

Proxy IP->MAC Advertisements in EVPNs

draft-rbickhart-evpn-ip-mac-proxy-adv-00

Ryan Bickhart (Juniper)

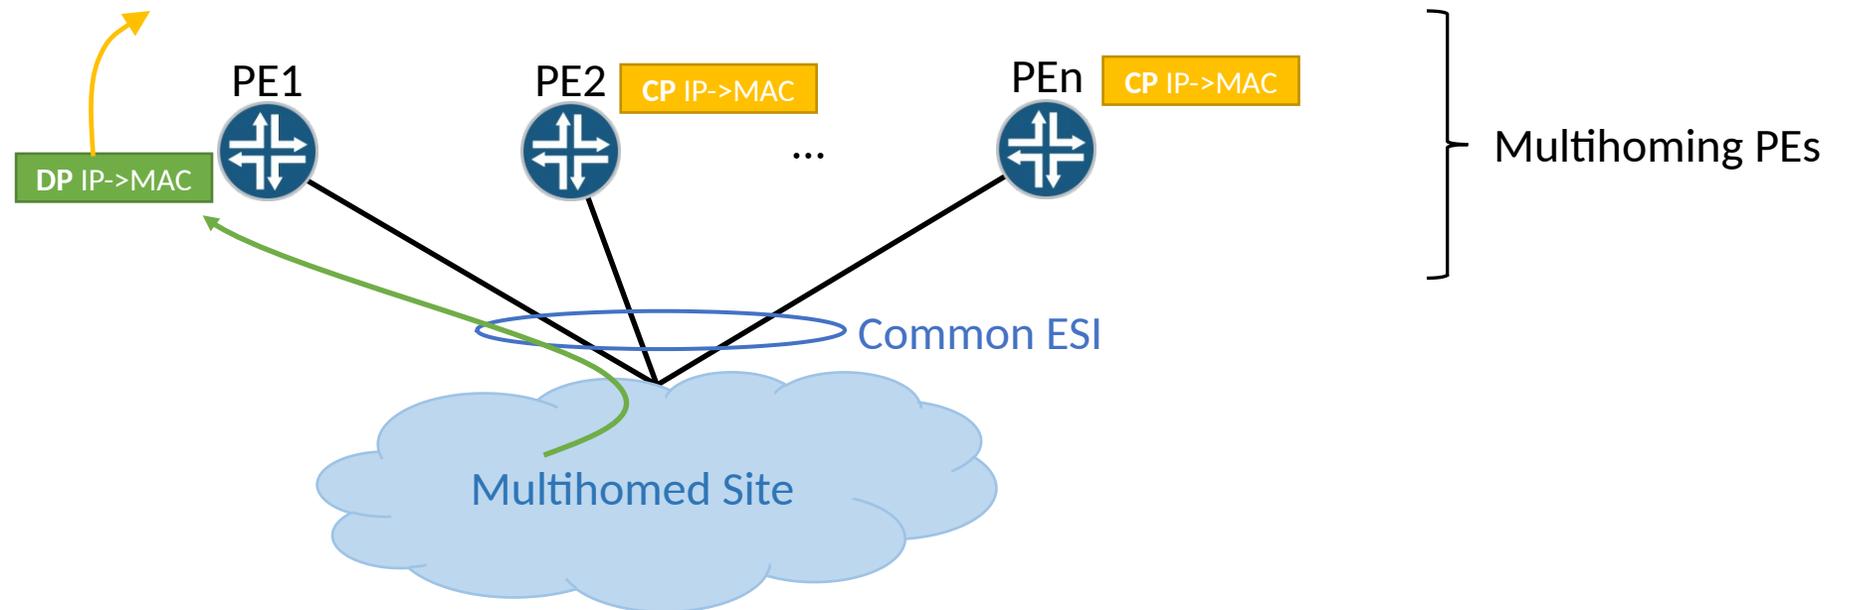
Wen Lin (Juniper)

John Drake (Juniper)

Jorge Rabadan (Nokia)

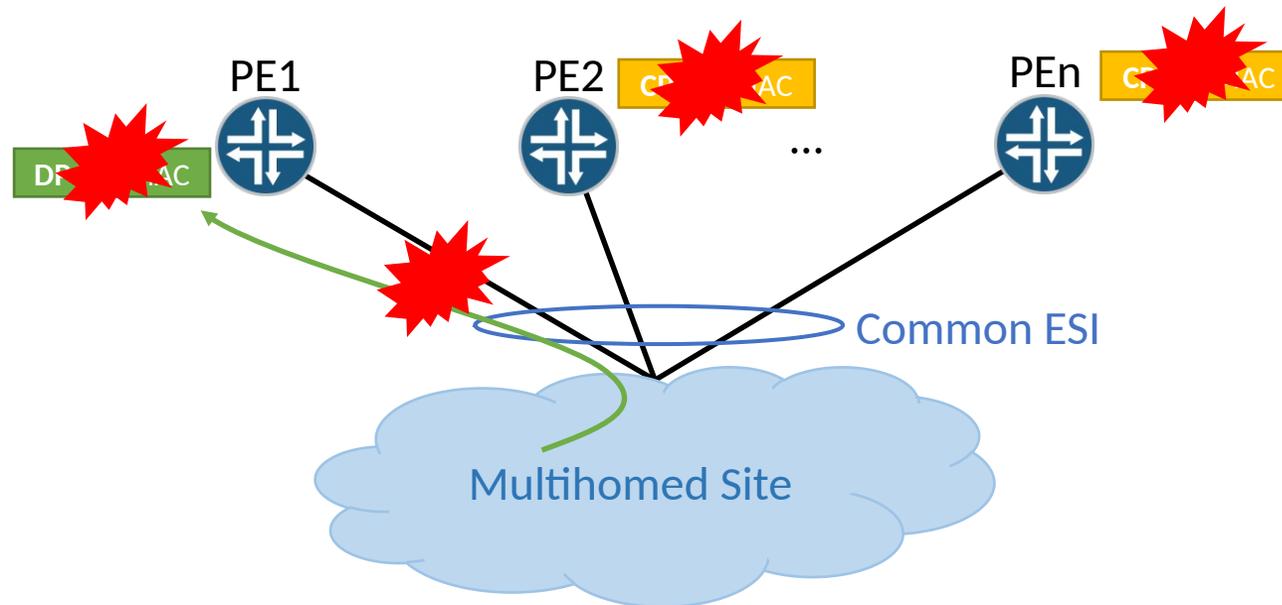
Problem

- IP->MAC bindings of hosts connected to multihomed sites in EVPNs may be **locally learned in the data plane** only on a subset of attached multihoming PEs
- Other PEs learn the IP->MAC through the **control plane** only



Problem

- If PE1 or the link from PE1 to the multihomed site fails, all traces of the IP->MAC binding are lost from the EVPN
- However, the MAC and IP are still reachable through PE2 ... PEn



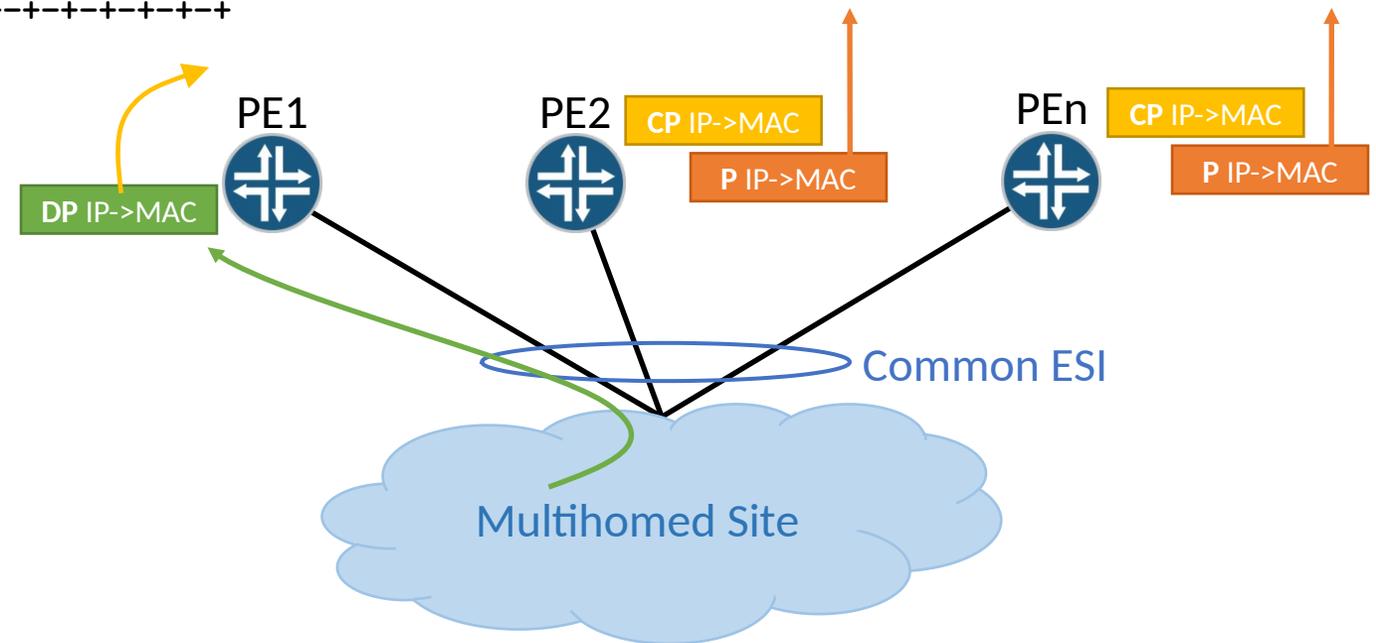
Solution – MH PE Behavior

- Upon receiving the control plane learned IP->MAC, PE2 ... PEn also originate their own advertisements for type-2 route: IP->MAC **marked with a proxy indication**
- EVPN ARP/ND Extended Community:

```

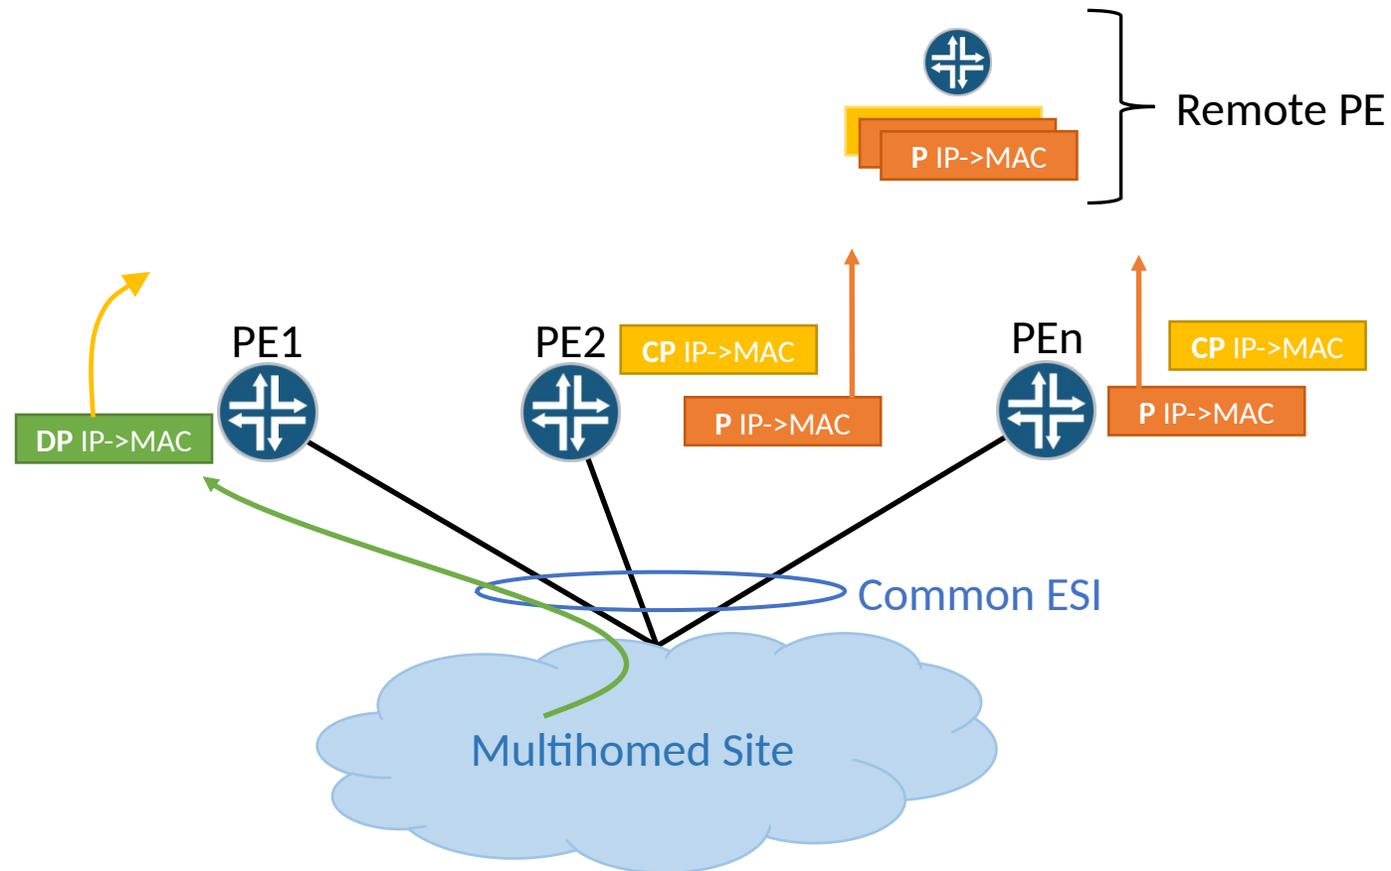
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type=0x06   | Sub-Type=0x08 |Reserved |P|O|R|   Reserved=0   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|                               Reserved=0                               |
|                                     |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
  
```

Bit Name: P -> Proxy IP->MAC advertisement



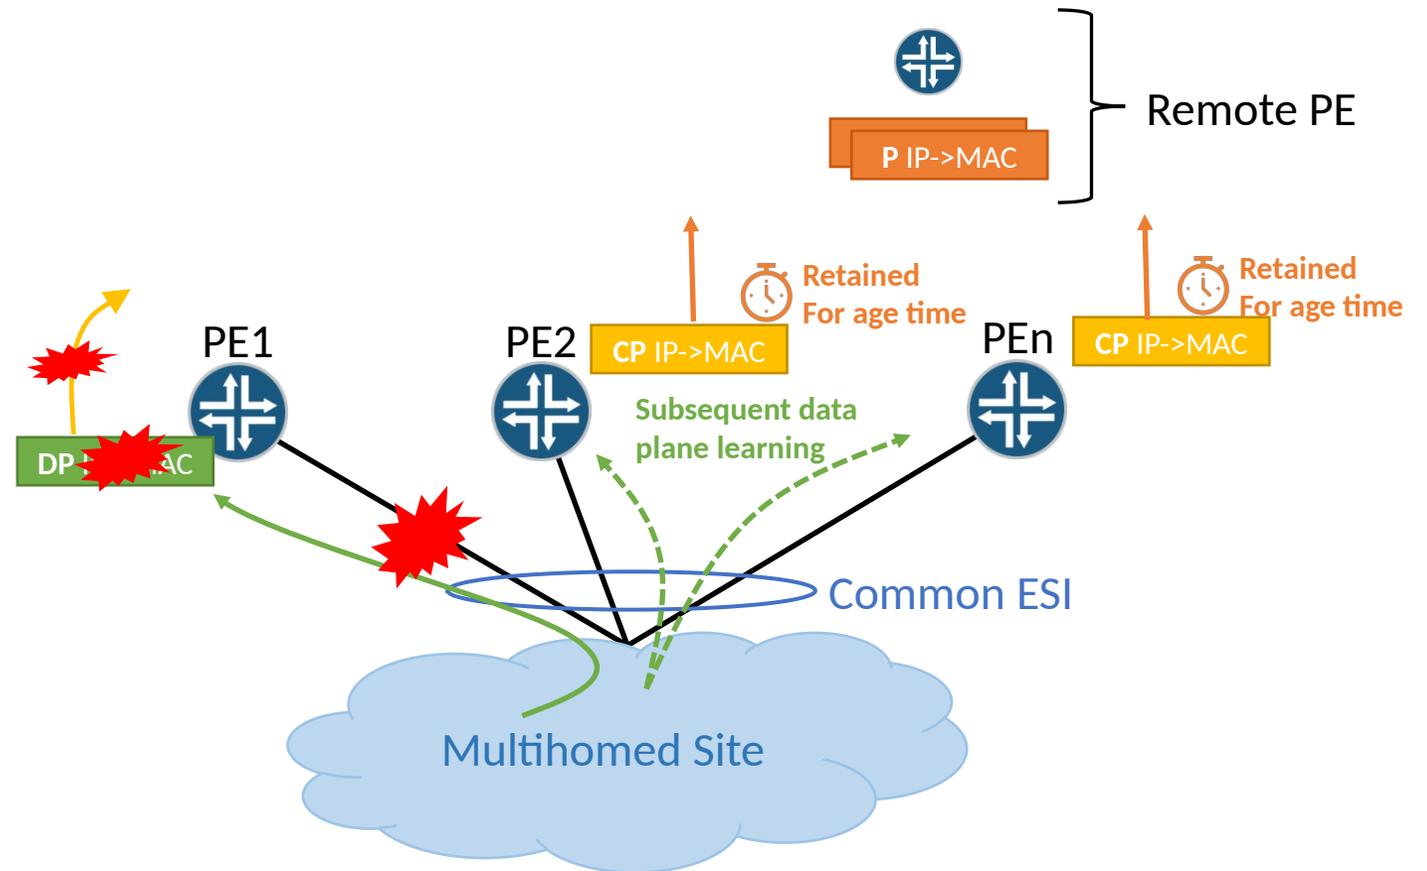
Solution – Remote PE Behavior

- Remote PEs receive the proxy IP->MACs and give them standard treatment because the remote PEs are not connected to the IP->MAC's ESI



Solution - Failure Handling

- If all non-proxy IP->MAC sources are withdrawn, a PE originating proxy IP->MAC keeps advertising it, but starts aging timer
- If IP->MAC learned from non-proxy source (DP or CP) before aging timer expires, aging timer cancelled
- When aging timer expires, proxy IP->MAC advertisement is withdrawn



Conclusion and Next Steps

- Close the gap between the link / node failure and the subsequent relearning of the IP->MAC on one of the remaining multihoming PEs.
- Avoid unnecessary churn in the network

- Would like to seek feedback from WG