BESS work on control planes for DC overlay networks
A short overview

Jorge Rabadan

IETF99, July 2017
Prague

Agenda

EVPN in a nutshell

BESS work on EVPN for NVO3 networks

EVPN in the industry today

Future work for NVO3

EVPN in a nutshell

A unified VPN control plane

- What is RFC7432 EVPN
  - IETF technology that allows multipoint layer-2 VPNs to be operated as RFC4364 IP-VPNs, where MACs and information to setup flooding trees are distributed by MP-BGP
- What were RFC7432's main objectives
  - Replace the old flood-and-learn behavior for BGP-based MAC learning (more control, mac-duplication, mac-protection, mac-mobility)
  - Efficient multi-destination delivery
  - All-active multi-homing
- How has EVPN evolved
  - EVPN is now a unified control plane protocol for E-LAN, E-Line, E-Tree, Layer-3 and Cloud services.
  - Transport agnostic
  - Advanced features

Ethernet Virtual Private Networks
RFC7432 main concepts

MAC-VRF
Virtual Routing Forwarding table for MACs

EVPN Instance (EVI)
Group of MAC-VRFs forming a VPN

Ethernet Tag
Broadcast Domain in the EVI

Single-active multi-homing
- Per-VLAN load-balancing
- Mass-withdraw provides a uniform failover irrespective of the number of MAC-VRFs in the ES

Data/Mgmt Plane Learning
Dynamic or Static (Provisioned)

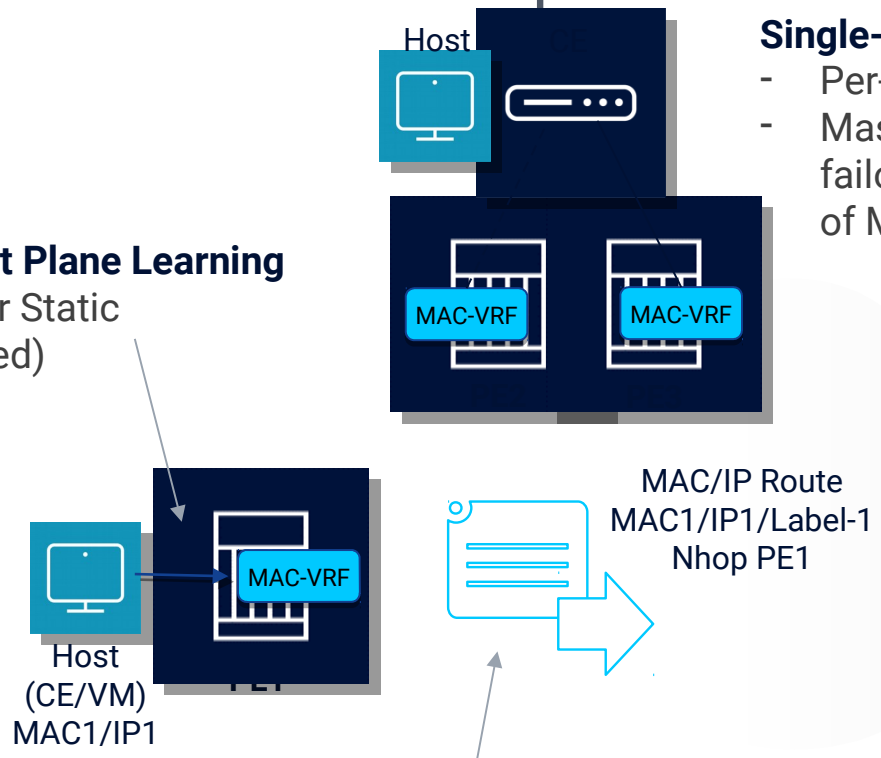

All-active multi-homing
- CE does per-flow load-balance to all PEs (LAG)
- Remote PEs do per-flow load-balance to the ES PEs (aliasing)
- Only the DF can send BUM to the CE
- Split-Horizon avoids echo'ed packets

Customer Edge (CE)
Host, VM, Router or Switch

Control Plane Learning
PEs Advertise MAC Addresses and Next Hops From Connected CEs Using MP-BGP

Data Plane Encapsulation
MPLS tunnels (Incl P2MP)
PBB over MPLS tunnels
NVO tunnels
(Encapsulation attribute indicates desired tunnel)

Ethernet Segment Identifier (ESI)
Link(s) that Connect the CE to PEs (ESIs are Unique Across the Network)

EVPN work in IETF

EVPN Application/Service	Standard document
E-LAN	RFC7432 (EVPN) RFC7623 (PBB-EVPN)
E-Line	draft-ietf-bess-evpn-vpws
E-Tree	draft-ietf-bess-evpn-etree
L3 VPN (Inter-subnet-forwarding)	draft-ietf-bess-evpn-inter-subnet-forwarding draft-ietf-bess-evpn-prefix-advertisement
EVPN for DC	draft-ietf-bess-evpn-overlay draft-ietf-bess-evpn-optimized-ir
EVPN for DCI	draft-ietf-bess-dci-evpn-overlay
Other applications or enhancements: - Multi-homing improvements - Proxy-ARP/ND and security - BUM optimizations - IP Multicast optimizations L2/L3 - EVPN <-> IPVPN integration - EVPN <-> VPLS integration	draft-ietf-bess-evpn-vpls-seamless-integ draft-ietf-bess-evpn-df-election draft-ietf-bess-evpn-ac-df draft-ietf-bess-evpn-pref-df draft-ietf-bess-evpn-proxy-arp-nd draft-ietf-bess-evpn-bum-procedure-updates draft-ietf-bess-evpn-igmp-mld-proxy

Current DC Overlay Network Deployments

Plus around 20 individual EVPN-related drafts !!

Why is EVPN used in DC Overlay Networks?

draft-ietf-bess-evpn-overlay

Modern DCs are based on:

- CLOS architecture and IP Fabrics
  - No loops, no flooding, fast convergence
  - ECMP
- Multi-tenancy with intra (L2) and inter-subnet (L3) connectivity
- IP Overlay tunnels are therefore needed (VXLAN is the most popular option)

Why do I need a control plane?

- Auto-discovery of the remote VTEPs
- Distribution of MAC/IP information in order to reduce/suppress flooding
- Other advanced options

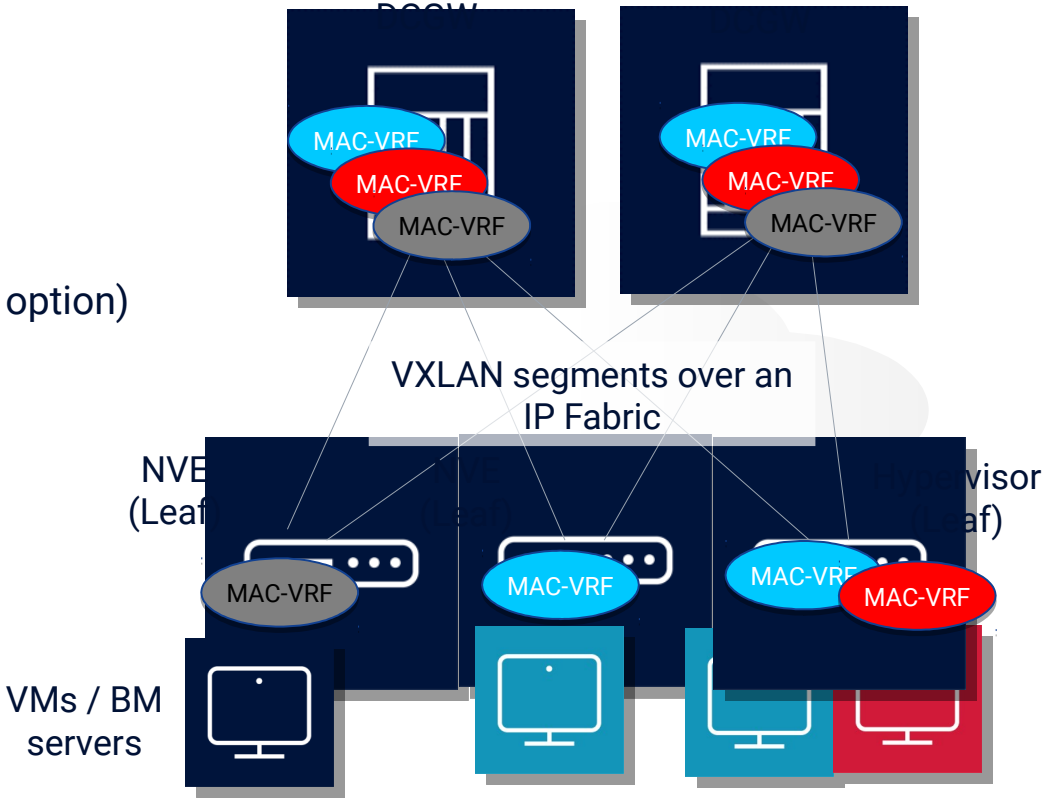

EVPN for NVO tunnels
Endless possibilities

EVPN provides the basic Control Plane needs…

- Auto-discovery of remote VTEPs through Inclusive Multicast Routes (IMET)
- Distribution of MAC/IP information in order to reduce/suppress flooding through MAC/IP routes

But also advanced options

- All-active multi-homing
- MAC/IP mobility
- MAC protection, duplication detection, loop protection
- Proxy-ARP/ND, IGMP and PIM proxy, to reduce/suppress the flooding in the DC
- Assisted-Replication
- Distribution of IP-Prefix/host routes and inter-subnet-forwarding so that not all the subnets must be defined in all the NVEs
- Etc

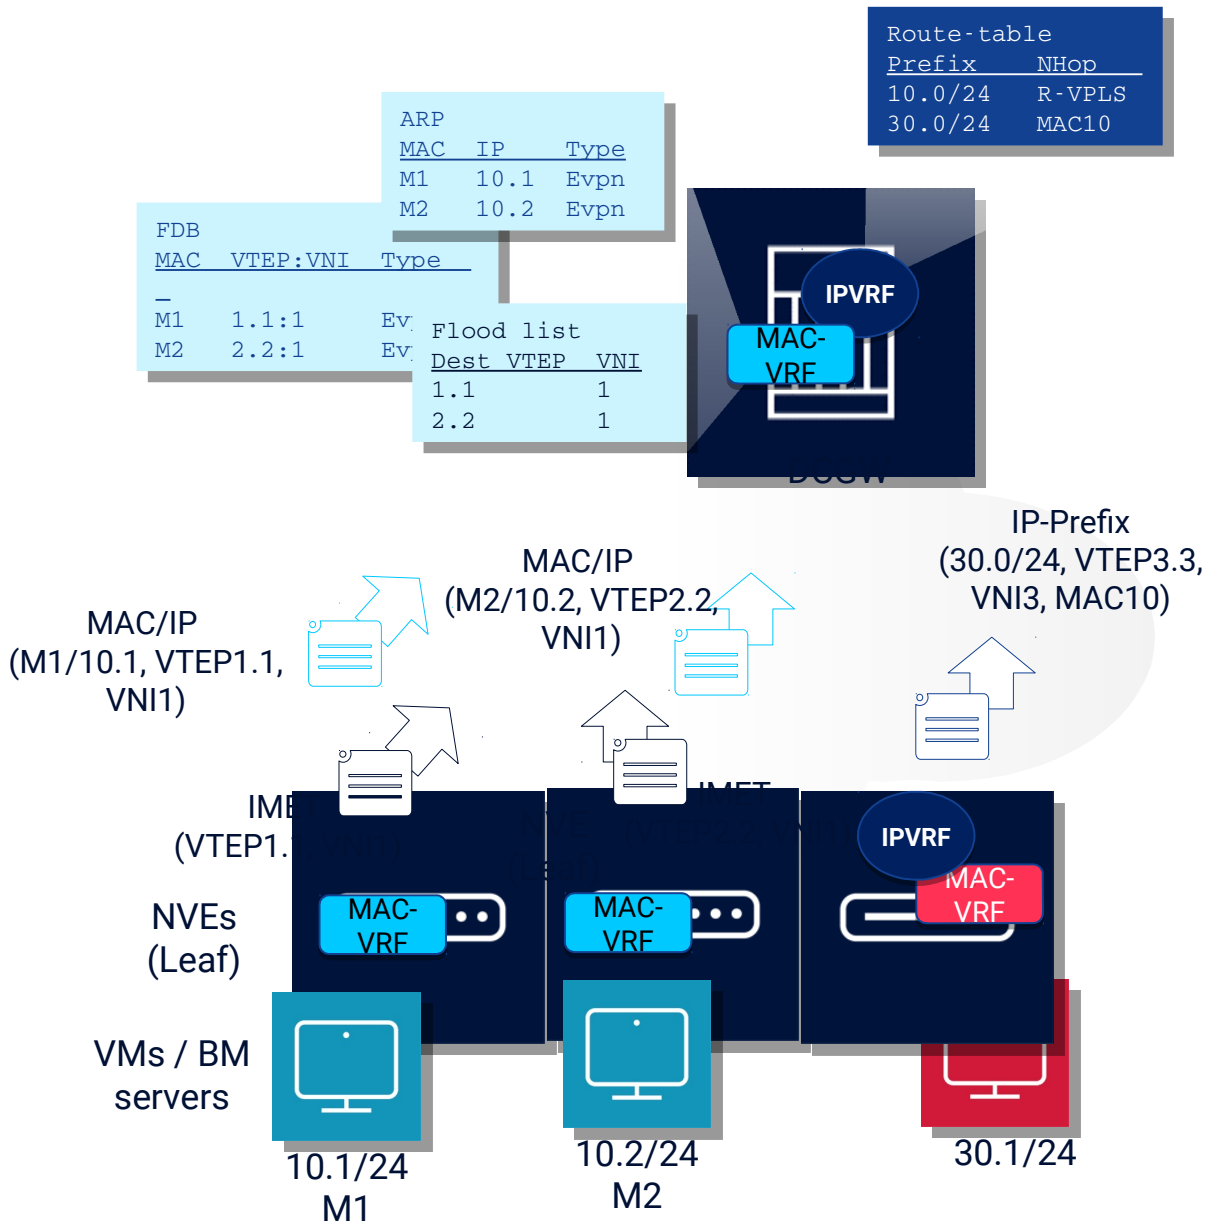

EVPN for VXLAN tunnels is widely deployed today in Data Centers

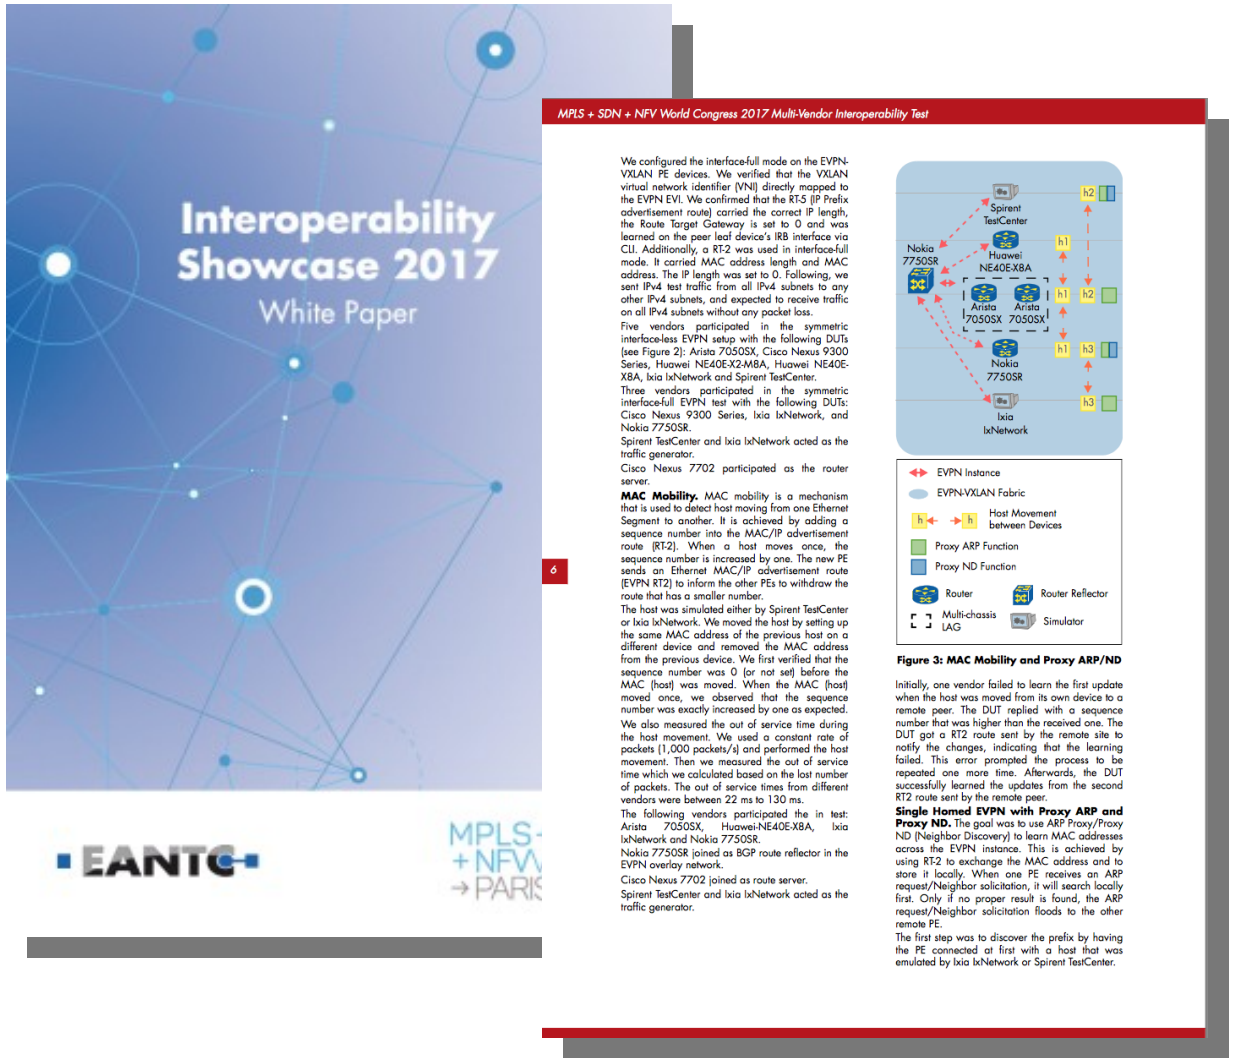

Multiple vendor implementations

- Hardware based NVEs
- Software based NVEs

Interoperability publicly demonstrated

www.eantc.de/en/showcases/mpls_sdn_2017

EVPN new/future work for NVO3

- Support new NVO3 data encapsulations
  - First attempt for Geneve - draft-boutros-bess-evpn-geneve-00
- Tunnel options/extensions negotiation
- Future extensions

Thank you