L3VPN End-System Requirements

draft-fang-l3vpn-end-system-requirements-01

Maria Napierala
mnapierala@att.com

Luyuan Fang
lufang@cisco.com

IETF 85, Atlanta
Nov 9 2012

Overview

- Service providers want to deploy service appliances on server COTS hardware.
- There is a requirement to extend MPLS/BGP VPNs to end-systems and associate virtual resources (Virtual Machines, applications, appliances) with VPNs.
- There is a requirement to decouple PE control and forwarding functionality.
  - Allows to implement the PE forwarding on multiple end-system devices, such as operating systems of application servers or network appliances.
  - Allows the PE control plane function to be itself virtualized and run as an application in end-system.

End-System is a computer or a server that sits at the edge of a network.

MPLS/BGP VPNs

- Standard for Network Virtualization in distributed enterprise, finance, B2B.
- L3VPNs are used to interconnect virtual networks across multiple administrative domains.
- L3VPN is also a reality in data-centers today:
  - Interconnection of data-centers (run by multiple admin-domains).
  - Interconnection of data-center virtual networks with existing L3VPN services.
- The draft defines requirements for extending L3VPNs to end-systems.

End-System Layer 3 VPN

- L3VPNs must be able to attach to application end-points and virtual machines that require IP connectivity.
- VPN traffic has to be routed end-to-end, not bridged.
- IP address space: public or private IPv4 and/or IPv6.
- IP service can be unicast, multicast, VPN broadcast.
- Support of a concept of IP-subnet.

End-System CE and PE Functions

Router-based PE

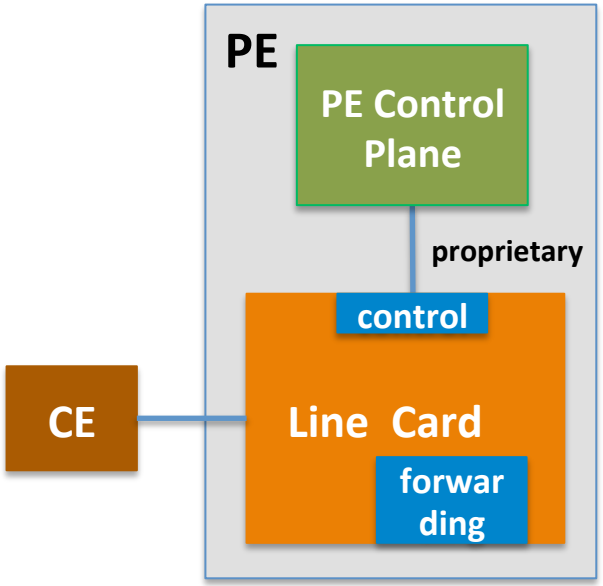

- CE is a physical device external to PE
- CE and PE are routing peers.
- PE device implements both L3VPN control and forwarding functions.

End-System PE

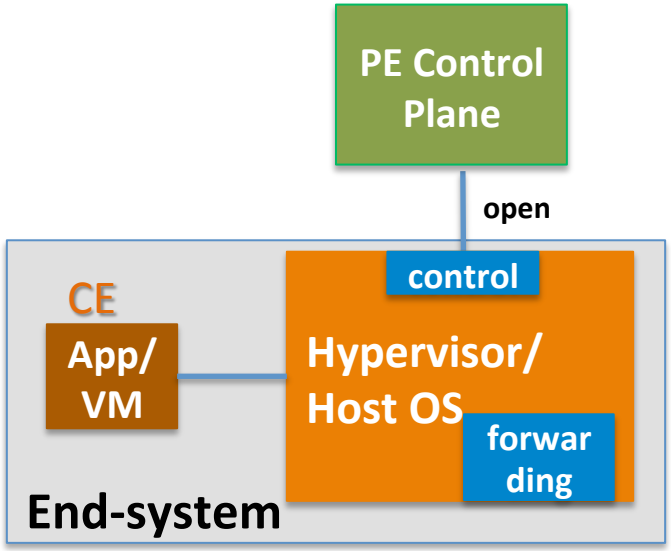

- CE is a non-routing host that resides in a Virtual Machine or is an application on the end-system
- End-system implements only PE forwarding function
- PE control plane (PE CP) function is implemented on an external device/appliance
- Scale: PE CP can control 1000+ of CE interfaces.

Remote End-System Attachments

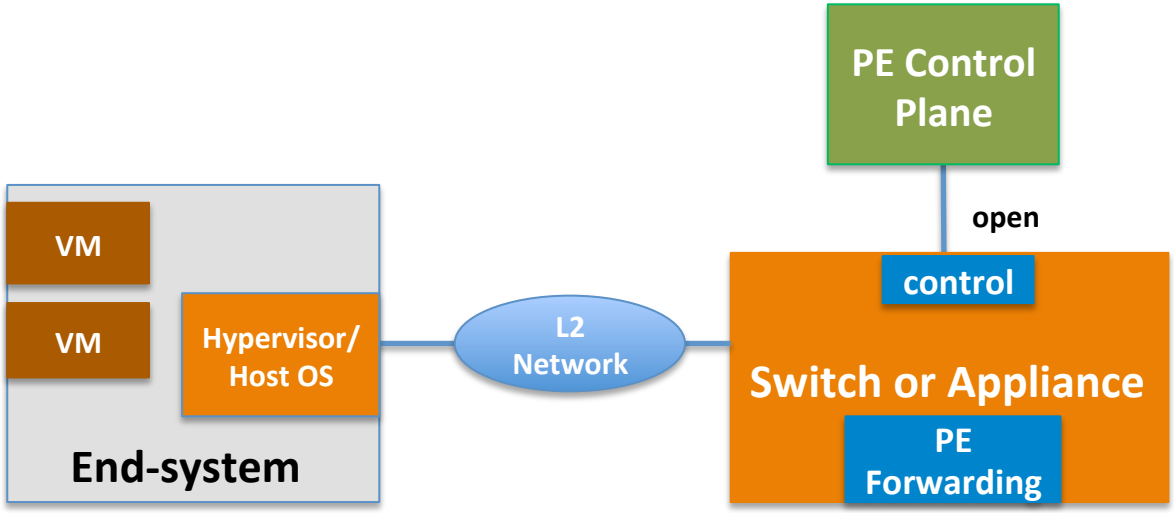

- End-system MPLS/BGP VPN solution must specify a method to convey virtual host information from the end-system to external device that implements PE forwarding
- Layer 2 infrastructure is transparent and unaware of virtual networking.
- The same network virtualization solution should support deployments with mixed, internal (co-located with CE) and external PE (i.e., remote CE) implementations.

Draft's Requirements

- Communication/Addressing (section 3)
- Multi-Tenancy (section 4)
- Decoupling from physical infrastructure (section 5)
- Decoupling from layer 2 topology (section 6)
- Encapsulation of virtual payloads (section 7)
- Optimal forwarding (section 8)
- IP Mobility (section 9)
- Inter-operability with router-based MPLS/BGP VPN (section 10)
- BGP in a Virtualized Environment (section 11)
- Security (section 12)

Next Steps

- Gather input from the WG.
- Asking WG to consider for WG adoption.