

BESS Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 28 January 2024

G. Mishra  
Verizon Inc.  
J. Tantsura  
Microsoft, Inc.  
M. Mishra  
Cisco Systems  
S. Madhavi  
Juniper Networks, Inc.  
Q. Yang  
Arista Networks  
A. Simpson  
Nokia  
S. Chen  
Huawei Technologies  
27 July 2023

IPv4-Only PE Design All SAFI  
draft-mishra-bess-ipv4-only-pe-design-all-safi-05

## Abstract

As Enterprises and Service Providers try to decide whether or not to upgrade their brown field or green field MPLS/SR core to an IPv6 transport, Multiprotocol BGP (MP-BGP) now plays an important role in the transition of their Provider (P) core network as well as Provider Edge (PE) Edge network from IPv4 to IPv6. Operators must be able to continue to support IPv4 customers when both the Core and Edge networks are IPv4-Only.

This specification details an important External BGP (eBGP) PE-CE Edge IPv4-Only peering design that leverages the MP-BGP capability exchange by using IPv4 peering as pure transport, allowing both IPv4 Network Layer Reachability Information (NLRI) and IPv6 Network Layer Reachability Information (NLRI) to be carried over the same (Border Gateway Protocol) BGP TCP session. The design change provides the same Dual Stacking functionality that exists today with separate IPv4 and IPv6 BGP sessions as we have today. With this design change from a control plane perspective a single IPv4 is required for both IPv4 and IPv6 routing updates and from a data plane forwarding perspective an IPv4 address need only be configured on the PE and CE interface for both IPv4 and IPv6 packet forwarding.

This document defines the IPv4-Only PE Design as a new PE-CE Edge and ASBR-ASBR PE-PE Inter-AS BGP peering Standard defined in this specification which is now extended to support to all AFI/SAFI ubiquitously. As service providers migrate to Segment Routing architecture SR-MPLS and SRv6, VPN overlay exists as well, and thus

Inter-AS options Option-A, Option-B, Option-AB and Option-C are still applicable and thus this extension of IPv4-Only peering architecture extension to Inter-AS peering is very relevant to Segment Routing as well.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 January 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- 1. Introduction . . . . . 4
- 2. Requirements Language . . . . . 9
- 3. Terminology . . . . . 9
- 4. Problem Statement . . . . . 10
- 5. IPv4-Only PE-CE Design ALL SAFI Solution . . . . . 11
  - 5.1. IPv4-Only Edge Peering Design . . . . . 14
    - 5.1.1. IPv4-Only PE Design ALL SAFI 6to4 Software IPv4-Only Core packet walk . . . . . 14
    - 5.1.2. IPv4-Only PE Design ALL SAFI 6to4 Software - IPv6 Edge over an IPv4-Only Core Packet Walk . . . . . 15

5.1.3.	IPv4-Only PE Design ALL SAFI 4to6 Software - IPv4 Edge over an IPv6-Only Core Packet Walk . . . . .	17
6.	RFC5549 and RFC8950 Applicability to IPv4-Only PE Design . .	19
6.1.	IPv4-Only Edge Peering design next-hop encoding . . . . .	19
6.2.	IPv4-Only PE Design Next Hop Encoding . . . . .	19
7.	IPv4-Only PE Design Edge E2E Test Cases . . . . .	21
7.1.	Test-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software . . . . .	22
7.2.	Test-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 Software . . . . .	23
7.3.	Test-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE), 4to6 Software . . . . .	23
7.4.	Test-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 Software . . . . .	24
7.5.	Test-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-B . . . . .	24
7.6.	Test-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-C . . . . .	25
7.7.	Test-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only, 6to4 software -Inter-AS Option-B . . . . .	25
7.8.	Test-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 software -Inter-AS Option-C . . . . .	26
7.9.	Test-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-B . . . . .	26
7.10.	Test-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-C . . . . .	27
7.11.	Test-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-B . . . . .	27
7.12.	Test-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-C . . . . .	28
7.13.	IPv4-Only PE-CE Operational Considerations Testing . . .	28
8.	IPv4-Only PE Design Edge E2E Design for ALL AFI/SAFI . . . .	29
8.1.	IPv4-Only PE Design All SAFI Case-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software . .	29
8.2.	IPv4-Only PE Design All SAFI Case-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 Software . . . . .	30
8.3.	IPv4-Only PE Design All SAFI Case-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE), 4to6 Software .	30
8.4.	IPv4-Only PE Design All SAFI Case-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 Software . . . . .	31
8.5.	IPv4-Only PE Design All SAFI Case-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-B . . . . .	31
8.6.	IPv4-Only PE Design All SAFI Case-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-C . . . . .	32
8.7.	IPv4-Only PE Design All SAFI Case-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only, 6to4 software -Inter-AS Option-B . .	32

8.8.	IPv4-Only PE Design All SAFI Case-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 software -Inter-AS Option-C . . . . .	33
8.9.	IPv4-Only PE Design All SAFI Case-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-B . . . . .	33
8.10.	IPv4-Only PE Design All SAFI Case-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-C . . . . .	34
8.11.	IPv4-Only PE Design All SAFI Case-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-B . . . . .	34
8.12.	IPv4-Only PE Design All SAFI Case-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-C . . . . .	35
9.	IPv6-Only PE Design ALL AFI/SFI Operational Considerations .	35
10.	IANA Considerations . . . . .	36
11.	Security Considerations . . . . .	36
12.	Acknowledgments . . . . .	36
13.	Contributors . . . . .	36
14.	References . . . . .	37
14.1.	Normative References . . . . .	37
14.2.	Informative References . . . . .	40
Appendix A.	SAFI LIST FOR PE-CE Edge Scenario . . . . .	42
Appendix B.	SAFI LIST FOR Inter-AS PE-PE Scenario . . . . .	42
Authors'	Addresses . . . . .	45

## 1. Introduction

As Enterprises and Service Providers upgrade their brown field or green field MPLS/SR core to an IPv6 transport such as MPLS LDPv6, SR-MPLSv6 or SRv6, Multiprotocol BGP (MP-BGP) now plays an important role in the transition of the Provider (P) core networks and Provider Edge (PE) edge networks from IPv4 to IPv6. Operators have a requirement to support IPv6 customers and must be able to support IPv6 address family and Sub-Address-Family Virtual Private Network (VPN)-IPv6, and Multicast VPN IPv6 customers.

This specification provides a IPv4-Only PE design solution for use cases where operators are not yet ready to migrate to IPv6 or SRv6 core and would like to stay on IPv4-Only Core short to long term and maybe even indefinitely. With this design, operators can now remain with an IPv4-Only Core and do not have to migrate to an IPv6-Only Core. From a technical standpoint the underlay can remain IPv4 and still transport IPv6 NLRI to support IPv6 customers, and so does not need to be migrated to IPv6-Only underlay. With this IPv4-Only PE Design solution, IPv4 addressing only needs to be provisioned for the IPv4-Only PE-CE eBGP Edge peering design, thereby eliminating

IPv6 provisioning at the Edge. This core and edge IPv4-Only peering design can apply to any eBGP peering, public internet or private, which can be either Core networks, Data Center networks, Access networks or can be any eBGP peering scenario.

This document provides a IPv4-Only PE design solution for use cases where operators are not yet ready to migrate to IPv6 or SRv6 core and would like to stay on IPv4-Only Core short to long term and maybe even indefinitely. With this design, operators can now remain with an IPv4-Only Core and do not have to migrate to an IPv6-Only Core. From a technical standpoint the underlay can remain IPv4 and still transport IPv6 NLRI to support IPv6 customers, and so does not need to be migrated to IPv6-Only underlay. With this IPv4-Only PE Design solution, IPv4 addressing only needs to be provisioned for the IPv4-Only PE-CE eBGP Edge peering design, thereby eliminating IPv6 provisioning at the Edge. This core and edge IPv4-Only peering design can apply to any eBGP peering, public internet or private, which can be either Core networks, Data Center networks, Access networks or can be any eBGP peering scenario.

This document provides vendor specific test cases for the IPv4-Only peering design as well as test results for the five major vendors stakeholders in the routing and switching industry, Cisco, Juniper, Arista, Nokia and Huawei. With the test results provided for the IPv4-Only Edge peering design, the goal is that all other vendors around the world that have not been tested will begin to adopt and implement this new Best Current Practice for eBGP IPv4-Only Edge peering.

This Best Current Practice IPv4-only eBGP peering design specification will help in use cases where operators are not yet ready to migrate to IPv6 or SRv6 core or for very large operator core with thousands of nodes where it maybe impractical to change the underlay infrastructure to IPv6, and can now keep the existing IPv4 data plane IP, MPLS or SR-MPLS underlay intact indefinitely.

This document also defines a new IPv4 next hop encoding for IPv6 NLRI over IPv4 Next Hop to use 4 byte IPv4 address for the next hop and not a IPv4 mapped IPv6 address. This encoding has been adopted by the industry but has not been standardized until now with this document.

This document details an important External BGP (eBGP) PE-PE Inter-AS IPv6-Only peering design that leverages the MP-BGP capability exchange by using IPv6 peering as pure transport, allowing all and any IPv4 Network Layer Reachability Information (NLRI) and IPv6 Network Layer Reachability Information (NLRI) to be carried over the same (Border Gateway Protocol) BGP TCP session for all Address Family

Identifiers (AFI) and Subsequent Address Family Identifiers (SAFI). The design change provides the same Dual Stacking functionality that exists today with separate IPv4 and IPv6 BGP sessions as we have today. With this IPv4-Only PE Design, IPv6 address MUST not be configured on the the Provider Edge (PE) - Customer Edge (CE), or Inter-AS ASBR (Autonomous System Boundary Router) to ASBR (Autonomous System Boundary Router) PE-PE Provider Edge (PE) - Provider Edge (PE). From a control plane perspective a single IPv4-Only peer is required for both IPv4 and IPv6 routing updates and from a data plane forwarding perspective an IPv4 address need only be configured on the PE to PE Inter-AS peering interface for both IPv4 and IPv6 packet forwarding. This document defines the IPv4-Only PE Design as a new PE-CE Edge and ASBR-ASBR PE-PE Inter-AS BGP peering Standard defined in this specification which is now extended to support to all AFI/SAFI ubiquitously. As service providers migrate to Segment Routing architecture SR-MPLS and SRv6, VPN overlay exists as well, and thus Inter-AS options Option-A, Option-B, Option-AB and Option-C are still applicable and thus this extension of IPv4-Only peering architecture extension to Inter-AS peering is very relevant to Segment Routing as well.

With this IPv4-only BGP peering design, only IPv4 is configured on the PE-CE interface, the Provider Edge (PE) - Customer Edge (CE), the IPv4 BGP peer is now used to carry IPv6 (Network Layer Reachability Information) NLRI over an IPv4 next hop using 4 byte IPv4 next hop encoding while continuing to forward both IPv4 and IPv6 packets. In the framework of this design the PE is no longer Dual Stacked. However in the case of the CE, PE-CE link CE side of the link is no longer Dual Stacked, however all other internal links within the CE domain may or maynot be Dual stacked.

MP-BGP specifies that the set of usable next-hop address families is determined by the Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI). Historically the AFI/SAFI definitions for the IPv4 address family only have provisions for advertising a Next Hop address that belongs to the IPv4 protocol when advertising IPv4 or VPN-IPv4. [RFC8950] specifies the extensions necessary to allow advertising IPv4 NLRI, Virtual Private Network Unicast (VPN-IPv4) NLRI, Multicast Virtual Private Network (MVPN-IPv4) NLRI with a Next Hop address that belongs to the IPv6 protocol. This comprises of an extended next hop encoding MP-REACH BGP capability exchange to allow the address of the Next Hop for IPv4 NLRI, VPN-IPv4 NLRI and MVPN-IPv4 NLRI to also belong to the IPv6 Protocol. [RFC8950] defines the encoding of the Next Hop to determine which of the protocols the address actually belongs to, and a new BGP Capability allowing MP-BGP Peers to discover dynamically whether they can exchange IPv4 NLRI and VPN-IPv4 NLRI with an IPv6 Next Hop.

With the IPv4-Only PE design, IPv6 NLRI will be carried over an IPv4 Next-hop. [RFC4798] and [RFC4659] specify how an IPv4 address can be encoded inside the next-hop IPv6 address field when IPv6 NLRI needs to be advertised with an IPv4 next hop. [RFC4798] defines how the IPv4-mapped IPv6 address format specified in the IPv6 addressing architecture [RFC4798] can be used for that purpose when the <AFI/SAFI> is IPv6-Unicast <2/1>, Multicast <2/2>, and Labeled Unicast <2/4>. [RFC4659] defines how the IPv4-mapped IPv6 address format as well as a null Route Distinguisher as ::FFFF:192.168.1.1 (RD) can be used for that purpose when the <AFI/SAFI> is VPN-IPv6 <2/128> MVPN-IPv6 <2/129>. This IPv4-Only PE specification utilizes IPv6 NLRI over IPv4 Next hop encoding adopted by the industry to not use IPv4 mapped IPv6 address defined above, and instead use 4 byte IPv4 address for the next hop which ultimately set the precedence for the adoption of [RFC8950] for 4to6 Softwire IPv4 NLRI over IPv6 next-hop using an IPv6 address for the next hop and not a IPv6 mapped IPv4 address. The IPv4 next hop encoding for cases where the NLRI advertised is different from the next hop encoding such as where IPv6 NLRI is advertised with IPv4 next hop for for <AFI/SAFI> is IPv6-Unicast <2/1>, Multicast <2/2>, and Labeled Unicast <2/4>. [RFC4659] defines Null(RD) for <AFI/SAFI> is VPN-IPv6 <2/128> MVPN-IPv6 <2/129> but now with a an official new IANA Capability code TBD as value 10 "IPv4 Next Hop Encoding". The IETF standards have not been updated with an IANA allocation Capability code for the IPv4 next hop encoding so this specification fixes that problem with an IANA allocated capability codepoint which will now be used for any eBGP or iBGP peering as well as the IPv4-Only PE design defined in this specification.

With this IPv4-Only PE Design, BGP peer session can now be treated as a pure TCP transport and carry both IPv4 and IPv6 NLRI at the Provider Edge (PE) - Customer Edge (CE) over a single IPv4 TCP session. This allows for the elimination of dual stack from the PE-CE peering point, and now enable the peering to be IPv4-ONLY. The elimination of IPv6 on the PE-CE peering points translates into OPEX expenditure savings of point-to-point infrastructure links as well as /127 address space savings and administration and network management of both IPv4 and IPv6 BGP peers. This reduction decreases the number of PE-CE BGP peers by fifty percent, which is a tremendous cost savings for operators. This also translates into Major CAPEX savings as now operators do not have to migrate their underlay to IPv6 and can remain indefinitely on IPv4-Only Core.

While the savings exists at the Edge eBGP PE-CE peering, on the core side PE to Route Reflector (RR) peering carrying <AFI/SAFI> IPv4 <2/1>, VPN-IPV4 <2/128>, and Multicasat VPN <2/129>, there is no savings as the Provider (P) Core is IPv4 Only and thus can only have an IPv4 peer standard 4 byte next hop encoding to carrying IPv4 NLRI IPV4 <2/1>, VPN-IPV4 <2/128>, and Multicasat VPN <2/129> over an IPv4 next hop.

This document defines the IPv4-Only PE Design Architecture details for External BGP (eBGP) PE-PE Inter-AS IPv6-Only peering design that leverages the MP-BGP capability exchange by using IPv6 peering as pure transport, allowing all and any IPv4 Network Layer Reachability Information (NLRI) and IPv6 Network Layer Reachability Information (NLRI) to be carried over the same (Border Gateway Protocol) BGP TCP session for all Address Family Identifiers (AFI) and Subsequent Address Family Identifiers (SAFI). The design change provides the same Dual Stacking functionality that exists today with separate IPv4 and IPv6 BGP sessions as we have today. With this IPv6-Only PE Design, IPv6 address MUST not be configured on the the Provider Edge (PE) - Customer Edge (CE), or Inter-AS ASBR (Autonomous System Boundary Router) to ASBR (Autonomous System Boundary Router) PE-PE Provider Edge (PE) - Provider Edge (PE). From a control plane perspective a single IPv4-Only peer MUST be configured for both IPv4 and IPv6 routing updates, and from a data plane forwarding perspective only an IPv4 address MUST be configured on the PE-CE Edge or ASBR-ASBR, PE to PE Inter-AS peering interface for both IPv4 and IPv6 packet forwarding for all AFI/SAFI. This document defines the IPv4-Only PE Design as a new Intra-AS PE-CE Edge and Inter-AS PE-PE BGP peering Standard which is described in detail in the test use cases of the design which is extended for applicability to to all AFI/SAFI ubiquitously. As service providers migrate to Segment Routing architecture SR-MPLS and SRv6, VPN overlay exists as well, and thus Inter-AS options Option-A, Option-AB and Option-C are still applicable and thus this extension of IPv4-Only peering architecture extension to Inter-AS peering is very relevant to Segment Routing as well as well as any other applicable AFI/SAFI is now as well relevant.

This core and edge IPv4-Only peering design paradigm change can apply to any eBGP peering, public internet or private, which can be either Core networks, Data Center networks, Access networks or can be any eBGP peering scenario. This document provides detailed vendor specific test cases and test results for the IPv4-Only peering design as well as successful test results between five major vendors stakeholders in the routing and switching industry, Cisco, Juniper, Arista, Nokia and Huawei. With the test results provided for the IPv4-Only Edge peering design, the goal is that all other vendors around the world that have not been tested will begin to adopt and implement this new best practice for eBGP IPv4-Only Edge peering.

This IPv4-Only PE ALL SAFI Design details an important External BGP (eBGP) PE-PE Inter-AS IPv4-Only peering design that leverages the MP-BGP capability exchange by using IPv6 peering as pure transport, allowing all and any IPv4 Network Layer Reachability Information (NLRI) and IPv6 Network Layer Reachability Information (NLRI) to be carried over the same (Border Gateway Protocol) BGP TCP session for all remaining Address Family Identifiers (AFI) and Subsequent Address Family Identifiers (SAFI) below as well that can be carried over IPv4-Only Inter-AS peerings: <AFI/SAFI> MCAST-VPN [RFC6514] <1/5>, NLRI Multi-Segment Pseudowires [RFC7267] <1/6>, BGP Tunnel Encapsulation SAFI [RFC9012] <1/7>, MCAST-VPLS [RFC7117] <1/8>, Tunnel SAFI [I-D.nalawade-kapoor-tunnel-safi] <1/6>, BGP MDT SAFI [RFC6037] <1/66>, BGP 4to6 SAFI [RFC5747] <1/67>, BGP 6to4 SAFI draft xx <1/8>, Layer 1 VPN Auto-Discovery [RFC5195] <1/69>, SR-TE Policy SAFI draft <1/73>, BGP 6to4 SAFI draft <1/8>, SDN WAN Capabilities draft <1/74>, Classful-Transport SAFI draftxx <1/76>, Tunneled Traffic FlowSpec draftxx <1/77>, MCAST-TREE SAFI draft xx <1/78>, Route Target Constraints [RFC4684] <1/132>, Dissemination of Flow Specification Rules [RFC8955] <1/133>, L3 VPN Dissemination of Flow Specification Rules [RFC8955] <1/1344>, VPN Auto-Discovery SAFI draftxx <1/140>

## 2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 3. Terminology

Terminology used in defining the IPv6-Only Edge specification.

AFBR: Address Family Border Router Provider Edge (PE).

Edge: PE-CE Edge Network Provider Edge - Customer Edge

Core: P Core Network Provider (P)

4to6 Softwire : IPv4 edge over an IPv6-Only core

6to4 Softwire: IPv6 edge over an IPv4-Only core

E2E: End to End

4. Problem Statement

This specification addresses a real issue that has been discussed at many operator with extremely large core networks around the world related migration to IPv6 underlay transport which can be put off indefinitely. Operators around the world are clamoring for a solution that can help solve issues related to IPv4 address depletion at these large IXP peering points. With this solution, infrastructure networks such as Core networks, DC networks, Access networks as well as any PE-CE public or private network can now utilize this IPv4-Only Edge solution and reap the benefits immediately on IPv6 address space saving and CAPEX and OPEX savings.

Problem Statement

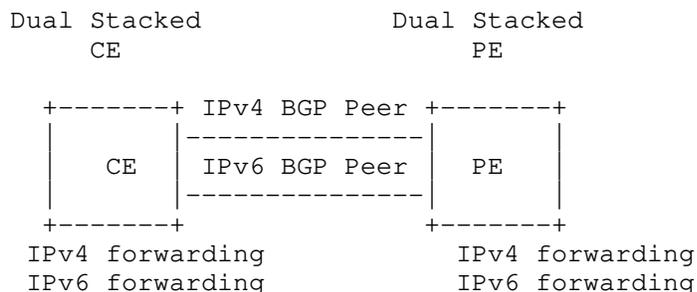


Figure 1: Problem Statement - Dual Stack Peering

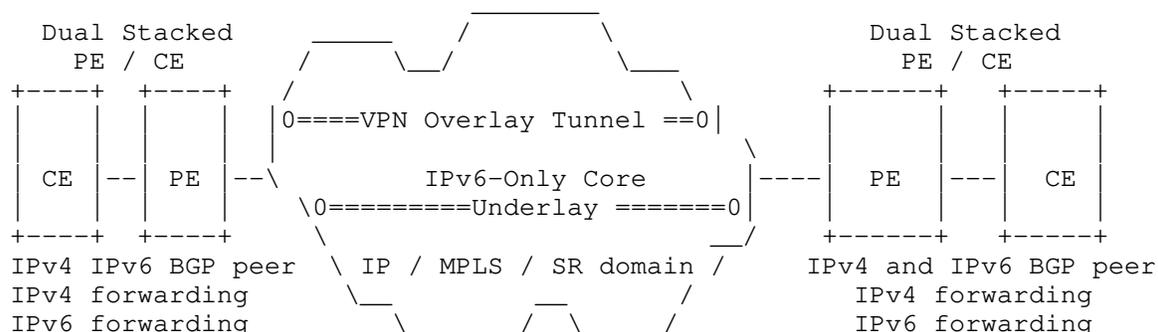


Figure 2: Problem Statement - E2E Dual Stack Edge

5. IPv4-Only PE-CE Design ALL SAFI Solution

The IPv4-Only Edge design solution applies to any and all IPv4 Network Layer Reachability Information (NLRI) and IPv6 Network Layer Reachability Information (NLRI) over an IPv4-Only BGP Peering session.

2 Main scenario groupings for the IPv4-Only PE Design ALL SAFI

- \* Group-1 PE-CE
- \* Group-2 PE-PE Inter-AS
- 1. Group-1 Edge Customer IPv4/IPv6 NLRI PE-CE AFI / SAFI grouping (CP-DP)
  - 1. Unicast
    - \* AFI/SAFI 1/1 IPv4 Unicast 2/1 IPv6 Unicast
  - 1. Multicast
    - \* AFI/SAFI 1/2 IPv4 Multicast 2/2 IPv6 Multicast
    - \* AFI/SAFI 1/78 IPv4 MCAST-TREE SAFI 2/78 MCAST-TREE SAFI
- 1. Group-2 ASBR-ASBR Inter AS Customer IPv4/IPv6 NLRI AFI/SAFI grouping (CP-DP)
  - a. Unicast

1. Global Table

- \* AFI/SAFI 2/4 6PE

1. L3 VPN

- \* AFI/SAFI 1/128 IPv4 VPN 2/128 IPv6 VPN

- \* AFI/SAFI 1/132 IPv4 RTC 2/132 IPv6 RTC

- \* AFI/SAFI 1/140 IPv4 VPN Auto Discovery 2/140 IPv6 VPN Auto Discovery

- a. Multicast

1. Global Table

- \* AFI/SAFI 1/2 IPv4 Multicast 2/2 IPv6 Multicast

1. L3 VPN

- \* AFI/SAFI 1/129 IPv4 MVPN 2/129 IPv6 MVPN

- \* AFI/SAFI 1/5 IPv4 MCAST-VPN 2/5 IPv6 MCAST-VPN

- \* AFI/SAFI 1/66 IPv4 BGP MDT SAFI 2/66 IPv6 BGP MDT SAFI

- \* AFI/SAFI 1/78 IPv4 MCAST-TREE SAFI 2/78 MCAST-TREE SAFI

1. L1 and L2 VPN

- \* L2 VPN related NLRI control plane in BGP

- \* AFI/SAFI 1/6 IPv4 Multi Segment PW 2/6 IPv6 Multi Segment PW

- \* AFI/SAFI 1/69 L1 VPN Auto Discovery 2/69 L1 VPN Auto Discovery

1. Multicast

- \* AFI/SAFI 1/8 IPv4 MCAST-VPLS 2/8 IPv6 MCAST-VPLS

1. Tunnel

- \* AFI/SAFI 1/64 Tunnel-SAFI 2/64 Tunnel-SAFI

- \* AFI/SAFI 1/67 BGP 4over6 Tunnel SAFI 2/67 BGP 4over6 Tunnel SAFI

- \* AFI/SAFI 1/68 BGP 6over4 Tunnel SAFI 2/68 BGP 6over4 Tunnel SAFI

- 1. BGP Policy and Flowspec Policy and Misc Policy
- \* AFI/SAFI 1/73 SR-TE Policy SAFI 2/73 SR-TE Policy SAFI
- \* AFI/SAFI 1/74 SD-WAN Capabilities 2/74 SD-WAN Capabilities
- \* AFI/SAFI 1/77 Tunneled Traffic Flowspec 2/77 Tunneled Traffic Flowspec
- \* AFI/SAFI 1/133 Dissemination of Flowspec Rules 2/133 133 Dissemination of Flowspec Rules
- \* AFI/SAFI 1/134 L3 VPN 133 Dissemination of Flowspec Rules 2/134 L3VPN Dissemination of Flowspec Rules
- \* AFI/SAFI 1/79 BGP-DPS Arista 2/79 BGP-DPS Arista
- \* AFI/SAFI 1/83 BGP CAR 2/83 BGP CAR
- \* AFI/SAFI 1/84 BGP VPN CAR 2/84 BGP VPN CAR
- \* AFI/SAFI 1/85 BGP MUP SAFI 2/85 BGP MUP SAFI

Solution applicable to all AFI/SAFI  
 AFI/SAFI 1/X 2/X Where X = ALL SAFI

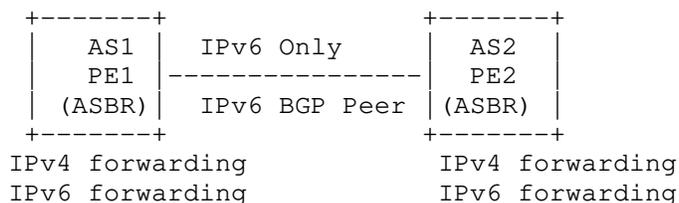


Figure 3: IPv4-Only PE Design Solution Applicability to ALL AFI/SAFI

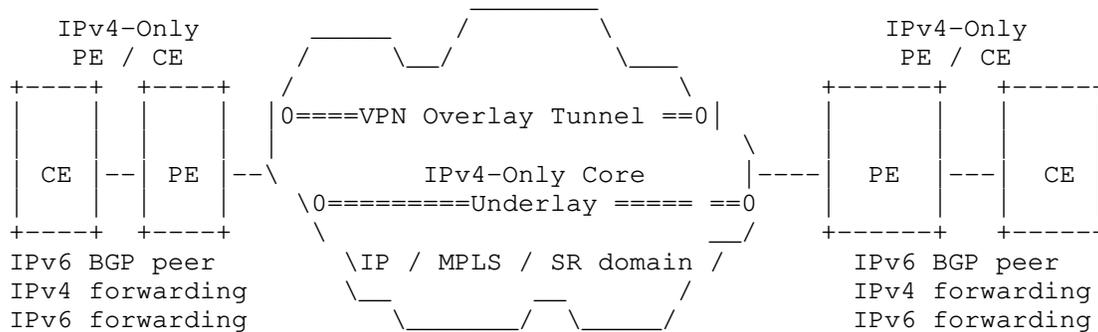


Figure 4: E2E VPN Solution

5.1. IPv4-Only Edge Peering Design

5.1.1. IPv4-Only PE Design ALL SAFI 6to4 Software IPv4-Only Core packet walk

The IPv4-Only Edge Peering design utilizes two key E2E Software Mesh Framework scenario's, 4to6 software and 6to4 software. The Software mesh framework concept is based on the overlay and underlay MPLS or SR based technology framework, where the underlay is the transport layer and the overlay is a Virtual Private Network (VPN) layer, and is the the tunneled virtualization layer containing the customer payload. The concept of a 6to4 Software is based on transmission of IPv6 packets at the edge of the network by tunneling the IPv6 packets over an IPv4-Only Core. The concept of a 4to6 Software is also based on transmission of IPv4 packets at the edge of the network by tunneling the IPv4 packets over an IPv6-Only Core.

This document describes End to End (E2E) test scenarios that follow a packet flow from IPv4-Only attachment circuit from ingress PE-CE to egress PE-CE tracing the routing protocol control plane and data plane forwarding of IPv4 packets in a 4to6 software or 6to4 software within the IPv4-Only or IPv6-Only Core network. In both secneario we are focusing on IPv4 packets and the control plane and data plane forwarding aspects of IPv4 packets from the PE-CE Edge network over an IPv4-Only P (Provider) core network or IPv6-Only P (Provider) core network. With this IPv4-Only Edge peering design, the Software Mesh Framework is not extended beyond the Provider Edge (PE) and continues to terminate on the PE router.

### 5.1.2. IPv4-Only PE Design ALL SAFI 6to4 Softwire - IPv6 Edge over an IPv4-Only Core Packet Walk

6to4 softwire where IPv4-Edge eBGP IPv4 peering where IPv6 packets at network Edge traverse a IPv4-Only Core

In the scenario where IPv6 packets originating from a PE-CE edge are tunneled over an MPLS or Segment Routing IPv4 underlay core network, the PE and CE only have an IPv6 address configured on the interface. In this scenario the IPv6 packets that ingress the CE from within the CE AS are over an IPv4-Only interface and are forwarded to an IPv6 NLRI destination prefix learned from the Pure Transport Single IPv4 BGP Peer. In the IPv4-Only Edge peering architecture the PE is IPv4-Only as all PE-CE interfaces are IPv4-Only. However, on the CE, the PE-CE interface is the only interface that is IPv4-Only and all other interfaces may or may not be IPv4-Only. Following the data plane packet flow, IPv4 packets are forwarded from the ingress CE to the IPv4-Only ingress PE where the VPN label imposition push per prefix, per-vrf, per-CE occurs and the labeled packet is forwarded over a 6to4 softwire IPv4-Only core, to the egress PE where the VPN label disposition pop occurs and the native IPv4 packet is forwarded to the egress CE. In the reverse direction IPv4 packets are forwarded from the egress CE to egress PE where the VPN label imposition per prefix, per-vrf, per-CE push occurs and the labeled packet is forwarded back over the 6to4 softwire IPv4-Only core, to the ingress PE where the VPN label disposition pop occurs and the native IPv4 packet is forwarded to the ingress CE. . The functionality of the IPv4 forwarding plane in this scenario is identical from a data plane forwarding perspective to Dual Stack IPv4 forwarding scenario.

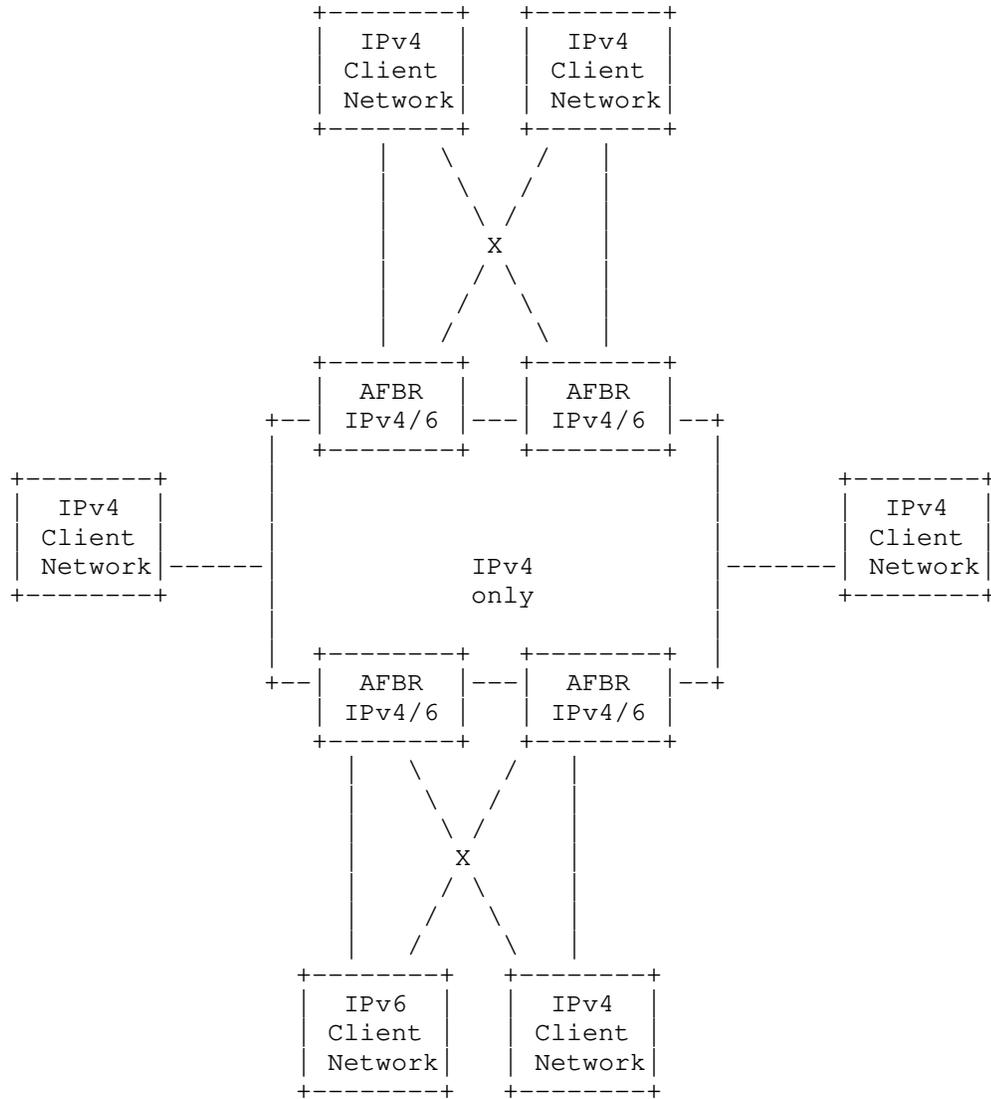


Figure 5: 6to4 Software - IPv6 Edge over an IPv4-Only Core Packet Walk

### 5.1.3. IPv4-Only PE Design ALL SAFI 4to6 Softwire - IPv4 Edge over an IPv6-Only Core Packet Walk

4to6 softwire where IPv4-Edge eBGP IPv4 peering where IPv6 packets at network Edge traverse a IPv6-Only Core

In the scenario where IPv6 packets originating from a PE-CE edge are tunneled over an MPLS or Segment Routing IPv4 underlay core network, the PE and CE only have an IPv4 address configured on the interface. In this scenario the IPv6 packets that ingress the CE from within the CE AS are over an IPv4-Only interface and are forwarded to an IPv6 NLRI destination prefix learned from the Pure Transport Single IPv4 BGP Peer. In the IPv4-Only Edge peering architecture the PE is IPv4-Only as all PE-CE interfaces are IPv4-Only. However, on the CE, the PE-CE interface is the only interface that is IPv4-Only and all other interfaces may or may not be IPv4-Only. Following the data plane packet flow, IPv6 packets are forwarded from the ingress CE to the IPv4-Only ingress PE where the VPN label imposition push per prefix, per-vrf, per-CE occurs and the labeled packet is forwarded over a 4to6 softwire IPv6-Only core, to the egress PE where the VPN label disposition pop occurs and the native IPv6 packet is forwarded to the egress CE. In the reverse direction IPv6 packets are forwarded from the egress CE to egress PE where the VPN label imposition per prefix, per-vrf, per-CE push occurs and the labeled packet is forwarded back over the 4to6 softwire IPv6-Only core, to the ingress PE where the VPN label disposition pop occurs and the native IPv6 packet is forwarded to the ingress CE. . The functionality of the IPv4 forwarding plane in this scenario is identical from a data plane forwarding perspective to Dual Stack IPv4 / IPv6 forwarding scenario.

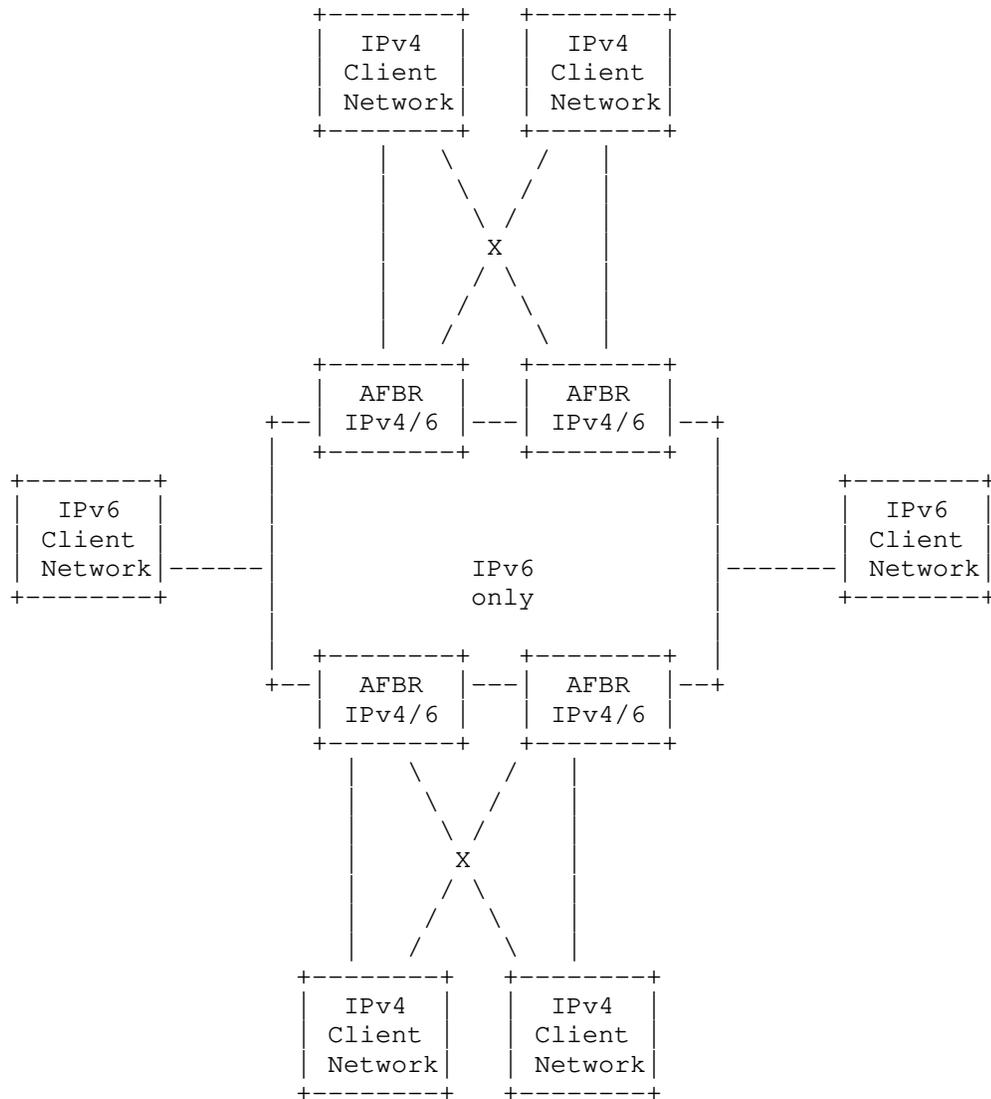


Figure 6: 4to6 Software - IPv4 Edge over an IPv6-Only Core Packet Walk

## 6. RFC5549 and RFC8950 Applicability to IPv4-Only PE Design

### 6.1. IPv4-Only Edge Peering design next-hop encoding

This section describes [RFC8950] next hop encoding updates to [RFC5549] applicability to this specification. IPv4-Only eBGP Edge PE-CE peering to carry IPv4 Unicast NLRI <AFI/SAFI> IPv4 <1/1> over an IPv6 next hop BGP capability extended hop encoding IANA capability codepoint value 5 defined is applicable to both [RFC5549] and [RFC8950] as IPv4 Unicast NLRI <AFI/SAFI> IPv4 <1/1> does not change in the RFC updates.

IPv4 packets over an IPv6-Only core 4to6 Software E2E packet flow is part of the IPv6-Only PE design and this same style next hop encoding applies to 6to4 Software IPv6 NLRI over IPv4 next hop with 4 byte Next hop encoding and not IPv4 mapped IPv6 address. [RFC8950] updates [RFC5549] for <AFI/SAFI> VPN-IPV4 <1/128>, and Multicasat VPN <1/129>

### 6.2. IPv4-Only PE Design Next Hop Encoding

This section describes IPv4 Next Hop Encoding for IPv6 NLRI over an IPv4 Next hop.

With the IPv4-Only PE design, IPv6 NLRI will be carried over an IPv4 Next-hop. [RFC4798] and [RFC4659] specify how an IPv4 address can be encoded inside the next-hop IPv6 address field when IPv6 NLRI needs to be advertised with an IPv4 next hop. [RFC4798] defines how the IPv4-mapped IPv6 address format specified in the IPv6 addressing architecture [RFC4798] can be used for that purpose when the <AFI/SAFI> is IPv6-Unicast <2/1>, Multicast <2/2>, and Labeled Unicast <2/4>. [RFC4659] defines how the IPv4-mapped IPv6 address format as well as a null Route Distinguisher as ::FFFF:192.168.1.1 (RD) can be used for that purpose when the <AFI/SAFI> is VPN-IPv6 <2/128> MVPN-IPv6 <2/129>. This IPv4-Only PE specification utilizes IPv6 NLRI over IPv4 Next hop encoding adopted by the industry to not use IPv4 mapped IPv6 address defined above, and instead use 4 byte IPv4 address for the next hop which ultimately set the precedence for the adoption of [RFC8950] for 4to6 Software IPv4 NLRI over IPv6 next-hop. The IPv4 next hop encoding for cases where the NLRI advertised is different from the next hop encoding such as where IPv6 NLRI is advertised with IPv4 next hop for for <AFI/SAFI> is IPv6-Unicast <2/1>, Multicast <2/2>, and Labeled Unicast <2/4>. [RFC4659] defines Null (RD) for <AFI/SAFI> is VPN-IPv6 <2/128> MVPN-IPv6 <2/129> but now with an official new IANA Capability code TBD for "IPv4 Next Hop Encoding". The IETF standards have not been updated with an IANA allocation Capability code for the IPv4 next hop encoding so this specification fixes that with an IANA allocated codepoint which will now be used for any eBGP or iBGP peering as well as the IPv4-Only PE design defined in this specification.

With this specification when VPN-IPv6 AFI/SAFI 2/128, MVPN-IPv6 AFI/SAFI 2/129 is used, the next-hop address is encoded as an IPv4 address with a length of 12 bytes. The next-hop address is now encoded for VPN-IPv6 AFI/SAFI with a length of 12 bytes. The 12 byte next hop includes 4 byte IPv4 address plus 8 byte Route Distinguisher. This document modifies how the next-hop address is encoded to accommodate all existing implementations and bring consistency with VPN-IPv6, MVPN-IPv6 and 6PE. As all known and deployed implementations are interoperable today and use the new proposed encoding, the change does not break existing interoperability. This change is applicable to all iBGP and eBGP peering as well as the IPv4-Only PE, PE-CE edge and Inter-AS peering design for the IPv4 next hop encoding E2E test case of IPv4 packets over and IPv4-Only core 6to4 Software. In this test case IPv6 Unicast NLRI <AFI/SAFI> IPv4 <1/1> is advertised over the PE to RR core peering 6to4 software in <AFI/SAFI> VPN-IPV6 <2/128> MVPN-IPv6 <2/129>. In this test cases label allocation mode comes into play which is discussed in a subsequent section.

This document defines with the new IANA BGP Capability codepoint allocation next hop encoding of MP\_REACH\_NLRI with:

- \* Specifically, this document allows advertising the MP\_REACH\_NLRI attribute [RFC2545] with this content:

Advertising with [RFC2545] MP\_REACH\_NLRI with:

- \* AFI = 2
- \* SAFI = 1, 2 or 4
- \* Length of Next Hop Address = 4
- \* Specifically, this document allows advertising the MP\_REACH\_NLRI attribute [RFC2545] with this content:

Advertising with [RFC2545] MP\_REACH\_NLRI with:

- \* AFI = 2
- \* SAFI = 128 or 129
- \* Length of Next Hop Address = 8
- \* Next Hop Address = VPN-IPv4 address of next hop with an 8-octet RD set to zero.

## 7. IPv4-Only PE Design Edge E2E Test Cases

Proof of concept interoperability testing of the 4 test cases between the 5 vendors Cisco, Juniper, Arista, Nokia and Huawei.

Cisco, Juniper, Arista, Nokia, Huawei, platform, code revision and test results for all use cases

Cisco: Edge Router- XR ASR 9910 IOS XR 7.4.1, Core Router- NCS 6000 7.2.2, CRS-X 6.7.4

Juniper: Edge Router- MX platform MX480, MX960, Core Router- PTX Platform PTX5000, PTC10K8 (JUNOS and EVO) Release 20.4R2

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

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

Arista:

Intra-AS tests PE-CE Edge Peering IPv4-Only Core, IPv6-Only Core, Global Table (GRT) and IP VPN

AFI/SAFI IPv4-Unicast SAFI IPv6-Unicast SAFI

IPv4 Core:

Test-1 Global table (6PE)

Test-2 IP VPN

Global table IPv6

IPv6 Core:

Test-3 Global table

Test-4 IP VPN

Inter-AS Options tests IPv4-Only Core, IPv6-Only Core, Global Table (GRT) and IP VPN

AFI/SAFI VPN and MVPN

IPv4-Only Core

Test-5 Global table 6PE Option-B (Segmented LSP stitched IPv4 Core - Inter-AS Link IPv6-Only PE - IPv4 Core)

Test-6 Global table 6PE Option-C (Redistribute IPv4 Loopbacks into BGP-LU AFI/SAFI 2/6)

Test-7 IP VPN Inter AS Option-B (Segmented LSP stitched IPv4 Core - Inter-AS Link IPv6-Only PE - IPv4 Core)

Test-8 IP VPN Inter AS Option-C (Redistribute IPv4 Loopbacks into BGP-LU AFI/SAFI 2/6)

IPv6-Only Core

Test-9 Global table Option-B

Test-10 Global table Option-C

Test-11 IP VPN Inter AS Option-B

Test-12 IP VPN Inter AS Option-C

7.1. Test-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software

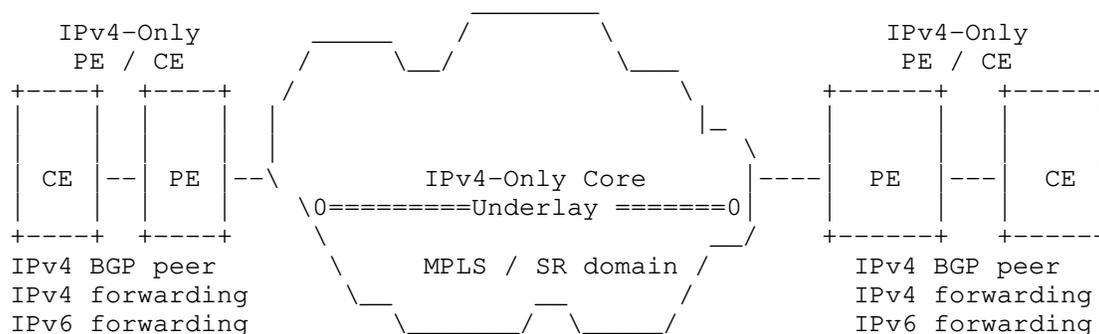


Figure 7: Test-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE)

7.2. Test-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 Software

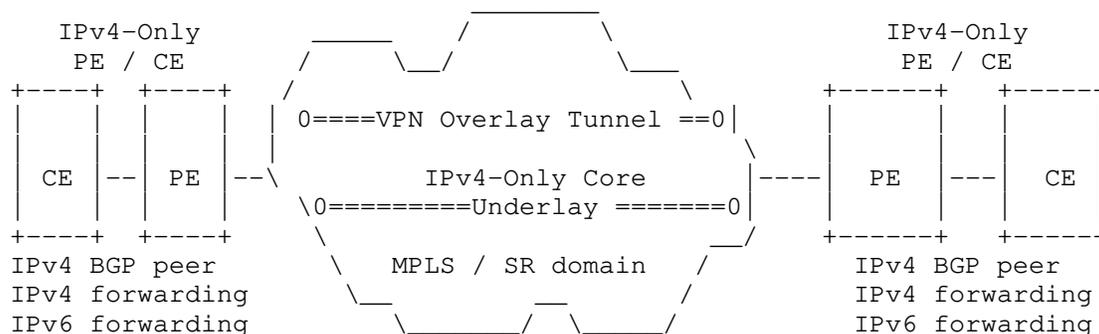


Figure 8: Test-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core

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

7.3. Test-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE), 4to6 Software

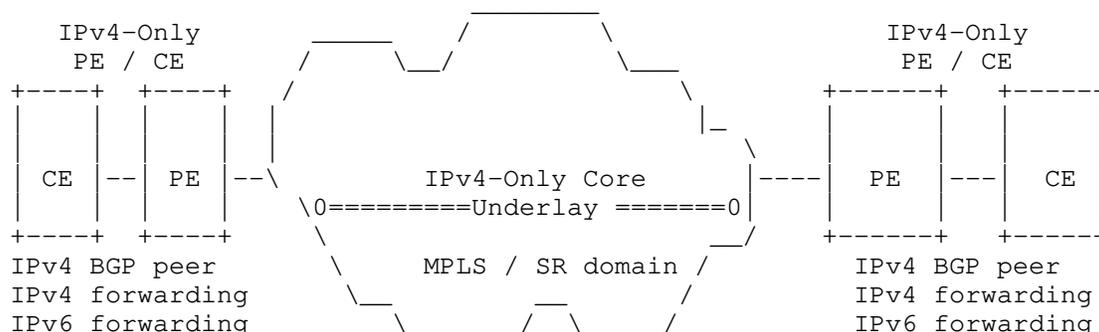


Figure 9: Test-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE)

7.4. Test-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 Software

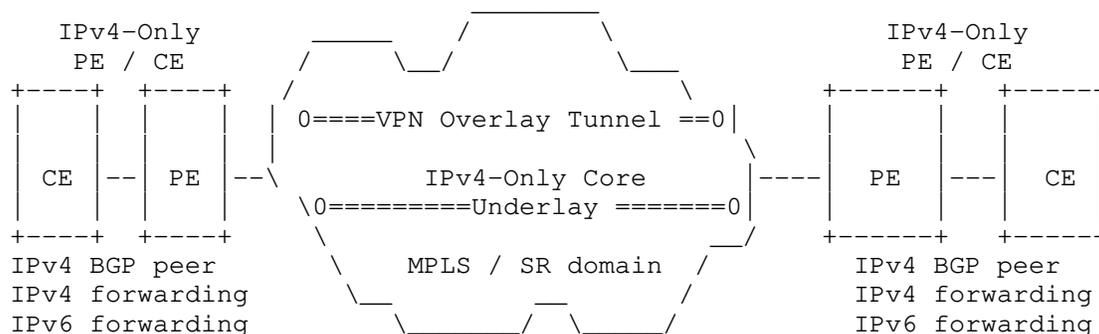


Figure 10: Test-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core

7.5. Test-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-B

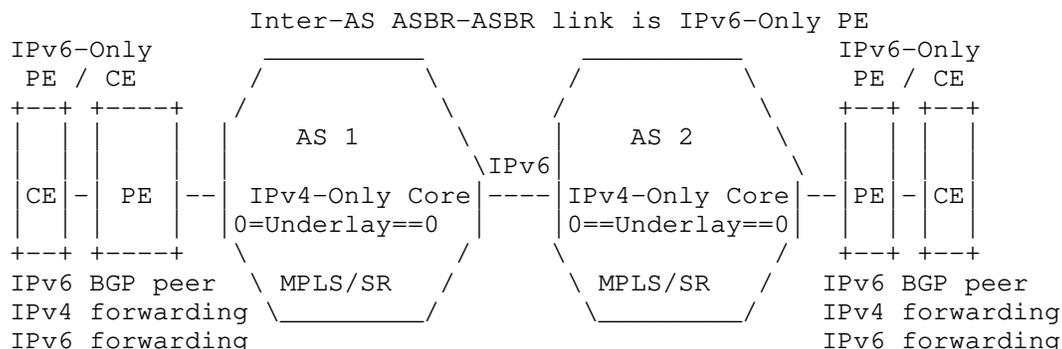


Figure 11: Test-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-B

7.6. Test-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-C

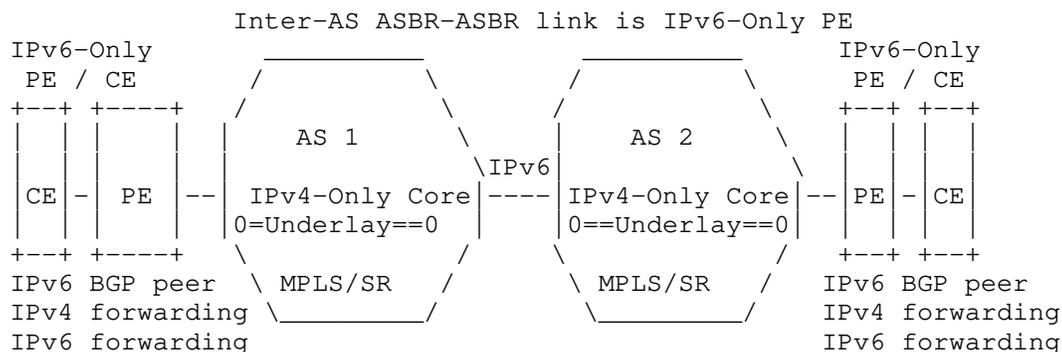


Figure 12: Test-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-C

7.7. Test-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only, 6to4 software - Inter-AS Option-B

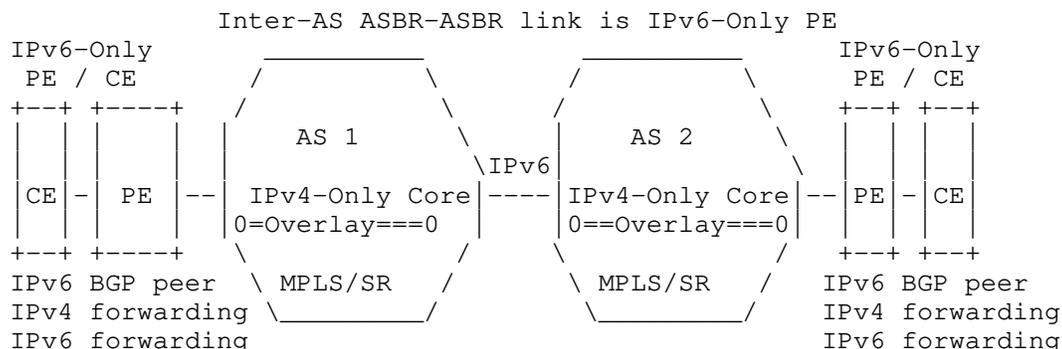


Figure 13: Test-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-B

7.8. Test-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 software -Inter-AS Option-C

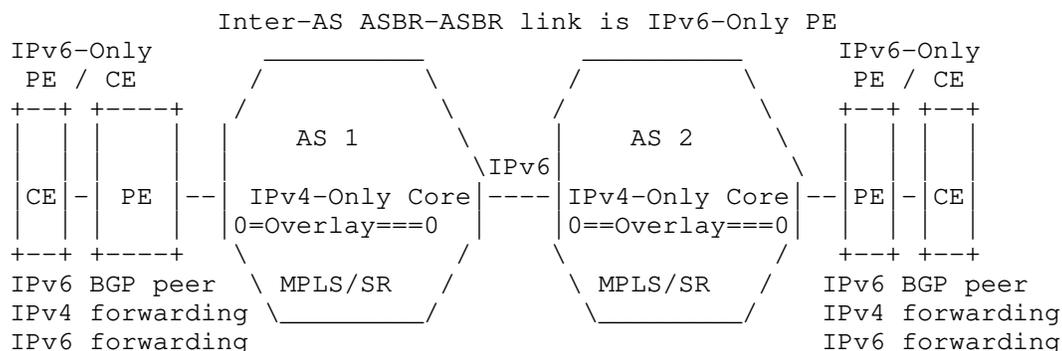


Figure 14: Test-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C

7.9. Test-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-B

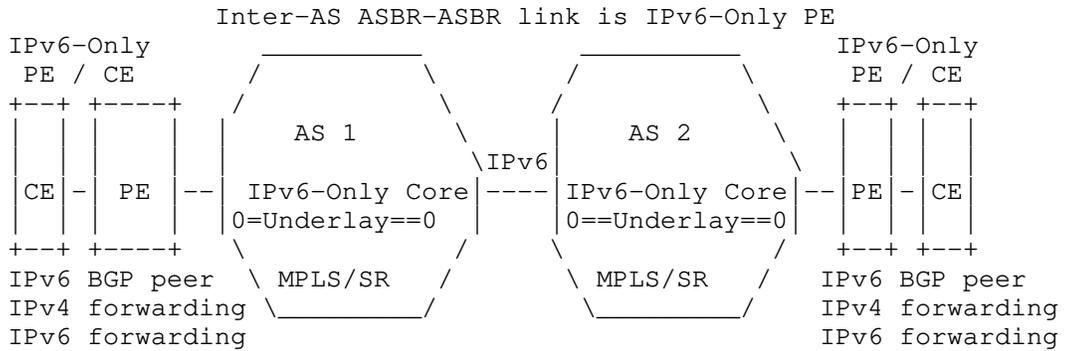


Figure 15: Test-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter- AS Option-B

7.10. Test-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-C

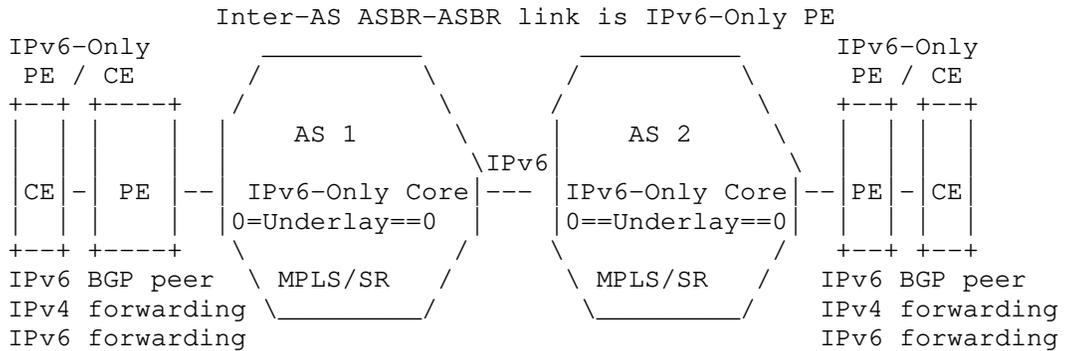


Figure 16: Test-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C

7.11. Test-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-B

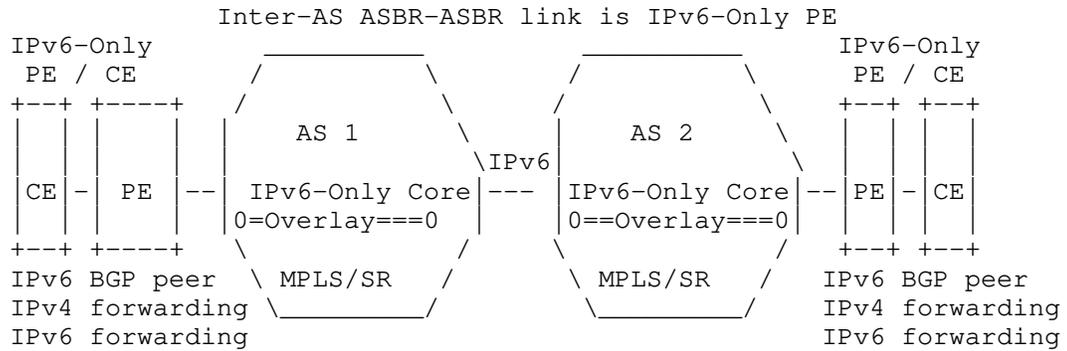


Figure 17: Test-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B

7.12. Test-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-C

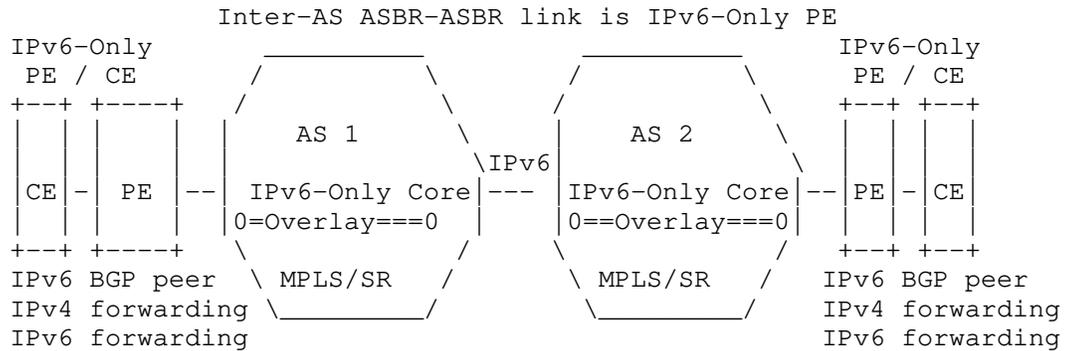


Figure 18: Test-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C

7.13. IPv4-Only PE-CE Operational Considerations Testing

Ping CE to PE when destination prefix is withdrawn  
Traceroute CE to PE and test all ICMPv4 and ICMPv6 type codes

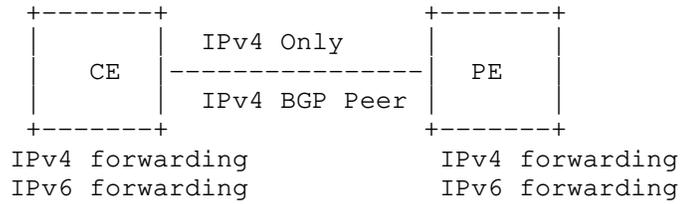


Figure 19: Ping and Trace Test Case

8. IPv4-Only PE Design Edge E2E Design for ALL AFI/SAFI

Listed below are the following IPv4-Only PE Design ALL SAFI design scenario's:

<AFI/SAFI> IPv4 Unicast <1/1>, IPv6 Unicast <2/1>, VPN-IPV4 <1/128>, VPN-IPV6 <2/128>, Multicasat VPN <1/129>, Multicasat VPN <2/129>,BGP-LU IPV4 (GRT) <1/4>

8.1. IPv4-Only PE Design All SAFI Case-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 softwire

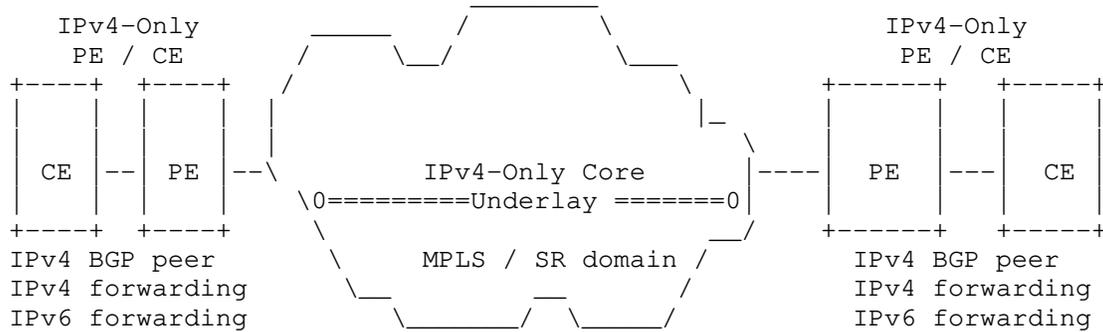


Figure 20: Design Solution-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE)

8.2. IPv4-Only PE Design All SAFI Case-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 Software

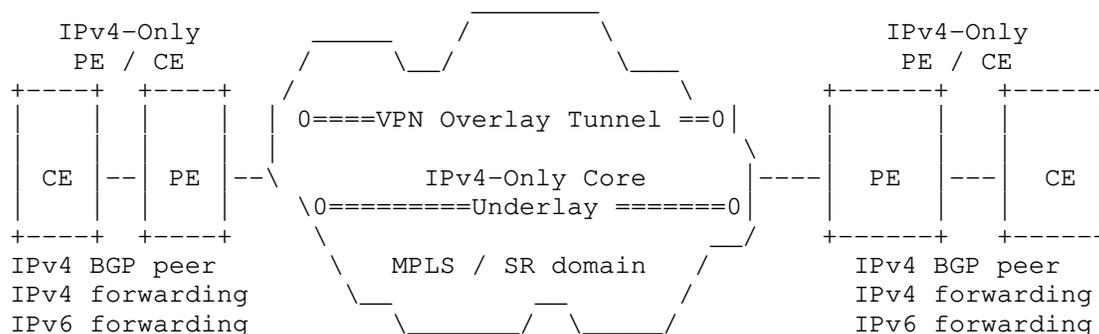


Figure 21: Design Solution-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core

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

8.3. IPv4-Only PE Design All SAFI Case-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE), 4to6 Software

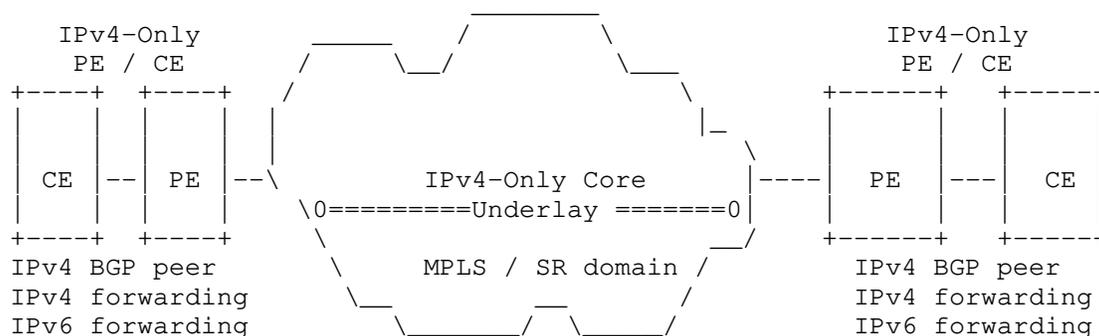


Figure 22: Design Solution-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE)

8.4. IPv4-Only PE Design All SAFI Case-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 Software

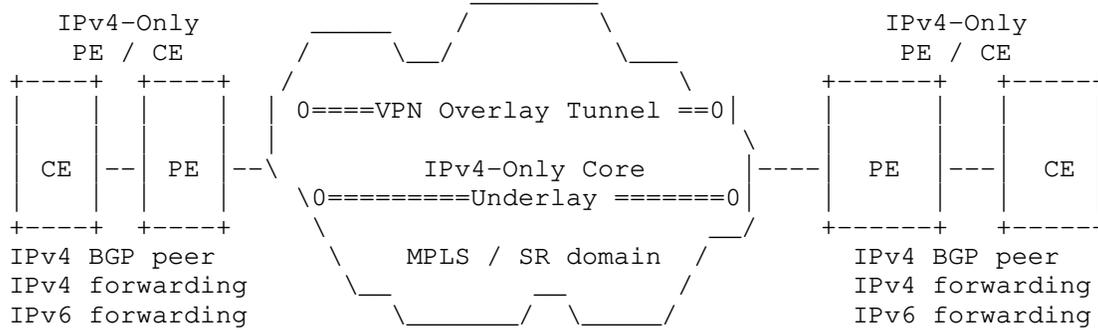


Figure 23: Design Solution-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core

8.5. IPv4-Only PE Design All SAFI Case-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-B

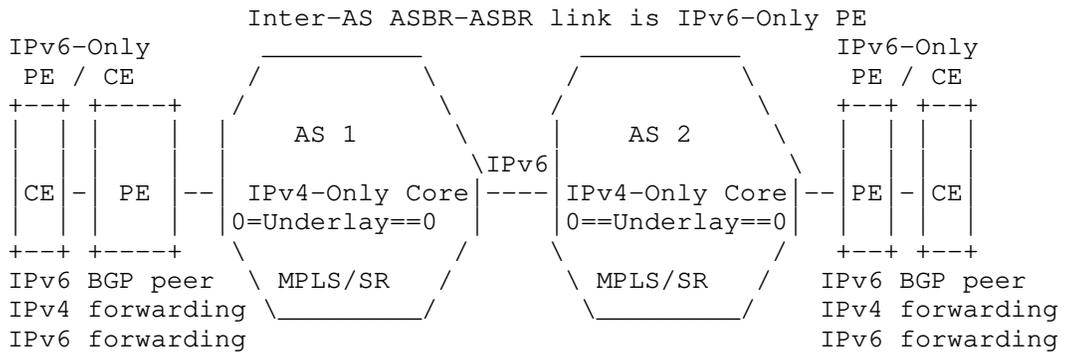


Figure 24: Design Solution-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-B

8.6. IPv4-Only PE Design All SAFI Case-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE), 6to4 software -Inter-AS Option-C

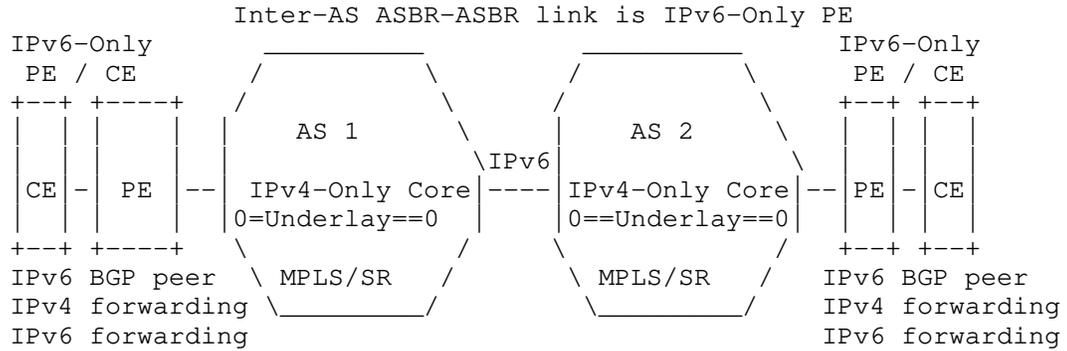


Figure 25: Design Solution-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-C

8.7. IPv4-Only PE Design All SAFI Case-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only, 6to4 software -Inter-AS Option-B

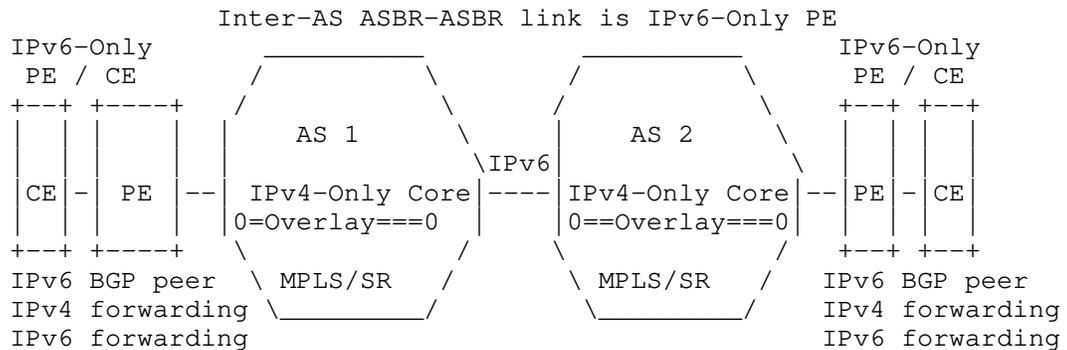


Figure 26: Design Solution-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-B

8.8. IPv4-Only PE Design All SAFI Case-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core, 6to4 software -Inter-AS Option-C

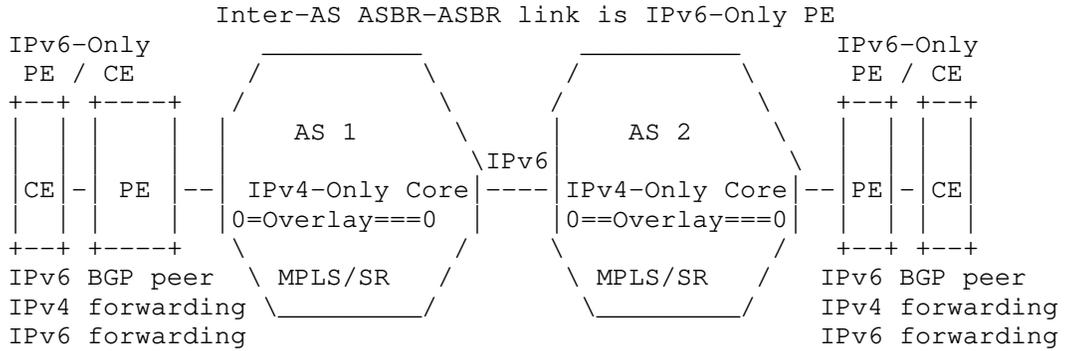


Figure 27: Design Solution-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C

8.9. IPv4-Only PE Design All SAFI Case-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-B

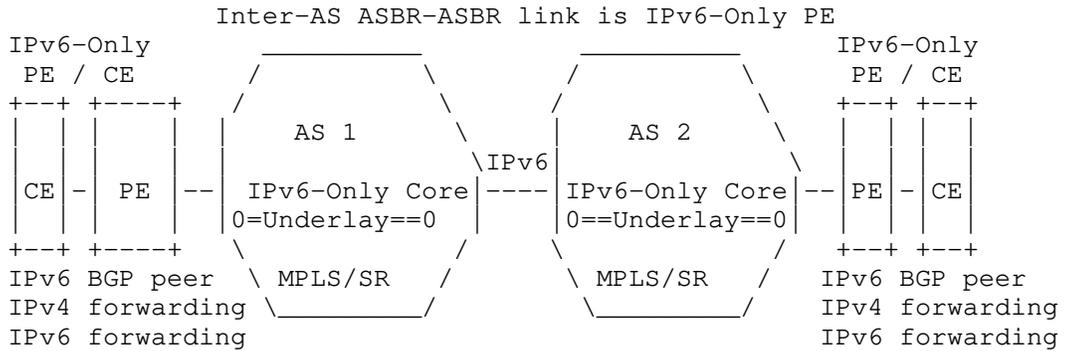


Figure 28: Design Solution-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B

8.10. IPv4-Only PE Design All SAFI Case-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core, 4to6 software -Inter-AS Option-C

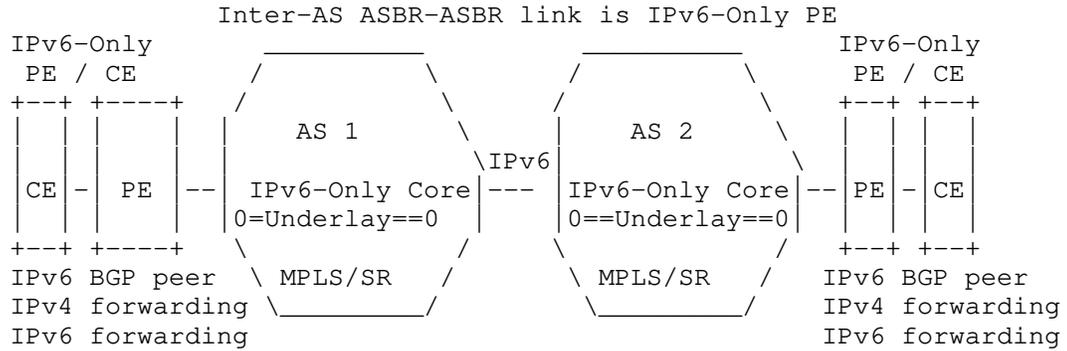


Figure 29: Design Solution-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C

8.11. IPv4-Only PE Design All SAFI Case-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-B

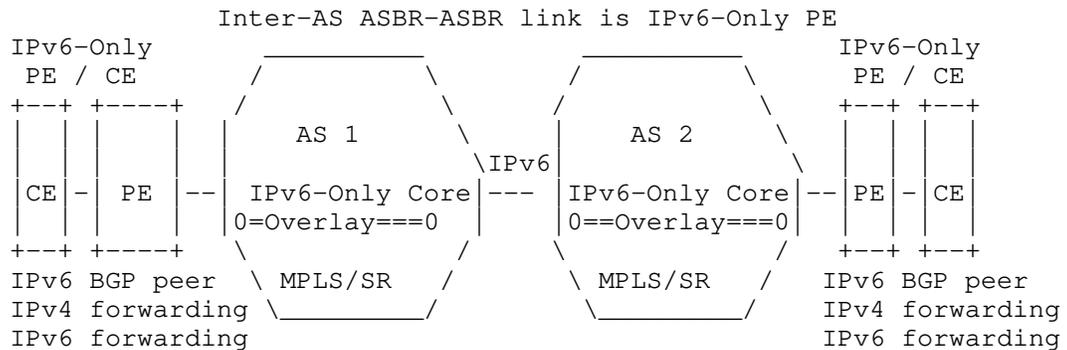


Figure 30: Design Solution-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B

8.12. IPv4-Only PE Design All SAFI Case-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core, 4to6 software -Inter-AS Option-C

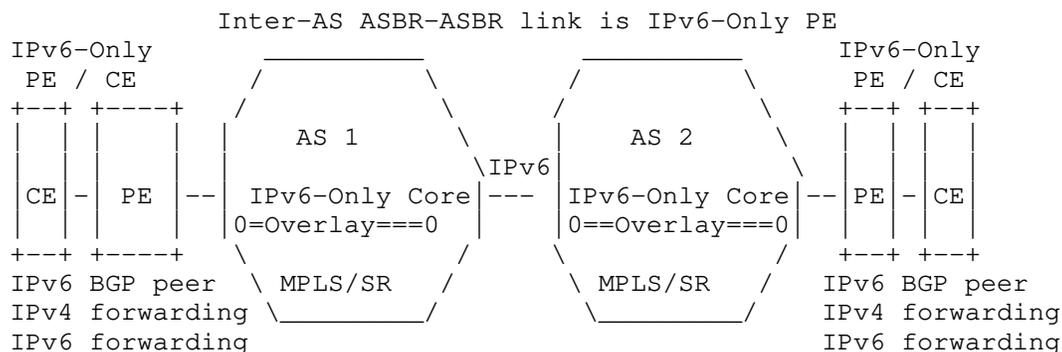


Figure 31: Design Solution-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C

9. IPv6-Only PE Design ALL AFI/SFI Operational Considerations

With a single IPv6 Peer carrying both IPv4 and IPv6 NLRI there are some operational considerations in terms of what changes and what does not change.

What does not change with a single IPv6 transport peer carrying IPv4 NLRI and IPv6 NLRI below:

Routing Policy configuration is still separate for IPv4 and IPv6 configured by capability as previously.

Layer 1, Layer 2 issues such as one-way fiber or fiber cut will impact both IPv4 and IPv6 as previously.

If the interface is in the Admin Down state, the IPv6 peer would go down, and IPv4 NLRI and IPv6 NLRI would be withdrawn as previously.

Changes resulting from a single IPv6 transport peer carrying IPv4 NLRI and IPv6 NLRI below:

Physical interface is no longer dual stacked.

Any change in IPv6 address or DAD state will impact both IPv4 and IPv6 NLRI exchange.

Single BFD session for both IPv4 and IPv6 NLRI fate sharing as the session is now tied to the transport, which now is only IPv6 address family.

Both IPv4 and IPv6 peer now exists under the IPv6 address family configuration.

Fate sharing of IPv4 and IPv6 address family from a logical perspective now carried over a single physical IPv6 peer.

From an operations perspective, prior to elimination of IPv4 peers, an audit is recommended to identify and IPv4 and IPv6 peering incongruencies that may exist and to rectify them. No operational impacts or issues are expected with this change.

With MPLS VPN overlay, per-CE next-hop label allcoation mode where both IPv4 and IPv6 prefixes have the same label in no table lookup pop-n-forward mode should be taken into consideration.

#### 10. IANA Considerations

There are not any IANA considerations.

#### 11. Security Considerations

The extensions defined in this document allow BGP to propagate reachability information about IPv4 prefixes over an MPLS or SR IPv6-Only core network. As such, no new security issues are raised beyond those that already exist in BGP-4 and the use of MP-BGP for IPv6. Both IPv4 and IPv6 peers exist under the IPv6 address family configuration. The security features of BGP and corresponding security policy defined in the ISP domain are applicable. For the inter-AS distribution of IPv6 routes according to case (a) of Section 4 of this document, no new security issues are raised beyond those that already exist in the use of eBGP for IPv6 [RFC2545].

#### 12. Acknowledgments

Thanks to Kaliraj Vairavakkalai, Linda Dunbar, Aijun Wang, Eduardfor Vasilenko, Joel Harlpern, Michael McBride, Ketan Talaulikar for review comments.

#### 13. Contributors

The following people contributed substantive text to this document:

Mohana Sundari  
EMail: mohanas@juniper.net

## 14. References

### 14.1. Normative References

[I-D.ietf-bess-bgp-multicast]

Zhang, Z. J., Giuliano, L., Patel, K., Wijnands, I., Mishra, M. P., and A. Gulko, "BGP Based Multicast", Work in Progress, Internet-Draft, draft-ietf-bess-bgp-multicast-05, 27 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-bgp-multicast-05>>.

[I-D.ietf-idr-bgp-car]

Rao, D., Agrawal, S., and Co-authors, "BGP Color-Aware Routing (CAR)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-car-02, 6 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-02>>.

[I-D.ietf-idr-flowspec-nvo3]

Eastlake, D. E., Weiguo, H., Zhuang, S., Li, Z., and R. Gu, "BGP Dissemination of Flow Specification Rules for Tunneled Traffic", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-nvo3-18, 5 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-nvo3-18>>.

[I-D.ietf-idr-rpd]

Li, Z., Ou, L., Luo, Y., Mishra, G. S., Chen, H., and H. Wang, "BGP Extensions for Routing Policy Distribution (RPD)", Work in Progress, Internet-Draft, draft-ietf-idr-rpd-17, 30 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-rpd-17>>.

[I-D.ietf-idr-sdwan-edge-discovery]

Dunbar, L., Hares, S., Raszuk, R., Majumdar, K., Mishra, G. S., and V. Kasiviswanathan, "BGP UPDATE for SD-WAN Edge Discovery", Work in Progress, Internet-Draft, draft-ietf-idr-sdwan-edge-discovery-10, 23 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sdwan-edge-discovery-10>>.

- [I-D.ietf-idr-segment-routing-te-policy]  
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-21, 23 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-21>>.
- [I-D.ietf-l3vpn-bgpvpn-auto]  
Ould-Brahim, H. H., "Using BGP as an Auto-Discovery Mechanism for VR-based Layer-3 VPNs", Work in Progress, Internet-Draft, draft-ietf-l3vpn-bgpvpn-auto-09, 25 April 2007, <<https://datatracker.ietf.org/doc/html/draft-ietf-l3vpn-bgpvpn-auto-09>>.
- [I-D.ietf-lsvr-bgp-spf]  
Patel, K., Lindem, A., Zandi, S., and W. Henderickx, "BGP Link-State Shortest Path First (SPF) Routing", Work in Progress, Internet-Draft, draft-ietf-lsvr-bgp-spf-26, 19 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsvr-bgp-spf-26>>.
- [I-D.mishra-bess-ipv4-only-pe-design]  
Mishra, G. S., Tantsura, J., Mishra, M. P., Madhavi, S., Yang, Q., Simpson, A., and S. Chen, "IPv4-Only PE Design for IPv6-NLRI with IPv4-NH", Work in Progress, Internet-Draft, draft-mishra-bess-ipv4-only-pe-design-02, 7 July 2022, <<https://datatracker.ietf.org/doc/html/draft-mishra-bess-ipv4-only-pe-design-02>>.
- [I-D.mpmz-bess-mup-safi]  
Murakami, T., Patel, K., Matsushima, S., Zhang, Z. J., Agrawal, S., and D. Voyer, "BGP Extensions for the Mobile User Plane (MUP) SAFI", Work in Progress, Internet-Draft, draft-mpmz-bess-mup-safi-02, 13 March 2023, <<https://datatracker.ietf.org/doc/html/draft-mpmz-bess-mup-safi-02>>.
- [I-D.nalawade-kapoor-tunnel-safi]  
Nalawade, G., "BGP Tunnel SAFI", Work in Progress, Internet-Draft, draft-nalawade-kapoor-tunnel-safi-05, 29 June 2006, <<https://datatracker.ietf.org/doc/html/draft-nalawade-kapoor-tunnel-safi-05>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, DOI 10.17487/RFC2545, March 1999, <<https://www.rfc-editor.org/info/rfc2545>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, DOI 10.17487/RFC5195, June 2008, <<https://www.rfc-editor.org/info/rfc5195>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/info/rfc5492>>.
- [RFC5747] Wu, J., Cui, Y., Li, X., Xu, M., and C. Metz, "4over6 Transit Solution Using IP Encapsulation and MP-BGP Extensions", RFC 5747, DOI 10.17487/RFC5747, March 2010, <<https://www.rfc-editor.org/info/rfc5747>>.
- [RFC6037] Rosen, E., Ed., Cai, Y., Ed., and IJ. Wijnands, "Cisco Systems' Solution for Multicast in BGP/MPLS IP VPNs", RFC 6037, DOI 10.17487/RFC6037, October 2010, <<https://www.rfc-editor.org/info/rfc6037>>.
- [RFC7117] Aggarwal, R., Ed., Kamite, Y., Fang, L., Rekhter, Y., and C. Kodeboniya, "Multicast in Virtual Private LAN Service (VPLS)", RFC 7117, DOI 10.17487/RFC7117, February 2014, <<https://www.rfc-editor.org/info/rfc7117>>.

- [RFC7267] Martini, L., Ed., Bocci, M., Ed., and F. Balus, Ed., "Dynamic Placement of Multi-Segment Pseudowires", RFC 7267, DOI 10.17487/RFC7267, June 2014, <<https://www.rfc-editor.org/info/rfc7267>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9015] Farrel, A., Drake, J., Rosen, E., Uttaro, J., and L. Jalil, "BGP Control Plane for the Network Service Header in Service Function Chaining", RFC 9015, DOI 10.17487/RFC9015, June 2021, <<https://www.rfc-editor.org/info/rfc9015>>.

#### 14.2. Informative References

- [I-D.ietf-idr-dynamic-cap]  
Chen, E. and S. R. Sangli, "Dynamic Capability for BGP-4", Work in Progress, Internet-Draft, draft-ietf-idr-dynamic-cap-16, 21 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-dynamic-cap-16>>.

- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.
- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, DOI 10.17487/RFC4798, February 2007, <<https://www.rfc-editor.org/info/rfc4798>>.
- [RFC4925] Li, X., Ed., Dawkins, S., Ed., Ward, D., Ed., and A. Durand, Ed., "Softwire Problem Statement", RFC 4925, DOI 10.17487/RFC4925, July 2007, <<https://www.rfc-editor.org/info/rfc4925>>.
- [RFC5549] Le Faucheur, F. and E. Rosen, "Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop", RFC 5549, DOI 10.17487/RFC5549, May 2009, <<https://www.rfc-editor.org/info/rfc5549>>.
- [RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", RFC 5565, DOI 10.17487/RFC5565, June 2009, <<https://www.rfc-editor.org/info/rfc5565>>.
- [RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, DOI 10.17487/RFC6074, January 2011, <<https://www.rfc-editor.org/info/rfc6074>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.

- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8950] Litkowski, S., Agrawal, S., Ananthamurthy, K., and K. Patel, "Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop", RFC 8950, DOI 10.17487/RFC8950, November 2020, <<https://www.rfc-editor.org/info/rfc8950>>.

Appendix A. SAFI LIST FOR PE-CE Edge Scenario

IPv4-Only PE Design listing of PE-CE Edge ALL applicable SAFI. Here we showing the catagorization grouping by columns of the SAFI into 2 use case categories.

\* Unicst

\* Multicast

SAFI Value	Description	Unicast	Multicast	Reference
1	Unicast	Yes	No	[RFC4760]
2	Multicast	No	Yes	[RFC2545]
78	MCAST-TREE SAFI	No	Yes	[RFC2545]

Table 1: IPv4-Only PE DESIGN ALL SAFI's PE-CE Edge Scenario

IPv4-Only PE design supports 25 / 32 IANA SAFI's of which the 7 not supported do not use AFI=1 IPv4 or AFI=2 IPv6.

Appendix B. SAFI LIST FOR Inter-AS PE-PE Scenario

IPv4-Only PE Design listing of Inter-AS PE-PE ALL applicable SAFI. Here we show the catagorization grouping by columns of the SAFI into 5 use case categories.

\* Unicst

\* Multicast

\* L1-L2 VPN

\* Tunnel

\* BGP Policy

SAFI Value	Description	Unicast	Multicast	L1-L2 VPN	Tunnel	BGP Policy	Reference
1	NLRI Unicast	Yes	No	No	No	No	[RFC4760]
2	NLRI Multicast	No	Yes	No	No	No	[RFC4760]
4	NLRI MPLS Labels	Yes	No	No	No	No	[RFC8277]
5	MCAST-VPN	No	Yes	No	No	No	[RFC6514]
6	Dynamic Multi Segment PW	No	No	Yes	No	No	[RFC7267]
8	MCAST-VPLS	No	Yes	No	No	No	[RFC7117]
9	BGP-SFC	No	No	No	No	No	[RFC9015]
64	Tunnel SAFI [I-D.nalawade-kapoor-tunnel-safi]	No	No	No	Yes	No	[I-D.nalawade-kapoor-tunnel-safi]
65	VPLS	No	No	No	No	No	[RFC4761]
66	BGP MDT SAFI	NO	Yes	No	No	No	[RFC6037]

67	BGP 4to6 SAFI	No	No	No	Yes	No	[RFC5747]
+-----+-----+-----+-----+-----+-----+-----+-----							
68	BGP 6to4 SAFI	No	No	No	Yes	No	[RFC5747]
+-----+-----+-----+-----+-----+-----+-----+-----							
69	L1 VPN Auto	No	No	Yes	No	No	[RFC5195]
	Discovery						
+-----+-----+-----+-----+-----+-----+-----+-----							
70	BGP EVPN	No	No	No	No	No	[RFC7432]
+-----+-----+-----+-----+-----+-----+-----+-----							
71	BGP-LS	No	No	No	No	No	[RFC7752]
+-----+-----+-----+-----+-----+-----+-----+-----							
72	BGP-LS-VPN	No	No	No	No	No	[RFC7752]
+-----+-----+-----+-----+-----+-----+-----+-----							
73	SR-TE Policy	No	No	No	No	Yes	[I-D.ietf-idr-segme
nt-routing-te-policy]	SAFI						
+-----+-----+-----+-----+-----+-----+-----+-----							
74	SD-WAN	No	No	No	No	Yes	[I-D.ietf-idr-sdwan
-edge-discovery]							

	Capabilities						
75	Routing	No	No	No	No	No	[I-D.ietf-idr-rpd]
	Policy SAFI						
77	Tunneled	No	No	No	Yes	No	[I-D.ietf-idr-flowspec-nvo3]
	Traffic						
	Flowspec						
78	MCAST-TREE	No	Yes	No	No	No	[I-D.ietf-bess-bgp-multicast]
	SAFI						
79	BGP-DPS	No	No	No	No	Yes	Arista Dynamic Path Selection
80	BGP-LS-SPF	No	No	No	No	No	[I-D.ietf-lsvr-bgp-spf]
83	BGP CAR	No	No	No	No	Yes	[I-D.ietf-idr-bgp-car]
84	BGP CAR VPN	No	No	No	No	Yes	[I-D.ietf-idr-bgp-car]
85	BGP MUP SAFI	No	No	No	No	Yes	[I-D.mpmz-bess-mup-safi]
128	MPLS VPN	Yes	No	No	No	No	[RFC4364]
129	MPLS	No	Yes	No	No	No	[RFC6513]
	Multicast VPN						
132	Route Target	Yes	No	No	No	No	[RFC4684]
	Constrains						

133	Dissemination	No	No	No	No	Yes	[RFC8955]
	of Flowspec						
	Rules						
134	L3VPN	No	No	No	No	Yes	[RFC8955]
	Dissemination						
	of Flowspec						
	Rules						
140	VPN Auto	No	No	No	No	Yes	[I-D.ietf-l3vpn-bgp
vpn-auto]	Discovery						

Table 2: IPv4-Only PE DESIGN ALL SAFI's Inter-AS PE-PE Scenario

IPv4-Only PE design supports 25 / 32 IANA SAFI's of which the 7 not supported do not use AFI=1 IPv4 or AFI=2 IPv6.

## Authors' Addresses

Gyan Mishra  
Verizon Inc.  
Email: gyan.s.mishra@verizon.com

Jeff Tantsura  
Microsoft, Inc.  
Email: jefftant.ietf@gmail.com

Mankamana Mishra  
Cisco Systems  
821 Alder Drive,  
MILPITAS  
Email: mankamis@cisco.com

Sudha Madhavi  
Juniper Networks, Inc.  
Email: smadhavi@juniper.net

Qing Yang  
Arista Networks  
Email: qyang@arista.com

Adam Simpson  
Nokia  
Email: adam.1.simpson@nokia.com

Shuanglong Chen  
Huawei Technologies  
Email: chenshuanglong@huawei.com