

draft-boutros-bess-evpn-geneve-02

Sami Boutros (Vmware)

Ali Sajassi (Cisco)

John Drake (Juniper)

Jorge Rabadan (Nokia)

IETF101, Mar 2018

London

Generic Network Virtualization Encapsulation (GENEVE) EVPN as control plane for GENEVE

- Based on draft-ietf-bess-evpn-overlay, this spec captures the EVPN specifics for GENEVE
- Presented in IETF 100, primarily focused on:
 - Signaling GENEVE encapsulation type in draft-ietf-idr-tunnel-encaps extended community or tunnel-encapsulation attribute
 - Communicating GENEVE tunnel option types (to the ingress NVE) in a new BGP Tunnel Encapsulation Attribute sub-TLV
- New in revision 02:
 - **NEW** Ethernet Option TLV
 - **CLARIFIES** GENEVE Tunnel Option Types

Ethernet Option TLV

Signaled by the egress NVE based on its capabilities

ETHERNET OPTION TLV

0	16	24			
Option Class=ETHERNET	Type=EVPN	B	L	R	Len=0x0

0	16	24			
Option Class=ETHERNET	Type=EVPN	B	L	R	Len=0x1
Rsvd	Source-ID				

Fields

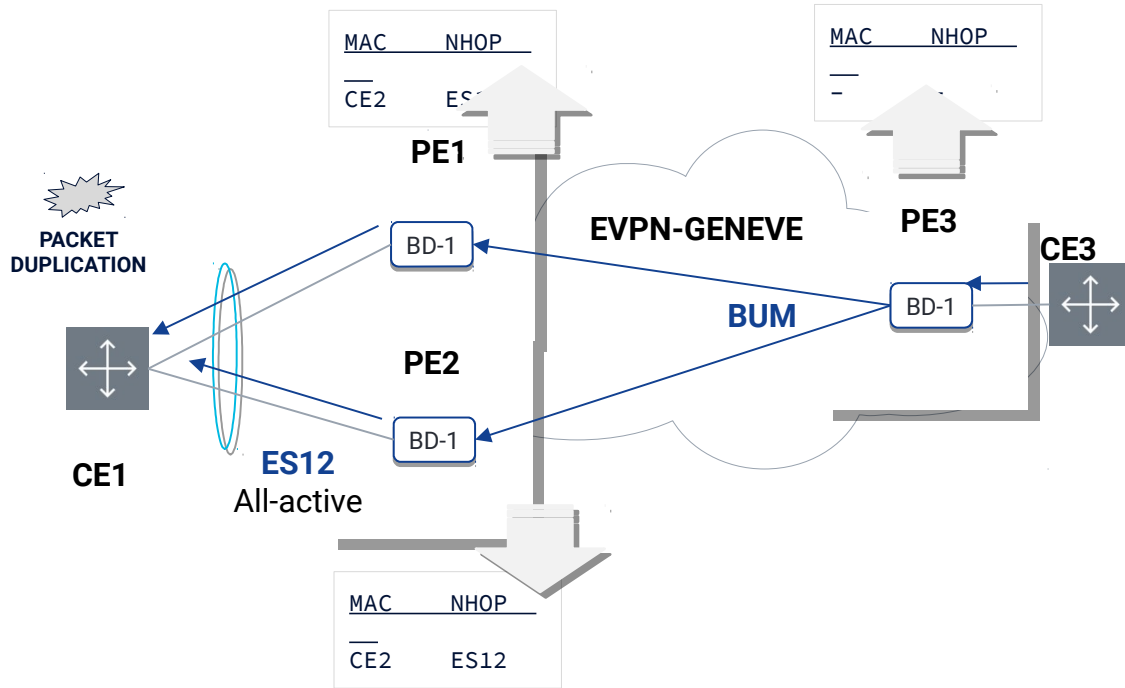
- Option Class Ethernet is a new class (requested to IANA)
- Type set to EVPN Option (requested to IANA)
- B == BUM traffic indication
- L == Leaf traffic indication
- Source-ID == encodes an **optional** source Ethernet Segment identifier that can be used for:
 - Multi-homing Split-horizon as in RFC7432
 - E-Tree as in RFC8317

Two possible lengths (4-bytes or 8-bytes)

- If no source-ID is needed, length is 4-bytes
- Else length is 8-bytes

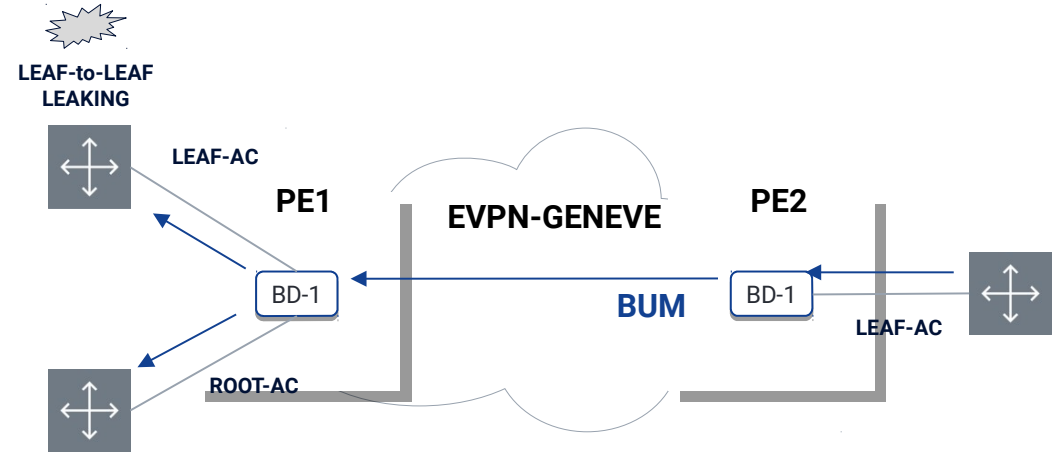
The B and the L bits

Why and when are they used



B-bit avoids transient packet duplication

When Ingress Replication is used PE3 needs to set B=1
PE1 and PE2 identify the packets as multicast and only the DF will forward.

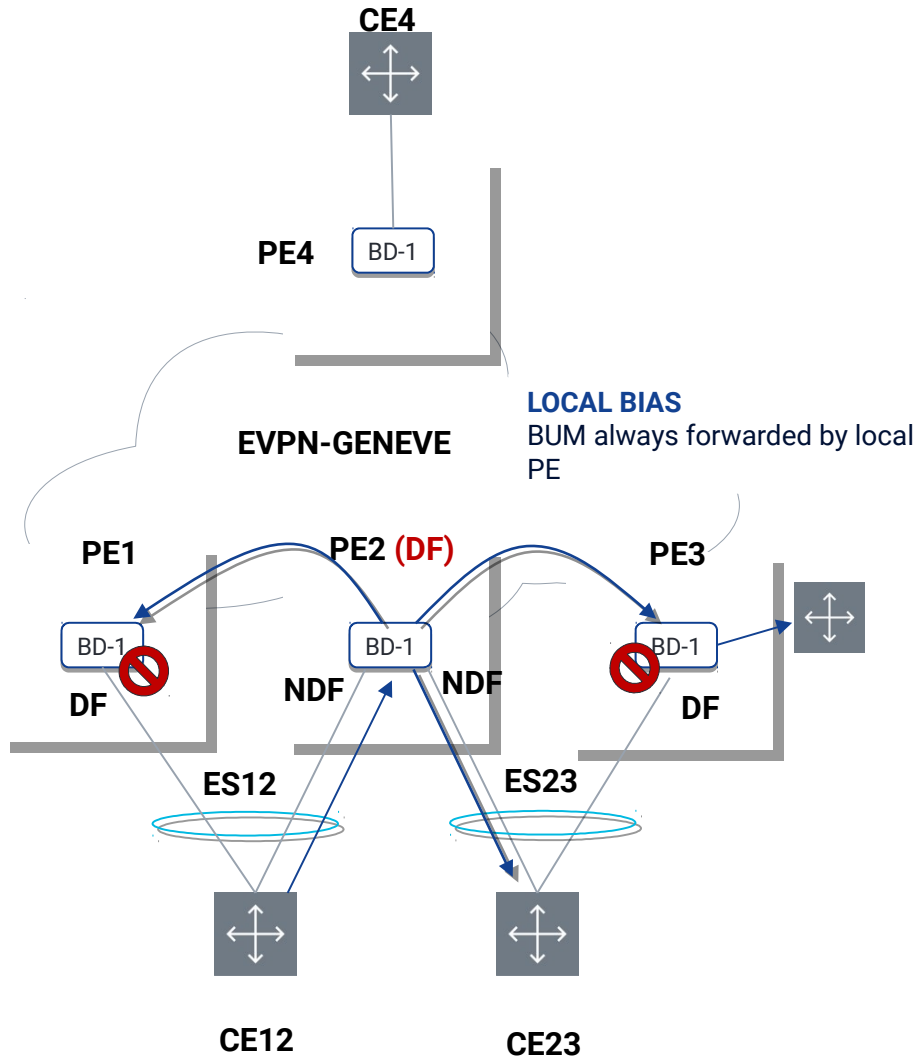


L-bit avoids BUM leaf-to-leaf leaking in EVPN E-Tree

When PE2 receives BUM from a leaf-ac, it needs to set L=1
PE1 identifies the packets as leaf-originated and will filter appropriately

Optimizing consistency with EVPN MPLS procedures

EVPN-GENEVE and Multi-Homing Split-Horizon



- **Split-Horizon in EVPN-GENEVE:**
 - [EVPN-OVERLAY] defines LOCAL-BIAS for multi-homing split-horizon
 - LOCAL-BIAS is mandatory for GENEVE too.
 - Source-ID based split-horizon is OPTIONAL
- **When should I use source-ID:**
 - Consistency with MPLS split-horizon procedures, e.g. the BD has a mix of GENEVE and MPLS PEs.

Conclusions and next steps

- Document defines extensions to close the gap between MPLS encaps and GENEVE
- NEXT STEPS
 - Clarify the use of the Ethernet Option TLV further
- Please provide feedback / comments