EVPN Fast Reroute

draft-burdet-bess-evpn-fast-reroute

L. A. Burdet
P. Brissette
Cisco
T. Miyasaka
KDDI Corporation
J. Rabadan
Nokia

draft-burdet-bess-evpn-fast-reroute

- EVPN convergence is control plane driven
  - Route-scale and topology dependant (route-reflector delays, network delays etc.)
- Using peer's service label (straightforward solution) :
  - May lead to loops in the network or extra dataplane work
  - Does not work for all EVPN load-balancing modes (A/A, S/A, P/A, SF/A)
- EVPN convergence cannot reasonably be expected to meet more stringent requirements (sub-second)
- Support for SRV6 underlay transport recently added:
  - new End.DT2U.Reroute and End.DX2.Reroute

No further reroute : Terminal disposition

- Downstream-allocated Service-Label, Reroute-Label
- Standard Service-Label disposition:
  AC state-based forwarding chain
  - AC-Up: send to CE
  - AC-Down: re-encap with peer's Reroute-Label

Failure

- Disposition of Reroute-Label is terminal (final)
  - Local knowledge. Applies regardless of AC state
  - Once-rerouted packets are not rerouted again
- No dataplane modifications required
  - No extra label on the stack to prevent further reroutes
  - Reroute-Label replaces Service-Label

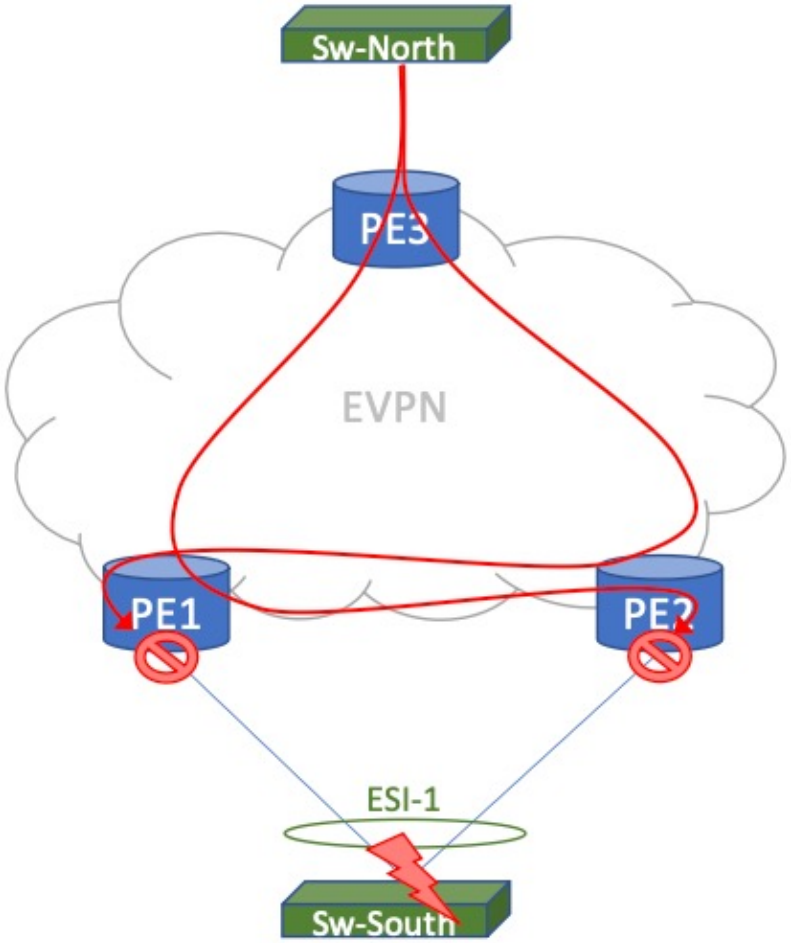

Bypassing Designated-Forwarder Election

- PE1's AC down: drop until Control plane DF-Elects PE2, and PE3->PE2 forwarding created
- Applying a Reroute at PE1 results in redirection to PE2 which is still NDF and drops.

- DF-Election Bypass on Reroute-Label
  - Traffic is rerouted much faster than EVPN DF-Election may unblock ports (control plane)
  - DF-Election results do not apply to Reroute-Label traffic (unidirectional bypass)
  - Especially applicable to **Single-Active load-balancing**

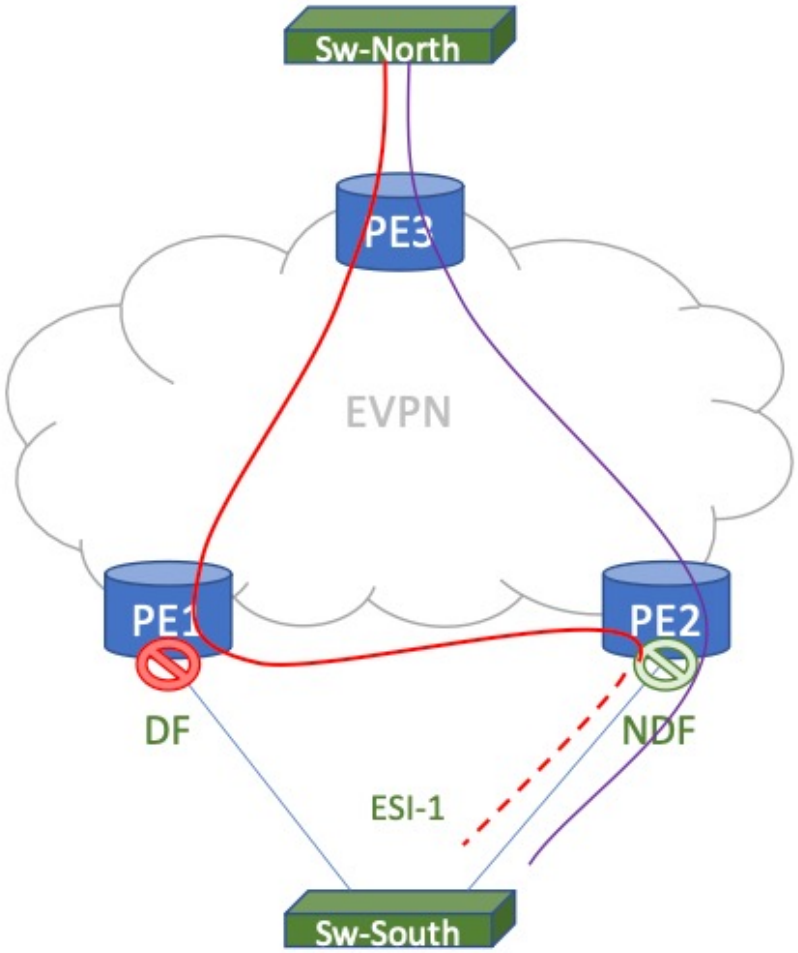

SRv6 Support

New *End.DT2U.Reroute* and *End.DX2.Reroute* signaled alongside base behaviour

```
BGP Prefix SID Attr:
  SRv6 L2 Service TLV:
    SRv6 SID Information sub-TLV:
      SID: 2001:db8:b:1:fbd1::
      Behavior: End.DT2U
      SRv6 SID Structure sub-sub-TLV:
        LBL: 48, LNL: 16, FL: 16, AL: 0, TPOS-L: 0, TPOS-O: 0
    SRv6 SID Information sub-TLV:
      SID: 2001:db8:b:1:fbd1:aaaa::
      Behavior: End.DT2U.Reroute
      SRv6 SID Structure sub-sub-TLV:
        LBL: 48, LNL: 16, FL: 16, AL: 16, TPOS-L: 0, TPOS-O: 0
```

SRv6 Support

- Backwards-compatibility is specifically addressed
  - Priority between 2 SIDs, Applicability at Node role, Unrecognized SIDs
- Supports per-Attachment Circuit granularity redirection with Arg
- Reachability (IGP) via single advertisement of the base End.DT2U or End.DX2 SIDs
  - Base End.DT2U / End.DX2 can be extracted when Arg=0 is forced

- Important draft representing easy way to achieve fast convergence for EVPN Services

Seeking Spring WG feedback on SRV6 procedures