

# EVPN and L2 Access Protocols: Single-Flow-Active load-balancing mode

draft-brissette-bess-evpn-l2gw-proto

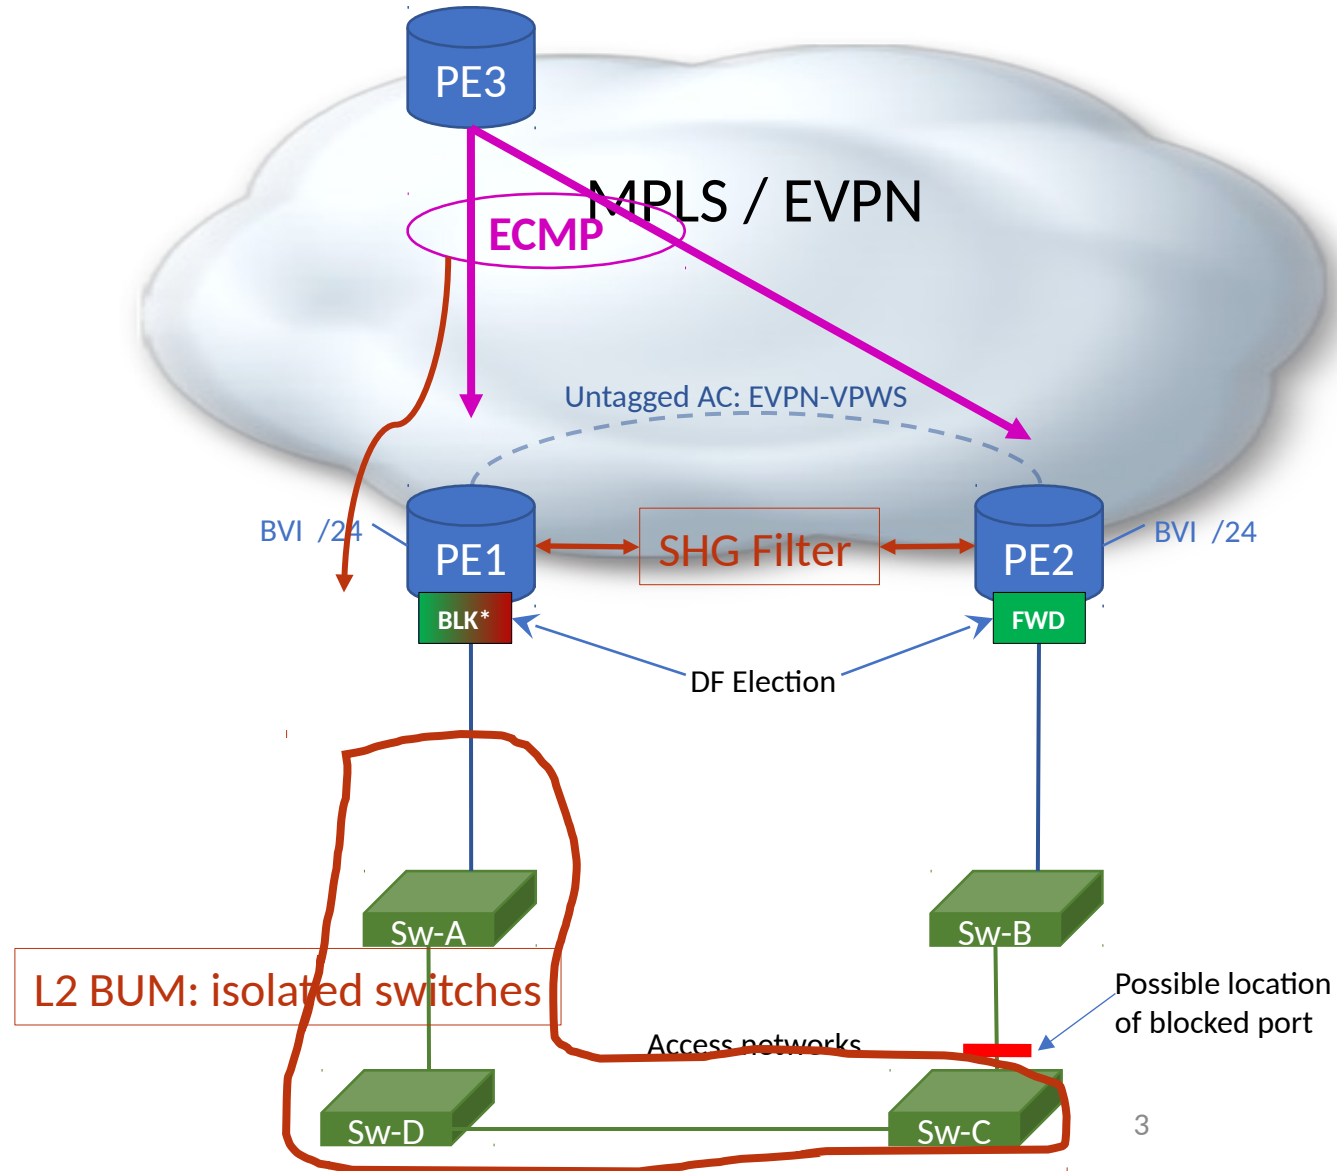
# Requirements

- Support L2 Access Protocols into EVPN core (incl. EVPN-IRB).
  - (M)STP 802.1d, G.8032, MPLS-TP, MST-AG, REP-AG, etc.
- L2 gateways largely unaware of the access network topology **but**
  - Access network is running an L2 protocol breaking loops
  - Access network is **segmented by design**
- Solution **MUST** support:
  - Optimal forwarding at L2 and L3
  - Fast-convergence
  - Scalability
  - Anycast IRB support

# Challenges with existing load-balancing modes

## All-Active

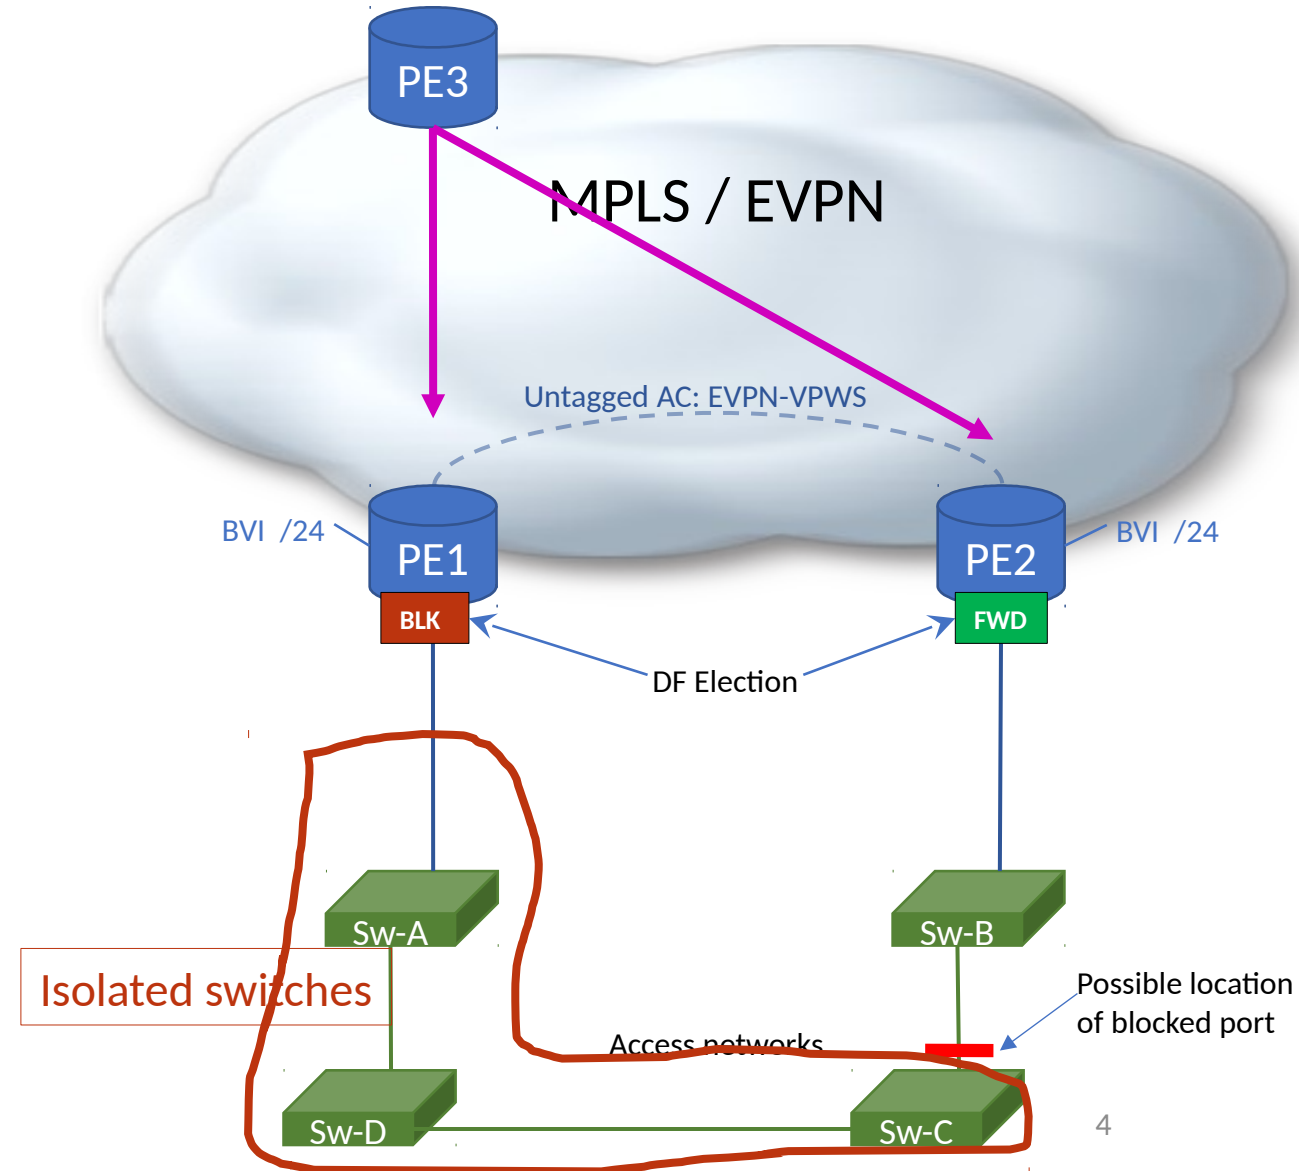
- creates L2 and L3 aliasing resulting in remote ECMP towards “wrong” PE
- L2 BUM isolated switches
- TCN propagation blocked by ESI split-horizon label filtering



# Challenges with existing load-balancing modes

## Single-Active

- Isolates network segments
- TCN propagation blocked by port state or ESI split-horizon label filtering



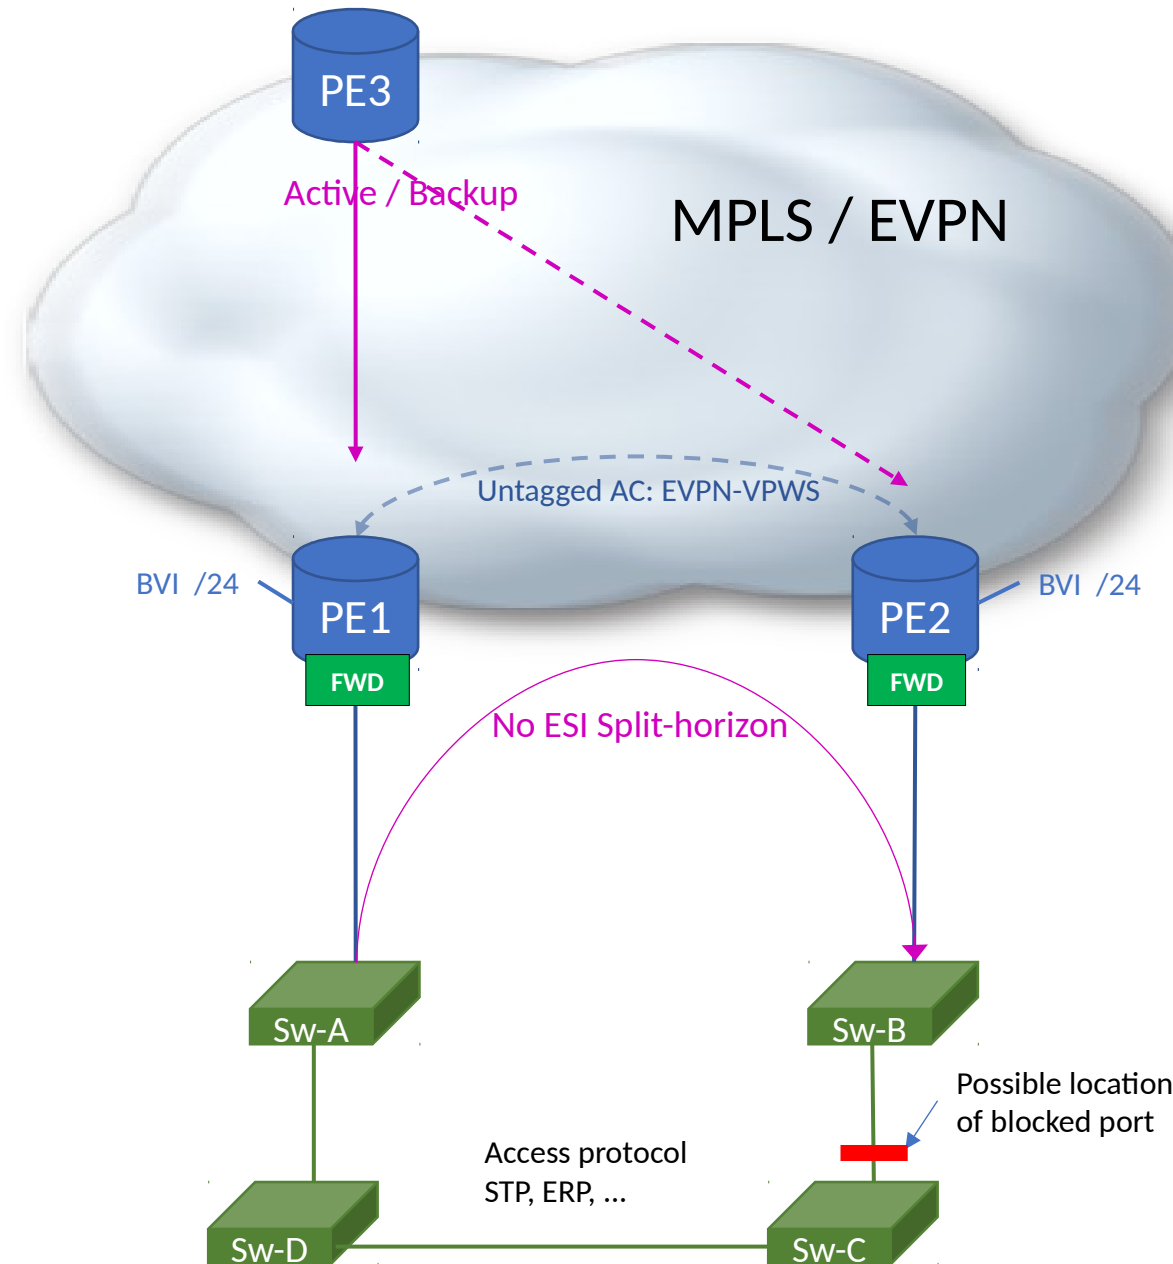
# Single-Flow-Active

- Same ESI between peering PEs:
  - Ethernet-segment procedures still apply
- ESI-Label BGP extcomm extension to signal

[2:0]- 000 = all-active,  
001 = single-active,  
010 = single-flow-active,  
other = unused

[7:3]- Reserved

- Peering PE in SFA mode, MUST:
  - not perform ESI Split-Horizon filtering
  - enable MAC-Mobility within same ESI
- Remote PE for remote ESI in SFA mode:
  - MUST disable aliasing (at Layer-2 and Layer-3)
  - SHOULD pre-program backup per MAC entry



# Single-Flow-Active: TCN

Peering PEs, upon receiving TCN from access

- Perform local MAC flush
  - MAY result in MAC RT-2 withdraws
- TCN propagation via dedicated PW
  - EVPN-VPWS well-suited for this
- Relay TCN between peering PEs and also forward back into access
- Peering access PEs MAY map TCN into EVI/EAD “Flush” extcomm with new Sequence ID

