

White Paper

# Virtual Private LAN Service (VPLS)

---

Scalable Ethernet-Based Enterprise Connectivity and  
Broadband Delivery



Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408.745.2000  
1.888 JUNIPER  
[www.juniper.net](http://www.juniper.net)

## Table of Contents

Executive Summary .....	3
Introduction .....	3
Virtual Private LAN Service (VPLS) .....	4
VPLS Implementations .....	5
Auto-Discovery .....	5
Signaling .....	6
Juniper Networks VPLS Solution .....	7
Conclusion .....	9
Glossary .....	9
About Juniper Networks .....	10

## Executive Summary

Virtual Private LAN Service (VPLS) allows different sites to communicate as if they are connected to the same LAN. Service providers offer simplified “any to any” (or “multipoint-to-multipoint”) VPLS service to enterprise customers, allowing enterprises to focus on their core business. In addition, broadband network operators can use VPLS to efficiently distribute “point-to-multipoint” traffic such as IPTV to multiple subscribers concurrently.

There are two VPLS implementations supported by the IETF, and both are implemented in Juniper Networks routers. RFC 4761 uses BGP signalling, while RFC 4762 uses LDP signalling (RFC 4762).

This paper highlights the key architectural difference between them and concludes that the BGP-based implementation provides the highest level of automation and operational efficiency.

## Introduction

Ethernet is the most widely deployed and ubiquitous local area network (LAN) technology in the world with over 100 million Ethernet clients deployed today. Over the past few years, there has been significant innovation around Ethernet standards, not only in the form of dramatic throughput increases from 10 Mbps all the way to 10 Gbps, but also protocol enhancements extending Ethernet’s physical reach to function as a wide area network (WAN) solution – commonly known as Metro Ethernet. In those areas today where Metro Ethernet service is offered by service providers, it is often point-to-point connections between multiple sites within the same metro. However, the ultimate vision held by Metro Ethernet proponents is the ability to move beyond point-to-point connectivity that is confined to a single metro area to deliver point-to-multipoint or multipoint-to-multipoint connectivity either within a single metro or spanning multiple metro areas. In other words, make all sites appear if they are connected to the same simple Ethernet LAN, irrespective of whether the sites are in the same metro area or spread across multiple metro areas. This is known as Virtual Private LAN Service (VPLS), which provides both intra- and inter-metro Ethernet connectivity over a scalable IP/MPLS service provider network.

The alternative to offering an MPLS-based VPLS service is to use stacked VLANs. This is otherwise known as “Q-in-Q” since VLANs are defined in the IEEE 802.1Q standard. However, there are two key challenges with this approach:

- Since there are 4096 unique VLANs, the number of customers is severely limited.
- If there is a network failure, Ethernet’s Spanning Tree Protocol (STP) takes tens of seconds to find an alternate path. Even the newer Rapid Spanning Tree Protocol can take multiple seconds in most situations, and convergence time increases as the network grows.

Therefore, this solution does not meet the needs of most service providers.

## Virtual Private LAN Service (VPLS)

VPLS delivers an Ethernet service that can span one or more metro areas and that provides connectivity between multiple sites as if these sites were attached to the same Ethernet LAN. In contrast to the current Ethernet service offering that is delivered upon a service provider infrastructure composed of Ethernet switches, VPLS uses the IP/MPLS service provider infrastructure. From the service provider's point of view, use of IP/MPLS routing protocols and procedures instead of the Spanning Tree Protocol, and MPLS labels instead of VLAN IDs, significantly improves the scalability of the VPLS service.

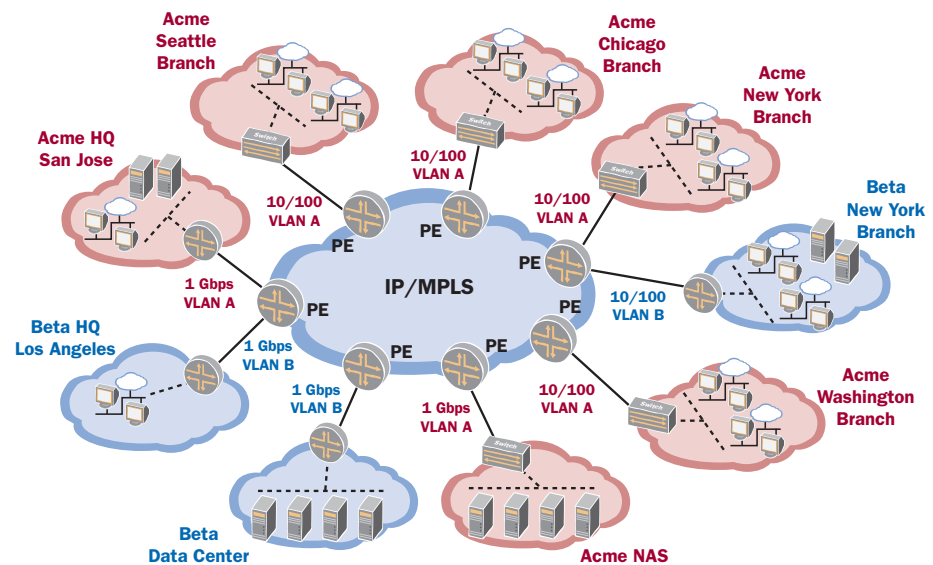


Figure 1: VPLS delivers a flexible Ethernet service that can span one or more metro areas

Each Provider Edge (PE) router at the edge of the service provider's IP/MPLS network is enhanced with special VPLS capabilities as defined by the IETF standards. There are one or more VPLS domains that will be associated with each enterprise that is using the service provider network as a virtual LAN. Each VPLS domain is composed of some number of PEs, each running a VPLS instance that participates in that particular VPLS domain. To keep the concept simple, assume that there is only one VPLS domain per enterprise such that a VPLS instance will run on each PE that is connected to a site belonging to that enterprise. A full mesh of LSPs must be built between all of the VPLS instances on each of the PEs in a particular VPLS domain.<sup>1</sup> Depending on the exact VPLS implementation, when a new PE or VPLS instance is added, the amount of effort to establish this mesh of LSPs can vary dramatically.

Once the LSP mesh is built, the VPLS instance on a particular PE is now able to receive Ethernet frames from the customer site and, based on the MAC address, switch those frames into the appropriate LSP. This is possible because VPLS enables the PE router to act as a learning bridge with one MAC table per VPLS instance on each PE. In other words, the VPLS instance on the PE router has a MAC table that is populated by learning the MAC addresses as Ethernet frames enter on specific physical or logical ports, exactly the same way that an Ethernet switch works today.

<sup>1</sup> As in the IP-VPN architecture based upon RFC 4364 (which updates the better-known RFC 2547), a) these LSPs are visible only to the PE routers – they are not visible to the other routers within the service provider and b) this is accomplished by using MPLS Label Stacking construct.

Once an Ethernet frame enters via a customer-facing ingress port, the destination MAC address is looked up in the MAC table and the frame is sent unaltered (as long as the MAC table contains the MAC address) into the LSP that will deliver it to the correct PE attached to the remote site. If the MAC address is not in the MAC address table, the Ethernet frame is replicated and flooded to all logical ports associated with that VPLS instance, except the ingress port where it just entered. Once the PE hears back from the host that owns that MAC address on a specific port, the MAC table is updated in the PE. Just like a switch, the MAC addresses that have not been used for a certain amount of time are aged out to control the MAC table size.

## VPLS Implementations

There are two standardized VPLS implementations supported by the IETF. The first is RFC 4761: Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling. The second is RFC 4762: Virtual Private LAN Service (VPLS) Using LDP Signaling. Each of these models can be described by two fundamental characteristics:

- Auto-Discovery—What method is used that enables multiple provider edge routers (PE) participating in a VPLS domain to find each other?
- Signaling—What protocol is used to set up MPLS tunnels and distribute labels between PEs for packet demultiplexing purpose?

VPLS Implementation Model	Discovery	Signaling
RFC 4761 (BGP-based VPLS)	BGP	BGP
RFC 4762 (LDP-based VPLS)	None	LDP

## Auto-Discovery

Auto-discovery is absolutely critical to enabling service providers to keep operational costs low, as it automates the creation of the LSP mesh. This is particularly important as it supports the automatic creation of the LSP mesh.

In order to understand auto-discovery further, assume a new PE is added by the service provider. Under RFC 4761 (BGP-based VPLS), a single BGP session is established between the new PE and a route reflector<sup>2</sup>. The new PE then joins a VPLS domain (for example, a new branch office is opened and needs connectivity) when the VPLS instance is configured on that PE, and one or more customer-facing ports on that PE are associated with that VPLS instance (each VPLS instance is identified by a particular Route Target BGP Extended Community, which is configured as part of configuring a VPLS instance.) Once this occurs, the PE advertises that it is part of the VPLS domain<sup>3</sup> via the route reflector to other PEs that are joined in that VPLS instance. Now all appropriate PEs are “aware” of the new PE and these PE members now have all of the information they need to establish LSPs with the new PE automatically.

Since RFC 4762 (LDP-based VPLS) does not specify auto-discovery, the service provider must know explicitly which PEs are part of the VPLS instance. For every VPLS instance present on a PE, the service provider will have to configure that PE with the addresses of all other PEs that are part of that VPLS instance. There are a number of ways this information can be stored: in a LDAP database, in a provisioning system, or even in a spiral notebook. However, all of these mechanisms are operationally intensive and subject to human error.

<sup>2</sup> For the purpose of redundancy, the PE may establish BGP sessions with more than one Route Reflector.

<sup>3</sup> The advertisement carries the BGP Route Target Extended Community that is configured for that VPLS, and this Community identifies the advertisement with a particular VPLS.

It is interesting to note that this same autodiscovery issue has been addressed many times in the past and the answer in most cases ultimately pointed towards BGP. Examples of this include autodiscovery for virtual routers and IP-VPNs (also known as Layer 3 VPNs). One reason often used against BGP is that it is “too complex”. The reality is that BGP is as simple to use by the service provider as other protocols such as OSPF, but BGP can be quite challenging to implement well by the router vendor, which creates a motivation for lobbying against BGP.

## Signaling

RFC 4761 advocates BGP for signaling (label distribution). Alternatively, RFC 4762 uses LDP as signaling mechanisms. The arguments against using BGP for signaling are typically 1) BGP is too complex to use and 2) BGP requires pre-block allocation of labels. However, these arguments against BGP are fairly trivial, since many providers are now deploying IP-VPNs which use BGP and the pre-definition of block sizes has no impact on resources until the actual labels within the block are assigned or configured.

On the other hand, the arguments against LDP signaling are significant. First, since LDP-based VPLS does not define autodiscovery, every time a PE joins a VPLS domain, the service provider must manually look up the other PEs that are part of that VPLS domain. Once this information is attained, they must then build a full mesh of LDP sessions between that PE and every other PE that is part of the VPLS domain. This tremendous overhead of a full mesh of LDP sessions is required because LDP does not have the advantage of BGP’s route reflector architecture. For a service provider offering VPLS for just a few enterprises with a very small number of sites in each enterprise, the burden of LDP may not be that noticeable. However, the burden becomes more and more significant with the growth of the service.

Secondly, this  $O(N^2)$  LDP sessions operational challenge becomes even more noticeable when a service provider chooses to authenticate LDP signaling sessions via MD5. With a full LDP mesh, MD5 keys need to be configured on either end of every LDP session. Thirdly, if the VPLS instance spans multiple Autonomous Systems (ASs), the globally significant 32-bit VCID used by LDP signaling requires operationally intensive manual coordination between ASs. In other words, if a VPLS instance spanned three ASs, all three providers would need to use the same LDP VCID for that VPLS. Finally, if RFC4762 (LDP-based VPLS) is extended to support autodiscovery, then BGP is the most likely mechanism which will be used to perform this function. This in turn requires synchronization of BGP and LDP. Even if LDP sessions already exist between PEs, BGP still needs to communicate which PEs need LSPs established.

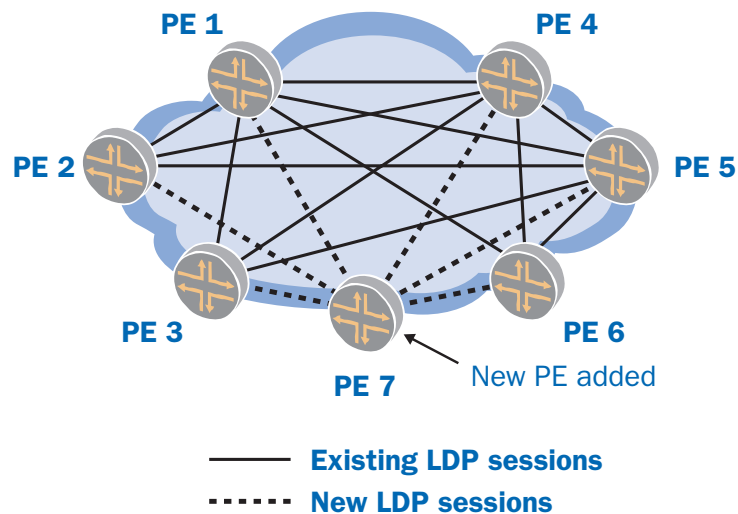
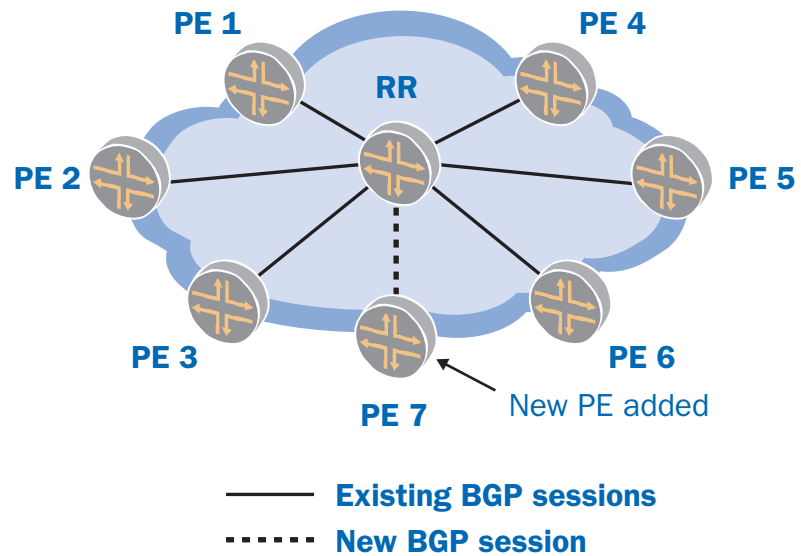


Figure 2: Using LDP requires a full mesh of LDP sessions be established when adding a new PE.

Contrast this approach to using BGP for signaling. When a PE is added, only a BGP session between it and the route reflector need be established. If the session is to be authenticated with MD5, then only keys for the two endpoints of that BGP session need be configured. When a new VPLS instance is configured on that PE, it then advertises its availability via the route reflector, making all other relevant PEs aware of its presence. At the same time, BGP signaling automatically builds the mesh of LSPs for that VPLS instance. Furthermore, if the VPLS instance needs to span multiple ASs (including the case of multiple providers) use of the Route Target for identifying a particular VPLS simplifies operations, as each AS can assign a particular Route Target to that VPLS on its own. This is possible because Route Target Extended Community embeds the Autonomous System Number, and these numbers are globally unique by virtue of assignment<sup>4</sup>.



**Figure 3: A simplified approach where BGP sessions only need to be set up between the new PE and route reflectors, the same way IP-VPNs are deployed.**

## Juniper Networks VPLS Solution

Juniper Networks has implemented VPLS based on both RFCs. BGP-based VPLS is the superior solution, but LDP-based VPLS is supported for those service providers which have already deployed this alternative. The common BGP framework approach shared by IP-VPNs, point-to-point VPNs and VPLS is production-proven in the world's largest networks.

Juniper also provides several important VPLS enhancements, including:

- Point to Multipoint (P2MP) LSP support. This provides efficient distribution of multicast traffic such as IP-based television (IPTV).
- Constrained Shortest Path First (CSPF). This allows VPLS to take advantage of MPLS's traffic engineering to dynamically determine the best path between VPLS endpoints. Different types of traffic, such as video and data, can follow different paths across the network.

<sup>4</sup> The syntax for Route Target is [00 02 xx xx II II II II], where xx xx is an Autonomous System Number, and II II II II is a number from a numbering space, which is administered by the organization to which the Autonomous System Number (xx xx) has been assigned by an appropriate authority.

- Multi-homing support. Juniper integrates BGP's path selection capability with VPLS to allow a customer edge (CE) Ethernet switch to have a back-up path across the network

Juniper Networks routers boast a number of characteristics that make them well suited to serving as PEs operating either at the central office location or the interexchange point of a Metro Ethernet network:

- High density of 10/100Mbps, 1Gps, 10Gbps Ethernet interfaces
- Rich QoS capabilities
- Deep packet processing for layering services with no performance compromise
- MPLS-optimized high performance architecture, proven in the world's largest networks:
  - Point-to-point pseudowire VPNs
  - L3 IP-VPNs
  - Traffic engineering (using RSVP TE)
  - Support for both BGP and LDP
  - G.MPLS
- Highly secure
  - Able to process large filter lists with no forwarding degradation
  - Source address verification for anti-spoofing
- Highly dependable
  - Built on modular JUNOS software
  - Ability to restart PE's control plane with no impact on the forwarding plane
- Rich service set and broad interface portfolio for point of presence (POP) consolidation
  - Reduces operational costs
  - Enables new business models, providers can support Metro Ethernet and private line with services

Networks built with Juniper Networks IP/MPLS solutions enable service providers to construct networks that provide a broad array of connectivity options, a rich set of VPN offerings, and packet processing with no compromise in forwarding performance. Once a Metro Ethernet network is constructed with Juniper Networks platforms, VPLS is simply one of many services that can be deployed. By the same token, service providers that have already deployed Juniper Networks platforms to support dedicated access via TDM or Frame Relay/ATM connectivity can now deploy a Metro Ethernet/VPLS offering. Juniper Networks ability to perform robust packet processing allows service providers to deliver and charge for premium services over and above simple connectivity. This includes services such as:

- Flexible set of VPN services based on a common BGP-based provisioning infrastructure, including VPLS, point-to-point VPNs and IP-VPNs
- High speed filtering on multiple fields for security, such as DOS attack mitigation
- Granular QoS for mission critical applications
- VoIP, including IP Centrex-type services
- Bandwidth on demand via user self-service interface
- Disaster recovery with prioritized traffic for recovery
- Video communications with guaranteed bandwidth
- Streaming information services via multicast
- Detailed accounting for granular billing

## Conclusion

VPLS delivers an Ethernet service that can span one or more metro areas and that provides connectivity between multiple sites as if these sites were attached to the same Ethernet LAN. In contrast to the current Ethernet service offerings that are delivered on a service provider infrastructure composed of Ethernet switches, VPLS uses the IP/MPLS service provider infrastructure, which provides the scalability needed. The use of IP/MPLS routing protocols and procedures instead of the Spanning Tree Protocol and MPLS labels instead of VLAN IDs results in significant improvements in the scalability of the VPLS service. However, all VPLS implementations do not deliver equal benefits. To deploy VPLS with the optimal operational efficiency, service providers should seriously consider using both BGP for autodiscovery and signaling, as specified in RFC 4761.

## Glossary

AS	Autonomous System
IP	Internet Protocol
LAN	Local Area Network
MD5	Message-Digest Algorithm 5
MPLS	Multiprotocol Label Switching
PE	Provider Edge Router
VCID	Virtual Circuit Identifier (used by LDP)
VPLS	Virtual Private LAN Service
VPLS Domain	The virtual LAN that is composed of [N] VPLS instances, each running on a unique PE
VPLS Instance	A software process that runs on a PE that creates a MAC table and enables the PE to act as a learning bridge and participate in a VPLS domain
WAN	Wide Area Network

## About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. The company offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at [www.juniper.net](http://www.juniper.net).

CORPORATE HEADQUARTERS  
AND SALES HEADQUARTERS FOR  
NORTH AND SOUTH AMERICA  
Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
Phone: 888.JUNIPER (888.586.4737)  
or 408.745.2000  
Fax: 408.745.2100  
[www.juniper.net](http://www.juniper.net)

EUROPE, MIDDLE EAST, AFRICA  
REGIONAL SALES HEADQUARTERS  
Juniper Networks (UK) Limited  
Building 1  
Aviator Park  
Station Road  
Addlestone  
Surrey, KT15 2PG, U.K.  
Phone: 44.(0).1372.385500  
Fax: 44.(0).1372.385501

EAST COAST OFFICE  
Juniper Networks, Inc.  
10 Technology Park Drive  
Westford, MA 01886-3146 USA  
Phone: 978.589.5800  
Fax: 978.589.0800

ASIA PACIFIC REGIONAL SALES HEADQUARTERS  
Juniper Networks (Hong Kong) Ltd.  
Suite 2507-11, 25/F  
ICBC Tower  
Citibank Plaza, 3 Garden Road  
Central, Hong Kong  
Phone: 852.2332.3636  
Fax: 852.2574.7803

Copyright 2007 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS and JUNOSe are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

**To purchase Juniper Networks solutions, please  
contact your Juniper Networks sales representative  
at 1-866-298-6428 or authorized reseller.**