

IEFT 108 BESS WG

IPv4 NLRI with IPv6 NH Use cases

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv4nlri-ipv6nh-use-cases/>



Gyan Mishra
July 28th 2020

With RFC 5549 NH encoding of IPv4 NLRI with IPv6 next hop used for iBGP PE-RR peering we can apply the same concept to all eBGP “Core” & “Edge” peering and utilize a “single protocol only” on all peering, IPv6 AFI=2.

Tremendous OPEX Saving to eliminate IPv4 eBGP peering. This basic concept will help alleviate IPv4 address depletion issues on IXP NAPs as well as both public & private peering pints for both Service Providers & Enterprises. This concept will also help pave the way as a stepping stone to eventual IPv6 Only customers as well as the eventual complete elimination of IPv4.

IPv4 AFI=1 eBGP Peering can be eliminated from all Edge PE-CE for both Enterprise and Service Provider networks.

BGP is a transport and so just as we stack SAFI’s on an AFI on our PE-RR iBGP peering in a typical MPLS or SR environment, we can do the same with stacking both AFI=1 & AFI=2 capability onto a **single IPv6 eBGP** peer. So basically the **IPv6 eBGP** peering has in the MP Reach capability exchange has advertised & received both AFI=1 & AFI=2 capabilities, and so now the IPv6 peer can advertise both IPv4 NLRI & IPV6 NLRI with IPv6 next hop encoding as an IPv6 address as defined in RFC 5549 update.

With IPv4 address depletion issues with Internet IXP exchange points which have come up in recent at NANOG 65 Montreal 2015, have proposed using RFC 5594 IPv4 NLRI encoding in IPv6 NH as a way to eliminate all IPv4 peering at the IXP POP’s to save on address space as well as OPEX expenditure in maintaining both IPV4 & IPV6 peering.

****Softwire mesh framework RFC 5565 tunneling v6 over v4 transport data plane using IPv4 Next Hop encoding over MPLS LDPv4 core** (6PE scenario)**

In an Enterprise Dual stacked customers edge environment over IPv4 MPLS “BGP free” core, where 6PE is used to connect IPv6 islands over an IPv4 core, PE-RR core peering AFI/SAFI 1/1 and 6PE BGP-LU IPv6 labeled unicast AFI/SAFI 1/4. **PE-RR 6PE requires IPv6 NLRI to be carried with IPv4 Next Hop encoding.**

****Softwire mesh framework RFC 5565 tunneling v6 over v4 transport data plane using IPv4 Next Hop encoding over MPLS LDPv4 core** (IP VPN scenario)**

In a Service provider Dual stacked customers edge environment over IPv4 MPLS “BGP free” core at the edge, we maintain the Dual stacked PE-CE IPv4 peering AFI/SAFI 1/1 (IPV4) & 2/1 (IPV6). In the core PE-RR peering we have AFI/SAFI 1/128 (VPN-IPV4) 1/129 (MVPN) 2/128 (VPN-IPV6) 2/129 (MVPN). **PE-RR peering requires 2/128 (VPN-IPV6) 2/129 (MVPN) NLRI to be carried over IPv4 Next Hop encoding.**

With Enterprise or Service provider Dual stacked customers edge environment in Green & Brown field deployments of IPv6 transport Core using MPLS LDPv6, SR-MPLSv6 or SRv6 IPv6 data plane with “BGP free” core, we can now eliminate separate v4 & v6 Edge PE-CE peering so all SAFI related to AFI=1 can now be carried over AFI=2.

****Softwire mesh framework RFC 5565 tunneling v4 over v6 transport data plane using RFC 5549 NH encoding over MPLS LDPv6 or SRv6 core** (4PE scenario)**

With Enterprise or Service provider Dual stacked customers edge environment in Green & Brown field deployments of IPv6 transport Core using MPLS LDPv6, SR-MPLSv6 or SRv6 IPv6 data plane with “BGP free” core, we can now eliminate separate v4 & v6 peering so all SAFI related to AFI=1 can now be carried over AFI=2.

In an Enterprise Dual stacked customers edge environment over an IPv6 MPLS LDPv6, SR-MPLSv6 or SRv6 IPv6 core, where 4PE is used to connect IPv4 islands over a IPv6 core, PE-RR peering ↔ BGP-LU IPV4 labeled unicast (4PE) AFI/SAFI 1/4 to connect IPv4 islands over IPv6 core, carrying AFI/SAFI 2/1 **would now also carry AFI/SAFI 1/1 & 1/4 4PE IPv4 Labeled unicast with RFC 5549 IPv6 Next Hop encoding.**

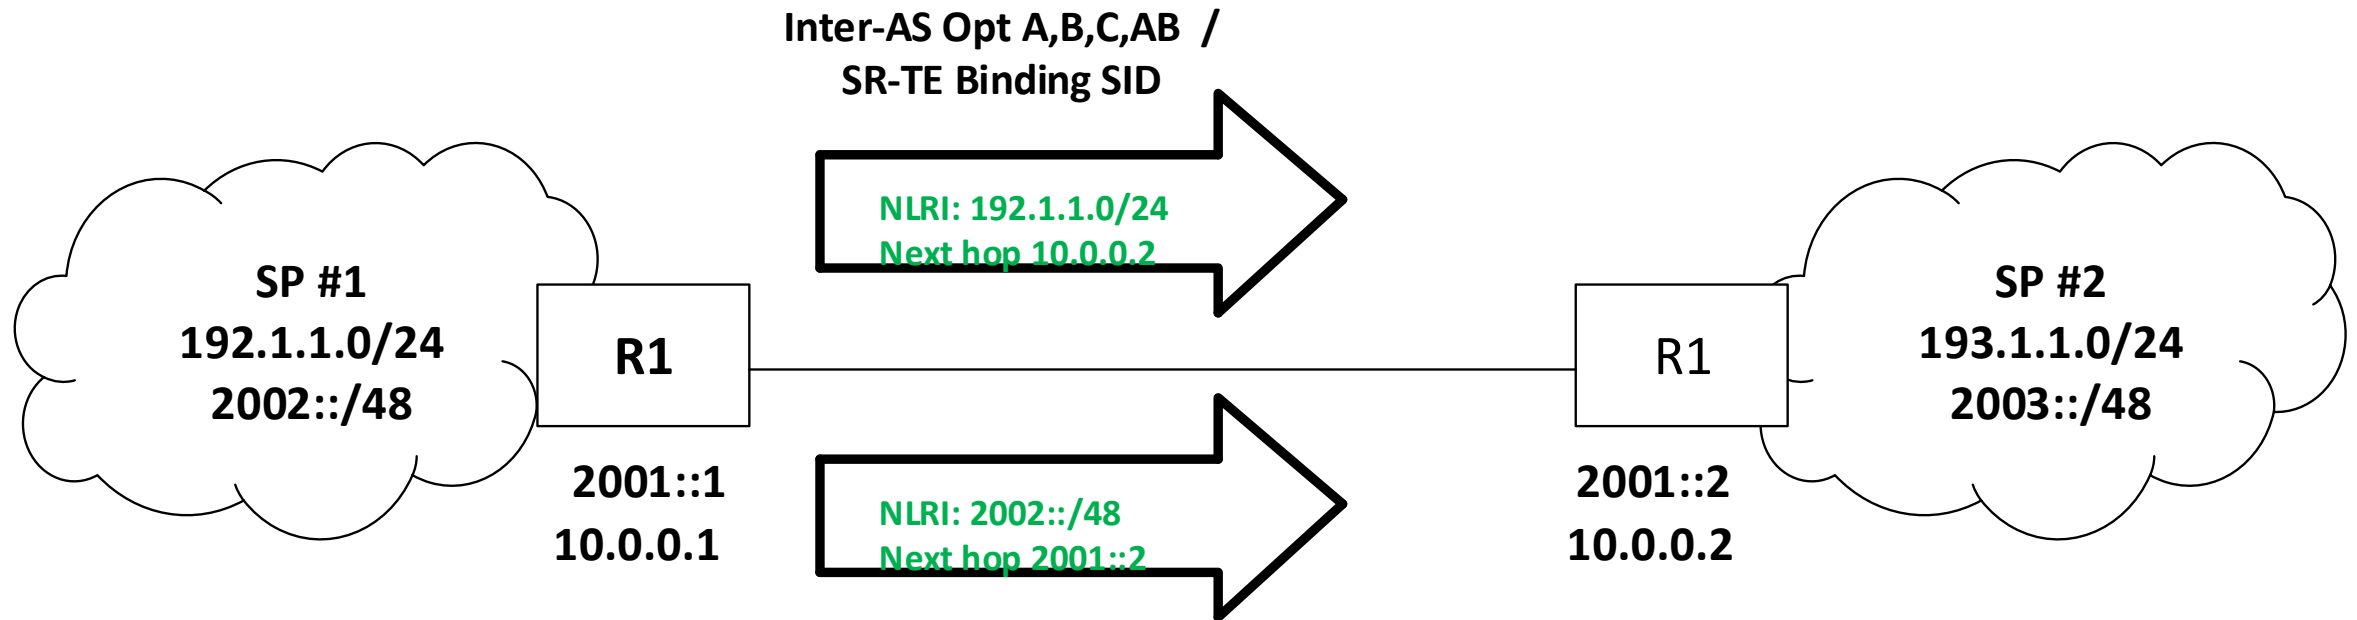
At the PE-CE **eBGP** customer edge now AFI/SAFI 1/1 & 2/1 would now both be carried by a single protocol **IPv6**, with IPv6 Next Hop defined in **RFC 5549 draft update for the IPv6 Next Hop encoding.**

****Softwire mesh framework RFC 5565 tunneling v4 over v6 transport data plane using RFC 5549 NH encoding over MPLS LDPv6 or SRv6 core** (IP VPN scenario)**

In a Service provider Dual stacked customers edge environment over an IPv6 MPLS LDPv6, SR-MPLSv6 or SRv6 IPv6 core, with core PE-RR peering, IPv6 peering PE-RR already carrying IPv6 NLRI - 2/128 (VPN-IPV6) 2/129 (MVPN) **would now carry as well 1/128 (VPN-IPV6) 1/129 (MVPN) with RFC 5549 draft updated IPv6 NH encoding.**

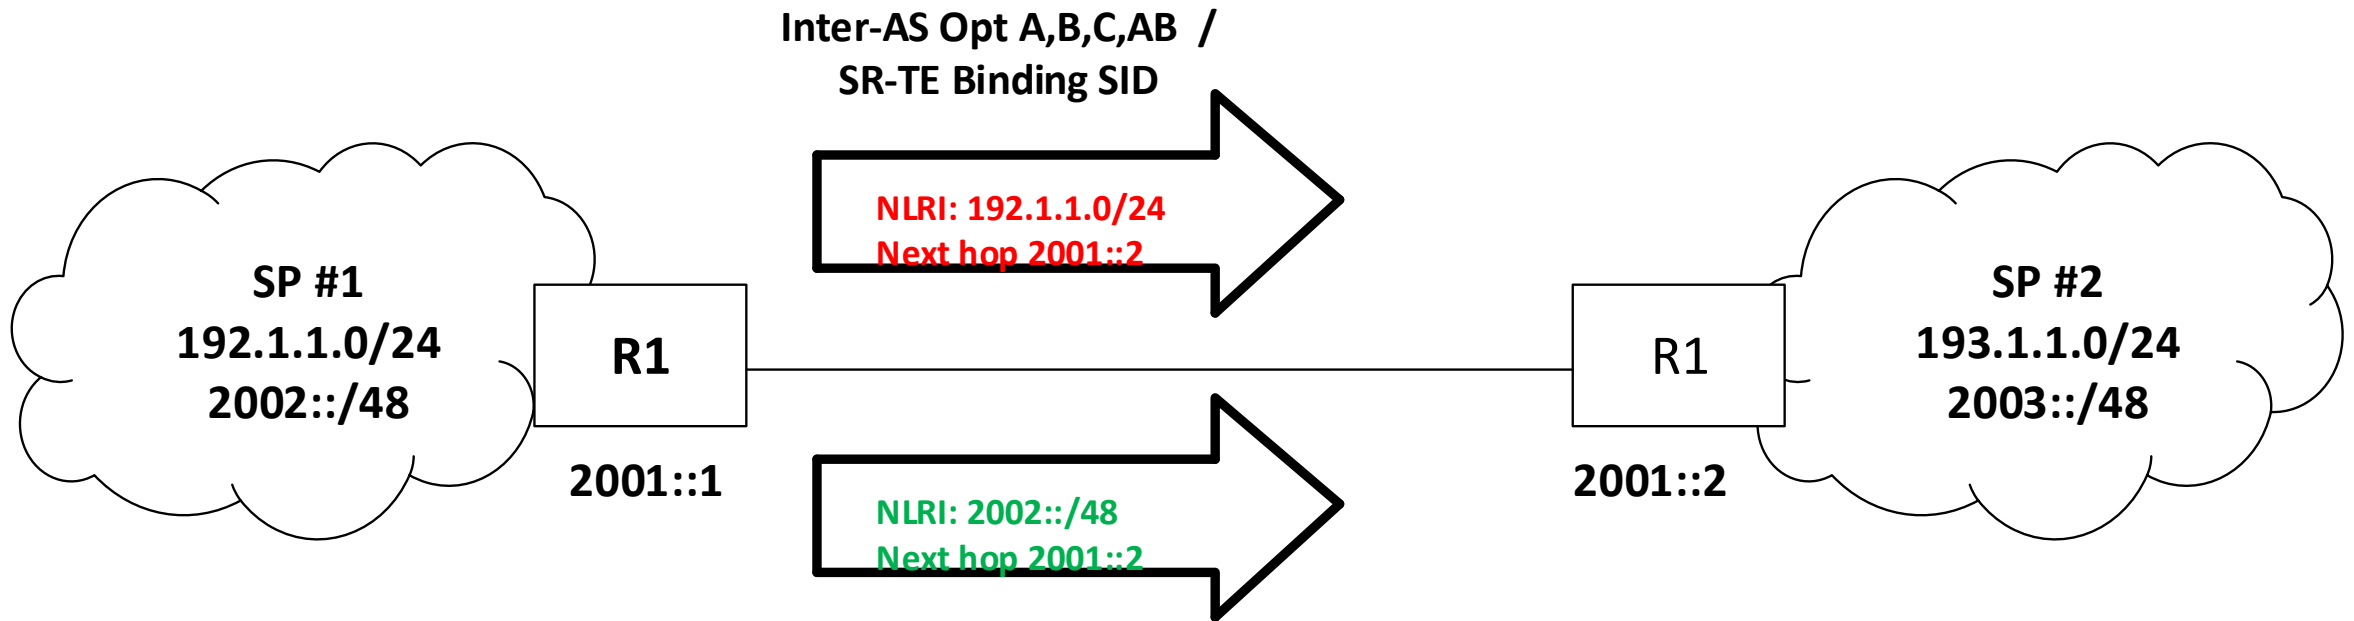
At the PE-CE **eBGP** customer edge now AFI/SAFI 1/1 & 2/1 would now both be carried by the single protocol **IPv6**, with IPv6 Next Hop defined in **RFC 5549 draft update for the IPv6 Next Hop Encoding.**

**IXP Peering has a separate IPv4 & IPv6 peer
with IPv4 NH peer carrying IPv4 NLRI & IPv6
NH peer carrying IPv6 NLRI**



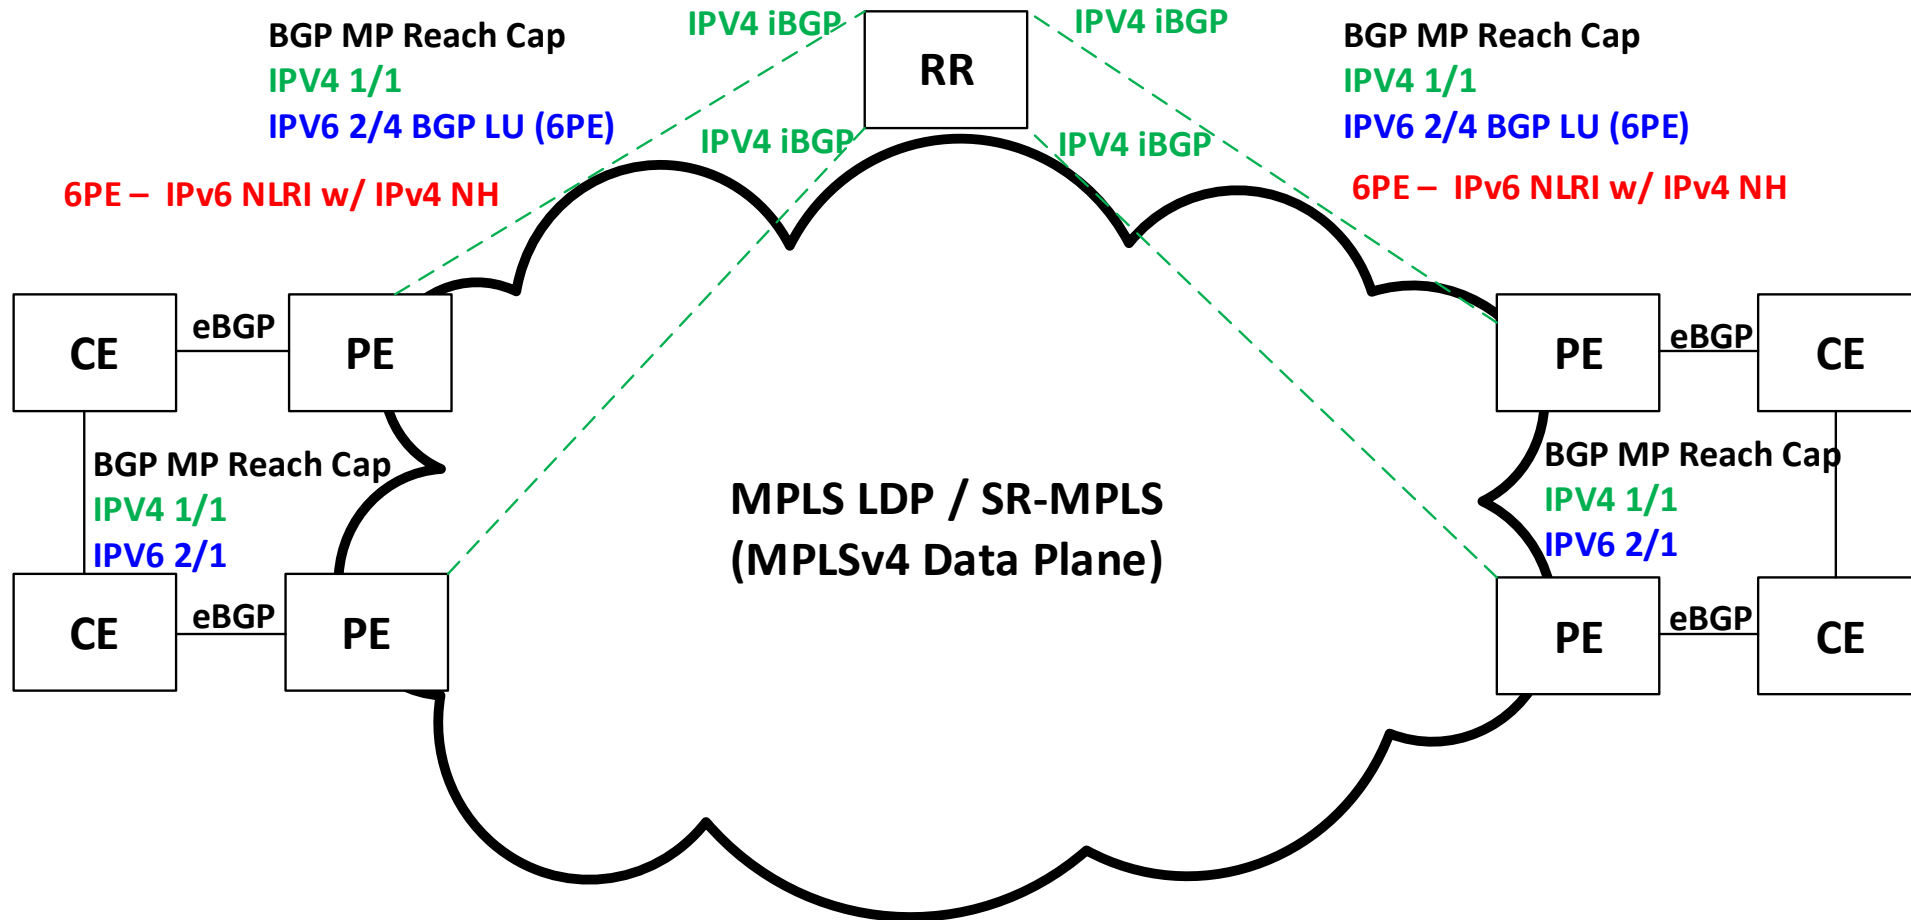
So now with RFC 5549 NH encoding schema of 16 / 32 byte IPv6 next hop both IPv4 & IPv6 NLRI can be advertised using a single IPv6 peer.

This basic concept can eliminate all IPv4 peering at the Edge and within the Core.

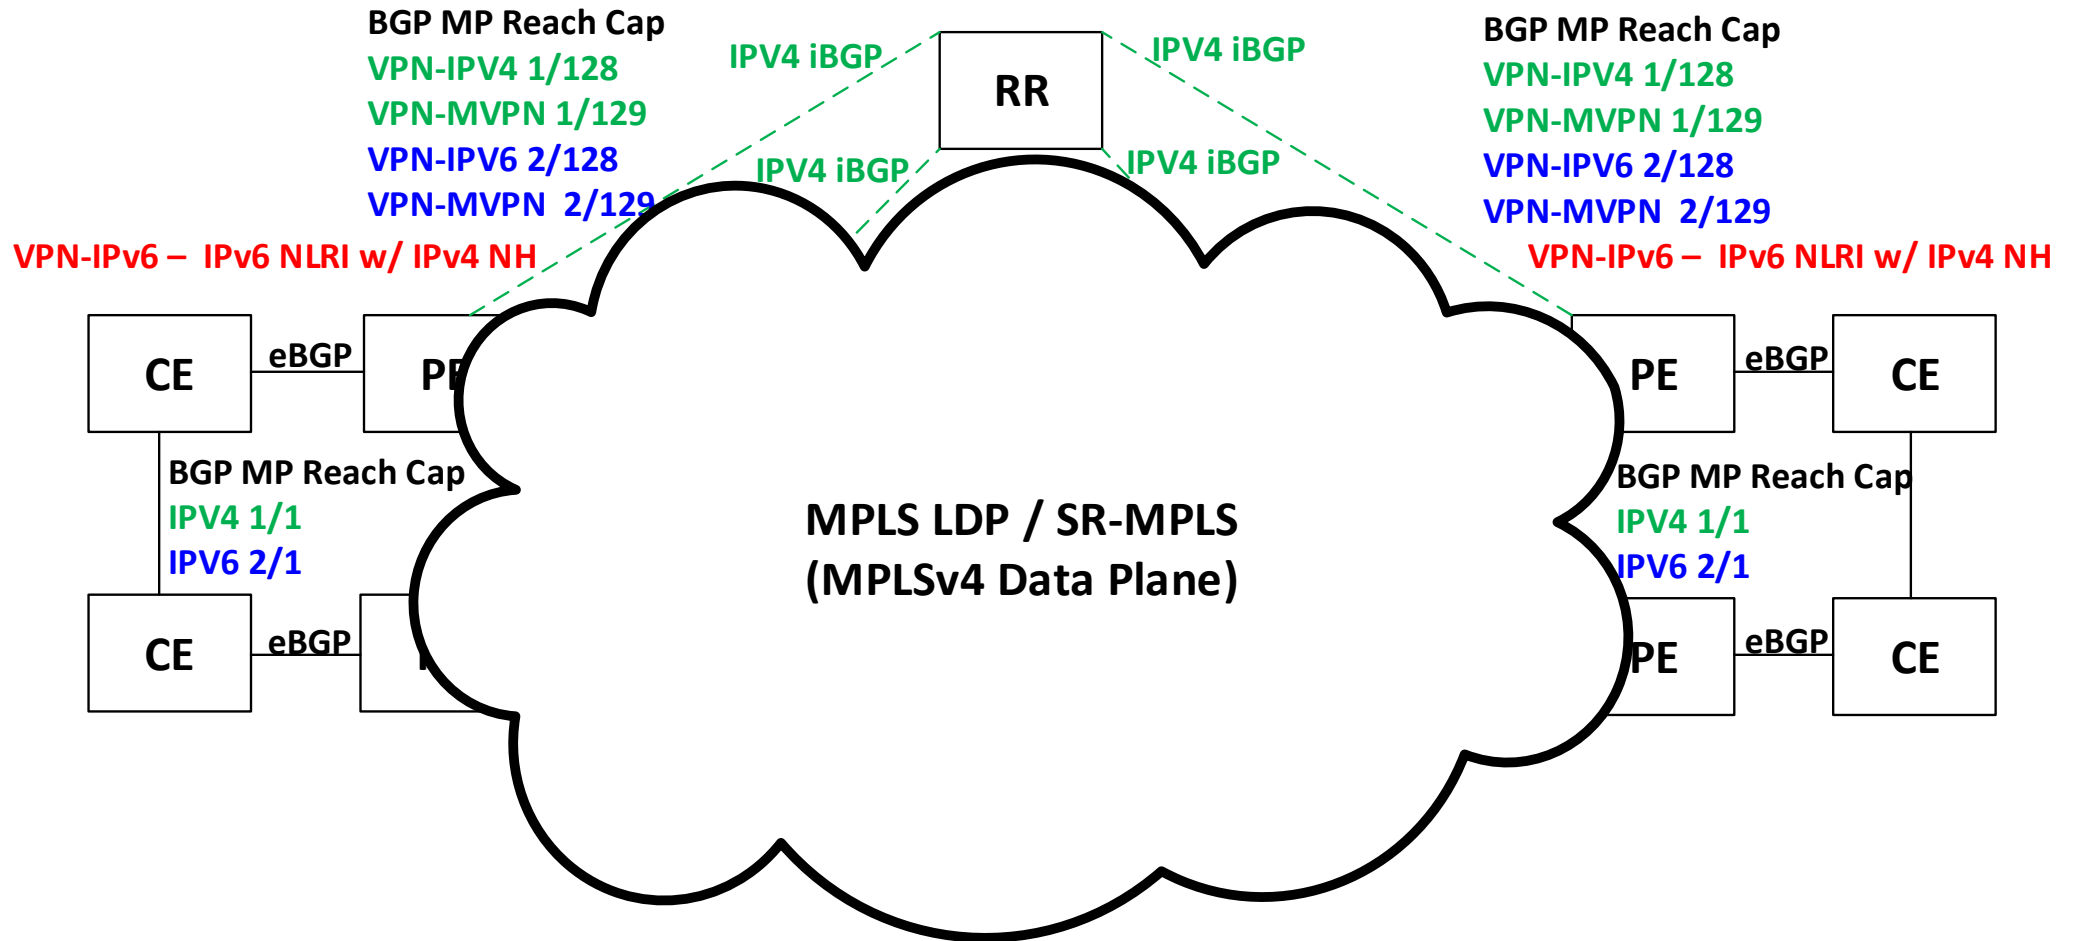


IXP Peering of all SPs would now be able to use a single IPv6 peer per SP adjacency

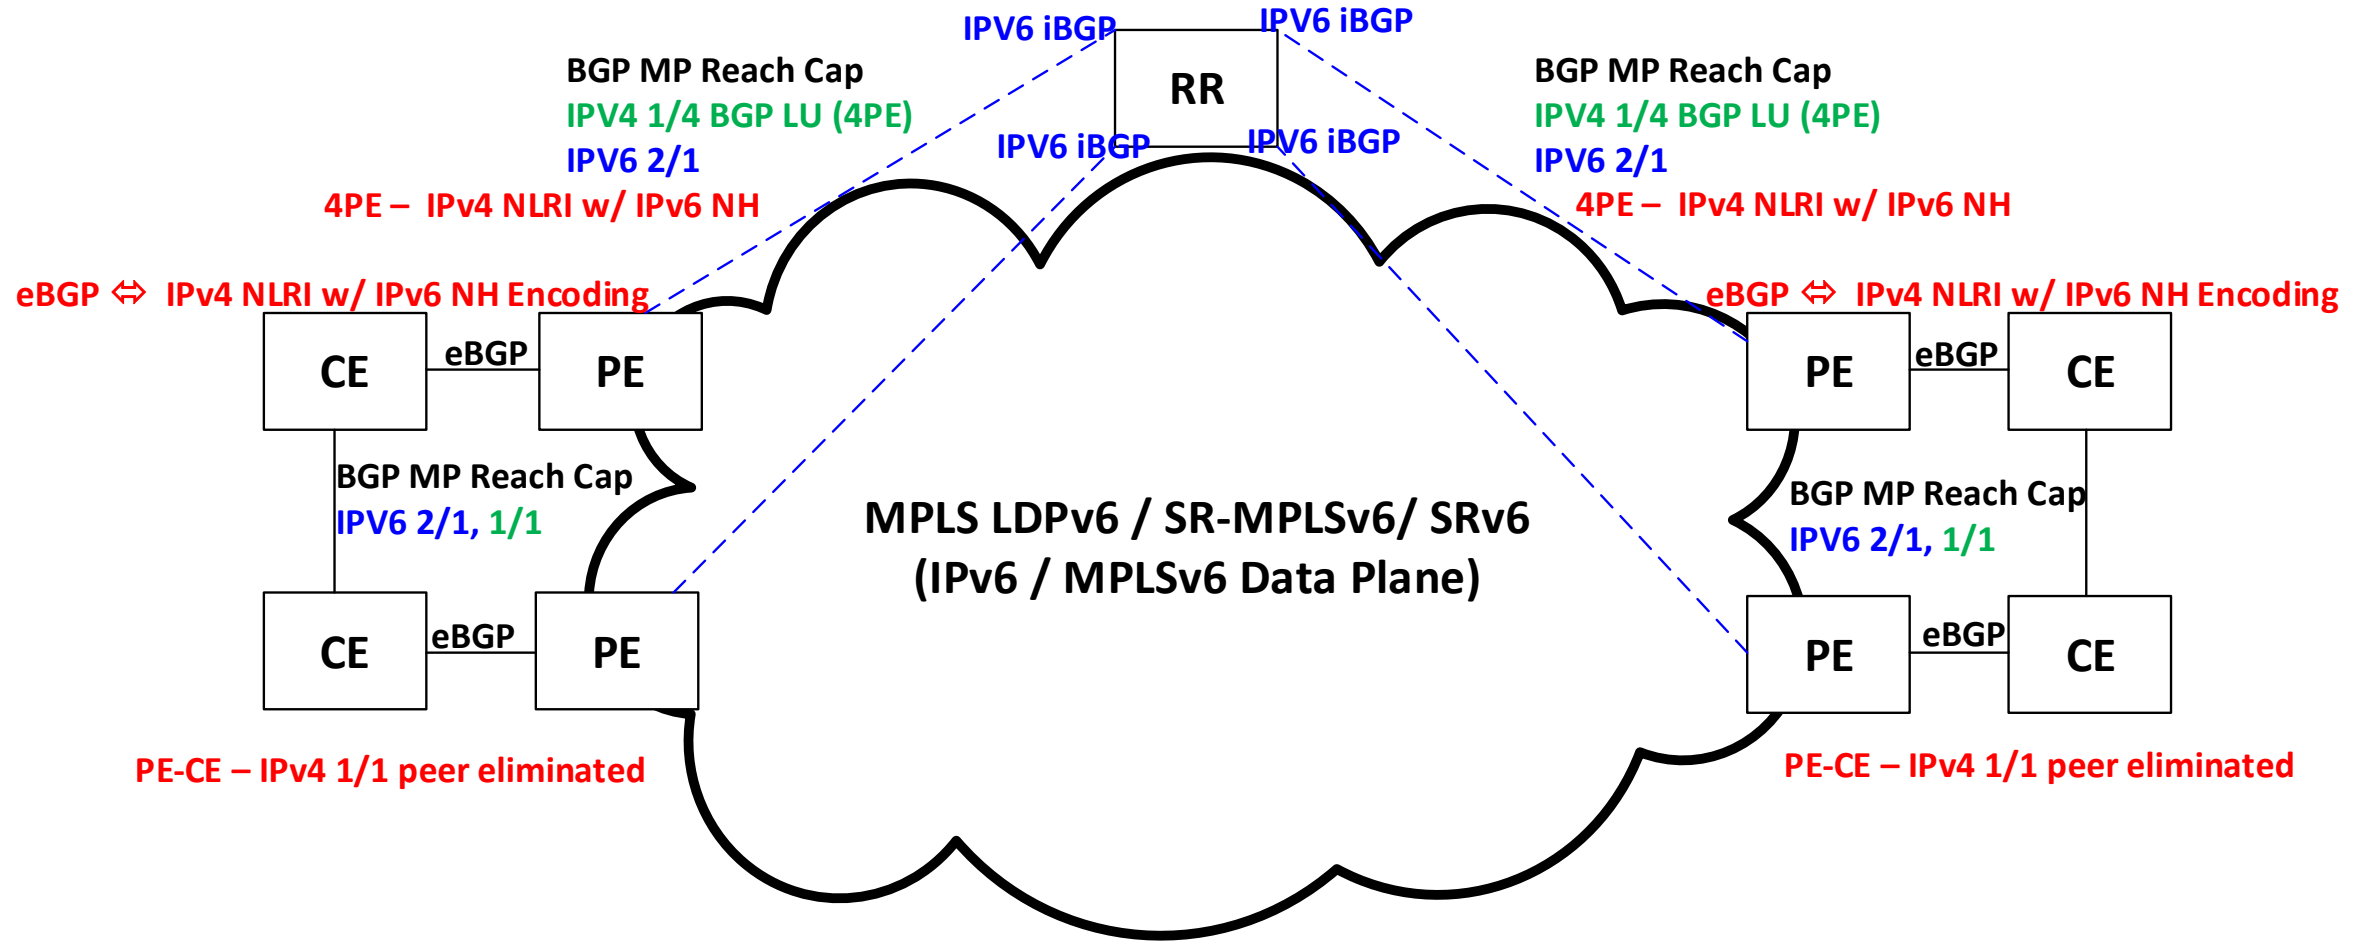
MPLS / SR-MPLS Core– SOFTWARE MESH FRAMEWORK 6to4 (6PE)



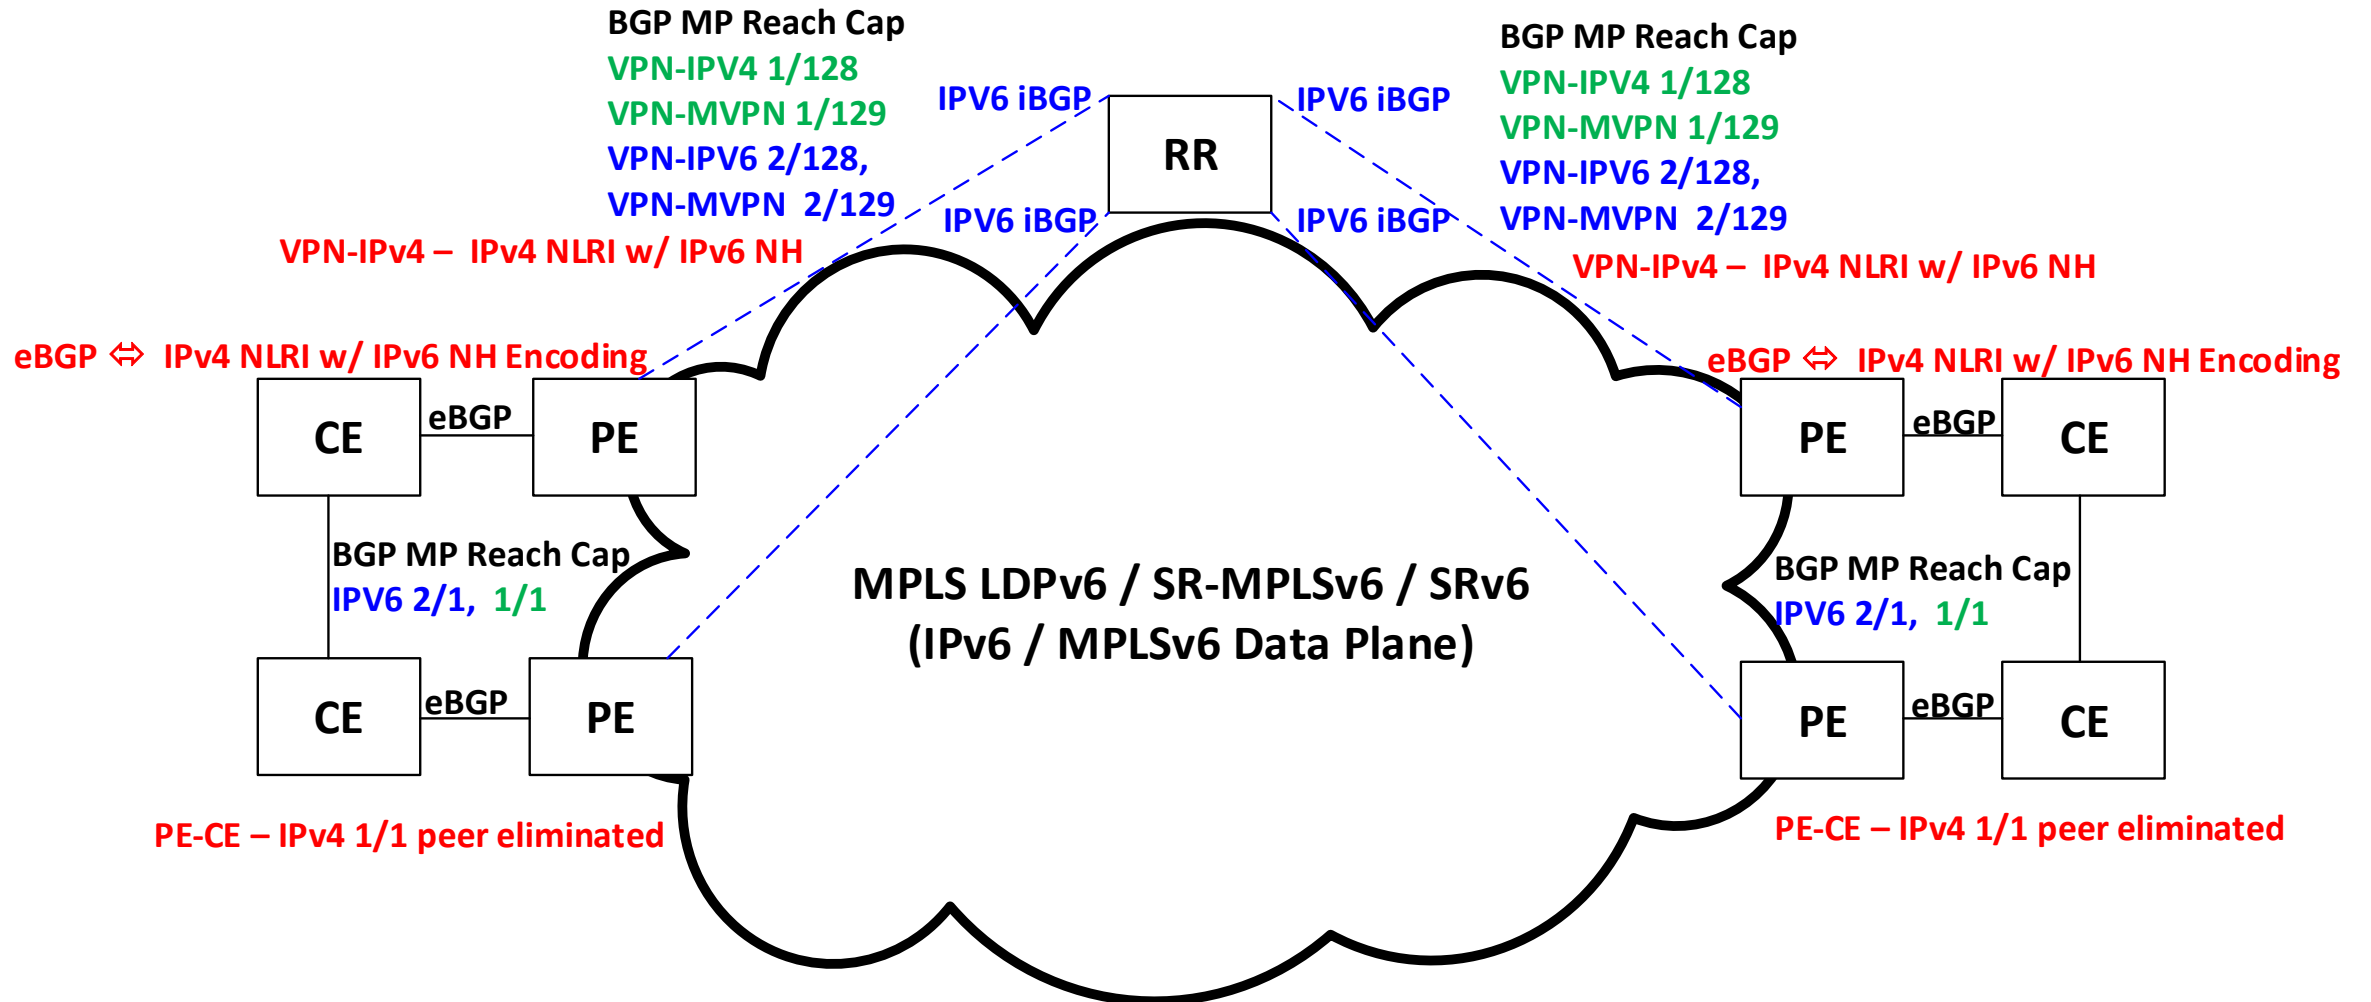
MPLS / SR-MPLS Core – SOFTWARE MESH FRAMEWORK 6to4 (IP VPN)



MPLS LDPv6 / SR-MPLSv6 / SRv6 Core – SOFTWARE MESH FRAMEWORK 4to6 (4PE)



MPLS LDPv6 / SR-MPLSv6 / SRv6 Core – SOFTWARE MESH FRAMEWORK 4to6 (IP VPN)



Q&A

THANK YOU