



Colocation Datacenter Interconnect with VXLAN and BGP-EVPN

OVERVIEW

Modern cloud applications rely on high performance interactions between services that may be hosted on premise, in the cloud or at partner ecosystem datacenters. Colocation datacenters are at the heart of the infrastructure that enables these interactions. Colocation interconnects have to be agile, high performance, reliable and cost efficient.

In this document, we present a disaggregated networking solution comprising of Mellanox Spectrum switches and Cumulus Linux that addresses the needs for the Colocation Datacenters.

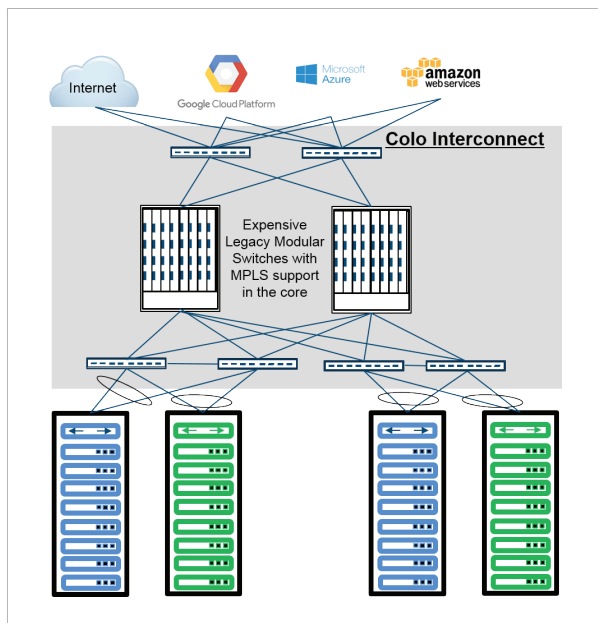


Figure 1. Legacy Colocation Interconnect

THE CHALLENGE WITH THE STATUS-QUO

Today, Colocation datacenters provide patch panel connectivity between customer racks. While the customers have direct access to their racks, they typically need to open a service ticket to setup a inter rack connections. It may take several days or even weeks by the time the connection gets manually provisioned.

In other Colocation datacenters, a more agile, programmable VPLS/VLL based interconnect (see Figure 1) is used. However, VPLS based switch infrastructure has several deficiencies from a cost and management perspective. VPLS is a legacy WAN L2 technology that relies on flood based MAC learning. With all the packet flooding, VPLS is slow to converge and uses bandwidth inefficiently. Additionally, the VPLS infrastructure uses clunky expensive modular systems with MPLS capability. Scalable variants of VPLS such as H-VPLS are too complex to be useful in a Colocation facility. In addition to the expensive hardware, several legacy vendors charge premium license charges for VPLS capabilities.

Colocation Data Center market is highly competitive. The efficiency and performance of the interconnect technology used can be the difference in making or breaking the Colocation Data Center business model.

THE SOLUTION – MELLANOX SPECTRUM WITH CUMULUS LINUX RUNNING VXLAN

Cloud datacenters are now adopting BGP EVPN with VXLAN for network virtualization. Disaggregated switches, 25GbE to server and high density 100GbE switches are being used to improve performance and cost efficiencies. The same technologies can be leveraged to build an agile and high performance Colocation interconnect.

Mellanox Solution comprises of Spectrum 100GbE Ethernet switches running Cumulus Linux and uses standards based VXLAN. The disaggregation of the platform provides customers an opportunity to use the best of breed hardware along with a modern linux based operating system to address their interconnect needs economically.

THE SOLUTION BENEFITS

Cumulus Linux

BGP EVPN with VXLAN scales better than VPLS: Cumulus Linux uses standard based messages to exchange information and automatically discover host IP and VXLAN Tunnel End Points. There is also an option to suppress ARP messages. This results in reduced packet flooding, more efficient bandwidth utilization and faster convergence.

Provides L2 and L3 services with multi-tenancy: Cumulus Linux BGP EVPN supports both L2 and L3 overlay services with VXLAN. Additionally, it allows for overlapping addresses between tenants with support for MP-BGP Route Destinations and Route Targets. With this, Colocation services support multi-tenancy while providing both L2 and L3 connectivity between customer racks.

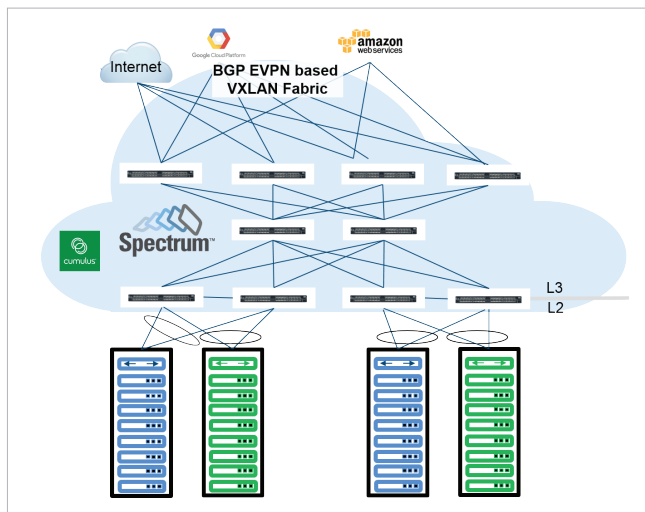


Figure 2. BGP-EVPN fabric as Colocation Data Center Interconnect

Supports simple point-point connectivity: For smaller interconnects, point to point connectivity can be setup using Q-in-Q and VXLAN. Customer ports can be mapped to unique VNI (outer VLAN) and VXLAN tunnels can be setup between the edge switches which effectively emulates point-to-point connectivity between racks.



Mellanox Spectrum Switch High performance 100GbE data plane:

The network is a natural point of convergence for traffic from all racks in the datacenter.

Network performance and fairness across different tenant traffic is critical. Spectrum supports non-blocking line rate 100GbE traffic for all packet sizes. With consistent low latency across all packet sizes, Spectrum switches are ideal for 100GbE high performance interconnect applications.



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

Line rate L3 service: VXLAN routing capability enables L3 interconnect services. Spectrum switches support line rate 100GbE VXLAN routing. Many switches in the market need a second pass of the packet for VXLAN routing. This second pass of the packet halves the bandwidth available and doubles the latency.

Best Horizontal Scale: Spectrum switches can support 10X more peering switches in the BGP EVPN VXLAN fabric than the nearest competitor. Spectrum can replicate broadcast, multicast and unknown unicast packets to upto 750 peer switches.

Flexible form factors: Spectrum switches are available in both half width and full width form factors. The half width switches are ideal for customer racks tighter cost and power constraints.

CONCLUSION

Legacy colocation datacenter interconnect technologies are difficult to provision, utilize bandwidth inefficiently and are expensive. Cumulus Linux BGP EVPN with VXLAN automates endpoint discovery process, minimizes flooding and supports multi-tenancy at scale. Mellanox Spectrum switches have the best VXLAN scale and can perform line rate 100GbE VXLAN bridging as well as routing. Cumulus Linux and Mellanox Spectrum based solution is open, agile, efficient, high performance and is ideal for colocation datacenter interconnect applications.

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high-performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services.

Learn more about Mellanox products and solutions:

www.mellanox.com

About Cumulus Networks

Cumulus Linux embodies native Linux networking. Supercharged versions of the kernel and other networking-related packages encompass the latest industry thinking in networking while retaining compatibility with the full range of software available in Debian. The SN2000 series running Cumulus Linux provides standard networking functions such as bridging, routing, VLANs, MLAGs, IPv4/IPv6, OSPF/BGP, access control, VRF and VXLAN overlays.

Learn more about Cumulus Networks operating system:

www.cumulusnetworks.com