





EVPN Multicast

Disha Chopra

Agenda

EVPN Multicast Optimizations

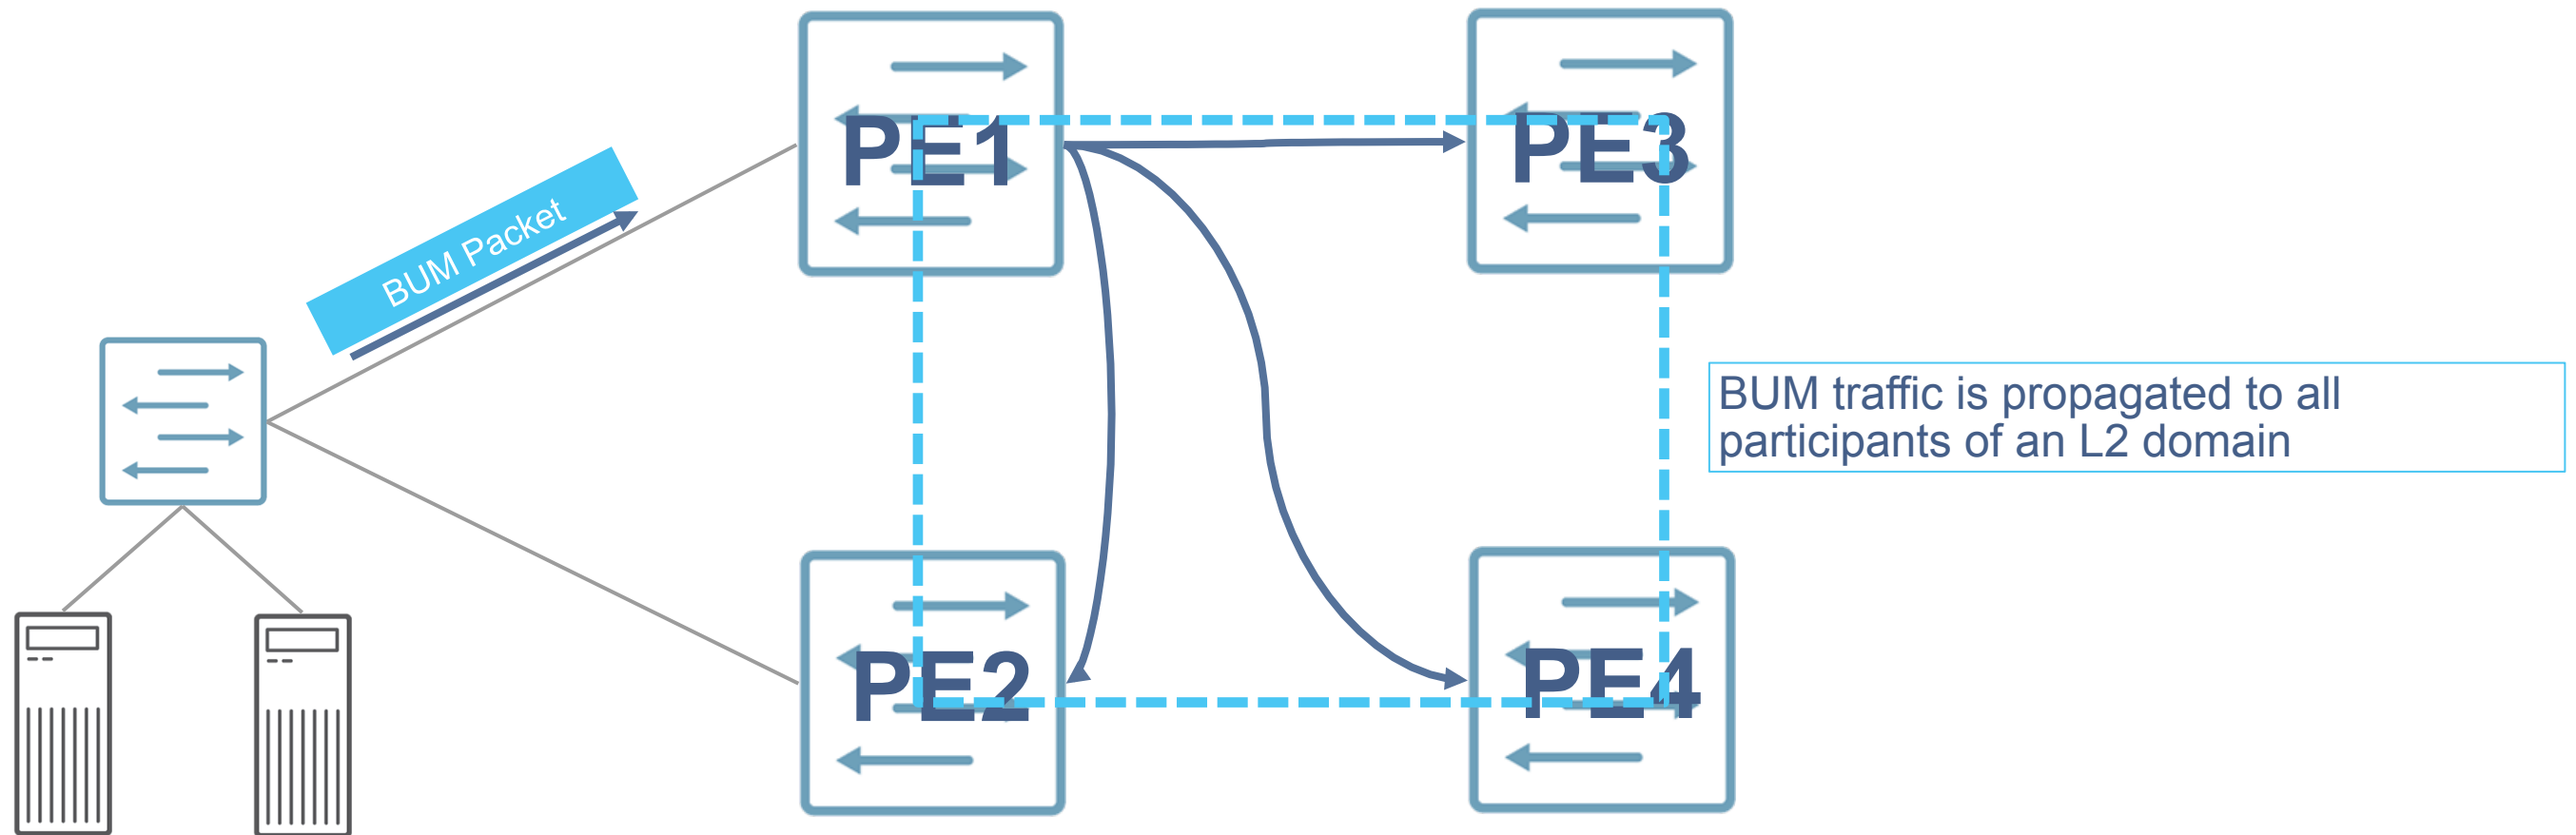
-  Introduction to EVPN Multicast (BUM)
-  IGMP Join/Leave Sync Routes
-  Selective Multicast Ethernet Tag Route
-  Use Case



EVPN BUM Traffic Basics


What is BUM traffic ?

Broadcast, Unknown Unicast, Multicast




What is a Bridge Domain (BD) ?


BUM traffic is sent to all PEs where the BD is configured



Modes of BUM Replication for MPLS

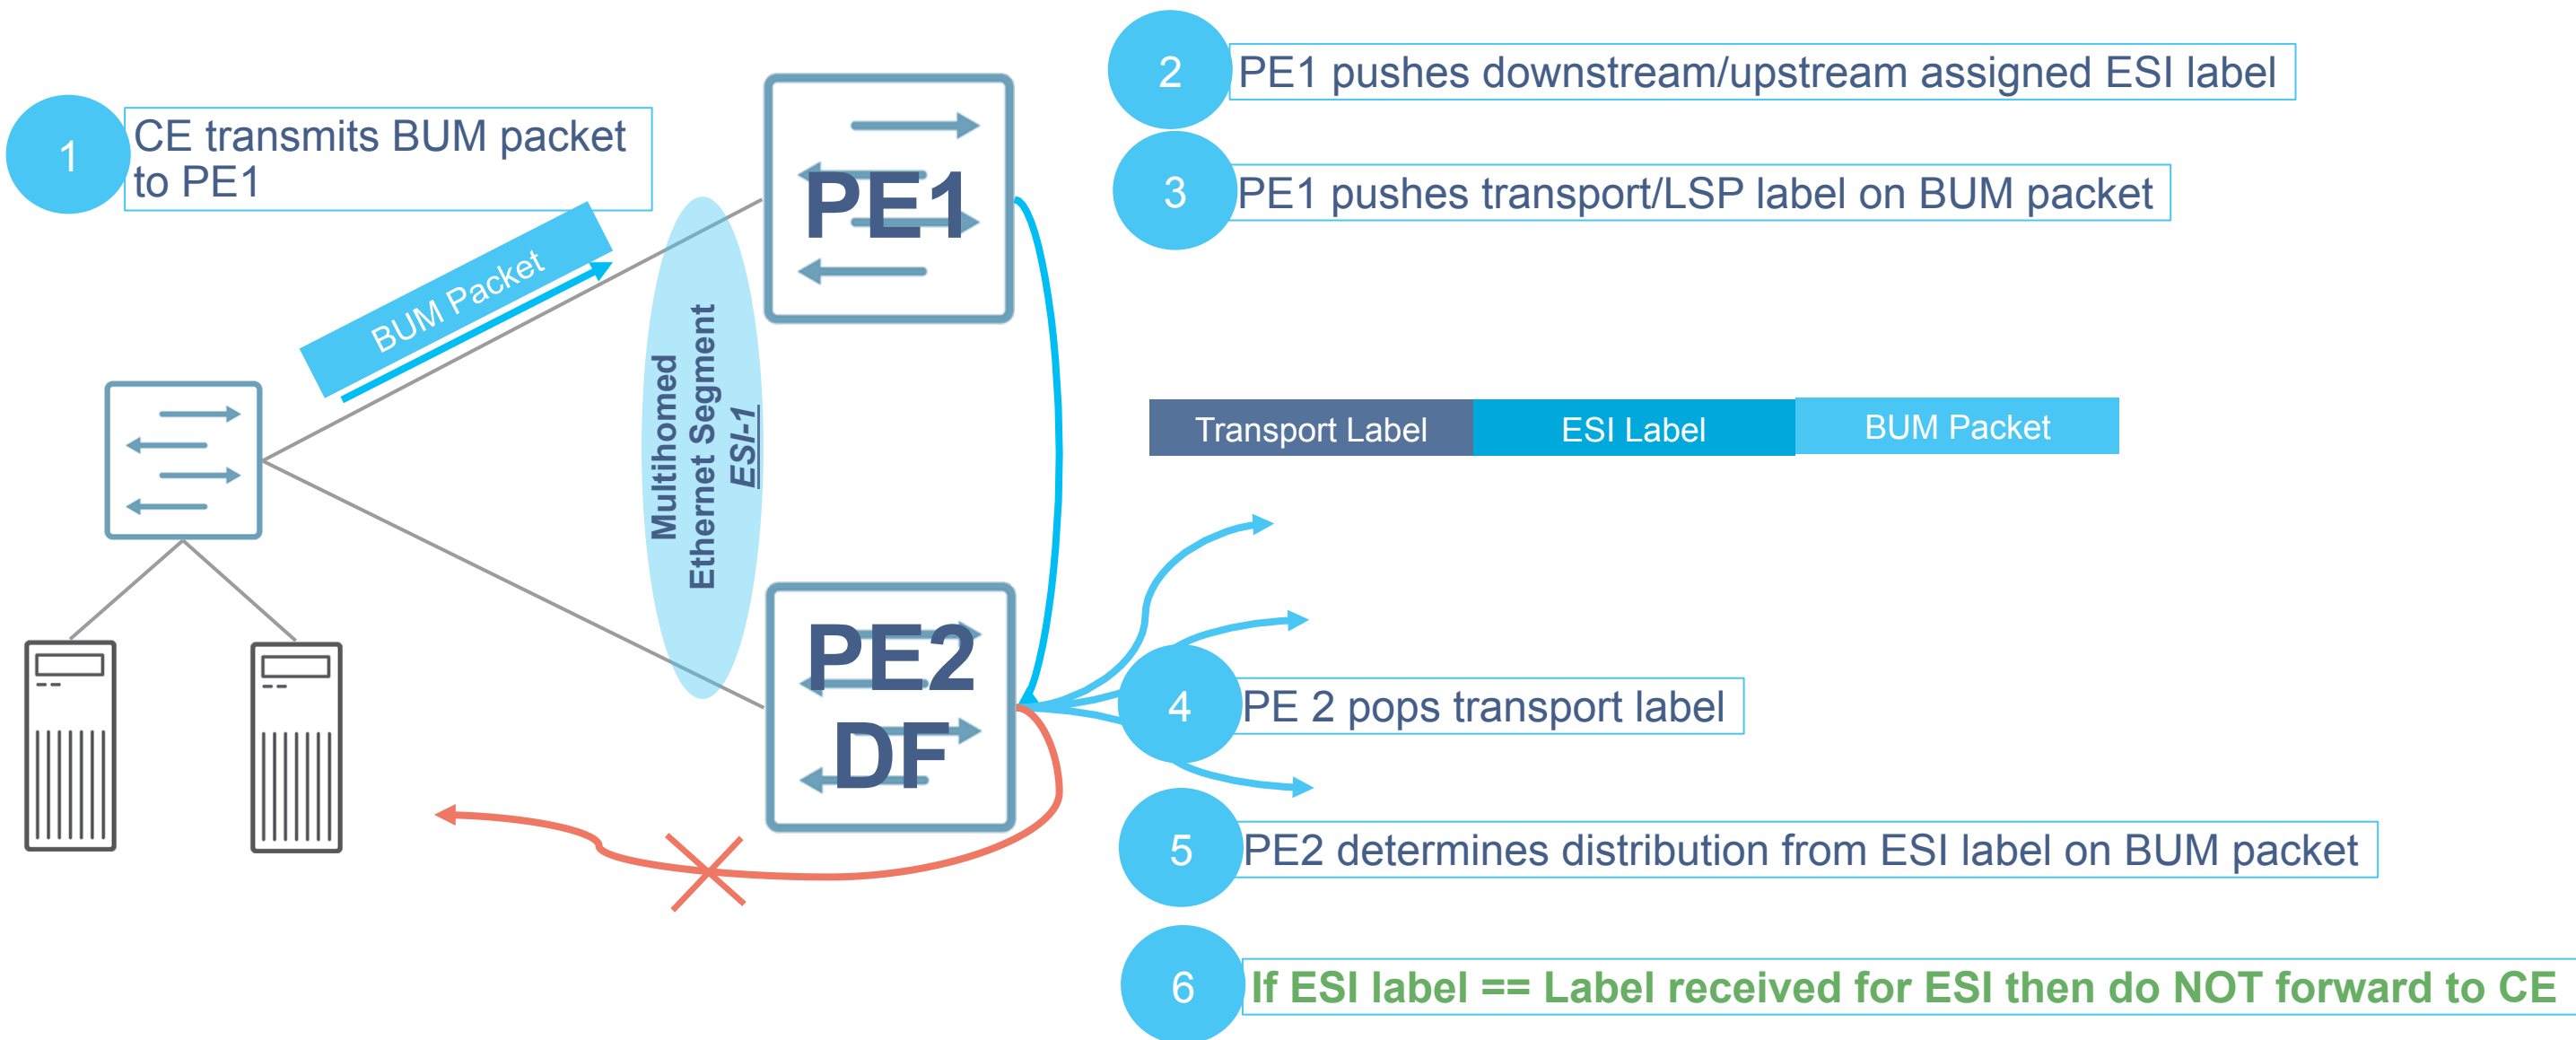


Ingress Replication in the Core



P2MP Replication in the Core

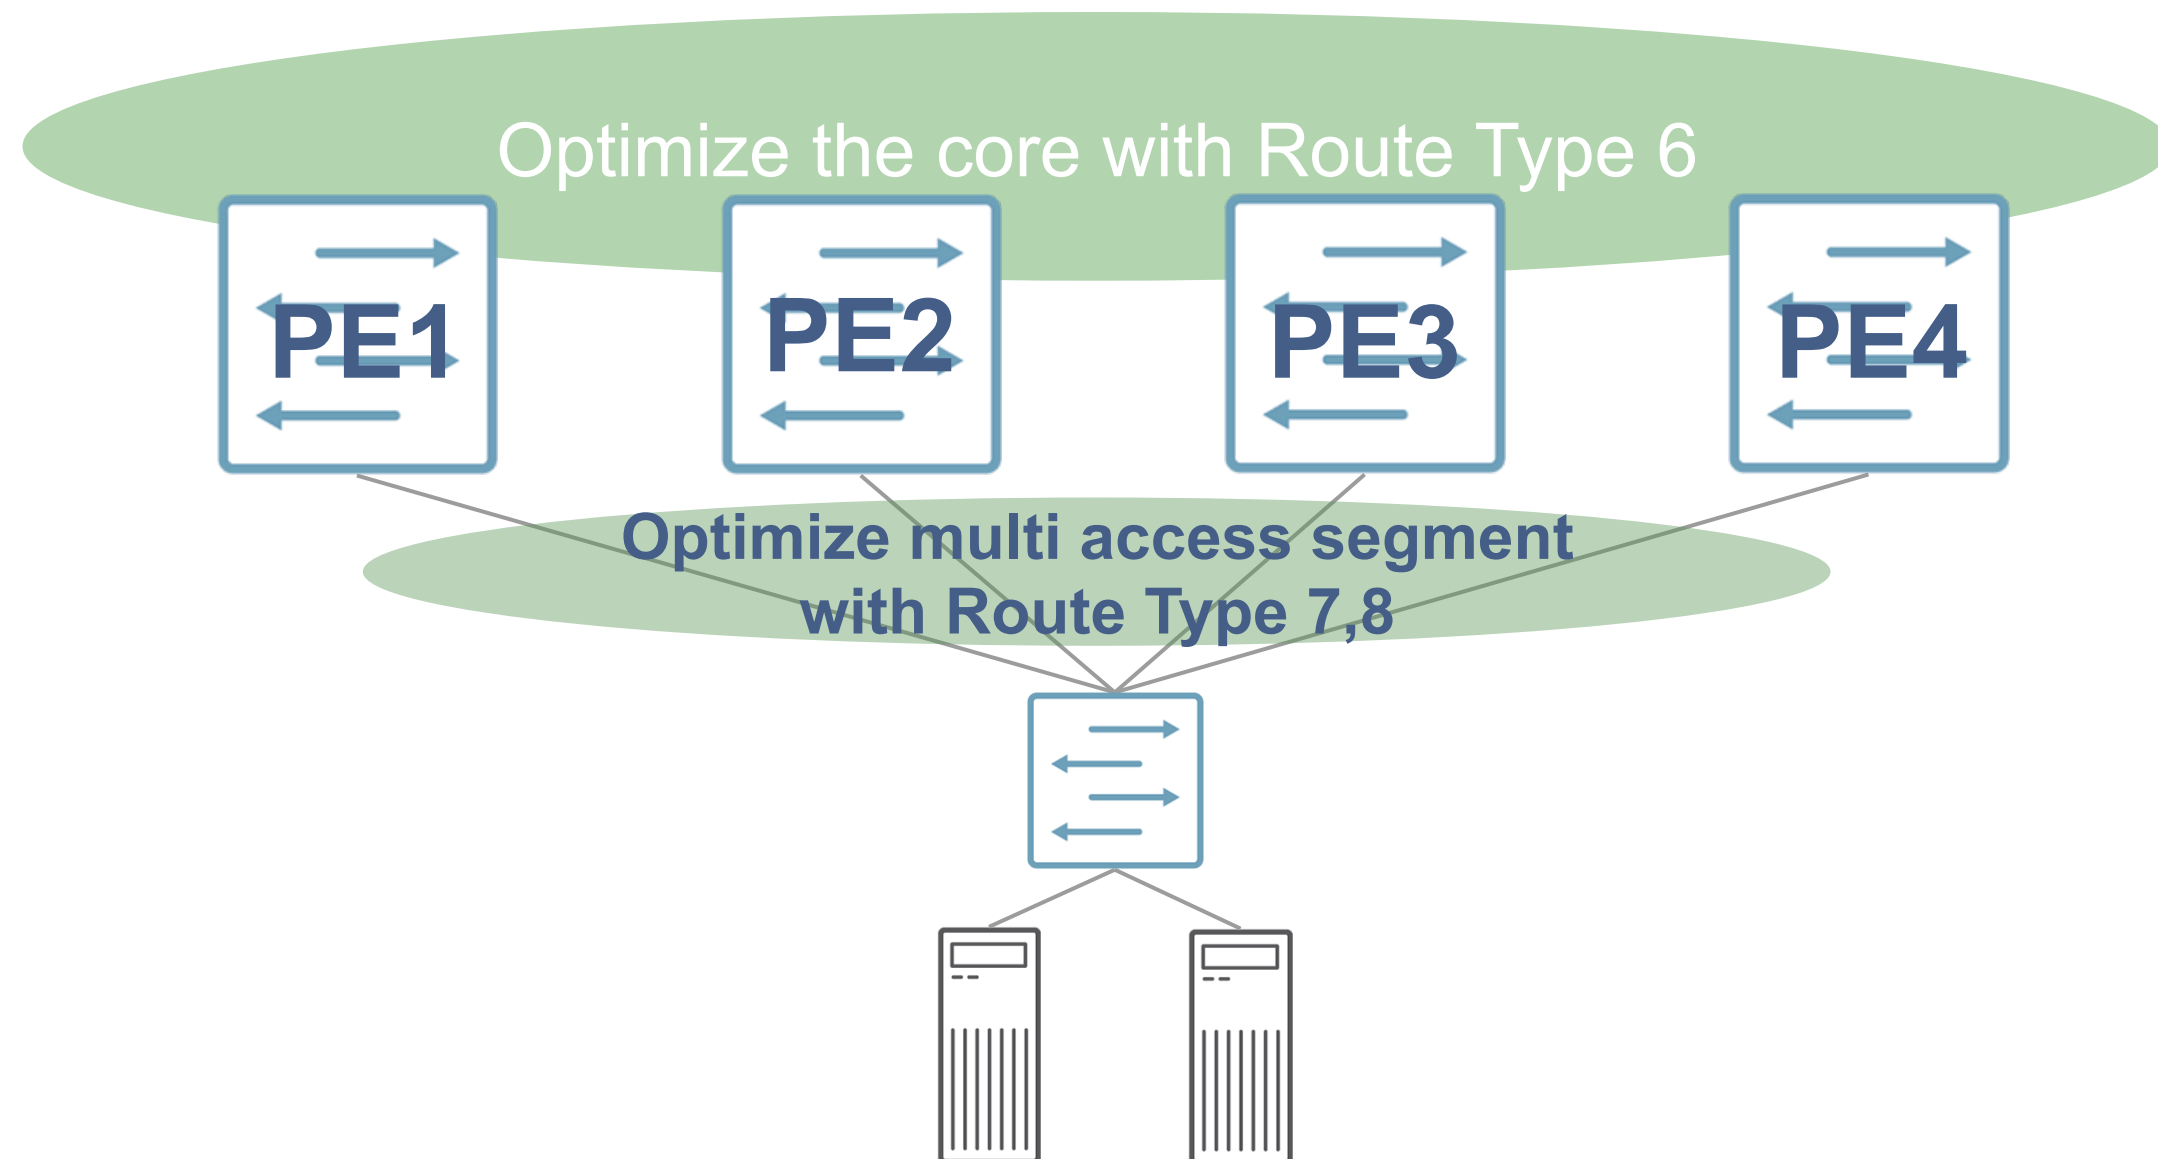
Bum Packet Flow on multi access segment





Multicast Optimizations

Multicast Optimizations in the Overlay



Multicast Optimizations with Route Types 6,7,8

- Route Type 6 (Selective Multicast)
 - To efficiently forward multicast traffic over an EVPN network; only PEs interested in the traffic should receive it.
- Route Type 7,8 (IGMP Proxy)
 - Reduce Bum state just like Proxy ARP. IGMP reports are summarized and propagated to BGP peers. Receiver BGP PE translates them back to IGMP messages. i.e IGMP Overlay is created
 - IGMP proxy is for
 - Querier
 - To enable the collection of EVPN PEs providing L2VPN service to act as distributed multicast router with Anycast IP address for all attached hosts/VMs in that subnet
 - To enable suppression of IGMP membership reports and queries over MPLS/IP core.
 - Especially useful for v1 case where there are no leaves and you need to rely on queries to timeout joins. Reduces Leave Latency
 - Join proxy
 - Leave Proxy

Multicast Flags Extended Community

0	1	2	3
Type=0x06	Sub-type = TBD	Flags	
Reserved = 0		Tunnel Type	

- PEs supporting IGMP Proxy attach this community to the IMET Route
- Absence of the community indicates IGMP Proxy not supported
- Enables efficient Multicast Tunnel Setup from the Source PE
- Especially useful for Ingress Replication
 - If an egress PE supports IGMP proxy it advertises this extended community but, if it does not have interest in a particular S,G then it does not advertise an SMET route for that S,G. In this case, the Ingress PE will NOT replicate traffic to the egress PE
 - Egress PEs that do not advertise this community but have no interest in an S,G will still receive traffic for the S,G
- Tunnel type support defined for RSVP-TE P2MP LSP, P2MP LSP , PIM-SSM, PIM-SM, BIDIR-PIM, mLDP MP2MP LSP

Packet Format : IGMP Join Sync Route : Type 7

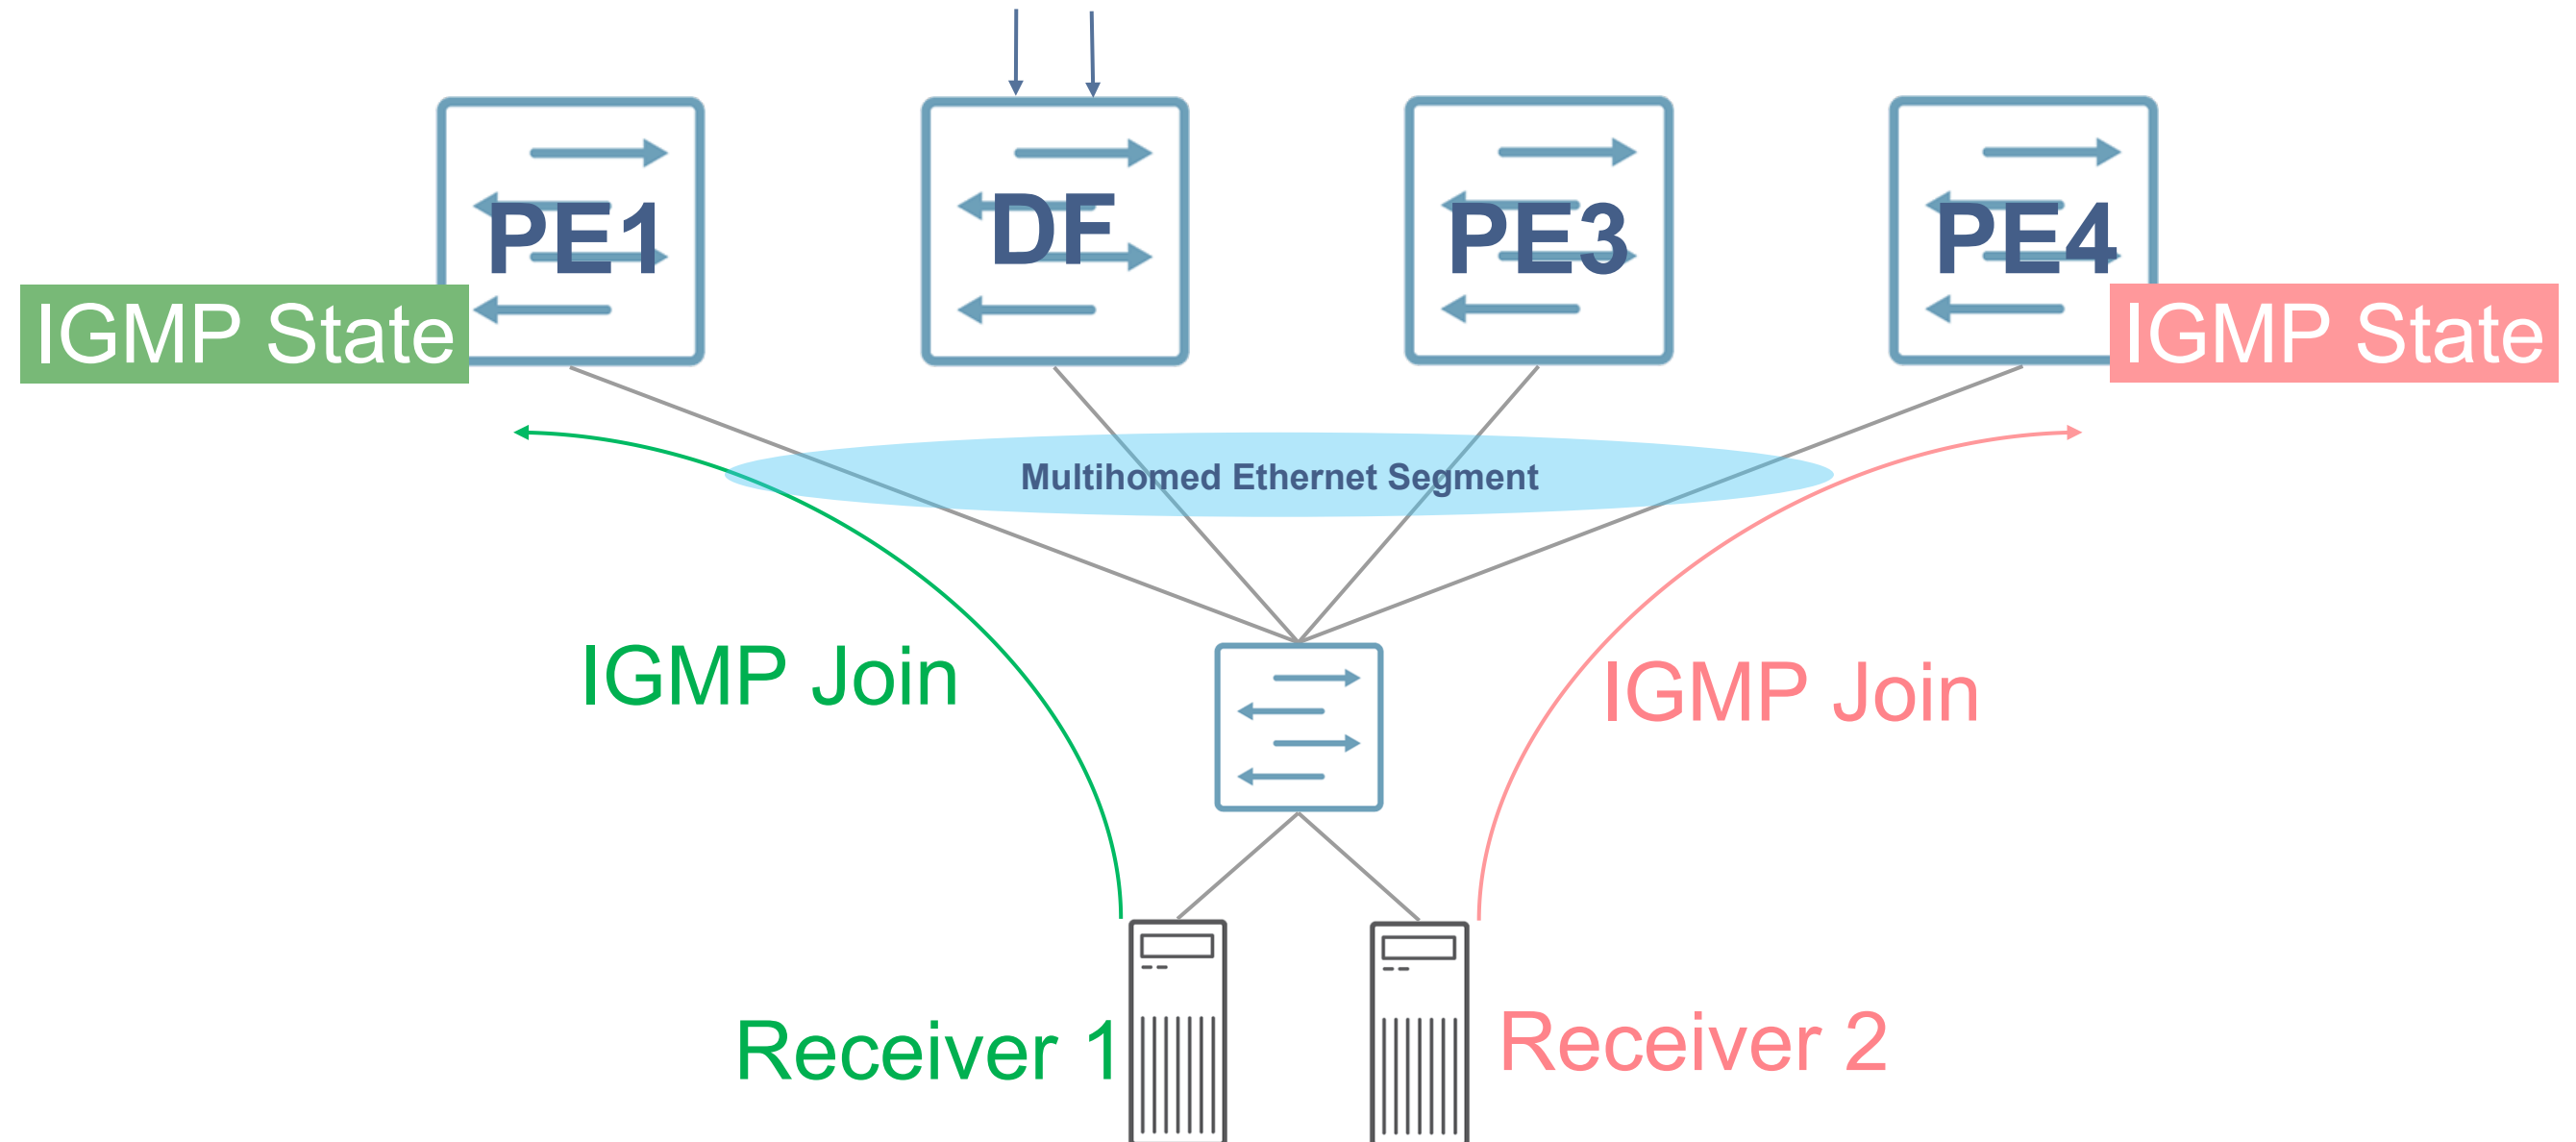
RD (8 octets)						
Ethernet Segment Identifier (10 octets)						
Ethernet Tag ID (4 octets)						
Multicast Source Length (1 octet)						
Multicast Source Address (variable)						
Multicast Group Length (1 octet)						
Multicast Group Address (Variable)						
Originator Router Length (1 octet)						
Originator Router Address (variable)						
Flags (1 octets) (optional)						
Flags						
	RESERVED		IE	V3	V2	V1

NOTES

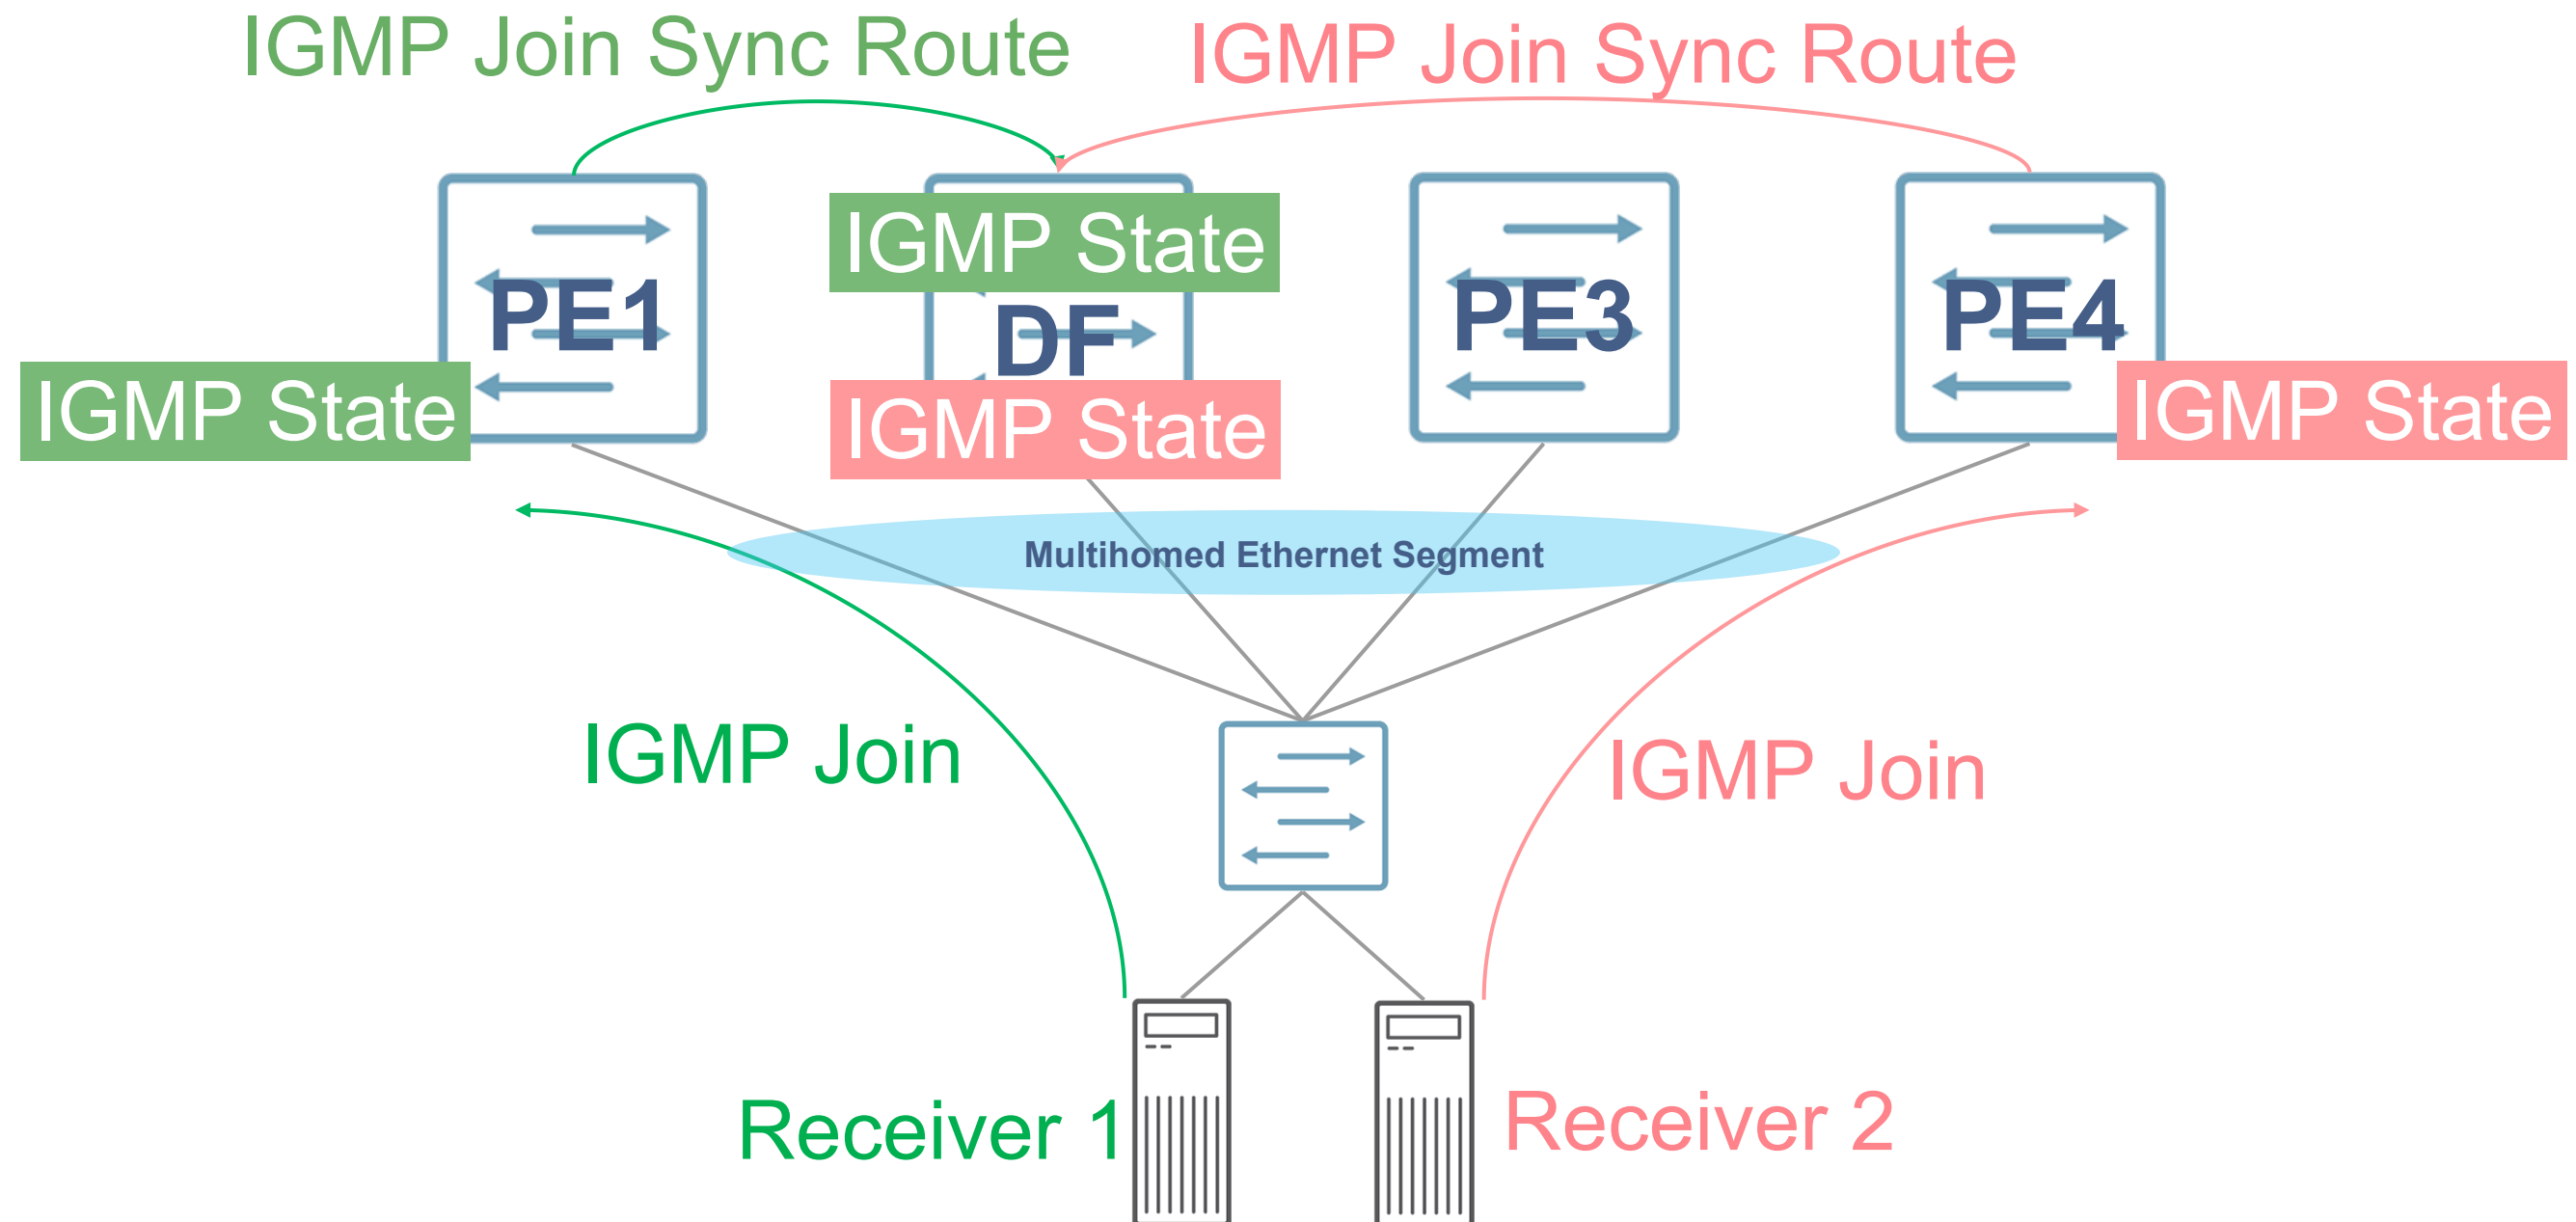
- This EVPN route type is used to coordinate IGMP Join (x,G) state for a given [EVI, BD] between the PEs attached to a given ES
- The Flags field assists in distributing IGMP membership interest of a given host/VM for a given multicast route.
- The Eth Tag ID is set as follows
 - EVI is VLAN-Based or VLAN Bundle service - set to 0
 - EVI is VLAN-Aware Bundle service without translation - set to the customer VID for the [EVI, BD]
 - EVI is VLAN-Aware Bundle service with translation - set to the normalized Ethernet Tag ID for the [EVI, BD]

IGMP Join State Propagation

Multicast Data arrives at EVPN DF from PIM Core



EVPN Route Type 7 : IGMP Join Sync Route



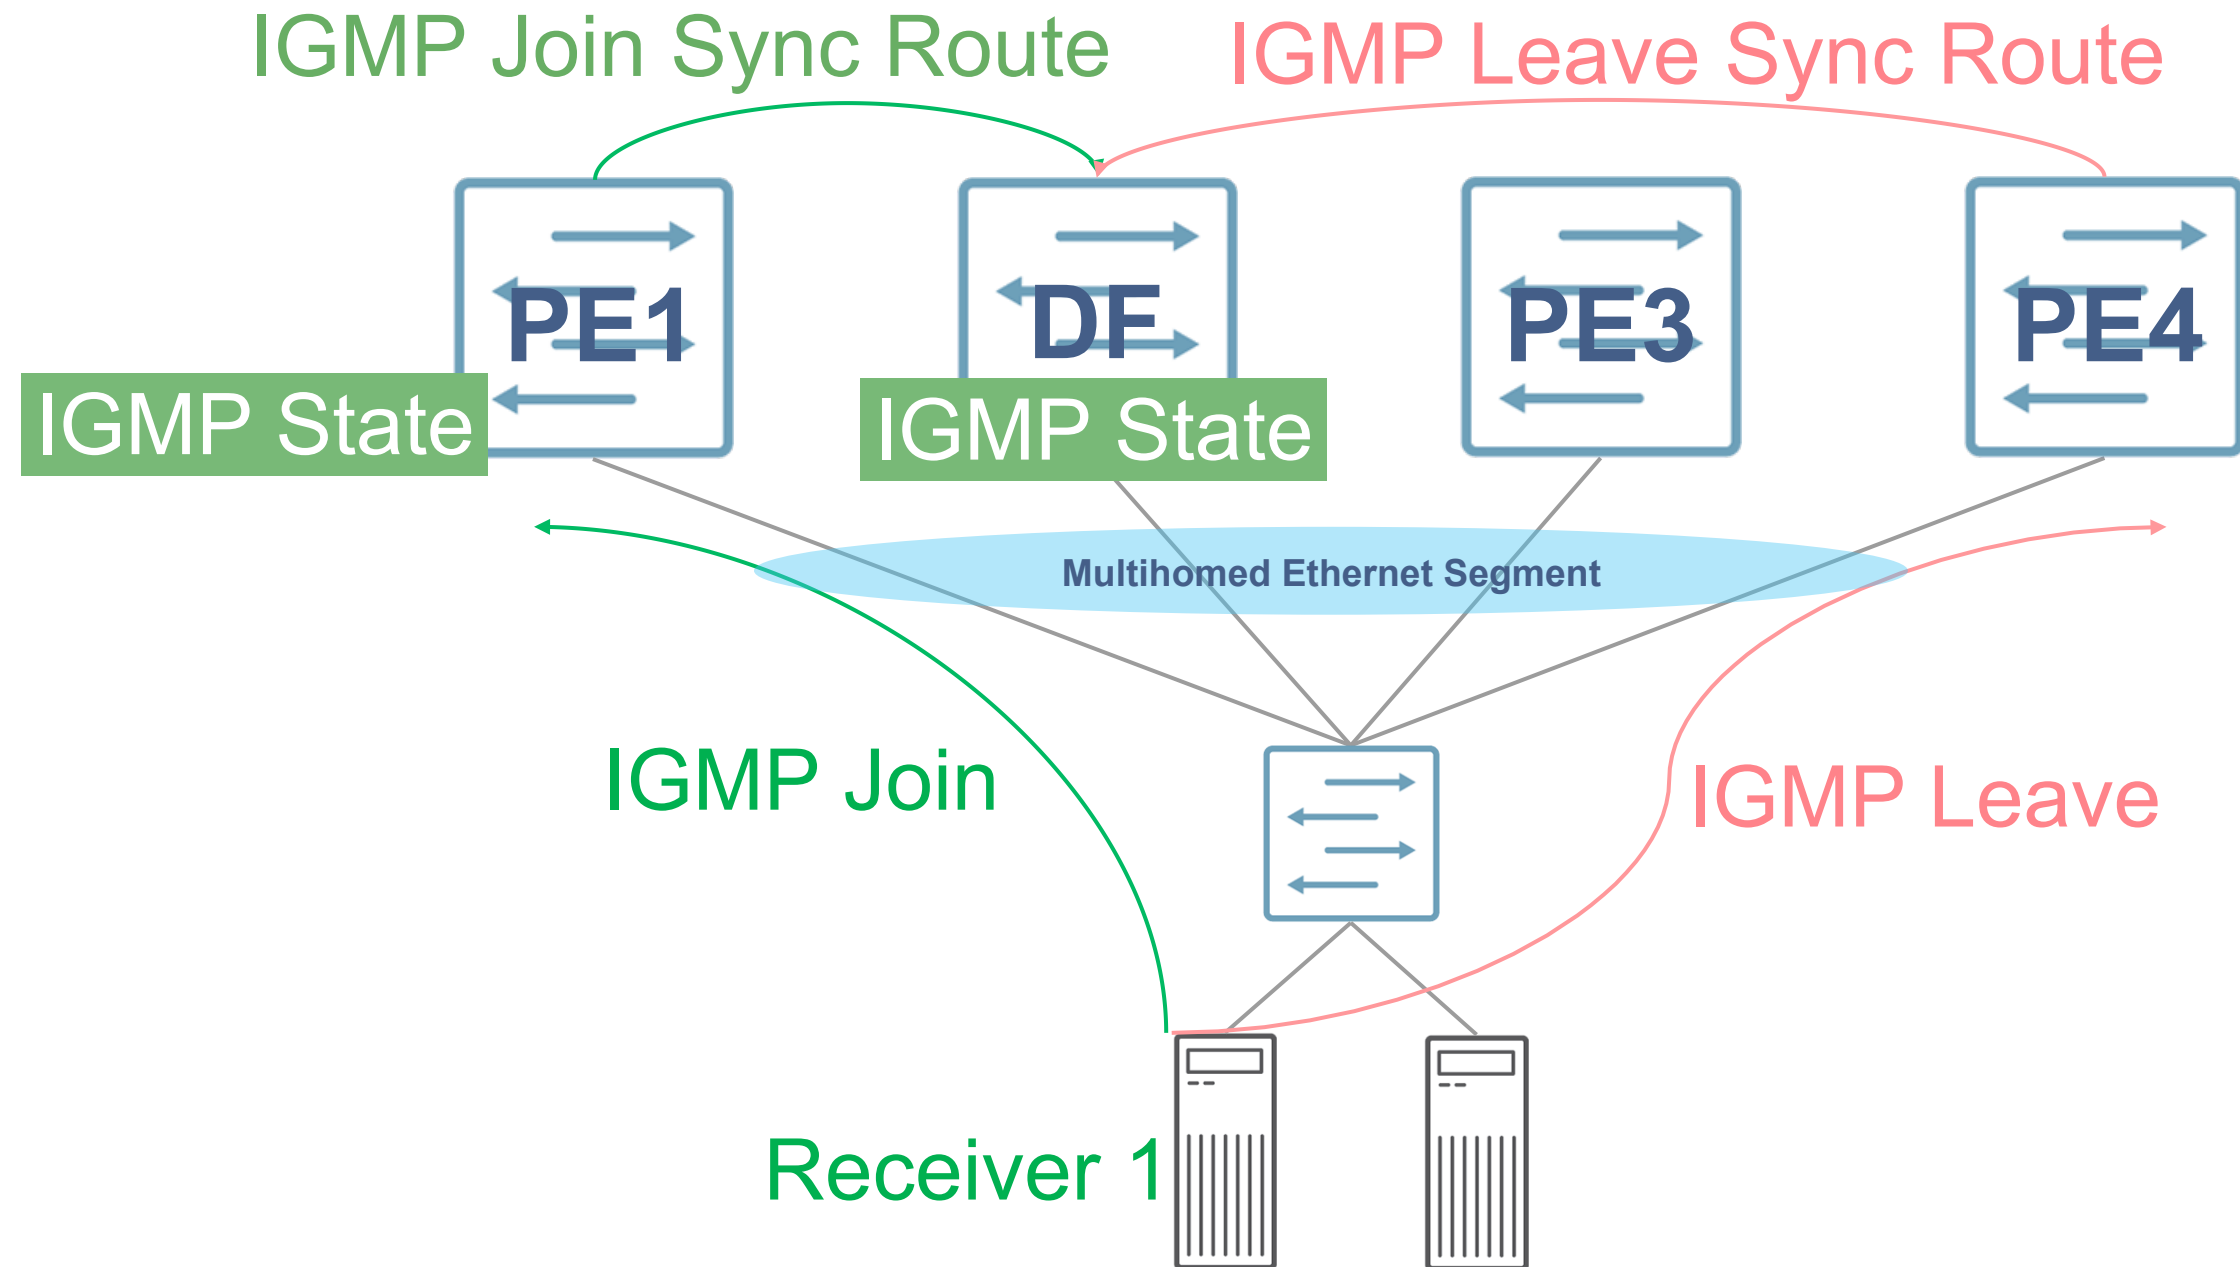
Packet Format : IGMP Leave Sync Route : Type 8

RD (8 octets)
Ethernet Segment Identifier (10 octets)
Ethernet Tag ID (4 octets)
Multicast Source Length (1 octet)
Multicast Source Address (variable)
Multicast Group Length (1 octet)
Multicast Group Address (Variable)
Originator Router Length (1 octet)
Originator Router Address (variable)
Leave Group Synchronization # (4 octets)
Max Response Time (1 octet)
Flags (1 octets) (optional)

NOTES

- This EVPN route type is used to coordinate IGMP Leave Group (x,G) state for a given [EVI, BD] between the PEs attached to a given ES
- The Flags field assists in distributing IGMP membership interest of a given host/VM for a given multicast route.
- The Eth Tag ID is set as follows
 - EVI is VLAN-Based or VLAN Bundle service - set to 0
 - EVI is VLAN-Aware Bundle service without translation - set to the customer VID for the [EVI, BD]
 - EVI is VLAN-Aware Bundle service with translation - set to the normalized Ethernet Tag ID for the [EVI, BD]

EVPN Route Type 8 : IGMP Leave Sync Route



IGMP Join/Sync Routes: Summary

- Efficient support for endpoints running IGMP in EVPN domain
- Terminate IGMP reports on the first hop EVPN PE and perform Proxy reporting performed at first hop to summarize (when possible) reports from downstream hosts or VMs. The BGP updates are translated back to IGMP messages by the EVPN PE on the other side.
- Reduce the flood of IGMP messages.

Route Type 6 - Selective Multicast Ethernet Tag route

- Efficiently forward multicast traffic over an EVPN network; only PEs interested in the traffic should receive it.
- On receiving SMET route, the receiving PE instantiates IGMP State with 'core' in the outgoing list.
- Multicast traffic not propagated to PEs which do not have interested receivers

Packet Format : SMET


RD (8 octets)							
Ethernet Tag ID (4 octets)							
Multicast Source Length (1 octet)							
Multicast Source Address (variable)							
Multicast Group Length (1 octet)							
Multicast Group Address (Variable)							
Originator Router Length (1 octet)							
Originator Router Address (variable)							
Flags (1 octets) (optional)							
Flags							
	RESERVED		IE	V3	V2	V1	

NOTES

- This EVPN route type is used to carry tenant IGMP multicast group information.
- The flag field assists in distributing IGMP membership interest of a given host/VM for a given multicast route.
- The version bits help associate IGMP version of receivers participating within the EVPN domain.
- The include/exclude bit helps in creating filters for a given multicast route

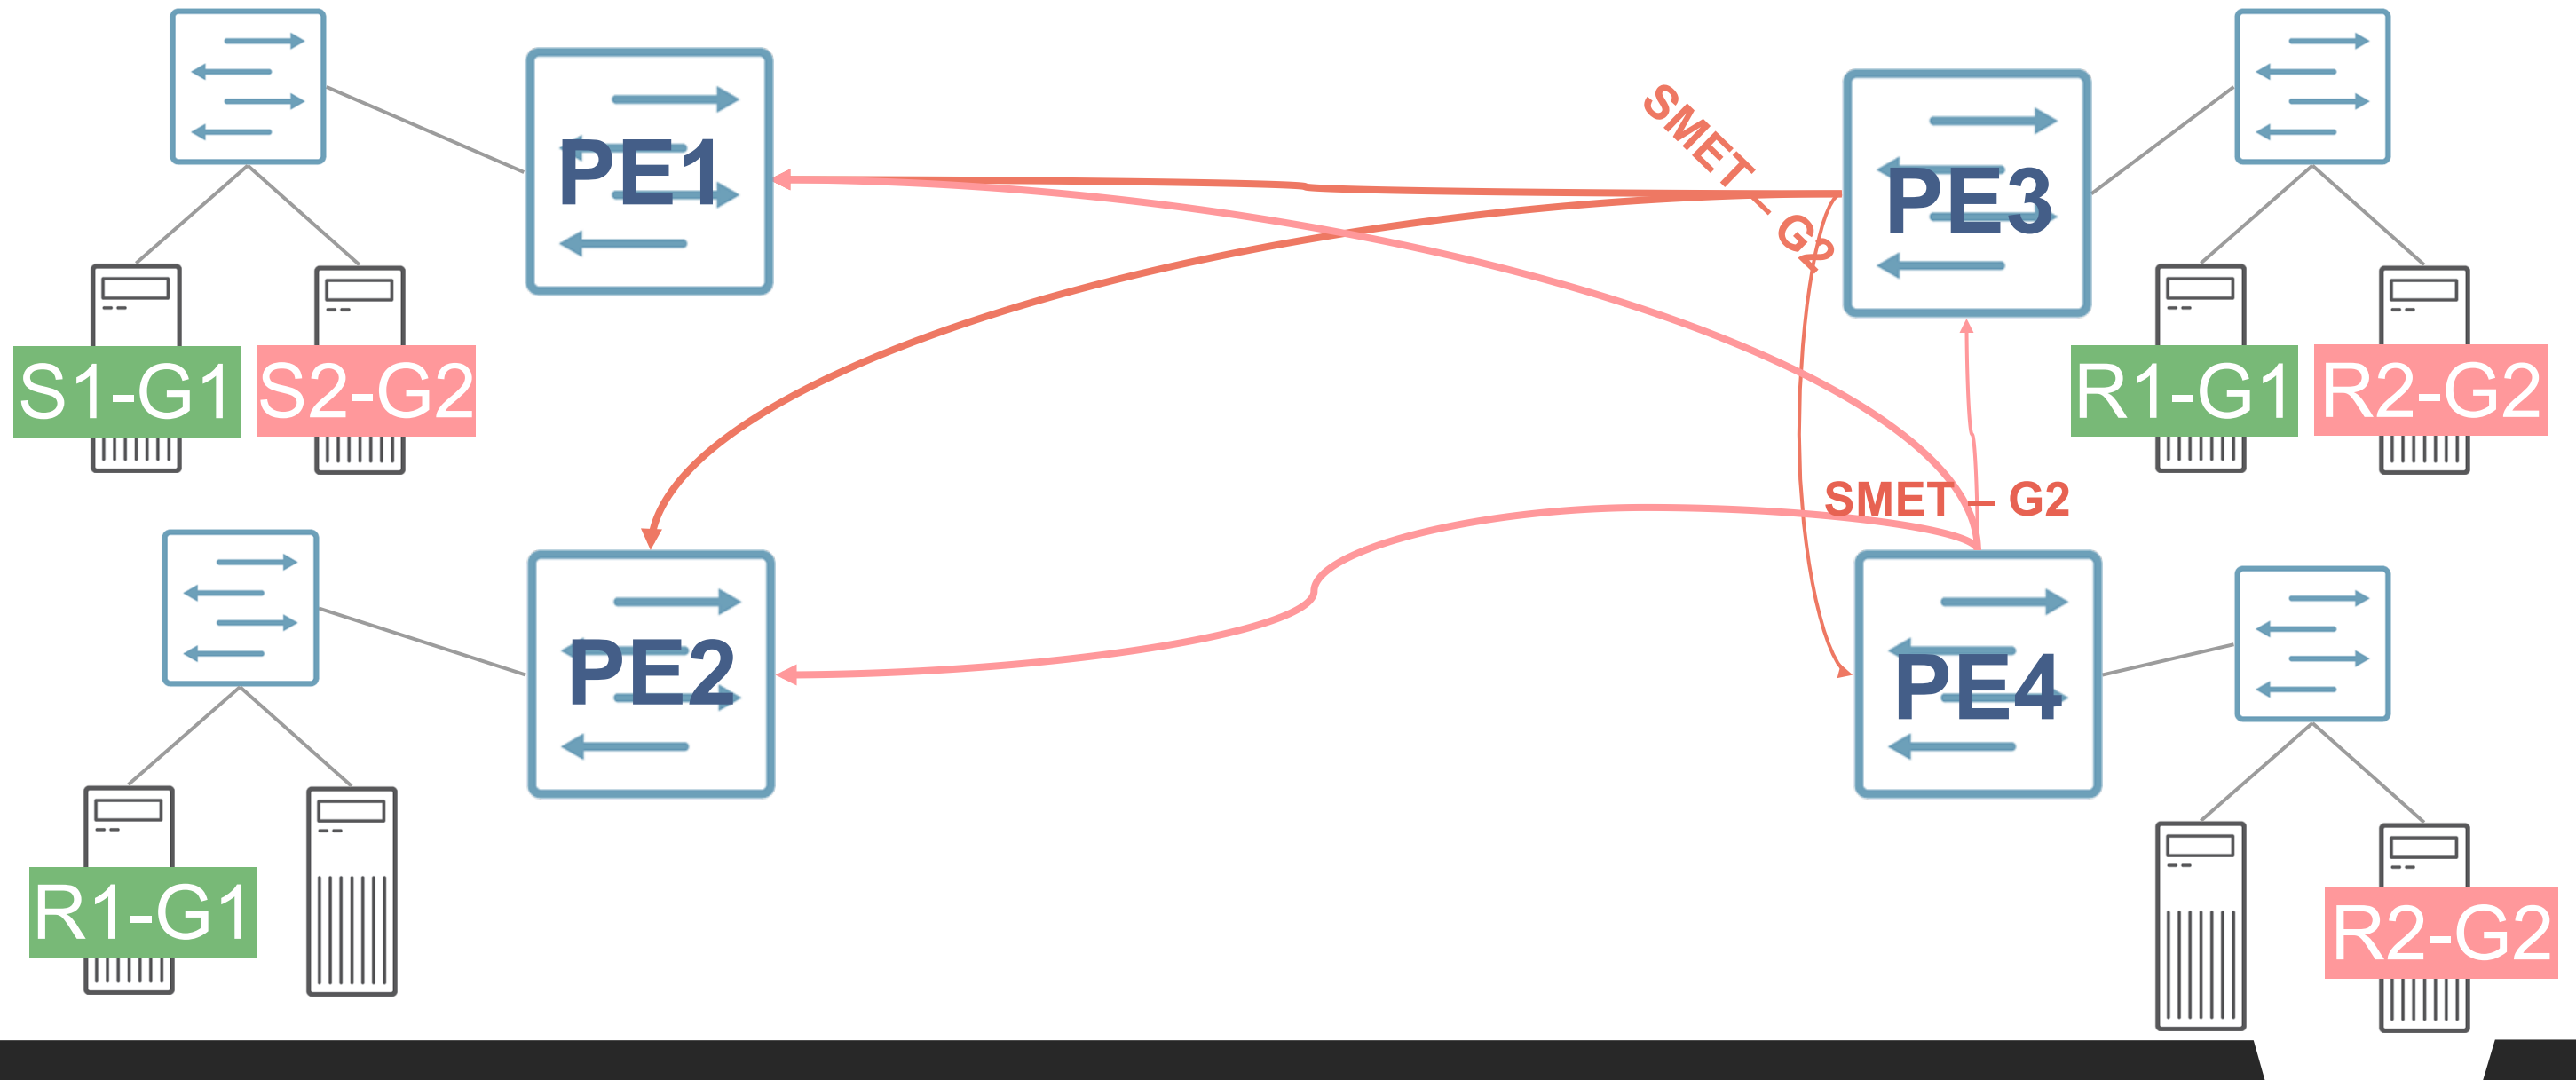
SMET Propagation for Group G1

PE2 and PE3 send SMET routes to all other PEs for group G1



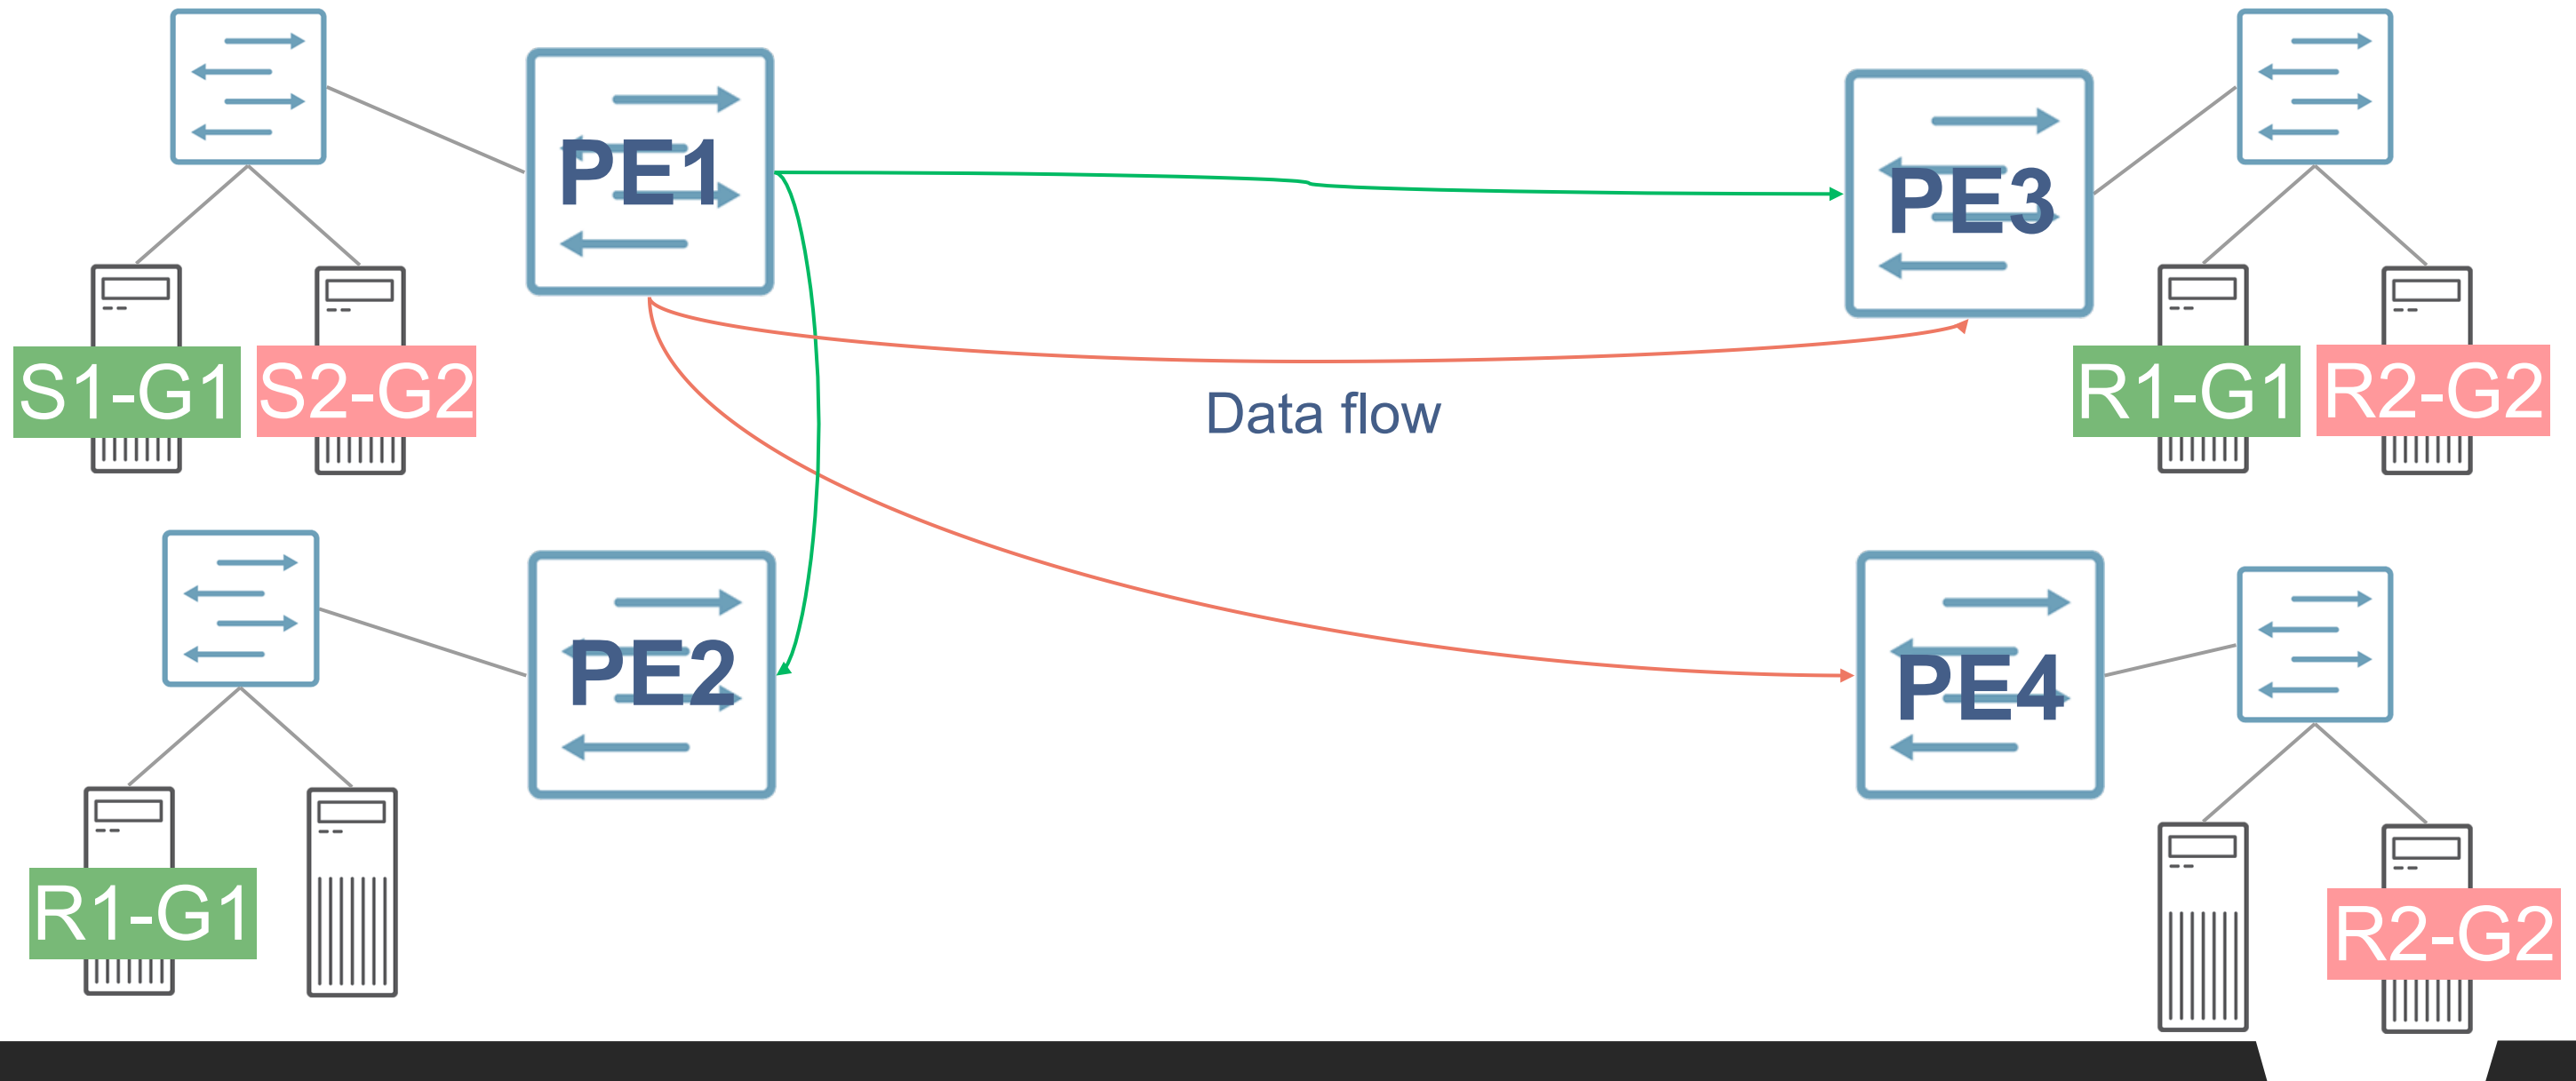
SMET Propagation for Group G2

PE3 and PE4 send SMET routes to all other PEs for group G2




Reduction of BUM in Core due to SMET

Source PE1 sends data to interested PEs ONLY for G1 and G2



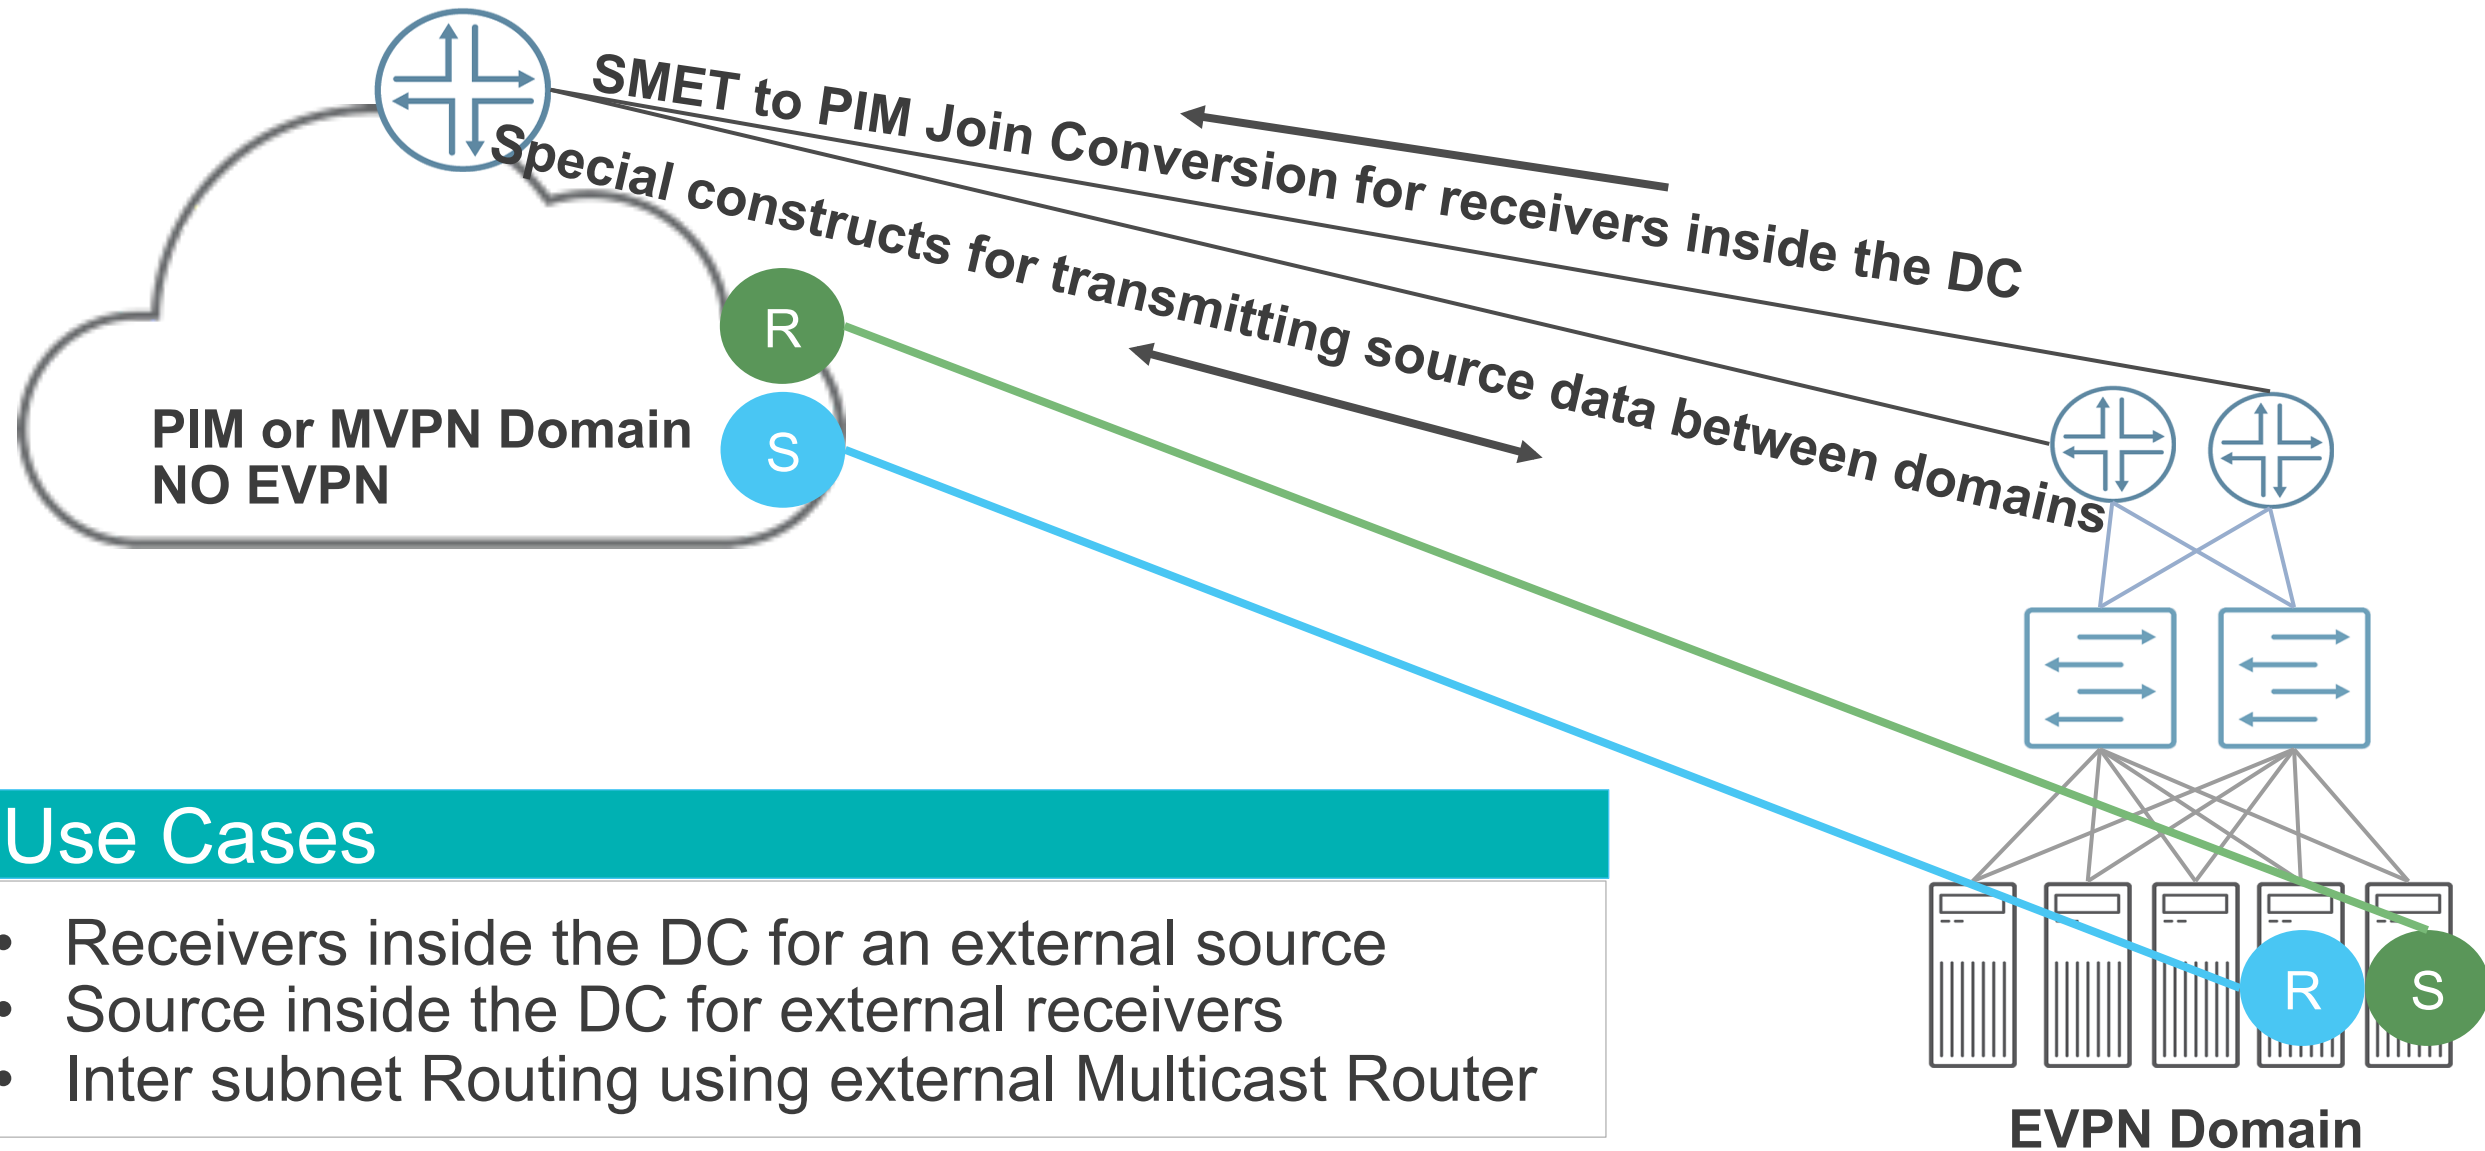
Selective P tunnels : Summary

- Most basic way to support replication is Ingress Replication
- EVPN integrates with P2MP LSPs to improve on Ingress Replication
- This process can be further improved by integration with SMET Route
- Ingress PE needs to discover the types of selective P-tunnels supported by the receiving PEs : Multicast Flags extended community defined for this purpose
- In a given EVI, a single broadcast domain may map to a single P Tunnel OR all broadcast domains may map to a single P Tunnel



Practical Applications

DC Fabric connection to external Multicast Domain





Final thoughts

Summary

- What we discussed today
 - <https://tools.ietf.org/html/draft-ietf-bess-evpn-igmp-mld-proxy-00>
- Optimization in Multi-access domain : Route Type 7 & 8 (IGMP Join/Leave)
- Optimizations in Core : Route Type 6 (SMET)
- Considerations for Inter subnet Routing for Multicast
 - <https://www.ietf.org/id/draft-lin-bess-evpn-irb-mcast-04.txt>



Thank you