

draft-sajassi-l2vpn-evpn-inter-subnet-switching-04.txt

A. Sajassi (Cisco), S. Salam (Cisco), Y. Rekhter (Juniper), John Drake (Juniper), Wim Henderickx (ALU), Florin Balus (Nuage), Lucy Yong (Huawei), Linda Dunbar (Huawei), Jorge Rabadan (ALU), Samir Thoria (Cisco), Dennis Cai (Cisco)

IETF 90, July 2014

Toronto

History

- Published & presented rev01 for IETF-86 (Florida)
- Published & presented rev02 for IETF-87 (Berlin)
- Published & presented rev03 for IETF-89 (London)
- Publishing & presenting rev04 for IETF-90 (Toronto)

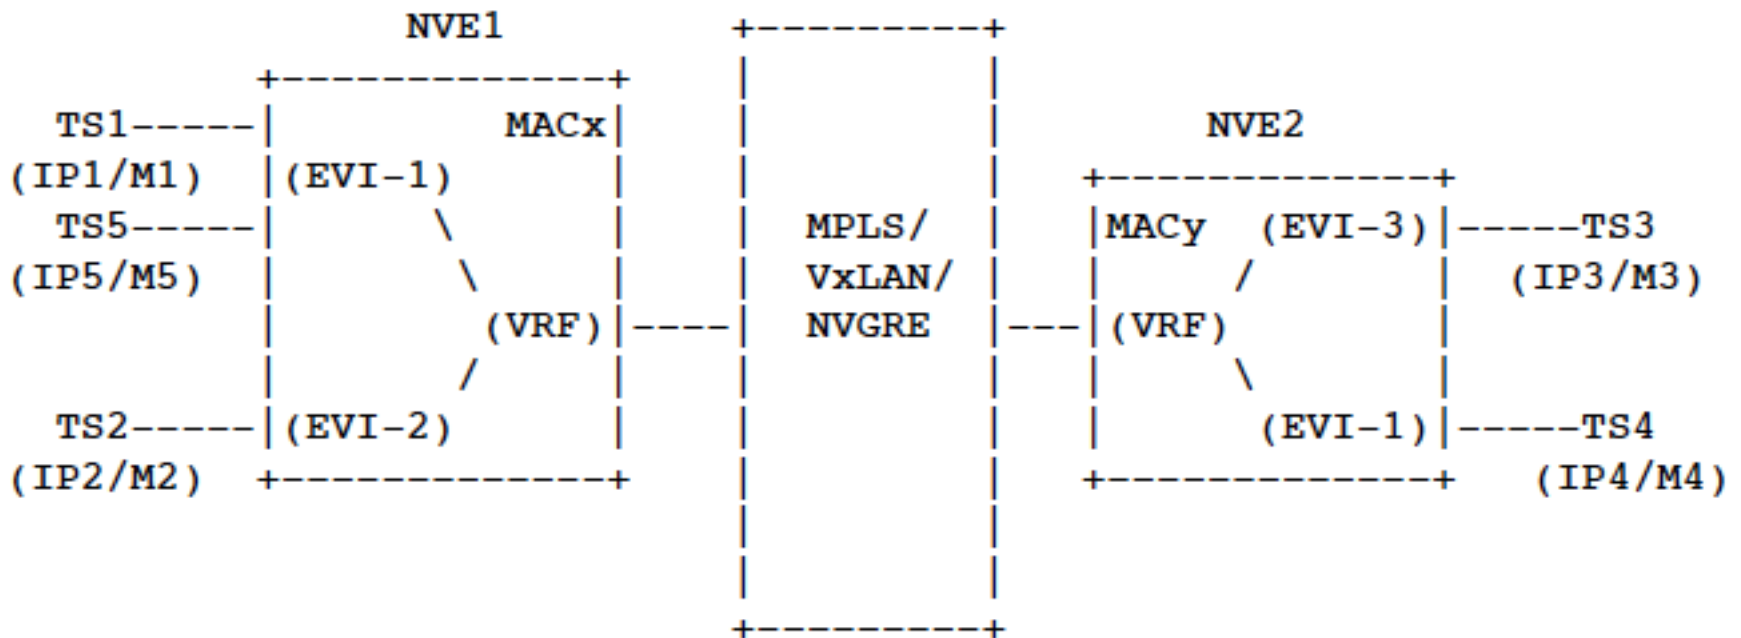
Industry Adoption

- Several vendors have implemented or are implementing this draft: Cisco, Juniper, ALU
- Several service providers support this draft

Changes from rev03

- Expanded on “Operational Models for Symmetric Inter-Subnet Forwarding” by adding two subsections:
 - 5.1: IRB forwarding on NVEs without core-facing IRB Interface
 - 5.2: IRB forwarding on NVEs with core-facing IRB Interface

IRB forwarding on NVEs without core-facing IRB Interface



Control Plane Operation

Each NVE advertises an RT-2 (MAC/IP Advertisement Route) for each of its TSeS with the following field set:

- RD and ESI per [EVPN]
- Ethernet Tag = 0; assuming VLAN-based service
- MAC Address Length = 48
- MAC Address = M_i ; where $i = 1,2,3,4$, or 5 in the above example
- IP Address Length = 32 or 128
- IP Address = IP_i ; where $i = 1,2,3,4$, or 5 in the above example
- Label-1 = MPLS Label or VNID corresponding to EVI
- Label-2 = MPLS Label or VNID corresponding to IP-VRF

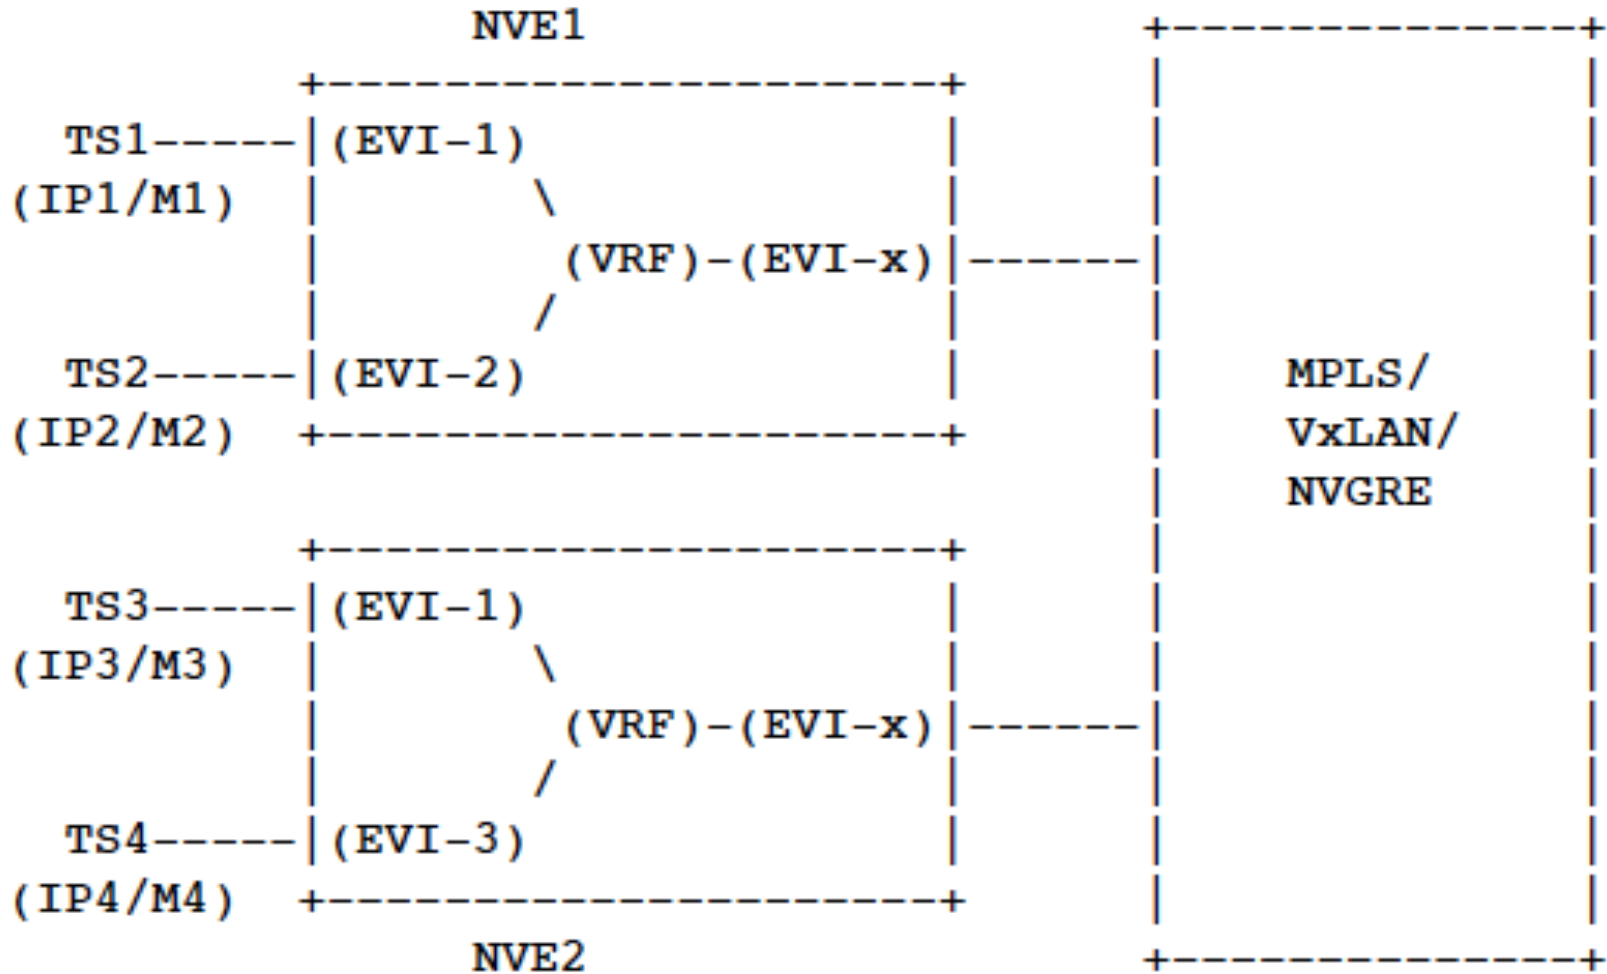
Each RT-2 route is advertised with two RTs (one corresponding to the EVI and the other corresponding to the IP-VPN) and with a new BGP attribute (section 6) that includes the tunnel type and the MAC address of the NVE (e.g., MAC_x for NVE1 or MAC_y for NVE2) .

Control Plane Operation – Cont.

Upon receiving this advertisement, the receiving NVE performs the following:

- It uses Route Targets corresponding to EVI and IP-VPN for importing this route into the corresponding MAC-VRF and IP-VRF tables.
- It imports the MAC address into the MAC-VRF with BGP Next Hop address as underlay tunnel destination address (e.g., VTEP DA for VxLAN encapsulation) and Label-1 as EVI VNID for VxLAN encapsulation or EVPN label for MPLS encapsulation.
- It imports the IP address into IP-VRF with NVE's MAC address (from the new BGP attribute) as inner MAC DA and BGP Next Hop address as underlay tunnel destination address (e.g., VTEP DA for VxLAN encapsulation) and Label-2 as IP-VPN VNID for VxLAN encapsulation or IP-VPN label for MPLS encapsulation.

IRB forwarding on NVEs with core-facing IRB Interface



Control Plane Operation

Each NVE advertises an RT-2 (MAC/IP Advertisement Route) for each of its TSeS and it also advertises a single RT-2 for core-facing IRB interface (which is per tenant or per IP-VPN).

The fields of RT-2 for each TS are set as follow:

- RD and ESI per [EVPN]
- Ethernet Tag = 0; assuming VLAN-based service
- MAC Address Length = 48
- MAC Address = M_i ; MAC address of TS
- IP Address Length = 32 or 128
- IP Address = IP_i ; IP address of TS
- Label-1 = MPLS Label or VNID corresponding to access-facing EVI

Control Plane Operation – Cont.

The fields of RT-2 advertised for core-facing IRB interface, are set as follow. This RT-2 is advertised with an RT corresponding to the core-facing EVI (e.g., EVI-x). This RT-2 is also advertised as a sticky MAC per section 15.2 of [EVPN] in order to ensure mis-configuration is caught quickly.

- RD per [EVPN]
- ESI = 0
- Ethernet Tag = 0
- MAC Address Length = 48
- MAC Address = Ma ; MAC address of core-facing IRB interface
- IP Address Length = 32 or 128
- IP Address = IPa ; IP address of core-facing IRB interface
- Label-1 = MPLS Label or VNID corresponding to core-facing EVI

Next Steps

- There seems to be good consensus behind this draft
- Requesting WG adoption of this draft