

draft-sajassi-bess-evpn-vpws-fxc-00.txt

A. Sajassi (Cisco), P. Brissette (Cisco), J. Uttaro  
(ATT), J. Drake (Juniper), W. Lin (Juniper), S.  
Boutros (VMWare), J. Rabadan (Nokia)

IETF 96, July 2016

Berlin

# Problem Statement

---

- EVPN-VPWS currently covers the following services interfaces (per section 2)
  - VLAN-based
  - VLAN-bundle
  - Port-based
- This draft defines a special type of VLAN-aware bundle service interface called Flexible Cross Connect (FXC)

# Current Situation with EVPN VPWS

---

- Provides P2P connectivity between two Attachment Circuits (ACs) on two different PEs
  - AC can be defined as <port> or <port, VLAN> or <port, VLAN group>
- Cannot aggregate ACs across different ports on the same PE
- Each AC is signaled via BGP to other PEs
  - Remote PE with the right service-id creates cross-connect between its local AC and the one received via BGP signaling
- Upon failure (AC or port or PE), the EVPN route associated with the AC is withdrawn

# What's more needed ?

---

1. Significantly reduce number of PWs by muxing VLANs across many interfaces
2. Reduce BGP signaling as much as possible (e.g., not to signal every VLAN/AC)
3. Support All-Active (LAG) and Single-Active multi-homing


Note: some of the above requirements conflict w/ each other  
– ie, to support multi-homing and use a single PW to mux VLANs on multiple interfaces, requires per-VLAN signaling

# Flex XC VLAN Unaware: Single-homed CEs

A-Leaf 1 Eth A-D Route
RD = 5500
ESI = 0
Eth.Tag ID = 10,000
Label = 123432

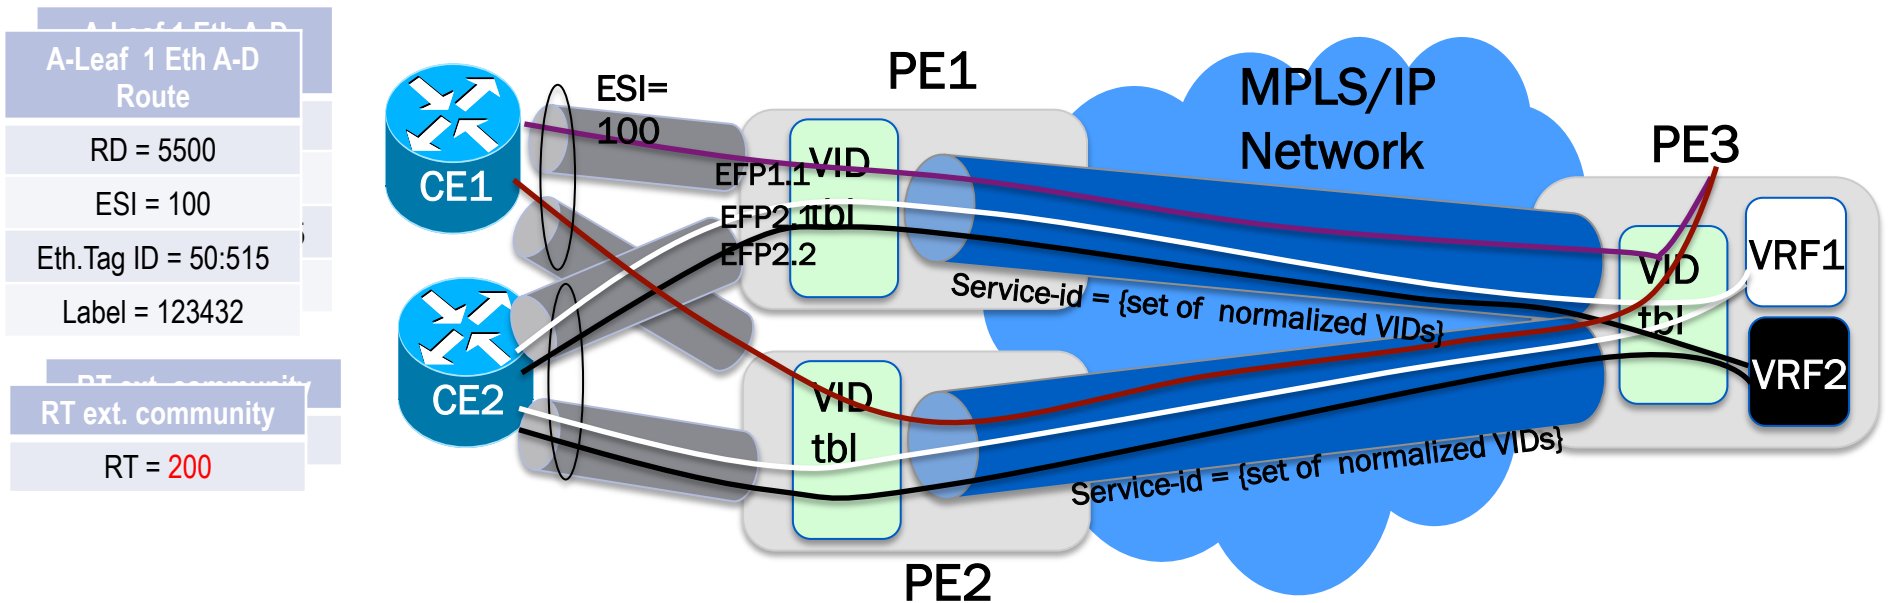
RT ext. community
RT = 100



## Service Tunnel Characteristics

- One single service tunnel between a pair of VID tables
- A single service-id represents the service tunnel
- Normalized VIDs are unique in context of VID table
- VLAN unaware - VLAN failure is NOT signaled over BGP
- Muxes many VLANs across several interfaces on the PE
- Because of muxing, VID table is needed at disposition
- C-VIDs are normalized into double-tag S-VIDs and then sent over service tunnel
- When all VLANs (across all interfaces) for a service tunnel are failed, then the PE withdraws the route for that PW

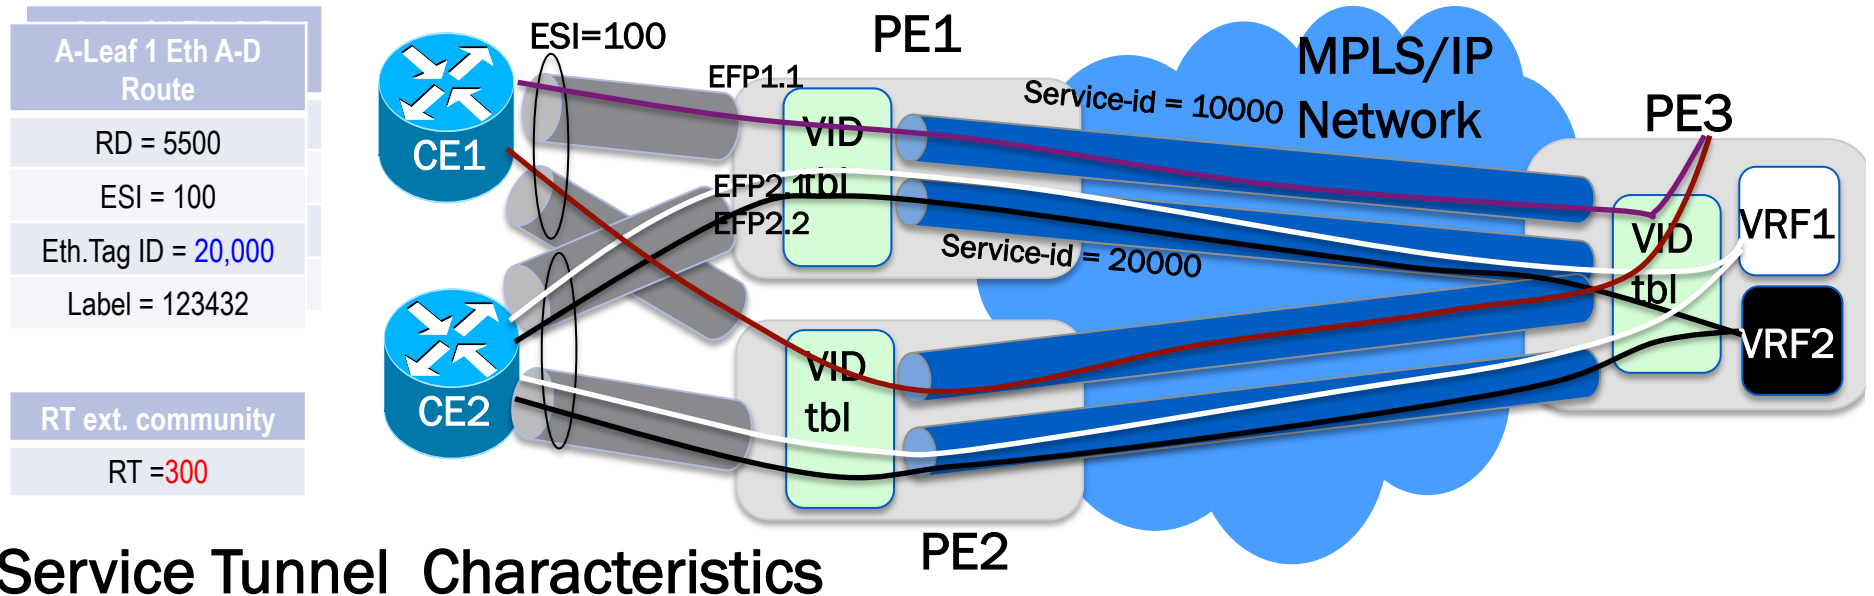
# Flex XC VLAN Aware: Multi-homed CEs



## Service Tunnel Characteristics

- One single service tunnel between a pair of VID tables
- Many service-ids (normalized VIDs) represent the tunnel
- Each EVI represent a single normalized VID space in this option
- Automatic cross-checking of normalized VIDs on PEs
- VLAN aware – VLAN failure is signaled over BGP
- Muxes many VLANs across several interfaces into a single service tunnel

# Flex XC VLAN Unaware: Multi-homed CEs



## Service Tunnel Characteristics

- One single service tunnel for vlan bundles between two PEs
- A single service-id represents the service tunnel
- VLAN unaware - VLAN failure is NOT signaled over BGP
- Muxes group of VLANs on a **single** interface on the PE
- VID table may not be used at disposition PE
- C-VIDs are normalized into double-tag S-VIDs and then sent over service tunnel
- When Ether1 on PE1 fails, then it is signaled via BGP to remote PEs so that they remove PE1 from their path list – e.g., traffic for black VLAN is sent to only PE2

# BGP Extension

---

```

+-----+
| Type(0x06)/Sub-type(TBD)(2 octet) |
+-----+
| Control Flags (2 octets)          |
+-----+
| L2 MTU (2 octets)                 |
+-----+
| Reserved (2 octets)               |
+-----+

```

```

      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   MBZ                               | V | M | C | P | B |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Name	Meaning
B,P,C	per definition in <a href="#">[EVPN-VPWS]</a>
M	00 mode of operation as defined in <a href="#">[EVPN-VPWS]</a> 01 VLAN-aware FXC 10 VLAN-unaware FXC
V	00 operating per <a href="#">[EVPN-VPWS]</a> 01 single-VID normalization 10 double-VID normalization



# Next Step

---

- Questions ?
- Discussions on the mailing list