TECHNOLOGY WHITE PAPER

Virtual Private LAN Service (VPLS) Technical Primer

The emergence of consolidated IP/MPLS networks, together with the development of new standards within the Internet Engineering Task Force (IETF) now enables service providers to offer both VPN services and Internet access from a single packet switched infrastructure. One of the most interesting of the new VPN services is a multipoint Ethernet VPN commonly referred to as virtual private LAN service (VPLS). This paper examines VPLS and describes the basic technical operation of the service.

ALCATEL

Table of Contents

VPN Choices 1

Frame Relay 1

IP-VPN (RFC2547bis) Services 1

A New Contender: VPLS 1

VPLS Over MPLS: Solution Overview 2

MAC Learning and Packet Forwarding 2

VPLS Packet Walkthrough 2

PE Router A 2

Core Router Switching 3

PE Router C 4

Common VPLS Deployment Scenarios 6

Industry Support For VPLS 6

Conclusion 6

List of Acronyms 7

VPN Choices

Today customers have a choice of subscribing to one of the following network-based VPN types to connect geographically dispersed offices:

Frame Relay

Frame relay has been by far the most prevalent service offering to date. However, since the service provider simply offers site-to-site links, customers have to design, manage and maintain their own WAN access equipment. These devices are typically routers that are configured and managed at every site by the customer. The customer designs their WAN architecture around a hub-and-spoke or full-mesh topology, or a hybrid of the two.

IP-VPN (RFC2547bis) Services

In the case of IP-VPN services based on border gateway protocol (BGP)/MPLS, such as RFC2547bis, customers connect routers at each site, but the service provider is responsible for routing between these sites. While this makes multisite connectivity easier, some customers are reluctant to relinquish control of their IP routing to the service provider. In addition, service providers are hesitant to become involved with a customer's routing plan, as this could mean a customer care call to the service provider even though the problem is within the customer's own network.

A New Contender: VPLS

Virtual private LAN service (VPLS) and described in Internet-Draft (I-D) draft-l2vpn-vpls-ldp-00.txt (formerly known as draft-lasserre-vkompella [lasserre-vkompella]), is a class of VPN that allows the connection of multiple sites in a single bridged domain over a provider managed IP/MPLS network. All customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface as the customer handoff, simplifying the LAN/WAN boundary and allowing for rapid and flexible service provisioning.

VPLS offers a balance between point-to-point frame relay service on one hand, and outsourced routed services (e.g., VPRN) on the other hand. In the case of VPLS, customers maintain complete control over their routing, and since all the customer routers in the VPLS are part of the same subnet (LAN), the result is a simplified IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The service provider also benefits from reduced complexity to manage the VPLS service since it has no awareness or participation in the customer's IP addressing space and routing. VPLS also offers some additional advantages:

- A transparent, protocol independent service
- LAN/WAN Ethernet interface on the customer router, which reduces complexity and total cost of ownership
- No Layer 2 protocol conversion between LAN and WAN technologies
- No need to train personnel on WAN technologies such as frame relay since there is no need to design, manage, configure and maintain separate WAN access equipment
- Complete customer control over their routing (unlike IP-VPNs), a clear demarcation of functionality between service provider and customer that makes troubleshooting easier
- No need for the service provider to train technicians to deal with customer routing issues
- Ability to add a new site without configuration of the service provider's equipment or the customer equipment at existing sites
- Faster provisioning, with potential for customer-provisioned bandwidth-on-demand
- Granular bandwidth from 64 Kb/s to 1 Gb/s (compared to frame relay 'step-function' in DS1/DS3 multiples).
- Ability to offer VPLS combined with a managed customer edge (CE) router as a fully managed alternative to IP-VPN services.

The remainder of this discussion will focus on the details of the VPLS solution as described in draft-l2vpn-vpls-ldp-00.txt.

VPLS Over MPLS: Solution Overview

The VPLS architecture proposed in [lasserre-vkompella] specifies use of a provider edge (PE) router that is capable of learning, bridging and replication on a per-VPLS basis. The PE routers that participate in the service are connected together by a full mesh of MPLS label switched path (LSP) tunnels. Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in I-D draft-martini-12circuit-trans-mpls-11.txt ([martini-sig]) is used to negotiate a set of ingress and egress virtual connection (VC) labels on a per-service basis. The VC labels are used by the PE routers for de-multiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

MAC Learning and Packet Forwarding

PE routers learn the source MAC addresses of the traffic arriving on their access and network ports. Each PE router maintains a forwarding information base (FIB) for each VPLS service instance and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating PE routers using the LSP tunnels. Unknown packets (i.e., the destination MAC address has not been learned) are forwarded on all LSPs to the participating PE routers for that service until the target station responds and the MAC address is learned by the PE routers associated with that service.

VPLS Packet Walkthrough

The following is a description of VPLS processing of a customer packet sent from site A, which is connected to PE router A, to site C, which is connected to PE router C.

PE Router A

Customer packets arriving at PE router A are associated to the appropriate VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN ID) in the packet. PE router A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the access port on which it was received.

The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

Figure 1 - VPLS Over MPLS

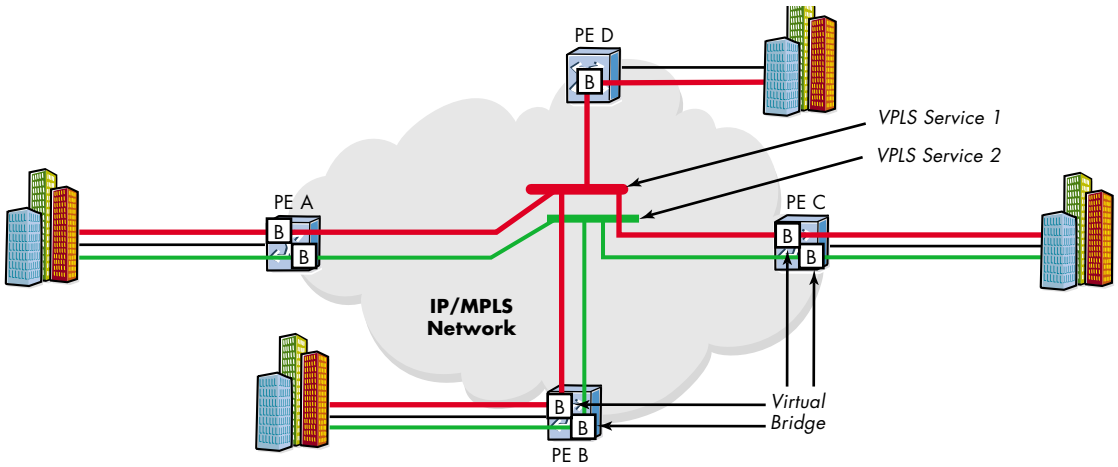

Known MAC address

If the destination MAC address has been previously learned by PE router A, an existing entry in the FIB table identifies the far end PE router and the service VC label (inner label) to be used before sending the packet to the far end PE router C.

PE router A chooses a transport LSP to send the customer packets to PE router C. The customer packet is sent on this LSP once the IEEE 802.1Q tag is stripped and the service VC label (inner label) and the transport label (outer label) are added to the packet.

Unknown MAC address

If the destination MAC address has not been learned, PE router A will flood the packet to both PE router B and PE router C. It does this using the VC labels that each PE router previously signaled for this VPLS instance. Note that the packet is not sent to PE router D since this VPLS service does not exist on that PE router.

Figure 2 - Packet Forwarding by the Ingress PE Router

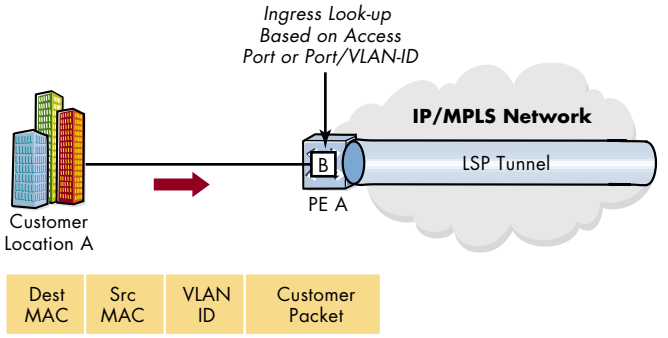

Core Router Switching

All the core routers ('P' routers in IETF nomenclature) are label switch routers (LSRs) that switch the packet based on the transport (outer) label of the packet until the packet arrives at the far end PE router. All core routers are unaware of the fact that this traffic is associated with a VPLS service.

Figure 3 - Packet Forwarding in the Core

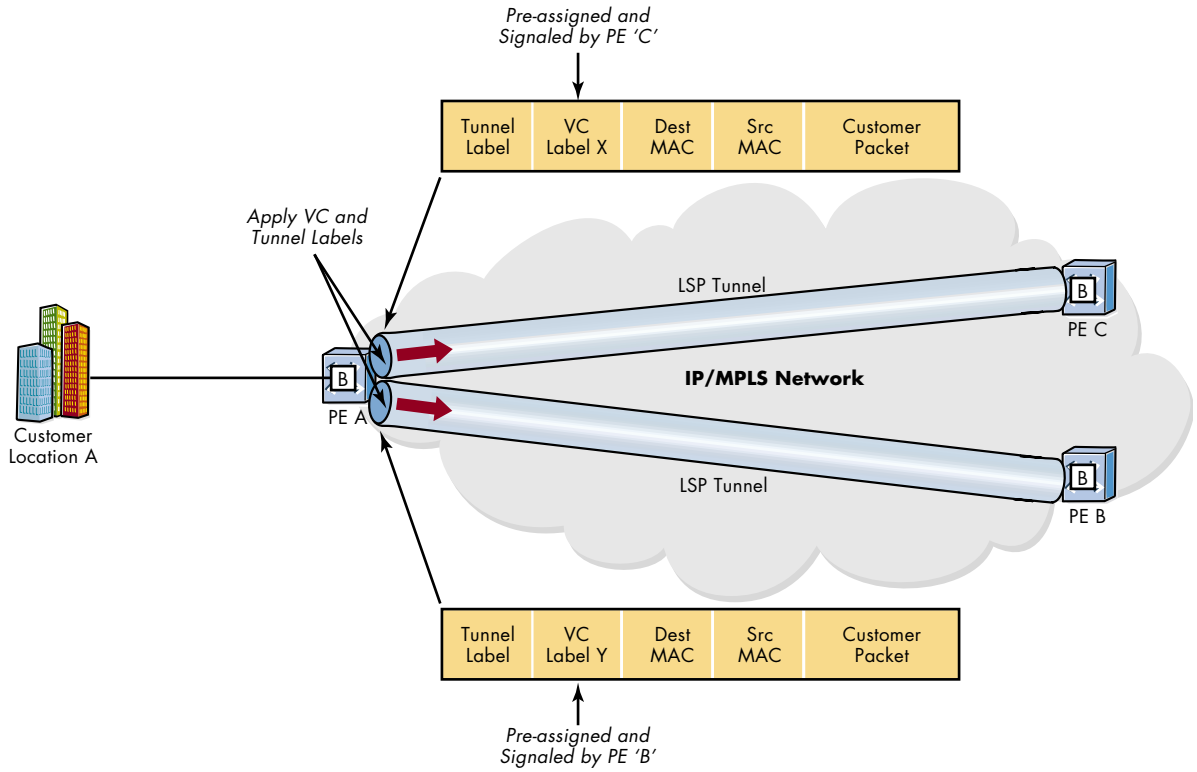

PE Router C

PE router C strips the transport label of the received packet to reveal the inner VC label. The VC label identifies the VPLS service instance to which the packet belongs. PE router C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address with PE router A and the VC label that PE router A previously signaled it for the VPLS service. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Once again there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of PE router C (unknown MAC address).

Known MAC address

If the destination MAC address has been learned by PE router C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag to be added before sending the packet to customer location C. (Note: The egress Q tag may be different from the ingress Q tag that was used on PE router A's access port.)

Unknown MAC address

If the destination MAC address has not been learned, PE router C will flood the packet to all its local access ports that belong to the same VPLS instance as the source MAC address.

Figure 4 - Packet Forwarding by the Egress PE Router

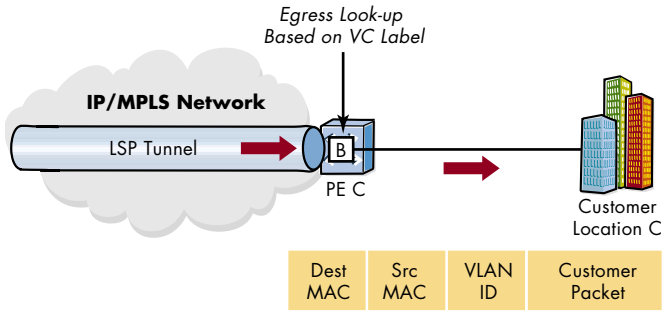

Hierarchical VPLS (H-VPLS)

The H-VPLS architecture initially introduced in draft-khandekar-hvpls-ppvpn-00.txt and subsequently incorporated into [lasserre-vkompella], builds upon the base VPLS solution to provide several scaling and operational advantages. The scaling advantages are gained by introducing hierarchy and eliminating the need for a full mesh of VCs between all participating devices. Hierarchy is achieved by augmenting the base VPLS core mesh of VCs (hub) with access VCs (spoke) to form two tiers, as shown below. Spoke connections are generally created between Layer 2 switches placed at the multitenant unit (MTU) and the PE routers placed at the service provider's point of presence (POP). This considerably reduces both the signaling and replication overhead on all devices.

Figure 5 - Hierarchical VPLS Architecture

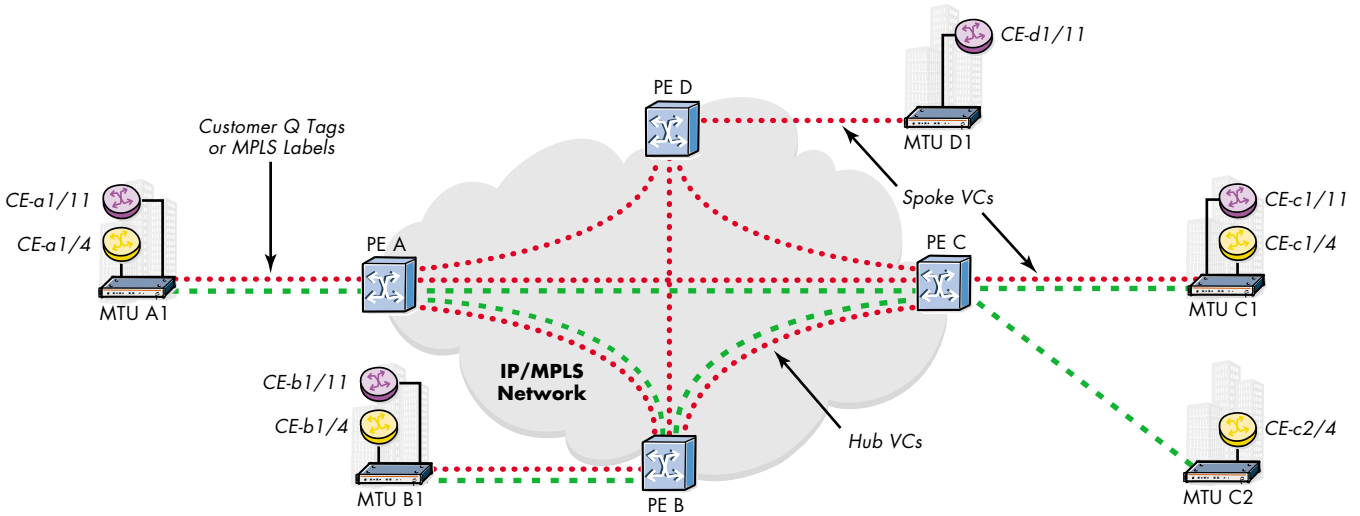

H-VPLS offers the flexibility of utilizing different types of spoke connections - either an IEEE 802.1Q tagged connection or an MPLS LSP with [martini-sig] signaling.

H-VPLS also offers several operational advantages by centralizing all the major functions in the POP PE routers, allowing the use of low cost, low maintenance MTU devices, reducing overall capital expenditures (CAPEX) and operational expenditures (OPEX) (since there are an order of magnitude more MTU devices than PE routers). Another operational advantage offered by H-VPLS is centralized provisioning with fewer elements to touch when turning-up service for a customer. Adding a new MTU device requires some configuration of the local PE router, but does not require any signaling of other PE routers or MTU devices, thus greatly simplifying the provisioning process.

Inter-Metro Services

H-VPLS also enables VPLS services to span multiple metro networks. A spoke connection is used to connect each VPLS service between the two metros and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels is exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

Figure 6 - Inter-Metro VPLS

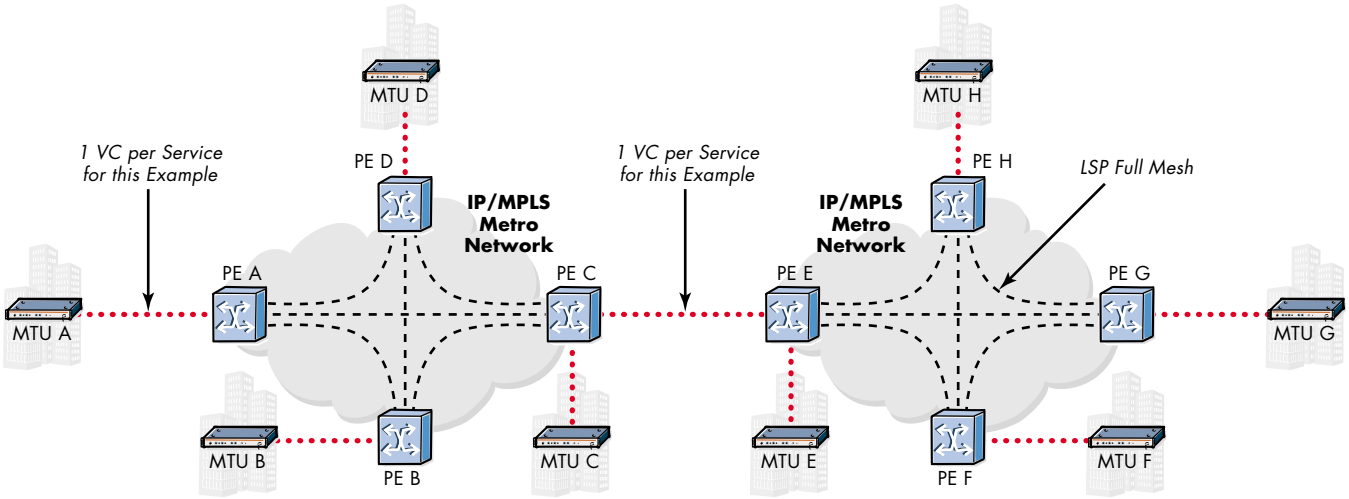

Common VPLS Deployment Scenarios

Regardless of the VPN service utilized, most enterprise customers use routers at the LAN/WAN boundary. Although VPLS is a Layer 2 VPN service and allows the use of Layer 2 switches as the CE device, most customers will use routers to interconnect their office LANs (just as they do for frame relay services).

The advantages to service providers offering VPLS primarily as a router interconnect service are as follows:

- Minimize MAC address exposure, improving scaling (router = one MAC address per site, per service, Layer 2 switch = potentially hundreds of MAC addresses per site, per service)
- Improve customer separation, e.g., unnecessary broadcast or multicast traffic from a badly designed customer LAN will be squelched by the CE router

Service providers may decide to offer a Layer 2 switch interconnect option for VPLS, in which case they could charge per block of MAC addresses. This would allow smaller customer offices to be connected using switches and larger offices to be connected using routers, with differentiated pricing.

Industry Support For VPLS

Alcatel was the first to introduce the VPLS architecture and solutions draft as an alternative VPN offering on IP/MPLS-enabled networks. Since then, VPLS has received widespread support from other vendors, as well as numerous U.S. and international service providers. Many service providers have helped in development of the VPLS solution and have indicated their intent to provide VPLS services as part of their VPN portfolio — an example would be Masergy, the first service provider to offer a commercial VPLS service in the U.S. and Europe. It should be noted that several large service providers that today offer TLS over ATM networks will wish to migrate this service to a packet switched network service such as VPLS because of the advantages outlined above.

Conclusion

VPLS is one of the most exciting emerging VPN services. It offers enterprise customers exactly what they need for intersite connectivity: protocol transparency, scalable and granular bandwidth from 64 Kb/s to 1 Gb/s, fast service activation and provisioning, and a simplified LAN/WAN boundary. VPLS also enables service providers to deliver a scalable VPN service offering that can be combined with Internet access on a consolidated IP/MPLS infrastructure, reducing OPEX. VPLS has received widespread industry support from both vendors and service providers, products such as the Alcatel 7750 Service Router (SR) are now available to enable scalable VPLS, and providers are already beginning to offer commercial VPLS services to their enterprise customers.

Figure 7 - VPLS as a Router Interconnect Service

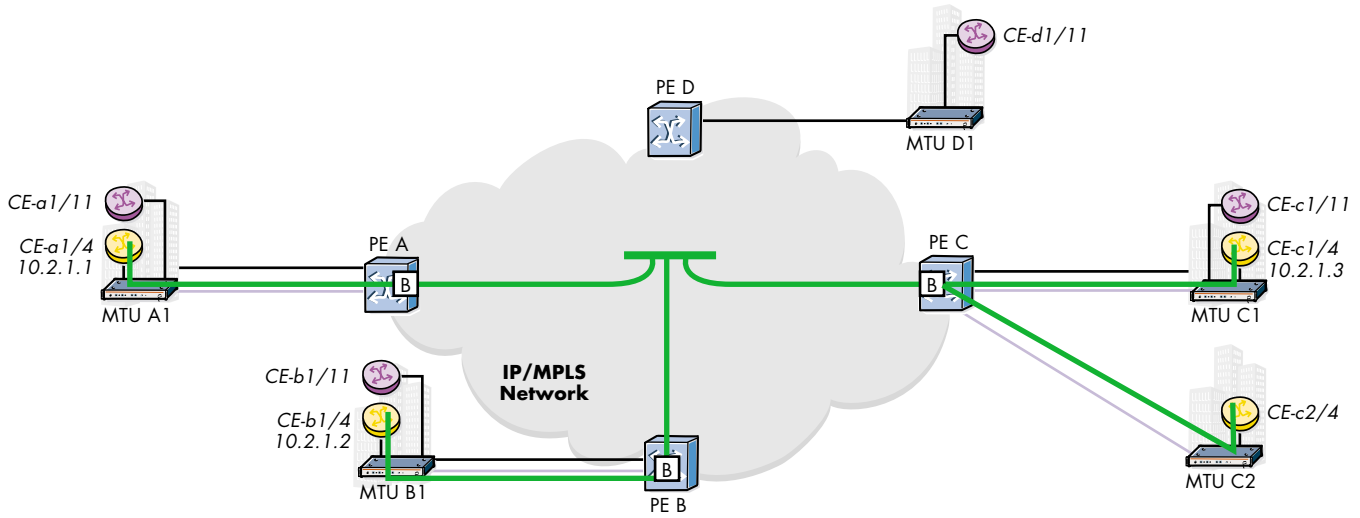

List of Acronyms

BGP	border gateway protocol
CAPEX	capital expenditures
CE	customer edge
FIB	forwarding information base
H-VPLS	hierarchical virtual private LAN service
I-D	Internet-Draft
IETF	Internet Engineering Task Force
IP	Internet protocol
LSP	label switched path
LSR	label switched router
MPLS	multiprotocol label switching
MTU	multitenant unit
OPEX	operational expenditures
PE	provider edge
POP	point of presence
PPVPN	Provider Provisioned Virtual Private Network (IETF Task Force)
TLS	transparent LAN service
VC	virtual connection
VPLS	virtual private LAN service
VPN	virtual private network

www.alcatel.com

3CL 00469 0484 TQZZA Ed.01 17305

ALCATEL