


Virtual Private LAN Service on Cisco Catalyst 6500/6800 Supervisor Engine 2T

Introduction to Virtual Private LAN Service

The Cisco® Catalyst® 6500/6800 Series Supervisor Engine 2T supports virtual private LAN service (VPLS) natively in the new PFC4. VPLS is a technology that allows Multiprotocol Label Switching (MPLS) networks to offer Layer 2 Ethernet services. It provides multipoint Ethernet service as compared to Ethernet over MPLS (EoMPLS) that is point to point. VPLS emulates a virtual IEEE Ethernet bridge network.

Unlike Layer 3 VPN, there is no routing interaction between customer and service provider networks. (See Figure 1.)

Figure 1. Virtual Bridges Linked with Virtual Ports, aka Pseudo Wires (PWs)



- Multipoint-to-multipoint configuration
- Forwarding of frames based on learned MAC addresses
- Uses virtual forwarding instances (VFI, like VLAN) for customer separation


VPLS Components

VPLS Concepts and Components Are Common for Enterprise and Service Providers Alike.

- **User-facing PE (U-PE):** The U-PE is the device to which the functions needed to take forwarding or switching decisions at the ingress of the provider network.
- **Network PE (N-PE):** The N-PE is the device to which the signaling and control functions are allocated when a VPLS-PE is distributed across more than one box.

- **Virtual switching instance (VSI):** Virtual switching instance that serves one single VPLS. A VSI performs standard LAN (that is, Ethernet) bridging functions, including forwarding done by a VSI based on MAC addresses and VLAN tags.
- **Pseudowire (PW):** PWE3 is a mechanism that emulates the essential attributes of a telecommunications service (such as a T1 leased line or Frame Relay) over a PSN.
- **Attachment circuit (AC):** The physical or virtual circuit attaching (AC) a CE to a PE. An attachment circuit may be, for example, a Frame Relay DLCI, an ATM VPI/VCI, an Ethernet port, a VLAN, or an MPLS LSP. One or multiple ACs can belong to same VFI.
- **VC (virtual circuit):** Martini-based data encapsulation, tunnel label is used to reach remote PE, VC label is used to identify VFI. One or multiple VCs can belong to same VFI (see Figure 2).
- **VFI (virtual forwarding instance):**
 - VFI creates L2 multipoint bridging among all ACs and VCs. It's an L2 broadcast domain such as VLAN.
 - Multiple VFIs can exist on the same PE box to separate user traffic such as VLANs.

Figure 2. VPLS Concepts and Components




Signaling

Signaling uses LDP to establish and tear down PWs. Using LDP as the signaling VPLS control plane does not have inherent support of auto-discovery. Therefore, LDP-VPLS relies on manual configuration to identify all PE routers.

MPLS in the core, normal LDP sessions per hop to exchange tunnel label or IGP label. Targeted or directed LDP session between PEs to exchange VC label. Tunnel label is used to forward packet from PE to PE. VC label is used to identify L2VPN circuit.

Emulated VC signaling is done using a directed LDP session between PEs. Information such as VC type, VC ID, interface parameter, and so on are negotiated using VC signaling. VPLS on c6500, c6800 platform supports both VC types: VC type 4 (Ethernet VLAN) and VC type 5 (Ethernet). 6500/6800 uses VC type 5 by default, but can negotiate to VC type 4 per peer's request. Similarly, CW is supported on c6500, c6800 platform but will negotiate to no-CW if peer platform does not support it. (See Figure 3.)

Figure 3. VPLS Signaling Using LDP



Data Forwarding

VPLS on Supervisor Engine 2T conducts data forwarding in the exact same way as switch would conduct its forwarding between switched ports:


- Flooding/forwarding:
 - Forwarding is based on VLAN, destination MAC address
 - Unknown unicast/multicast/broadcast is flooded to all ports (IGMP snooping can be used to limit multicast flooding; storm control can be used to limit other types of flooding)
- MAC learning/aging/withdrawal:
 - Dynamic learning based on source MAC and VLAN
 - Refresh aging timers with incoming packet
 - MAC withdrawal upon topology changes

Loop Prevention

VPLS uses split-horizon (Figure 4) to avoid loops (Spanning Tree is possible but not desirable to avoid loops):

- Packet received on VPLS VC can only be forwarded to ACs, not the other VPLS VCs
- Require full mesh VCs among all PEs
- For PE redundancy, active/active VSS provides loop prevention (no active/backup scheme such as STP, EEM, or BGP is required)

Figure 4. VPLS Split-Horizon to Prevent Loops in the Network



H-VPLS

The Cisco Catalyst 6500/6800 Supervisor Engine 2T will support hierarchical VPLS (H-VPLS) natively in the new PFC4. H-VPLS reduces both signaling and replication overhead by using both full-mesh as well as hub-and-spoke configurations. Hub-and-spoke configurations operate with split horizon to allow packets to be switched between pseudowires (PWs), effectively reducing the number of PWs between PEs. (See Figure 5.)

- Minimizes signaling overhead
- Full PW mesh among core devices only
- Packet replication done in the core only

Figure 5. H-VPLS Provides VPLS Scaling




Table 1 shows VPLS IETF standards compliance.

Table 1. VPLS IETF Standards Compliance

RFC	Category	Description
RFC4026	Informational	Provider-provisioned VPN terminology
RFC3809	Informational	Requirements for Layer 2 provider-provisioned VPNs

RFC	Category	Description
draft-martini-l2circuit-trans-mpls-19.txt	Draft	Transport of Layer 2 frames over MPLS
RFC3985	Informational	Pseudowire emulation edge-to-edge (PWE3) architecture
RFC4385	Standards track	Pseudowire emulation edge-to-edge (PWE3) control word for use over an MPLS PSN
RFC4447	Standards track	Pseudowire setup and maintenance using the Label Distribution Protocol (LDP)
RFC 4448	Standards track	Encapsulation methods for transport of Ethernet over MPLS networks

VPLS Support on Cisco Catalyst Supervisor Engine 2T

Cisco Catalyst 6500/6800 running as a PE can run as many as 4K VPLS instances with Supervisor Engine 2T. It can have multiple local Layer 2 ports, and as many as 256 VCs in its flooding domain. If remote (peer) MAC is learned, packet is sent using unicast VC. Incoming Layer 2 frame can be dot1q tagged for VLAN mode or both tagged frame and native frame for port mode.

In CFC mode the 67XX series cards will perform the VPLS encapsulation on the Supervisor Engine 2T. In DFC mode the VPLS encapsulation is performed by the downlink 67XX, 68XX or 69XX (also applicable to all new LCs) Series cards and not in the uplink-facing cards such as with Supervisor Engine 720, which also requires SIP or ES+ to perform VPLS functionality. (See Figure 6.)

VPLS Packet forwarding on the Catalyst SUP2T works in the following manner:

Ingress: A lookup on the MAC table will get the peer ID (if no MAC address is present). We then do a bridge domain and peer ID lookup on PFC4, which points to PW rewrite. For peer ID being 0x3ff (flooding case), TCAM will point to the multicast replication table (MET) to replicate packets first, then rewrite for each PW.

Egress: On egress, the first lookup is like a regular Layer 3 MPLS lookup; label lookup in TCAM points to EoM deencapsulation adjacency, which removes the MPLS label and control word. The inner packet starting from inner DMAC is recirculated. Layer 2 lookup in the second pass will provide the outgoing port.

Figure 6. Cisco Catalyst/6800 SUP2T



For more details please refer to the [Catalyst 6500 Supervisor 2T architecture](#) white paper.

For more details please refer to the [Catalyst 6880-X architecture](#) white paper.

Configuring and Deploying VPLS on Cisco Catalyst 6500/6800 Series Supervisor Engine 2T

VPLS on Supervisor Engine 2T uses the same Cisco IOS® Software code as Supervisor Engine 720, adding additional scalability benefits; the configuration process remains the same on Supervisor Engine 2T as it was with Supervisor Engine 720.

```
12 vfi vpls-300 manual
  vpn id 300
  neighbor 130.0.0.2 encapsulation mpls
  neighbor 130.0.0.3 201 encapsulation mpls

interface Vlan300
  no ip address
  xconnect vfi vpls-300
  !
! Access port
interface GigabitEthernet2/1
  switchport
  switchport access vlan 300
  switchport mode access

! Trunk port
interface GigabitEthernet6/3
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 300,3001
  switchport mode trunk
```

VPLS over GRE

The Cisco Catalyst 6500/6800 Supervisor Engine 2T will allow a flexible transport option for VPLS deployment by supporting VPLS over GRE deployment models. This allows the user to configure VPLS over an IP network, and thus it can be used for a phased migration to MPLS or Layer 2 multipoint encryption.

```
!
Interface tunnel 1
  tunnel mode gre
  mpls ip
  tunnel source 11.11.11.11
  tunnel destination 22.22.22.22
!
Interface tunnel 2
  tunnel mode gre
  mpls ip
  tunnel source 11.11.11.12
  tunnel destination 33.33.33.33
!
Interface TenGigabitEthernet1/1/3/0
  ip address 10.1.1.1 255.255.255.0
```

!

Integrated Routing and Bridging or Routed PW with VPLS

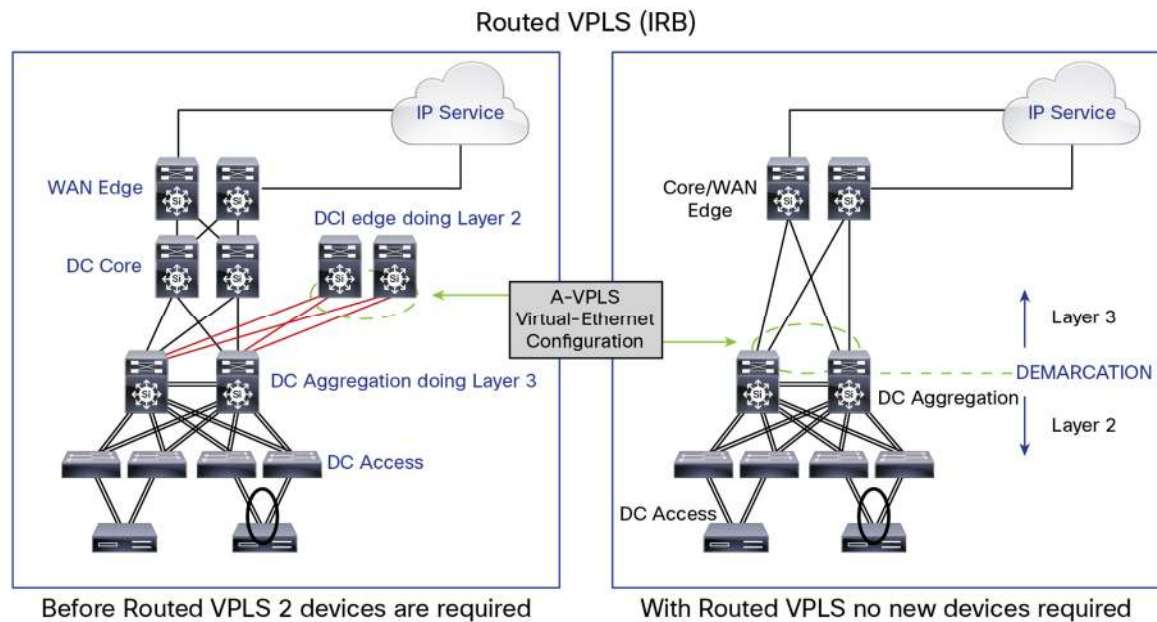
Cisco Catalyst 6500/6800 Supervisor Engine 2T supports configuration of IP address and VRFs on interfaces where Xconnects are configured with the support of the integrated routing and bridging for VPLS feature. This allows more flexibility for VPLS to be configured in the aggregation layer of campus and data center networks.

!

```
Interface VLAN 10
ip vrf forwarding vrf_1
ip address 12.12.12.1 255.255.255.0
xconnect vfi vpls-300
```

!


Figure 7. Routed VPLS Removed the Need for Dedicated DCI Devices or Extending Layer 2 to the WAN Layers



VPLS PE Redundancy with Supervisor Engine 2T

Cisco Catalyst 6500/6800 Supervisor Engine 2T will support VPLS in VSS mode with complete support for NSF/SSO for PE redundancy. This will provide for MCEC-based dual homing into VPLS. (See Figure 8.)

Figure 8. Simplified N-PE VPLS Redundancy with VSS*



* With software release 15.1(1)SY1 VSS Quad Supervisor SSO also supported for all MPLS services, for more information please refer to:

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-729039.html.

Deployment Use Cases

Cisco Catalyst Supervisor Engine 2T provides tremendous scale and flexibility with VPLS, thus allowing it to be deployed in many different deployments use cases:

- Extending Layer 2 domains in data centers:
 - Within different pods in the same single large data center
 - Across geographically dispersed data center over an MPLS and Layer 3 boundary
- Reducing STP domains and extending Layer 2 in large campus networks
- Service providers offering multipoint Layer 2 Ethernet service

Extending Layer 2 Domains Within Different Pods in Same Single Large-Scale Data Center

In large data centers, especially multitenancy data center hosting providers, new data center pods are added when the number of clients grows. This brings challenges of overlapping VLANs, and 4000 VLANs aren't enough to scale.

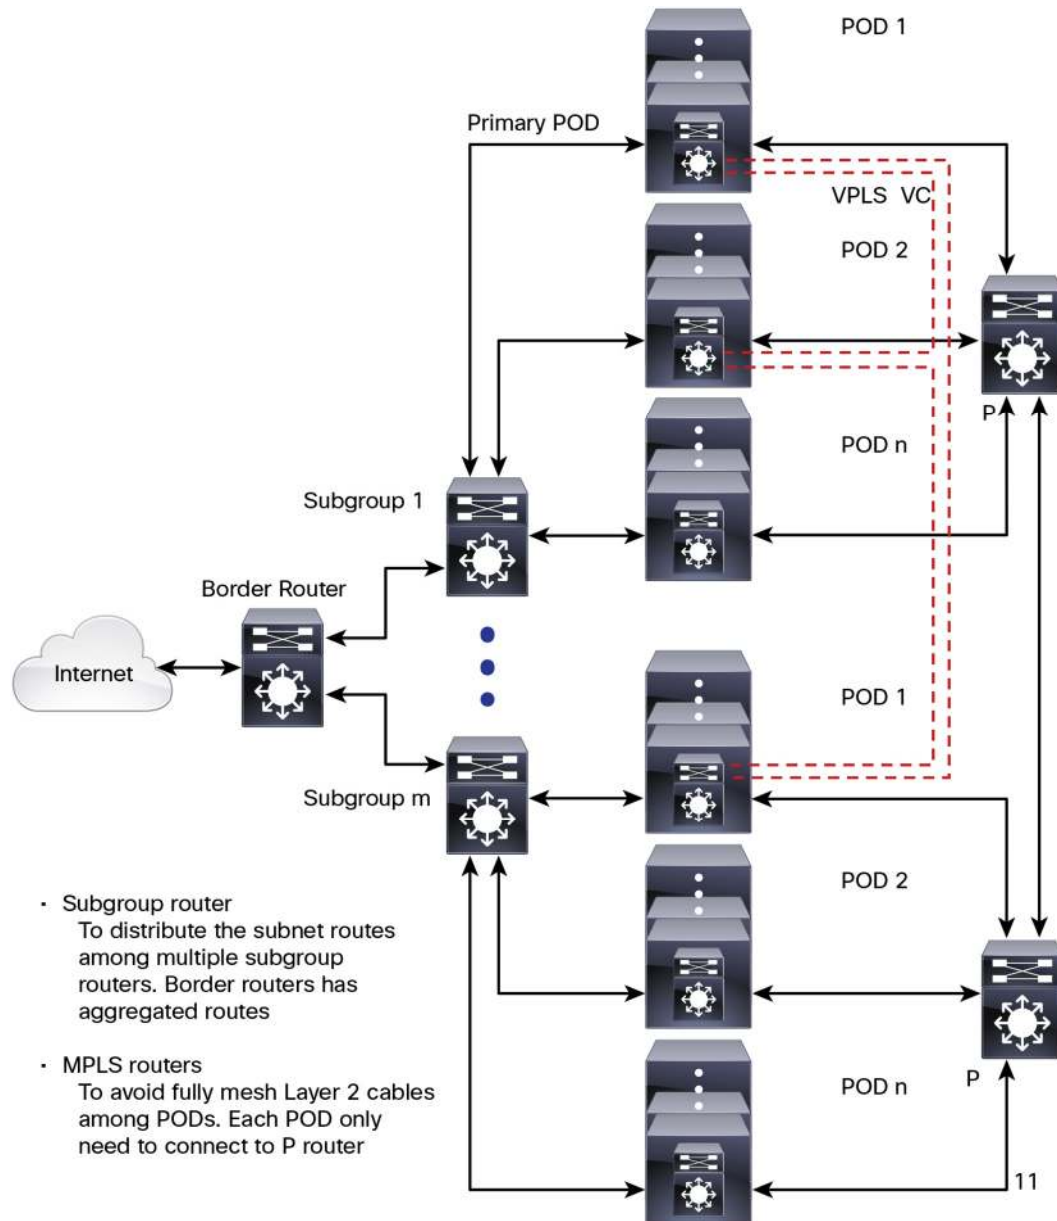
VPLS allows customers to connect discontinuous LAN segments across MPLS/IP core. The VLAN number on different LAN segments can be different: 20 bits of label space to map to bridge domain vs. 12 bits of VLAN space.

As shown in Figure 9, VPLS can be used to scale Layer 2 domains in data centers:

- Each subgroup can support 256K clients:
 - 256K subnets in the subgroup router (256K comes from routing table size, not bridging limitation)
 - Each pod supports up to 4K clients
 - 64 or more pods within each subgroup
- Whole data center supports 256K x m clients
 - m is the number of subgroups
- Clients can move/expand from any pod to any other pod (within/across subgroups)

- VPLS (VFI/VC) is configured on pod Cisco Catalyst 6k when client is added to the pod for the first time

Figure 9. Large Scale Intra DC VLAN Scaling with VPLS

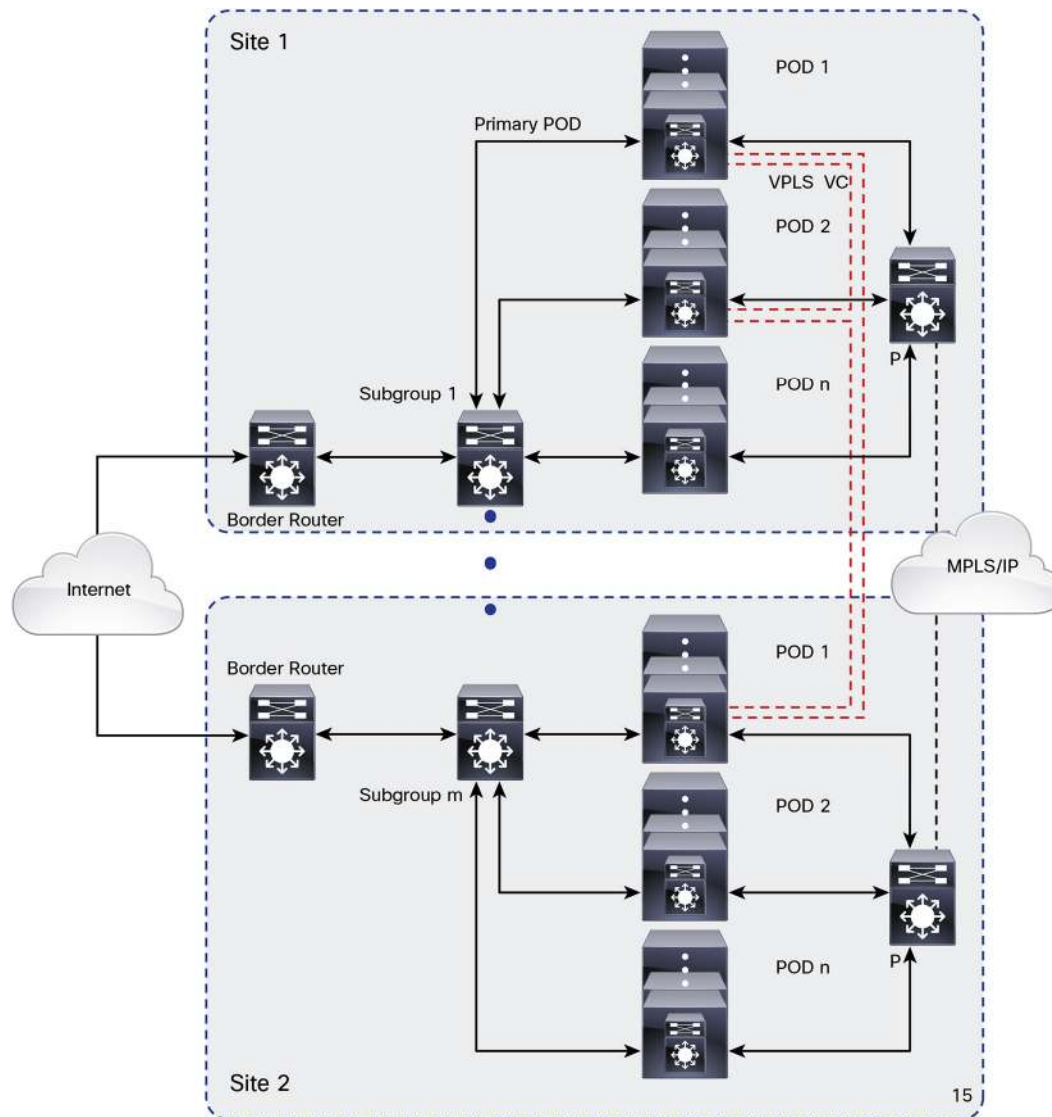


The following section about inter-DC will explain how the same concept is used between data centers.

Extending Layer 2 Domains Across Geographically Dispersed Large Data Center or Across Data Center over a Layer 3 Boundary

For customers who need to extend Layer 2 to more than two sites and point-to-point connectivity is not sufficient, VPLS provides multipoint connectivity over MPLS or IP infrastructure. VPLS allows customers to extend Layer 2 domains while isolating STP in the individual data centers. (See Figure 10.)

Figure 10. Data Center Interconnect Using VPLS




For more information please refer to the [Data Center Interconnect](#) White paper.

Reducing STP Domains and Extending Layer 2 in Large Campus Networks

Based on application requirements, certain customers require network virtualization at Layer 2 between campus sites. It's an alternate and complementary solution to L3VPNs, where application needs drive the Layer 2 extension. These applications could be certain Layer 2 multicast (trading or market data applications) or legacy client server applications with fixed embedded IP addressing, which requires no interaction to a Layer 3 gateway. (See Figure 11.)

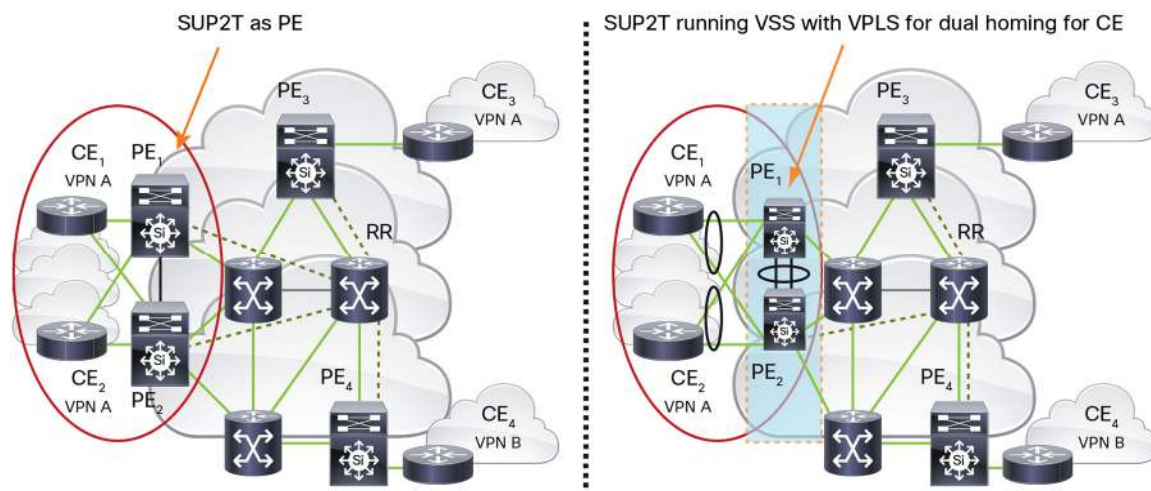
Figure 11. VPLS in Campus Deployments



Service Providers Offering Multipoint Layer 2 Ethernet Service

Service providers can offer L2VPN multipoint connectivity over their own MPLS core. It allows expanding MPLS to the edge or in the access network. This allows service providers to not worry about customer routing policies or IP addressing schemes. It can also serve as a transport network aggregation. Some other deployment scenarios for VPLS are mobile operator backhaul transport using VPLS or VPLSoGRE and multi-tenant DC service providers providing Layer 2 data center interconnect for their customers. (See Figure 12.)

Figure 12. Service Provider Ethernet Offering Using VPLS



Conclusion

The Cisco Catalyst 6500/6800 Supervisor Engine 2T provides the following benefits:

- Native hardware forwarding; no need for SIP or ES+ cards
- No STP required, looping prevented by VPLS split horizon
- VLAN scalability
- MAC scalability: host MAC entries are only in aggregation boxes, which serves this customer. Not exposed in core routers and other aggregation boxes
- VLAN interoperability: same customer can have different VLANs for different servers on different sites
- No MPLS needed in Core Routers if VPLSoGRE is used. VPLS neighbor discovery can be simplified through BGP auto-discovery
- Hardware learning of MAC address; no packet loss due to software MAC learning
- Hardware flooding for unknown destination is rate limited
- Integrated routing bridge (IRB) for VPLS or Routed VPLS
- VPLS PIM and IGMP Snooping for efficient Multicast
- VSS Quad Supervisor SSO support for VPLS adds an additional level of Redundancy



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)