WHITE PAPER

CROSS-DOMAIN VPLS DEPLOYMENT STRATEGIES

Scaling and Extending VPLS with LDP-BGP VPLS Interworking

Table of Contents

Table of Figures

Executive Summary

There are several methods for using BGP virtual private LAN service (VPLS) to scale VPLS in an existing LDP-VPLS network. These methods do not require any changes on LDP-VPLS provider edge (PE) routers, but rather enable VPLS to scale by using BGP VPLS and using LDP-BGP VPLS interworking. BGP VPLS can also be used to extend the reach of VPLS from a single LDP-VPLS metro domain to the intermetro WAN. This scheme allows service providers to offer regional or national VPLS in an economical, efficient, and scalable manner.

Introduction

VPLS is a key technology in the delivery of multipoint Ethernet service. Service providers are using VPLS to offer transparent LAN service to enterprise customers. With the emergence of metro Ethernet networks, VPLS is also being used as an infrastructure technology. Service providers have shown significant interest in this technology, as measured by real VPLS deployments, including interprovider offerings. In general, the MPLS VPN market has seen explosive growth, and is forecast to continue growing at a compound annual growth rate (CAGR) of 12.8 percent to US$6.4 billion by 2011.[1] This growth will be further fueled by the efforts of major multiple service operators (MSOs), such as Comcast, and large service providers, such as Verizon (with its fiber optic service, FiOS), to target small- and medium-sized businesses (SMBs). Most major service providers already employ an IP/MPLS backbone and offer Layer 3 MPLS VPN services, and they are beginning to offer Layer 2 VPN services to both SMBs and large enterprises. VPLS appeals to enterprises in particular because it allows them to extend their reach beyond their LANs with the same Layer 2 Ethernet connectivity paradigm.

This increasing demand for VPLS requires a highly scalable VPLS network that supports many VPLS customers having multiple sites spread across geographically dispersed regions. A critical factor in growing a VPLS network is how well the underlying VPLS control plane scales. There are two standards for VPLS control planes: one uses BGP for autodiscovery and signaling of pseudowires (RFC 4761), and the other uses LDP for signaling of pseudowires (RFC 4762). The scaling characteristics of these two types of control plane differ vastly. BGP-VPLS signaling offers scaling advantages over LDP-VPLS signaling, and the BGP-VPLS autodiscovery feature offers superior operational scaling as well. These benefits have been confirmed in production deployments of BGP VPLS.

VPLS Overview

VPLS is a Layer 2 multipoint VPN that emulates LAN service across a WAN. VPLS enables service providers to interconnect several customer sites (each being a LAN segment) over a packet-switched network, effectively making all the customer LAN segments behave as one single LAN. A service provider's network appears as an Ethernet bridge to the service provider's customers using VPLS. With VPLS, no routing interaction occurs between the customer and service providers, and the customer can run any type of Layer 3 protocols between sites.

The Internet Engineering Task Force (IETF) Layer 2 VPN Working Group has produced two separate VPLS standards: RFC 4761 and RFC 4762. These two RFCs define almost identical approaches with respect to the VPLS forwarding plane, but very different approaches to the VPLS control plane.

VPLS Control Plane Choices

The VPLS control plane has two primary functions: autodiscovery and signaling.

- Discovery refers to the process of finding all the PE routers that participate in a given VPLS instance. A PE router can be configured with the identities of all the other PE routers in a given VPLS instance, or the PE router can use a protocol to automatically discover the other PE routers. The latter method is called *autodiscovery*.
- After discovery occurs, each pair of PE routers in a VPLS must be able to establish and tear down pseudowires to each other. This process is known as *signaling*. Signaling is also used to transmit certain characteristics of the pseudowire that a PE router sets up for a given VPLS.

[1]Source: IDC.

BGP-VPLS Control Plane

The BGP-VPLS control plane, as defined by RFC 4761, is similar to that for Layer 2 and Layer 3 VPNs. It defines a means for a PE router to know which remote PE routers are members of a given VPLS (autodiscovery), and for a PE router to know the pseudowire label expected by a given remote PE router for a given VPLS (signaling). The BGP network layer reachability information (NLRI) contains enough information to provide the autodiscovery and signaling functions simultaneously.

As in the BGP scheme for Layer 2 and Layer 3 VPNs, on each PE router a route target is configured for each VPLS. The route target is the same for a particular VPLS across all PE routers and is used to identify the VPLS to which an incoming BGP message pertains.

LDP-VPLS Control Plane

The LDP signaling scheme for VPLS is similar to the LDP scheme for point-to-point Layer 2 connections. LDP is used for signaling the pseudowires that are used to interconnect the VPLS instances of a given customer on the PE routers. In the absence of an autodiscovery mechanism, the identities of all the remote PE routers that are part of a VPLS instance must be configured manually on each PE router.

LDP VPLS defines the hierarchical VPLS (H-VPLS) scheme in which, instead of a PE router being fully meshed with LDP sessions, a two-level hierarchy is created involving hub PE routers and spoke PE routers. The hub PE routers are fully meshed with LDP sessions, whereas the spoke PE router has a pseudowire only to a single hub PE router or to multiple hub PE routers for redundancy. Spoke pseudowires can be implemented using any Layer 2 tunneling technology.

VPLS Deployment Trends and Issues

Both LDP-VPLS and BGP-VPLS signaling protocols are widely deployed in current service provider networks. Some service providers initially adopted LDP-based implementations for applications within a metro area network not because it was the technology of choice, but rather because LDP VPLS was the only option available in the vendor's implementation. The scope of VPLS deployments has expanded in terms of the number of PE routers and connected sites. As well, VPLS deployments have expanded in their need to carry new services, such as multicast over VPLS, and to traverse autonomous system (AS) boundaries. Thus, service providers are now looking for solutions to scale and extend VPLS.

In the past, BGP VPLS has been deployed within WANs, whereas LDP VPLS has been deployed in metro networks, for several reasons:

- Ethernet equipment offered only LDP VPLS
- Greater familiarity with BGP and IP VPN in WANs, and correspondingly less familiarity with BGP in metro networks

Now, service providers are seeking to interconnect LDP metro islands with BGP in the core, and some are even considering migrating LDP to BGP within the metro network. With this need to scale and extend VPLS, service providers are now motivated to carefully evaluate and choose their control planes.

Scaling VPLS

Table 1 compares the scaling characteristics of LDP VPLS and BGP VPLS.

Table 1: Comparison of LDP-VPLS and BGP-VPLS Scaling Characteristics

	LDP VPLS	BGP VPLS
Full-mesh requirement	Alleviated only somewhat by H-VPLS, though at the expense of introducing changes and additional overhead in the forwarding plane	Solved by the use of BGP route reflector hierarchy
Provisioning task of adding or removing PE router	Only somewhat simplified by H-VPLS	Highly simplified by the use of route reflector, a technique already proven for VPNs
Provisioning task of adding or changing VPLS customer sites	Manual or through provisioning	Automated by BGP's autodiscovery
VPLS with point-to-multipoint label switched path (LSP) integration to scale forwarding and data planes	Currently not supported in any commercial implementation	Supported
Signaling overhead	Each LDP-VPLS signaling update establishes only one pseudowire; signaling updates increase in proportion to the total number of pseudowires in the network	Each BGP-VPLS signaling update establishes pseudowires with multiple PE router peers; signaling updates can be used to establish multiple pseudowires

The scaling characteristics can determine the scope of a VPLS deployment in the context of a metro network and beyond. The following sections discuss several deployment strategies that optimize the use of BGP VPLS based on its scaling characteristics.

Extending VPLS

Inter-AS

Service providers are seeking mechanisms to extend the VPLS out of an AS. In some cases, a service provider has multiple autonomous systems because of acquisitions and mergers; in other cases, inter-AS operation is a consequence of interprovider VPLS.

LDP VPLS faces the following challenges in providing VPLS that spans multiple ASs:

- Inter-AS LDP VPLS may require the setup of LDP sessions between PE routers that are in different ASs and potentially different administrative domains, or it may require the use of multisegment pseudowires, which has its own complexities.
- The globally significant 32-bit virtual circuit identifier (VCID) used by LDP signaling requires operationally intensive manual coordination between ASs. BGP VPLS also requires site-identifier coordination between ASs.

In contrast, in BGP VPLS, exchange of control traffic between autonomous systems is localized to AS border routers (the so-called option B) or route reflectors (option C), thus facilitating tight control over the information exchanged and such factors as peer authentication. In addition, the use of BGP communities and route target filtering further simplifies the task of determining which VPLS crosses the AS boundary (and to where), and which VPLS remains within the AS.

Juniper Networks® has already helped service providers extend their VPLS across AS boundaries. In fact, in 2004 the first intercarrier virtual private LAN (VPLAN) service, between Hong Kong Hutchison Global Communications and Korea Telecom, was implemented with VPLS in Juniper Networks JUNOS® Software and was extended to include five carriers.

Intermetro Network

Many VPLS deployments are limited to a single metro network, primarily for fear of scaling to "MPLS in the large"—that is, dealing with a large number of PE and provider routers in a single domain and making these PE routers work together to offer VPLS. One potential challenge is finding good methods of scaling the overall network: how to organize the interior gateway protocol (IGP), how to partition control and administrative domains, and so on. Other challenges include the lack of adequate policy mechanisms to control which VPLS remains within a metro network and which should be in intermetro networks.

One useful approach to connecting multiple metro networks and a wide-area backbone in a scalable fashion is to consider each metro network as a separate AS and the backbone as yet another AS. This approach enables independent management of each metro network and the backbone, while enabling the inter-AS control plane (BGP) to work seamlessly. It also provides both intermetro connectivity and intermetro VPLS.

VPLS Deployment Strategies

Several strategies can be used to scale within and extend the VPLS beyond domain boundaries.

Cap and Grow

LDP VPLS works well for small-scale deployments. However, as the VPLS network expands, the LDP-VPLS control plane does not scale. The BGP-VPLS control plane, on the other hand, more easily enables the VPLS network to scale as it expands to support new VPLS customers or to include more sites for existing VPLS customers.

However, many service providers want to scale VPLS in an existing LDP-VPLS deployment. One way to achieve this is to abandon the LDP-VPLS control plane entirely and transition the VPLS network to a BGP-based control plane. However, this option may not be feasible for a number of reasons, including the following two constraints:

- Legacy PE routers deployed in the network may not support a BGP-based signaling mechanism. To protect the current investment, replacement of the legacy PE routers may not be feasible.
- Operating overhead and possible disruption of VPLS for existing customers may undermine a transition from an LDP-VPLS control plane to a BGP-VPLS control plane.

One solution that allows an LDP-VPLS network to scale is to cap the existing LDP-VPLS deployment and to extend the VPLS network using the BGP-based control plane. With this approach, the LDP- and BGP-VPLS control planes, including the signaling mechanisms, both must coexist in the network; therefore, this solution requires interworking between these two mechanisms.

Localize and Extend

To offer regional or national VPLS, service providers are seeking a scalable way to extend the reach of the VPLS of a single LDP-VPLS metro domain. One way to do this is to use BGP VPLS in a WAN core to interconnect multiple LDP-VPLS metro domains. This approach requires a new interdomain technique because the currently defined solutions for multiple ASs require that all domains be running the same signaling protocol.

Again, a critical requirement for such a deployment model is ensuring that legacy PE routers running in metro domains that use the LDP-VPLS control plane do not require any changes or upgrades. A second requirement is that extension of the VPLS be achieved without placing additional burden on the control plane of the LDP-VPLS PE router in the metro network.

To implement either of these strategies, interworking is required between the LDP- and BGP-VPLS domains.

Interworking Scenarios

There are two main scenarios for LDP-BGP VPLS internetworking, with several variations of each.

BGP VPLS and LDP VPLS in Different Domains

In scenarios using BGP and LDP VPLS in different domains, the domains can be metro specific or AS specific, each with several uses. Figure 1 shows a simple case for illustrative purposes.

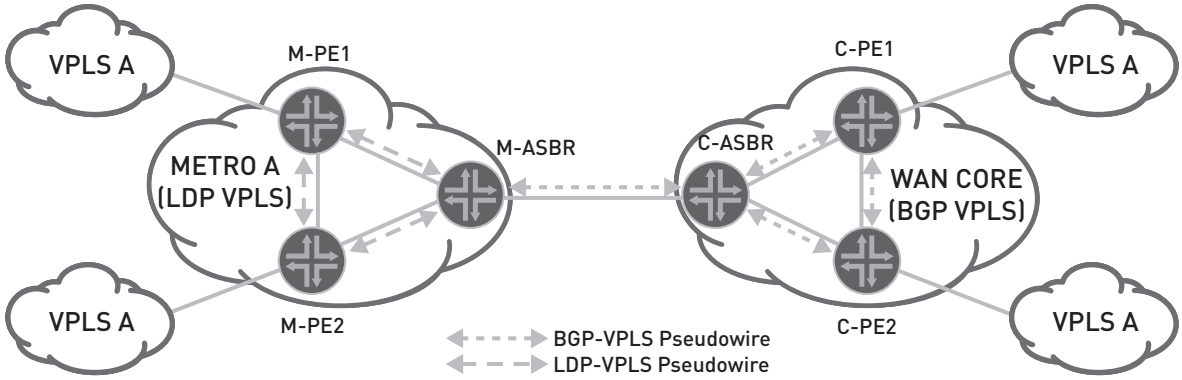

Figure 1: BGP and LDP VPLS in different domains

However, in most scenarios, there would be multiple LDP-VPLS domains that would be connected using a BGP-VPLS domain as shown in Figure 2.

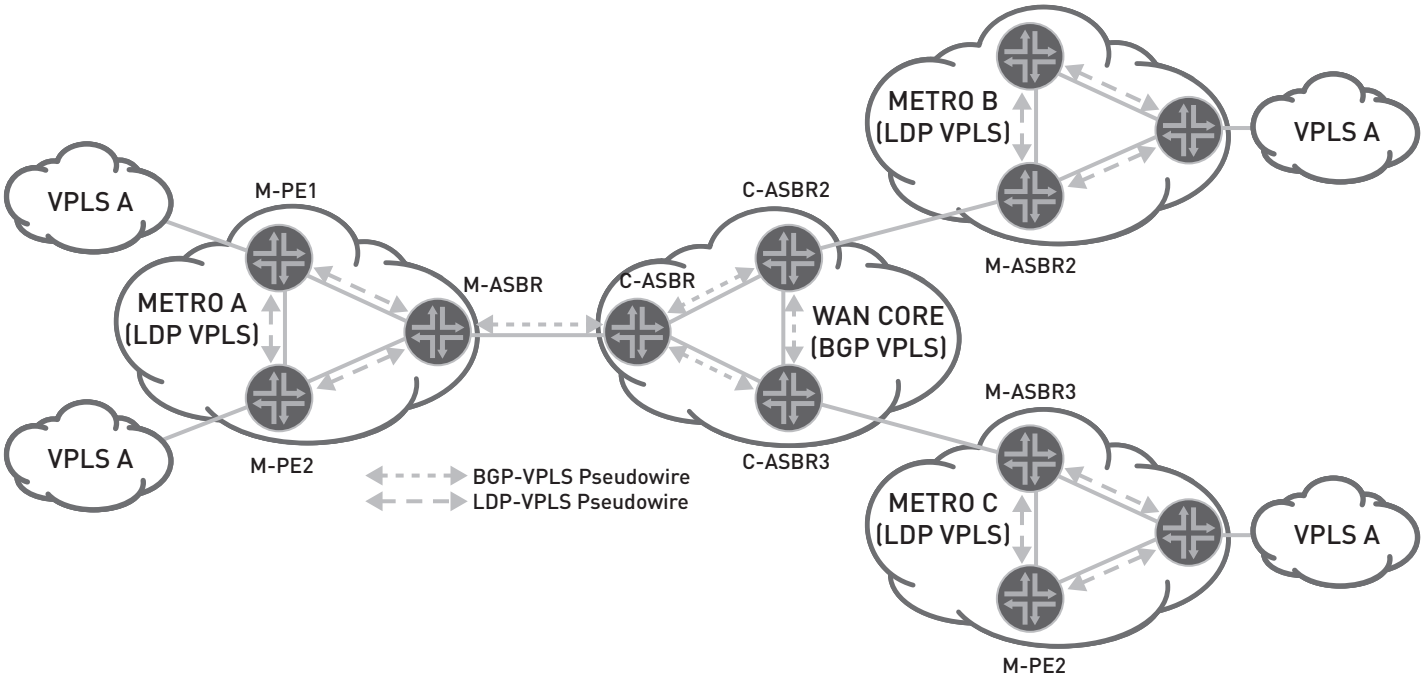

Figure 2: Multiple LDP-VPLS domains connected using a BGP-VPLS domain

Several factors may determine which border PE router in the network performs the interworking function:

- PE router's capability to perform the interworking function
- Administrative control of the routers: whether or not control falls within the AS boundary
- Migration strategy: either cap and grow to address LDP scaling within the metro network, or localize and expand to extend VPLS beyond the metro network

Interworking at M-ASBR

In this scenario (Figure 3), interworking occurs on metro-ASBR (M-ASBR) between LDP VPLS and BGP VPLS. This scenario could employ inter-AS VPLS between M-ASBR and core-ASBR (C-ASBR). It requires a metro device for BGP VPLS.

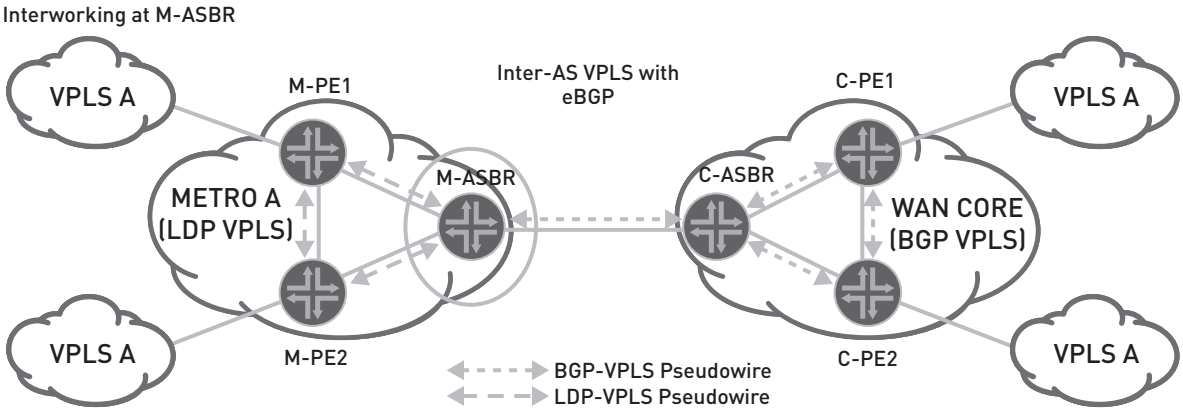

Figure 3: Interworking on M-ASBR between LDP VPLS and BGP VPLS

Interworking at C-ASBR

In this scenario (Figure 4), BGP-VPLS interworking occurs on C-ASBR, between LDP VPLS and BGP VPLS. Minor changes are required on metro devices, and the core devices do most of the heavy work.

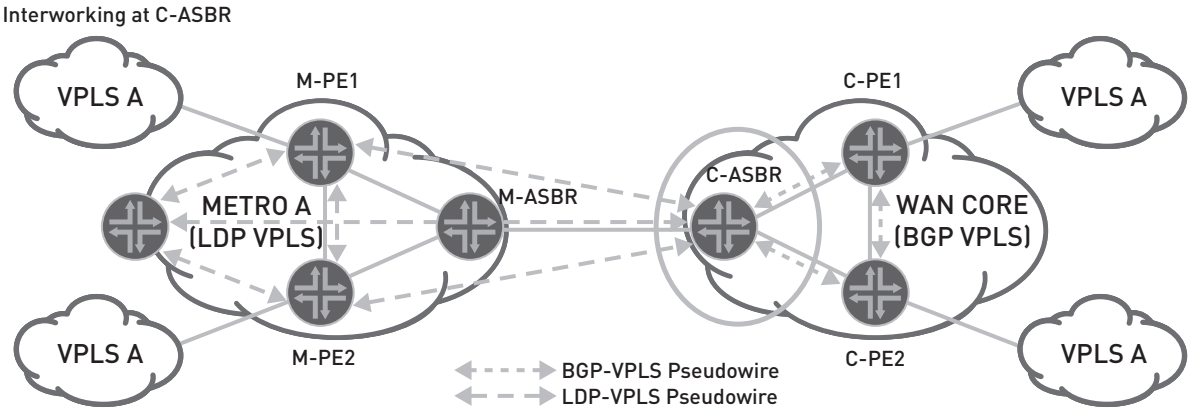

Figure 4: Interworking on C-ASBR between LDP VPLS and BGP VPLS

AS Border Router Supporting Interworking and Functioning as PE Router

Service providers need to have the flexibility to use existing AS border routers both for the interworking function and to terminate customer edge (CE) devices (for revenue-generating customers) directly. The JUNOS Software implementation enables this flexibility (Figure 5).

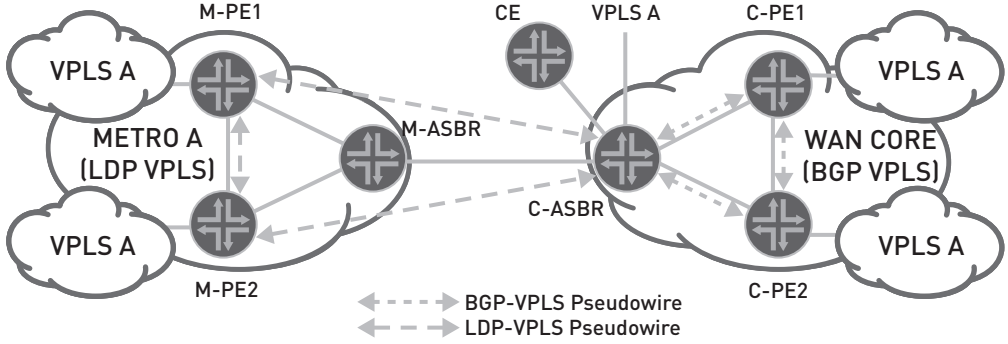

Figure 5: Interworking and PE services on the same ASBR

BGP VPLS and LDP VPLS Within a Single Domain

Scenarios using BGP and LDP VPLS within a single domain require coexistence of both types of control plane on each of the PE routers (Figure 6). In practice, both types of signaling are unlikely to be deployed within the same domain since doing so can create troubleshooting challenges. Nevertheless, such a scheme can be implemented if a service provider chooses to do so.

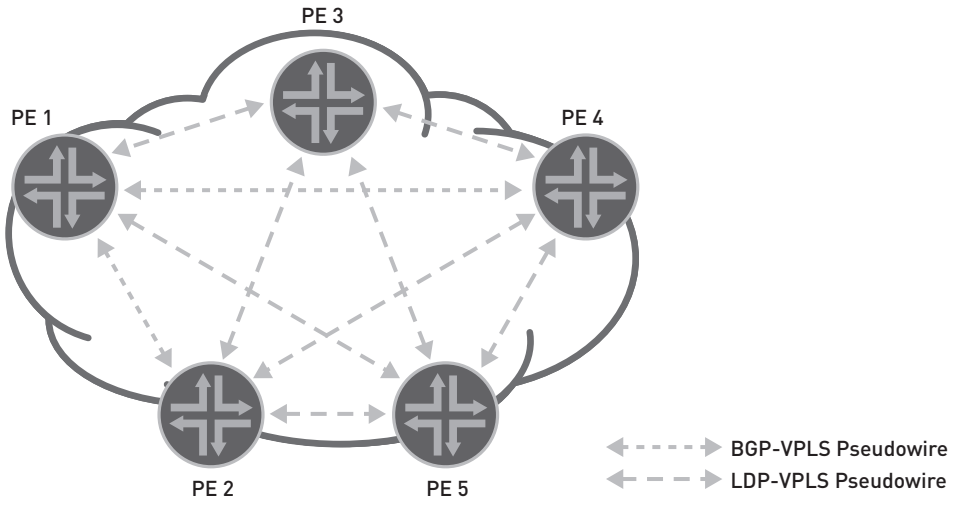

Figure 6: BGP and LDP VPLS within a single domain

Redundancy Using BGP-VPLS Multihoming

Since the interworking function is performed locally on a single node (AS border router), redundancy in the event of failure of the AS border router is critical. BGP VPLS inherently offers redundancy using BGP-based multihoming. PE routers in both LDP-VPLS and BGP-VPLS domains can be multihomed to multiple AS border routers. Each LDP-VPLS domain must be explicitly associated with a BGP-VPLS site that is marked as multihomed. The BGP-VPLS multihoming state machine on C-ASBR elects a single designated forwarder for each multihomed site. PE routers in the LDP-VPLS domain are not even aware of this redundancy, which is completely transparent (Figure 7).

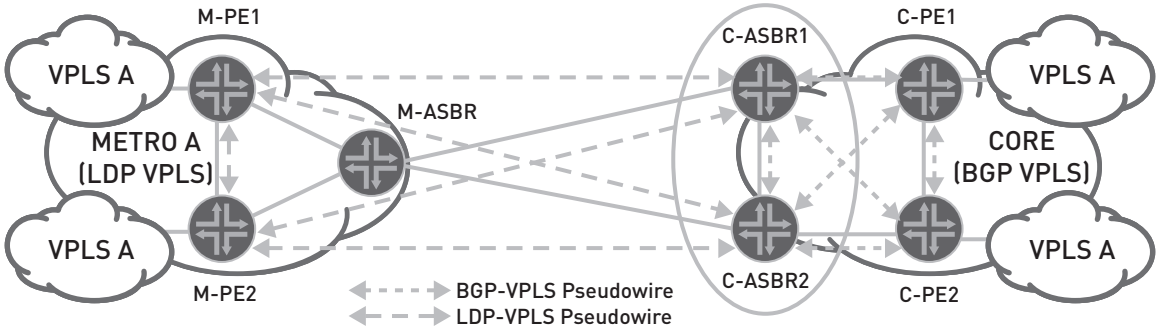

Figure 7: Redundancy using BGP-VPLS multihoming

Interworking Operation on a JUNOS Software Router

JUNOS Software offers a feature-rich implementation of VPLS, as well as the flexibility to deploy either type of VPLS. It is also the first implementation to provide seamless interoperability of both types of VPLS, offering maximum flexibility.

In the PE border router running JUNOS that performs the LDP-BGP VPLS interworking, no changes to existing LDP-VPLS or BGP-VPLS control plane mechanisms are needed. The PE border router running JUNOS can reach to all other (LDP and BGP) PE routers that are part of this VPLS instance.

The concept of mesh groups is central to the implementation of LDP-BGP VPLS interworking in JUNOS Software. Bidirectional pseudowires are created with each PE router for a VPLS instance using either BGP or LDP signaling. From the data plane perspective, each PE router mesh group can be viewed as a virtual pseudowire LAN. All virtual pseudowire LANs (also called a *PE router mesh group*) belonging to a single VPLS instance are stitched together using a common MAC address table (Figure 8).

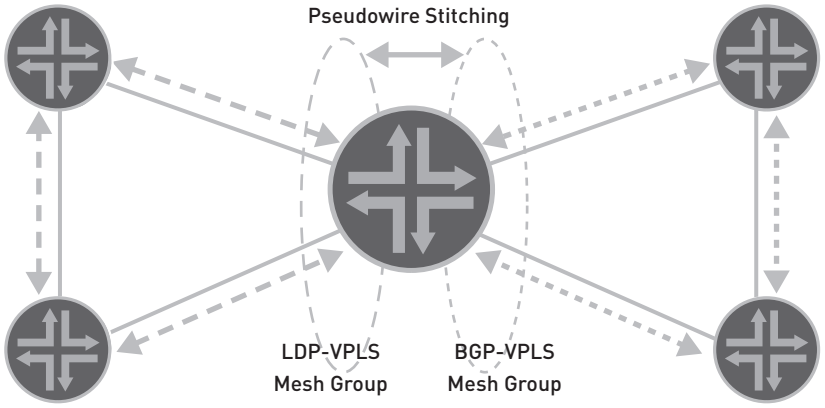

Figure 8: Interworking operation on a JUNOS router

Additionally, mesh groups provide the flexibility to create a two-level hierarchy, as defined by H-VPLS. For each VPLS instance, on a hub PE router one mesh group can be configured for each connected spoke PE router, and another single mesh group can be configured for all the fully meshed remote hub PE routers. On the hub PE router, after these mesh groups are defined according to mesh group forwarding semantics, traffic is bridged between all the spoke and hub mesh groups. Customers then have the choice of using either LDP or BGP for connecting the hubs. Note also that mesh groups can be easily used to support VPLS spanning even multiple LDP domains.

Case Study: Tier-1 Provider

After merging with two other large provider networks, a Tier-1 ISP had these requirements to offer inter-city VPLS across the multiple backbones:

- Integration of existing LDP-VPLS metro services
- Inter-AS VPLS
- Capping LDP VPLS in a legacy Metropolitan Area Network (MAN) with BGP VPLS in the WAN
- Moving to BGP VPLS in the MAN over time

The Tier-1 provider set the requirement to implement BGP VPLS in the common backbone as part of its long-term strategic plan based on its superior scaling characteristics. With the introduction of VPLS PE acting as an ASBR, the model gives them the ability to offer interworking function, ability to directly connect CEs for VPLS and implement BGP multihoming-based redundancy (Figure 9).

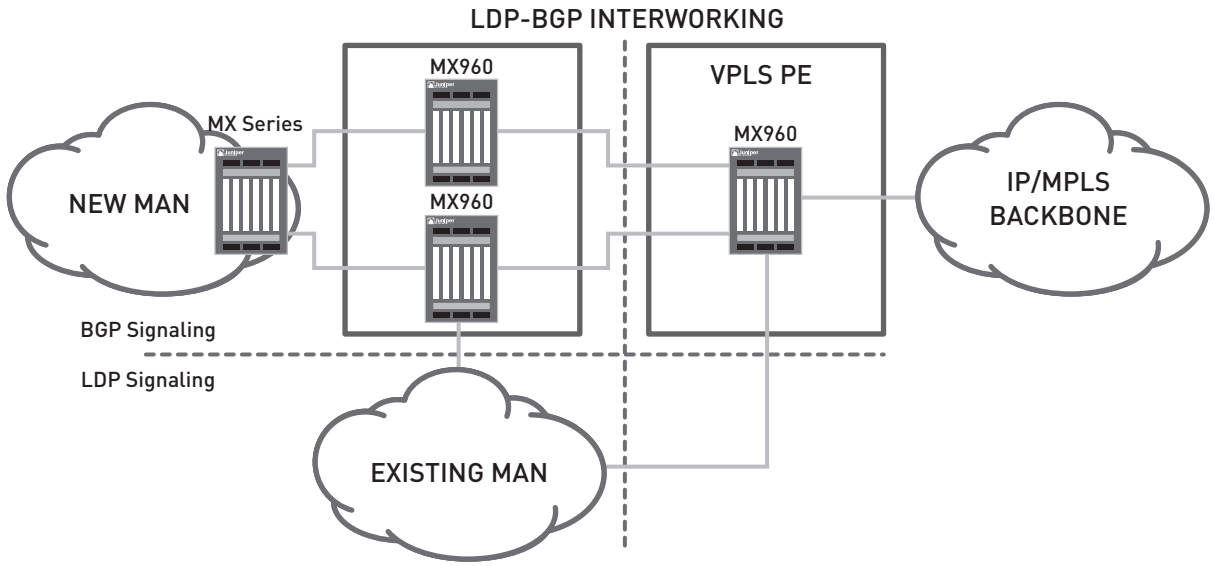

Figure 9: Example of LDP-BGP interworking deployment in a network

Conclusion

Service providers now can choose the type of VPLS they implement based on their requirements and applications, and obtain that solution from a single vendor. Scaling challenges are more effectively addressed by BGP VPLS than by LDP VPLS, and BGP VPLS is also ideal for inter-AS VPLS. Service providers with existing LDP-VPLS deployments can either adopt a cap-and-grow strategy or move toward a localize-and-extend model.

Juniper Networks has a serious commitment to VPLS development and innovation and to meeting the needs of customers who need to interoperate and interwork with other VPLS implementations and who need to make VPLS configuration simpler and less error prone. For information on advanced topics such as multihoming, redundancy, and inter-AS options, refer to other documents on **www.juniper.net.**

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate And Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER
(888.586.4737)
or 408.745.2000
Fax: 408.745.2100

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin,
Ireland
Phone: 35.31.8903.600
Fax: 35.31.8903.601

2000279-001-EN Feb 2009