

draft-burdet-bess-evpn-fast-reroute-00


- 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 load-balancing modes
- EVPN convergence cannot reasonably be expected to meet more stringent requirements (sub-second)

Desired outcome

- Local detection at PE1
- Local restoration at PE1
- *Then* control-plane propagation in EVPN to restore optimal routing/forwarding PE3 ↔ PE2


Draft addresses rerouting EVPN traffic and in particular issues with using peer's service label:

- Continuous reroutes at simultaneous failures
- Rerouting to an NDF interface



Continuous reroutes for simultaneous failures

- PE1 and PE2 detection in same time window
Common for EVPN All-active with failure at CE-side of LAG
- Using regular EVPN Service-Label for redirection
- PE1 and PE2 reroute to each other continuously:
 - traffic storm
 - congestion on PE1-PE2 core links.

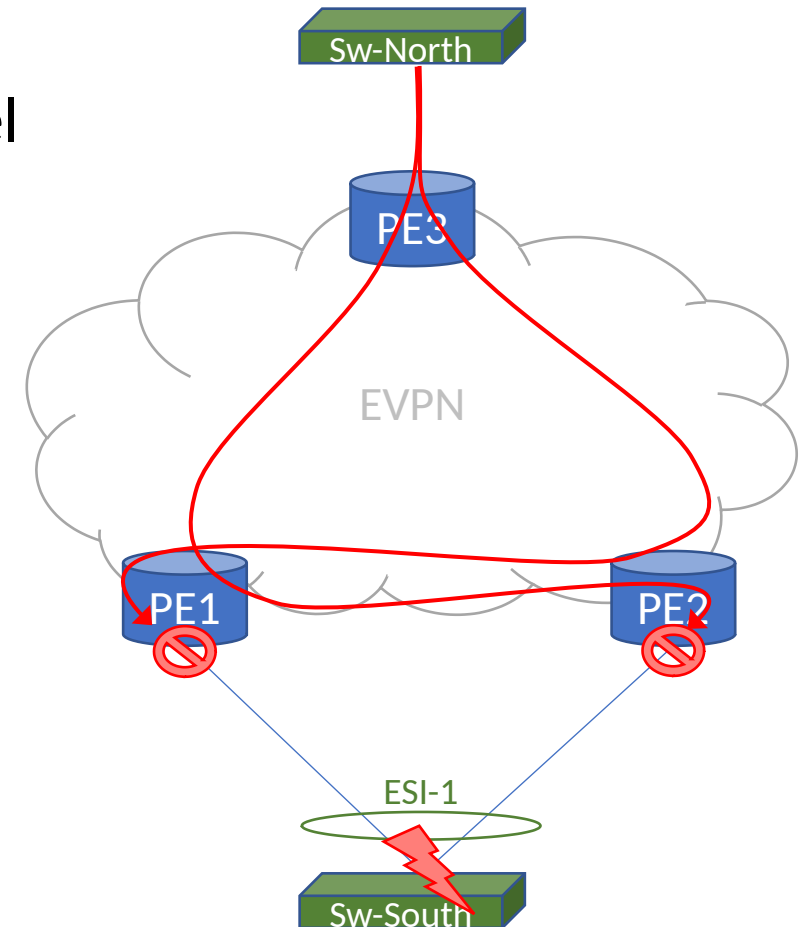


Reroute-Label : Terminal disposition

- 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



Reroute-Label : Bypassing DF-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**



draft-burdet-bess-evpn-fast-reroute-00

Implementation Results

- Sub-50msec traffic loss for AC (edge) failures
- Applicable to most load-balancing modes
- Complements other procedures
 - Orthogonal to draft-ietf-bess-fast-recovery

- num_of_flows:128000
- num_of_BD:14
- mac_per_device:20
- number_of_mac:560
- flows_per_stream:457

TRIGGER NAME	BEST N-S FLOW	WORST N-S FLOW	AVG N-S FLOW	BEST S-N FLOW	WORST S-N FLOW	AVG S-N FLOW
Access Shut Recovery	2424	2424	2424	2427	2427	2427
Access Shut	17	17	17	52	52	52

(milliseconds)

For recovery improvement, see:
[draft-ietf-bess-evpn-fast-df-recovery](#)