

**FOUNDRY**  
**NETWORKS**

*IronClad Network Performance*



## **Implementing Virtual Leased Lines Using MPLS**

# Implementing Virtual Leased Lines Using MPLS



## Table of Contents

1. Objective .....	3
2. Target Audience .....	3
3. Pre-Requisites.....	3
4. Introduction: MPLS and IP-Based VPNs .....	3
5. The Promise of MPLS Layer-2 VPNs.....	5
6. Tunneling Layer-2 Frames from One Point to Another.....	6
7. Providing an Emulated Leased Line .....	9
8. Provisioning the Service Using Foundry Routers .....	9
9. Service Flexibility .....	10
10. Improving Network Bandwidth Utilization .....	11
11. Providing Quality of Service (QoS).....	12
12. Creating Multiple Service Offerings .....	13
13. VLL Traffic Protection for High Service Availability.....	15
14. Sample Application Scenarios .....	16
14.1. Offering a Leased Line Service.....	16
14.2. Offering Multiple Service Levels.....	18
14.3. Providing Inter-MAN Connectivity .....	20
14.4. Replacing DSL/ATM or DSL/Frame Relay Infrastructures .....	22
14.5. Offering Transparent LAN Services (TLS) .....	24
15. Looking Forward.....	26

# Implementing Virtual Leased Lines Using MPLS

---

## 1. Objective

The objective of this paper is to illustrate the use of MPLS Layer-2 VPNs, specifically, the point to point type, in order to implement a Virtual Leased Line service capable of carrying Ethernet frames across an MPLS cloud.

## 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 and IP-Based VPNs

The virtual private network (VPN) as a concept simply means delivering private network services over a public infrastructure; in other words, creating the illusion of being connected to a private network even though the packets are being carried over a public infrastructure.

VPNs have been around for quite some time. X.25 or Frame Relay, for instance, have been used for a long time, and were used by service providers in order to provide subscribers with such an illusion of being connected through a private network, even though the packets from different customers were actually carried by a common public infrastructure managed by the service provider. Hence, it could be said that X.25 or Frame Relay provided VPN services.

Over the past few years there has been an increasing interest in IP-based VPNs. Typically, solutions for IP-based VPNs were deployed on the side of ISP subscribers, like enterprise customers, in order to tunnel their traffic from one site to the other across the public Internet. The availability of the Internet almost everywhere made IP-based VPNs that leverage the existing Internet infrastructure an attractive solution to enterprises that addresses some of the enterprise connectivity needs. Typical applications were the implementation of extranets, and telecommuting. Several technologies – both proprietary and standard – like PPTP, IPoverIP, L2TP, IPSec, etc. evolved to address the market needs.

# Implementing Virtual Leased Lines Using MPLS

---

There has also been increasing interest from service providers in IP-based VPNs. As a means of driving revenue, service providers are always looking for new services that they could offer their subscribers. Hence, solutions that offer site to site connectivity or multiple site connectivity are always in demand by service providers.

To distinguish VPNs implemented by subscribers – using their own edge devices – from those VPNs implemented and managed by service providers, the latter type is called Provider Provisioned Virtual Private Networks (PPVPNs).

When implementing a PPVPN service, the service provider's goals might include some or all of the following:

- Delivering private network services to their subscribers over the provider's public infrastructure.
- Being able to support global as well as private – non-unique – IP addresses on the subscriber side.
- Avoiding large protocol overheads.
- Being able to prioritize frames in order to deliver QoS.
- Being able to offer subscribers more than one service level to choose from.

In order to have investment protection, and to avoid being locked in one vendor's environment, a service providers would prefer to go with standards based technologies. IPSec, for instance, being a standards based technology, seems like an attractive technology to use for the implementation of PPVPNs. However, IPSec in this context presents too much additional overhead and complexity. This is due to the fact that IPSec not only provides connectivity through tunneling traffic across an IP cloud, but also confidentiality (customer data encryption). Indeed, it is the confidentiality portion of the technology that presents a management challenge for service providers. The management of encryption keys, and interaction with the public key infrastructure (PKI) might complicate service provisioning and management for a service provider, which translates into a higher operational cost. Also, the encapsulation overhead of IPSec is relatively high, which translates into wasted bandwidth for the service provider, if confidentiality is not an essential requirement of the service. Furthermore, IPSec has no built-in mechanisms for traffic engineering or providing QoS, it simply inherits the capabilities of IP.

Another example of a candidate technology would be Generic Routing Encapsulation (GRE), currently being standardized<sup>1</sup>, specifically, the use of GRE over IP<sup>2</sup>. However, like IPSec, the encapsulation overhead is high, and it does not offer its own mechanisms

---

<sup>1</sup> Described in RFC 2784, with extensions in RFC 2890.

<sup>2</sup> GRE was designed as a generic protocol that could be run over many protocols, called delivery protocols.

# Implementing Virtual Leased Lines Using MPLS



for traffic engineering<sup>1</sup> or QoS, it simply relies on the capabilities of the underlying delivery protocol – IP in this context.

A question that might pop up in the reader's mind is, why MPLS for VPNs? The answer is quite simple, MPLS seems like an attractive technology for the following reasons:

- The MPLS Label Switched Paths (LSPs) inherently provide tunneling of traffic from one point to the other.
- Since MPLS switches packets based on their labels, it inherently masks the IP address, and hence, could be used to isolate the IP addresses on the subscriber side from those on the service provider side. In other words, it wouldn't matter, then, whether the subscriber is using global or private IP addresses; MPLS is capable of supporting both.
- The overhead of MPLS encapsulation is small when compared to other encapsulation technologies. MPLS labels are only four octets long.
- The EXP bits in the MPLS shim header could be used to prioritize MPLS frames – a feature that is available in most MPLS implementations today.
- Traffic engineered Label Switched Paths (LSPs) could be deployed in order to offer multiple service levels to subscribers, or to avoid network congestion points.
- The MPLS approach allows for creating highly scalable VPNs.

## 5. The Promise of MPLS Layer-2 VPNs

MPLS Layer-2 VPNs offer a service provider the capability of carrying subscriber layer-2 frames from one site to another (belonging to the same subscriber) in a manner that is totally transparent to the subscriber edge devices.

The advantage of such a technology is that it allows the provider to offer a transport service that is independent of the layer-3 protocol used by the subscriber. Offering a layer-2 service means that customer layer-2 frames regardless of their payload, IPv4, IPv6, IPX, NetBEUI, etc., could take advantage of the service, which makes the service more usable and more attractive to customers.

Another advantage to the service provider is that by offering the service at layer-2, the provider does not have to keep and manage any layer-3 reachability information on their routers pertaining to the individual customer sites that they serve. This would not be the case, if a provider attempts to implement a layer-3 VPN service. The decreased operational complexity of the layer-2 approach translates into less operational cost in the end.

---

<sup>1</sup> The original informational RFCs 1701 and 1702 specified a basic source routing mechanism that was not included in the new standards track RFCs. Yet, source routing does not offer the level of intuitiveness and flexibility required for traffic engineering applications, and could translate into a performance hit since the source route has to be checked and updated for each GRE packet by each router the packet arrives at.

# Implementing Virtual Leased Lines Using MPLS

The MPLS layer-2 VPNs technology addresses, primarily, two problems:

- Providing point to point connectivity between two sites, in a manner similar to connecting both sites with a leased line or a Frame Relay Virtual Circuit.
- Providing multi-point connectivity between multiple site, in a manner that resembles connecting the customer edge (CE) devices at those sites to a LAN segment, in other words, a switch.

The mechanisms needed to address the first problem are the most mature ones, and indeed, they are the focus of this paper. Mechanisms for addressing the second problem are still under development and standardization by the IETF, and shall be covered in a future white paper.

## 6. Tunneling Layer-2 Frames from One Point to Another

As mentioned above, the goal here is to allow service providers to offer a service that connects two subscriber sites at layer-2, in a manner that resembles the connection via a leased line.

The de facto standard for accomplishing this is described in the Martini drafts, written by Luca Martini *et al.* LSPs are still used under the Martini approach. However, 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 subscriber layer-2 frames.




Figure 1 Relationship between physical links, tunnel LSPs, and VCs.

# Implementing Virtual Leased Lines Using MPLS

Indeed, a VC is nothing more than another LSP within the original transport tunnel LSP. The introduction of VCs leverages an MPLS capability known as “MPLS Label Stacking” that allows an MPLS frame to carry multiple MPLS labels. Therefore, looking at a Martini encapsulated frame as it traverses an MPLS cloud, it would have two MPLS labels:

- A label pertaining to the LSP tunnel carrying the frame. This label is referred to as the “Tunnel Label”. Of course, that label gets swapped at each intermediate node (transit LSR) in the MPLS domain.
- A label pertaining to the virtual circuit carrying the frame. This label is referred to as the “VC Label”. This label gets attached at the ingress point, and remains unchanged till the frame reaches the egress point.

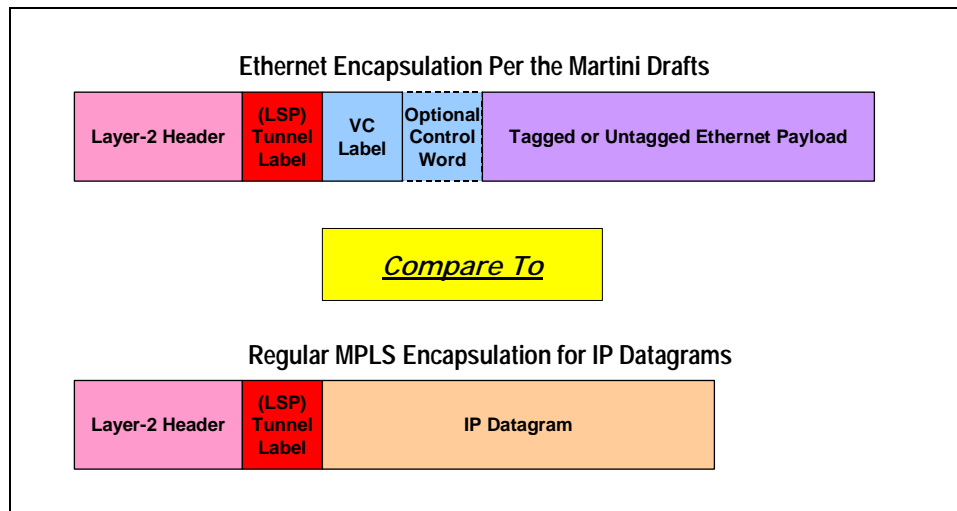


Figure 2 Martini encapsulated Ethernet frame compared with normal MPLS encapsulation for IP

At the edge of the MPLS cloud, a provider edge (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.

# Implementing Virtual Leased Lines Using MPLS




Figure 3 Sending layer-2 frames over a Virtual Circuit

The Martini drafts are:

- “draft-martini-l2circuit-trans-mpls-08.txt”
  - Describes the role of Tunnel LSPs and the new concept of Virtual Circuits.
  - Specifies the exchange of VC labels via LDP in downstream unsolicited mode.
  - Introduces a new VC FEC Element, to be used within LDP Label Mapping messages for the exchange of VC labels.
- “draft-martini-l2circuit-encap-mpls-04.txt”
  - Defines encapsulations for layer-2 frames of the following types: Ethernet, Ethernet tagged frames, PPP, HDLC, ATM, and Frame Relay. These are the frame types that could be transported over the VC.

Transporting Ethernet frames, either tagged or untagged, is the application that has gained large interest from network equipment vendors. This could be attributed to the fact that traffic at many customer sites is predominantly Ethernet. Also, Ethernet is gaining increasing momentum in the metropolitan area networks space each day, which makes the transport of Ethernet frames as per the Martini drafts a viable solution for carrying customer traffic from one Metro zone to another.

It is worth noting that the IETF Pseudo Wire Emulation Edge to Edge (PWE3) working group adopted the Martini drafts and will continue expanding on them. The PWE3 working group regards the Martini approach as one of the paths they will pursue in order to provide service providers with frameworks for implementing emulated services – like emulated Frame Relay, ATM, etc. – over a common IP infrastructure.

# Implementing Virtual Leased Lines Using MPLS

## 7. Providing an Emulated Leased Line

Based on the discussion above, one could realize that the VC, being an LSP within the tunnel LSP, is unidirectional just like any normal LSP. Therefore, in order to provide bi-directional communication for a subscriber, two VCs, one in each direction, are needed.

In order to ease the provisioning and management of the service, Foundry uses the concept of the Virtual Leased Line (VLL) for all the configuration tasks on its devices. The VLL is a more abstract entity that represents two VCs serving a certain subscriber. The VCs are established between two provider edge devices as one VC in the forward direction, another in the backward direction.

Hence, in operations, the administrator applies a configuration to a VLL, and the Foundry devices automatically establish/reconfigure the two VCs – associated with the VLL – as necessary.




Figure 4 A Virtual Leased Line represents two VCs, one in each direction.

## 8. Provisioning the Service Using Foundry Routers

The Foundry routers allow for the creation of VLLs capable of transporting subscriber Ethernet frames as per the Martini drafts. The implementation is quite mature, and has been tested successfully for interoperability with many major vendors. For more information on interoperability, please refer to Foundry's "Statement of Multivendor Interoperability" white paper.

Basically, the configuration of a VLL requires the following three pieces of information:

- A unique ID for the VLL, called the VLL ID. This ID is assigned by the service provider.
- The IP addresses of the PE routers that will establish the VLL.
- The port designators of the ports facing the customer (and, optionally, VLAN IDs) and acting as the end points of the virtual leased line.

# Implementing Virtual Leased Lines Using MPLS

---

To provision the service:

- The VLL is defined on both PE routers with the assigned VLL ID.
- On each PE router, the address of the other PE router – called VLL Peer – is configured.
- On each PE router the customer facing port (and, optionally, VLAN ID) is configured. This is the port from which traffic is accepted and forwarded over the VLL, and to which traffic coming over the VLL is forwarded.

The PE routers automatically select the appropriate tunnel LSPs to carry the subscriber VCs. However, the administrator could force the selection of a certain tunnel LSP, or an LSP out of a group of LSPs using techniques described later in this paper.

The PE routers, then, use LDP to automatically negotiate the VC labels to be used when forwarding traffic over the VCs.

## 9. Service Flexibility

Foundry's NetIron series routers offer a high degree of flexibility for configuring the service. They allow any of the following customer frames to be transported over the VLL:

- Untagged customer frames from a certain customer port
- Tagged customer frames with a certain VLAN ID from a certain customer port
- Customer frames from a certain customer port, regardless of their status: tagged or untagged.

The last option is very useful for service providers who do not want to enforce a service tag, or to place any restrictions on the type of frames (tagged/untagged) to be transported.

Taking a close look at the way tagged frames are handled above, one could see that VLAN IDs now have link scope only, not box scope, nor network scope. This allows the service provider to reuse the VLAN IDs on different ports on the same router, and hence, in essence, break the barrier of 4096 VLANs commonly encountered in Ethernet-based networks.

When transporting tagged frames, the NetIron routers allow the service provider to have non-matching VLAN IDs at both ends of the VLL. The VLL is capable of converting the VLAN ID to the appropriate value at the egress port to the customer.

Taking this a bit further, one could have tagged frames going into one side of the VLL and emerging as untagged on the other side, and vice versa. The NetIron routers add or discard the 802.1Q tag according to the configuration of the customer facing side.

This decoupling of both sides of the VLL allows the service provider to easily provision the service while avoiding the reassignment of VLAN IDs to customers, or changing the

# Implementing Virtual Leased Lines Using MPLS

service terms with the customer so that tagged frames would be exchanged instead of untagged or vice versa.

For service providers who use stacked 802.1Q tags in a certain zone but not in the other, the Foundry routers are capable of taking single tagged frames on one side of the VLL and submitting them with an additional tag on the other side, and vice versa. Hence, the VLL service implemented via Foundry routers is capable of performing Super VLAN Aggregation<sup>1</sup> in a manner similar to the Foundry layer-2 switches. This leads to better integration of functions, and eliminates the need for additional layer-2 switches to perform the function of 802.1Q tag stacking.

As for tunnel LSPs, the NetIron series routers allow the use of either LDP created LSPs, or RSVP/TE created LSPs – for scenarios where traffic engineering is needed – so that service providers might choose the combination that best suits their application.

## 10. Improving Network Bandwidth Utilization

Network bandwidth is usually a costly resource, and not an easily attainable one. Through continuous monitoring of the network for performance, a service provider might observe certain spots of their network experiencing congestion, commonly known as congestion points, while other links in the network remain under-utilized.

Increasing bandwidth at the congestion points might be a solution, but a costly one. Smarter use of bandwidth using traffic engineering might be a more viable solution.

Using MPLS traffic engineering, traffic could be rerouted around network congestion points, making use of the less utilized/under utilized links in a service provider's network. Hence, avoiding the cost of acquiring and deploying new faster line cards, or acquiring additional fiber.

In the context of MPLS layer-2 VPNs and VLLs, should the service provider observe the emergence of such a congestion point that impacts on the VLLs of some customers, the provider could easily reroute one or more tunnel LSPs using traffic engineering to avoid the congestion point. Rerouting a tunnel LSP means that all traffic from all VCs going over that LSP will be shifted to the new path, hence avoiding the congestion point. This could be regarded as one of the benefits of the Martini approach to transporting layer-2 frames, where traffic is carried by VCs that get assigned to tunnel LSPs in a many to one relationship.

---

<sup>1</sup> A Foundry term referring to the ability to support a hierarchy of VLANs via stacking the 802.1Q tags.

# Implementing Virtual Leased Lines Using MPLS

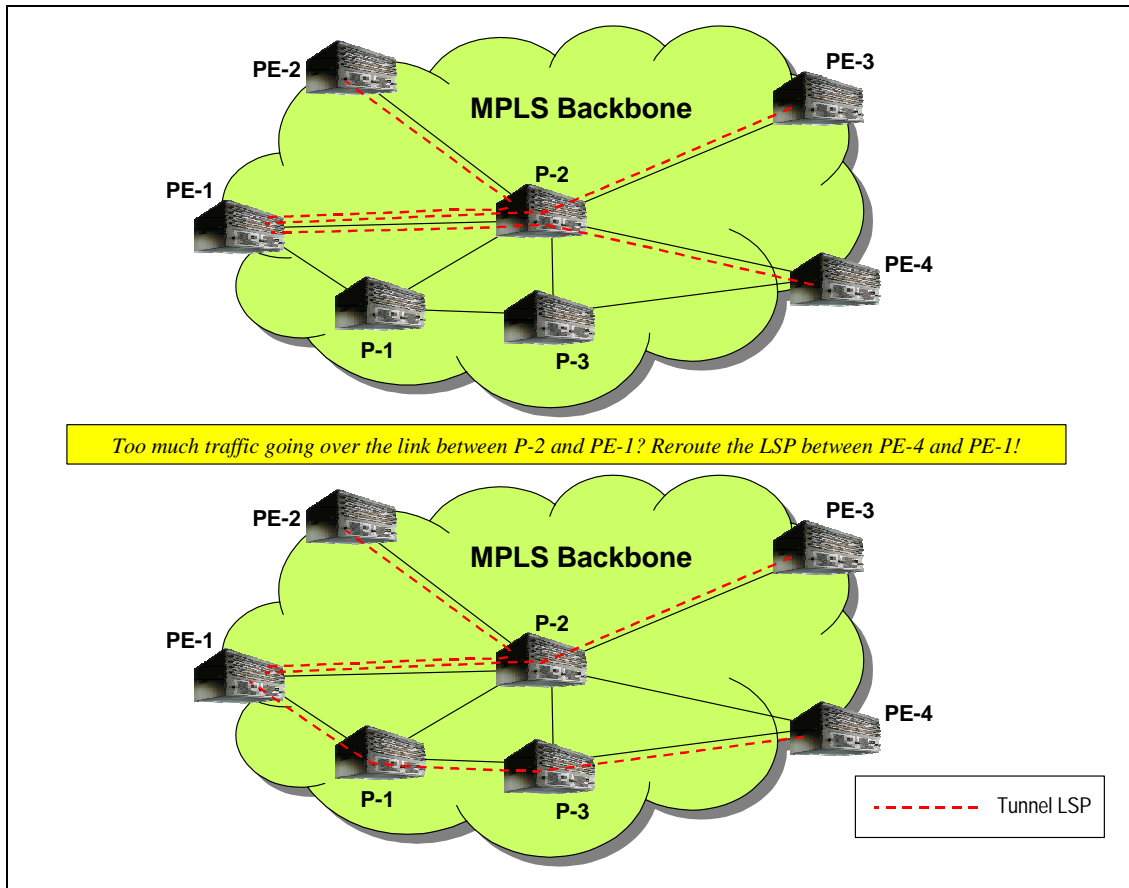


Figure 5 Shifting traffic to under-utilized links in the network via traffic engineering

## 11. Providing Quality of Service (QoS)

For a service provider to provide QoS, they must implement and run an infrastructure that is capable of prioritizing customer frames. When such an infrastructure is available, higher priorities could be given to the more critical flows, or to the more demanding customers with higher service level agreements (SLA).

In order to give higher priority to the more critical customer flows, a mechanism is needed for identifying those flows and prioritizing them. For a service provider carrying customer Ethernet frames across the provider's network the 802.1Q priority bits seem like a good candidate mechanism. The NetIron series routers are capable of prioritizing customer frames based on the 802.1Q priority bits, end to end.

When a customer frame is received by a NetIron router acting as a PE device, the frame is encapsulated and the 802.1Q priority bits are copied into the EXP bits of both the tunnel label and the VC label to ensure that the resulting MPLS frame gets prioritized

# Implementing Virtual Leased Lines Using MPLS

appropriately all the way along the path to its destination. The frame is then assigned to one of four priority queues (per port) based on its EXP bits:

- Queue 3: for priority 6 or 7
- Queue 2: for priority 4 or 5
- Queue 1: for priority 2 or 3
- Queue 0: for priority 0 or 1

Foundry's architecture serves these queues either in weighted fair queuing (WFQ) mode, or in strict priority mode.

NetIron routers acting as transit LSRs are capable of prioritizing MPLS frames based on the value contained in the EXP bits of the tunnel label. The value of the EXP bits is preserved when doing label swapping.

A NetIron PE router acting as an egress LSR for a Martini encapsulated frame, also, prioritizes the transmission of the decapsulated customer frame out the customer facing port using the EXP bits of the VC label. This is by design to cover the case where penultimate hop popping (PHP) is used, and hence, the tunnel label is discarded before the MPLS frame reaches the egress LSR.

Prioritization in all cases is accomplished by assigning the packet to the appropriate queue out of the four queues (per port) mentioned above, and based on one of the algorithms mentioned above – user configurable.

Alternatively, in order to enforce a certain priority level on a certain customer's traffic, the NetIron series routers allow for the manual assignment of a value to the EXP bits of tunnel label as well as the VC label. These values will be inserted in the EXP field of the tunnel label and the VC label, regardless of whatever 802.1Q priority values the customer is sending in their Ethernet frames.

To offer the service provider additional flexibility, a NetIron edge LSR allows the administrator to assign priorities to the customer facing ports. This priority value gets copied into the EXP bits of both the tunnel label and the VC label, and hence, prioritization occurs end to end as mentioned above. This is intended for scenarios where a subscriber is sending untagged Ethernet frames – and hence, the frames have no priority field – and the provider prefers not to assign an explicit priority value for the VLL traffic, but to have the traffic inherit the priority value from that priority configured on the customer facing port.

## 12. Creating Multiple Service Offerings

In order to increase their revenues, service providers need to be able to address the needs of a larger customer base. A larger customer base encompasses customers with slightly different needs. Some customers might need high data transfer rate because they might be

# Implementing Virtual Leased Lines Using MPLS

---



transferring large volumes of data per day. Others might need average data transfer rates but with low end to end delay and low jitter, because they might be running delay sensitive applications like voice applications, video conferencing, etc.

A service provider needs to offer a service that is appealing to all those customers. Hence, the concept of the “one package fits all” might not be practical for a service provider. Service providers need to be able to deliver the same service in more than one package, i.e., they need to be able to support differentiated services for traffic from different customers.

To offer multiple service packages that guarantee to the customer certain ranges for the data transfer rate or the end to end delay could be rather challenging for a service provider. To accomplish this, control over the paths taken by the customers’ traffic is needed in order to have some degree of predictability in the service. Unless their infrastructure is ATM or Frame Relay based, maintaining that level of predictability with layer-3 IP switching, or layer-2 Ethernet switching would always be a challenge. However, ATM and Frame Relay are expensive technologies, and meanwhile, service providers are always considering more cost efficient alternatives like IP over PoS or Ethernet.

MPLS is a viable solution for service providers looking forward to implementing new services – like the VLL service – using cost efficient components: IP, PoS, Ethernet, while maintaining a level of control similar to what they had with ATM or Frame Relay. In addition, unlike ATM or Frame Relay, MPLS as a technology is hardware independent, i.e., it was designed to work over Ethernet, PoS (PPP, HDLC), ATM, and Frame Relay. Therefore, using MPLS, a service provider is able to maintain control across their new PoS/Ethernet clouds, as well as their legacy ATM and Frame Relay clouds.

In our case, offering a virtual leased line service, a service provider could use MPLS traffic engineering to control the paths taken by customer traffic, and hence, implement a more predictable service. For instance, if a service provider needs to offer the VLL service in three different flavors – three different packages –, each one designed with a certain type of customer in mind. Instead of creating one pair of tunnel LSPs – one LSP in each direction – between each pair of edge LSRs receiving customer traffic, the provider could create three, each optimized to meet the requirement of a certain package offered.

When a customer subscribes to a certain package, all that needs to be done is assign that customer’s VLL to the tunnel LSP pair of that package. Should the customer upgrade to a higher package, their VLL just needs to be moved to the higher package tunnel LSP pair. The converse is true.

The MPLS implementation on the NetIron routers was designed with these scenarios in mind. As mentioned before, a NetIron edge LSR is capable of automatically selecting an

# Implementing Virtual Leased Lines Using MPLS

appropriate tunnel LSP to carry a configured VC. However, the selection of the tunnel LSP could be influenced by the administrator in order to guarantee that the VC will use a certain tunnel LSP or one of a group of LSPs

A NetIron edge LSR allows the administrator to assign a CoS value for a configured LSP as well as a CoS value for a configured VLL. In this case, when the software of an edge LSR searches for a suitable LSP to carry the outbound VC, it looks for an LSP with a CoS value that is equal to the CoS value of the VC. If none is found, the device selects the next best LSP, i.e., LSP with highest CoS among the group of LSPs with CoS less than that of the VC<sup>1</sup>. By configuring the appropriate values for the LSP CoS and VC CoS, customer VCs could be assigned or reassigned to the suitable tunnel LSPs, according to their packages.

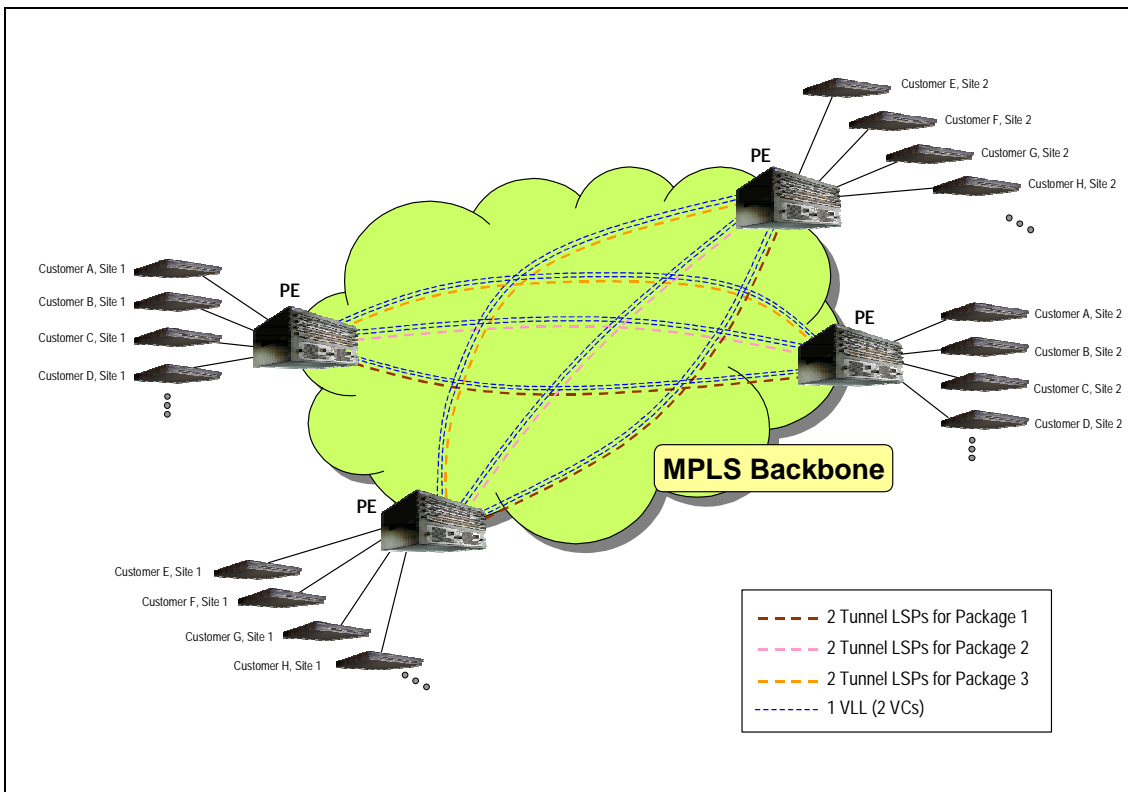


Figure 6 Offering Virtual Leased Lines in 3 different flavors (packages).

## 13. VLL Traffic Protection for High Service Availability

In order to provide higher VLL service availability, the NetIron series routers offer two types of protection for the network architect to choose from:

<sup>1</sup> Intended for a fault tolerance mechanism to be described later.

# Implementing Virtual Leased Lines Using MPLS

---

- Protection at the VLL level
- Protection at the tunnel LSP level

Protection at the VLL level could be used in any scenario where multiple LSPs exist between two edge LSRs bearing the VLL service. In case of a failure, and in case CoS values are configured for the LSPs and the VLLs, the edge LSR searches for another tunnel LSP leading to the VLL peer with a CoS value equal to that of the VLL, and assigns the outbound VC to it. If such a tunnel LSP is not found, the device falls back to the next best LSP: the LSP with the highest CoS among the group of LSPs whose CoS is less than that of the VLL. The idea here is to prevent the VLL from falling back to a tunnel LSP that is used by subscriber's with a more demanding service level agreement.

As a reoptimization mechanism, the NetIron edge LSR periodically checks the available tunnel LSPs to see if there is any room for optimizing their use by VLLs: reassigning some VLLs to tunnel LSPs to achieve better load distribution, or to assign a VLL to a tunnel LSP with a better matching CoS. Hence, in the case where the VLL falls back to a tunnel LSP with a less CoS, then the original tunnel LSP with the matching CoS comes up again, the reoptimization algorithm will automatically reassign the VLL to that matching tunnel LSP without any user intervention.

A simpler scenario would be when a failure occurs, and the CoS values were not configured. In this case, the NetIron router just finds an alternate tunnel LSP and assigns the outbound VC to it.

Alternatively, protection at the LSP level could be used. In this case, the network architect relies on one LSP only that has been designed to be fault tolerant. In other words, traffic protection mechanisms of the LSP are put to use here. For instance, in a traffic engineering scenario using Foundry routers, the network architect could design the LSPs to have primary paths and secondary paths. The LSPs would be established over the primary path. In case the primary path fails, the LSP falls back to the secondary path, i.e., the ingress LSR starts signaling for the secondary path. Should faster convergence be required for certain applications, the NetIron provides the option of placing the secondary path in hot-standby mode. In this case, the secondary path is pre-established, so that when the primary path fails, traffic is switched immediately to the secondary path without having the delay caused by path setup.

## 14. Sample Application Scenarios

### 14.1. Offering a Leased Line Service


A service provider might leverage an existing IP infrastructure in order to offer point to point layer-2 connectivity to their customers, i.e. a Virtual Leased Line service. It is believed that by adopting IP, PoS, and Gigabit Ethernet in the service provider network,

# Implementing Virtual Leased Lines Using MPLS

this service could be offered at very generous data transfer rates, and at very competitive prices, which would make it a quite attractive service offering to the customer.

This service could be used by a customer – e.g., enterprise customers – to connect their sites as an alternative to the more expensive TDM leased lines or frame relay PVCs, that are typically used to build an enterprise IT infrastructure.

For the customer premises equipment, be it a router or a switch, all that is needed is a free Ethernet port on the customer device to be hooked up, since the service emulates a direct Ethernet link between the customer devices. This simplicity on the customer edge side, makes the service even more attractive, since it allows the customer's IT staff to avoid the complexities of dealing with Frame Relay or TDM line problems.



**Figure 7 Service provider's view of the VLL service.**

# Implementing Virtual Leased Lines Using MPLS

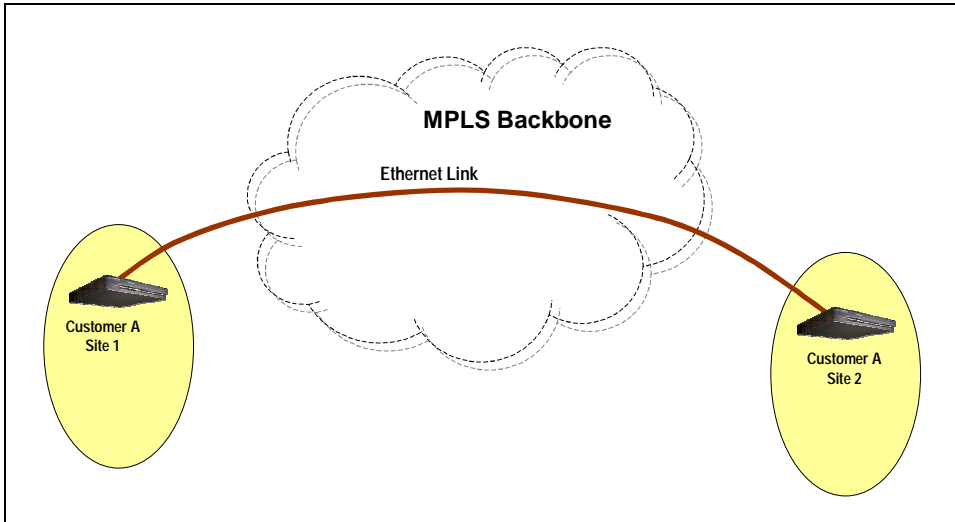


Figure 8 Customer's view of the service.

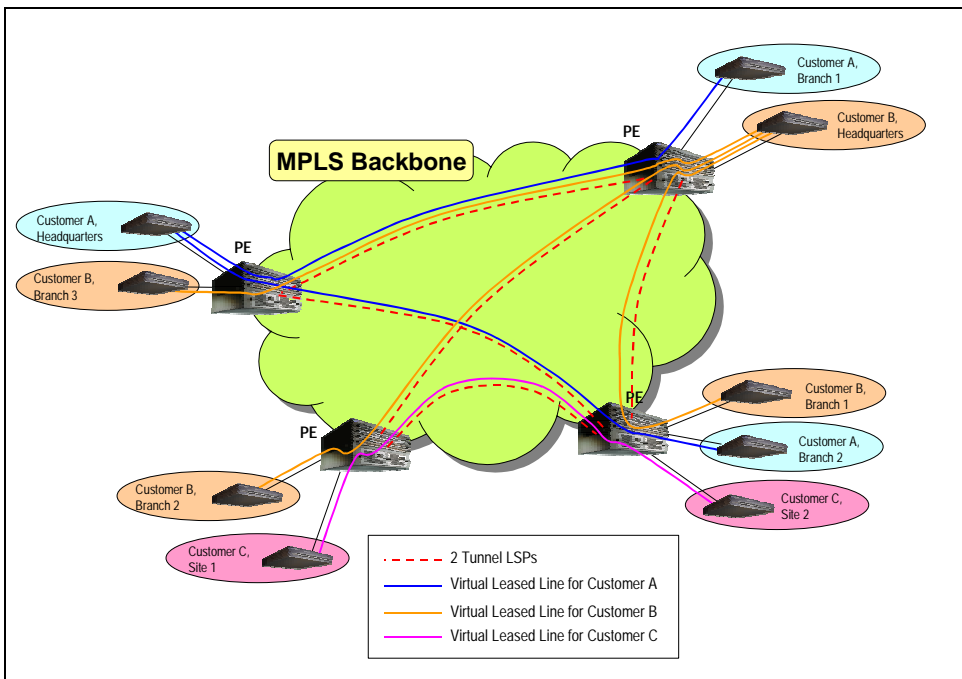


Figure 9 Customers use VLLs in a manner similar to TDM lines or Frame Relay PVCs.

## 14.2. Offering Multiple Service Levels

As mentioned before, the VLL service could leverage MPLS traffic engineering capabilities in order to allow the provider to offer multiple service levels.

# Implementing Virtual Leased Lines Using MPLS



More than one service package could be created with different characteristics – data transfer rate, end to end latency, availability, etc. – and pricing to address the needs of a wider base of customers.

When a customer subscribes to a package, all that needs to be done is to create the customer’s VLL and make sure it gets assigned to the appropriate tunnel LSP pair – the one specifically created to serve subscribers of that package.

Furthermore, a customer would be able to request more than one VLL between any two customer sites. For instance, a customer might request one VLL with a high data transfer rate for transferring data, and another VLL with a moderate data transfer rate and a low end to end latency for video conferencing or VoIP applications. The traffic of both VLLs could be multiplexed over the same access line – between the customer edge (CE) device and the provider edge (PE) device – by tagging the traffic to/from each VLL with a distinct VLAN ID.

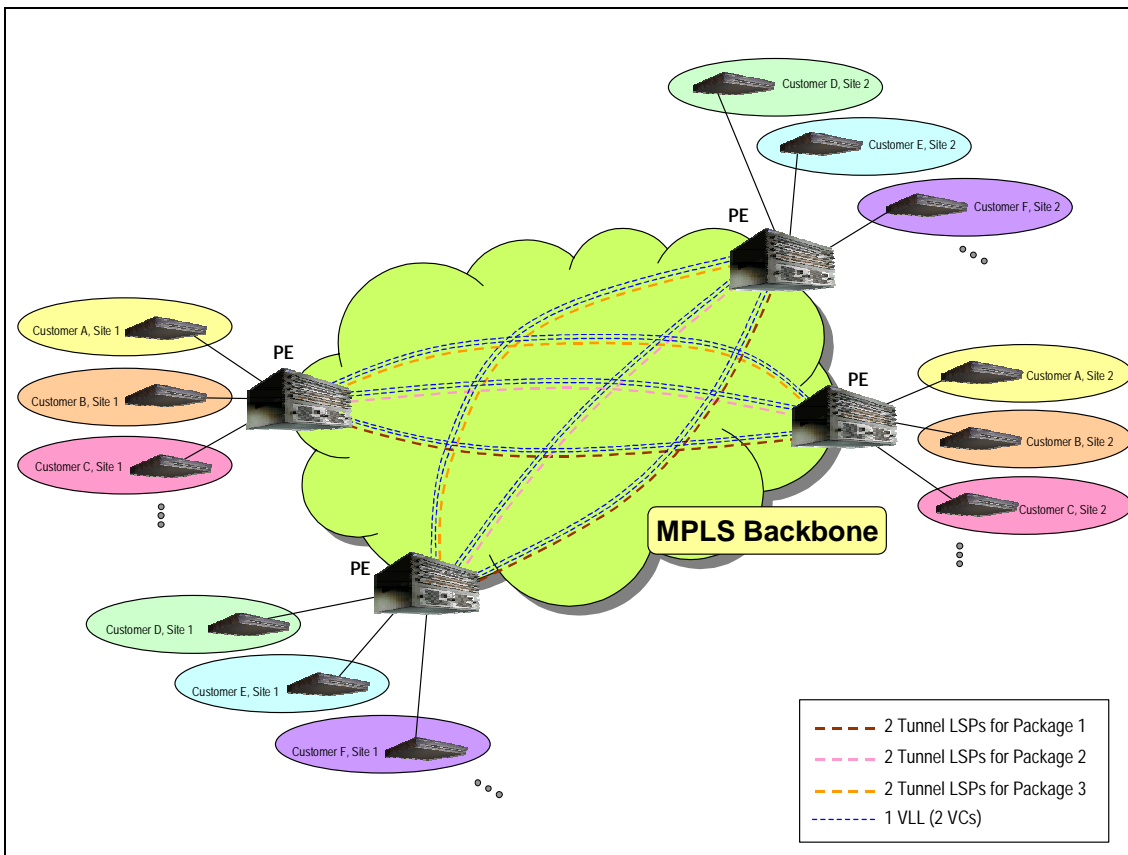


Figure 10 Offering the VLL service in more than one flavor (package).

# Implementing Virtual Leased Lines Using MPLS

## 14.3. Providing Inter-MAN Connectivity

VLLs could be used to carry layer-2 traffic between two customer sites across a MAN boundary, over an IP/MPLS backbone. In the diagram below, the metro zones are pure layer-2 domains where each customer gets a unique VLAN ID. Traffic within the metro zone is layer-2 switched, and STP is used to control the active topology. Assuming that the service providers needs to connect site pairs for customers C1, C2, C3, and C4, the service provider deploys an edge LSR to interface with the IP/MPLS backbone, creates the tunnels LSPs and the required VLLs. For each customer site pair, a VLL is created that carries traffic tagged with a certain VLAN ID from a certain metro zone to another metro zone, with/without changing the VLAN ID<sup>1</sup>.

A VLL realized by two edge LSRs creates the illusion for the core switches – connected to the two edge LSRs – that they are directly connected. This, also, means that both VLANs – in both metro zones – belonging to a given customer will be virtually merged together into one VLAN, which is the desired effect in order to provide such layer-2 connectivity. Another implication is that STP instances that used to run independently on both sides will be merged into one STP instance that controls the new merged VLAN. However, this should not cause any additional complexity. Regardless of which metro zone the root bridge will be in, the VLL will always be in use by the core switches, i.e., no core switch will put the port going to the edge LSR into blocking mode.

To enhance the availability of the VLL, either of the techniques mentioned before could be used: multiple candidate tunnel LSPs for carrying the VLL, or a fault tolerant tunnel LSP that has a secondary path.

---

<sup>1</sup> Foundry's implementation offers the flexibility to do that.

# Implementing Virtual Leased Lines Using MPLS

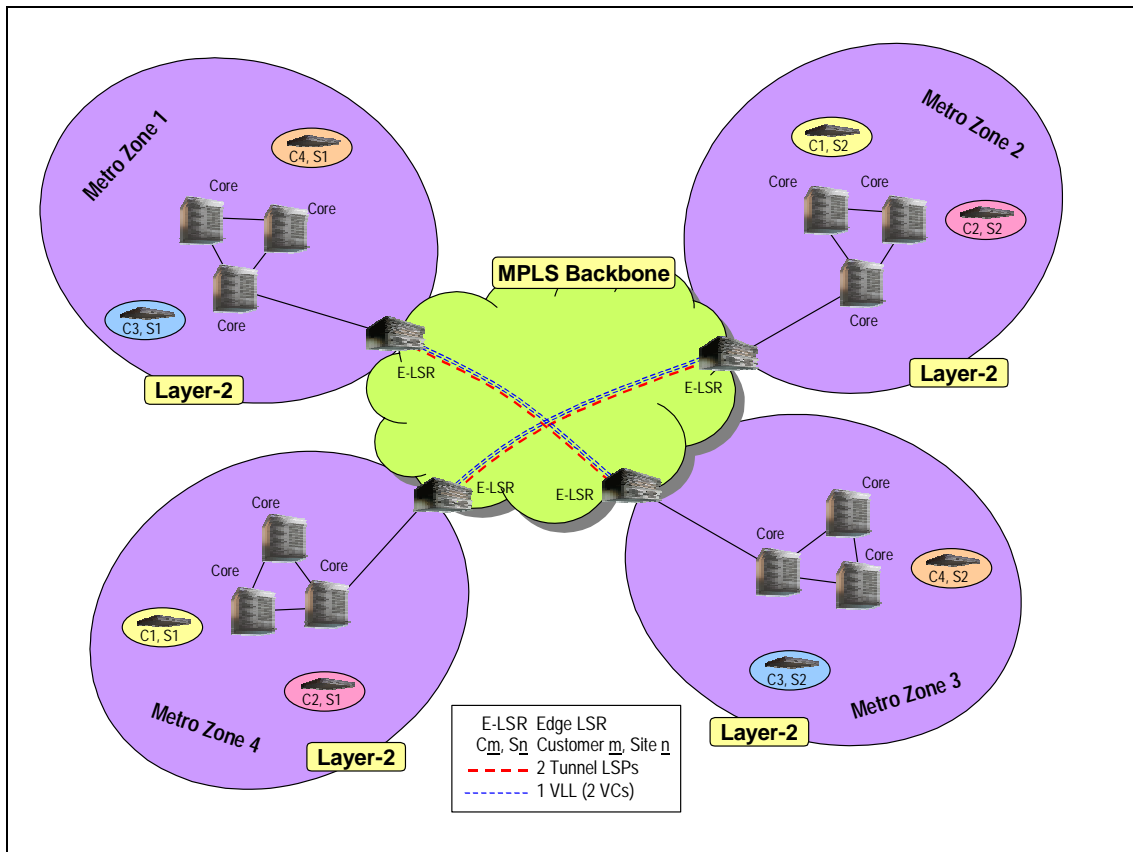


Figure 11 Connecting customer sites across the metro zone boundaries.

For service providers interested in increased availability, or concerned about the failure of the edge LSR, or the link between the core switch and the edge LSR, a scenario like the one shown in the diagram below could be used. This scenario duplicates the edge LSR in each metro zones, as well as the tunnel LSPs and the VLLs. This has an effect equivalent to connecting two pairs of core switches – where a pair is two switches not in the same zone –, such that each pair would have its own direct link.

As mentioned above, VLANs and STP instances will be merged. However, this shouldn't add any considerable complexity to STP management. By proper tweaking of the port cost at the core switch ports that connect the core switches to the edge LSR, for each pair of VLLs serving a certain customer, one could be made active, the other standby.

Using Foundry's STP grouping, one could create – for example – two STP instances, one that converges on a topology that uses the VLL(s) on a certain tunnel LSP, the other instance uses the VLL(s) of the other tunnel. This has the effect of enhancing scalability, since this solution avoids running too many STP instances. Also, it allows for load distribution, by distributing the VLANs of the subscribers over the created STP groups.

# Implementing Virtual Leased Lines Using MPLS

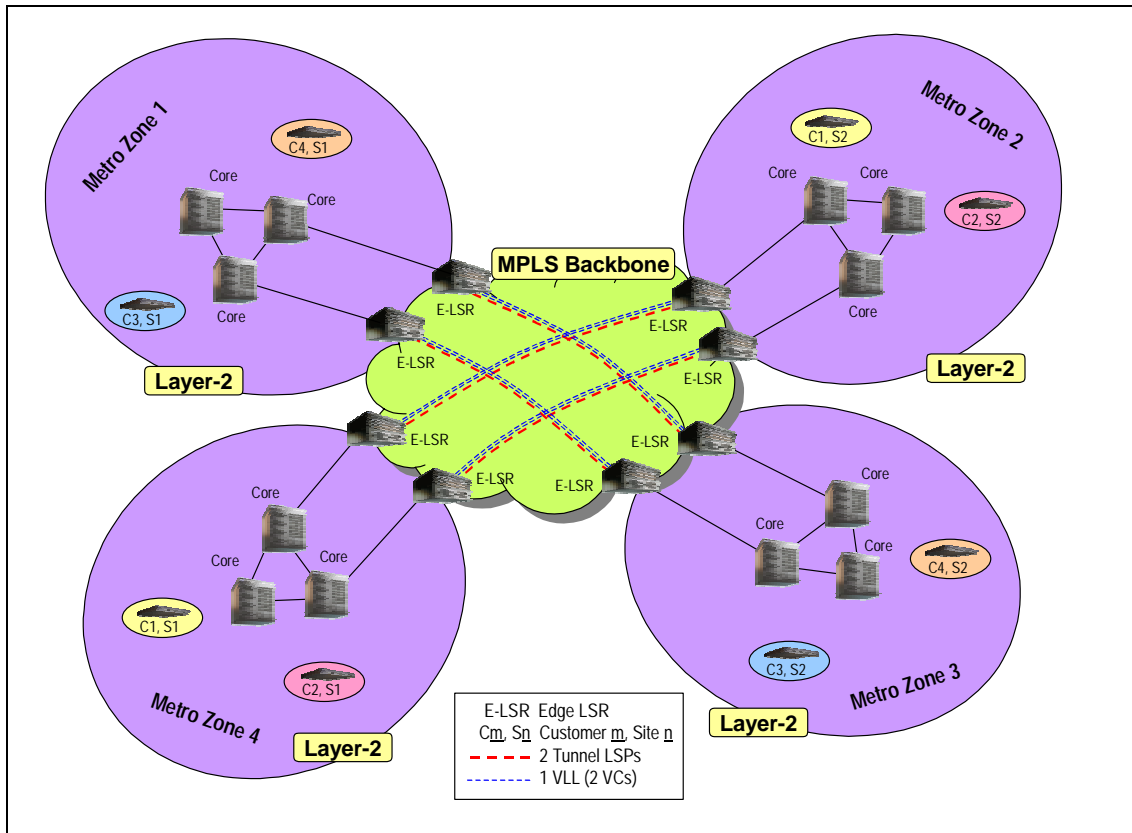


Figure 12 Connecting customer sites across the metro zone boundaries, with edge LSR redundancy.

## 14.4. Replacing DSL/ATM or DSL/Frame Relay Infrastructures

High speed access providers typically relied on an ATM or a Frame Relay infrastructure in order to deliver DSL services to business customers, residential broadband customers, and ISPs.

With MPLS, high speed access providers have another option to consider when building their networks. Using a more cost efficient infrastructure that relies on IP, Ethernet, and PoS, they could provide similar connectivity to subscribers and their respective Internet service providers.

VLLs provide the required transport functionality, to carry subscriber layer-2 frames from the subscriber's premises, over the access provider's infrastructure, to the ISP of choice.

Such an IP/MPLS infrastructure would give the access provider control similar to that achievable with either ATM or Frame Relay. In other words, congestion point avoidance,

# Implementing Virtual Leased Lines Using MPLS

---



distributing flows over the available network links, and creating multiple service levels would all be possible with this approach.

Figure 13 illustrates an example of the use of IP, MPLS, and VLLs in order to implement a high speed access network. As shown in the diagram, for each customer who wants to be connected to a certain ISP, a VLL is created between the customer facing PE device and the PE device facing the desired ISP.

Tunnel LSPs could be pre-provisioned at the time a new customer facing PE is installed, so that whenever a new customer is connected to a customer facing PE, all that would be needed is a VLL definition.

Customers frames are handed off to an ISP as 802.1Q tagged frames; the tag identifies the frames belonging to a certain customer (VLL). This is the equivalent of handing off customer packets tagged with a DLCI – for Frame Relay based DSL – or with a VPI/VCI – for ATM based DSL.

Using NetIron routers as PE devices, the VLLs could be configured such that customer frames would go untagged – normal frames -- into the customer facing PE, and emerge tagged on the other side, or frames could go tagged into the customer facing PE, and emerge tagged with a different VLAN ID on the other side. This decoupling of ISP facing edge and the customer facing edge gives the access provider high flexibility when provisioning the service. It also allows for the use of layer-2 MTUs – for instance, for residential customers – that multiplex traffic from multiple customers over the same uplink using different VLAN IDs.

Note that VLAN IDs have link scope only, not network scope, and hence, the same VLAN ID used over one link, could be re-used over another.

# Implementing Virtual Leased Lines Using MPLS

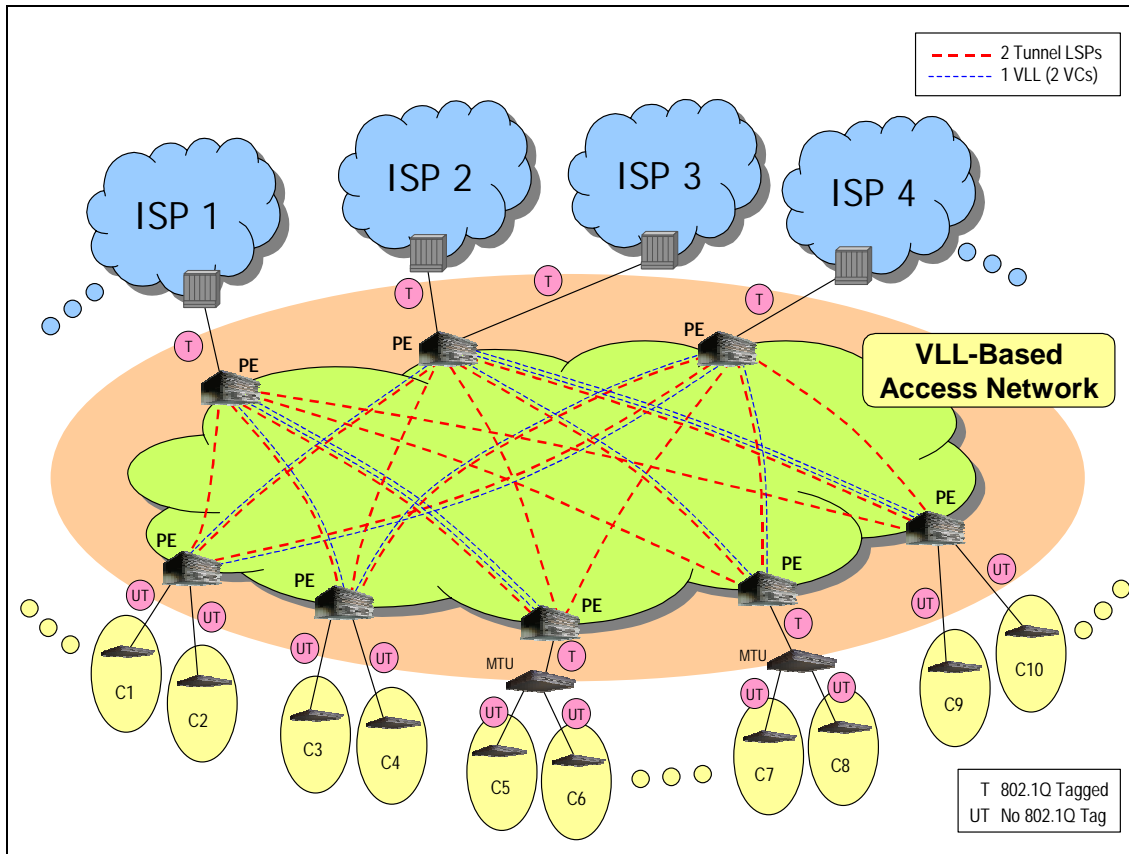


Figure 13 Providing high speed Internet access using an IP/MPLS Infrastructure.

## 14.5. Offering Transparent LAN Services (TLS)

Transparent LAN services means providing multiple site connectivity at layer-2, while switching frames between these sites in a manner that is transparent to the user (customer). Under such scenario, the service provider creates the illusion for the customer edge (CE) devices that they are connected to a layer-2 switch.

Figure 14 illustrates the possible implementation of such services using an MPLS-based network in a metro zone. An external layer of layer-2 switches surrounding the MPLS cloud could be deployed.

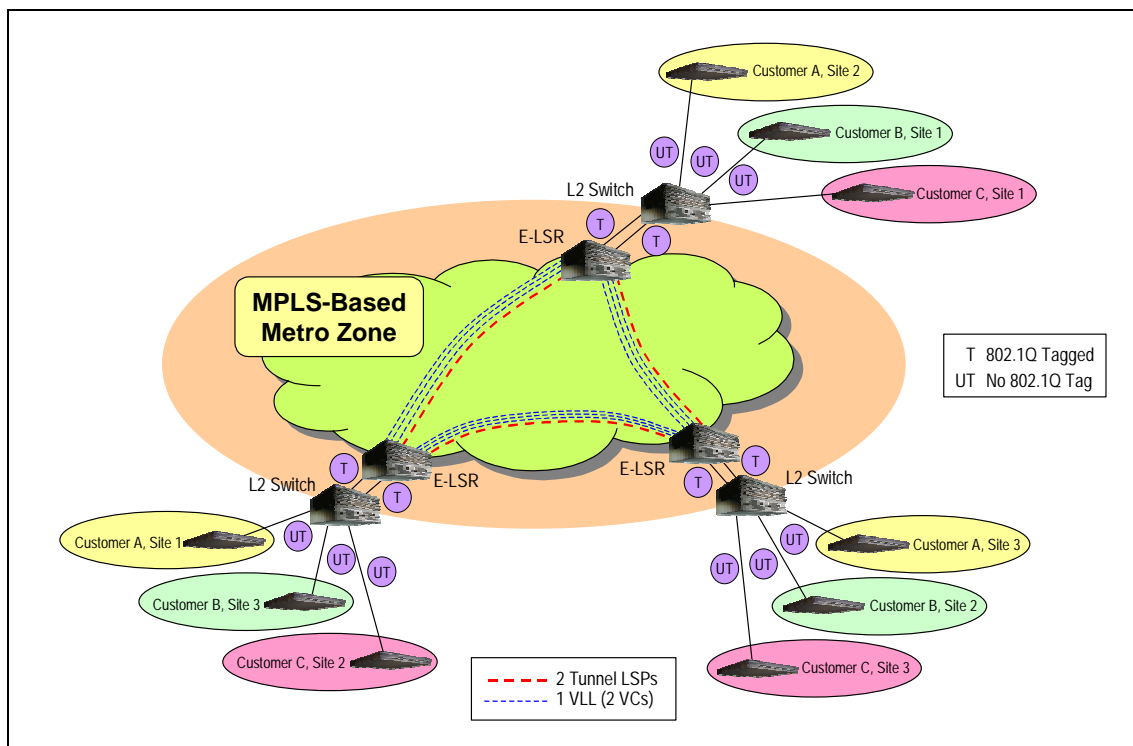
The MPLS cloud provides connectivity between these customer facing switching through the use of VLLs. For the set of sites of a given customer, VLLs are provisioned in order to connect the layer-2 switches facing that customer's sites. The VLLs provide the vehicle for carrying customer frames over the provider's network while offering the ability to distribute flows over the available network bandwidth, and the ability to offer multiple service levels.

# Implementing Virtual Leased Lines Using MPLS

The layer-2 switches perform the following functions:

- Connect the customer to the provider's network.
- Forward customer layer-2 frames – based on their destination MAC address – to the appropriate customer site, i.e., they switch the frames to the appropriate VLLs.
- Running STP with STP grouping<sup>1</sup> to prevent loops in the logical topology.

Customer frames on the links between the layer-2 switches and the edge LSRs (E-LSRs) are identified by a unique VLAN ID. The VLAN ID together with the ingress port on the edge LSR get mapped to a customer's VLL.



**Figure 14** Providing TLS within a MAN based on VLLs.

Also, VLLs could be used to provide TLS services across MAN boundaries. In Figure 15, the provider is running pure layer-2 MANs. Should a customer need connectivity for multiple sites that reside in multiple metro zones, the service provider connects the core switches of the zones where the customer sites are located via VLLs that belong to that customer. The diagram, illustrates the interconnection of the core switches in three different metro zones, in a partial mesh topology in order to provide a fault tolerant

<sup>1</sup> A feature of Foundry switches that allows the creation of a single instance of STP for a group of VLANs in order to avoid launching a separate STP instance for each VLAN. This has the effect of dramatically reducing the number of STP instances required, and hence, enhances scalability.

# Implementing Virtual Leased Lines Using MPLS

service. Several variations of this scheme could be deployed depending on the service provider's needs.

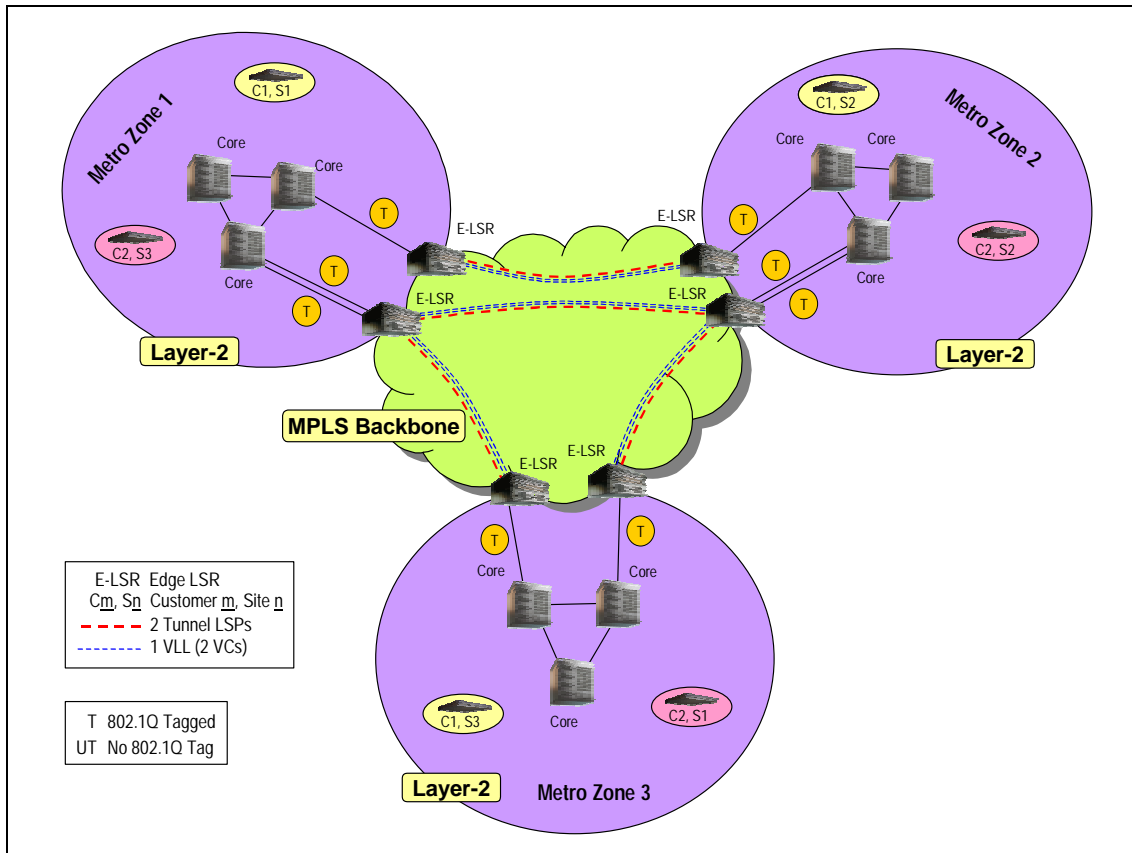


Figure 15 Providing TLS across MAN boundaries.

## 15. Looking Forward

To address the problem of multiple site connectivity at layer-2 using MPLS-based VPNs, another approach is currently being developed by the IETF Provider Provisioned Virtual Private Networks (PPVPN) working group. The new approach is called Virtual Private LAN Segment/Services (VPLS). VPLS is the term of choice when referring to TLS in the context of IP-based VPNs.

The VPLS approach aims at better integration between the pure layer-2 functions required for layer-2 frame switching (source MAC address learning and layer-2 switching) and the MPLS functions of the edge LSR (label exchange, label pushing/popping). These functions could be integrated either by building them into one box, or by distributing them over two boxes one facing the customer, the other acting as

# Implementing Virtual Leased Lines Using MPLS

---

an edge LSR with a special protocol run between both boxes to coordinate the required functions.

Other objectives of the VPLS approach are:

- Minimize/eliminate the use of VLAN IDs as a customer frame identification scheme, in order to achieve a solution that is scalable beyond the limitations of classical Ethernet-based MANs
- Eliminating the use of STP to avoid running into the same scalability issues experienced with classical Ethernet-based designs as to the maximum number of STP instances that could be launched per switch.

Topics like auto discovery, whereby an edge LSR discovers the other edge LSRs that make the customer VPN(s) that it carries, are yet to be researched and standardized.

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.