

Offering scalable Layer 2 services with VPLS and VLL

Introduction

Over the last few years, there has been an increasing interest in deploying scalable, end-to-end Layer 2 services. Several developments in the area of Layer 2 VPNs have helped to convert this into reality. Chief among these are two MPLS-based technologies, Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) that offer the power of delivering scalable multi-point and point-to-point services respectively. This paper gives an overview of these technologies, the relative benefits of these technologies, and Foundry Networks' solutions to implementing them.

Applications

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

Background

Three trends have contributed to the increasing interest in implementing scalable, end-to-end Layer 2 services. Historically, geographically separated networks such as 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. 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 costs for both the service provider and the end-user.

At the same time, Ethernet has made rapid strides in the last decade increasing its viability as an effective means of WAN communication. The high volume of Ethernet ports allows the price points

of Ethernet ports to be far lower than other technologies. For instance, Gartner estimates that worldwide, close to 55 million Ethernet ports were shipped in just the 3rd quarter of 2005 alone. This phenomenal volume allows unmatched economies of scale.

A third trend has been the increasing adoption of MPLS. Although MPLS was initially developed as a means for rapid switching of packets in an IP network with 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 this goal.

Service providers offering VPLS and VLL services enjoy the benefits of lowered operational and capital costs 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.

Difference between VPLS and VLL

VLL (Virtual Leased Line) is a point-to-point Ethernet VPN service that emulates the behavior of a leased line between 2 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

VLL

VLL is the simpler of the two technologies. It effectively offers a pseudo-wire between two points, and emulates a leased line service between those two end-points. A VLL is ideal in situations when 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 may carry traffic corresponding to several customers between the same peering routers. Further, because multiple LSPs may exist on the same physical wire, a means of multiplexing and de-multiplexing traffic is needed. The pseudo-wire

encapsulation technique uses 2 labels to accomplish this—the outer label is used for LSP tunnel identification and the second (inner) label is used for customer identification (technically called Virtual Circuit). These 2 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 (e.g. either by LDP or by RSVP-TE if traffic engineering is desired).

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

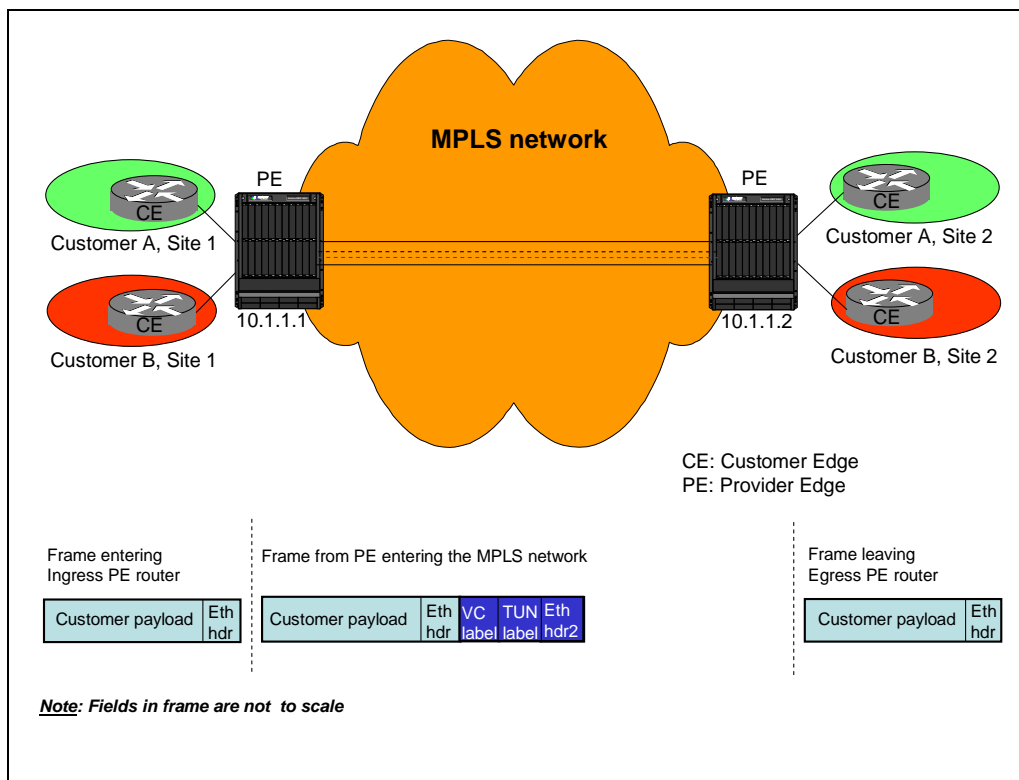


Figure 1: Example of VLL services in an MPLS network

Figure 1 above gives an example of 2 separate VLL services being offered to two customers, Customer A and Customer B over an MPLS network.

For a more detailed description of deploying VLL services, please refer to the white paper on this topic [1].

Application note:
Offering scalable Layer 2 services with VPLS and VLL

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.

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

- Maintenance of 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 should also perform identical functions. Figure 2 below is a pictorial representation of a typical VPLS network.

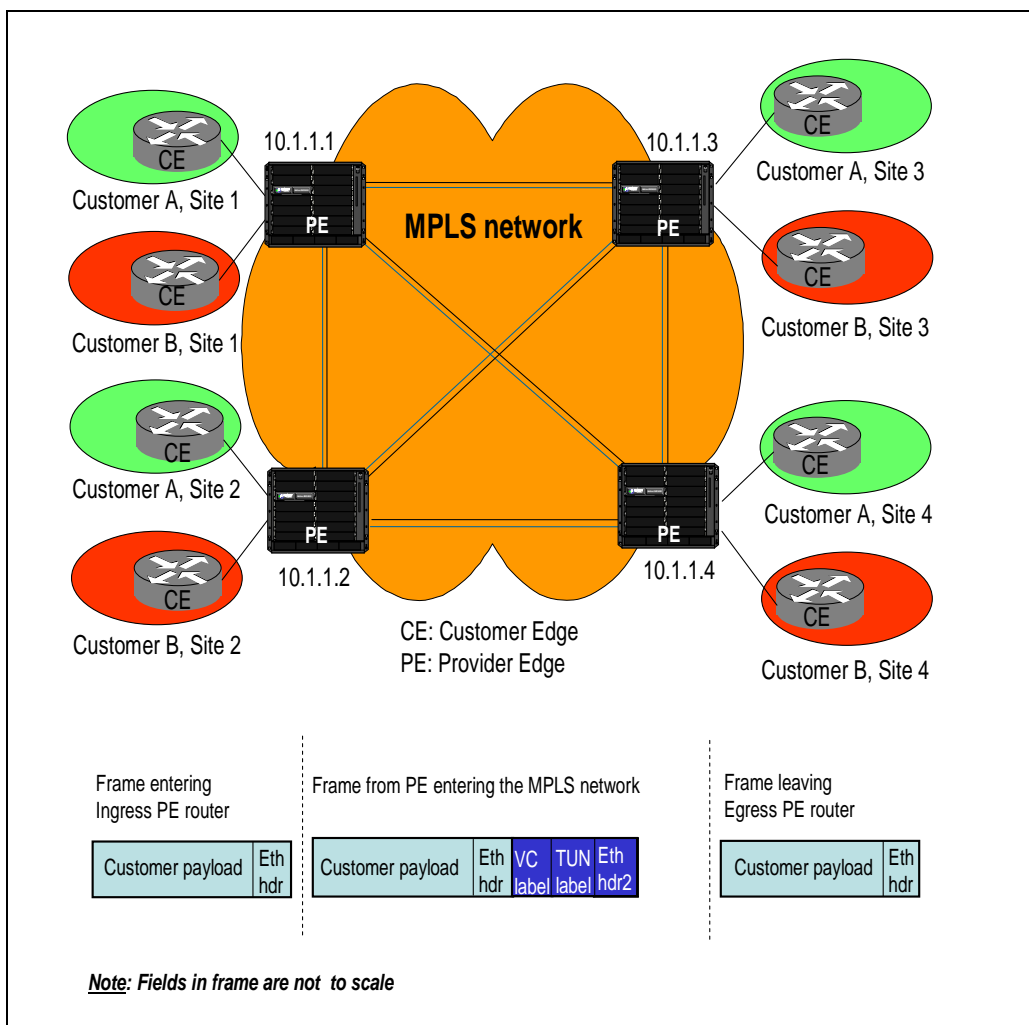


Figure 2: A VPLS network

In Figure 2, there are 2 customers, customer A and customer B that are being served on the

same VPLS network. Each customer has 4 physical sites that need to be interconnected.

Application note:
Offering scalable Layer 2 services with VPLS and VLL

Because the 2 customers are distinct, each customer belongs to a unique VPLS instance. In the figure above, there are thus 2 unique VPLS instances. Each border node in the VPLS network of Figure 2 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 premise 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) that is established between the PE routers in the VPLS network. The LSPs that form a pseudo-wire are unidirectional in nature. Therefore, a pair of LSPs is created, one in each direction between a pair of

PE routers. In the above figure, there are a total of 12 LSPs that are created. In general, for a VPLS network with n PE routers, there will be $n*(n-1)/2$ pseudo-wires or $n*(n-1)$ LSPs that are created.

Note that in Figure 2, the PE routers are shown to be directly connected for illustrative purposes. In practice, the PE routers would be connected via transit MPLS routers (also termed "P" routers).

Figure 3 below shows a logical view of the same network shown in Figure 2. The 2 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 may 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 i.e. the Virtual Circuit carrying a customer's traffic between two PEs. These 2 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.

Application note:
Offering scalable Layer 2 services with VPLS and VLL

In order to provide multi-point connectivity, the VPLS should support flooding of a frame received with an unknown destination MAC address (i.e. 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—it can be an untagged Ethernet interface, a tagged Ethernet interface or even a pseudo-wire over a non-Ethernet physical 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 first case, the tag is used to map to a VPLS instance. Its use in delimiting a service, therefore, qualifies it as a “service-delimiting” tag. In the second case, any tag in the incoming frame is transported transparently through the network. The PE router does not process the tag in the incoming frame in this case.

When service-delimiting tagged Ethernet interfaces are used, the VLAN-ID that is 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.

Scaling VPLS with a hub-and-spoke architecture

The architecture described hitherto involves setting up a full mesh between all participating PE routers. For scalability reasons, it is often advisable to limit the number of nodes that are added to this full mesh. An alternative architecture involves a hub-and-spoke architecture where “smart” MTU (Multi-Tenant Unit) switches are placed by the service provider at the edge of the network and connected to the PE router via a spoke connection. The PE thus acts as a hub device aggregating traffic from multiple MTU devices. Figure 4 gives an example of such an application.

In terms of capabilities of the MTU, the MTU is a switch that is capable of performing functions such as Q-in-Q or VLLs. 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. Alternatively, a VLL can also be deployed between the MTU and the PE router.

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 or VLL tunnels may 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 conserve network costs.

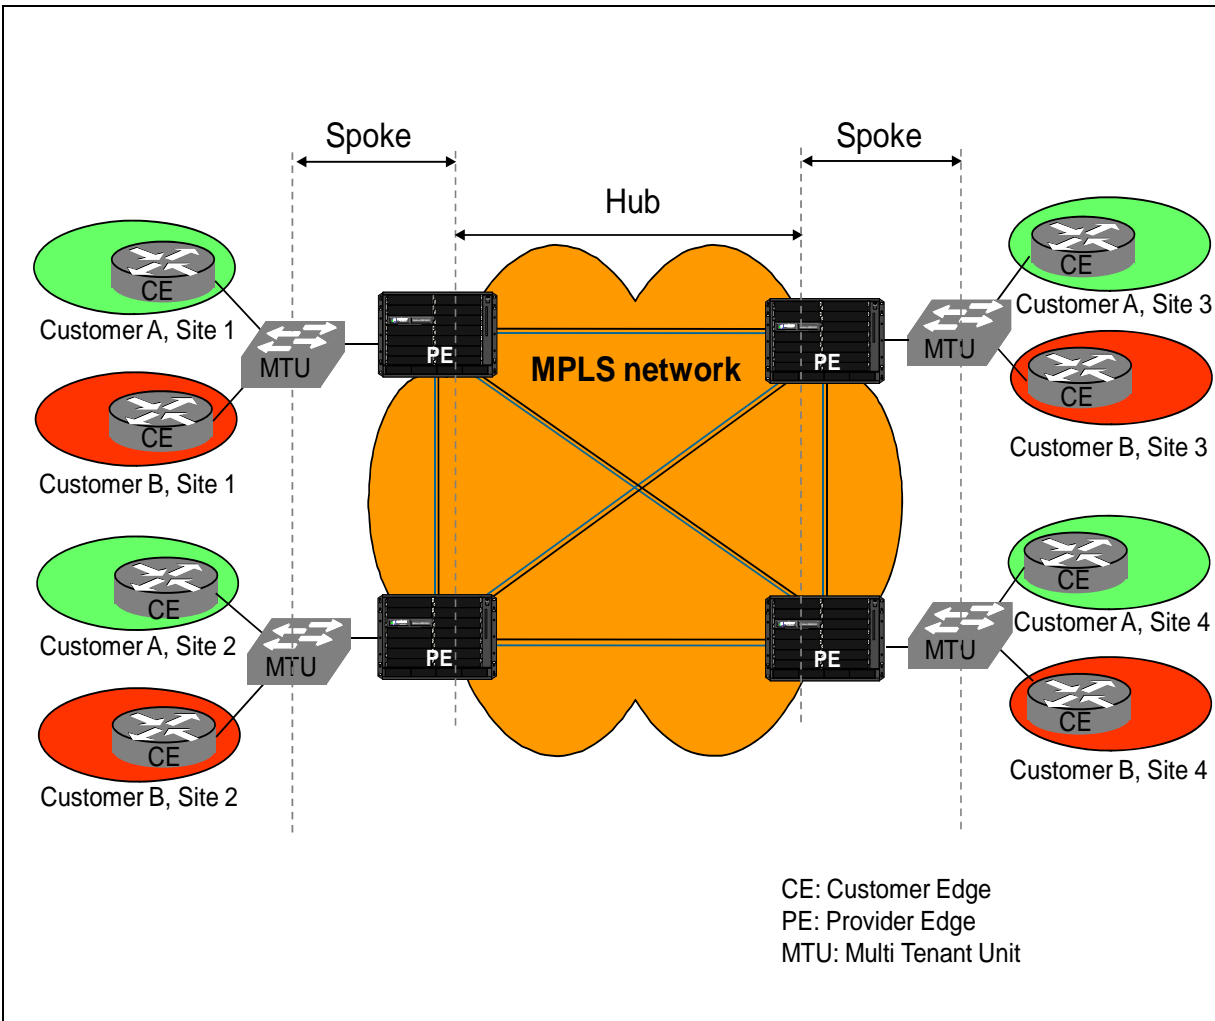


Figure 4: Hub and spoke model for scalability in the network

An alternative model of delivering scalable Carrier Ethernet services to the network edge is by combining 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
- Reduces 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

Such a combination of VPLS with PBB is depicted in Figure 5.

Application note:
Offering scalable Layer 2 services with VPLS and VLL

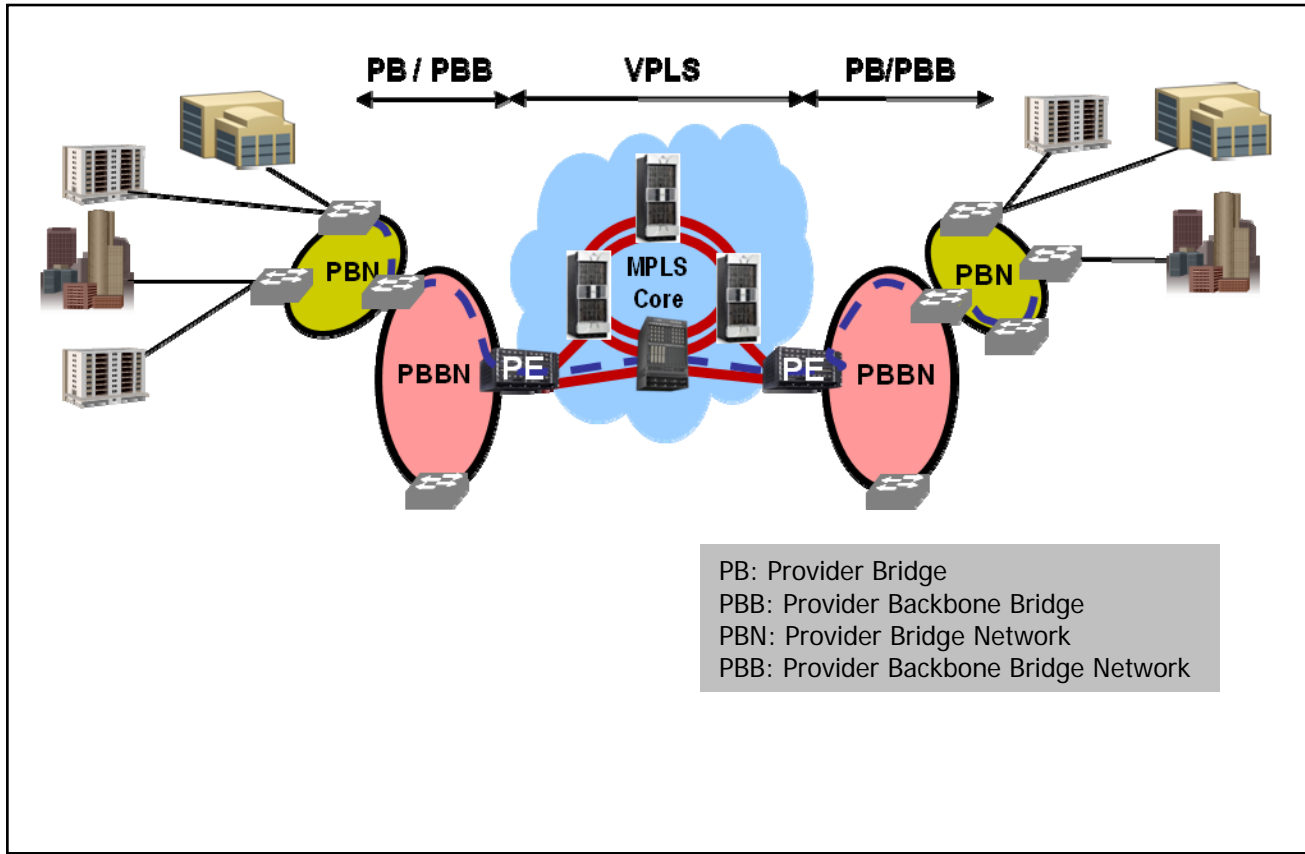


Figure 5: Connecting VPLS with PB/PBB for scalable Carrier Ethernet service delivery

The spoke model is also invaluable in creating large Layer 2 networks that span multiple provider domains. Given the geographical spread of many enterprises today, it is often necessary for large Layer 2 networks to span multiple provider

domains. By having spoke connections that interconnect the VPLS domains, such networks can easily be built. An example of this approach is shown in Figure 6.

Application note:
Offering scalable Layer 2 services with VPLS and VLL

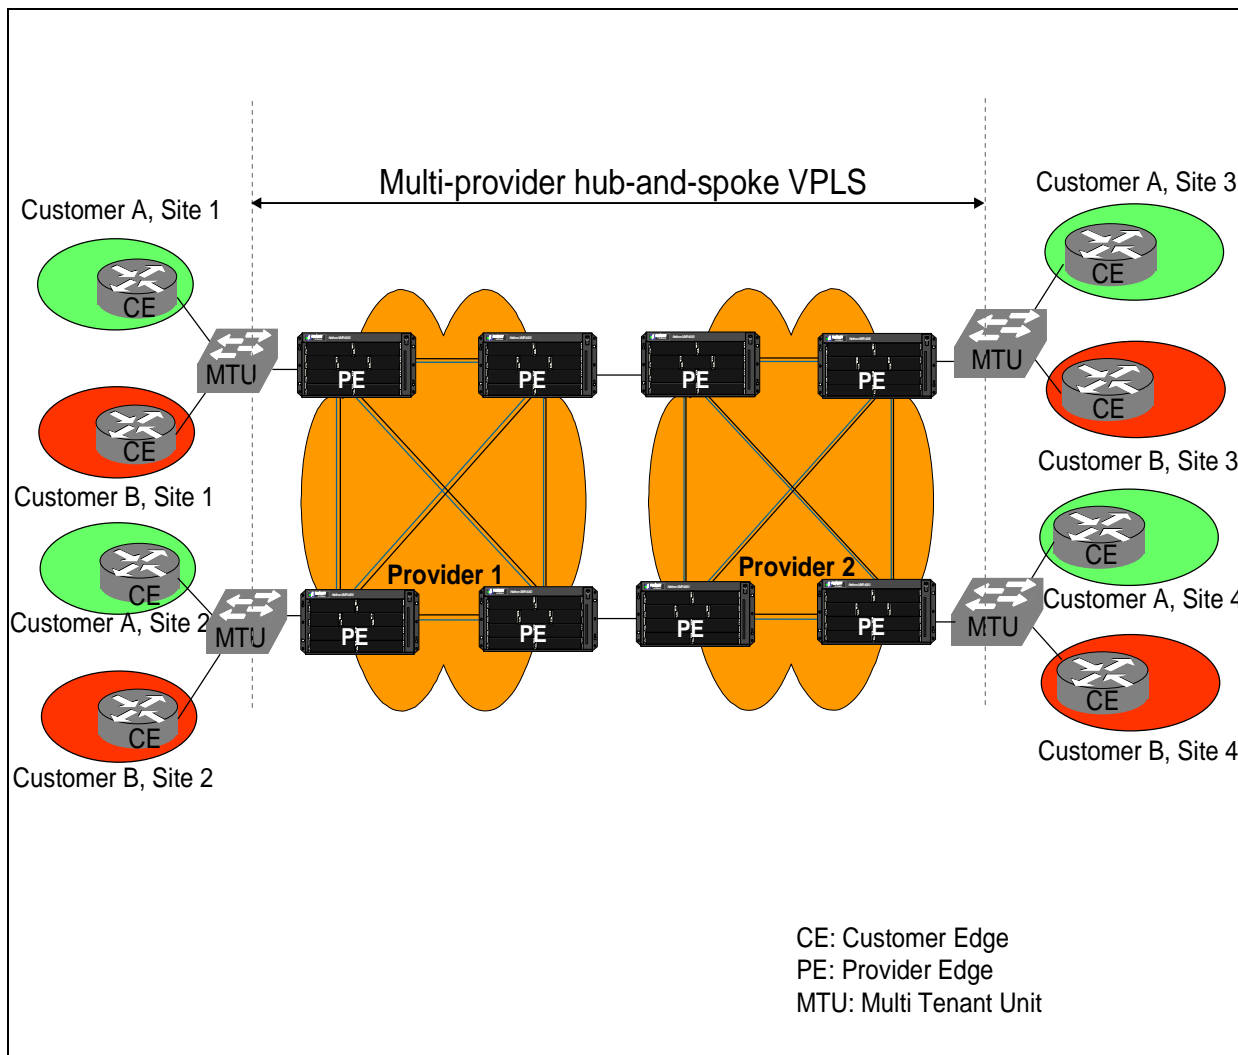


Figure 6: Building a multi-provider Layer 2 VPN service

Multi-homed spoke VPLS

An important prerequisite for offering very high service level agreements (SLA) to a provider's customers involves elimination of any single point of failure. This is particularly important at the edge where several customers may be served by the same MTU device. Foundry's switches that are used as MTUs support a capability called "protected link" where 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 the discussion in the preceding 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 in performing the arduous tasks of managing large

Application note: Offering scalable Layer 2 services with VPLS and VLL

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 lead 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. Foundry's VPLS-capable routers such as the 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, be able to have a high rate of establishing LSPs to ensure acceptable performance. Foundry's Multi-Service IronWare® operating system has a modular design with multiple internal priorities to handle various 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 handled efficiently. Given the large number of instances that are typical within a VPLS network, receipt of several frames with unknown MAC addresses should not lead to system overload. Foundry's VPLS-capable routers such as the NetIron XMR provide hardware-support for handling such flooding, which ensures that replication of such frames is done without any intervention by the CPU.

QoS control in VPLS and VLL

So far, the discussion in this paper has centered on scalability of the Layer 2 service. An area where VPLS and VLL strongly differentiate themselves from other Layer 2 VPN technologies is in the support for QoS. There are primarily two methods of controlling QoS in the network:

- Assign a 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 2 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 copying the 802.1p CoS value from the incoming tagged Ethernet frame. 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.

Within the MPLS network, the transit and 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 (a.k.a. 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 the 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, VLL with other Layer 2 technologies

Layer 2 VPNs are increasing in popularity today for several reasons. Price erosion in Ethernet and

Application note:
Offering scalable Layer 2 services with VPLS and VLL

the ubiquity of Ethernet has made it possible to rapidly slash both operational and capital expenditure costs. Technically, VPLS and VLL technologies offer many benefits over competing alternatives such as Ethernet-in-Ethernet, L2TP and Q-in-Q. Note that some of these technologies

can be used in combination (e.g. Q-in-Q in combination with VPLS or Ethernet-in-Ethernet in combination with VPLS).

Figure 7 below compares VPLS and VLL with other 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's T.E. features	Very rich: makes use of underlying MPLS network's T.E. features	Poor	Limited: can be accomplished by manipulation of 802.1p bits	Limited: can be 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
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

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

Application note:
Offering scalable Layer 2 services with VPLS and VLL

	VPLS	VLL	Ethernet PW over L2TPv3	Provider Backbone Bridge (PBB)	Provider Bridge (PB)
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 Layer 2 techniques such as MRP or RSTP	Low; requires the use of other Layer 2 techniques such as MRP or RSTP

Figure 7: Comparison of different L2 VPN technologies

VPLS compared to BGP/MPLS VPNs

A natural question that comes to mind is how VPLS fares compared to BGP/MPLS VPNs. 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 PE 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 learnt from a peer PE router are propagated by a PE to the attached CE router. Neither technology is the magic wand for all problems. Both approaches have their own merits, and which technology to use depends on several criteria.

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 — 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 that the customer chooses 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 towards transport of IP traffic only. Although IP is nearly ubiquitous, there could be niche applications that require IPX or AppleTalk or other non-IP protocols. While L3 VPNs cannot be used in such cases, Layer 2 VPNs can be used. 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 Layer 3 VPNs) or the number of MAC addresses that are maintained by the PE router (in Layer 2 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 Layer 3 VPNs. Similarly, the use of hub and spoke architectures or limiting the number of MAC addresses per customer should be used in Layer 2 VPNs when scalability is a concern.

Cost-wise, the cost of each solution varies widely by vendor. In many cases, the cost of a L3 VPN-enabled PE router is more than the cost of a L2 VPN-enabled PE router due to the requirement to maintain multiple VPN Routing/Forwarding (VRF)

Application note: Offering scalable Layer 2 services with VPLS and VLL

tables. Note that this is not true in the case of Foundry Networks' NetIron XMR or NetIron MLX series routers since the same router is cost- and performance-tuned to support both L2 and L3 VPNs.

For a detailed discussion on the topic of Layer 2 VPNs vs. Layer 3 VPNs, please refer to the application note on this subject [4].

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's L2VPN group is responsible for setting standards for these two technologies. VPLS using LDP as the signaling protocol is now standardized as RFC 4762. VLL is now standardized as RFC 4448.

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

Procedures for OAM, performing multicast within a VPLS are still under discussion in the working groups.

VPLS and VLL support on Foundry products

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

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

NetIron MLX-4, MLX-8, MLX-16 and MLX-32 routers.

- NetIron M2404, a 1 RU fixed-form factor compact MTU intended for extending MPLS deeper towards a network edge.

In addition, the 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 as described in this paper. The NetIron CES 2000 Series of compact switches supports 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.

Foundry's support for these technologies allows unparalleled scalability to be achieved by a service provider. The NetIron XMR series of routers for instance 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 particularly be viewed in the context of the industry-leading density of 1-GigE and 10-GigE ports on the NetIron XMR and NetIron MLX platforms. The low-latency, wire-speed performance for VPLS and VLL services allows very high performance to be achieved. Foundry's VPLS/VLL routers also allow 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 the management of such networks. Foundry Networks together with Oracle 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 Foundry's routers/switches. The service activation platform can also be used to manage a multi-vendor network, further simplifying the management of such a network.

Foundry Networks 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.

Application note:
Offering scalable Layer 2 services with VPLS and VLL

Accordingly, the products allow multiple services such as VPLS, VLL and BGP/MPLS VPNs to be enabled concurrently on the same port.

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 described before, 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, etc. are available; for details, please refer to the configuration guide of the appropriate product.

Configuration examples

Configuring a VPLS on a Foundry device is quite straight-forward. Figure 8 below shows the commands required for configuring the 2 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

<p>On PE1:</p> <pre>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</pre>	<p>On PE2:</p> <pre>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</pre>
--	--

Figure 8: Configuration example for VPLS

Configuring a VLL on a Foundry device is similarly very simple. Figure 9 illustrates the configuration for a VLL. These configuration commands are for the network shown in Figure 1. This example assumes that Customer A wishes to use VLAN-ID 300 on the PE-CE interface at site 1 and 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.

<p>On PE1:</p> <pre>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</pre>	<p>On PE2:</p> <pre>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</pre>
--	---

Figure 9: 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

Application note:
Offering scalable Layer 2 services with VPLS and VLL

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. Foundry Networks' solutions for these technologies allow service providers to rapidly and cost-effectively deploy such services over a converged MPLS infrastructure.

References

- [1] "Implementing Virtual Leased Lines using MPLS", Foundry Networks Application Note, http://www.foundrynet.com/solutions/appNotes/PDFs/DM_VLL.pdf
- [2] IETF's Pseudo-Wire Emulation Edge to Edge (pwe3) working group, <http://www.ietf.org/html.charters/pwe3-charter.html>
- [3] IETF's Layer 2 Virtual Private Networks (l2vpn) working group, <http://www.ietf.org/html.charters/l2vpn-charter.html>
- [4] "IP/MPLS based VPNs: Layer 3 vs. Layer 2", Foundry Networks Application Note, http://www.foundrynet.com/solutions/appNotes/PDFs/L3_vsL2.pdf

- [5] IEEE P802.1ah, "Virtual Bridged Local Area Networks: Provider Backbone Bridges", Work in progress

Author: Ananda Rajagopal
Document version 2.5

Foundry Networks, Inc.
Headquarters
4980 Great America Parkway
Santa Clara, CA 95054-1200
U.S. and Canada Toll-free: (888) TURBOLAN
Direct telephone: +1 408.207.1700
Fax: +1 408.207.1900
Email: info@foundrynet.com
Web: <http://www.foundrynet.com>

Foundry Networks, AccessIron, BigIron, EdgeIron, FastIron, IronPoint, IronView, IronWare, JetCore, NetIron, ServerIron, Terathon, TurboIron, and the "Iron" family of marks are trademarks or registered trademarks of Foundry Networks, Inc. in United States and other countries. All other trademarks are the properties of their respective owners.

Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

©2006-2008 Foundry Networks, Inc. All Rights Reserved.