

IP Prefix Advertisement in EVPN

draft-rabadan-l2vpn-evpn-prefix-advertisement-00

Jorge Rabadan

Wim Hendericks

Florin Balus

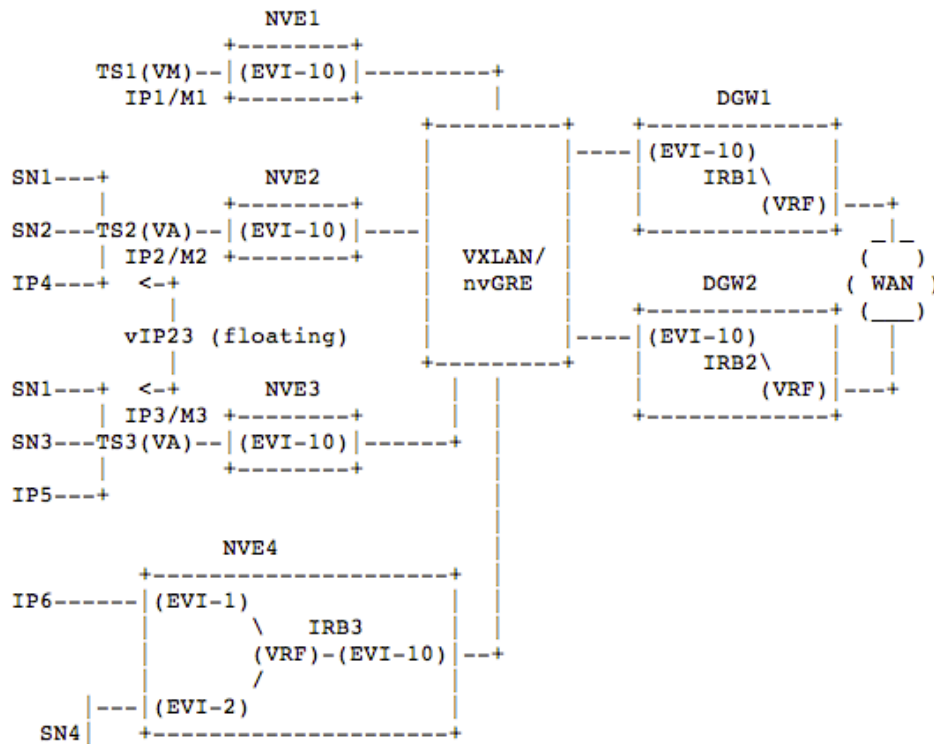
Senad Palislamovic

Aldrin Isaac

IETF 87, July 2013

Berlin, Germany

Problem statement: flexible inter-subnet connectivity is required within the DC



TS = Tenant System

EVI = EVPN instance

VM = Virtual Machine

IRB = Integrated Routing and Bridging interface

VA = Virtual Appliances


NVE = Network Virtualization Edge

SN = Subnet

- Subnets can be located:
 - Behind IRB interfaces
 - Behind VA (Virtual Appliances) such as Firewalls, Load Balancers, NAT routers
- Virtual Appliance characteristics:
 - They do not run dynamic routing protocols
 - They support basic redundancy mechanisms based on a floating IP
 - They can move
- EVPN NVE/CP requirements:
 - Need to advertise IP Prefixes on behalf of the VAs (but not route)
 - IP Prefix next-hop is hosted in the NVE and can be an IRB, an overlay floating IP, a MAC or an ESI

NVEs must advertise prefixes from different overlay next-hops

Solution: The prefix-advertisement route (route-type 5)




Route-types:

- 1- Ethernet A-D route
- 2- MAC-advertisement route
- 3- Inclusive Multicast route
- 4- Ethernet Segment route

5- IP Prefix route (OPTIONAL)

Route-type 5 specific content



- RT 2 (MAC-advertisement) used for advertising MACs and MAC-IP for ARP resolution
 - Also used for MAC Mobility
- RT 5 (IP-prefix) used to advertise Prefixes independently of the mac-advertisement routes
- Possible Prefix next-hops in the overlay topology
 - ESI, MAC address, IRB IP address, Floating IP address

The prefix-advertisement route benefits

- Clean identification of a prefix
- NVEs running EVPN but not supporting the OPTIONAL prefix-advertisement route can easily identify it and ignore it without processing the route
- No MAC information is compared by BGP when selecting routes
- Flexible overlay next-hop (IRB, floating IP, MAC, ESI) addressing different end-point types
- Decouples prefix-advertisement from mac mobility procedures
- Supports VA resiliency procedures

Next steps

- Feedback from WG
- Merge or keep it as separate draft?