

SERVICE PROVIDER

Technical Brief: Offering Scalable Layer 2 Services with VPLS and VLL

An overview of Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) and Brocade solutions for deploying them.

BROCADE

CONTENTS

Introduction	3
Technology Overview	4
Virtual Leased Line	4
VPLS	5
VPLS Reference Architectures	8
Scaling VPLS with a Hub-and-Spoke Architecture.....	8
VPLS with Provider Backbone Bridge	9
Other Considerations	10
Multi-homed Spoke VPLS	10
Scalability of PE Routers	10
QoS Control in VPLS and VLL	11
High Availability of the Network	11
Comparison of VPLS and VLL with Other Layer 2 Technologies	11
Comparison of VPLS to BGP/MPLS VPNs	13
Standards Update	14
VPLS and VLL Support in Brocade Products	14
Configuration Examples.....	15
Summary	16

INTRODUCTION

During the last few years, there has been an increasing interest in deploying scalable, end-to-end Layer 2 services. Developments in the area of Layer 2 Virtual Private Networks (VPNs) have helped to bring this solution to market. Chief among these are two Multiprotocol Label Switching (MPLS)-based technologies, Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL), which offer the power of delivering scalable multi-point and point-to-point services, respectively. This paper provides an overview of these technologies, their relative benefits, and Brocade® solutions for implementing them.

Service providers can use these technologies to offer their customers advanced services such as managed VPN services, efficient metro aggregation, and so on. Similarly, large enterprises can use these technologies to do virtual segmentation of the network based on business needs and across geographical boundaries.

Three trends have contributed to the increasing interest in implementing scalable, end-to-end Layer 2 services. Historically, geographically separated networks such as the branch offices of a large corporation have been connected by leased T1 lines or more recently by Frame Relay and ATM connections. Leased lines offer the benefit of a private line, but they typically come at a prohibitive cost. Frame relay and ATM in contrast offer the ability to set up “virtual connections” over a statistically multiplexed network, with definitive Service Level Agreements (SLAs). While the service itself is relatively inexpensive compared to a leased line, the lack of ubiquity and limited volumes of frame relay and ATM equipment have often translated into higher operational and Capital Expenditure (CapEx) costs for both the service provider and the end user.

At the same time, Ethernet has made rapid strides—increasing its viability as an effective means of WAN communication. The high volume of Ethernet ports allows their price points to be far lower than other technologies. For example, according to the Dell’Oro Group, about 237,524,000 Ethernet ports were shipped in CY 2008. This phenomenal volume allows unmatched economies of scale.

A third trend has supported the increasing adoption of MPLS. Although MPLS was initially developed as a means for rapid switching of packets in an IP network for more than just “best-effort” service, the technology itself was soon adapted to scale Ethernet networks and offer VPN services. Both VPLS and VLL leverage the foundation and underlying power of MPLS to accomplish these goals.

Service providers offering VPLS and VLL services enjoy the benefits of lowered OpEx and CapEx as they use a common infrastructure for delivering these services. In addition, the familiarity of Ethernet dramatically reduces the training costs needed for operating such a network.

VLL and VPLS

VLL (Virtual Leased Line) is a point-to-point Ethernet VPN service that emulates the behavior of a leased line between two points. In the industry, the technology is also referred to as Virtual Private Wire Service (VPWS) or EoMPLS (Ethernet over MPLS). VLL uses the pseudo-wire encapsulation for transporting Ethernet traffic over an MPLS tunnel across an IP/MPLS backbone.

VPLS (Virtual Private LAN Service) is a multi-point Ethernet VPN service that leverages the underlying benefits offered by MPLS. In other words, it emulates the behavior of a traditional IEEE 802.1D bridge over an MPLS network. A VPLS service creates a complete Layer 2 broadcast domain for a set of users and is capable of learning and forwarding traffic across a “virtual bridge” based on destination Ethernet MAC addresses.

TECHNOLOGY OVERVIEW

Virtual Leased Line

VLL is the simpler of the two technologies. It effectively offers a pseudo-wire between two points and emulates a leased line service between the two end-points. A VLL is ideal in situations in which traffic patterns in a geographically distributed enterprise are predominantly between select locations.

A Label Switched Path (LSP) is a sequence of MPLS nodes that connects peering routers within the MPLS network. It should be noted that an LSP can carry traffic corresponding to several customers between the same peering routers. Further, because multiple LSPs can exist on the same physical wire, a means of multiplexing and de-multiplexing traffic is needed. The pseudo-wire encapsulation technique uses two labels to accomplish this:

- The outer label is used for LSP tunnel identification.
- The second (inner) label is used for customer identification (technically called Virtual Circuit).

These two labels are inserted as part of the MPLS header after the Ethernet header in the outgoing frame¹. For the payload, the entire incoming Ethernet frame from the customer is encapsulated within the payload of the outgoing MPLS frame, thereby increasing the resulting size of the frame to be transported in the VLL network. The adjacent routers along an LSP need to agree on the tunnel label to be used; this is accomplished as part of MPLS signaling (by either LDP or RSVP-TE if traffic engineering is desired).

Figure 1 shows an example of two separate VLL services being offered to two customers, Customer A and Customer B over an MPLS network.

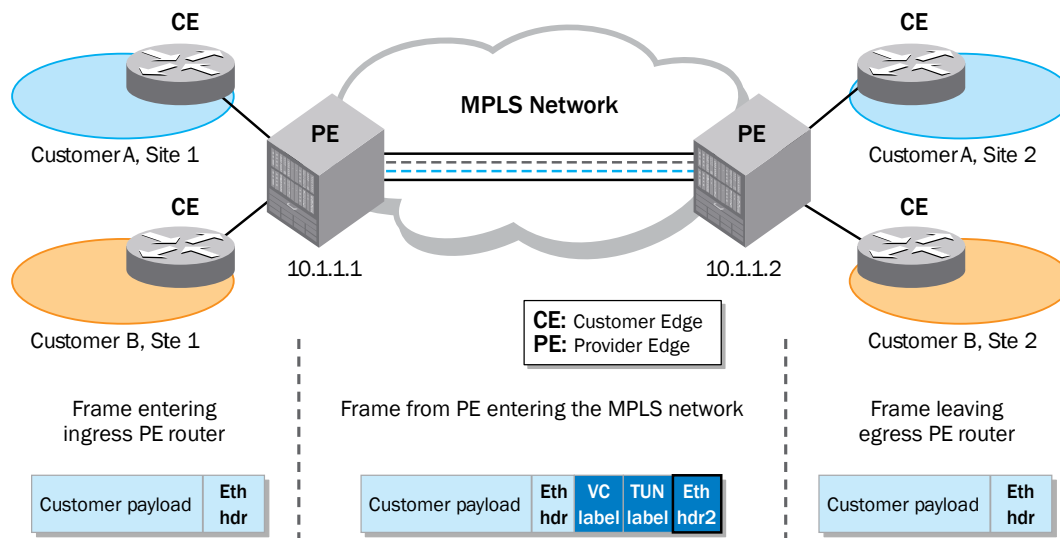


Figure 1. Example of VLL services in an MPLS network (fields in frame are not to scale)

¹ It is assumed here that the interface used to connect to the MPLS network is an Ethernet interface. In general, the MPLS header is inserted after the data link header of the interface.

VPLS

The VPLS architecture defines a mechanism to offer multi-point Ethernet services over a shared MPLS infrastructure. To achieve this, every VPLS instance in the network simulates the behavior of an IEEE 802.1D bridge. This is done by setting up point-to-point pseudo-wires (PW) between a node and every other node in the VPLS instance, thereby creating a full mesh of PWs between all nodes in the VPLS instance. The full mesh of PWs that is created ensures that any node can reach any other node in the VPLS using a single pseudo-wire hop.

Recall that some of the main characteristics of an 802.1D bridge are:

- Maintain a MAC address table that contains the MAC addresses learnt on a port
- Use the MAC address table to determine the destination port for incoming frames
- Flood frames with unknown destination MAC address, broadcast MAC address, or multicast MAC address to all ports in that bridge instance.
- Use a mechanism to prevent loops in the Layer 2 bridged network

Because a VPLS instance emulates the behavior of a bridge, it also performs identical functions. Figure 2 shows a typical VPLS network.

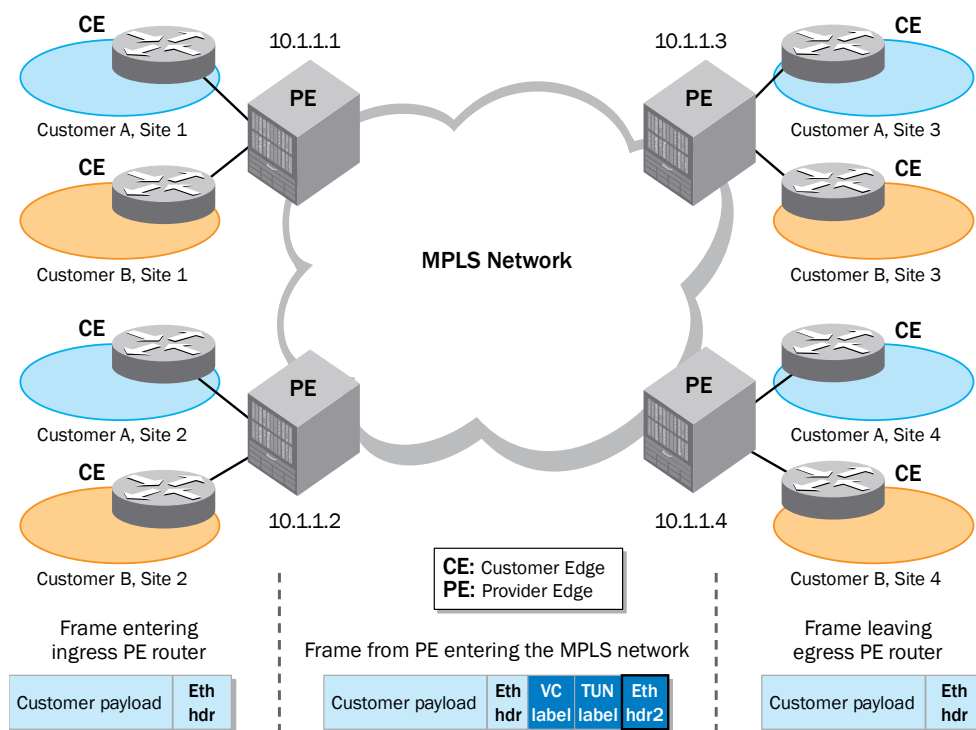


Figure 2. VPLS network (fields in frame are not to scale)

Because the two customers are distinct, each customer belongs to a unique VPLS instance, as shown in Figure 2. Each border node in the VPLS network is called a Provider Edge (PE) router since that router denotes the demarcation point of a service provider's network. The switch/router on the customer site that communicates with the PE router is designated the Customer Edge (CE) device. The connection between the PE and CE routers is also referred to as the Attachment Circuit (AC). Within the VPLS network, there is a full mesh of pseudo-wires (PWs) established between the PE routers in the VPLS network. The LSPs that carry pseudo-wires are unidirectional in nature. Therefore, a pair of LSPs is created, one in each direction between a pair of PE routers. In Figure 2, there are a total of 12 LSPs created. In general, for a VPLS network with n PE routers, there will be $n*(n-1)/2$ pseudo-wires and $n*(n-1)$ LSPs created.

Figure 3 shows a logical view of the same network shown in Figure 2. The two VPLS instances are distinct. Both customer A and customer B share the same MPLS infrastructure, but the traffic on each VPN is completely isolated from that on the other VPN.

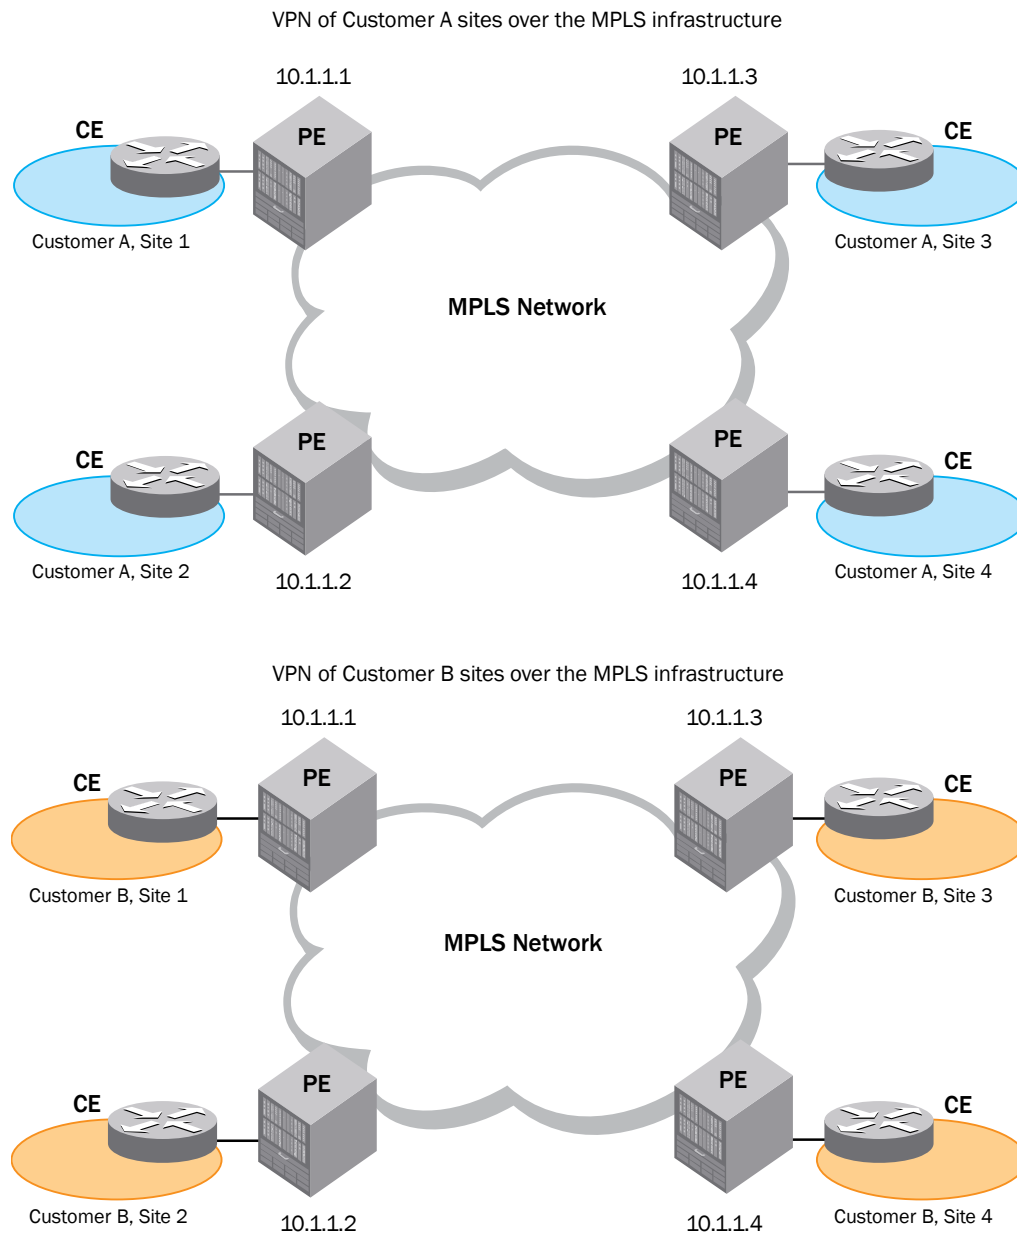


Figure 3. Virtual private network view of customer A and customer B sites

As in the case of VLLs, multiple LSPs can exist on the same physical wire. Therefore, a means of multiplexing and de-multiplexing traffic is needed. Two labels are used on each pseudo-wire—the outer label is used for LSP tunnel identification and the second (inner) label is used for customer VPLS identification, that is, the inner label represents a Virtual Circuit carrying a customer’s traffic between two PEs. These two labels are inserted in the MPLS header after the Ethernet header. The adjacent routers along an LSP need to agree on the outer label to be used. This is accomplished as part of MPLS signaling (either by LDP or by RSVP-TE if traffic engineering is desired). The inner label is inserted at the ingress PE router and used for de-multiplexing of traffic only by the egress PE router; transit routers do not look at the inner label.

In order to provide multi-point connectivity, the VPLS should support flooding of a frame received with an unknown destination MAC address (that is, a MAC address that has not yet been learnt) or a frame received with a broadcast/multicast destination MAC address. Because of the presence of a full mesh of PWs, the VPLS should ensure that loops are not formed in the process of flooding traffic. This is accomplished by using a “split-horizon” technique. A frame that is received from a pseudo-wire and requires flooding is not sent over any other pseudo-wire in the VPLS. On the other hand, a frame that requires flooding and was originally received from an attachment circuit by the PE is sent to all pseudo-wires that are part of the VPLS.

The PE-CE interface can support a variety of configurations:

- An untagged Ethernet interface
- A tagged Ethernet interface

When the VPLS is configured, the PE node is configured with the information of whether the tag in the incoming frame is used for service delimiting by the provider or whether the tag is local to the customer, and hence of no significance for service delimiting purposes by the network. In the former, the tag is used to map to a VPLS instance, so its use in delimiting a service qualifies it as a “service-delimiting” tag. In the latter case, any tag in the incoming frame is transported transparently through the network, and the PE router does not process the tag in the incoming frame.

When service-delimiting tagged Ethernet interfaces are used, the VLAN-ID used is of local significance only. This leads to many interesting scaling properties in a VPLS network—the two PEs can be configured such that when a frame exits the VPLS instance, a different VLAN ID can be inserted. In other words, the VPLS can be used to do VLAN translation. On well-designed PE routers, the entire range of 4K VLAN IDs can be configured on each port of the PE router. These concepts can also be extended to Q-in-Q to achieve even greater scalability.

VPLS REFERENCE ARCHITECTURES

Scaling VPLS with a Hub-and-Spoke Architecture

The architecture previously described involves setting up a full mesh among all participating PE routers. For scalability, it is often advisable to limit the number of nodes that are added to this full mesh. An alternate architecture involves a hub-and-spoke architecture in which “smart” Multi-Tenant Unit (MTU) switches are placed by the service provider at the edge of the network and connected to the PE router via spoke connections. The PE acts as a hub device aggregating traffic from multiple MTU devices, shown in Figure 4.

The MTU is a switch that is capable of performing functions such as Q-in-Q. Therefore, traffic from each customer can be assigned a unique provider VLAN tag and sent over a Q-in-Q tunnel to the PE device. The PE router should be capable of mapping the provider VLAN tag (outer VLAN tag) to a distinct VPLS instance.

The hub-and-spoke scheme has the advantage of separating manageability of the core network from that of the network edge. In addition, the MTU can also be used to switch local traffic without sending traffic to the PE.

This hierarchical separation between the MTU and the PE router also allows different types of MTU devices to be used by a service provider based on the needs of the customers being serviced. For example, more advanced techniques such as service-based Q-in-Q can also be used to map traffic from customer edge devices to the spoke connection. By right-sizing the capabilities in the MTU, the MTU can be offered at a much lower price point compared to a PE router and thus decrease network costs.

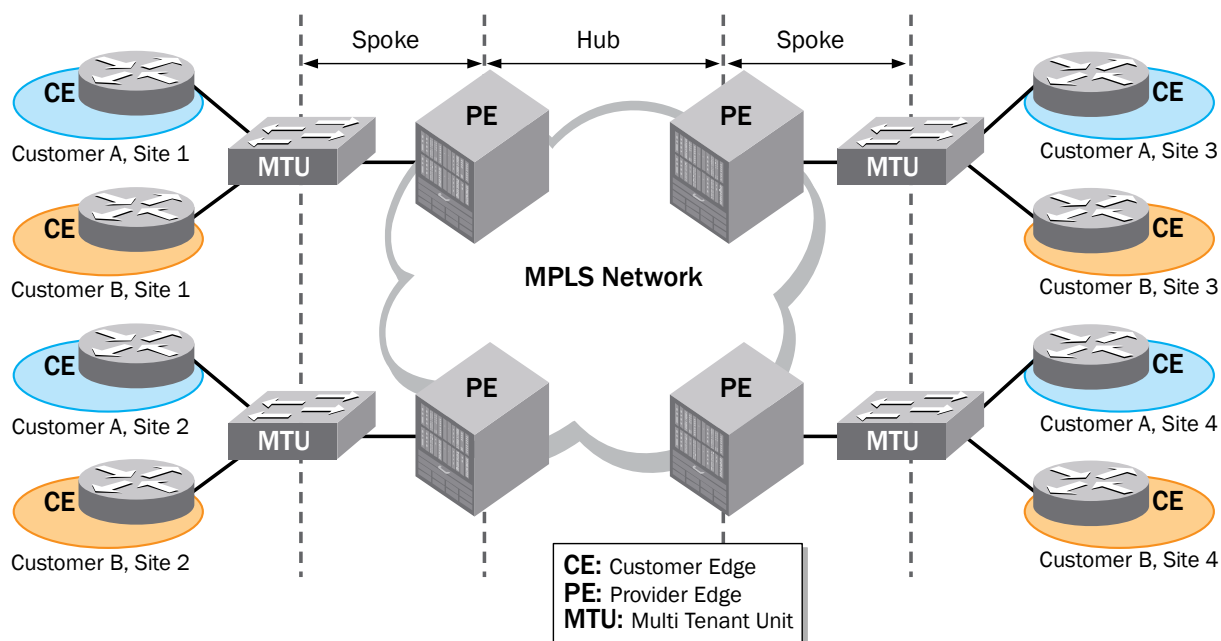


Figure 4. Hub-and-spoke model for scalability in the network

VPLS with Provider Backbone Bridge

An alternate model of delivering scalable Carrier Ethernet services to the network edge combines VPLS in the core with Provider Backbone Bridge (PBB) at the network edge. Combining VPLS and PBB in this manner offers several benefits:

- An efficient, secure architecture while preserving investment in existing MPLS core
- Reduced number of PE routers that need to establish a full mesh for VPLS
- Lower signaling overhead on the PE: enables Carrier Ethernet service to scale further
- MAC hiding, isolates the MPLS PE device from learning customer MAC addresses
- Simplified management at the edge

A combination of VPLS with PBB is shown in Figure 5.

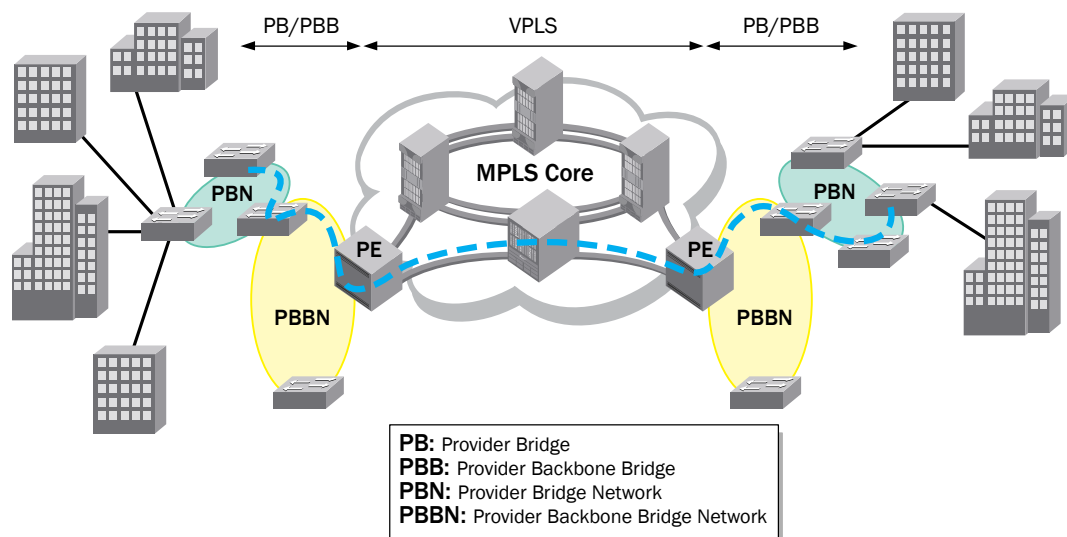


Figure 5. Connecting VPLS with PB/PBB for scalable Carrier Ethernet service delivery

OTHER CONSIDERATIONS

Multi-homed Spoke VPLS

An important prerequisite for offering very high Service Level Agreements (SLAs) to a provider's customers involves elimination of any single point of failure. This is particularly important at the edge where several customers can be served by the same MTU device. Brocade switches that are used as MTUs support a capability called "protected link," with which a standby link can protect the failure of an active link. This capability allows very fast switchover (in less than 100 milliseconds) from the active to the protected link in the event of failure on the active link.

The protected link capability on an edge MTU device allows an MTU to be multi-homed to different PE routers, thereby offering the benefits of redundant connectivity as well as very fast recovery during a failover. Alternatively, Rapid Spanning Tree Protocol (RSTP) can also be used on the MTU if dual-homing is deployed, but the convergence time with RSTP after failover is longer compared to protected links.

Scalability of PE Routers

As might be apparent from information in previous sections, routers that support VPLS should have efficient support for the underlying capabilities required to offer a VPLS service. Poorly designed VPLS devices have trouble performing the arduous tasks of managing large MAC address tables, performing signaling for a large number of LSPs, and efficiently handling broadcasts of frames. In extreme cases, these stressful scenarios could lead to heavy CPU utilization, which eventually leads to router crashes.

The number of customers, number of customer sites supported, and the number of devices at each customer site have a direct impact on the size of the MAC address table. In large networks, hundreds of thousands of MAC addresses may have to be maintained, leading to the need for large MAC address tables on each PE device. Brocade VPLS-capable routers such as the Brocade NetIron® XMR can concurrently support up to 1 million MAC addresses in its VPLS MAC table.

As mentioned in the "Technology Overview" section, the number of LSPs that have to be set up and maintained can lead to a heavy load on the control plane. The PE router should, therefore, have a high rate of establishing LSPs to ensure acceptable performance. The Brocade Multi-Service IronWare® operating system has a modular design with multiple internal priorities to handle task scheduling and inter-process communication messaging. Ironware operating system runs on ultra-fast CPUs on the management and interface modules that distribute the processing load to achieve superior performance.

Finally, handling frames with unknown MAC addresses must be done efficiently. Given the large number of instances typical in a VPLS network, receipt of several frames with unknown MAC addresses should not lead to system overload. Brocade VPLS-capable routers such as the NetIron XMR provide hardware support for handling this flooding, which ensures that replication of these frames is done without any intervention by the CPU.

QoS Control in VPLS and VLL

So far, this paper has focused on scalability of the Layer 2 service. Another area in which VPLS and VLL strongly differentiate themselves from other Layer 2 VPN technologies is support for Quality of Service (QoS). There are two primary methods of controlling QoS in the network:

- Assign a Class of Service (CoS) value for the VPLS instance at configuration time. In this method, the ingress PE router overrides the priority of an incoming frame and selects a tunnel LSP with a CoS value that matches what was configured for the VPLS instance.
- Assign a distinct CoS value for each MPLS frame entering the network. In this method, a common LSP tunnel is used for traffic between two PEs, irrespective of the VPLS instance to which it belongs. Discrimination between different traffic types in the MPLS network is done by using different EXP bits in the MPLS frame. The EXP bits are set by mapping the 802.1p CoS value from the incoming tagged Ethernet frame to desired EXP bits. If the incoming frame is untagged, the router can also be configured to assign a priority to the port to determine the CoS value for an incoming untagged Ethernet frame.

In an MPLS network, transit PE routers and ingress/egress PE routers give an appropriate priority to frames marked with a higher CoS value when making scheduling and congestion-handling decisions.

High Availability of the Network

A powerful advantage of VPLS is the inherent reliability in the service that can be achieved from the use of the underlying MPLS network. Capabilities such as load balancing, fast detection of failures, and adaptive re-routing (also called “Fast Reroute”) around node or link failures all contribute to high resiliency of the services that are offered on the network.

To achieve load balancing, multiple LSPs can be used between two PE devices. The multiple LSP paths can actively share the traffic carried between the PEs. The LSPs can also be configured as active or standby LSPs, with failure of the primary LSP triggering failover to the standby LSP. Failure detection can be done using RSVP messages. If very fast detection is desired, Bidirectional Failure Detection (BFD) can also be used.

COMPARISON OF VPLS AND VLL WITH OTHER LAYER 2 TECHNOLOGIES

Layer 2 VPNs are increasing in popularity today for several reasons. Price erosion in Ethernet and the ubiquity of Ethernet has made it possible to rapidly slash both CapEx and OpEx. Technically, VPLS and VLL technologies offer many benefits over competing alternatives such as MAC-in-MAC (also known as Provider Backbone Bridges), L2TPm and Q-in-Q. Note that some of these technologies can be used in combination (for example, Q-in-Q in combination with VPLS or MAC-in-MAC in combination with VPLS).

Table 1 compares VPLS and VLL with other Layer 2 VPN technologies.

Table 1. Comparison of different Layer 2 VPN technologies

	VPLS	VLL	Ethernet PW over L2TPv3	Provider Backbone Bridge (PBB)	Provider Bridge (PB)
Connectivity	Multi-point	Point-to-point	Predominantly point-to-point	Both point-to-point and multi-point	Both point-to-point and multi-point
Technology maturity	Advanced state	Advanced state	Infancy	Infancy	Mature
Standardization efforts	Standardized	Standardized	Infancy	Standardized	Standardized
Protocol overhead in data plane	22 bytes of additional information in every frame	22 bytes of additional information in every frame	Variable: can be up to 44 bytes depending on how the L2TP messages are transported	22 bytes of additional overhead ²	Additional 8 bytes of data on top of a traditional Ethernet frame
Scalability	Highly scalable particularly with hub-and-spoke architecture	Highly scalable	Highly scalable	Highly scalable	Limited to at most 4K instances in the network
Multicast support	Yes	No	No	Yes	Yes
Traffic engineering and QoS	Very rich: makes use of underlying MPLS network T.E. features	Very rich: makes use of underlying MPLS network T.E. features	Poor	Limited traffic engineering. QoS accomplished by manipulation of 802.1p bits	Limited traffic engineering. QoS accomplished by manipulation of 802.1p bits
Vendor support	Wide-spread	Wide-spread	Limited; many vendor implementations are still at L2TPv2 that requires PPP	Currently limited	Wide-spread
Cost	Varies by vendor	Low	Typically high	Varies by vendor	Low
OAM	Via MPLS OAM	Via MPLS OAM	Not defined	Via 802.1ag	Via 802.1ag

² Assumes full Ethernet encapsulation (aligned with MPLS Martini headers) per IEEE 802.1ah

	VPLS	VLL	Ethernet PW over L2TPv3	Provider Backbone Bridge (PBB)	Provider Bridge (PB)
Provisioning complexity	Very simple, particularly in comparison to L3 VPNs but more involved than simpler Q-in-Q techniques. Also supported by several provisioning systems	Very simple. Also supported by several provisioning systems	Complex. Limited support by provisioning systems	Relatively simple. Limited support by provisioning systems.	Very simple. Widespread support by provisioning systems
Complexity in troubleshooting	May be involved based on the extent of capabilities in the underlying MPLS network that are used	Fairly simple	High	May be involved	Very simple
Solution robustness and resiliency	Very high; with use of procedures such as MPLS fast re-route and hot standby LSPs, a high level of SLA can be guaranteed	Very high; with use of procedures such as MPLS fast re-route and hot standby LSPs, a high level of SLA can be guaranteed	Moderate: can detect connectivity losses but difficult to have efficient sub-second recovery	High; requires the use of other L2 techniques such as MRP or RSTP	Low; requires the use of other L2 techniques such as MRP or RSTP

COMPARISON OF VPLS TO BGP/MPLS VPNs

How does VPLS perform compared to BGP/MPLS VPNs? Both approaches have their own merits and which technology to use depends on several criteria.

BGP/MPLS VPNs, also known as 2547bis VPNs, make use of BGP to propagate route information from various Layer 3 VPNs to the relevant peer PEs that also host the same Layer 3 VPN. The P (Provider – transit) routers are oblivious to routing protocol exchanges between the PE routers. The PE routers use LSP tunnels to forward Layer 3 VPN packets from one PE to the other through the MPLS network. Finally, the Layer 3 VPN routes that are learned from a peer PE router are propagated by a PE to the attached CE router. Neither technology is the “magic wand” for all problems.

VPLS is very simple to administer. It is ideal when a provider does not desire—and customers do not require—the administration of customer routing protocols within the provider network. Configuration is very simple and only the peer PE routers for a VPLS instance need to be specified. BGP VPNs require sound knowledge of routing protocols in order to correctly administer them. As the number of instances increase, service provisioning systems are often recommended in both cases to ease the burden on the administrator, particularly for L3 VPNs.

Layer 2 VPNs also enjoy a clear separation between the customer's network and the provider's network—a fact that has contributed heavily to its increasing popularity. Each customer is still free to run any routing protocol and that choice is transparent to the provider. It is also not necessary to run any Spanning Tree Protocol within the provider network (even though it emulates a Layer 2 service).

Layer 3 VPNs are geared toward transport of IP traffic only. Although IP is nearly ubiquitous, niche applications exist that require IPX or AppleTalk or other non-IP protocols. While L3 VPNs cannot be used in such cases, L2 VPNs can. Even a transition from IPv4 to IPv6 in a customer's network is completely seamless to the provider when Layer 2 VPNs are used.

Unlike a Layer 3 VPN, there is no separate control plane protocol in a Layer 2 VPN that is used to exchange reachability information. Rather, the data plane itself is used to build information related to reachability using standard MAC address learning procedures.

In terms of scalability, both approaches are highly scalable. Critics of each approach often point to the number of routes that are maintained by the PE router (in L3 VPNs) or the number of MAC addresses that are maintained by the PE router (in L2 VPNs) as deficiencies of the other approach. In practice, there are techniques available to limit the impact of large tables in both cases. For instance, route summarization is often recommended in L3 VPNs. Similarly, the use of hub-and-spoke architectures or limiting the number of MAC addresses per customer should be used in L2 VPNs when scalability is a concern.

The cost of each solution varies widely by vendor. In many cases, the cost of a Layer 3 VPN-enabled PE router is more than the cost of a Layer 2 VPN-enabled PE router due to the requirement to maintain multiple VPN Routing/Forwarding (VRF) tables. Note that this is not true for Brocade NetIron XMR or NetIron MLX series routers, since the same router is cost- and performance-tuned to support both L2 and L3 VPNs.

STANDARDS UPDATE

The standardization process for VPLS and VLLs has involved the active participation of the industry, both in development of the standards and development of products to conform to those standards. The IETF L2VPN group is responsible for setting standards for these two technologies. VPLS using LDP as the signaling protocol is standardized as RFC 4762. VLL is standardized as RFC 4448.

For VPLS, most industry participants agree on the use of LDP as the signaling protocol for establishing the PWs required for VPLS. An alternate standard (RFC 4761) that uses BGP for such signaling is also available, but has limited acceptance by vendors.

Procedures for OAM and multicast transport optimizations within a VPLS are still topics of discussion in the working groups.

VPLS AND VLL SUPPORT IN BROCADE PRODUCTS

Brocade has a broad range of products that support VPLS and VLL. From MTU to provider edge to aggregation to provider core applications, these products offer the complete range needed for service providers to build scalable Layer 2 networks:

- Brocade NetIron XMR Series is a family of high-end, carrier-class, MPLS backbone routers offered in various configurations, including the NetIron XMR 4000, XMR 8000, XMR 16000, and XMR 32000 routers.
- Brocade NetIron MLX Series is a family of MPLS-enabled switching routers with unique scalability for Layer 2 metro applications, including the NetIron MLX-4, MLX-8, MLX-16, and MLX-32 routers.

In addition, the Brocade NetIron CES 2000 Series, FastIron® Edge Switch (FES), and Fast Iron Edge Switch Series-X (FES-X) are compact MTUs with high port density for use in the network architectures described in this paper. The NetIron CES 2000 Series of compact switches support Provider Backbone Bridging (PBB) technology and allows scalable extension of VPLS services to the edge of the network by combining VPLS in the core with PBB at the network edge.

Brocade support for these technologies allows unparalleled scalability for a service provider. For example, the NetIron XMR series of routers allows 1 million MAC addresses to be maintained in its MAC tables when using VPLS. With the ability to support up to 4K active VLAN-IDs per port on all ports of a system simultaneously, large-scale VPLS networks can easily be deployed. These figures should be considered in the context of the industry-leading density of 1 Gigabit Ethernet (GbE) and 10 GbE ports on the NetIron XMR and NetIron MLX platforms. The low-latency, wire-speed performance for VPLS and VLL services enables very high performance. Brocade VPLS/VLL routers also enable sFlow sampling on VPLS and VLL endpoints, thereby simplifying troubleshooting of customer traffic in the event of service failures.

For large-scale Layer 2 services, service provisioning solutions are needed to simplify network management. Brocade and Oracle jointly deliver a cohesive end-to-end service delivery solution. Using Oracle Communications IP Service Activator (IPSA) suite, VPN and Ethernet services can be reliably and rapidly activated on Brocade routers and switches. The service activation platform can also be used to manage a multi-vendor network, further simplifying the management of such a network.

Brocade believes in using the right technology for the right application and giving service providers the freedom to select services for deployment as needed by their target market. Accordingly, the products allow multiple services such as VPLS, VLL, and BGP/MPLS VPNs to be enabled concurrently on the same port.

Configuration Examples

Configuring a VPLS on a Brocade device is quite straightforward. Figure 7 shows the commands required for configuring the two VPLS instances for the network in Figure 2. The configuration commands for only PE1 and PE2 are shown; the commands for PE3 and PE4 are similar. This example assumes that Customer A wishes to use VLAN-ID 500 on the PE-CE interface at site 1 and VLAN-ID 400 on the PE-CE interface at site 2. As previously described, the VPLS effectively does a VLAN translation. It is also possible to retain the same VLAN-ID as shown in the configuration statements for customer B. The values 10000 and 10100 are the Virtual Circuit Identifiers (VC-ID); the VC-ID should be same on all the PE routers for the same VPLS instance. Additional options to control CoS, load balancing, and so on, are available; for details, refer to the configuration guide of the appropriate product.

```
On PE1:
vpls customerA 10000
 vpls-peer 10.1.1.2 10.1.1.3 10.1.1.4
 vlan 500
  tag ethernet 1/1
vpls customerB 10100
 vpls-peer 10.1.1.2 10.1.1.3 10.1.1.4
 vlan 500
  tag ethernet 1/2

On PE2:
vpls customerA 10000
 vpls-peer 10.1.1.1 10.1.1.3 10.1.1.4
 vlan 400
  tag ethernet 1/1
vpls customerB 10100
 vpls-peer 10.1.1.1 10.1.1.3 10.1.1.4
 vlan 500
  tag ethernet 1/2
```

Figure 7. Configuration example for VPLS

Configuring a VLL on a Brocade device is similarly very simple, as shown in Figure 8. These configuration commands are for the network in Figure 1. This example assumes that Customer A wishes to use VLAN-ID 300 on the PE-CE interface at site 1 and the VLAN-ID 200 on the PE-CE interface at site 2. On the other hand, Customer B wishes to use the same VLAN-ID (300) at both sites.

```
On PE1:
vll customerA 10000
  vll-peer 10.1.1.2
  vlan 300
    tag e 1/1
vll customerB 10100
  vll-peer 10.1.1.2
  vlan 300
    tag e 1/2

On PE2:
vll customerA 10000
  vpls-peer 10.1.1.1
  vlan 200
    tag e 1/1
vll customerB 10100
  vll-peer 10.1.1.1
  vlan 300
    tag e 1/2
```

Figure 8. Configuration example for VLL

SUMMARY

Both VPLS and VLL are excellent mechanisms to offer scalable Layer 2 services over a converged infrastructure. These technologies offer the benefits of a standards-based approach to offer scalable Layer 2 services. With the benefits of the underlying MPLS network's advanced traffic engineering, QoS, and resiliency properties, sophisticated service level agreements can be offered by a service provider to end customers. In addition, deploying such services over a converged network allows service providers to slash capital and operating costs. Brocade's solutions for these technologies allow service providers to rapidly and cost-effectively deploy such services over a converged MPLS infrastructure.

For an overview of PBB, common PBB deployment scenarios, and benefits when deployed with a core MPLS network supporting VPLS, visit the Brocade Web site and see "Leveraging the Benefits of Provider Backbone Bridges" www.brocade.com/forms/getFile?p=documents/white_papers/Leveraging_Benefits_PBB_WP-00.pdf

© 2009 Brocade Communications Systems, Inc. All Rights Reserved. 12/09 GA-TB-177-01

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronPoint, IronShield, IronView, IronWare, JetCore, NetIron, SecureIron, ServerIron, StorageX, and Turbolron are registered trademarks, and DCFM, Extraordinary Networks, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.