Connect IPv4 Islands over IPv6 Core (4PE)

draft-mishra-idr-v4-islands-v6-core-4pe-05

IETF 117

Gyan Mishra, Verizon
Jeff Tantsura, Microsoft
Mankamana Mishra, Cisco
Sudha Madhavi, Juniper
Adam Simpson, Nokia
Shuanglong Chen, Huawei

4PE Draft Updates

- Customer IPv4 prefixes must be labeled when tunneled over IPv6 LSP. If not labeled on the PHP node when the topmost label is popped the native IPv4 prefix is exposed and is not routable and will be dropped unless RFC 3270 Pipe mode explicit null signaling is enabled, which may not be always the case for customers wanting to use PHP signaling implicit null. For IPv6 prefixes over an IPv6 core or IPv4 prefixes over IPv4 core at PHP node the IPv4 or IPv6 prefix can still be routed since the protocol of the tunneled prefix matches underlay protocol. Not the case for 4PE ⇔ protocol mismatch between 2 level label stack topmost IPv6 and BOS S bit IPv4.
- Label stack MUST be 2 Level Label Stack is only supported. This is for interoperability as additional labels could be added for flexibility to the specification but that could break interop. It is possible for the IPv4 prefixes to not be labeled but that is not supported.
- RFC 7948 “6PE” states that the label bound to the IPv4 prefixes may be an arbitrary value or explicit null label which has led to vendor interoperability issues in the past. 4PE draft states that the IPv4 prefixes must use MPLS QOS RFC 3270 Pipe mode explicit null label bound to the IPv4 prefix and MUST be used to signal from egress 4PE to ingress 4PE router that the packet is an IPv4 packet to identify the IPv4 routing context or outgoing interface to forward the packet.
- When RFC 7948 “6PE” was written when Segment Routing did not exist. The 4PE draft provides a detailed interworking of how 4PE is implemented with Segment Routing both SR-MPLS & SRv6. I have cleaned up the related text in the draft on Segment Routing support to make it more clear.
- Additional text clarity added related RFC 8950 next hop encoding interaction with 4PE and the importance of 4PE procedures and that RFC 8950 is strictly about the next hop encoding of IPv4 NLRI over an IPv6 next hop peer. Also acknowledged comments related to alternatives to 4PE that exist to connect IPv4 islands over an IPv6 core and why a standardized 4PE specification is the desired solutions as compare to alternatives that exist today.
- Cleaned up the draft to make it more readable.
- Updated contributors section and many thanks for all the feedback on the draft.

WG Adoption??

Thank You!

Motivation for a Standards Track 4PE Specification

IETF standard exists for connecting IPv6 islands over an IPv4 core **(RFC 4798)**, however a standard does **NOT exist** for connecting IPv4 islands over an IPv6 core.

This draft provides the specification for connecting **IPv4 islands over an IPv6 core**.

As operators migrate to a single protocol IPv6-Only core per RFC 5565 **Softwire Mesh Framework** which involves **6to4 tunnel** of IPv6 packets over an IPv4 core called **“6PE”**, and now with this draft **4to6 tunnel** of IPv4 packets over an IPv6 core now called **“4PE”**.

The name **“6PE”** termed to define the tunneling of **IPv6 labeled packets** over an **IPv4 core** and now the name **“4PE”** is termed to define the tunneling of **IPv4 labeled packets** over an **IPv6 core**.

Recap of 4PE

- 4PE routers exchange IPv4 reachability **transparently tunneled** over an IPv6 core using MP-BGP IPv6 RFC 2545 using the BGP next hop field to convey the IPv6 address of the 4PE router so that the dynamically established **IPv6 signaled MPLS LSP** can be utilized without explicit tunnel configuration. **(Signal Topmost transport label LSP)**
- 4PE uses RFC 8950 for the **16 or 32 byte next hop** encoding.
- Ingress & Egress 4PE routers must **bind a label to all the IPv4 prefixes** per RFC 8277 BGP-LU (Service Label BOS (Bottom of Stack) S bit set. **2 Level label stack**.
- 4PE supports **Explicit Null Signaling** for Diff-Serv PIPE mode model.
- 4PE design supports RFC 4364 Inter AS Option A, B, C, AB.
- 4PE design supports MPLS, SR-MPLS & SRv6 data planes.

RFC 4798 –Connecting IPv6 islands over IPv4 MPLS using IPv6 Provider Edge Routers (6PE)

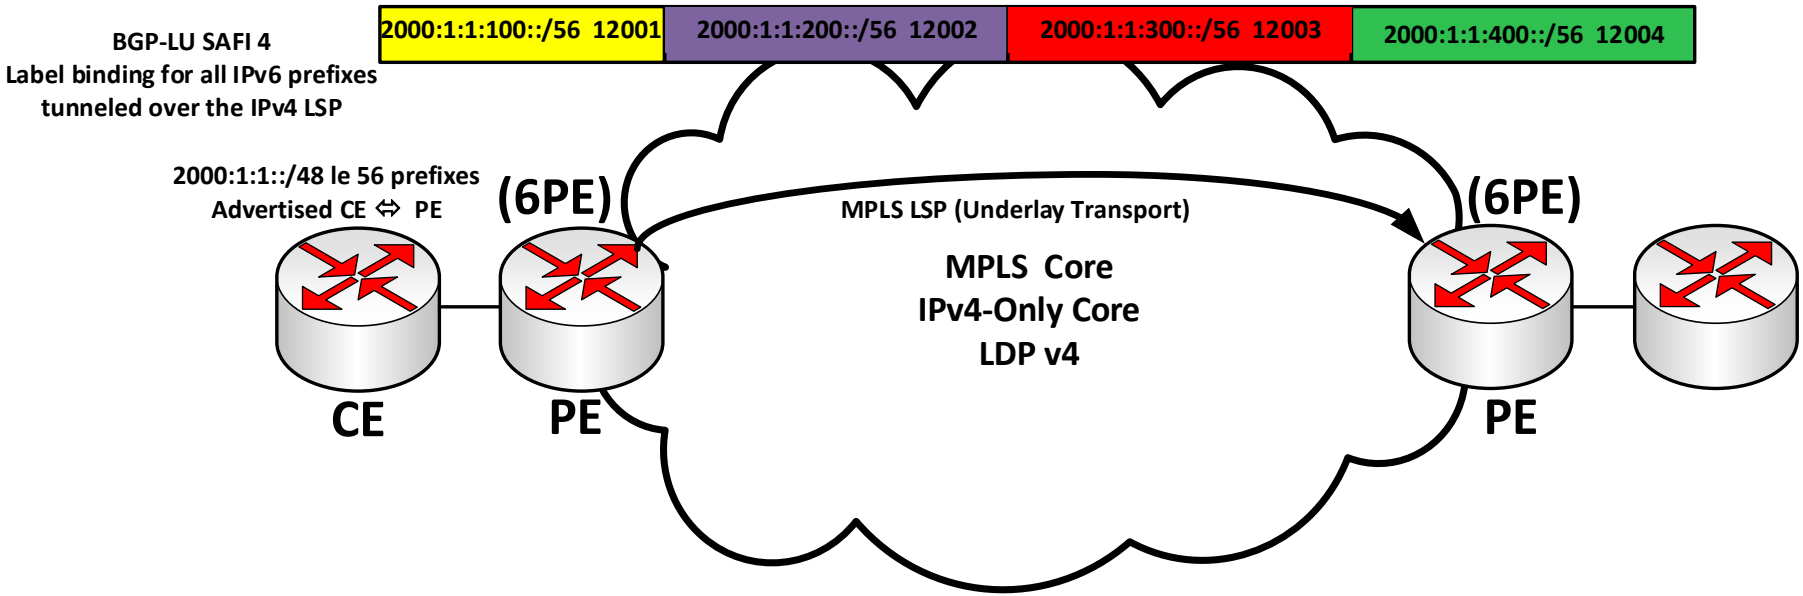

(4PE) –Connecting IPv4 islands over IPv6 MPLS using IPv4 Provider Edge Routers

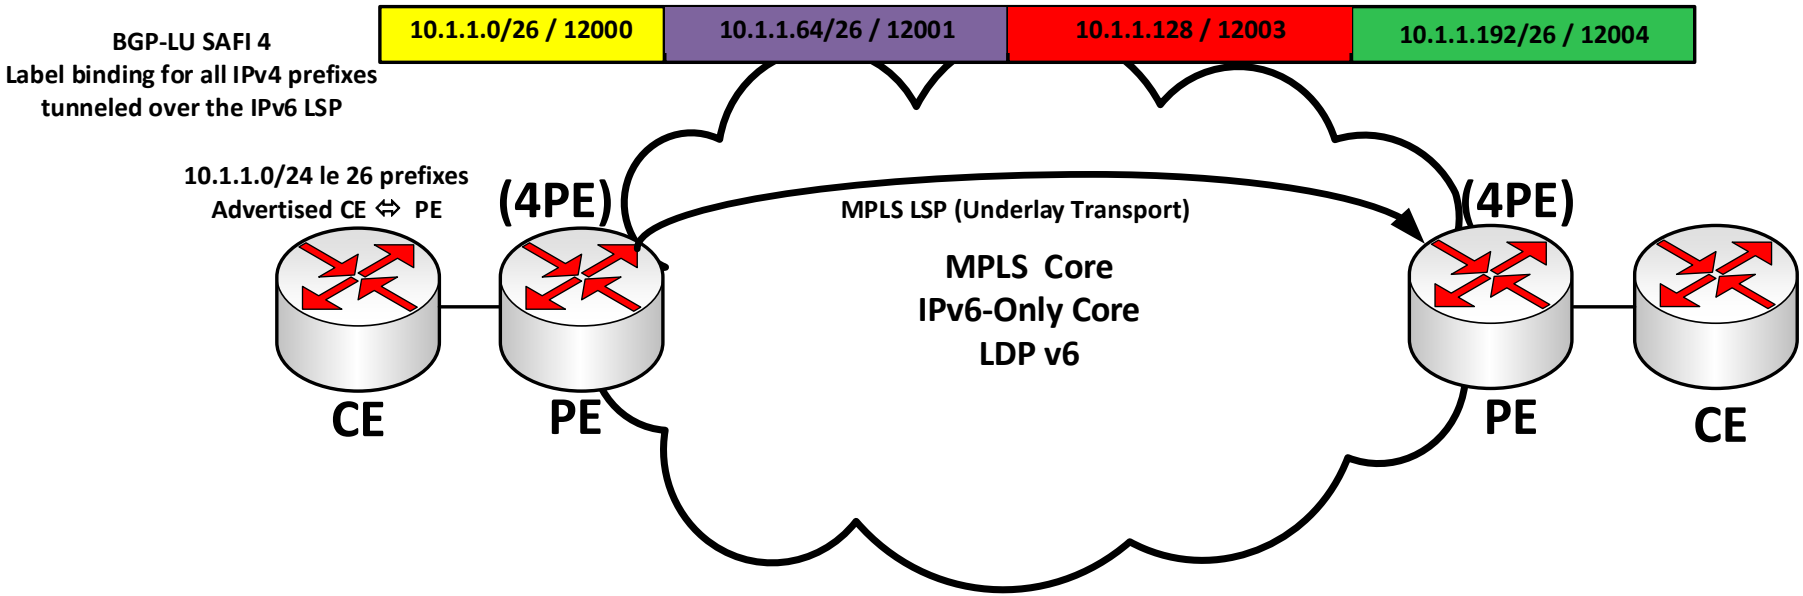

(4PE) –Connecting IPv4 islands over IPv6 SR-MPLS using IPv4 Provider Edge Routers

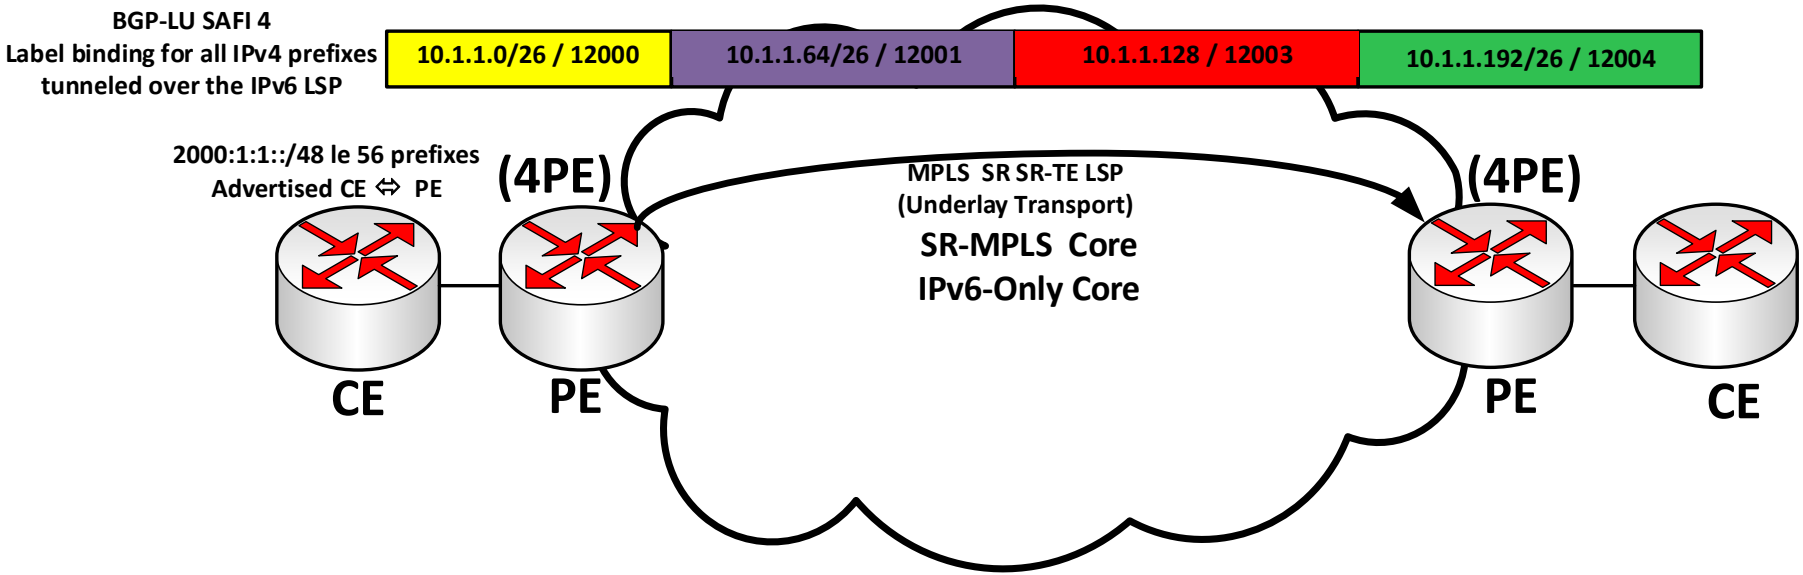

(4PE) –Connecting IPv4 islands over IPv6 SR-MPLS using IPv4 Provider Edge Routers

RFC 9252 BGP SRv6 Service
BGP Prefix-SID Attribute
SRv6 L3 Service TLV
MPLS labels Transposition
to Func/Arg
Section 5.3

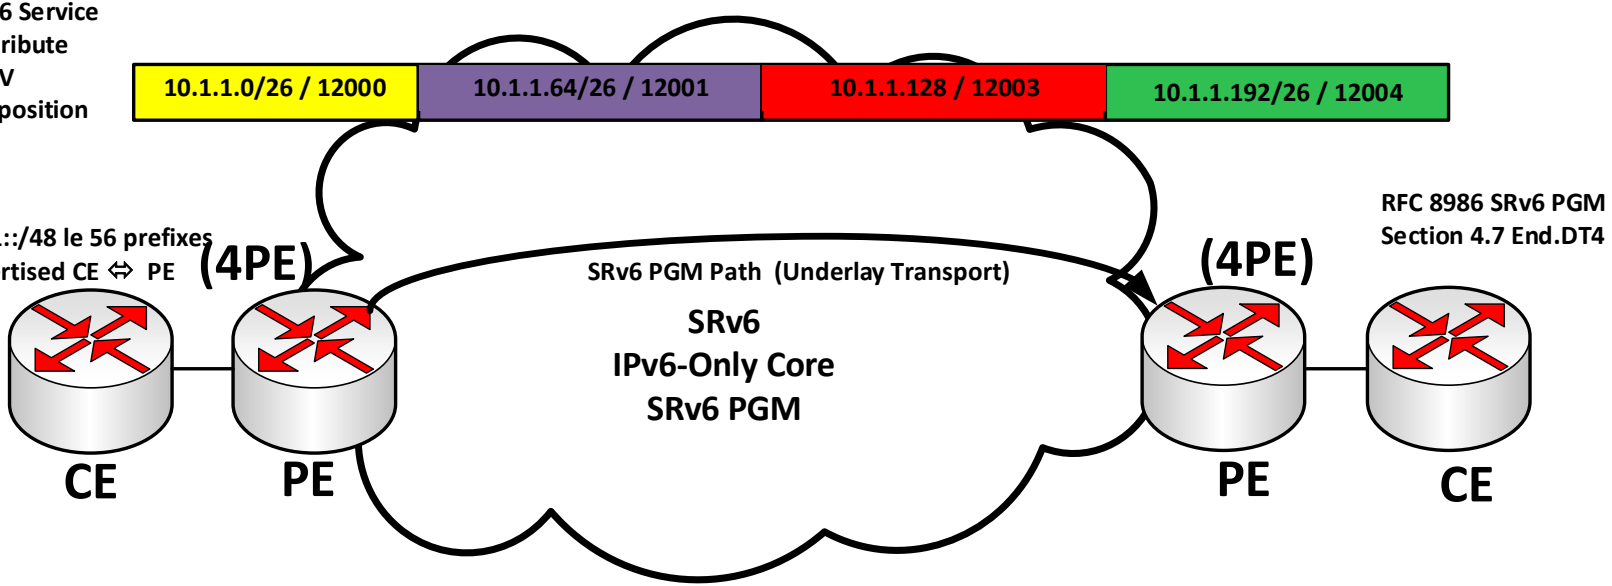

RFC 8986 SRv6 PGM
Section 4.7 End.DT4

(4PE) –Control Plane & Data Plane Intra-AS

MPLS LDPv6 / SR-MPLSv6 / SRv6 – Softwire Mesh Framework 4to6 (4PE)
(4PE Control Plane)
“

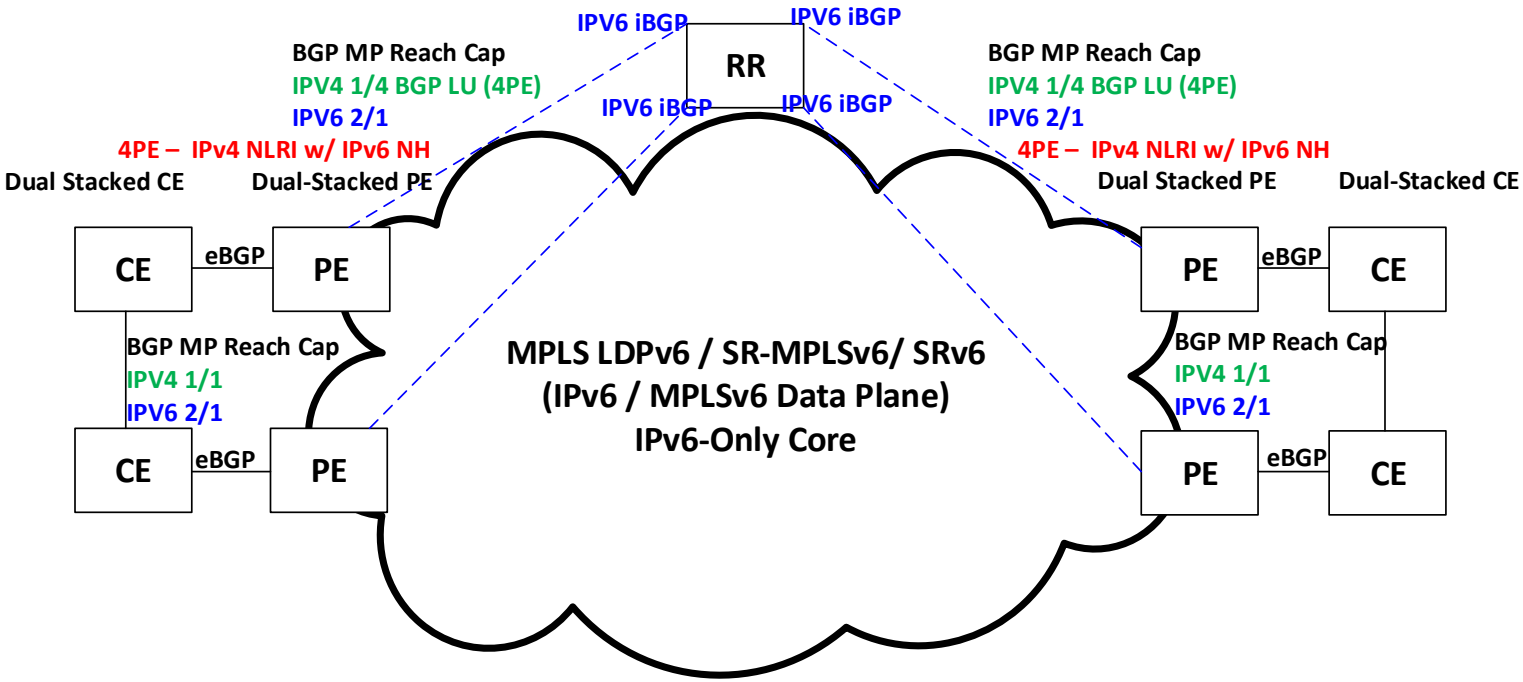

(4PE) –Control Plane & Data Plane Inter-AS Option A Procedure

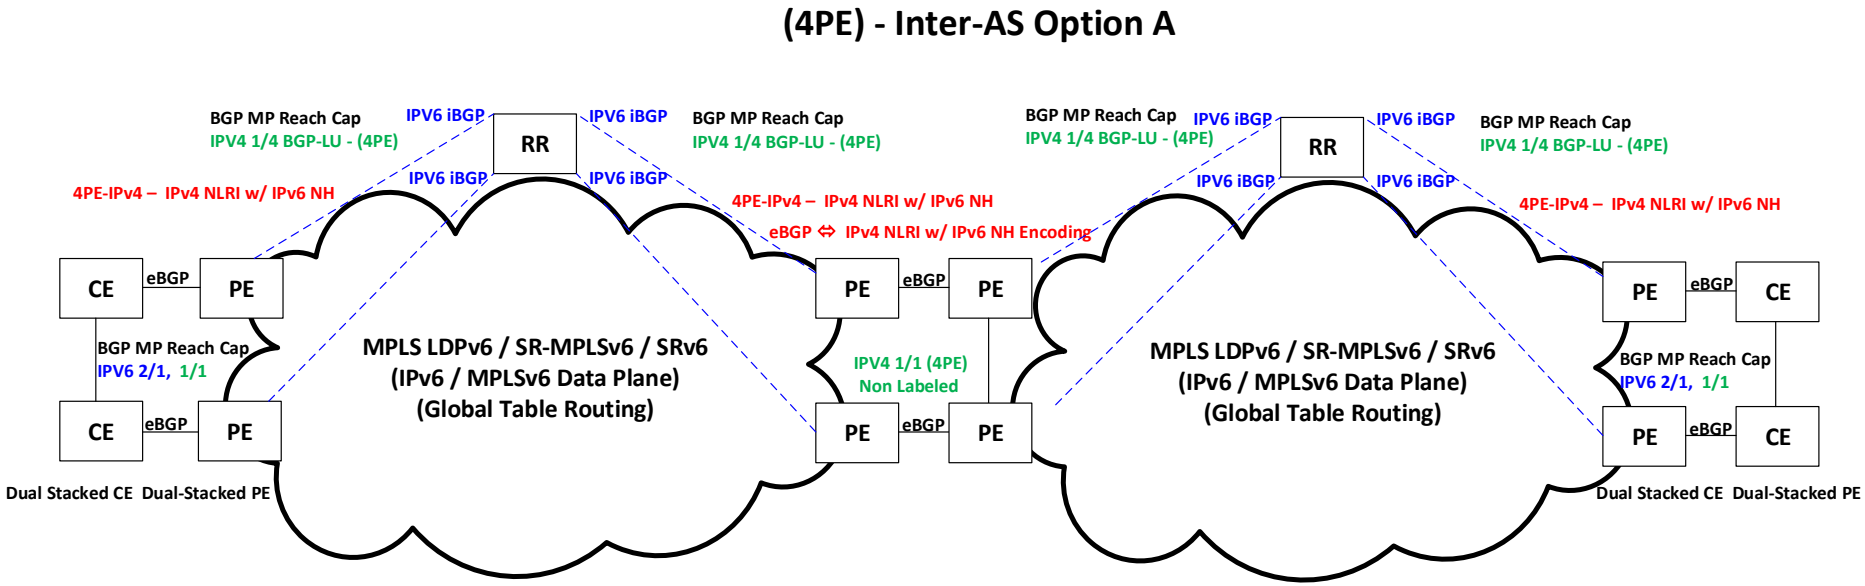

(4PE) –Control Plane & Data Plane Inter-AS Option B Procedure

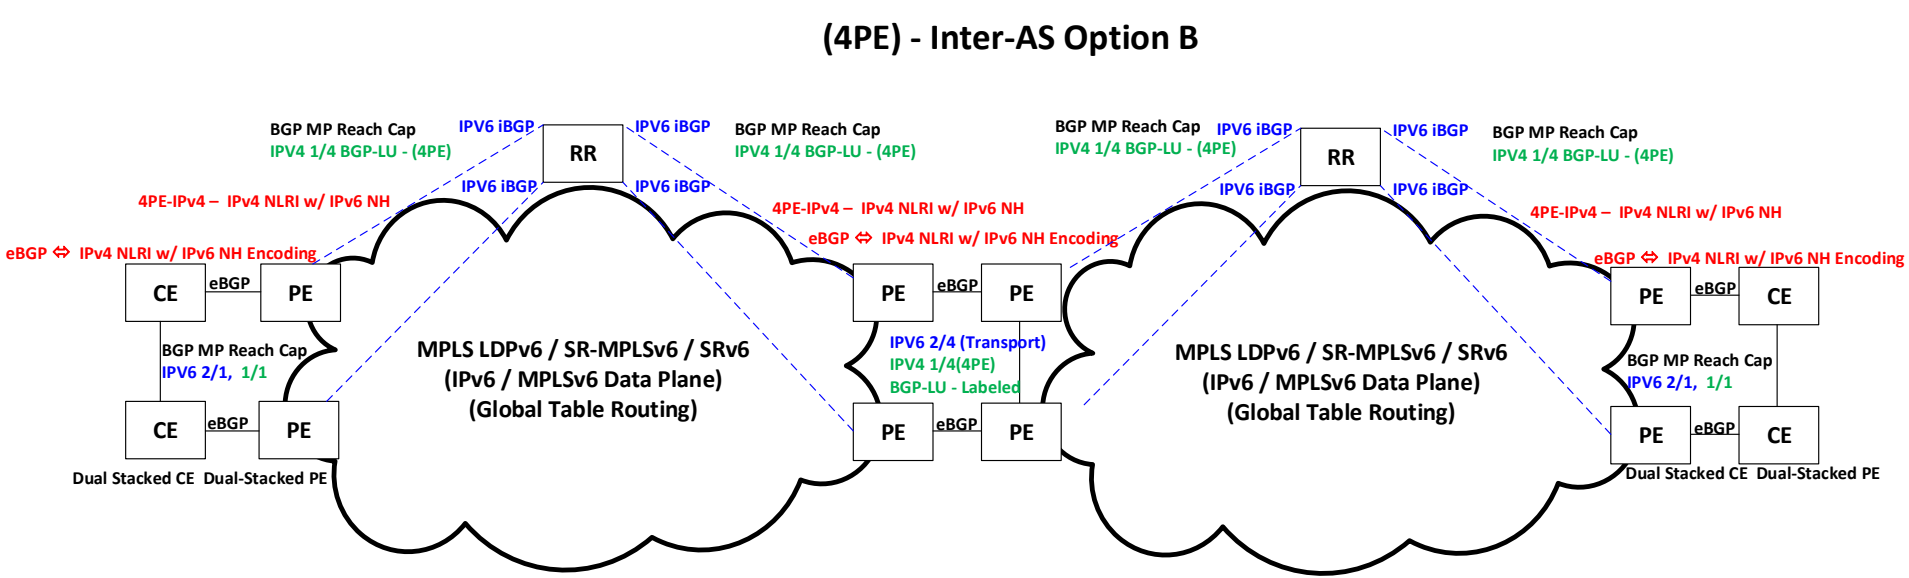

(4PE) –Control Plane & Data Plane Inter-AS Option C Procedure

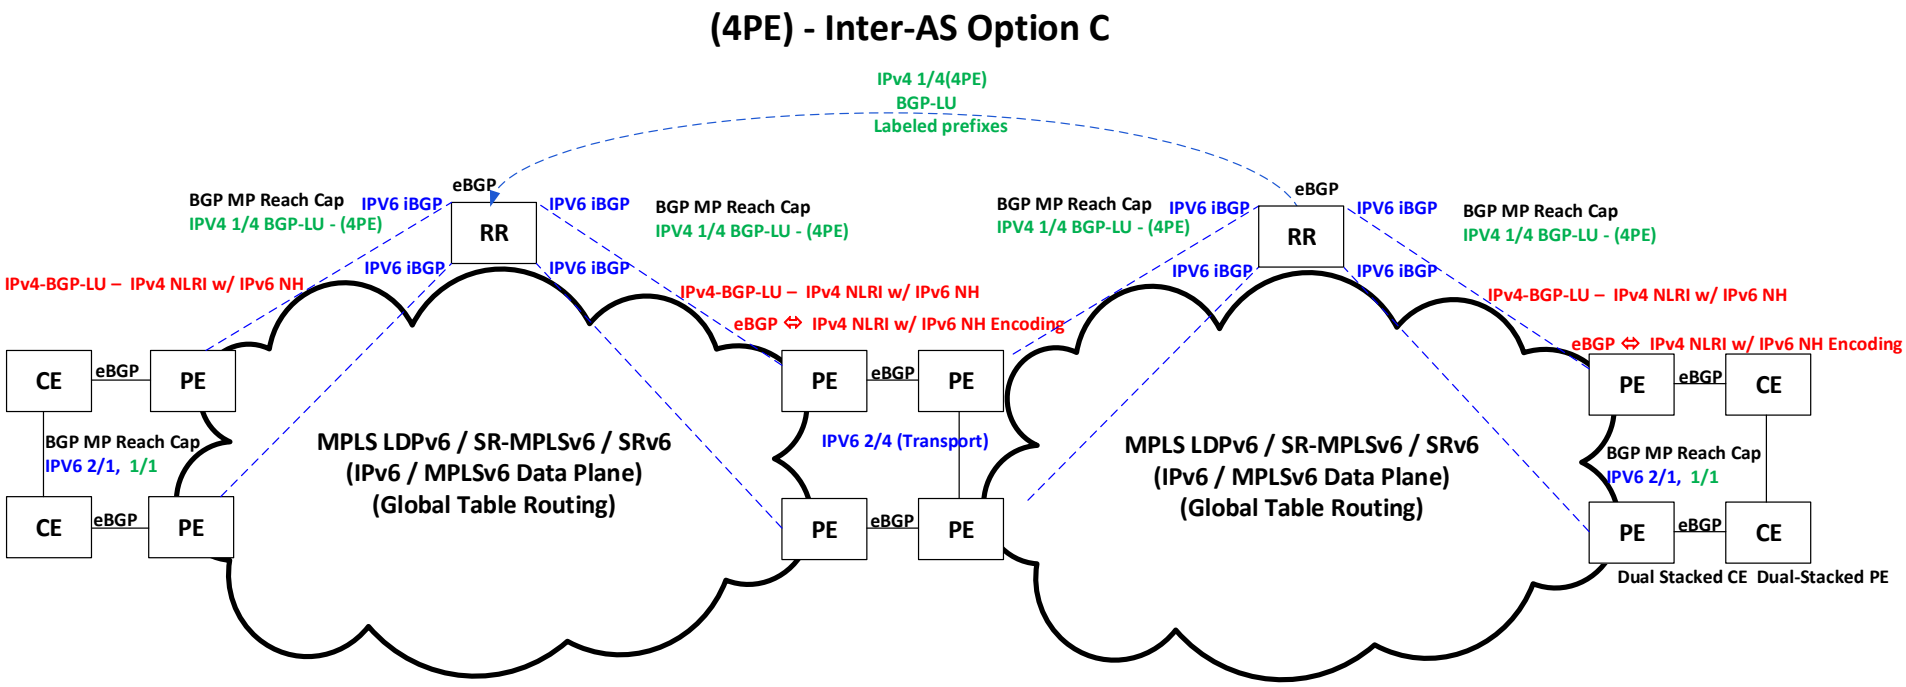

(4PE) –Control Plane & Data Plane Inter-AS Option AB Procedure