

IPv6-Only PE Design & IPv4-Only PE design & ALL SAFI Supported

IPv4-Only & IPv6-Only PE Design ALL SAFI:
`draft-ietf-bess-v4-v6-pe-all-safi-00`

Gyan Mishra (Verizon)
Jeff Tantsura (Microsoft)
Mankamana Mishra (Cisco)
Sudha Madhavi (Juniper)
Mohana Sundari (Juniper)
Adam Simpson (Nokia)
Qing Yang (Arista)
Shaunglong Chen (Huawei)

IPv6-Only PE Design - ALL SAFI

(RECAP)

IPv6-Only PE Design (BCP): Adopted April 28th 2021 (Original Draft)

<https://datatracker.ietf.org/doc/draft-ietf-bess-ipv6-only-pe-design/>

Focus of this draft is on the Proof of Concept testing BCP for vendor implementation & operator deployment

IPv6-Only PE Design PE-CE Edge Peering – Supports IPv4-Unicast 1/1, IPv6-Unicast 2/1 over IPv6 Next Hop

This draft had VPN & MVPN in the original versions so we are just adding back in – so net-net no change (Vendor Testing Draft)

IPv4-Unicast 1/1, IPv6-Unicast 2/1, IPv4-VPN 1/128 IPv4-MVPN 1/129 IPv6-VPN 2/128 IPv6-MVPN 2/129

IPv6-Only PE Design All SAFI (Standards Track): (Presented IETF 114) -Feedback on combining the Drafts @ IETF 114

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv6-only-pe-design-all-safi/>

This draft is a super set of the original draft above.

Focus of this draft is on the IPv6-Only PE Design from a design change perspective and normative language pertaining to the new design procedural, technological as well as paradigm change from traditional “Dual Stacking” to new IPv6-Only PE design for “ALL” SAFI

This draft defines the new procedure for Edge & Inter-AS peering control plane/data plane (CP-DP) & control plane (CP) controller peering optimization to traditional Dual Stacking providing the identical functionality with the IPV4 addressing & IPV4 BGP peering savings, OPEX Saving and its design applicability to all AFI & SAFI and all eBGP peering use cases.

IPv6-Only PE Design ALL SAFI – Supports All AFI/SAFI over **IPv6 Next Hop**.

IPv4-Only PE Design - ALL SAFI

(RECAP)

IPv4-Only PE Design (Standards Track): (New Draft) (This draft has been combined into ALL SAFI Draft) – Feedback from IETF 114

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv4-only-pe-design/>

Focus of this draft is on the Proof of Concept testing BCP for vendor implementation & operator deployment and BGP capability codepoint for next hop encoding similar to RFC 8950.

IPv4-Only PE Design PE-CE Edge Peering, PE-PE inter-as options peering

IPv4-Unicast 1/1, IPv6-Unicast 2/1, IPv4-VPN 1/128 IPv4-MVPN 1/129 IPv6-VPN 2/128 IPv6-MVPN 2/129

IPv4-Only PE Design All SAFI (Standards Track): (New draft) □ Above draft was combined into this draft below after IETF 114

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv4-only-pe-design-all-safi/>

This draft is a super set of the original draft above.

Focus of this draft is on the IPv4-Only PE Design from a design change perspective and normative language pertaining to the new design procedural, technological as well as paradigm change from traditional “Dual Stacking” to new IPv4-Only PE design for “ALL” SAFI

This draft defines the new procedure for Edge & Inter-AS peering control plane/data plane (CP-DP) & control plane (CP) controller peering optimization to traditional Dual Stacking providing the identical functionality with the IPV6 addressing & IPv6 BGP peering savings, OPEX Saving and its design applicability to all AFI & SAFI & all eBGP use cases as well as CAPEX Savings to remain on IPv4 indefinitely if desired

IPv4-Only PE Design ALL SAFI – Supports All AFI/SAFI over **IPv4 Next Hop**. New IANA BGP capability codepoint for Next Hop encoding

Combining ALL SAFI drafts & WG adopted draft into ~~new WG Draft~~

IPv4-Only PE Design All SAFI (Standards Track): (New Draft) (#1)

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv4-only-pe-design-all-safi/>

IPv6-Only PE Design All SAFI (Standards Track): (New Draft) (#2)

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv6-only-pe-design-all-safi/>



Combine the above two drafts into the WG document below (BCP) (#3)

<https://datatracker.ietf.org/doc/draft-ietf-bess-ipv6-only-pe-design/>

ALL 3 DRAFTS ABOVE ARE NOW COMBINED INTO DRAFT BELOW



I polled the WG on July-27th & received consensus from WG to combine on August 7th & chairs approval on August 17th to move forward replace 3 with 1 new draft

New name for the combined draft: (Standards Track)

<https://datatracker.ietf.org/doc/draft-ietf-bess-v4-v6-pe-all-safi/>

Why combine the drafts?

- Both IPv4-Only PE design & IPv6-Only PE design have the BCP testing component to POC test the SAFI, 1,128,129 which can now be done in parallel
- Both drafts have the same design semantics & concepts for single peering design eliminating the other protocol peer alternative dual stack approach
- Both drafts are as well extensible to support ALL SAFI
- Saves Authors & WG effort to progress 3 drafts instead of one draft
- With combined draft a single draft has the entire combined v4 & v6 PE Design architecture one specification for the entire design.

IPv4-Only PE Design & IPv6-Only PE Design Timelines & Caveats

IPv4-Only PE Design & IPv6-Only PE Design POC timeline & Caveats:

With the combined Draft for IPv4 & IPv6 PE Design we will test both in parallel!! (Saves time)

We will Only test SAF 1, 128, 129 as outlined in both the IPV4-Only PE Design & IPv6-Only PE Design draft

The IPv6-Only PE Design & IPv6-Only PE Design are extensibility to support ALL SAFI with the drafts as they provide a design paradigm shift and procedures and process standardization with the Standards Track draft to support “All SAFI” that use IPv4 or IPv6 AFI that can be carried on a single IPv4 peer or IPv6 peer. Further testing related to the ALL SAFI will be deferred to the development teams of each vendor to supported based on roadmaps & timelines in the future after the draft is published for each vendor and can be tested by operators as they plan to start using deployment using the single peer concept for all other SAFI.

Applies to IPv4-Only PE design:

- The new IPv6 NLRI over IPv4 next hop, next hop encoding we will complete all the testing using the existing next hop encoding whatever that may be if its IPv4 mapped IPv6 address or RFC 5549 / RFC 8950 style next hop encoding.
- We will work with each vendors development team to get their support and agreement on the change to the RFC 5549 / RFC 8950 next hop style encoding using the new standard 4 Byte IPv4 address next hop encoding for SAF 1, 2, 4 and SAF 128 & 129 use RD filed added 8 bytes RD 0 12 byte next hop encoding standard.
- Once each vendor has developed the new next hop encoding we will test as well on the new next hop encoding method.
- After all testing is completed with the support of all the vendors we will progress the draft to RFC.

IPv6-Only PE Design Testing update / IPv4-Only PE Design Testing Update:

Vendors: Cisco, Juniper, Nokia, Arista, Huawei

Hardware Platforms, Router Code Revision & Testing Updates:

Vendors: Cisco, Juniper, Nokia, Arista, Huawei

Code & Platform chosen for testing platform by Vender

Cisco: Edge Router- XR ASR 9910 IOS XR 7.4.1, Core Router- NCS 6000

7.2.2, CRS-X 6.7.4

Update: Lab hardware requested in queue for testing and plan to **complete by EOY 2023**

Vendor specific knob for forwarding IPv4 packets without IPv4 address configured on interface which is required for IPv6-Only PE design.

Juniper: Edge Router- MX platform MX480, MX960, Core Router- PTX

Platform PTX5000, PTC10K8 (JUNOS and EVO) Release 20.4R2

Update: All unicast testing is completed & just multicast testing remaining plan to **complete by EOY 2023**

Vendor specific knob exists for forwarding IPv4 packets without IPv4 address configured on interface required for IPv4-Only PE design & IPv6-Only PE design.

Nokia: Edge and Core-7750 Service Router, Release R21

Update: Lab hardware requested in queue for testing and plan to **complete by EOY 2023**

Vendor specific knob exists for forwarding IPv4 packets without IPv4 address configured on interface which is required for IPv6-Only PE design.

Huawei: Edge and Core-VRPv8, Release VRP-V800R020C10

Update: Huawei supports RFC 5549 & RFC 8950 and all 12 test cases confirmed with R&D and supports some but not all test cases and plans to **complete by 2023**.

Vendor specific knob exists for forwarding IPv4 packets without IPv4 address configured on interface which is required for IPv6-Only PE design.

Arista: No ETA

Q&A

THANK YOU

IPv4-Only / IPv6-Only PE ALL SAFI Design □ Benefits to all Operators Worldwide

Scope

- IPv4-Only PE design / IPv6-Only PE design ALL SAFI
- Advertise both IPv4 and IPv6 routes (NLRI) ANY SAFI over a single IPv6 BGP peer or single IPv4 BGP peer
- Supports PE-CE Edge & Inter-AS peering & controller peering

Benefits

- OpEx savings through elimination of IPv4 BGP peering and IPv4 address usage for IPv6-Only PE design
- OpEX savings through elimination of IPv6 BGP peering & IPv6 address for IPv4-Only PE design
- Simplified configuration
- Reduce network resource consumption
- Will function in all BGP AFI / SAFI scenarios **listed**
- Ubiquitous use cases
- Allows operators “core” underlay network to remain IPv4-Only indefinitely if desired
- **CAPEX & OPEX Savings \$\$ for Operators!!!**

IPv4-Only PE Design & IPv6-Only

PE Design

IPv4-Only PE Design where Only an IPv4 address is configured on the PE-CE interface:

Vendor specific knob for IPv6 processing and how it works w/o IPv6 address configured on the interface:

- The interface acts as a L3 interface as normal layer 3 hop w/o an IPv6 address configured on the interface and so the TTL is decremented in the packet and the L2 header is stripped and new L2 header is added when forwarded? Yes
- So we have 100% L3 functionality w/o having a L3 IPv6 address enabled on the interface? Yes
- Ping & trace is as if it's a L3 hop and so we just need IPv6 address on the loopback interface? Yes
- IPFIX, QOS and ACL processing works as it normally would work with Dual Stack interface? Yes

IPv6-Only PE Design where Only an IPv6 address is configured on the PE-CE interface:

Vendor specific knob for IPv4 processing and how it works w/o IPv4 address configured on the interface:

- The interface acts as a L3 interface as normal layer 3 hop w/o an IPv4 address configured on the interface and so the TTL is decremented in the packet and the L2 header is stripped and new L2 header is added when forwarded? Yes
- So we have 100% L3 functionality w/o having a L3 IPv4 address enabled on the interface? Yes
- Ping & trace is as if it's a L3 hop and so we just need IPv4 address on the loopback interface? Yes
- IPFIX, QOS and ACL processing works as it normally would work with Dual Stack interface? Yes

IPv4-Only PE Design Next Hop Encoding Standardization

- RFC 4798 (6PE) section 2 defines how the next hop should be encoded for IPv6 NLRI over an IPv4 next hop using IPv4 mapped IPv6 address ::FFFF:192.168.1.1. RFC 4659 BGP MPLS VPNs section 3.2.1.2 defines VPN SAFI next hop encoding of IPv4 mapped IPv6 address ::FFFF:192.168.1.1.
- RFC 5549 and now updated by RFC 8950 defines the IPv6 next hop encoding to carry IPv4 NLRI over an IPv6 next hop. The IPv6 next hop encoding defined is not an IPv6 mapped IPv4 address. The IPv6 next hop encoding is 16/32 byte (RFC 2545 – NH address = IPv6 address + Link Local address) for Unicast SAFI 1, Multicast SAFI 2 and BGP-LU SAFI 4, and 24/48 byte (RFC 2545 – NH address / IPv6 address + link local address) for VPN SAFI 128, MVPN SAFI 129. The IANA BGP Capability codepoint defined with RFC 5549 is value 5 for Extended Next hop encoding.
- The industry implementation uses a mix of IPv4 mapped IPv6 address for IPv6 NLRI carried over an IPv4 address next hop and uses 4 byte field for IPv4 next hop address for Unicast SAFI 1, Multicast SAFI2 and BGP-LU SAFI 4, and 12 byte next hop field, 4 byte IPv4 address plus 8 byte RD (Route Distinguisher) set to 0 for VPN SAFI 128, MVPN SAFI 129.
- This draft standardizes the encoding to use an IPv4 address next hop and uses 4 byte field for IPv4 next hop address for Unicast SAFI 1, Multicast SAFI2 and BGP-LU SAFI 4, and 12 byte next hop field, 4 byte IPv4 address plus 8 byte RD (Route Distinguisher) set to 0 for VPN SAFI 128, MVPN SAFI 129.
- This draft standardizes that encoding to ensure interoperability with IANA BGP Capability codepoint allocation thus providing parity between the RFC 5549/RFC 8950 IPv6 next hop encoding where the next hop address follows the underlay core protocol which is an IPv6 core and how the next hop here being an IPv6 address and not following the NLRI protocol with IPv6 mapped IPv4 address. Now with this draft the next hop encoding follows the underlay core which is an IPv4 core and so now the next hop being an IPv4 address and not following the NLRI with an IPv4 mapped IPv6 address. So this parity between IPv4 next hop encoding and IPv6 next hop encoding savings in OPEX and operations troubleshooting as well as interoperability that all vendor implementations now use the same IPv4 next hop encoding is the reason the encoding must be standardized.
- This IPv4 next hop encoding is applicable for IPv6 NLRI for both iBGP control plane (CP) peering as well as eBGP PE-CE, PE-PE in-line control / data plane (CP-DP) peering which is used for IPv4-Only PE design as well as any IPv4 peering. The IPv4 Next hop encoding updates both RFC 4271 next hop path attribute and RFC MP-BGP RFC 4760 NLRI path attribute.
- I would like to add information to IDR Wiki on which vendors support the IPv4 mapped IPv6 address and which support the RFC 5549/RFC 8950 style encoding. As well if the vendor does not support the new next hop encoding, it would continue to use the IPv4 mapped IPv6 address format until the P2P peering both neighbors MP-BGP MP_REACH BGP capability exchange is for the new IPv4 Next hop encoding codepoint.

IPv4-Only PE Design ALL SAFI & IPv6-Only PE ALL SAFI Design

<https://datatracker.ietf.org/doc/draft-ietf-bess-ipv6-only-pe-design-all-safi/>

<https://datatracker.ietf.org/doc/draft-mishra-bess-ipv4-only-pe-design-all-safi/>

IPv4-Only PE Design ALL SAFI & IPv6-Only PE Design ALL SAFI

- Supports Advertising **ALL SAFI**
- Advertise both IPv4 and IPv6 routes (NLRI) ALL SAFI over a single IPv6 BGP peer or single IPv4 BGP Peer
- Supports PE-CE Edge, Inter-AS peering & every type of peering relationships
- Supports **Control Plane (CP-Only)** such as to a controller and **Control Plane / Data Plane (CP-DP)** Scenario
- Supports Dual control plane during migration such as from SR-MPLS(IPv4) to SRv6(IPv6)

Functional use cases for ALL AFI/SAFI Design:

- Edge Customer NLRI IPv4 or IPV6 related AFI/SAFI (PE-CE) **(CP-DP)**
 - 1/1 2/1 (Unicast), 1/2 2/2 (Multicast)
- Inter-AS Customer NLRI IPv4 or IPV6 related AFI/SAFI (ASBR-ASBR) **(CP-DP)**
 - 1/1 2/1 (Unicast), 1/2 2/2 (Multicast), 1/128 2/128 (VPN), 1/129 2/129 (MVPN), 1/4 2/4 BGP-LU (6PE/4PE), 1/140 2/140 (BGP VPN Auto Discovery)
- Inter-AS Multicast NLRI IPv4 or IPV6 related AFI/SAFI (ASBR-ASBR) **(CP-DP)**
 - 1/5 2/5 (MCAST-VPN), 1/8 2/8 (MCAST-VPLS), 1/66 2/66 (BGP MDT-SAFI), 1/78 2/78 (MCAST-TREE)
- PE to Controller NLRI IPv4 or IPV6 related AFI/SAFI **(CP-Only)**
 - 1/80 2/80 BGP-LS-SPF
- Inter-AS L1 VPN, L2 VPN NLRI IPv4 or IPV6 related AFI/SAFI (ASBR-ASBR) **(CP-DP)**
 - 1/69 2/69 (L1 VPN) (ASBR-ASBR)
- Inter-AS BGP FlowSpec, Optimizations & SFC NLRI IPv4 or IPV6 related AFI/SAFI (ASBR-ASBR) **(CP-DP)**
 - 1/132 2/132 (RTC), 1/133 2/133 (BGP FlowSpec), 1/134 2/134 (VPN BGP FlowSpec)
- Inter-AS BGP Policy - SR-TE Policy, SD-WAN Policy NLRI IPv4 or IPV6 related AFI/SAFI (ASBR-ASBR) **(CP-DP)**
 - 1/73 2/73 (SR-TE), 1/74 2/74 (SD-WAN Capabilities)

IPv4-Only PE Design ALL SAFI & IPv6-Only PE ALL SAFI Design

IANA BGP AFI SAFI □ ALL SAFI's listed below support the IPv6-Only PE Design framework

<https://www.iana.org/assignments/safi-namespace/safi-namespace.xhtml>

Value	Description	Reference
0	Reserved	[RFC4760]
1	Network Layer Reachability Information used for unicast forwarding	[RFC4760]
2	Network Layer Reachability Information used for multicast forwarding	[RFC4760]
3	Reserved	[RFC4760]
4	Network Layer Reachability Information (NLRI) with MPLS Labels	[RFC8277]
5	MCAST-VPN	[RFC6514]
6	Network Layer Reachability Information used for Dynamic Placement of Multi-Segment Pseudowires	[RFC7267]
7	Encapsulation SAFI (OBSOLETE)	[RFC9012]
8	MCAST-VPLS	[RFC7117]
9	BGP-SFC	[RFC9015]
10-63	Unassigned	
64	Tunnel SAFI	[Gargi_Nalawade][draft-nalawade-kapoor-tunnel-safi-01]
65	Virtual Private LAN Service (VPLS)	[RFC4761][RFC6074]
66	BGP MDT SAFI	[RFC6037]
67	BGP 4over6 SAFI	[RFC5747]
68	BGP 6over4 SAFI	[Yong_Cui]
69	Layer-1 VPN auto-discovery information	[RFC5195]
70	BGP-EVPNs	[RFC7432]
71	BGP-LS	[RFC7752]
72	BGP-LS-VPN	[RFC7752]

IPv6-Only PE ALL SAFI Design

IANA BGP AFI SAFI (Continued)

<https://www.iana.org/assignments/safi-namespace/safi-namespace.xhtml>

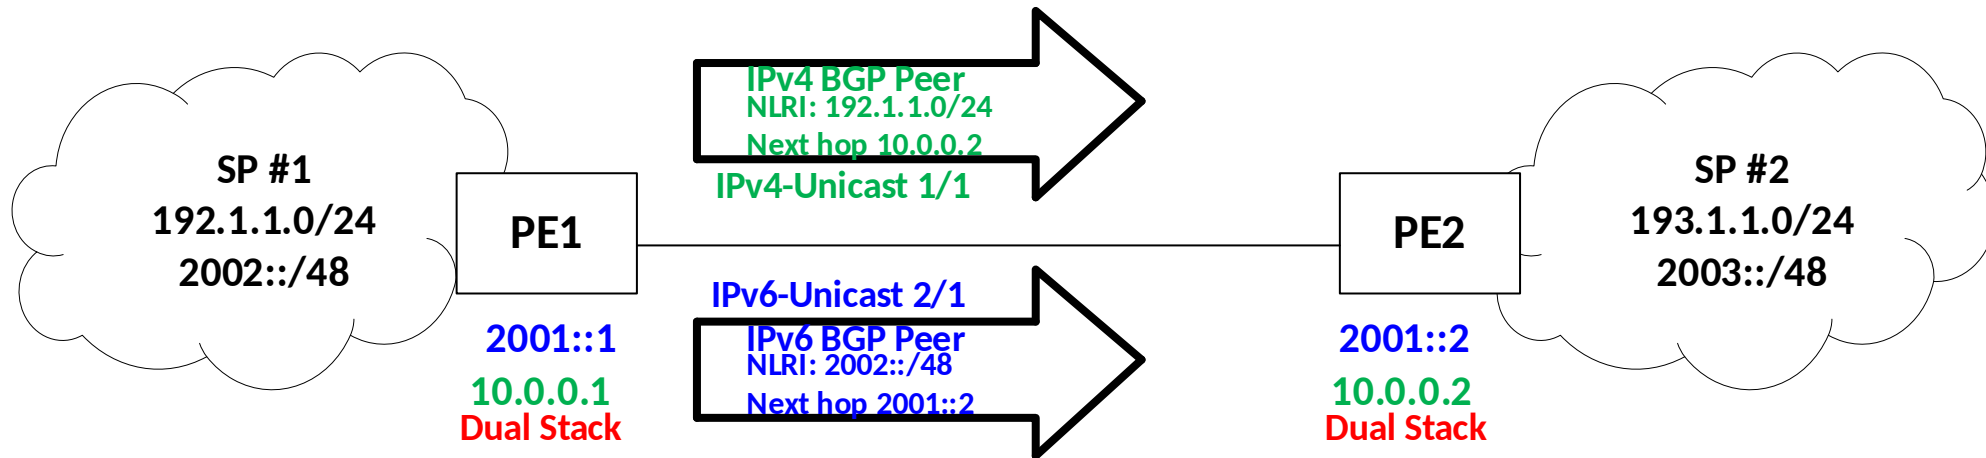
Value	Description	Reference
73	SR TE Policy SAFI	[draft-previdi-idr-segment-routing-te-policy]
74	SD-WAN Capabilities	[draft-dunbar-idr-sdwan-port-safi]
75	Routing Policy SAFI	[draft-ietf-idr-rpd-02]
76	Classful-Transport SAFI	[draft-kaliraj-idr-bgp-classful-transport-planes-00]-- Transport Only
77	Tunneled Traffic Flowspec	[draft-ietf-idr-flowspec-nvo3-10]
78	MCAST-TREE	[draft-ietf-bess-bgp-multicast-03]
79	BGP-DPS (Dynamic Path Selection)	[https://eos.arista.com/eos-4-26-2f/dps-vpn-scaling-using-bgp] [Venkit_Kasiviswanathan]
80	BGP-LS-SPF	[draft-ietf-lsvr-bgp-spf-15][Victor_Kuarsingh]
81-127	Unassigned	
128	MPLS-labeled VPN address	[RFC4364][RFC8277]
129	Multicast for BGP/MPLS IP Virtual Private Networks (VPNs)	[RFC6513][RFC6514]
130-131	Reserved	[RFC4760]
132	Route Target constrains	[RFC4684]
133	Dissemination of Flow Specification rules	[RFC8955]
134	L3VPN Dissemination of Flow Specification rules	[RFC8955]
135-139	Reserved	[RFC4760]
140	VPN auto-discovery	[draft-ietf-l3vpn-bgpvpn-auto]
141-240	Reserved	[RFC4760]
241-254	Reserved for Private Use	[RFC4760]
255	Reserved	[RFC4760]

Traditional PE-CE Dual Stacked Peering with Separate IPv4 & IPv6 Peer □ (IPv6-Only PE Design Slide Set)

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

PE-CE Native IP (GRT) Routing

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering



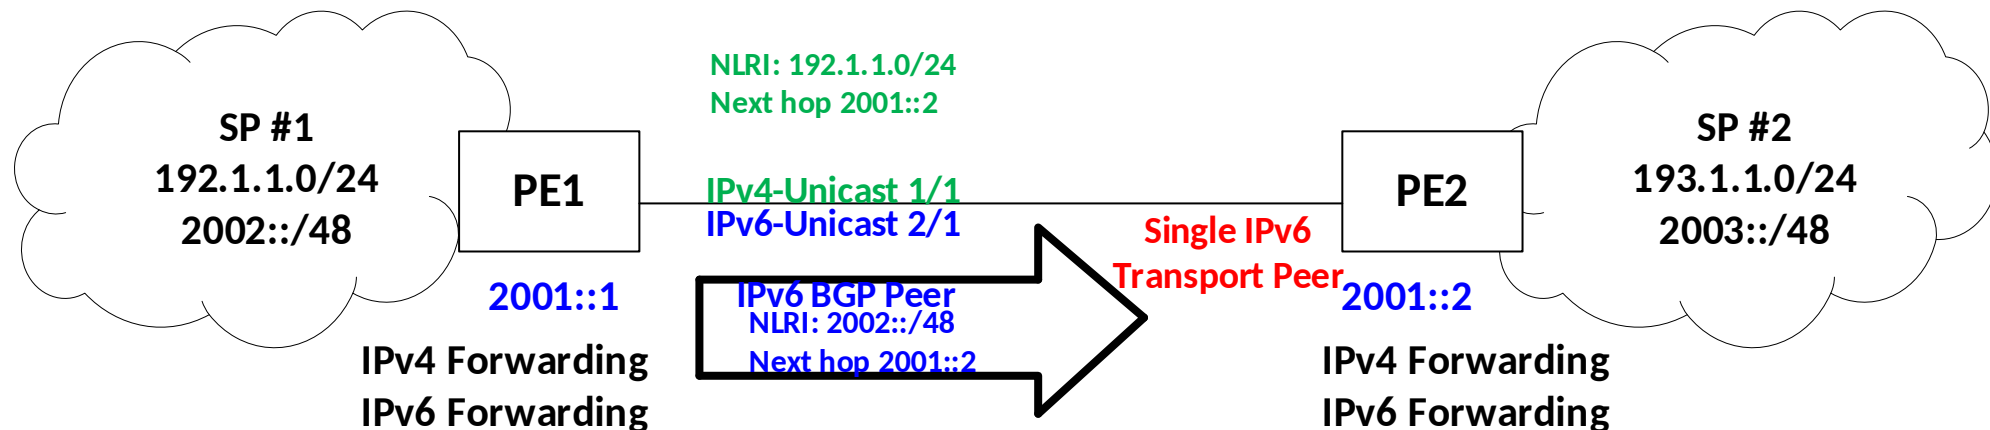
IPv6-Only PE Design ALL SAFI □ PE-CE Edge Single IPv6 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

So now with RFC 8950 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.

PE-CE Native IP (GRT) Routing

IPv6-Only PE Design ⇔ Single IPv6-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI



IXP Peering of all SPs would now be able to use a single IPv6 peer per SP adjacency
Eliminate IPv4 Address depletion issues

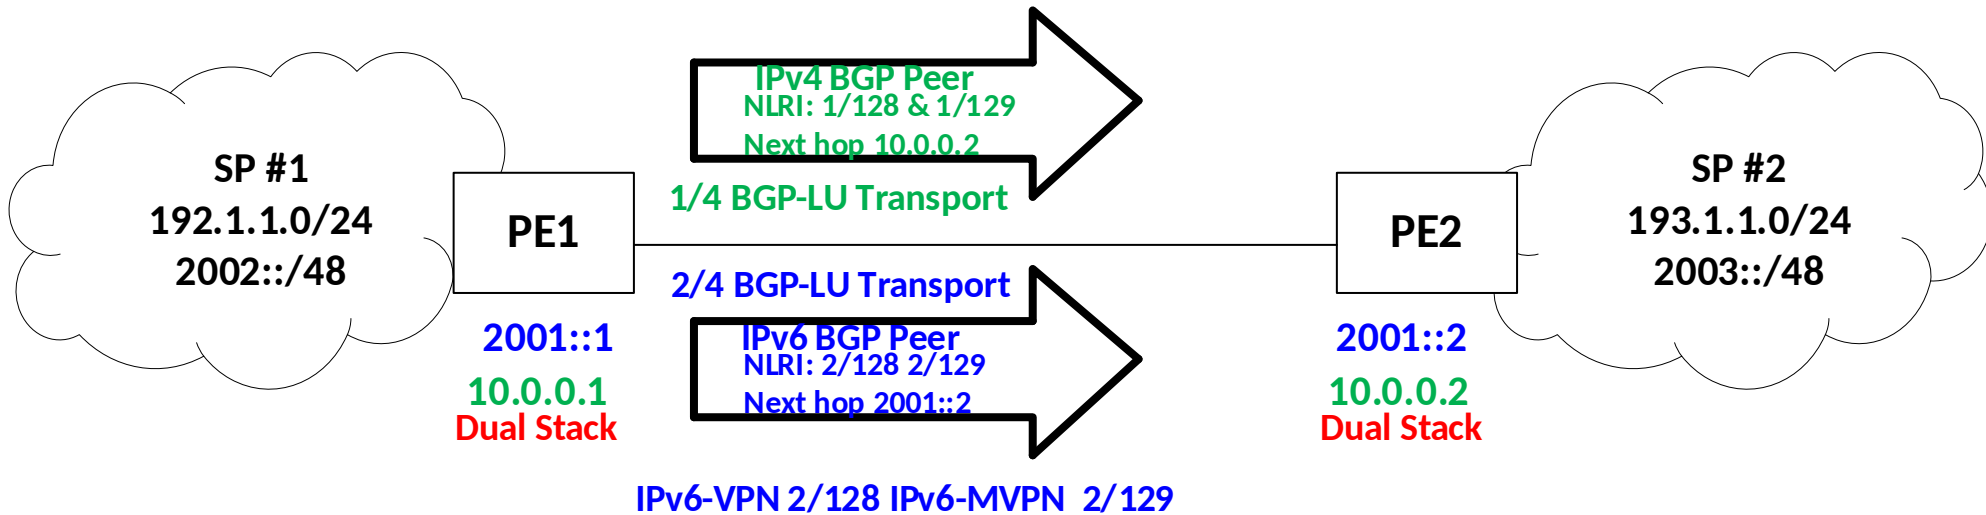
Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv6-Only PE Design Slide Set)

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

Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering

IPv4-VPN 1/128 IPv4-MVPN 1/129

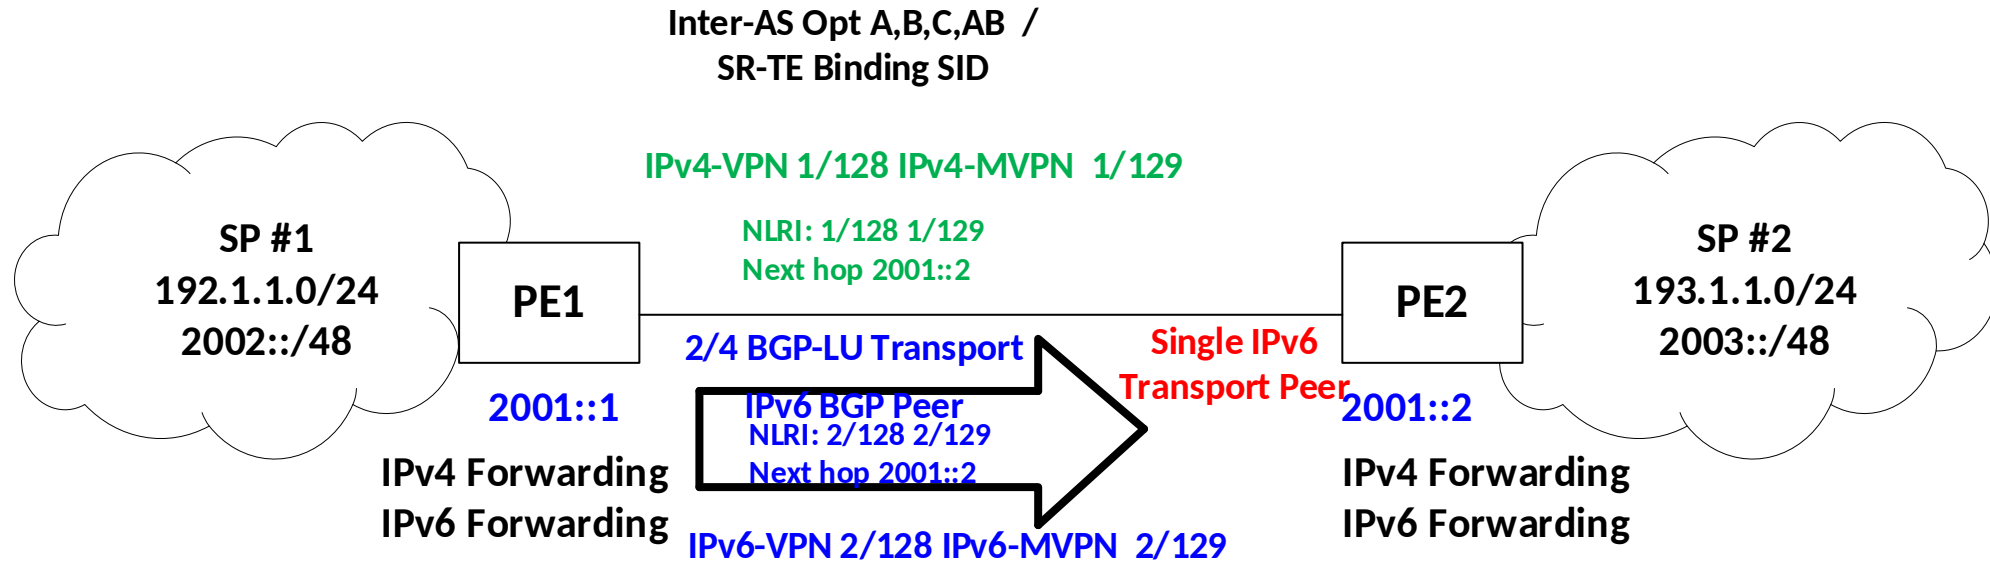


IPv6-Only PE Design ALL SAFI □ Inter-AS - Single IPv6 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

So now with RFC8950- NH encoding schema of 24/48 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.

IPv6-Only PE Design ⇔ Single IPv6-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI



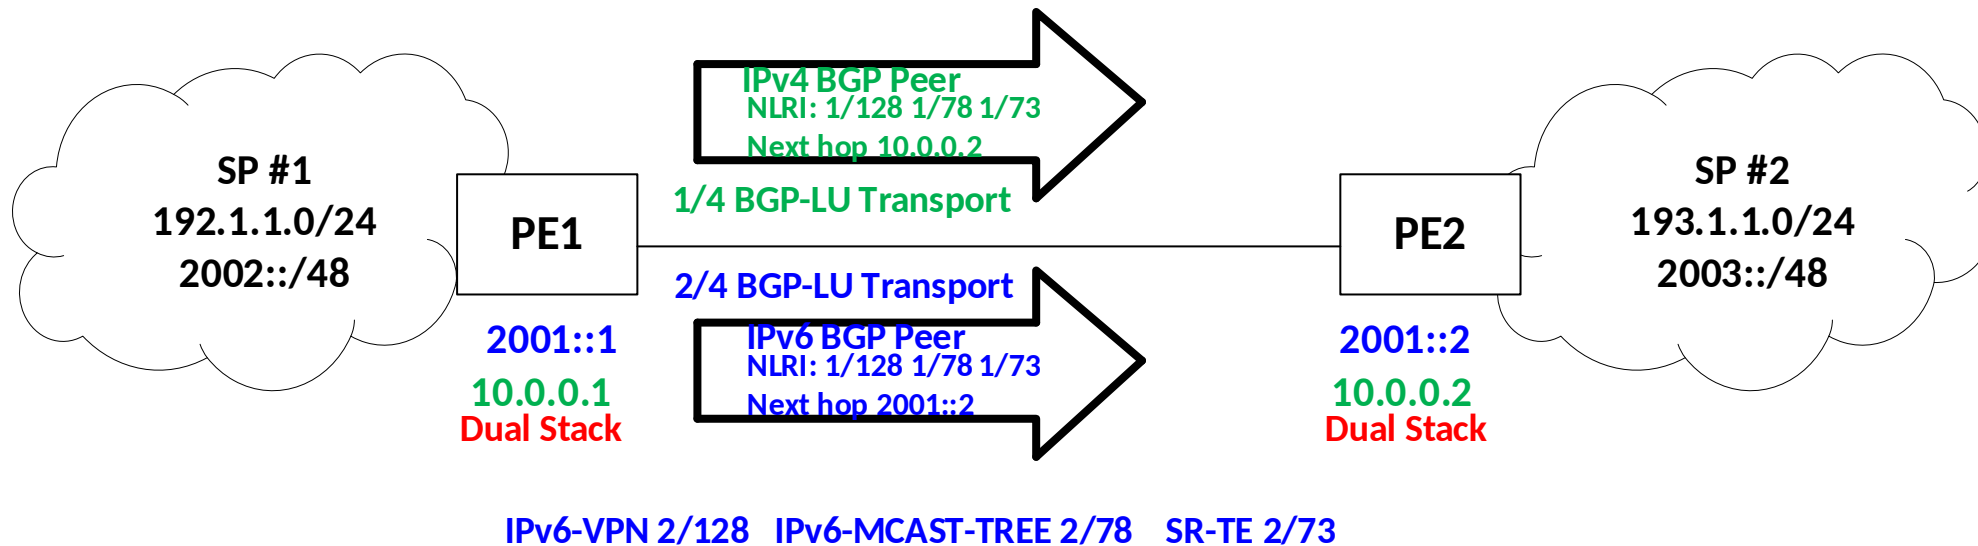
IXP Peering of all SPs would now be able to use
a single IPv6 peer per SP adjacency
Eliminate IPv4 Address Depletion Issues

Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv6-Only PE Design Slide Set)

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

Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering
IPv4-VPN 1/128 IPv4-MCAST-TREE 1/78 SR-TE 1/73



IPv6-Only PE Design ALL SAFI □ Inter-AS - Single IPv6 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

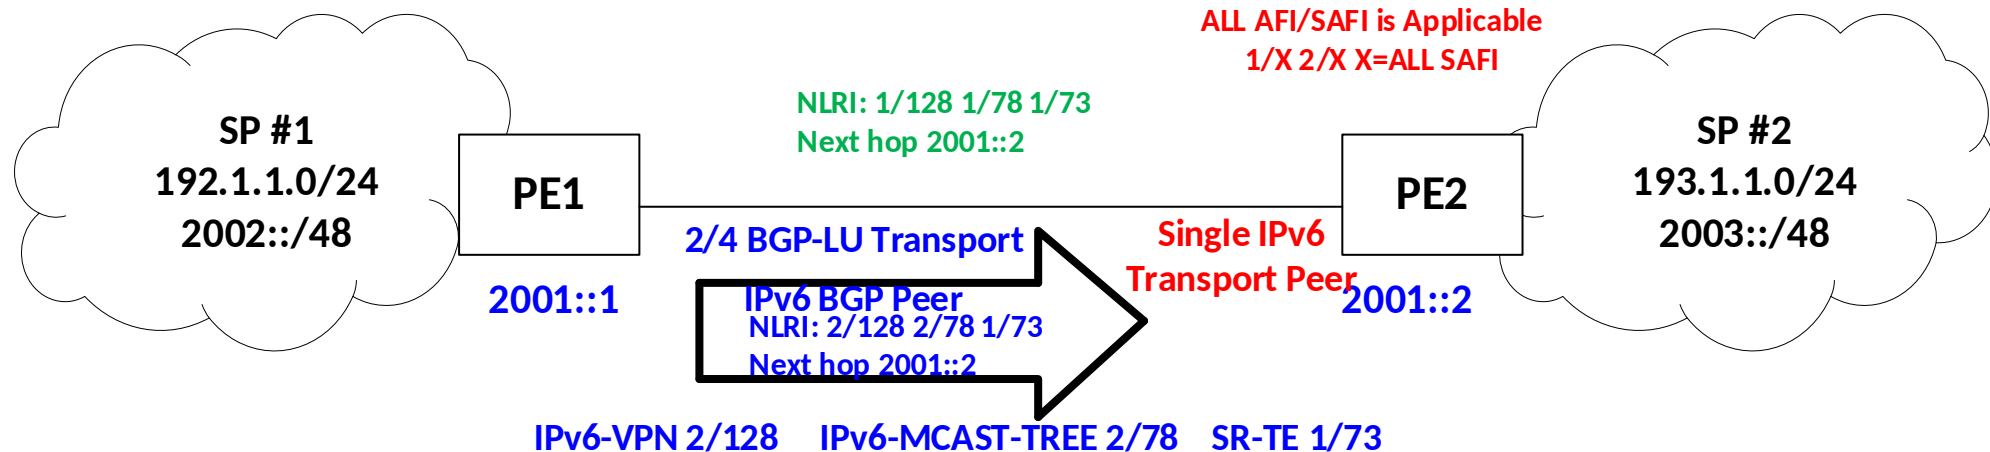
So now with RFC8950- NH encoding schema of 24/48 byte IPv6 next hop both IPv4 & IPv6 NLRI can be advertised using a single IPv6 peer.

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

IPv6-Only PE Design ⇔ Single IPv6-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI

Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

IPv4-VPN 1/128 IPv4-MCAST-TREE 1/78 SR-TE 1/73

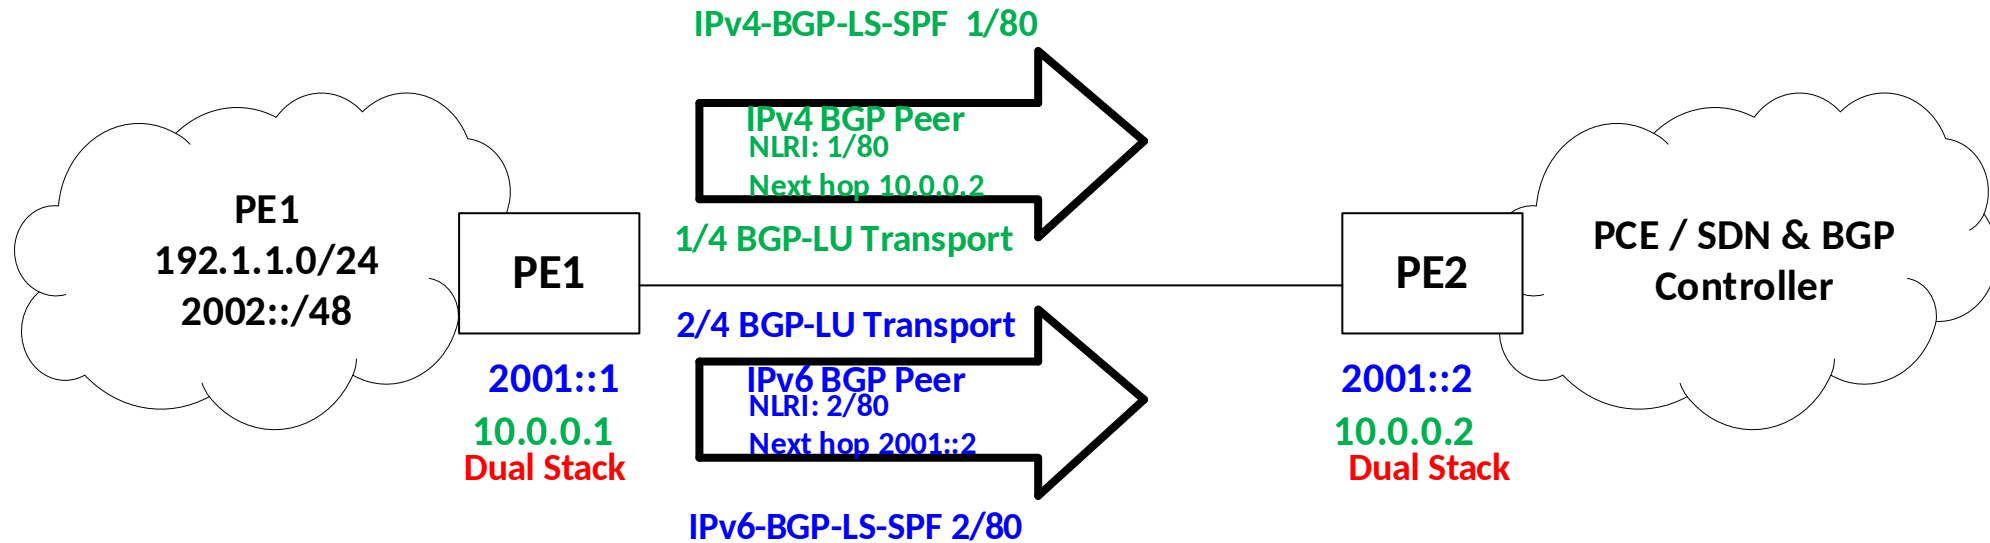


IXP Peering of all SPs would now be able to use
a single IPv6 peer per SP adjacency
Eliminate IPv4 Address Depletion Issues

Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv6-Only PE Design Slide Set)

PE to PCE/SDN or BGP Controller has a separate IPv4 & IPv6 peer with IPv4 NH peer carrying IPv4 NLRI & IPv6 NH peer carrying IPv6 NLRI
Inter-AS Opt A,B,C,AB / SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering



IPv6-Only PE Design ALL SAFI ↔ Inter-AS - Single IPv6 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

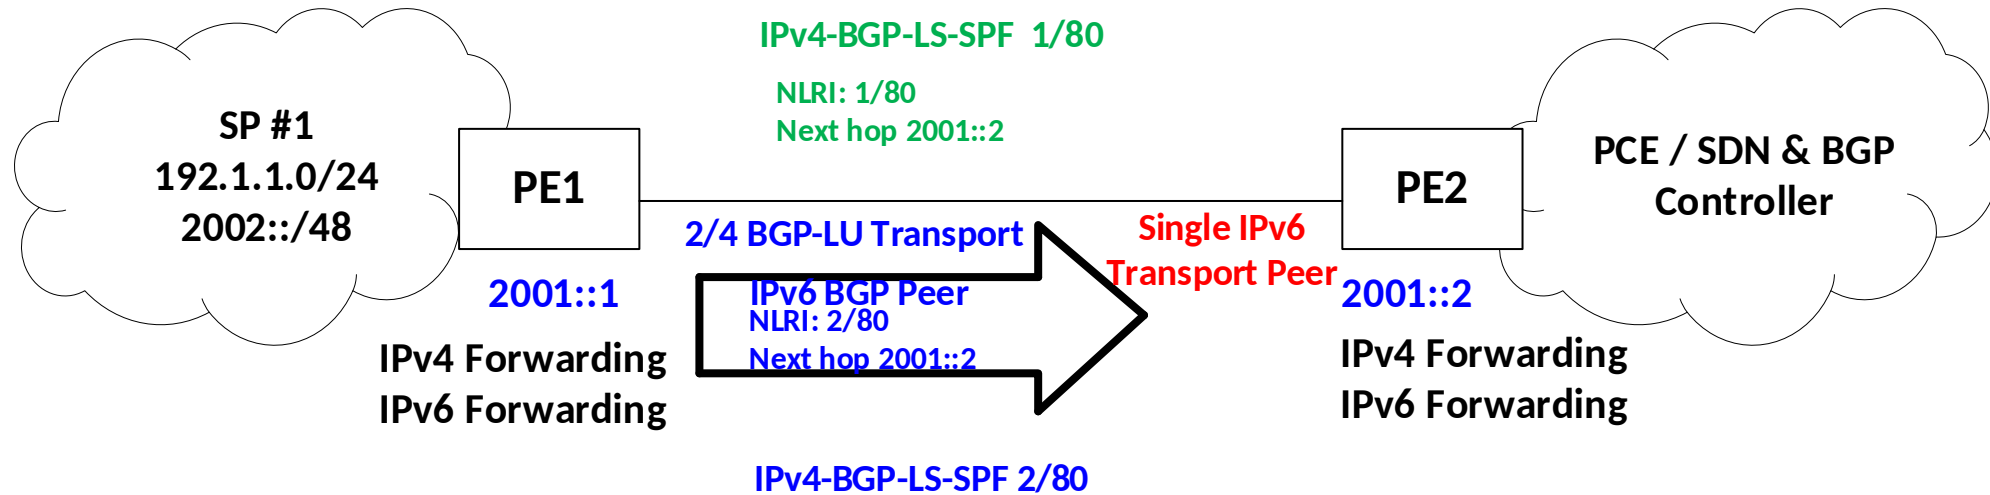
PE to PCE/SDN or BGP Controller single IPv6 Peer

So now with RFC8950- NH encoding schema of 24/48 byte IPv6 next hop both IPv4 & IPv6 NLRI can be advertised using a single IPv6 peer.

This basic concept can eliminate all PE to Controller peering at the Edge and within the Core.

IPv6-Only PE Design ↔ Single IPv6-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI

ALL AFI/SAFI is Applicable
1/X 2/X X=ALL SAFI

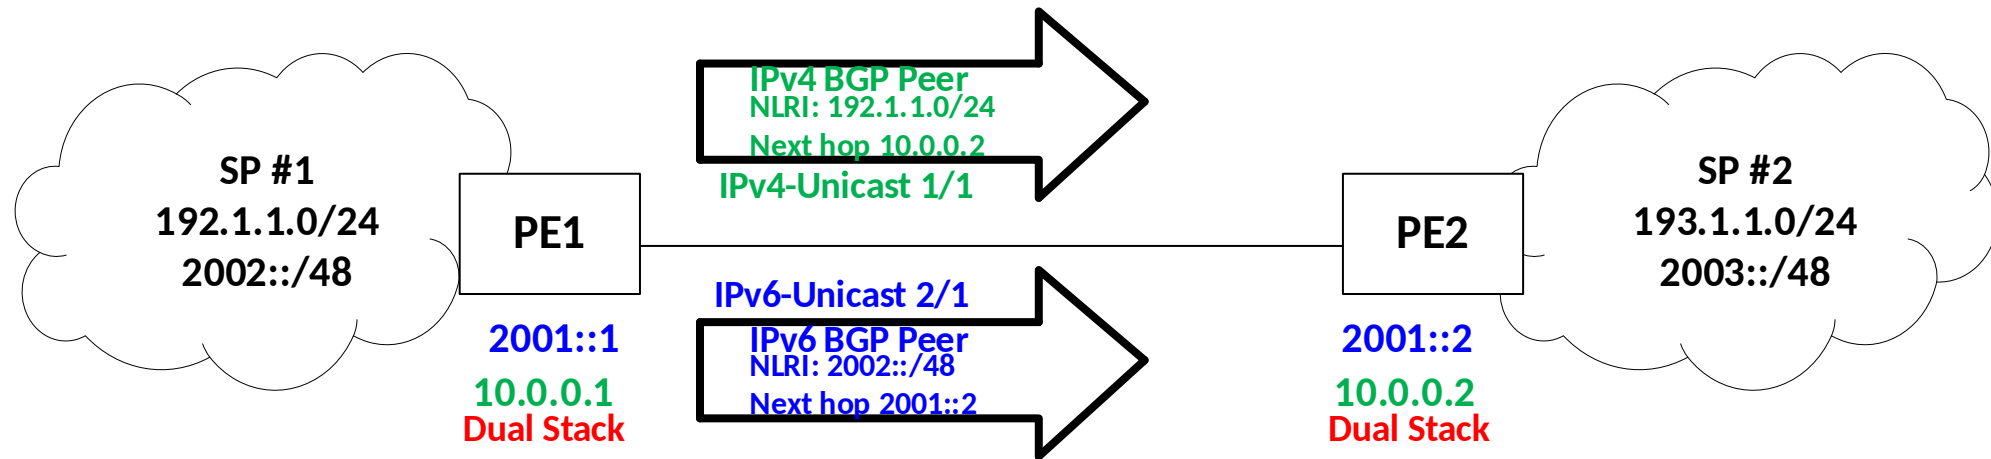


Traditional PE-CE Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv4-Only PE Design Slide Set)

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

PE-CE Native IP (GRT) Routing

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering



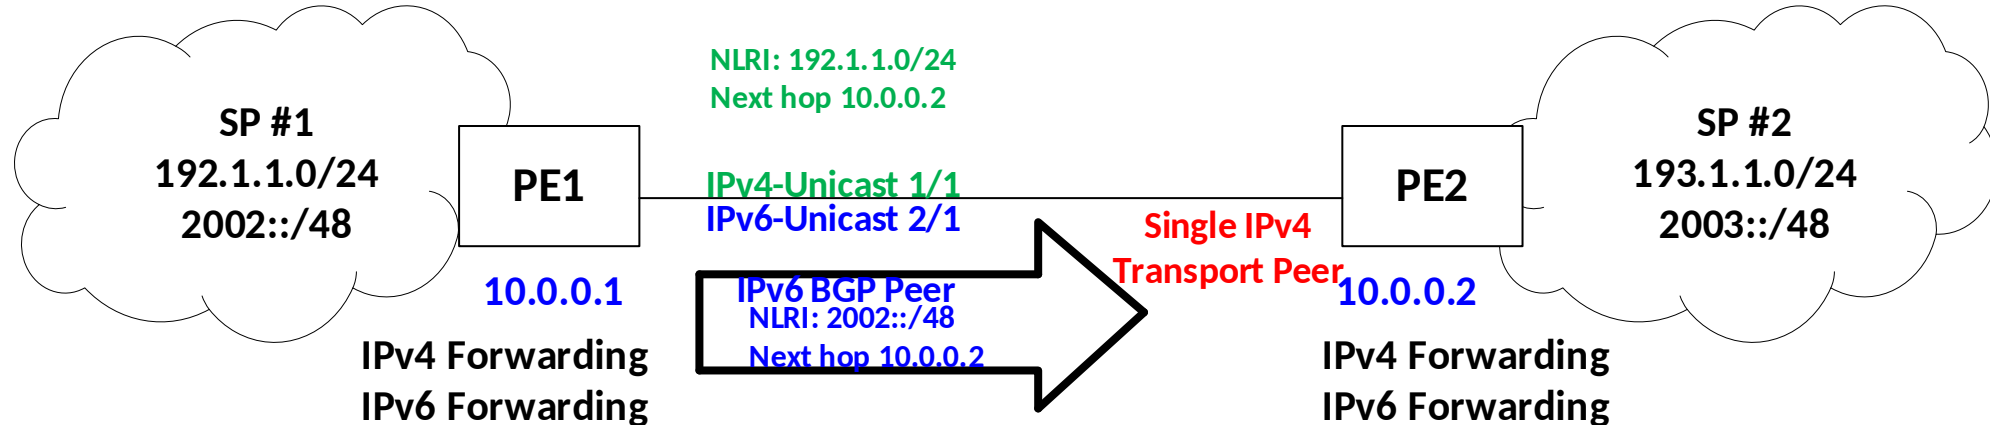
IPv4-Only PE Design ALL SAFI □ PE-CE Edge Single IPv4 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

With this drafts standardized next hop encoding schema of 4 byte IPv4 next hop encoding for IPv6 Unicast SAFI & 12 byte IPv4 next hop encoding for VPN SAFI, both IPv4 & IPv6 NLRI can be advertised using a Single IPv4 peer.

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

PE-CE Native IP (GRT) Routing

IPv4-Only PE Design ⇔ Single IPv4-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI



IXP Peering of all SPs would now be able to use a Single IPv4 peer per SP adjacency
Eliminates resource issues & provides CAPEX & OPEX Savings

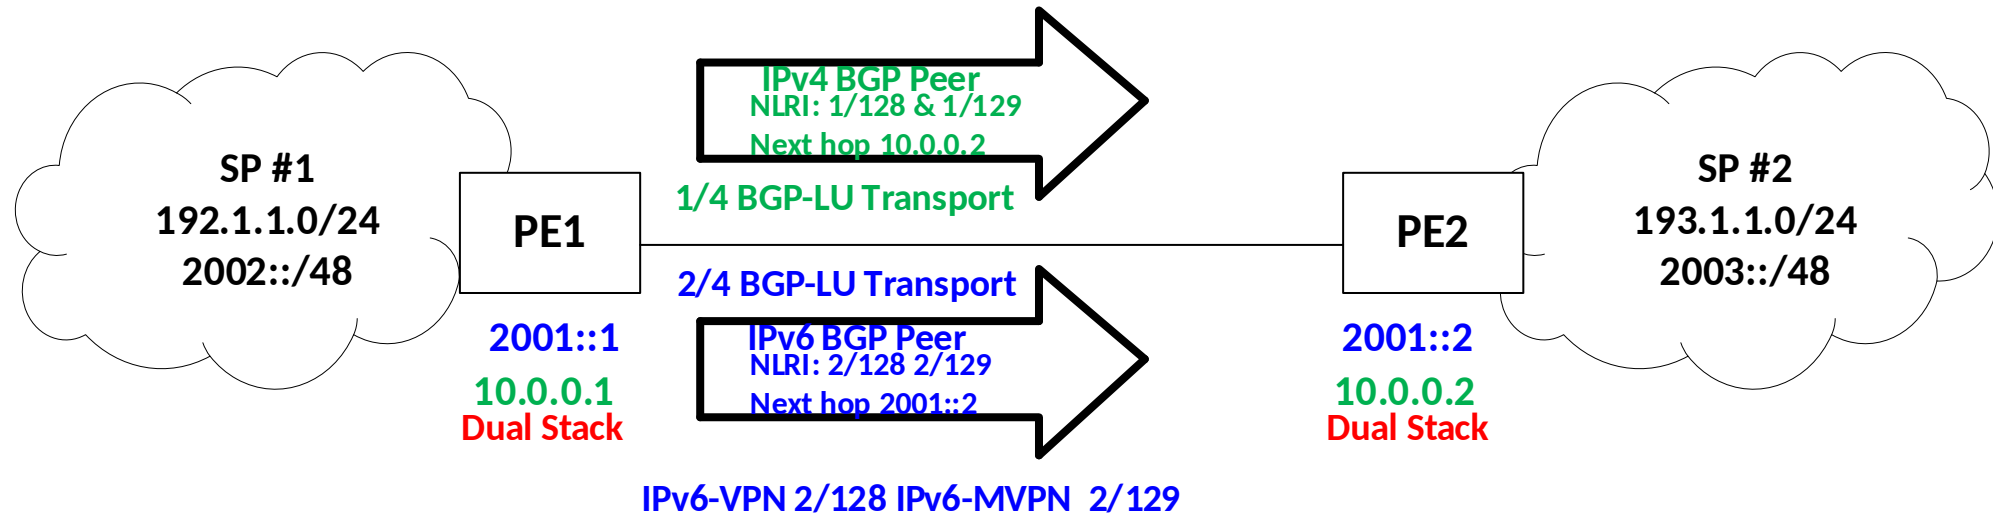
Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv4-Only PE Design Slide Set)

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

Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering

IPv4-VPN 1/128 IPv4-MVPN 1/129

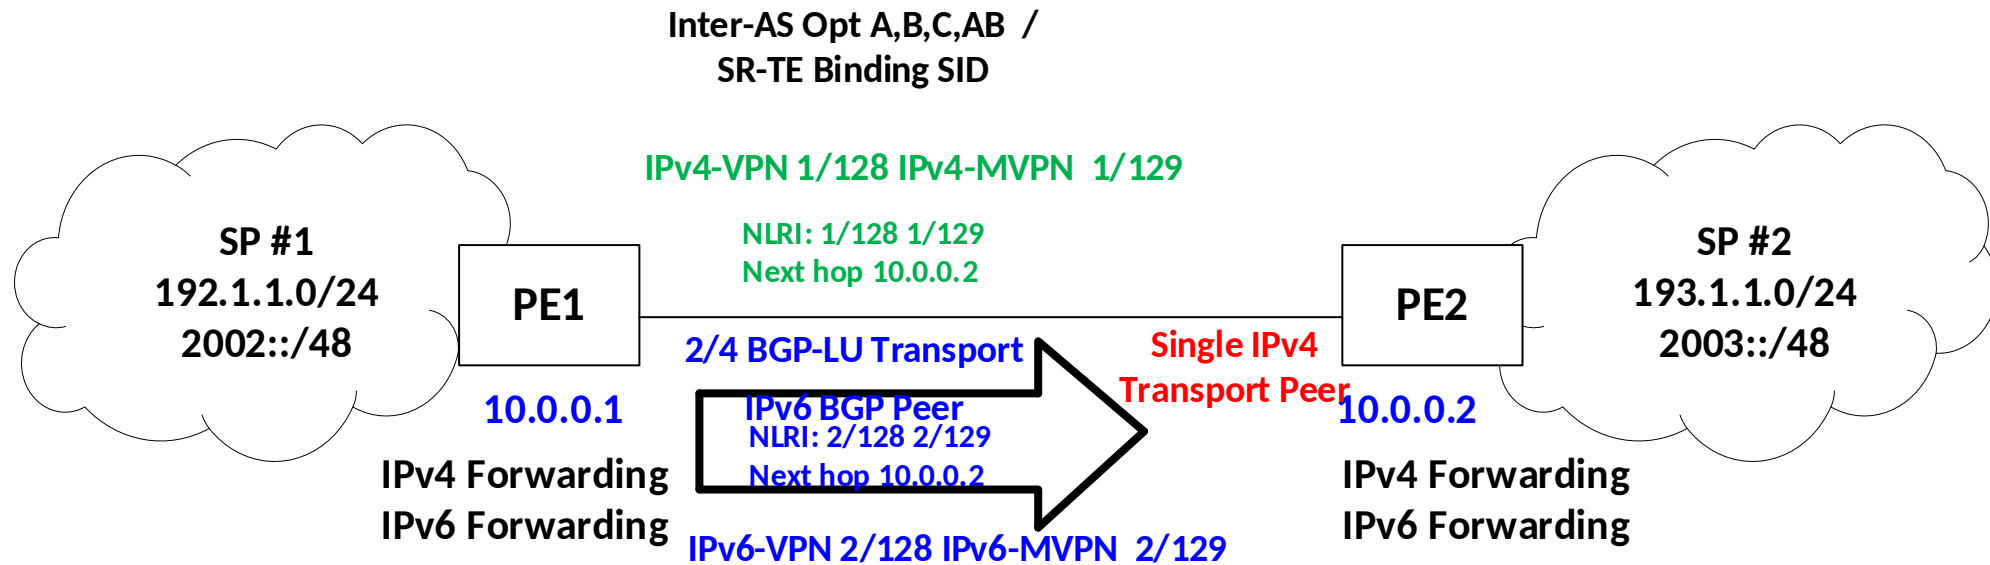


IPv4-Only PE Design ALL SAFI ☐ Inter-AS - Single IPv4 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

With this drafts standardized next hop encoding schema of 4 byte IPv4 next hop encoding for IPv6 Unicast SAFI & 12 byte IPv4 next hop encoding for VPN SAFI, both IPv4 & IPv6 NLRI can be advertised using a Single IPv4 peer.

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

IPv4-Only PE Design ⇔ Single IPv4-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI



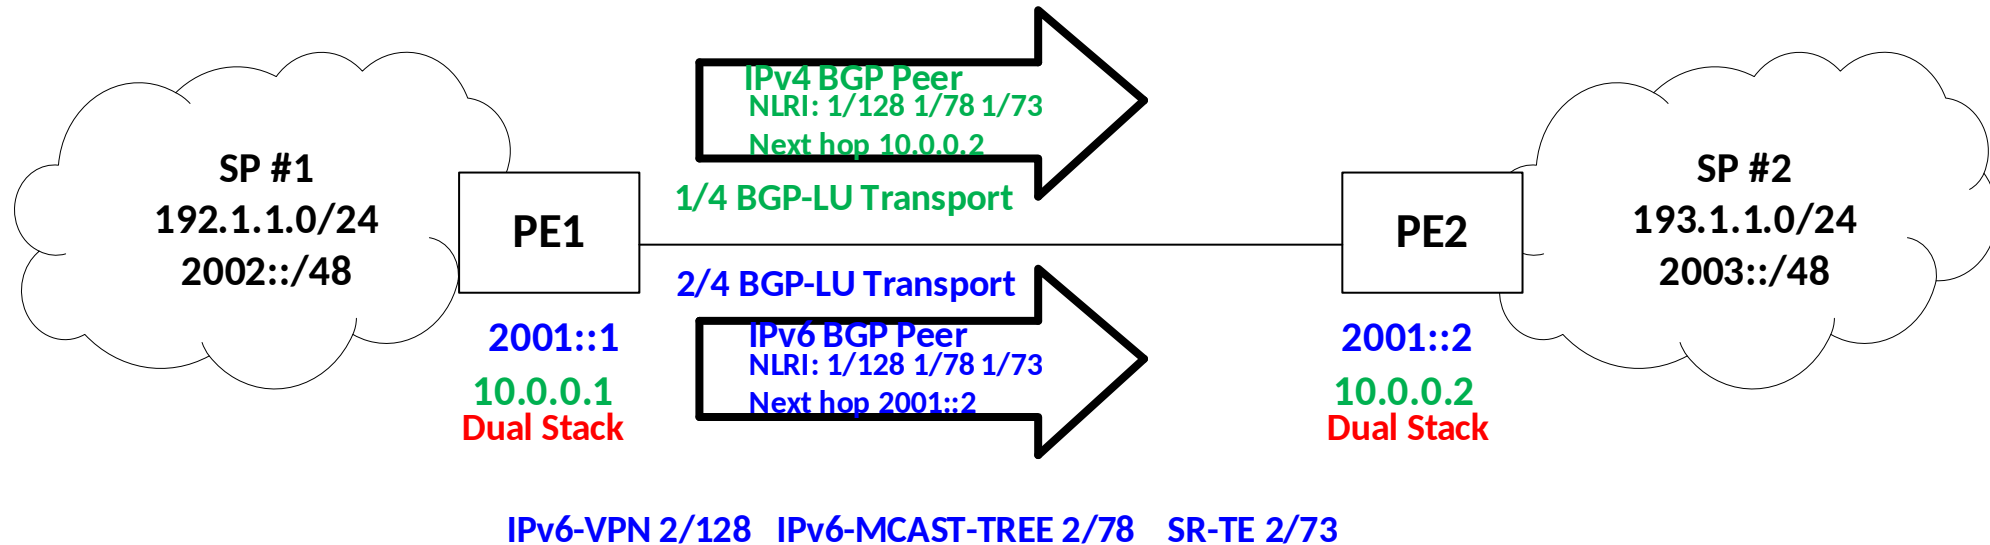
IXP Peering of all SPs would now be able to use a Single IPv4 peer per SP adjacency
Eliminates resource issues & provides CAPEX & OPEX Savings

Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv4-Only PE Design Slide Set)

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

Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering
IPv4-VPN 1/128 IPv4-MCAST-TREE 1/78 SR-TE 1/73

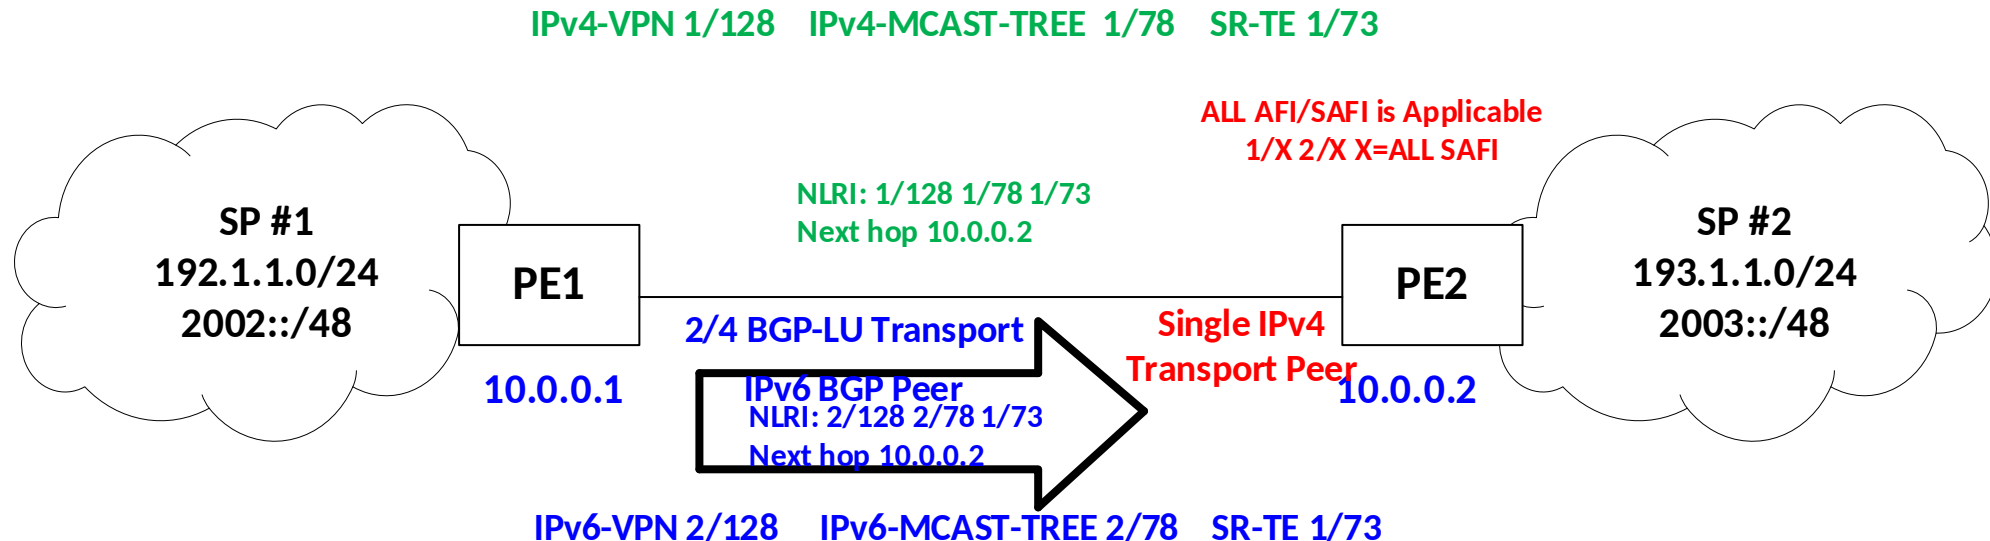


IPv4-Only PE Design ALL SAFI □ Inter-AS - Single IPv4 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

With this drafts standardized next hop encoding schema of 4 byte IPv4 next hop encoding for IPv6 Unicast SAFI & 12 byte IPv4 next hop encoding for VPN SAFI, both IPv4 & IPv6 NLRI can be advertised using a Single IPv4 peer.

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

IPv4-Only PE Design ↔ Single IPv4-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI
Inter-AS Opt A,B,C,AB /
SR-TE Binding SID

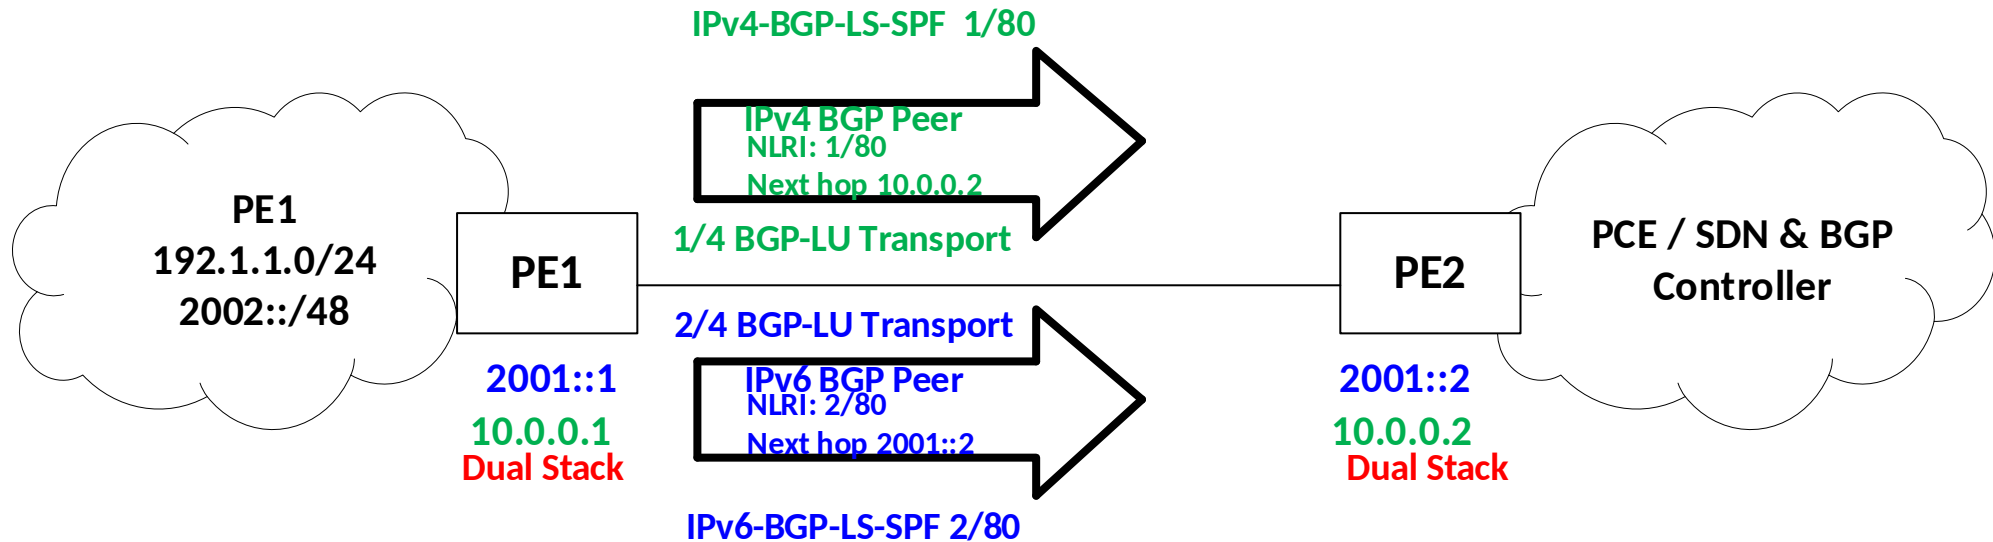


IXP Peering of all SPs would now be able to use a Single IPv4 peer per SP adjacency
Eliminates resource issues & provides CAPEX & OPEX Savings

Traditional PE-PE Inter-AS Dual Stacked Peering with Separate IPv4 & IPv6 Peer (IPv4-Only PE Design Slide Set)

PE to PCE/SDN or BGP Controller has a separate IPv4 & IPv6 peer with IPv4 NH peer carrying IPv4 NLRI & IPv6 NH peer carrying IPv6 NLRI
Inter-AS Opt A,B,C,AB / SR-TE Binding SID

Traditional Dual Stack ↔ Separate IPv4 & IPv6 peering



IPv4-Only PE Design ALL SAFI ☐ Inter-AS - Single IPv4 Peer carrying IPv4 & IPv6 NLRI (Same Dual Stack Functionality)

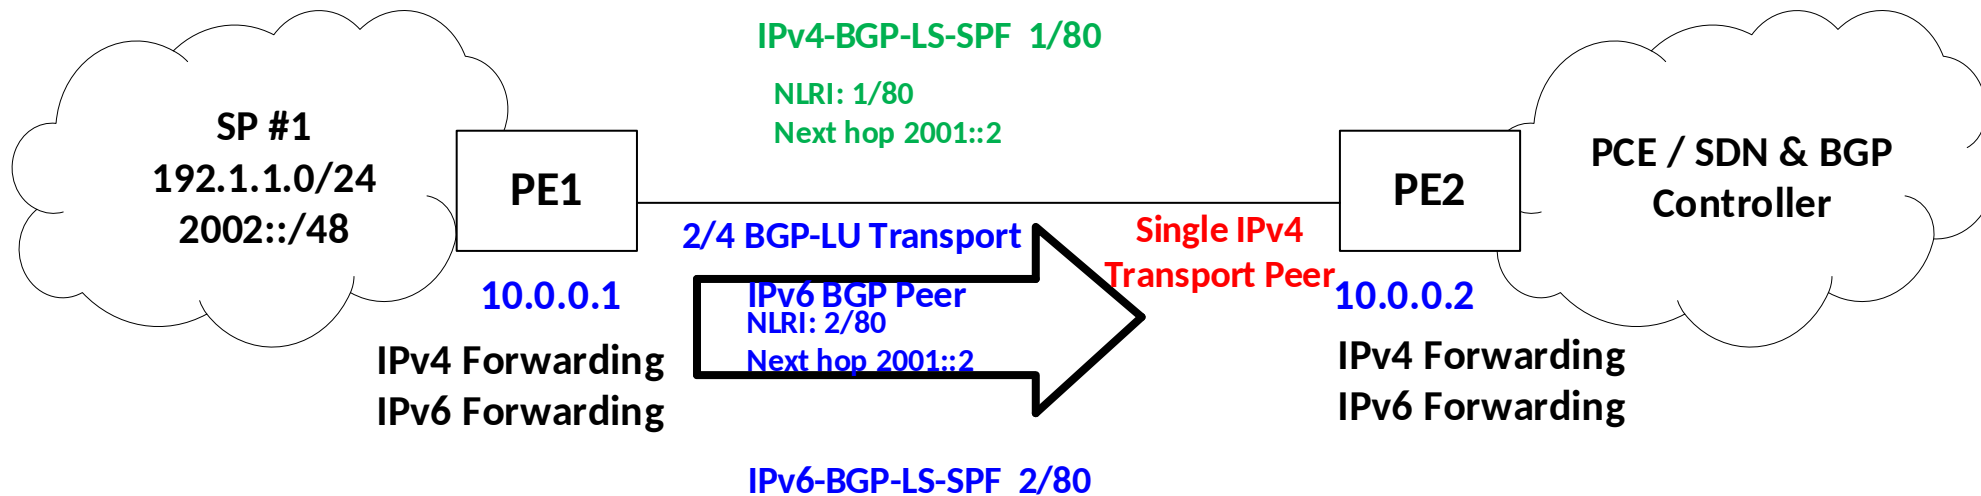
PE to PCE/SDN or BGP Controller Single IPv4 Peer

With this drafts standardized next hop encoding schema of 4 byte IPv4 next hop encoding for IPv6 Unicast SAFI & 12 byte IPv4 next hop encoding for VPN SAFI, both IPv4 & IPv6 NLRI can be advertised using a Single IPv4 peer.

This basic concept can eliminate all PE to Controller peering at the Edge and within the Core.

IPv4-Only PE Design ⇔ Single IPv4-Only Pure Transport Peer to carry both IPv4 & IPv6 NLRI

ALL AFI/SAFI is Applicable
1/X 2/X X=ALL SAFI



IXP Peering of all SPs would now be able to use a Single IPv4 peer per SP adjacency
Eliminates resource issues & provides CAPEX & OPEX Savings