Technology Comparison: Cisco Overlay Transport Virtualization and Virtual Private LAN Service as Enablers of LAN Extensions

What You Will Learn

Geographically dispersed data centers provide added application resiliency and workload allocation flexibility. To gain these benefits, the network must provide Layer 2 and 3 and storage connectivity between data centers. Connectivity must be provided without compromising the autonomy of data centers or the stability of the overall network.

The attributes of the Cisco® Overlay Transport Virtualization (OTV), Virtual Private LAN Service (VPLS), and Cisco's VPLS enhancements are compared in the context of the challenges posed when providing LAN extensions for enterprises. Solutions should:

- Be nondisruptive (transparent to the core and sites)
- Be transport agnostic
- Be multihomed and multipathed
- Preserve failure isolation between data centers

Technology decision makers, IT managers, and network architects will find this document useful in understanding the merits of OTV and VPLS.

The Need for LAN Extensions

Businesses face the challenge of providing very high availability for applications while maximizing infrastructure utilization and keeping operating expenses low. Applications must be available any time and anywhere with optimal response times.

The deployment of geographically dispersed data centers allows the IT designer to put in place effective disaster-avoidance and disaster-recovery mechanisms that increase the availability of the applications. Geographic dispersion also enables optimization of application response through improved facility placement and allows flexible mobility of workloads across data centers to avoid demand hotspots and fully utilize available capacity.

To enable all the benefits of geographically dispersed data centers, the network must extend Layer 2 connectivity across the diverse locations. LAN extensions may be required at different layers in the data center to enable the resiliency and clustering mechanisms offered by the different applications at the web, application, and database layers. Also of importance are the Layer 3 and storage connectivity requirements. This document focuses on the Layer 2 connectivity requirements and how they are best met.

LAN Extensions Compared to Layer 2 Transport Services

Enterprises and service providers have different views on the use of Layer 2 virtual private networks (VPNs). Service providers have requirements that derive from their need to offer a very large number of Layer 2 VPNs as a transport service to a multitude of customers. The technical requirements of service providers are therefore very different from those of the enterprise seeking LAN extensions between data centers. OTV was specifically designed to address the challenges of LAN extensions between data centers. To meet the technical challenges of a provider network, Cisco continues to provide innovative and industry-leading technology, offering transport services for which Multiprotocol Label Switching (MPLS)–based technologies are optimized.

Just as service providers and enterprises face different challenges, they also require different solutions.

Note that some enterprises are structured like service providers and may face some of the same challenges as service providers. Because some challenges are shared, Cisco has designed all MPLS and IP-based transport solutions to be compatible and complementary. For provider- like enterprises, the use of MPLS technologies across the organization can be beneficial, yet certain services may be better addressed by using an IP-based solution such as OTV. For example, an enterprise may have an MPLS backbone that provides Layer 3 VPNs and traffic engineering services, while providing inter–data center LAN extensions with OTV. All traffic, including OTV traffic, will benefit from MPLS-based services in the backbone (traffic engineering fast reroute [TE-FRR]) while optimal LAN extensions are provided by OTV.

Challenges of LAN Extensions

Extending the LAN across multiple data centers creates a series of challenges that are different from the challenges faced by service providers providing transport services:

- **Maintaining site independence:** The extension of Layer 2 domains across multiple data centers can cause the data centers to share protocols and failures that would normally have been isolated when interconnecting data centers over an IP network. These failures propagate freely over the open Layer 2 flood domain. A solution that provides Layer 2 connectivity yet restricts the reach of the flood domain is necessary to contain failures and preserve the resiliency achieved by the use of multiple data centers.
- **Transport independence:** The nature of the transport between data centers varies depending on the location of the data centers and the availability and cost of services in the different areas. A cost-effective solution for the interconnection of data centers must be transport agnostic and give the network designer the flexibility to choose any transport between data centers based on business and operational preferences. An IP-capable transport is the most generalized offering and provides flexibility and enables long-reach connectivity. A solution capable of using an IP transport is expected to provide the most flexibility.
- **Multihoming and end-to-end loop prevention:** LAN extension techniques should provide a high degree of resiliency, and therefore multihoming of the Layer 2 sites onto the VPN is required. Mechanisms must be provided to prevent loops that may be induced when connecting bridged networks that are multihomed.
- **Bandwidth utilization with replication, load balancing, and path diversity:** When extending Layer 2 domains across data centers, the use of available bandwidth between data centers must be optimized to obtain the best connectivity at the lowest cost. Balancing the load across all available paths while providing resilient connectivity between the data center and the transport network requires added intelligence above and beyond that available in traditional Ethernet switching and Layer 2 VPNs. Multicast and broadcast traffic should also be replicated optimally to reduce bandwidth consumption.
- **Scalability and topology independence:** As LAN extensions are deployed in the data center, it is important to provide solutions that do not affect the network design and can therefore be deployed at any point in the topology. This flexibility will usually demand high scalability of the LAN extension solution as the number of edge devices increases as capabilities are pushed toward the data center access.
- **VLAN and MAC address scalability:** The extension of LANs between data centers requires the simultaneous extension of multiple VLANs. Furthermore, in some applications, duplicate VLAN IDs will be in use, and these must be carried independently of each other yet on a common LAN extension. As sites are interconnected, the number of MAC addresses involved will grow, since the MAC address space cannot be summarized; this can become a problem and limit the reach of the solution if not handled correctly.
- **Complex operations:** Layer 2 VPNs can provide extended Layer 2 connectivity across data centers, but will usually involve a mix of complex protocols, distributed provisioning, and an operationally intensive

hierarchical scaling model. A simple overlay protocol with built-in capability and point-to-cloud provisioning is crucial to reducing the cost of providing this connectivity.

How OTV and VPLS Meet the Challenges of LAN Extensions

Table 1 summarizes the different ways that OTV and VPLS meet the challenges of LAN extensions.

Table 1. Comparison of OTV and VPLS

Preservation of Site Independence	
OTV	**VPLS**
OTV conveys MAC address reachability information in a control protocol. The flooding of unknown unicast traffic can be suppressed from the overlay as these are not required to trigger MAC address learning. Flooding anomalies are contained within a single site.	VPLS relies on flooding to propagate MAC address reachability information. Therefore, flooding cannot be prevented.
The OTV control protocol can carry MAC address to IP mappings and use them to populate the Address Resolution Protocol (ARP) cache on the different edge devices. Edge devices will serve as ARP proxies and allow the suppression of ARP broadcasts across the overlay. ARP storms will not propagate across sites.	VPLS does not have a control protocol capable of associating information with particular MAC addresses at an appropriate scale. Controlling ARP traffic and other network events is not practical without the addition of a control protocol.[1]
OTV has built-in filtering capabilities to localize the most common link-local networking protocols (Spanning Tree Protocol, VLAN Trunking Protocol [VTP], Hot Standby Router Protocol [HSRP], etc.) and prevent them from traversing the overlay. This feature prevents protocol failures from propagating across sites. The localization of First-Hop Resiliency Protocols (HSRP, Virtual Router Redundancy Protocol [VRRP], etc.) both isolates failures and helps ensure optimal routing.	VPLS allows the suppression of Spanning Tree Protocol and VLAN distribution protocols such as VTP and Generic VLAN Registration Protocol (GVRP). VPLS does not provide integrated mechanisms to maintain First Hop Resiliency Protocols such as HSRP, VRRP, or Gateway Load-Balancing Protocol (GLBP) localized.
Transport Independence	
OTV	**VPLS**
The overlay nature of OTV allows it to work over any transport as long as this transport can forward IP packets. Any optimizations performed for IP in the transport will benefit the OTV encapsulated traffic.	VPLS requires a label-switched transport to function. This approach is best when an MPLS transport is available. When an MPLS transport is not available, variants such as VPLS over Generic Routing Encapsulation (GRE) allow the deployment of a VPLS solution over a mesh of IP GRE tunnels.
Multihoming and End-to-End Loop Prevention	
OTV	**VPLS**
As part of the OTV control protocol, automatic detection of multihoming is included. This feature enables the multihoming of sites without requiring additional configuration or protocols.	VPLS requires the addition of specific protocols to provide multihoming. Some examples of protocols that must be added to VPLS include Border Gateway Protocol (BGP) with multihoming extensions, Interchassis Communication Protocol (ICCP) and Multichassis Link Aggregation Control Protocol (MLACP), Cisco IOS® Software Embedded Event Manager (EEM), and Multiple Spanning Tree (MST).
	Cisco VPLS eliminates the need for these protocols by using device clustering solutions that allow multihoming that is transparent to the VPLS cloud. With the use of virtual switching system (VSS) technology, a pair of provider edge devices can appear as a single device for the purpose of providing dual-active multihoming of VPLS sites without the need for any additional protocols.
OTV provides per-VLAN single active edge device multihoming by default. When combined with virtual PortChannel (vPC), Cisco Layer 2 multipathing, or Transparent Interconnection of Lots of Links (TRILL) technology, OTV can offer all-active multihoming. Up to 16 active edge devices can be used per site in OTV, allowing continuity of Cisco Layer 2 multipathing and TRILL clouds as they are extended over OTV.	All multihoming schemes for VPLS focus on reducing the multiple provider-edge devices on a site to a single active device.
	Cisco VPLS benefits from the capability to consolidate two provider-edge devices into a single device using VSS to provide active-active dual homing.
Bandwidth Utilization: Replication, Load Balancing, and Path Diversity	
OTV	**VPLS**
OTV uses native IP multicast to help ensure optimal replication of	VPLS uses a full mesh of point-to-multipoint (P2MP) tunnels to prevent

[1] IP-only LAN service (IPLS), because of its static nature, is not appropriate for the volume of hosts that must be handled per site in a LAN extension service.

multicast, broadcast, and signaling traffic.	head-end replication of multicast traffic.
OTV headers are defined to allow the core to hash traffic based on five-tuples of Layer 2, 3, and 4 information and distribute traffic over multiple paths to avoid polarization of encapsulated traffic.[2]	The addition of FAT-pseudowire (FAT-PW) gives Cisco VPLS (and all MPLS services) an effective mechanism for distributing the load over multiple paths in the backbone based on Layer 2 through 4 information and thus avoiding tunnel polarization.
OTV allows effective load balancing of flows across the multiple edge devices available in an all-active multihomed deployment. Load balancing follows equal-cost multipath (ECMP) rules based on the information provided by the OTV control protocol. Hashing is performed based on Layer 2 through 4 information.	Since VPLS is intrinsically single homed for all active-path forwarding purposes, per-flow load balancing is not possible in VPLS.
	Cisco VPLS overcomes this limitation by using VSS to provide dual-active provider-edge multihoming.
Scalability and Topology Independence	
OTV	**VPLS**
OTV is designed to scale to a relatively large density of edge devices (in the hundreds). This capability is a critical element in deciding where in the data center the functions should be deployed. A convenient place to deploy the required edge devices is the aggregation layer in the data center network. This placement simplifies network design and operations by taking the existing Layer 2 domains directly into an OTV overlay when required. Positioning the edge devices in the aggregation layer requires a solution capable of supporting a high number of edge devices, and OTV provides the required scalability. Placing the edge devices elsewhere in the network (at the WAN edge, for example) would require additional hardware to extend the Layer 2 domains from the aggregation layer to the edge devices, increasing the network complexity and operational burden.	VPLS is designed to include a few provider edge devices (40 to 60). If the number of provider edge devices is large, schemes like Hierarchical VPLS (H-VPLS) are required. Use of H-VPLS is equivalent to placing the provider edge devices at the data center edge and adding Ethernet over MPLS (EoMPLS) or IEEE QinQ to aggregate traffic to the provider edge devices. Clearly, there are many elements to be maintained in such a model, and this poses strict topological restrictions on the deployment of VPLS. The use of BGP signaling in VPLS enables greater scalability, does not require H-VPLS, and removes the topological constraints discussed.
VLAN and MAC Address Scalability	
OTV	**VPLS**
OTV intrinsically carries traffic for multiple VLANs over a single overlay.	VPLS can carry a single VLAN per VPLS instance. To multiplex multiple VLANs on a single instance, VPLS uses IEEE QinQ.
OTV has built-in hierarchical identifiers that allow the scaling of the VLAN ID space beyond the 4K VLANs possible in a single 802.1Q domain.	To scale beyond 4K VLANs using a single VPLS instance, VPLS requires the assistance of external devices in the site to provide additional IEEE QinQ encapsulation that can be transported transparently in the single instance.
	Cisco VPLS is enhanced to support more than 4K per instance without the assistance of external devices and with a simplified configuration model.
For the first release of OTV, MAC addresses in the overlay are learned on all sites. However, VLANs can be scoped to the sites where they are relevant, which reduces the size of the tables based on the site. OTV also can provide conversational programming of the forwarding tables to save valuable hardware memory space.[3]	VPLS relies on flooding, and therefore all MAC addresses are learned at all sites. Since multiple VLANs must be hidden from VPLS with IEEE QinQ encapsulation, VLAN scoping is not really an option and all MAC addresses are learned everywhere.
	Cisco VPLS enables much more efficient handling of the IEEE 802.1Q headers, allowing VLAN pruning and benefiting MAC address scalability.
Complex Operations	
OTV	**VPLS**
OTV provides a single protocol to address the different requirements posed by LAN extensions. OTV provides autodiscovery mechanisms that are built into the single protocol and allow point-to-cloud provisioning with zero impact on existing sites.	VPLS requires many protocols to address the different LAN extension challenges. BGP for autodiscovery; Label Distribution Protocol (LDP) for pseudowire establishment; BGP, ICCP, MLACP, and Cisco IOS EEM for multihoming; P2MP LDP for multicast distribution; BGP and Next-Hop Resolution Protocol (NHRP) for GRE tunnel establishment if VPLS is used over GRE; etc.
	Cisco VPLS simplifies much of this complexity by eliminating the need for some of these add-on protocols and concealing others behind an enterprise-class provisioning model.
OTV has been designed with automated processes and little need to use the command-line interface (CLI).	Due to the proliferation of protocols, the VPLS CLI can be rather busy.
	Cisco VPLS has many enhancements to consolidate CLI operations and simplify operations.
Deployment of OTV does not affect the existing network, and therefore OTV can be transparently overlaid without losing site independence or altering the behavior of core or site protocols.	VPLS needs to be carefully designed into the network, which limits flexibility.

Conclusion

Cisco provides significant enhancements to VPLS and continues to invest in the development of innovative technologies for VPLS. Simultaneously, Cisco continues to achieve innovation with new technologies like OTV that

[2] User Datagram Protocol (UDP) headers for OTV are not available at first customer shipment (FCS).
[3] Feature will be available after FCS.

are the product of years of experience. Both technologies have their advantages and disadvantages for the specific application of LAN extensions in the data center interconnect (DCI) space, but OTV provides a much simpler approach.

For More Information

Cisco Nexus 7000 Series Switches: http://www.cisco.com/go/nexus7000.

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.

Printed in USA C11-574984-00 01/10