level, such as a maximum number of MAC moves allowed in a specific time interval. If that threshold is exceeded, EVPN withdraws that MAC route completely to protect the network.

Using the optimal default gateway is another challenge with VM mobility. Because a VM is unaware that it has moved, it does not flush its ARP table and it continues to send inter-VLAN packets to its configured default gateway. However, after the move this default gateway might not be the optimal, or even local, default gateway. Juniper has implemented a new technique called Virtual Machine Traffic Optimization (VMTO) to aid efficient VM mobility. As shown in Figure 4, VMTO enables a VM to migrate from one server to another within the same or between different data centers without traffic "tromboning." Network operators benefit from increased flexibility in VM location while maintaining efficient use of links, and users benefit from uninterrupted access to applications and services.

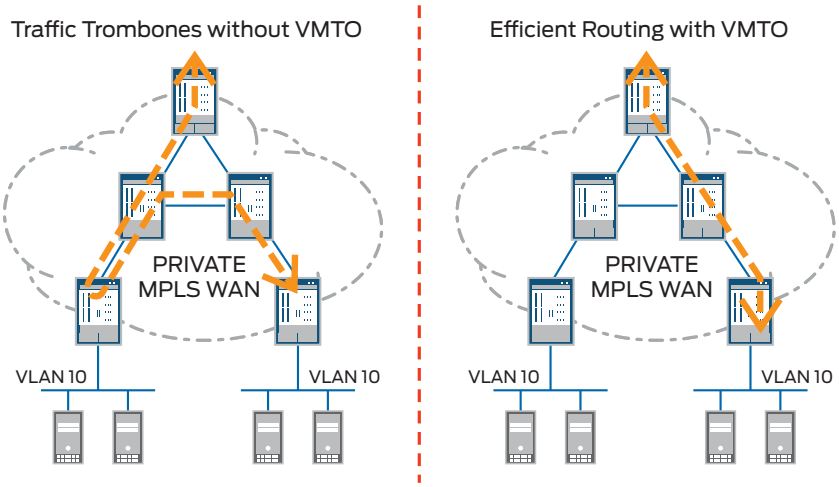

Figure 4: Optimizing inter-VLAN traffic flows

VMTO comes into play when a VM in a data center seeks to send packets to a server in another data center, but that server's active default gateway for its registered VLAN is in yet another data center due to the VM having moved. That traffic must travel from the first data center to the second data center to reach the VLAN's active default gateway in order to be routed to the destination data center. This results in duplicate traffic on WAN links and suboptimal routing. As shown in Figure 5, VMTO fixes this problem by virtualizing and distributing the default gateway so that it is active on every router that participates in the VLAN. The effect is that egress packets can be sent to any router on the VLAN, allowing the routing to be done in the local data center.

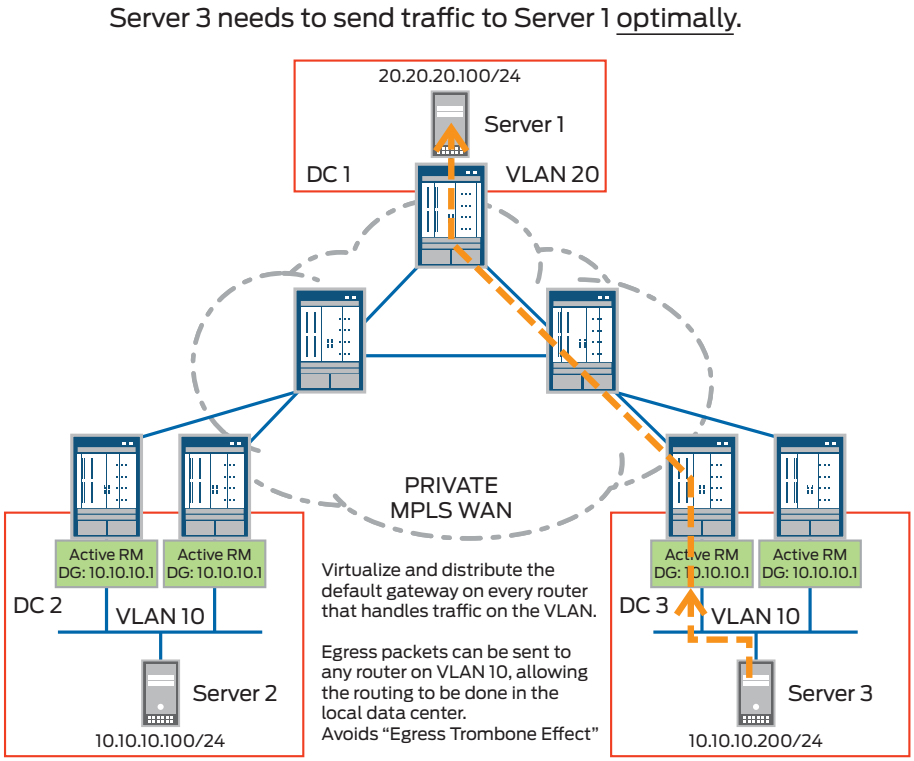

Figure 5: Optimizing egress traffic with VMTO

VMTO also addresses a similar problem with ingress routing. As shown in Figure 6, consider the situation when, after a VM moves, a server in one data center seeks to send packets to a server in another data center. Due to the VM move, the edge router has no knowledge of the host IP's new location. Consequently, it routes the traffic across the WAN to the original data center due to a lower-cost route for the subnet. Then the edge router at that original data center sends the packet to the destination data center.

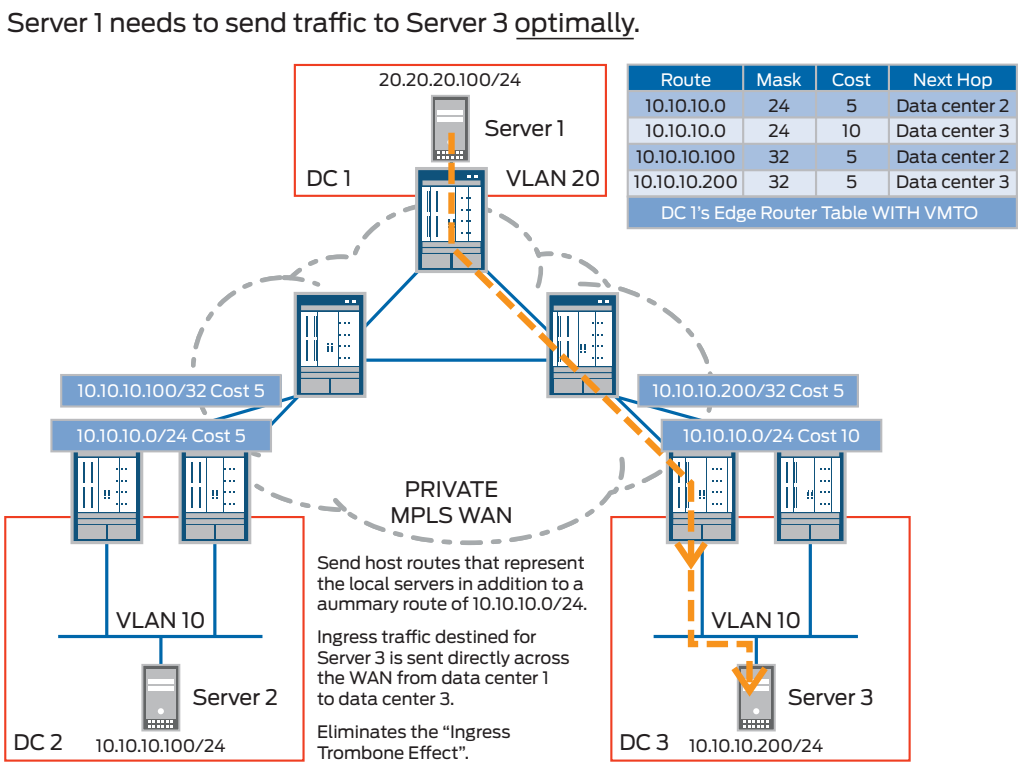

Figure 6: Optimizing ingress traffic with VMTO

VMTO addresses this inefficiency. In addition to sending a summary route of the subnet in question to the data center edge routers, VMTO also sends host routes that represent the location of local servers. As a result, ingress traffic destined for the servers is sent directly across the WAN from the first data center to the destination data center.

Provisioning

Although VPLS supports BGP-based autodiscovery, provisioning VPLS still requires that network operators specify various network parameters on top of the access-side Ethernet configuration. This provisioning overhead adds to operational expenses. In addition, VPLS has no provision for administrative control on MAC routes, which limits an operator's ability to maximize route efficiency.

EVPN's support for BGP-based policies gives operators better administrative control, including fine-grained, policy-driven control on route advertisement and consistent policy-based forwarding for L2 traffic. EVPN's IP VPN-like policy control also allows for more efficient, feature-rich E-LAN and E-LINE services and enhanced, customized services for end users. Additionally, by providing a single technology for both L2 and L3 VPNs, EVPN simplifies deployment and can seamlessly interoperate between L2 VPNs and IP VPNs.

All of the previously mentioned characteristics ideally suit EVPN technology for data center interconnection, L2 VPNs, and integrated L2/L3 VPN applications.

EVPN Technical Overview

To support its range of capabilities, EVPN introduces several new concepts such as new route types and BGP communities. It also defines a new BGP network layer reachability information (NLRI), called the EVPN NLRI. The following sections highlight key EVPN protocol building blocks.

Ethernet Segment Identifier

With EVPN, an Ethernet "segment" consists of the set of Ethernet links used to multihome a CE device to two or more PE devices. Each Ethernet segment has an identifier called the "Ethernet Segment Identifier" (ESI), which has a 10-octet value.

An Ethernet segment must have a non-reserved ESI that is unique network wide (that is, across all EVPN instances on all the PE devices). The EVPN protocol supports autodiscovery of Ethernet segments for ease of use and the election of a designated forwarder (DF) for better control over BUM traffic forwarding.

While EVPN PE devices have special semantics and processing for Ethernet segments, from the CE device side, a multihomed Ethernet segment appears as a regular LAG interface.

Ethernet TAG Identifier

An Ethernet tag ID is a 32-bit field containing either a 12-bit or a 24-bit string that identifies a particular broadcast domain (for example, a VLAN) in an EVPN instance. The 12-bit identifier is used for a VLAN ID (VID). A 24-bit identifier is used for the VNID for EVPN with VXLAN data plane and for the I-SID for PBB-EVPN.

An EVPN instance consists of one or more broadcast domains (one or more VLANs). VLANs are assigned to a given EVPN instance by the provider of the EVPN service. A CE device acknowledges the multiple PE devices to which it's connected as one switch. This allows the CE device to aggregate links that are attached to different PE devices in the same LAG bundle.

A given VLAN can itself be represented by multiple VIDs. In such cases, the PE devices participating in that VLAN for a given EVPN instance are responsible for performing VLAN ID translation to/from locally attached CE devices. However, if a VLAN is represented by a single VID across all PE devices participating in that VLAN for that EVPN instance, then there is no need for VID translation at the PE devices.

Each EVPN instance requires a Route Distinguisher (RD), which is unique per PE device, as well as one or more globally unique Route-Targets (RTs). In deployments where VIDs are unique across all EVPN instances, the RT(s) for each EVPN instance can be derived automatically from the corresponding VID.

In addition to the aforementioned constructs, the EVPN state engine operates by exchanging the following routes between EVPN peers.

EVPN Route Type 1: Ethernet Autodiscovery Route

Ethernet autodiscovery route type can be used on a per ESI and per EVI basis, each of which support different capabilities. Both Ethernet AD route per EVI and Ethernet AD route per ESI are required for multihoming.

Ethernet AD route per ESI: This route type is essential for multihoming. It uses ESI information to learn about multiple paths to the same CE and to signal all-active or single-active mode of operation for a multihomed CE device. In addition, Ethernet AD routes per ESI are used to advertise split-horizon labels for loop avoidance of L2 BUM traffic—withdrawal of this message helps fast convergence (mass withdrawal of MAC addresses).

Since EVPN learns MAC address reachability via the BGP control plane over the MPLS network, network convergence time depends on the number of MAC advertisement routes that must be withdrawn by the PE device experiencing a failure. In very large environments, convergence can be slow. EVPN addresses this problem using a mechanism that efficiently and quickly signals remote PE nodes to update their forwarding tables in the event of a failure to connect to an Ethernet segment. That is, the PE device withdraws the corresponding set of Ethernet AD per ES routes, which in turn triggers all PE devices that receive the withdrawal message to update their next-hop adjacencies for all MAC addresses associated with the Ethernet segment in question, resulting in rapid convergence.

Juniper supports the use of EVPN mass withdrawal for single-homed sites as well as multihomed sites. This allows an ingress PE device to discard a packet with an unknown destination MAC address rather than flooding it. In addition, when multihoming is used in all-active mode, Juniper supports local repair in the event of a failure. Specifically, if a PE device loses connectivity to the CE device on a given ES, that PE device forwards unicast traffic to other PE devices in the same ES rather than sending it to a black hole. Juniper has implemented this local repair capability without any protocol extensions to ensure interoperability.

Ethernet AD route per EVI: This route type enables load balancing for L2 and L3 traffic and supports the concept of "aliasing," which enables a remote PE device to load-balance known unicast traffic toward all PE devices multihomed to the same ESI. Aliasing lets a PE device signal that it has reachability to an EVPN instance on a given ES even when it has not learned MAC addresses from that EVI/ES.

For example, if a CE device is multihomed to multiple PE nodes using a LAG with all-active redundancy, it's possible that only a single PE device learns the set of MAC addresses associated with the traffic transmitted by that CE device. Consequently, remote PE nodes receive MAC advertisement routes for these addresses from a single PE device, even though multiple PE devices are connected to the multi-homed segment. This route provides load balancing to all active CE devices, even when the MAC address is learned by only one PE device. Ethernet AD per EVI also supports a related backup path function, which is used in single-active redundancy mode.

Juniper supports DF election per EVI. For single-active multihoming, Juniper precomputes which PE device becomes the DF for a given [ES, EVI] if the current DF fails, and that PE device advertises a per-EVI Ethernet AD route for that [ES, EVI]. This enhancement allows an ingress PE device to immediately send traffic to the new DF as soon as it learns that the current DF has failed.

EVPN Route Type 2: MAC/IP Advertisement Route

EVPN extends BGP to advertise MAC and IP addresses using the MAC/IP advertisement route type in the EVPN NLRI. Key uses of this route type include advertising host MAC and IP reachability with a "service label," allowing control plane-based MAC learning for remote PE devices, minimizing flooding across a WAN, and allowing PE devices to perform proxy-ARP locally for remote hosts.

By using the same label for both MAC advertisement and per-EVI Ethernet AD routes, Juniper improves MPLS label usage and therefore scalability.

EVPN Route Type 3: Inclusive Multicast Route

This route type sets up paths for BUM traffic on a per-VLAN, per-EVI basis. Specifically, it allows a PE device to send BUM traffic from a CE device on a VLAN in a given EVI to all the other PE devices that span that VLAN in that EVPN instance. This route type uses existing multicast VPN-defined constructs for signaling and transport, supports ingress replication as well as P2MP LSPs, and enables a PE device to carry the traffic of more than one EVPN instance on the same tree using aggregation.

EVPN Route Type 4: Ethernet Segment (ES) Route

This route type allows PE devices with the same ESI to discover each other. That is, PE devices connected to the same multihomed CE device—with the same Ethernet segment identifier value—discover each other by exchanging Ethernet Segment route messages. ES route messages are typically used with the ES-import extended community so that only PE devices with a given ES import that route, thereby protecting PE devices from processing ES routes that the PE devices do not need to process. This ES route type is used for DF election, which designates a router to carry BUM traffic from the core to the CE device. The non-DF router does not forward BUM traffic on that Ethernet segment, thereby saving the CE device from receiving multiple copies of the same traffic.

Extended Community Tags

In addition to new route types, EVPN also defines several new BGP communities. These attribute tags include:

- **ESI Label**—This is a new transitive extended community that can be advertised along with Ethernet autodiscovery routes. It enables split-horizon procedures for multihomed sites. An ESI Label is used in split-horizon filtering by PE devices connected to the same multihomed Ethernet segment.
- **MAC Mobility Extended Community**—This new transitive extended community is used to ensure that PE devices retain the correct MAC advertisement route when multiple updates occur for the same MAC address. This community supports VM mobility for workload migration between data centers, for example, by ensuring that PE devices learn the new MAC route and withdraw the old one. In addition, this community protects against rogue MAC addresses. For example, when a PE device learns a new MAC route with the MAC mobility community, it starts a timer (this is configurable but has a default of 180 seconds). If there are more MAC moves in the configured time interval than are specified, the PE device blocks that MAC address.
- **Default Gateway Extended Community**—This extended community is carried by a MAC/IP advertisement route to indicate that the route is associated with a default gateway, which is useful for IP routing for inter-VLAN traffic. For example, EVPN PE devices use this community to advertise all local IRB MAC and IP addresses in MAC routes with the default gateway community string. It can also be used to optimize L3 egress traffic forwarding.

EVPN Service Offerings – Customer Handoff Options

From a deployment perspective, EVPN offers service providers three customer connectivity options for delivering rich service offerings to their customers. Moreover, these service offerings are consistent with Metro Ethernet Forum-defined services and offer easy migration of these services onto new EVPN infrastructure for even richer service offerings.

Option 1. VLAN-Based Service Interface—This service interface supports a single broadcast domain or VLAN/Q-in-Q per EVPN instance. Providers use this service interface to offer an E-LAN or E-LINE service for a single broadcast domain, with customer VLANs/Q-in-Q having local significance. If a customer VLAN is represented by different VIDs on different CE-PE links, then each egress PE device seeks to perform VID translation for frames destined to its attached CE devices.

Option 2. VLAN Bundle Service Interface—This service interface supports a bundle of VLANs over one EVPN instance, but only a single bridge domain for each EVI. That is, each customer's traffic is represented by a single EVI, which corresponds to a virtual switch that can support 4000 VLANs per tenant as well as provides L2 and L3 connectivity on a single customer interface and all-active or single-active connectivity. Because it supports an N:1 mapping between VLAN ID and EVI, this service interface requires that MAC addresses be unique across VLANs and disallows VLAN translation.

This service type also supports a special case known as a port-based service interface, where all of the VLANs on a port are part of the same service and map to the same bundle.

Option 3. VLAN-Aware Bundle Service Interface—This is the most common deployment option for DCI, providing customers with an E-LAN or E-LINE service for multiple broadcast domains. With this service interface, an EVPN instance consists of multiple broadcast domains or VLANs, with each VLAN having its own bridge domain. Each customer's traffic is represented by a single EVI, which corresponds to a virtual switch that can support 4,000 VLANs per tenant as well as provide L2 and L3 connectivity on a single customer interface and all-active or single-active connectivity. Like the VLAN bundle service interface, this interface supports N:1 mapping between VLAN ID and EVI. However, since bridge domains are separate, it allows for local VLAN/Q-in-Q translation on CE-PE links. It also supports a port-based option.

EVPN Deployment Use Cases: Meeting the Need for Next-Generation Connectivity

EVPN delivers a wide range of benefits to service providers and enterprises alike, including greater network utilization, reliability, scalability, VM mobility, and policy control. The following section highlights some deployment use cases where EVPN offers significant business value.

Data Center Interconnect Use Case

Juniper Networks® MX Series 3D Universal Edge Routers have been de-facto data center/WAN edge gateways for many cloud and telecom providers. As a data center gateway, MX Series routers connect the network—either L2 or L3—on one side and WAN networks with IP/MPLS, L3VPN, or L2VPN on the other.

In many deployments, data centers running on different intra-data center technologies—for example, one on pure L2 technology and another on some overlay technology like VXLAN—need to be interconnected in a seamless manner.

As seen in Figure 7, EVPN technology in general, and a feature-rich MX Series implementation in particular, support all combinations of edge technologies for the data center to the WAN gateway.

Table 1: EVPN for DCI Use Cases

EVPN Properties	Benefit for DCI Use Case
Support for all active multihoming	Data center operators serve hundreds of gigabytes of content from their facilities. All-active support allows operators to efficiently use all links at the same time.
Support for VM mobility	Data center operators can allow workload migration from one data center to another without compromising the end-user experience.
Integrated L2/L3 services	Data center operators can support all types of customer connectivity demands, whether at Layer 2 or at an IP or IP VPN level. Additionally, different VLAN handoff options allow rich customizable services.
High availability	Data center operators serve hundreds of gigabytes of content from their facilities. Built-in high availability protects them from traffic loss caused by link or node failure.

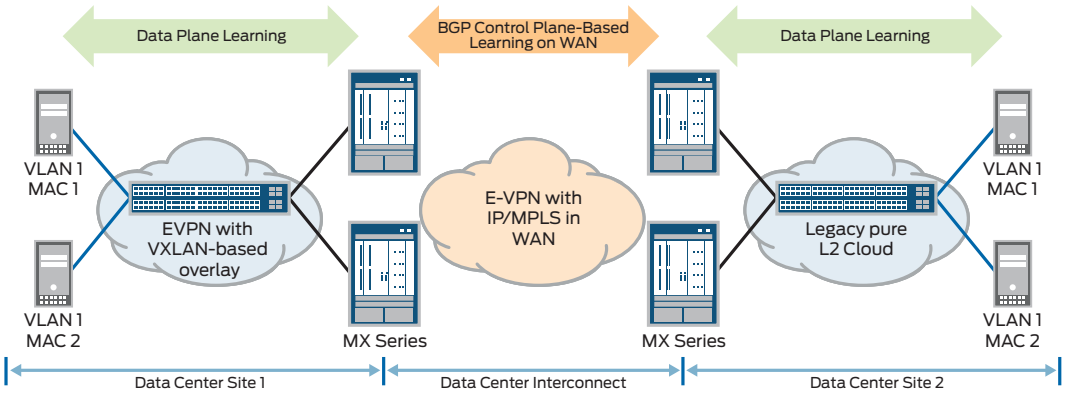

Figure 7: EVPN for data center interconnect

Next-Generation E-LINE/E-LAN/E-TREE Use Case

EVPN overcomes the shortcomings of current E-LINE, E-LAN, and E-TREE offerings by providing integrated L2/L3 connectivity, native support for multihoming, and network resiliency between edge nodes. EVPN builds on widely deployed VPLS and IP VPN technologies, protecting investments in MPLS infrastructure, as well as the existing knowledge base.

Table 1: EVPN for Next-Generation L2VPN Use Cases

EVPN Properties	Benefit for Next-Generation L2VPN Use Case
Support for all active multihoming	VPN operators are rolling out 100 Gbps access circuits to their customers. All-active support allows operators to efficiently use all links at the same time.
Rich policy-based services	EVPN's support for BGP-based policies gives operators better administrative control, including granular, policy-driven control on route advertisement and consistent policy-based forwarding for L2 traffic.
Integrated L2/L3 services	Efficient integration of L2 services with IP VPN and Internet services allows more connectivity options for applications. Different VLAN handoff options allow rich customizable services.
Common provisioning for all services	By providing a single, consistent technology for both L2 and L3 VPNs, EVPN simplifies deployment and can seamlessly interoperate between L2 VPNs and IP VPNs.
High availability	Given enterprises' reliance on cloud-based services, high availability at the network level has become table stakes for next-generation E-LINE, E-LAN, and E-TREE services. Built-in high availability protects them from traffic loss caused by link or node failure.

Conclusion: EVPN is Built for Business – From Technology Efficiencies to Business Effectiveness

For service providers and enterprises alike, EVPN's features and benefits translate directly to the bottom line:

- **Reduced CapEx**—EVPN fully utilizes each CE-PE link to offer higher bandwidth as well as improved resiliency, thus reducing the need for standby connections. For example, all-active multihoming and load balancing maximize bandwidth while features such as IRB and MAC learning help boost efficient utilization of network resources. By building on existing technologies, EVPN lets operators upgrade their current MPLS and VPLS infrastructure to the newer EVPN-based service infrastructure with a simple software upgrade—completely protecting customers' capital investments and thereby reducing CapEx.
- **Lower OpEx**—EVPN reduces operations overhead by making it possible to use a single VPN technology to support both L3 and L2 VPNs, simplifying deployment. In addition, EVPN's IRB functionality reduces configuration overhead and simplifies data center operations by efficiently handling L3VPN interoperability and supporting VM mobility at both L2 and L3. Likewise, EVPN gives operators fine-grained, policy-driven administrative control for greater efficiency. Additionally, by building on existing IP VPN and VPLS technologies, EVPN allows operators to utilize their knowledge base without having to retrain people on completely new technology. All these improvements and reuse options help reduce operators' operational expenditures.
- **Revenue**-enhancing services and greater service flexibility—EVPN enables service providers to offer feature-rich E-LAN and E-LINE services and easily expand their service offerings. For example, providers who already offer L2VPN/VPLS and L3VPN services over an IP/MPLS network can easily use EVPN's IRB feature to provide cloud, storage, and other services. Customers benefit from the ability to get advanced services—including integrated L2/L3 services, sophisticated service topologies via BGP policies, multihoming, and load balancing—at prices more comparable to VPLS than IP VPNs.
- **Improved customer satisfaction**—EVPN's HA features, such as link-level and node-level redundancy, ensure fast failover and convergence times so customers get uninterrupted access to applications and services. Likewise, EVPN lets operators ensure customer privacy. With EVPN, network operators can carefully control how network information is distributed and processed and isolate groups of devices, ensuring that the traffic sharing their network remains private.

By addressing the shortcomings of current L2VPN offerings, EVPN provides an optimal solution for data center interconnection, private cloud, E-LINE/E-LAN, and other uses cases. In addition, Juniper Networks' implementation of this technology extends these benefits to further boost network efficiency and ease deployment while maintaining standards-based compliance and interoperability.

Juniper supports several models of EVPN configuration to meet the needs of data center, cloud, and E-LINE/E-LAN customers. With Juniper's EVPN solution, customers benefit from an easily deployable next-generation VPN solution that allows network and data center operators to effectively meet their business demands.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

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

2000596-001-EN Jan 2015