

Domain Path (D-PATH) for Ethernet VPN (EVPN) Interconnect Networks

draft-sr-bess-evpn-dpath-00

Jorge Rabadan (Nokia)
Senthil Sathappan (Nokia)

IETF112, Nov 2021

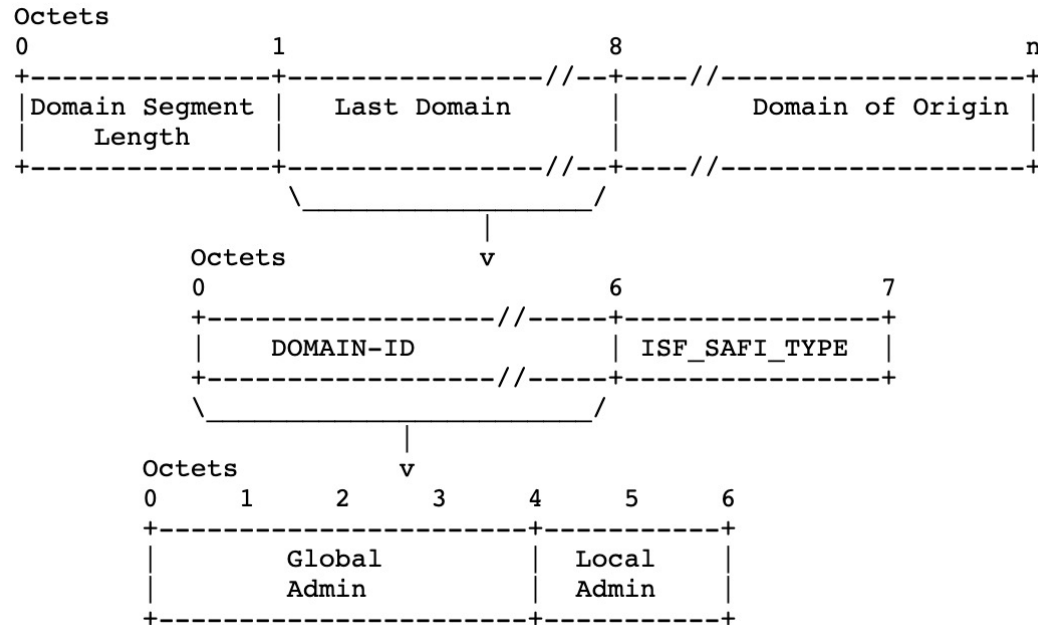
Online

Agenda

1. D-PATH Attribute – Refresh
2. Extension of D-PATH for EVPN Layer-2 routes
3. Next-steps

D-PATH Attribute Refresh

I-D.ietf-bess-evpn-ipvpn-interworking



Domain PATH Attribute (D-PATH) definition

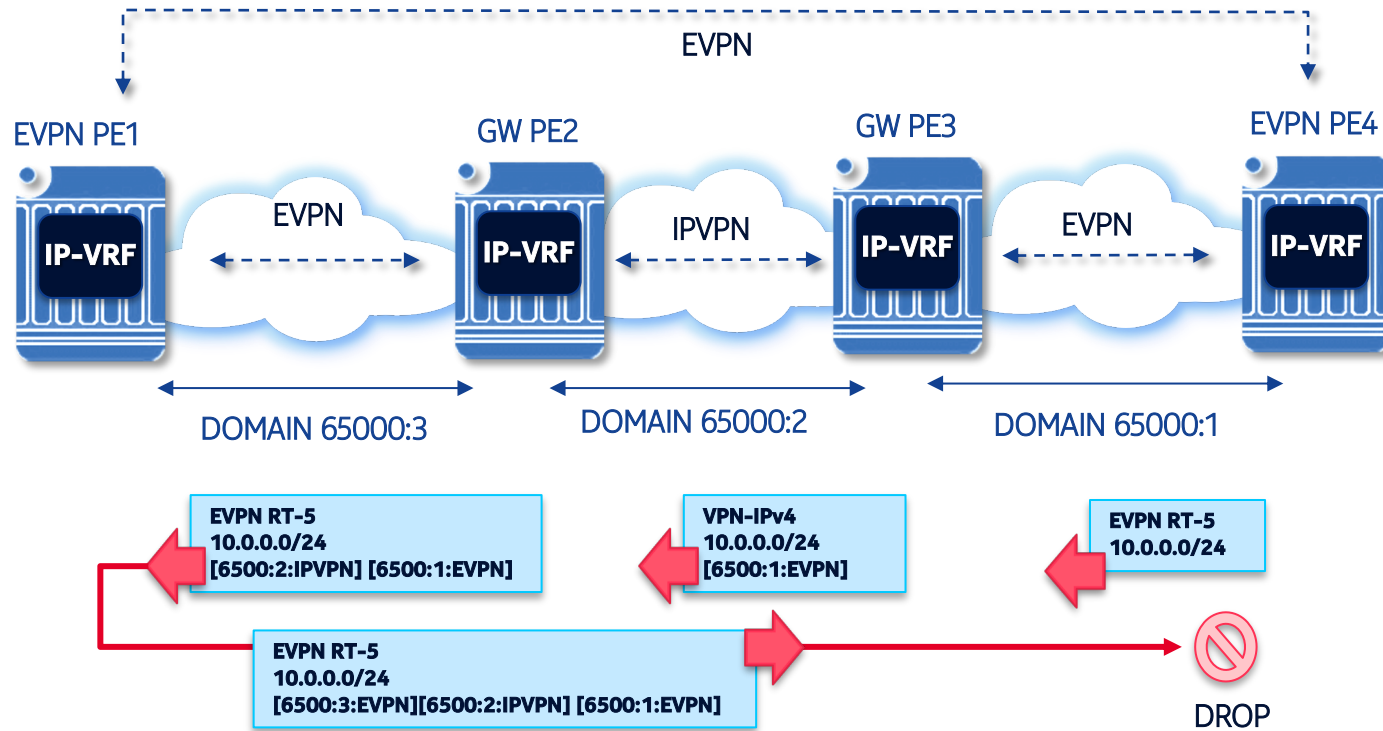
- Optional, transitive
- Composed of a sequence of domain segments
- Each domain segment defined by length and a sequence of Domains
- Domain is represented by <DOMAIN-ID:ISF_SAFI_TYPE>, where the ISF_SAFI_TYPE can be evpn, ip, vpn-ip, zero
- Each service Gateway with an IP-VRF connecting two domains appends the <DOMAIN-ID:ISF_SAFI_TYPE> of the domain of origin before re-advertising into another domain

D-PATH impacts on BGP procedures

- Control Plane loop protection
- Best path selection

D-PATH Attribute Refresh

Use-case examples



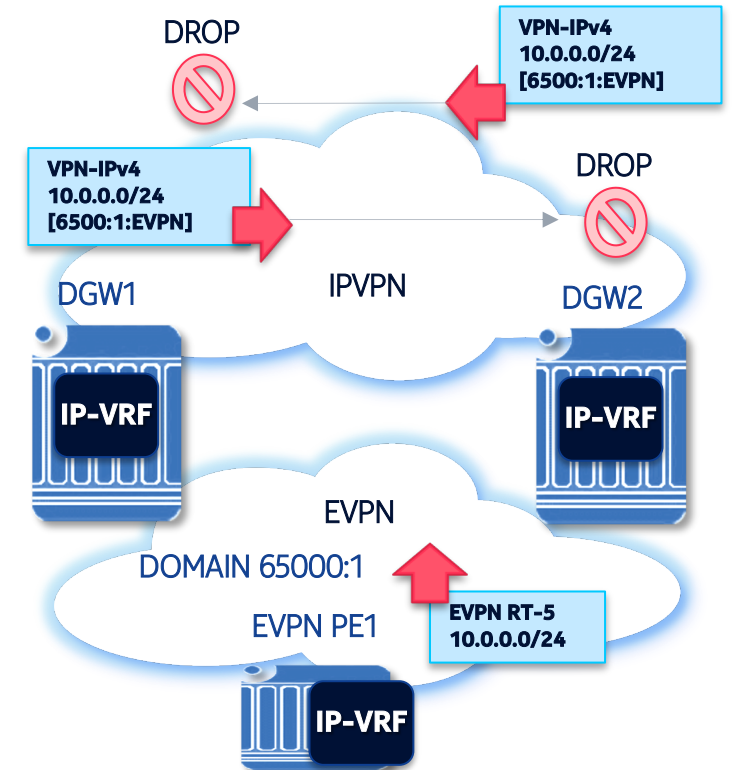
Use-Case 1

Multi-Domain propagation of routes

Traceability, loop protection, best path selection

Use-Case 2

DCGW loop protection



D-PATH used with non-Inter-Subnet-Forwarding EVPN

MAC/IP Advertisement routes

When used with MAC/IP routes D-PATH is a sequence of Layer-2 Domains

- DOMAIN-ID is a Layer2-Domain identifier configured in a MAC-VRF and SAFI type is set to either 70 (EVPN) or 0 (local route).
- D-PATH identifies the sequence of Layer2-Domains the route has gone through
- Added/modified by a Layer2-Domain Gateway PE that re-advertises the route and may be added by a PE that originates the route

Use Cases

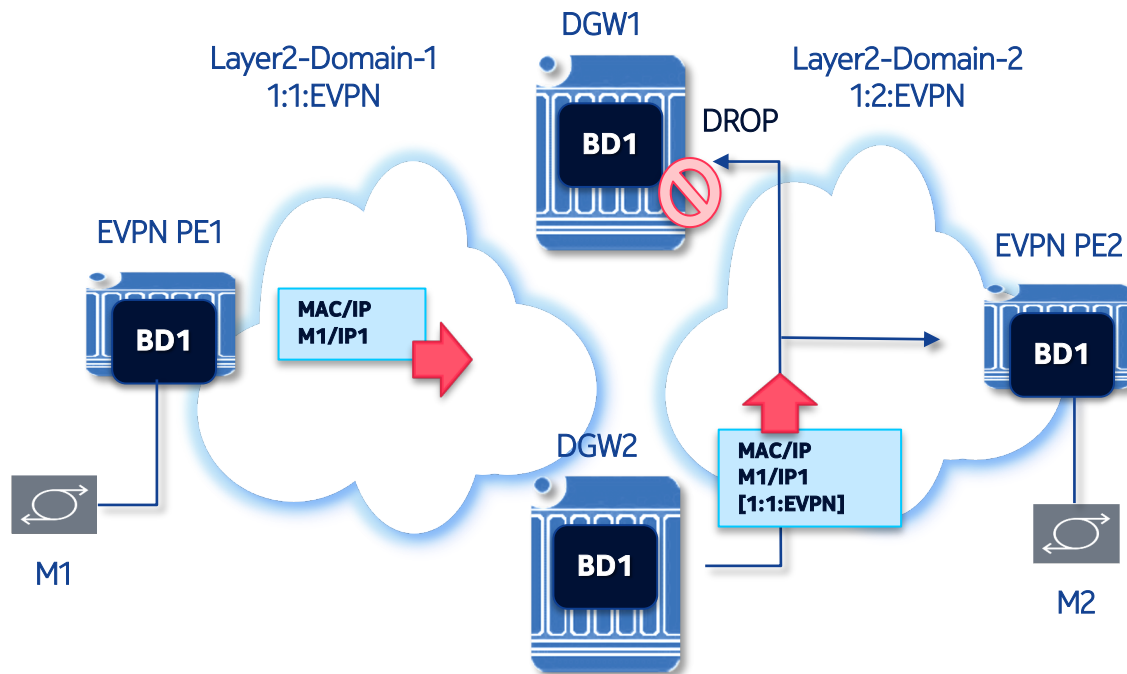
- Control Plane loop detection/protection on RFC9014 Inter-connect gateways
- MAC/IP Advertisement route traceability
- Best path selection

Best Path Selection

- Section 4.2 summarizes the current best path selection for MAC/IP routes including the impact of D-PATH in such selection

Examples in EVPN Interconnect Networks (RFC9014)

Loop Protection for re-advertised EVPN routes



Loop Protection on the DCGWs

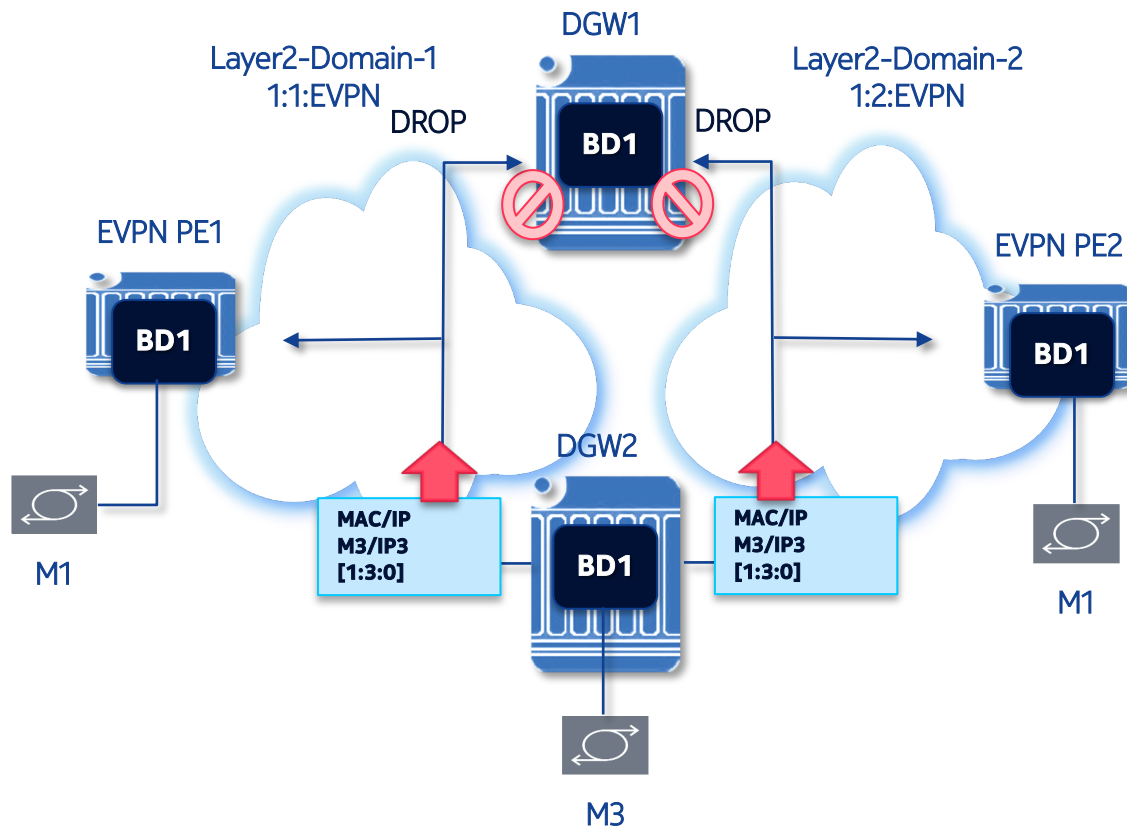
- DGW1 and DGW2 can compare the D-PATH of the incoming routes with their local list of Layer2-Domain-IDs, and detect a loop if any of the local Layer2-Domain-IDs matches a domain in the received D-PATH.
- This procedure prevents the re-advertisement of the route back into Layer2-Domain-1.

Traceability on PE2

- PE2 has the visibility of the Layer2-Domains through which the route has gone, and
- PE2 can also use the D-PATH for best path selection in case PE2 receives a MAC/IP Advertisement route for M1/IP1 by some other means.

Examples in EVPN Interconnect Networks (RFC9014)

Loop Protection for DCGW local routes



Loop Protection on the DCGWs for local routes with shared local domain-id (shown in diagram)

- DGW2 advertises local M3/IP3 with a local domain-id, i.e. 1:3
- If local domain 1:3 is configured on both DCGWs, DGW1 will identify 1:3 as local and will declare the route as looped

Loop Protection with non-shared local domain-ids (not shown)

- If DGW2 and DGW1 do not share local domain-ids, DGW1's MAC3/IP3 route will NOT be dropped at DGW1
- DGW1 will select one (lowest domain-id) and readvertise M3/IP3 into domain-2 with {1:1:EVPN, 1:3:0}
- The route from DGW1 will be:
 - Detected as looped in DGW2
 - Not selected by PE2 vs the route from DGW2, due to the D-PATH length in best path selection

Next steps

Authors seeking feedback

Thank you