

FOUNDRY
NETWORKS

IronClad Network Performance



IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

Table of Contents

1. Objective	3
2. Target Audience	3
3. Pre-Requisites.....	3
4. Introduction	3
5. MPLS Layer-3 VPNs	4
6. MPLS Layer-2 VPNs	7
6.1. Point-to-Point Connectivity	8
6.2. Multi-Point Connectivity.....	9
7. Which Way to Go: The Layer-3 or The Layer-2 Way	12
8. Summary	15

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

1. Objective

To give the reader an insight into the pros and cons of both the layer-3 and layer-2 approaches to implementing IP/MPLS-based VPNs.

2. Target Audience

Anyone working in the service provider space, or anyone interested in the MPLS technology in general.

3. Pre-Requisites

For the purpose of this paper, it is assumed that the reader is familiar with the basic concepts of MPLS label switching.

4. Introduction

MPLS has gained increasing interest from service providers over the past few years. It was originally used for traffic engineering purposes. Now, the latest application of MPLS is implementing provider provisioned VPNs. Using MPLS for implementing VPNs is a viable alternative to using a pure layer-2 solution, a pure layer-3 solution, or any of the tunneling methods commonly used for implementing VPNs.

When deciding on implementing an IP/MPLS-based VPN, the service provider has two choices:

- A layer-3 approach, commonly referred to as MPLS Layer-3 VPNs
- A layer-2 approach, commonly referred to as MPLS Layer-2 VPNs

Evaluating the merits of a given approach should be based on – but not necessarily restricted to – the following aspects of the approach:

- Type of traffic supported.
- VPN connectivity scenarios that could be offered to the customer using this approach.
- Scalability.
- Deployment complexity.
- Service provisioning complexity.
- Complexity of management and troubleshooting.
- Deployment cost.
- Management and maintenance costs.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

It can not be claimed that one approach is better than the other, since each approach attacks the problem from a different angle. Hence, what might be the best choice for a given provider does not necessarily have to be the best choice for another.

5. MPLS Layer-3 VPNs

The layer-3 approach to creating MPLS-based VPNs offers a routed solution to the problem. The de facto standard for implementing such VPNs is described in “RFC 2547”, with a new version, currently, under development referred to as 2547bis which is described in “draft-ietf-ppvpn-rfc2547bis-01.txt”. The approach is also referred to as BGP/MPLS VPNs.

The approach relies on taking customer IP datagrams from a given site, looking up the destination IP address of the datagram in a forwarding table, then sending that datagram to its destination across the provider’s network using an LSP.

In order for the service provider routers to acquire reachability information about a given customer’s networks, the provider edge (PE) routers exchange routes with the customer edge (CE) routers. Hence, the BGP/MPLS VPNs approach follows the peer to peer model of VPNs. These routes are propagated to other PE routers carrying the same VPN(s) via BGP. However, they are never shared with the provider’s core routers (P), since the PEs use LSPs to forward packets from one PE to the other. P routers do not need to know about the customer’s networks in order to perform their label switching functions. A PE router receiving routes of a given VPN site from another PE, propagates the routes to the CE router of the connected site belonging to that same VPN, so that the CE will also learn about the networks in the remote site.

The mechanisms behind BGP/MPLS VPNs were designed to address some of the shortcomings of the pure layer-3 VPNs (without tunneling) that preceded it. Some of the main goals were:

- Supporting globally unique IP addresses on the customer side, as well as private non-unique – and hence, overlapping – addresses.
- Supporting overlapping VPNs, where one site could belong to more than one VPN.

Since this type of VPNs relies on routing, achieving the abovementioned goals could be a challenge. To address the problem of overlapping address spaces in customer VPNs, multiple routing and forwarding tables, referred to as VPN Routing and Forwarding (VRF) tables, are created on each PE router, in order to separate the routes belonging to different VPNs on a PE router.

A VRF table is created for each site connected to the PE, however, if there were multiple sites belonging to the same VPN connected to the same PE, these sites might share a

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

single VRF table on that PE. A site that is a member of multiple VPNs is not a candidate for VRF table sharing with other sites that are not members of exactly the same set of VPNs. Such a site must have its own VRF table, which includes routes from all the VPNs it is a member of.

Another implication of the overlapping address spaces problem is that a PE router receiving BGP updates from its neighbors might receive conflicting or overlapping routes – belonging to different VPNs. In order to identify the advertised routes as belonging to different VPNs, and hence, prevent the BGP process from selecting one – the best – and ignoring the rest, an 8 octet Route Distinguisher (RD) is prepended to each prefix advertised. This is used to distinguish routes belonging to different VPNs on the BGP receiver side. The result of prepending the RD to the 4 octet IP prefix is a 12 octet address for which a new special address family was defined, the VPN-IPv4 family. Hence, to be precise, multi-protocol BGP is used to carry such prefixes.

Route Distinguishers provide nothing more than a way of differentiating routes. They play no role in controlling route distribution. An RD is assigned to a VRF, so that prefixes advertised from that VRF will have that RD prepended to them. Typically, it makes sense to assign the same RD to the VRFs of sites belonging to the same VPN, so that all the routes of that VPN will have the same distinguisher. So, it could be said that RDs are typically assigned uniquely to each VPN. However, this should not mean that VRFs of sites that belong to multiple VPNs get multiple RDs. VRFs of such sites need only one RD. For those sites, as well as those that are members of only one VPN, controlling the distribution of routers is performed as described below.

To prevent a PE router from accepting routes of VPNs that it doesn't carry, and hence, waste its own resources, BGP extended communities are put to use in order to control the distribution of routes within the provider's network. The extended community attribute Route Target is included with the advertised route(s) to indicate which VPN – or the group of sites in certain topologies – the route belongs too. A unique value for this attribute is assigned to each customer VPN. A PE router keeps track of those Route Target values associated with the VPNs that it carries. Upon receipt of an advertised route, the BGP process checks the Route Target to see if it is equal to the Route Target value of one of the VPNs that it carries. In case of a match, the route is accepted, if not, the route is ignored. This is to avoid having all the PE routers carrying all the routes of all the customer VPNs, which might severely limit the scalability of the solution.

Figure 1 illustrates the main concepts behind the BGP/MPLS VPNs approach.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

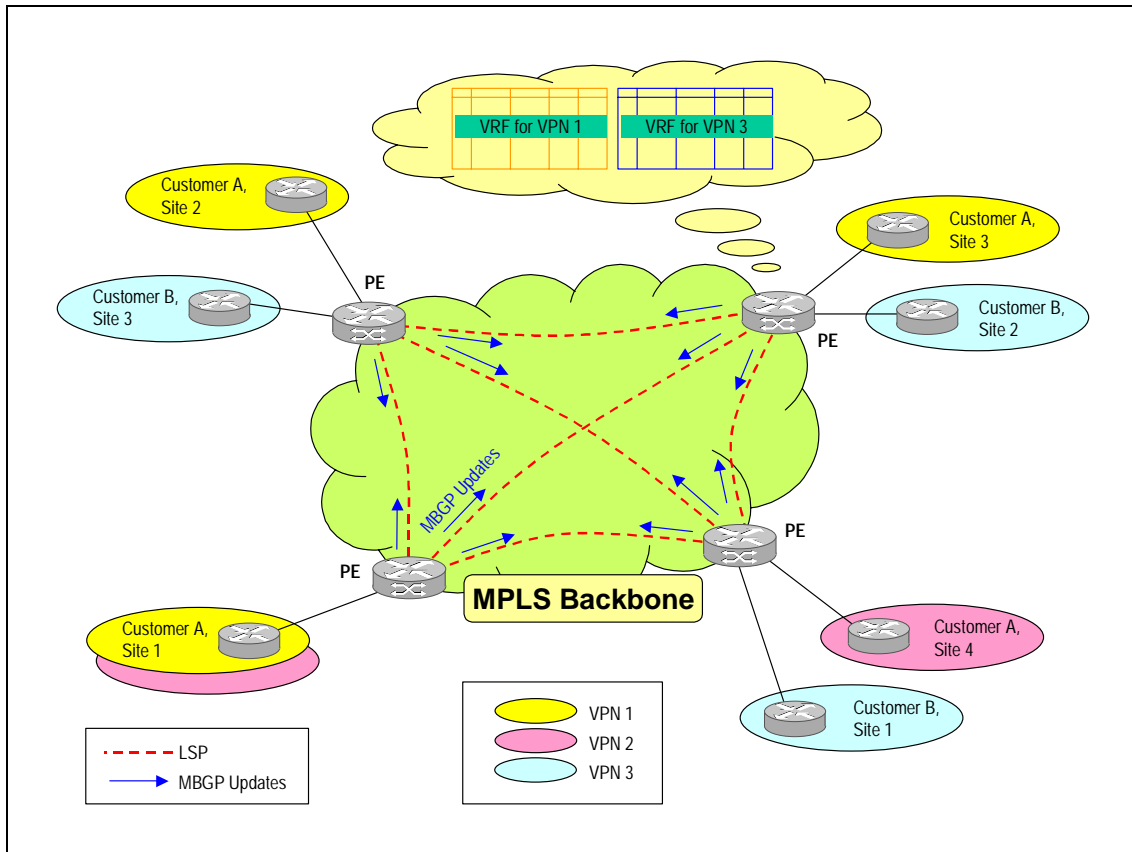


Figure 1 The BGP/MPLS VPN approach.

From the discussion above, it could be seen that the approach allows for creating overlapping VPNs. This is intended for scenarios like when a customer needs a VPN for their intranet, and another for their extranet with a different set of routes advertised in each to control the accessibility of resources. Such a customer would rely on the service provider to perform the required route control, i.e., route control is shifted from the CE router and delegated to the PE router. In Figure 1, Customer A, Site 1, lies in both VPN 1 and VPN 2. The routes of that site are advertised by the connected PE router with one RD, however, with two Route Target extended community attributes: one for VPN 1, the other for VPN 2. The connected PE router, also, accepts routes from the other PE routers, only if the routes have Route Target values equal to that value of either VPN 1 or VPN 2 – since these are the only VPNs carried by this router in this example.

When advertising a VPN-IPv4 route, the PE also includes an MPLS label – representing the route – in the BGP message, and it sets the BGP NEXT_HOP equal to its own address. The provider network is MPLS enabled, and each PE router should be capable of reaching any of the other PEs via an LSP. Those LSPs could be created by any protocol like LDP or RSVP/TE.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

When a PE receives a packet with a destination in a remote site, it attaches two MPLS labels to the packet in order to forward it to its destination. The outer label is for the LSP leading to the BGP NEXT_HOP. The inner label is the label associated with that destination, learned previously from a BGP update received from a peer. The PE, then, sends the frame out the port associated with that LSP. The frame gets label switched all the way to the remote PE, which then, pops the outer label, and examines the inner label. The inner label, in most cases, uniquely identifies the destination, therefore, it is popped and the packet is forwarded to its destination. In some cases, where route summarization is done on the PE, the receiving PE uses the inner label to determine which VRF to look into in order to know where to send the packet.

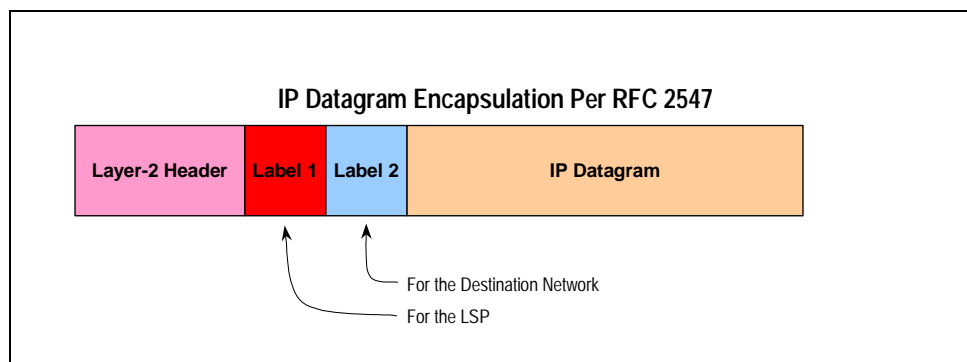


Figure 2 Two labels are attached to an IP datagram to be forwarded to its destination.

6. MPLS Layer-2 VPNs

The layer-2 approach is the newer approach to implementing MPLS-based VPNs, and it offers a layer-2 switched solution. The layer-2 approach provides complete separation between the provider's network and the customer's network, i.e., there is no route exchange between the PE devices and the CE devices. Hence, the approach follows the overlay model of VPNs.

The separation between the provider's network and the customer's networks provides simplicity. MPLS layer-2 VPNs provide emulated services capable of carrying customer layer-2 frames from one site to the other. This is done in a manner that is totally transparent to the CE devices. Handling customer layer-2 frames allows the service provider to offer a service that is independent of the layer-3 protocols in use by the customers, i.e., the provider would be able to carry IPv4, IPv6, IPX, DECNet, OSI, etc.

The layer-2 approach addresses two connectivity problems:

- Providing Point-to-Point connectivity
- Providing Multi-Point Connectivity

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

6.1. Point-to-Point Connectivity

The de facto standard for establishing point-to-point connectivity in MPLS layer-2 VPNs is described in the Martini drafts:

- “draft-martini-l2circuit-trans-mpls-08.txt”
- “draft-martini-l2circuit-encap-mpls-04.txt”

In order to carry layer-2 frames across an MPLS cloud, the Martini drafts introduce the concept of Virtual Circuits (VCs). An LSP acts as a tunnel carrying multiple VCs, whereas a VC acts like the actual circuit carrying customer layer-2 frames.

A VC, actually, is just another LSP within the original tunnel LSP. The tunnel LSP provides the tunnel between two PE routers, while the VC carries frames of a given customer only. VCs are uni-directional just like normal LSPs. Hence, for bi-directional communication, a pair of VCs – one in each direction – is need.

In order to create this hierarchy, an encapsulated customer frame traversing the service provider network has two labels attached to it:

- A label pertaining to the tunnel LSP leading to a destination PE. This is called the “tunnel label”.
- A label pertaining to the VC that carries the frame and leads to a certain site attached to the destination PE. This is called the “VC label”.

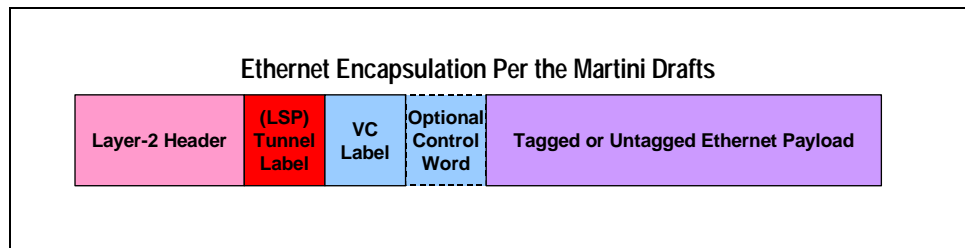


Figure 3 A Martini encapsulated Ethernet frame gets two labels attached to it.

Tunnel LSPs between the PE routers could be created using any protocol like RSVP/TE or LDP. PE routers exchange the VC labels via LDP in downstream unsolicited mode.

At the edge of the provider network, the PE router encapsulates the subscriber layer-2 frame as per the Martini drafts, attaches a VC label and a tunnel label, then sends the frame over the tunnel LSP.

At the other end of the tunnel LSP, the receiving PE router pops the tunnel label, determines which customer port the packet should go to based on the VC label, extracts the original layer-2 frame, and sends it out the port determined above.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

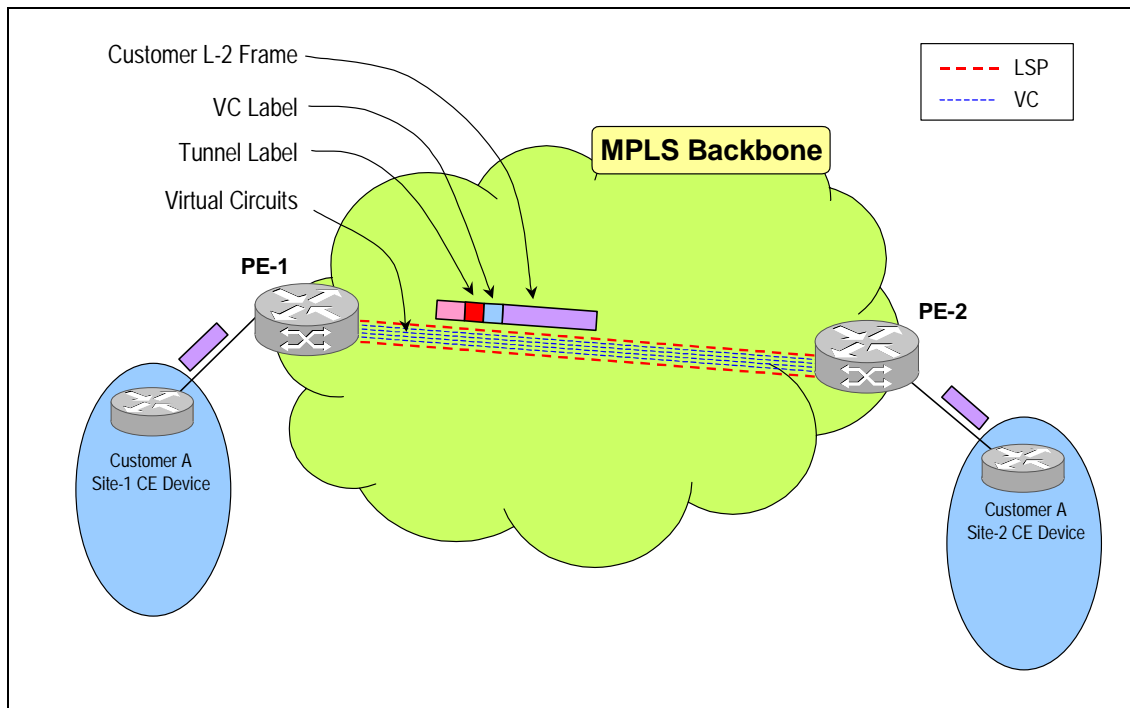


Figure 4 A tunnel LSP carries multiple VCs, a VC carries a given customer’s traffic.

Using this approach, a service provider could offer a service that resembles leased lines or Frame Relay PVCs, while using cheaper building blocks in the infrastructure: IP, PoS, Ethernet, etc.

6.2. Multi-Point Connectivity

Currently, there are several proposals within the IETF that address the problem of multiple site connectivity at layer-2. The goal here is a solution that facilitates carrying customer layer-2 frames – specifically, Ethernet – over the service provider’s IP/MPLS network from and to multiple sites that belong to a given VPN (customer). For efficient use of the provider’s network bandwidth, a frame should be sent only to the PE that connects to the target site of the frame whenever possible, instead of being flooded. This is accomplished by switching the customer frames based on their destination MAC address. The end result is a simple service that emulates connecting the sites constituting the VPN via a layer-2 switch.

The popular approach to implementing such a solution is called Virtual Private LAN Services (VPLS). The core of the technology is described in “draft-lasserre-vkompella-ppvpn-vpls-00.txt”, with several enhancements described in other drafts.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

The VPLS approach expands on the concepts introduced by the Martini drafts that were used for establishing point-to-point connectivity. It builds the VPN by creating a full mesh of VCs between the PEs facing the sites that make the VPN. Note that VCs are uni-directional, therefore, between any pair of PEs there should be a pair of VCs to carry bi-directional traffic. VPLS as described in the aforementioned draft relies on LDP for the exchange of VC labels between the PE routers. However, other methods of signaling could be used, and are described in other drafts.

Customer VPNs are identified via a unique VPN ID, currently, a 32 bit value. Several proposals exist for expanding that ID to a 56 or a 64 bit value. Another proposal was made to use simple descriptive text strings as VPN IDs that can be stored in the DNS system to ease provisioning.

Note that, even though this is a layer-2 service, VLAN IDs play no role in identifying the VPN a customer frame belongs to. Only the labels attached to the encapsulated customer frames have significance, since the provider's routers are just label switching the frames. Hence, the 4095 VLAN limitation of the VLAN ID, doesn't cause a similar limitation on the number of VPNs that could be supported by the VPLS approach.

PE routers perform source MAC address learning just like a normal transparent switch, except that they perform it on frames received over the VCs. For instance, if PE1 receives a frame with a source MAC X over the VC M, it creates an entry in its layer-2 forwarding table (MAC table) that associates MAC X with VC N, which is the other VC in the opposite direction of M. When PE1 receives a frame from an attached customer site with a destination MAC X, it looks the MAC up in its layer-2 forwarding table, and finds the associated VC N. Hence, it encapsulates the frame as per the Martini drafts, and sends it over VC N to its destination. Should PE1 receive a frame from an attached customer site with a destination MAC Y that has no entry in its layer-2 forwarding table, then it simply floods it over all the VCs belonging to that customer's VPN, i.e., it floods to the other sites of the VPN. Of course, like in normal layer-2 switching, once a response from the remote system arrives at PE1 in the form of a packet with source MAC Y, PE1 will create a forwarding entry for it, and all subsequent packets targeting MAC Y will not be flooded.

A PE router maintains a separate layer-2 forwarding table, called Virtual Forwarding Instance (VFI), for each VPN that it carries. Figure 5 illustrates the basic concepts behind the VPLS approach.

Note that, due to the learning scheme mentioned above, a PE router does not learn all the MAC addresses in all the VPNs carried by the provider network. A PE router learns MAC addresses related only to the VPNs that it carries. P routers do not learn any MAC addresses, they just perform label switching.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

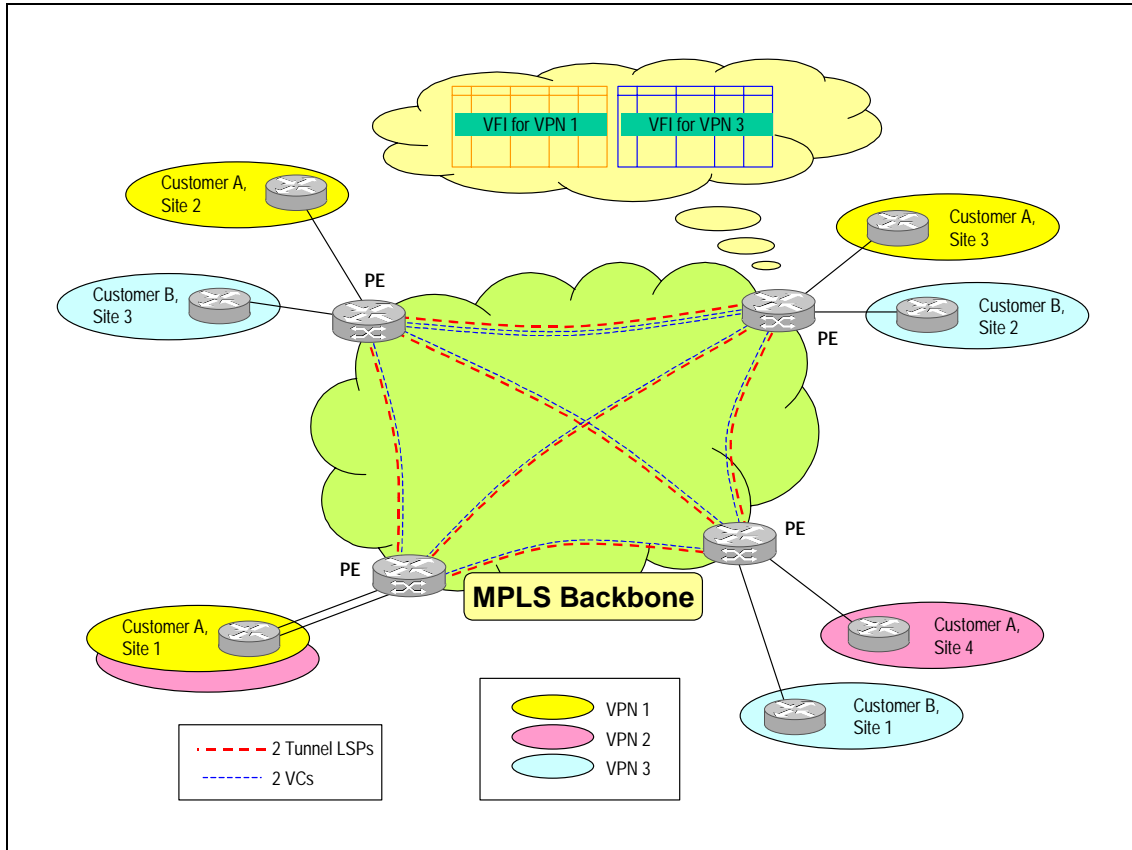


Figure 5 The VPLS approach.

Unlike normal layer-2 switches, PE routers do not run STP within the provider's network in order to implement fault tolerance and loop avoidance. Since VPLS is based on MPLS, it leverages MPLS' traffic protection abilities in order to implement a fault tolerant service. Also, since VPLS relies on a full mesh of VCs for a given VPN, i.e., each PE could reach any other PE within a VPN in exactly one hop without any transit PEs in between, the VPLS PEs apply a simple split horizon forwarding rule when forwarding customer frames:

If a customer frame is received over a VC within a VPN, that frame could only be forwarded to an attached customer site, not back to the same VPN (over another VC). This simple rule together with the full mesh topology of VCs addresses the issue of loop avoidance without using STP. Avoiding the use of STP allows the PE routers to avoid STP scalability issues commonly encountered in pure layer-2 networks. The intention here is to make VPLS more scalable.

As seen in Figure 5, overlapping VPNs could be implemented using VPLS. Customer A, Site 1 lies in both VPN 1 and VPN 2. To separate traffic belonging to each VPN, the customer site could be connected to the PE router using two access links, one for each VPN. Alternatively, traffic belonging to both VPNs could be multiplexed over the same access link using two different VLAN IDs, where one VLAN ID maps to VPN 1, the

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

other ID maps to VPN 2. The use of more than one 802.1Q tag within a frame helps the service provider and the customer use the required service tag (VLAN ID) without having any impact on the customer's choice of their own VLAN IDs.

In contrast with the layer-3 approach, the task of controlling the routes that get advertised in each VPN remains the customer's responsibility, since the PE router does not handle any customer routes.

7. Which Way to Go: The Layer-3 or The Layer-2 Way

From the discussion above, the reader could realize that each approach has its strengths and its weaknesses. A wise choice of an approach to adopt would consider those strengths and weaknesses, in addition to the current and future requirements of the service to be implemented, the existing infrastructure, and the costs involved.

Type of Traffic Supported

Comparing both approaches described above, it is clear that the layer-3 approach offers transport of IP traffic only. On the other hand, the layer-2 approach allows transporting any customer layer-3 protocol packets: IPv4, IPv6, IPX, DECNet, OSI, etc. Many enterprise customers still use other protocols than IP in their IT infrastructure, hence, a layer-2 service is less restricting for them. Also, with IPv6 on the horizon, some organizations are already experimenting with IPv6, and in the near future, many will be migrating to it. To continue providing connectivity for those organizations using a layer-3 solution would require some enhancement to the current standard – like creating a VPN-IPv6 address family – and might require some upgrades to the provider's routers. A layer-2 solution could continue to serve those organizations, even when the provider network has not yet been upgraded to use IPv6 internally.

Possible Connectivity Scenarios

Several connectivity scenarios for customer sites could be implemented using both approaches. Both approaches could be used to implement the following connectivity scenarios:

1. Point-to-Point.
2. Hub and Spoke.
3. Partial Mesh.
4. Full Mesh.
5. Overlapping VPNs.

The layer-3 approach performs well at implementing scenarios 1, 4, and 5 in a manner that is transparent to the CE devices. However, the layer-3 approach could get a bit more complicated when implementing scenarios 2 and 3.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

The layer-2 approach performs well at implementing scenarios 1, 2, 3, and 4. It is worth noting that when implementing scenarios 2 and 3, it is more straight forward to build the topology using VCs as in the layer-2 approach, than to build the topology by controlling BGP routes as in the layer-3 approach. Scenario 5 is also possible using the layer-2 approach, however, it requires some involvement from the CE device at the site where the overlap occurs: the CE device would have to control which routes get advertised in which VPN, i.e., it is not as transparent as in the layer-3 approach.

Scalability

When considering the scalability of a layer-3 solution vs. a layer-2 solution, one could find some similarities. A limiting factor for both solutions would be the maximum number of LSPs and/or VCs that could be supported by a given LSR.

Another limiting factor that is common to both is the maximum configuration file size that could be stored, specifically, on a PE router. This is due to the fact that the configuration file contains all the information related to the customers' VPNs. For a layer-3 solution, the configuration file contains definitions for the VRFs, RDs, extended communities, and route filtering policies. For a layer-2 solution, the configuration file contains definitions for the VPN peer PEs, and the ports associated with the customer VPNs. The use of auto-discovery in conjunction with a layer-2 solution obviates the explicit configuration of the VPN peer PEs, and hence, decreases the impact of the maximum configuration file size on the scalability of the solution.

For a layer-3 solution, the maximum number of routes that could be stored on a given PE is also a constraint. This is due to the fact that a PE router stores routes from all the VPNs that it carries. To alleviate the impact of this factor on the scalability of the solution, route summarization could be used whenever possible. For a layer-2 solution, the maximum number of layer-2 forwarding table entries supported on a PE routes is also a constraint. The PE router has to create those entries in order to be able to perform its layer-2 switching functionality. The impact of this factor on scalability could be alleviated by requiring that CE devices be routers, and/or applying limits to the number of (MAC) entries created for each VPN – to avoid having a customer VPN overwhelm the PE routers with a large number of source MAC addresses.

Deployment

Deployment of a layer-3 solution usually requires high end LSRs capable of handling multiple routing and forwarding tables at the provider edge. It also requires that BGP peering be set up between the these routers. If the service provider is already using BGP so extensively throughout there network, as in the case of ISPs or large IP carriers, then they might prefer going with a layer-3 solution since it allows them to take advantage of the already available BGP sessions, and the already available BGP know how. Then, of

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

course, LSPs between the PEs have to be set up for carrying traffic between the PEs. When leveraging the existing BGP peering session, however, some changes to route reflection clusters might be required, so that no route reflector would be overwhelmed by too many routes from too many VPNs. Should the provider be using a confederation, then the problem becomes similar to the inter-provider (inter-AS) problem, where the VPNs have to span multiple autonomous systems. Also, similar to the route reflection case, the provider needs to carefully consider what could be done in order to avoid having the routers connecting the member-ASes overwhelmed by too many routes.

A layer-2 solution, typically, would require simpler PE routers, and without the requirement of having BGP peering sessions set up between the PEs. For service providers who don't rely on BGP or are unwilling to deploy BGP for the new VPN service to avoid the complexity, the layer-2 solution might be more attractive. Use of BGP for VPN signaling between the PEs remains as an option¹ to the provider, in case they already have BGP deployed and would like to take advantage of it. As in the layer-3 case, LSPs between the PEs have to be set up for carrying traffic from one PE to the other.

Service Provisioning

For a layer-3 solution, service provisioning would clearly require designing routing for the specific VPN topology requested by the customer. This means designing the VRFs that are going to contain the customer routes, and deciding on how RDs and Route Target communities are going to be assigned. Note that the service provider has to decide whether a VRF should be shared by multiple customer facing interfaces, or whether a VRF should collect routes from multiple VPNs as in the case of overlapping VPNs. Also, RDs and Route Target communities need to be allocated for the VPN(s) to be provisioned. Then, the PE routers connected to the customer sites that make the VPN need to be configured for the required VRFs, RDs, Route Targets, and any additional options that might be needed for certain topologies. Peering between the PE routers and the customer's CE routers needs to be set up in order to allow the route exchange required for the operation of layer-3 VPNs.

Provisioning a VPN using a layer-2 solution is simpler, and more straight forward. Each PE router carrying the VPN needs to know the other PEs to establish VCs with in order to form the desired VPN. Then the PE ports connected to the customer sites are mapped to the VPN. Note that the use of auto-discovery eliminates the need to explicitly configure peer PEs that carry the same VPN. Currently, there are several ideas within the IETF for performing auto-discovery. When standardized, service provisioning using a layer-2 solution would be even simpler.

¹ Other variants of the VPLS approach described here utilize BGP for VPN signaling.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

Management and Maintenance

When managing a layer-3 solution, doing configuration changes, or troubleshooting problems, the service provider engineers would mainly be dealing with BGP peering sessions, BGP routes with different extended communities, their propagation, and selection by the PE, peering with customer CE routers, etc. As in many large scale IP networks, route reflection clusters or a confederation with multiple member-ASes might be in use which could contribute to the complexity of the task at hand. Also, dealing with a large number of routes belonging to multiple routing and forwarding table in addition to the global table is certainly more demanding than dealing with a single table. Finally, configuration files on the PE routers could grow so large which makes it harder to spot a misconfigured statement.

A layer-2 solution is simpler since the provider does not retain any customer routes, control their distribution, or peer with any customer CE routers. Also, since BGP is not required, management and troubleshooting become even simpler – unless the provider is using BGP for VPN signaling as in some variants of the VPLS approach described above. When performing management or troubleshooting, the service provider engineers deal with the simpler concepts of the VCs making the VPN, and the ports assigned to the VPN. On a given PE, the engineers deal with only one routing table while the VFI tables get dynamically populated via source MAC address learning. As in the layer-3 case, when the configuration file grows so large it becomes more challenging to recognize misconfigurations. As mentioned before, the use of auto-discovery will help keep the size of the configuration file to an absolute minimum.

Costs

Comparing deployment costs, it is more likely that a layer-3 solution would cost slightly more than a layer-2 solution, due to the fact that the layer-3 approach relies on more sophisticated routers capable of handling multiple VRFs.

Management and maintenance costs of a given solution are directly related to the complexity that solution. A layer-3 solution is more likely to cost more due to its higher complexity. The complexity of the solution demands a certain level of technical know-how, and might translate into more man hours required to accomplish any task related to the solution.

8. Summary

Currently, there are two main approaches to implementing IP/MPLS-based VPNs:

- The layer-3 approach with the de facto standard BGP/MPLS VPNs.
- The layer-2 approach with the de facto standards defined in the Martini drafts for point-to-point connectivity, and in the VPLS drafts for multi-point connectivity.

IP/MPLS-Based VPNs

Layer-3 vs. Layer-2

The layer-3 approach offers an IP only, routed solution which is capable of supporting multiple VPN topologies by leveraging the advanced route distribution control capabilities of BGP.

The layer-2 approach is the newer approach to the problem. It offers a layer-2 switched solution for transporting customer layer-2 frames which makes it independent of the layer-3 protocol in use by the customer. It is, also, capable of supporting multiple VPN topologies through the use of virtual circuits (VCs).

A wise choice of an approach to adopt would consider the strengths and weaknesses of each approach, in addition to the current and future requirements of the service to be implemented, the existing infrastructure, and the costs involved.

Ahmed Abdelhalim
Product Marketing Engineer
Service Provider Group
Foundry Networks, Inc.

Headquarters
2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
U.S. and Canada Toll-free: 1 (888) TURBOLAN
Direct Telephone: +1 408 586-1700
Fax: +1 408 586-1900
Web: <http://www.foundrynet.com>

© 2002 Foundry Networks, Inc. All Rights Reserved.