draft-sajassi-l2vpn-evpn-ipvpn-interop-01.txt

A. Sajassi (Cisco), S. Salam (Cisco), K. Patel (Cisco), Nabil Bitar (Verizon), W. Henderickx (Alcatel-Lucent), John Drake (Juniper), Y. Rekhter (Juniper), A. Issac (Bloomber), J. Uttaro (ATT)

IETF 85, November 2012

Atlanta

Objectives

- This draft describes how E-VPN can be used as part of an IRB solution to perform optimum unicast and multicast forwarding for both L2 and L3 traffic
- Why an IRB solution?
- Why not just and L2 or L3 solution?

Why Not an L2 Solution?

- No optimum forwarding of inter-subnet traffic
  - Even when the traffic is local – e.g., both subnets are on the same server
- IRB allows for optimum forwarding of both intra-subnet as well as inter-subnet traffic

Why Not an L2 Solution? – Cont.

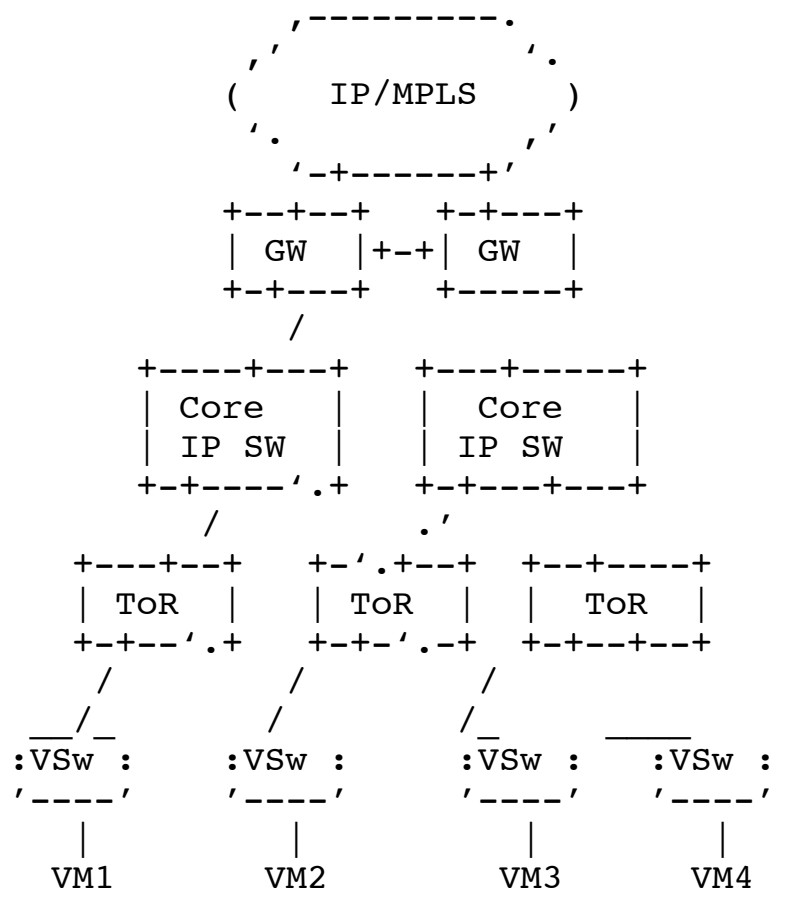

Figure 1: A typical DC network

Why Not an L3 Solution?

- May ran into the following issues for intra-subnet traffic
  - MAC address aliasing issue and not being able to detect duplicate MAC addresses
  - TTL issue for applications that use TTL=1 to confine traffic to within a subnet
  - IPv6 link-local addressing and duplicate address detection – it relies upon L2 connectivity
  - L3 forwarding cannot support the forwarding semantics of a subnet broadcast
  - Support of non-IP applications that require L2 forwarding

E-VPN-based IRB Solution

- An E-VPN-based IRB solution can provide optimum unicast and multicast forwarding for both intra and inter subnets
  - Both within a DC as well as between DCs (East-West traffic)
- But need to inter-operate with IP-VPN PEs as well (North-South traffic)
  - IP-VPN client sites accessing cloud services
  - Communication with IP-VPN ToRs/vSwitch
  - Communication with IP-VPN GWs

E-VPN-based IRB Solution – Cont.

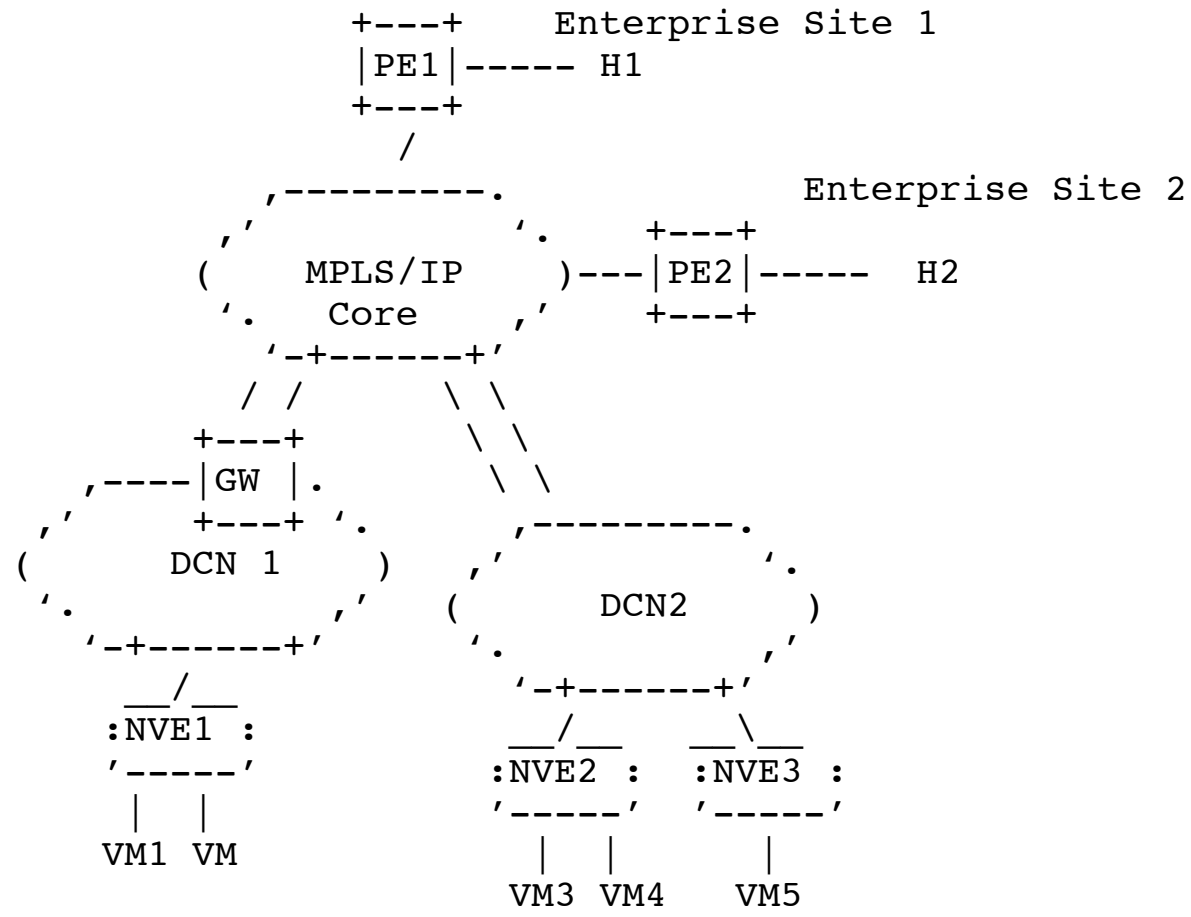

Figure 2: Interoperability Use-Cases

Characteristics of Seamless Interop

- Be completely transparent to the operation of IP-VPN PE
- Be optimal from data-plane forwarding perspective – not need to terminate the encapsulation (no need to look at client MAC/IP addresses)

E-VPN based IRB Solution provides

- Optimal forwarding for intra-subnet (L2) traffic
- Optimal forwarding for inter-subnet (L3) traffic
- Support for both ingress replication as well as P2MP tunnels for multicast traffic
- Support for multi-homing with active/active redundancy and per-flow load balancing
- Support for network-based as well as host-based overlay models
- Support for consistent policy-based forwarding for both L2 and L3 traffic

E-VPN PE Model for Seamless Interoperability

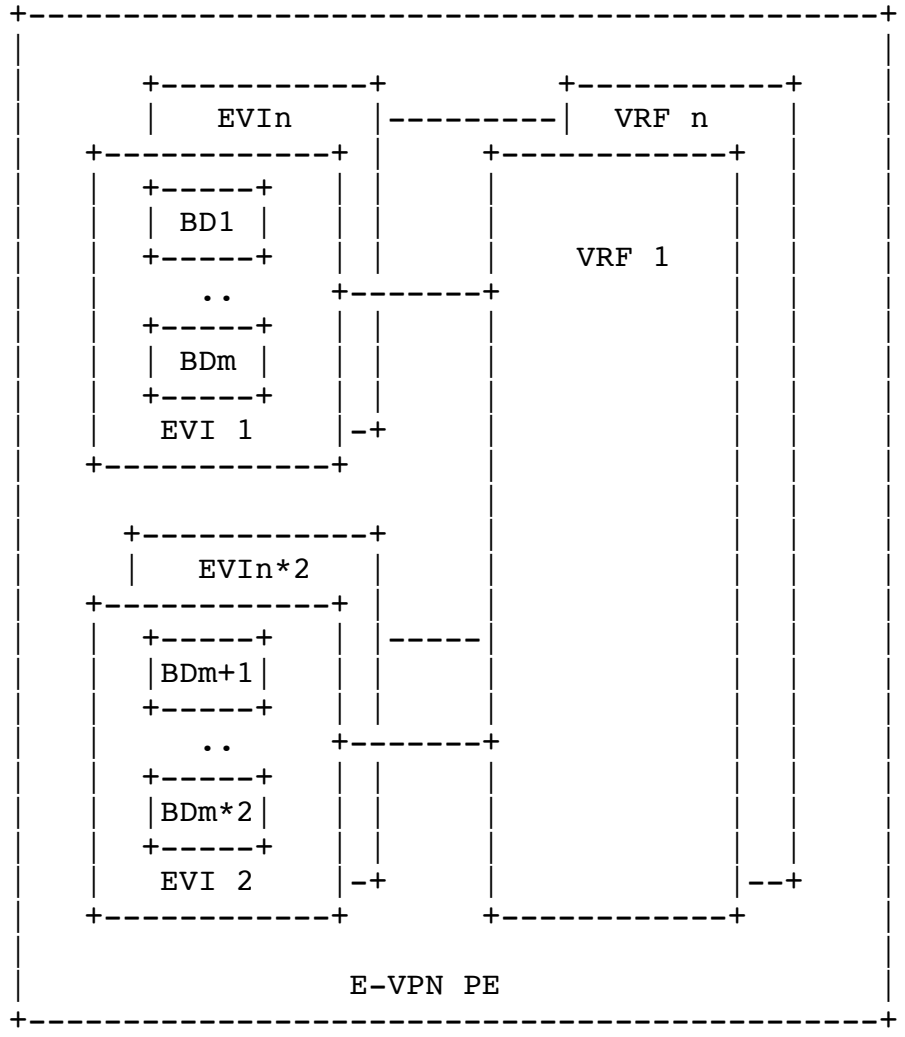

Figure 3: E-VPN PE Model for Seamless Interoperability with IP-VPN

Operation

- E-VPN PEs are bilingual – they advertise both E-VPN and IP-VPN routes
- When an E-VPN PE receive a MAC route, it uses client MAC address to populate the BD/EVI table and it uses the client IP address to populate the VRF table
- When an E-VPN PE receives a packet over MPLS/IP network, it uses client MAC address to decide whether IP forwarding is required or not
  - If MAC address correspond to that of its BVI, then it lookup the VRF table