

VXLAN Pseudowires

Consumers are now heavily invested in mobile access for applications and content. This shift in consumption models is driving new business requirements and creating new challenges for Telcos, from increasingly high bandwidth Over-The-Top (OTT) traffic and competition from the cloud providers. Instead of competing directly with public cloud offerings, Telcos is instead adopting cloud principles to deliver its network services in a more efficient manner. Some are providing cloud connection services to their existing customers to provide secure VPN access to the public cloud. Others are leveraging the user proximity of the Central Office (CO) to offer CDN services or to house 3rd party CDN caches from content providers.

The key drivers for the consumption models are due to new applications taking advantage of virtualization in the compute and storage environment. The network architecture, as a result of virtualized application clusters and workload mobility requirements, now needs to provide extension of logical layer-2 domains across the entire data center, or even between (multiple) data centers. Existing solutions to extend layer-2 domains include Ethernet Pseudowire over MPLS transport (EoMPLS). Pseudowires were developed to emulate a point-to-point connection for service delivery. An EoMPLS Pseudowire carries Ethernet traffic between any two nodes connected over an IP/MPLS network. As with any emulated service, there are shortcomings in the area of multi-pathing, resiliency, multicast traffic optimization, broadcast traffic suppression, as well as protocol overhead.

VXLAN was developed to address the current gaps in addressing virtualization requirements in data center fabrics, i.e. VLAN namespace exhaustion, and transport of layer-2 frames over layer-3 networks in order for workload mobility and for inter-cluster communications of disaggregated workloads. As a benefit, VXLAN is now rolled into designs of data center interconnects, connecting layer-2 domains between (multiple) data centers that are geographically separate to make them look like a single point of delivery, like a Pseudowire solution.




Figure 1: Current MPLS Pseudowire deployment

The Arista VXLAN Pseudowire Solution

Along with providing VXLAN for data center fabrics, Arista EOS extends VXLAN capabilities to provide point-to-point emulated service for connecting enterprise application domains needing layer-2 connectivity across locations - a "VXLAN Pseudowire". Pseudowires primarily used for such a service are delivered over MPLS transport networks today. Arista EOS VXLAN Pseudowire eliminates much of the complexity by providing connectivity for point-to-point Ethernet over a layer 3 IP network, along with multipathing, resiliency, and reducing protocol overhead.

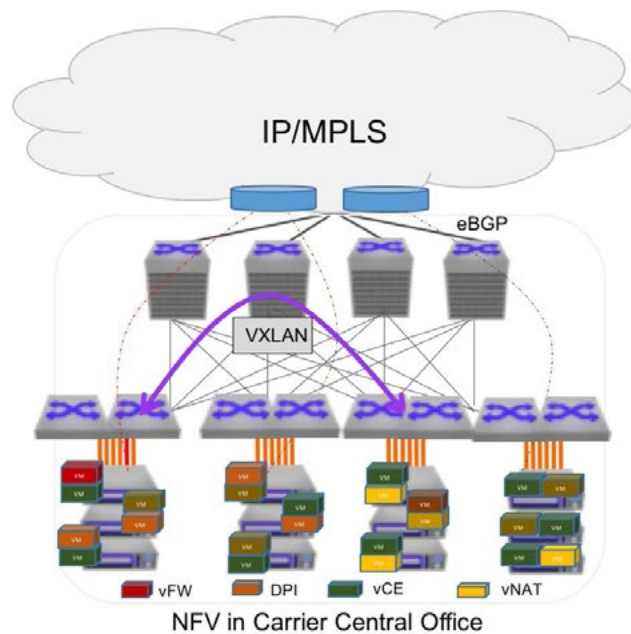


Figure 2: VXLAN Pseudowire for Virtualized Central Office Services

1. VXLAN Pseudowire for NFV Deployments in Carrier Central Office

Service provider Central Offices (CO) are undergoing rapid architectural changes due to the adoption of SDN and NFV architectures. With the deployment of network functions as a virtualized solution, the MPLS PE function (L2VPN or VPLS or L3VPN), is now disaggregated as a VM (virtual PE, or vPE) in the server, along with other functions such as virtual firewall (vFW) and deep packet inspection (DPI), etc. As a result, the CO architecture now becomes a virtualized cloud data center, with network functions delivered in VMs in the compute layer, at scale. To achieve cloud scale, the network design needs a leaf-spine architecture, with resiliency provided from the server to switch via MLAG and Layer-3 ECMP between the leaf and spine layer to provide a high performance, highly available network.

In this new CO, customer's incoming Ethernet circuits will be terminated on the leaf layer, instead of terminating at the PE router. There is, therefore, a need for cross connects/Pseudowires to extend these customer circuits onto the vPEs. There is a need for point-to-point wires between the incoming Ethernet circuit and the virtual PE. Furthermore, it is imperative to disable learning of MAC addresses in this scenario and transparently carry L2 BPDUs.

These requirements can be seamlessly delivered using Arista's VXLAN Pseudowire solution over Arista's Universal Cloud Network architecture. VXLAN Pseudowires with MLAG also provide active-active path redundancy. Thus, the solution is ideal for NFV deployments in the new CO deployments.

2. Simplifying the Cloud Edge




Figure 3: Service Edge Simplification with VXLAN Pseudowire

Another architectural transformation in progress is the cloud edge. Enterprises are rapidly adopting cloud first strategy, whereby significant application workloads are beginning to reside in the cloud. This is fueling bandwidth demand to connect to cloud providers. Amazon's DirectConnect and Microsoft's ExpressRoute offer services to establish private connection between their cloud and the enterprise data centers, or at colocation environments, such as an Exchange provider facility or ISPs. There is a growing trend of ISPs providing gateways to connect enterprise customers to public clouds over secure L3VPN connections.

Given the explosive growth of bandwidth, cloud providers are forced to rely on VRF-heavy design at cloud edge, which increases dependency on expensive, scale-up router platforms. Existing solutions are unable to scale fast enough to cater to the required scale, density and performance.

The cloud edge transformation can be scaled out by taking advantage of virtualization. As stated, the VRF functions can now be delivered as routing instances on compute engines on a server. This allows for scale out of the VRF functions and simplification of the edge to provide scale, density and performance, provided by edge switches like the Arista DCS- 7500 and DCS-7280 series. With the disaggregation of the VRF function from the PE, the customer's Ethernet circuit(s) will now need to extend to the software routing instance. Arista's VXLAN Pseudowire solution provides the perfect solution, whereby the incoming 802.1Q tagged frames can be mapped to the VXLAN VNI and can extend the Ethernet circuit, over layer-3 fabric, all the way to the application rack. In the example deployment above, customer traffic is encapsulated using VXLAN and delivered to the application PoD of interest. This simplification of operations drives agility in the services deployment and can help deliver new services for the operators.

3. Connecting Enterprise Layer-2 Islands




Figure 4: VXLAN Pseudowire for Enterprise Edge

To connect remote sites or have data center interconnects, enterprises have relied on expensive provider provisioned Pseudowire services. With transforming business models, Enterprises are under pressure to rapidly adapt to market changes in productive and cost-effective ways. Cloud adoption, agile services delivery and DevOps/Automation initiatives are driving the transformation. This is leading to changing traffic profiles led by virtualization and bandwidth hungry applications, impacting network costs, due connectivity requirements between data center, branch and cloud.

Seeking to gain more control and cost efficiency, enterprises are looking to own the network services edge, from service provider provisioned to self-provisioned model. Arista EOS VXLAN Pseudowire helps deliver high performance, high scale emulated service between data centers or remote locations, over any IP network. It is an ideal solution for workload mobility, application clustering across locations, multi-site connectivity, live migration between data centers and other use cases services needing emulated services. Enterprises can transform their service edge without the need to buy or expand the footprint of expensive MPLS circuits.

Summary

Every enterprise is adopting or looking to adopt new business models around cloud and mobility. This is driving new datacenter and cloud architectures. To cater to these new models, network architects have to look at new designs for extending application domains across data centers or connect virtualized services at cloud scale over layer-3 fabrics. Arista Networks' VXLAN Pseudowire solution delivered on market leading high-density 10G/25G/40G/100GbE switches, enables organizations to leverage cost effective and efficient methods to meet a variety of requirements and modernize the networks.

Arista's Universal Cloud Network (UCN) spine-leaf architecture with the 7500 Series and 7280 Series deliver the best scale-out network architecture for the SP and enterprise networks. The solution is built on Arista EOS, world's most robust and extensible switch operating system, with powerful capabilities for network automation, network visibility and rich telemetry in combination with EOS CloudVision®, all backed up by a world-class support organization.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. 07/16