EVPN VPWS Flexible Cross-Connect Service

draft-ietf-bess-evpn-vpws-fxc-02

- Updates since version–01

- Many editorials and cross-ref fixes for consistency.
  - Clarification of Normalised-VID and VPWS service identifiers

- Cleanup extension to *EVPN Layer-2 attribute* extcomm
  - Modified fields alignment
  - Add IANA-request section

EVPN VPWS Flexible Cross-Connect Service

draft-ietf-bess-evpn-vpws-fxc-02

- Clarification of Mode (M) and VID-normalisation (V) fields
  - More precise language for Error conditions
    (Alarming vs. preventing tunnel instantiation)
  - Optional vs. Required behaviour clarified
  - New Section: Service instantiation upon reception of mismatched V-

- Next steps
  - Implemented draft
  - Take any comments and proceed to WGLC

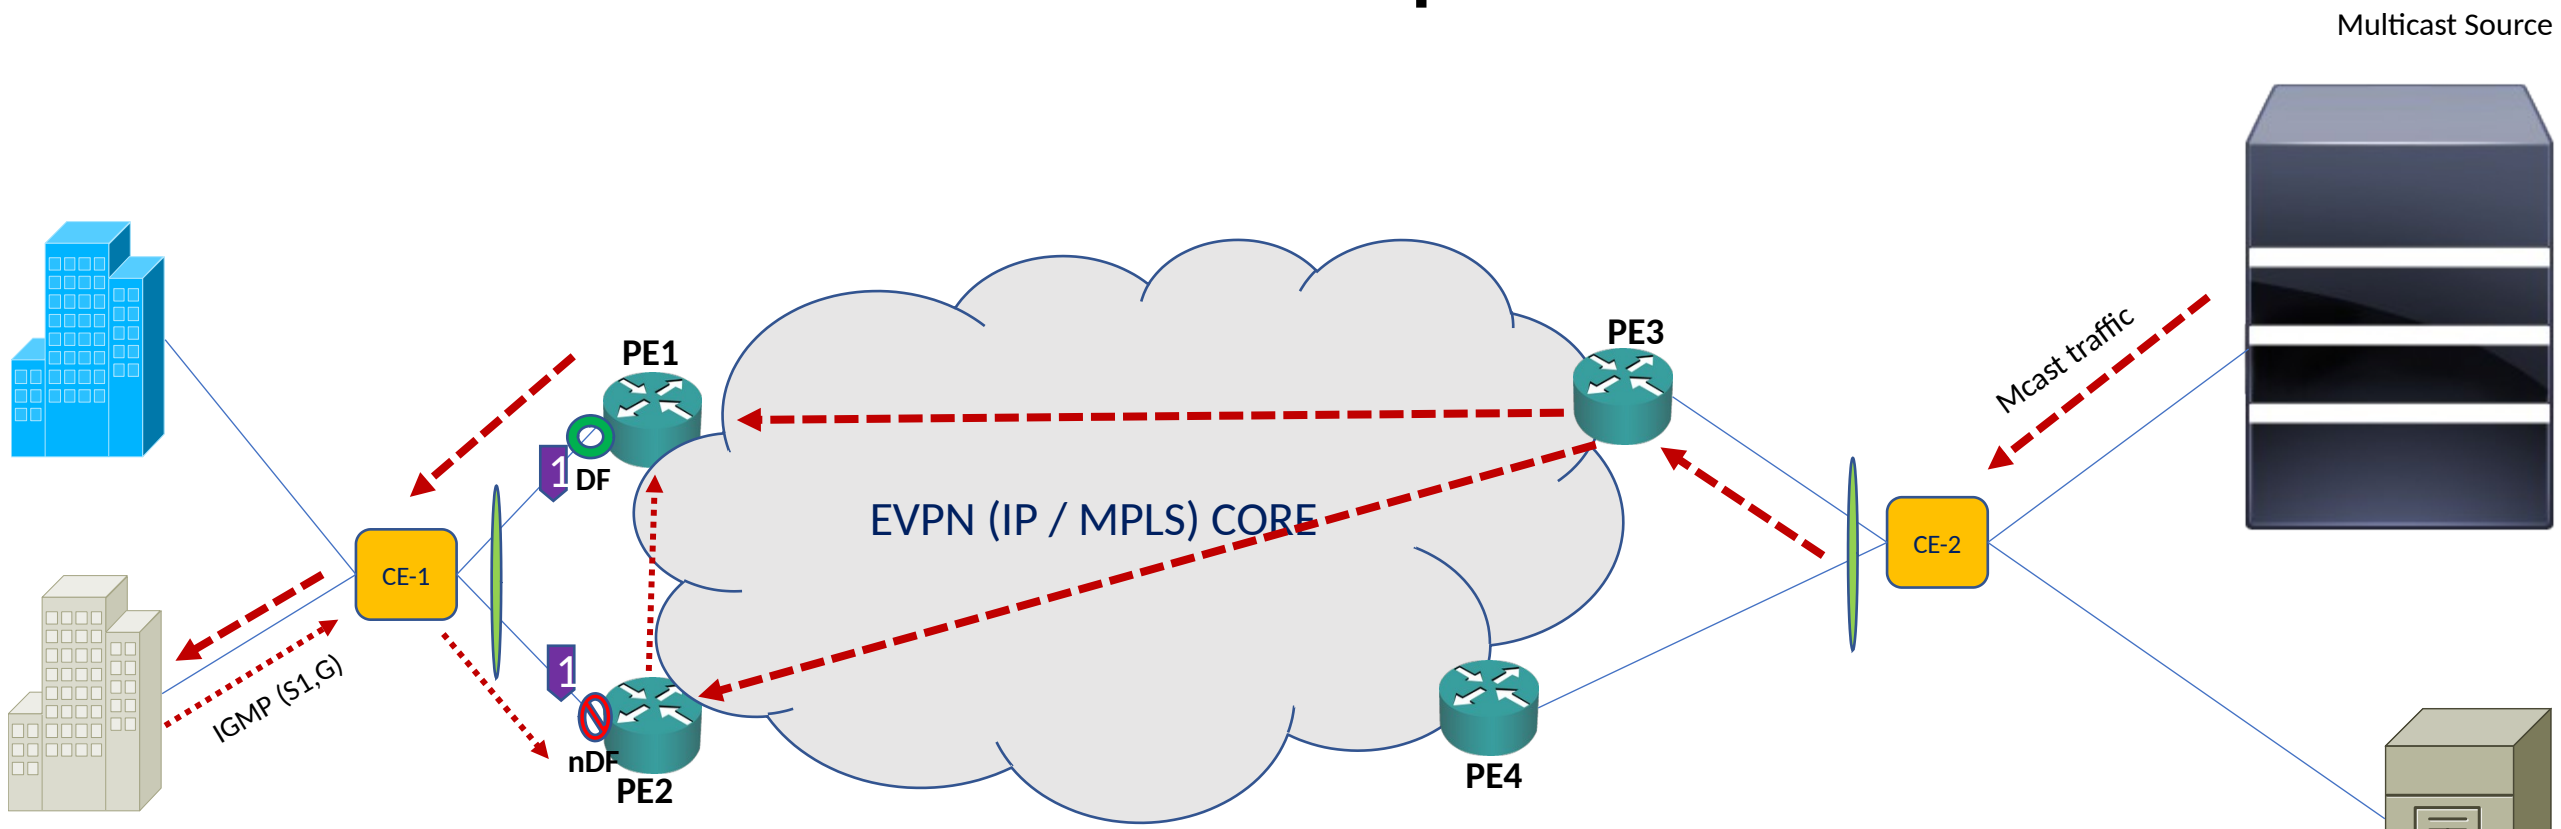# EVPN per multicast flow DF election

Ali Sajassi

Mankamana Mishra

Samir Thoria

Jorge Rabadan

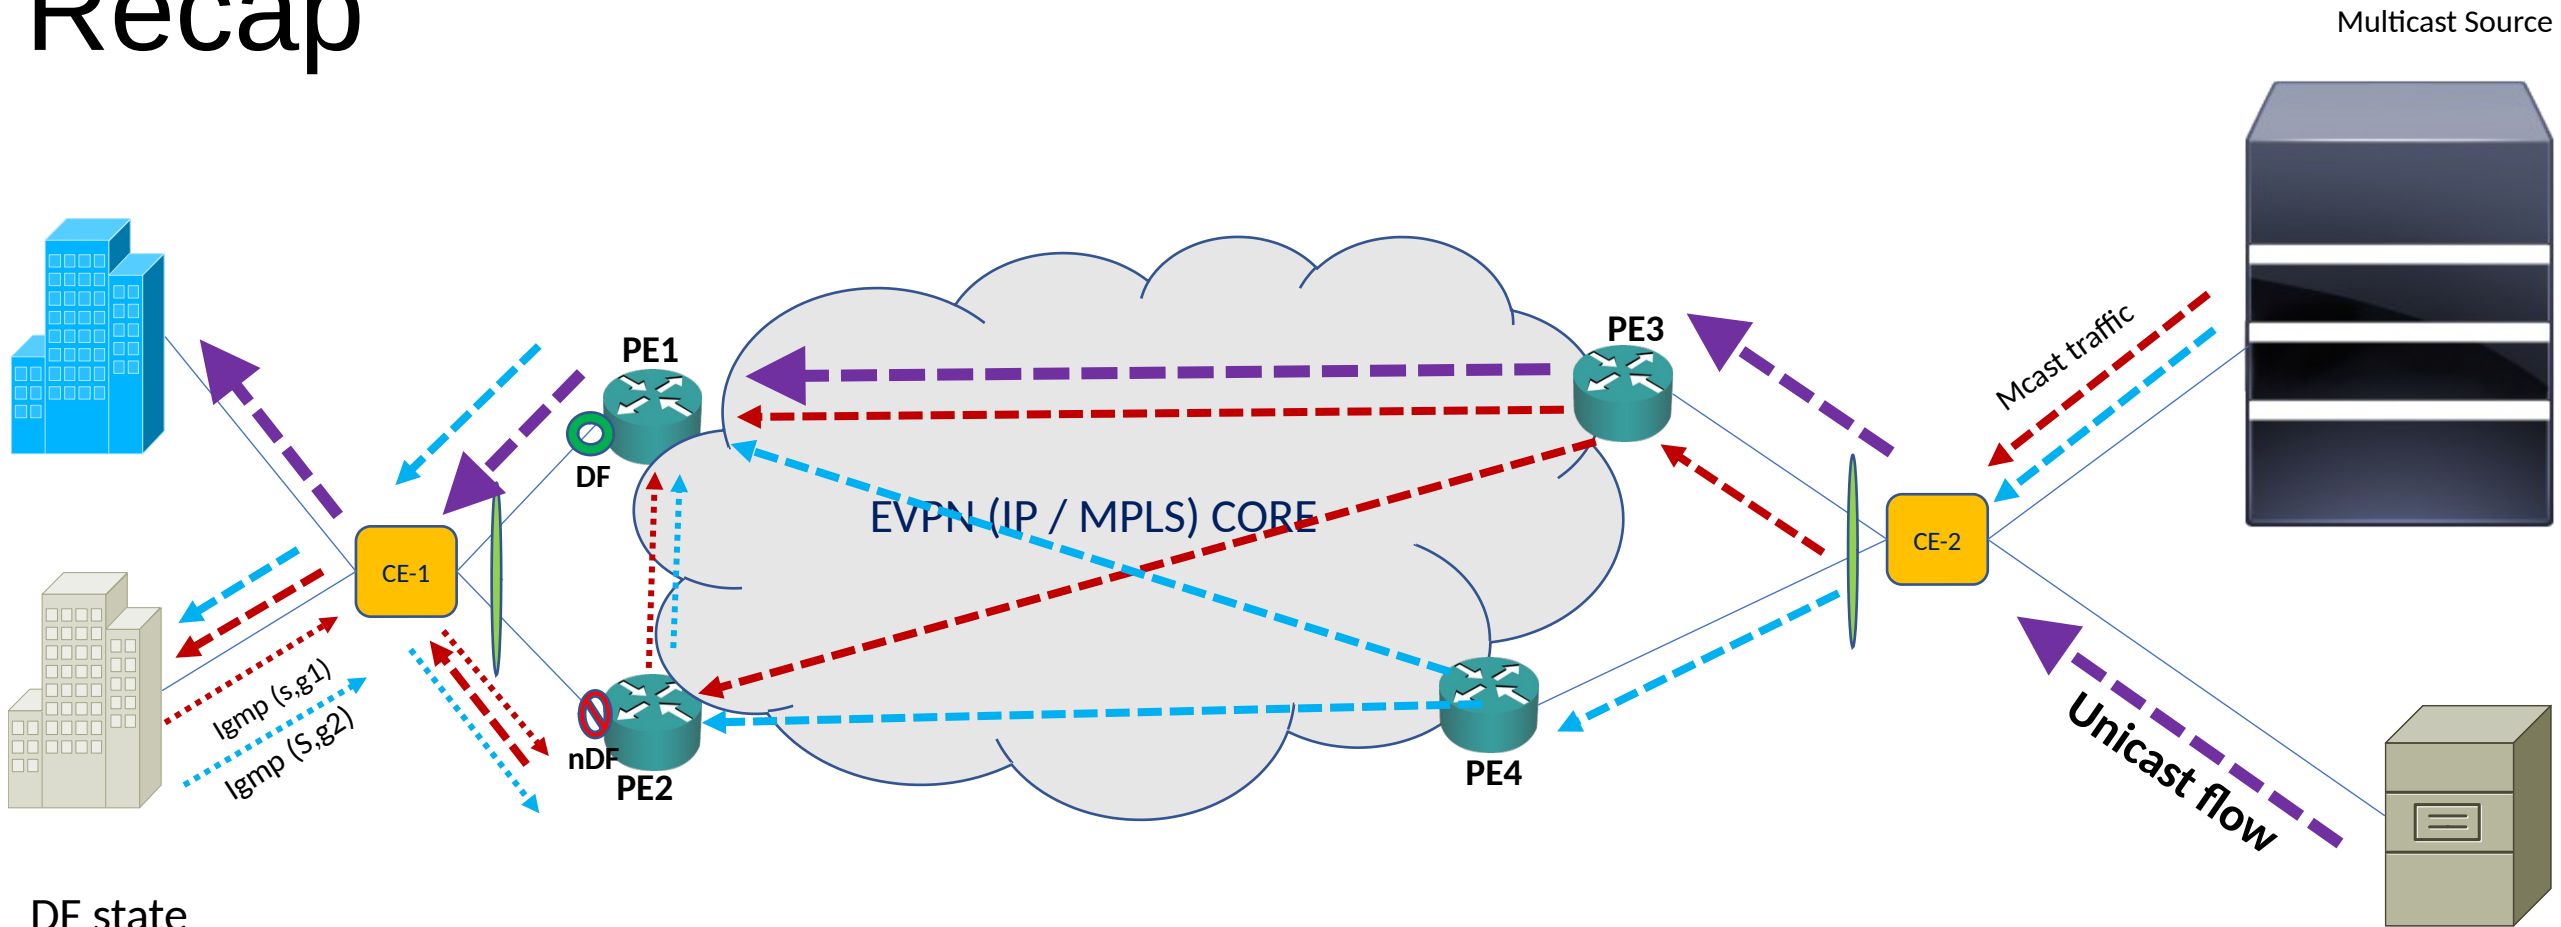John Drake

Problem Statement - Recap

Existing DF election procedure	Which Use case does not covered by existing DF election procedures
• <u>Default DF election procedure as per RFC 7432</u>: It describes mechanism to elect designated forwarder (DF) at the granularity of (ESI, EVI) which is per VLAN (or per group of VLANs in case of VLAN bundle or VLAN-aware bundle service) • <u>Enhanced DF election</u>: Framework for EVPN Designated Forwarder Election Extensibility enhances default DF election procedure further and adds HRW hash algorithm for better distribution	There are service providers where residential deployment requires all of multicast flows in one Vlan. For those case each of existing DF election mechanism does not help, As both of the existing procedure going to pick only one of the PE as DF for all of the multicast flow. So there is need for new mechanism which can still load balance multicast flows between redundancy group of PE's.

Proposed solution - Recap

Multicast Source

PE1

DF

PE3

EVPN (IP / MPLS) CORE

CE-1

CE-2

Mcast traffic

Igmp (s,g1)
Igmp (S,g2)

nDF
PE2

PE4

Unicast flow

File/
Application
Server

DF state

PE1	PE2
(s,g2):DF	(s,g2):nDF
(s,g1): nDF	(s,g1):DF
Default : DF	Default: nDF

Next step

- Draft has been implemented.
- Please provide comments if there is any .
- Plan to be ready for WGLC before IETF-110

AC-Aware Bundling Service Interface in EVPN

Ali Sajassi

Mankamana Mishra

Samir Thoria

Patrice

Jorge Rabadan

John Drake

Requirements

- Some of the deployments require capability to have multiple subnets within single Bridge Domain.

- Each of subnet are distinguished by Vlan in Bridge Domain

None of existing service interface from RFC 7432 fulfill requirement to have single IRB (Bridge Domain) with multiple Subnets (Distinguished by VLAN)
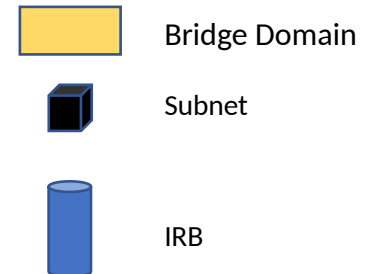
Existing Service interface

1. **VLAN-Based Service Interface**: Each Bridge Domain has only one Vlan. It means Bridge domain to subnet mapping is 1:1
2. **VLAN Bundle Service Interface:** An EVPN instance corresponds to multiple VLANs, but single Bridge Table is maintained. MAC address is key to lookup. The MPLS-encapsulated frames MUST remain tagged with the originating VID
3. **VLAN-Aware Bundle Service Interface:** With this service interface, an EVPN instance consists of multiple broadcast domains (e.g., multiple VLANs) with each VLAN having its own bridge table. Which would map to multiple IRB interface. MAC+VLAN are learnt, and lookup is based on MAC + VLAN to get appropriate port.

Problem Statement:
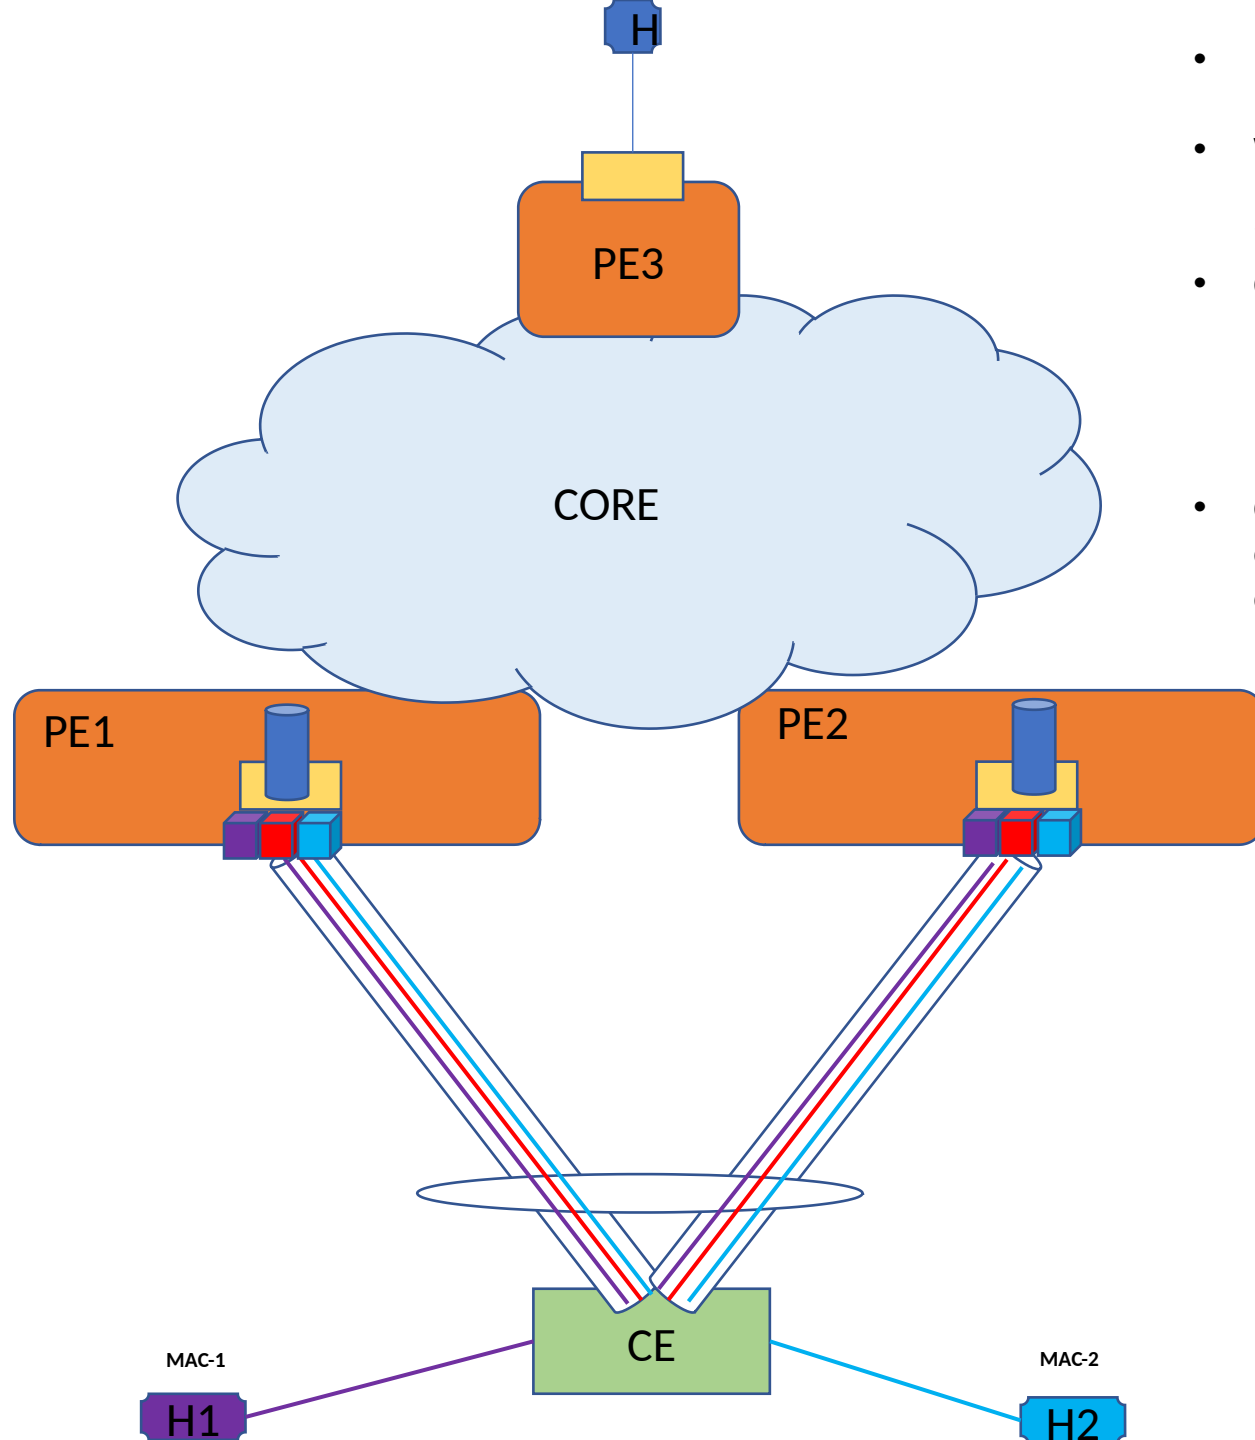
- PE1, PE2 are multi-homed PE where single Bridge Domain has multiple subnets .
- When PE1 learns a MAC address from any of subnet, it Need to advertise MAC route to peers in same EVPN instance.
- Currently if PE1 originates new MAC route,
  - PE3 need to create next hop entry to reach MAC.
  - PE2 need to assign appropriate Bridge port to forward packet destined to MAC.
- Current Service Interface defined in RFC7432 does not cover this requirement, as PE2 would not have any context of which VLAN this MAC was learnt.

AC-aware Bundling Service Interface

- New Service interface type is defined to fulfill the requirement which was described earlier.

- With AC-aware bundling service interface , MAC advertisement would have Extended Community which would carry Attachments Circuit ID.

- New Extended community MUST be carried with EVPN Route Type 2 (MAC Route), 7&8 (Multicast Route)

- Attachment Circuit extended community MUST be ignored by Non Multi-homed peers.

Control plane operations

Local Route:

1. When a PE learn MAC address, it MUST attach attachment circuit ID extended community with MAC route

2. If a PE receive IGMP membership notification (join / leave) it MUST attach attachment circuit ID extended community with multicast route

Remote Route:

1. If recipient PE of MAC route is non multi-homed peer, it must process MAC route with procedure defined in RFC7432
2. If recipient PE of MAC route is multi-homed peer, It MUST get AC information from Extended community and associate MAC address with appropriate AC / Subnet / Vlan
3. Multicast route MUST also programmed with appropriate AC / Subnet / Vlan, where AC information is derived from extended community

Data Plane operations at non multi-homed peer

- This service interface does not change any procedure at non-multihomed remote peer of EVPN instance. It MUST follow procedure defined in RFC 7432

Data Plane operations at multi-homed peer

- Data received from CE MUST follow procedure defined in RFC7432. New proposal does not add any extra processing.

- Unknown unicast packet MUST follow procedure defined in RFC7432. New proposal does not add any extra processing.

- With new defined Service interface for known unicast packet MAC lookup would provide Associated Vlan and port information. Appropriate Vlan tag must be pushed and forwarded on port.

- Multicast packet MUST be forwarded based on multicast state.

Next Step

- We already have implementation
- Will address comments in list in next revision and ask for adoption call