

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

G. Mishra
Verizon Inc.
S. Madhavi
Juniper Networks, Inc.
A. Simpson
Nokia
M. Mishra
Cisco Systems
J. Tantsura
Microsoft, Inc.
S. Chen
Huawei Technologies
20 October 2025

IPv4-Only and IPv6-Only PE Design DESIGN ALL SAFI
draft-ietf-bess-v4-v6-pe-all-safi-03

Abstract

As Enterprises and Service Providers 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) Inter-AS peering network from IPv4 to IPv6. Operators must have flexibility to continue to support IPv4 customers when both the Core and Edge networks migrate to IPv6. As well as must be able to support IPv6 networks in cases where operators decide to remain on an IPv4 network or during transition.

This document details the External BGP (eBGP) PE-PE Inter-AS and PE-CE Edge peering IPv4-Only PE design where both IPv4 and IPv6 all supported SAFI NLRI can be advertised over a single IPv4 peer and IPv6-Only PE Design where all supported SAFI NLRI can be advertised over a single IPv6 peer.

This document also defines a new IPv4 BGP next hop encoding standard that uses an IPv4 address as the next hop and not an IPv4 mapped IPv6 address.

This document also provides vendor specific test cases for the IPv4-Only peering design and IPv6-Only PE design as well as test results for the four major vendors stakeholders in the routing and switching industry, Cisco, Juniper, Nokia and Huawei. With the test results provided for the IPv6-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 the design.

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 23 April 2026.

Copyright Notice

Copyright (c) 2025 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	8
3. Terminology	8
4. Problem Statement	9
5. IPv6-Only PE-CE Design ALL SAFI Solution	10
6. IPv4-Only PE-CE Design ALL SAFI Solution	12
7. IPv6-Only Edge Peering Design ALL SAFI	15
7.1. IPv6-Only Edge Peering Packet Walk ALL SAFI	15
7.2. IPv6-Only PE Design ALL SAFI over IPv4-Only Core Packet Walk	16
7.3. IPv6-Only PE Design ALL SAFI over IPv6-Only Core Packet Walk	17
8. IPv4-Only Edge Peering Design	20
8.1. IPv4-Only Edge Peering Packet Walk ALL SAFI	20

8.2.	IPv4-Only PE Design ALL SAFI over IPv4-Only Core Packet Walk	20
8.3.	IPv4-Only PE Design ALL SAFI over IPv6-Only Core Packet Walk	22
9.	IPv4 Next Hop Encoding	25
10.	IPv6-Only PE Design / IPv6-Only PE Design Test Results . . .	26
11.	IPv6-Only PE Design Testing Details	28
11.1.	Test-1 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE)	28
11.2.	Test-2 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only Core	29
11.3.	Test-3 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core (4PE)	29
11.4.	Test-4 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core	30
11.5.	Test-5 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-B	30
11.6.	Test-6 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-C	31
11.7.	Test-7 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only - Inter-AS Option-B	31
11.8.	Test-8 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C	32
11.9.	Test-9 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B	32
11.10.	Test-10 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C	33
11.11.	Test-11 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B	33
11.12.	Test-12 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C	34
11.13.	IPv6-Only PE-CE Operational Considerations Testing . . .	34
12.	IPv4-Only PE Design Testing Details	35
12.1.	Test-1 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE)	35
12.2.	Test-2 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only Core	36
12.3.	Test-3 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE)	36
12.4.	Test-4 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core	37
12.5.	Test-5 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-B	37

12.6.	Test-6 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-C	37
12.7.	Test-7 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only - Inter-AS Option-B	38
12.8.	Test-8 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C	38
12.9.	Test-9 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B	39
12.10.	Test-10 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C	39
12.11.	Test-11 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B	40
12.12.	Test-12 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C	40
12.13.	IPv4-Only PE-CE Operational Considerations Testing	41
13.	SRv6 Applicability to IPv6 Only PE Design	41
14.	Typed and Untyped SAFI	43
15.	IPv4-Only PE Design and IPv6-Only PE Design ALL AFI/SFI Operational Considerations	43
16.	IANA Considerations	44
17.	Security Considerations	45
18.	Acknowledgments	45
19.	Contributors	45
20.	References	45
20.1.	Normative References	45
20.2.	Informative References	49
Appendix A.	SAFI LIST FOR PE-CE Edge Scenario	51
Appendix B.	SAFI LIST FOR Inter-AS PE-PE Scenario	51
Authors'	Addresses	54

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.

IXP are also facing IPv4 address depletion at their peering points, which are large Layer 2 transit backbones that service providers peer and exchange IPv4 and IPv6 Network Layer Reachability Information (NLRI). Today, these transit exchange points are Dual Stacked. With this IPv6-only BGP peering design, only IPv6 MUST be configured on the PE-PE inter-as peering interface, the Inter-AS Provider Edge (PE) - Provider Edge (PE), the IPv6 BGP peer is now used to carry IPv4 (Network Layer Reachability Information) NLRI over an IPv6 next hop using IPv6 next hop encoding defined in [RFC8950], while continuing

to forward both IPv4 and IPv6 packets. With this IPv6-Only PE Design, ASBRs providing Inter-AS options peering PE to PE extending L3 VPN services is now no longer Dual Stacked and as well can support ALL AFI/SAFI.

This document also provides a 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.

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 new extensions defined in [RFC8950] supporting NLRI and next hop address family mismatch, the 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 IPv6 TCP session. This allows for the elimination of dual stack from the PE-PE Inter-AS peering point, and now enable the Inter-AS peering to be IPv6-ONLY. The elimination of IPv4 Inter Provider ASBR tie point, PE-PE Inter-AS peering points translates into OPEX expenditure savings of point-to-point infrastructure links as well as /31 address space savings and administration and network management of both IPv4

and IPv6 BGP peers. This reduction decreases the number of PE-PE Inter-AS options BGP peers by fifty percent, which is a tremendous cost savings for operators.

This document details an important External BGP (eBGP) PE-CE Edge and PE-PE Inter-AS IPv4-Only PE Design and IPv6-Only PE Peering Design that leverages the MP-BGP capability exchange by using single IPv4 peering or IPv6 peering as pure transport, allowing all 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 supported Subsequent Address Family Identifiers (SAFI).

The design change provides two new alternatives to traditional Dual Stacking implemented today while providing the same Dual Stacking functionality and capabilities that exist today with separate IPv4 and IPv6 BGP sessions, but now with this paradigm shift now only requires a single IPv4 transport peer "IPv4-Only PE Design" or single IPv6 transport peer "IPv6-Only PE Design".

IPv6-Only PE Design entails that an IPv4 address MUST not be configured on the PE-CE or PE-PE interface and with the IPv6-Only PE Design an IPv4 address must not be configured on the PE-CE or PE-PE interface.

IPv4-Only PE Design entails that an IPv6 address MUST not be configured on the PE-CE or PE-PE interface and with the IPv4-Only PE Design an IPv6 address must not be configured on the PE-CE or PE-PE interface.

From a control plane perspective with the IPv6-Only PE design a single IPv6-Only peer is required for both IPv4 and IPv6 routing updates and from a data plane forwarding perspective an IPv6 address need only be configured on the PE-CE Edge or PE-PE Inter-AS peering interface for both IPv4 and IPv6 packet forwarding.

From a control plane perspective with the IPv4-Only PE design 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-CE Edge or PE-PE Inter-AS peering interface for both IPv4 and IPv6 packet forwarding.

This document defines the IPv6-Only PE Design and IPv4-Only PE Design as a new PE-CE Edge and ASBR-ASBR PE-PE Inter-AS BGP peering Standard to support all IPv4 and IPv6 AFI AFI and corresponding 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 paradigm shift to IPv4-Only or IPv6-Only peering architecture is still very relevant to Segment Routing architecture both SR-MPLS and SRv6.

With both the IPv4-Only PE Design and IPv6-Only PE Design, while the savings exists at the Edge eBGP PE-PE Inter-AS peering, on the core side iBGP PE to Route Reflector (RR) peering carrying <AFI/SAFI> IPv4 <1/1>, VPN-IPv4 <1/128>, and Multicast VPN <1/129>, there is no savings as the Provider (P) Core is IPv6 Only or IPv4-Only, thus can only have an IPv6 peer and must use [RFC8950] extended next hop encoding to carrying IPv4 NLRI IPv4 <2/1>, VPN-IPv4 <2/128>, and Multicast VPN <2/129> over an IPv4 or IPv6 next hop.

The IPv4-Only PE Design and IPv6-Only PE ALL SAFI Design supports the following IPv4 and IPv6 AFI and their corresponding SAFI below: <AFI/SAFI>, NLRI Multi-Segment Pseudowires [RFC7267] <1/6>, BGP Tunnel Encapsulation SAFI [RFC9012] <1/7>, 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] <1/134>, VPN Auto-Discovery SAFI draftxx <1/140>

This document provides proof of concept test results for the IPv4-Only PE Design and IPv6-Only PE design for 12 of the most common use cases with 3 of the most commonly used SAFI <AFI/SAFI> IPv4 <1/1>, VPN-IPv4 <1/128>, and Multicast VPN <1/129>, with four major vendors stakeholders in the routing and switching industry, Cisco, Juniper, Nokia and Huawei. With the test results provided for the IPv6-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 IPv6-Only Edge peering. This will give confidence to operators to start the proliferation of the IPv4-Only PE Design and IPv6-Only PE design worldwide. This document provides a detailed analysis of all IPv4 and IPv6 address family and related SAFI that is supported with the IPv4-Only PE Design and IPv6-Only PE Design in Section 4. Thus this draft provides a test use case basis with the three SAFI tested, giving extensibility to all of the other SAFI that are supported for future testing and deployment by operators.

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 as the new standard. Today the IPv4

next hop encoding has mix of 4 octet IPv4 address for the next hop as well as IPv4 mapped IPv6 address. This is discussed in detail in section 10.

The Major benefit from the IPv4-Only PE design is IPv6 address space savings and the ability to support IPv6 NLRI without configuring an IPv6 address at the PE-CE edge and PE-PE inter-as boundary and avoid having to upgrade to support IPv6 as well elimination of provisioning of IPv6 addressing and Network Operations and monitoring costs of traditional Dual Stacked interfaces with IPv4 and IPv6 BGP peering which now translating into CAPEX and OPEX Savings.

The Major benefit from the IPv6-Only PE design is IPv4 address space savings solving IPv4 address depletion issues and the ability to support IPv4 NLRI without configuring an IPv4 address at the PE-CE edge and PE-PE inter-as boundary, as well elimination of provisioning of IPv4 address and Network Operations and monitoring costs of traditional Dual Stacked interfaces with IPv4 and IPv6 BGP peering which now translates into CAPEX and OPEX Savings.

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 now be put off indefinitely. Operators around the world are clamoring for a solution as well that can help solve issues related to IPv4 address depletion at these large IXP peering points.

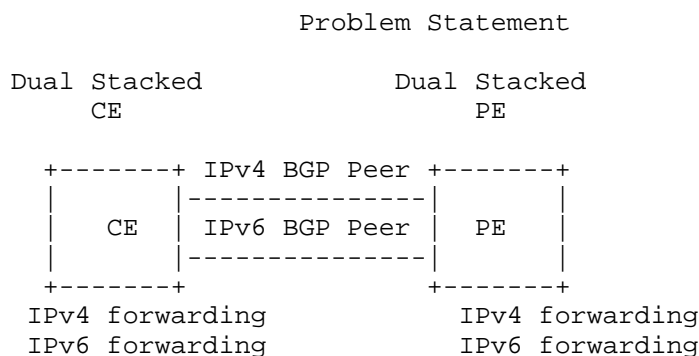


Figure 1: Problem Statement - Dual Stack Peering

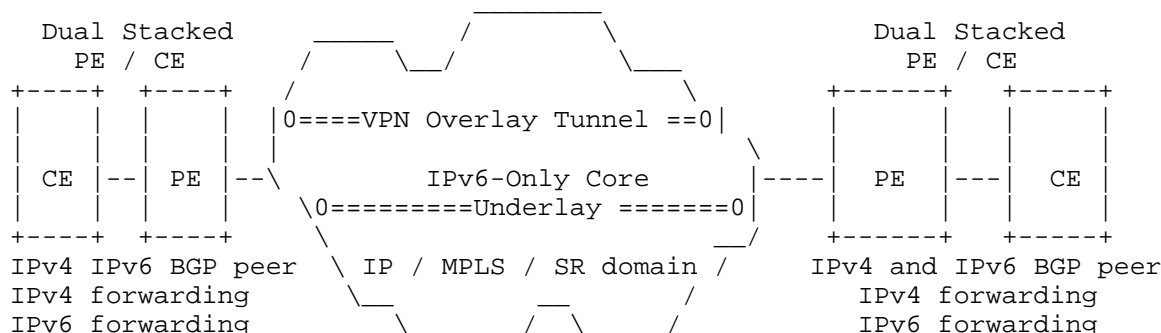


Figure 2: Problem Statement - E2E Dual Stack Edge

5. IPv6-Only PE-CE Design ALL SAFI Solution

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

IPv6-Only PE Design is applicable to infrastructure networks such as Core networks, DC networks, Access networks as well as any PE-CE public or private network can now utilize this IPv6-Only Edge solution and reap the benefits immediately on IPv6 address space saving and CAPEX and OPEX savings.

Six Groupings of AFI/SAFI Use Case Scenario for the IPv6-Only PE Design ALL SAFI

- * Group-1 PE-CE
- * Group-2 PE-PE Inter-AS
- * Group-3 L1 and L2 VPN
- * Group-4 Multicast
- * Group-5 Tunnel
- * Group-6 BGP, BGP Flowspec, BGP Misc Policy

Group-1 Edge Customer IPv4/IPv6 NLRI PE-CE AFI / SAFI grouping (CP-DP) - "UNICAST"

- * AFI/SAFI 1/1 IPv4 Unicast 2/1 IPv6 Unicast

Group-1 Edge Customer IPv4/IPv6 NLRI PE-CE AFI / SAFI grouping (CP-DP) - "MULTICAST"

- * AFI/SAFI 1/2 IPv4 Multicast 2/2 IPv6 Multicast
- * AFI/SAFI 1/78 IPv4 MCAST-TREE SAFI 2/78 MCAST-TREE SAFI

Group-2 ASBR-ASBR Inter-AS Customer IPv4/IPv6 NLRI AFI/SAFI grouping (CP-DP) - "UNICAST"

Global Table

- * AFI/SAFI 1/4 4PE
- 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

Group-2 ASBR-ASBR Inter-AS Customer IPv4/IPv6 NLRI AFI/SAFI grouping (CP-DP) - "MULTICAST"

Global Table

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

L3 VPN

- * 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

Group-3 - 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

Group-4 - 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

Group-5 - BGP, BGP Flowspec, BGP 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 1/85 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: IPv6-Only PE Design Solution Applicability to ALL AFI/SAFI

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

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

IPv4-Only PE Design is applicable to 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 of CAPEX and OPEX savings.

Six Groupings of AFI/SAFI Use Case Scenario for the IPv6-Only PE Design ALL SAFI

- * Group-1 PE-CE
- * Group-2 PE-PE Inter-AS
- * Group-3 L1 and L2 VPN
- * Group-4 Multicast
- * Group-5 Tunnel

* Group-6 BGP, BGP Flowspec, BGP Misc Policy

Group-1 Edge Customer IPv4/IPv6 NLRI PE-CE AFI / SAFI grouping (CP-DP) - "UNICAST"

* AFI/SAFI 1/1 IPv4 Unicast 2/1 IPv6 Unicast

Group-1 Edge Customer IPv4/IPv6 NLRI PE-CE AFI / SAFI grouping (CP-DP) - "MULTICAST"

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

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

Group-2 ASBR-ASBR Inter-AS Customer IPv4/IPv6 NLRI AFI/SAFI grouping (CP-DP) - "UNICAST"

Global Table

* AFI/SAFI 1/4 4PE

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

Group-2 ASBR-ASBR Inter-AS Customer IPv4/IPv6 NLRI AFI/SAFI grouping (CP-DP) - "MULTICAST"

Global Table

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

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

Group-3 - 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

Group-4 - Multicast

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

Group-5 - 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

Group-6 - BGP, BGP Flowspec, BGP 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 4: IPv4-Only PE Design Solution Applicability to ALL AFI/SAFI

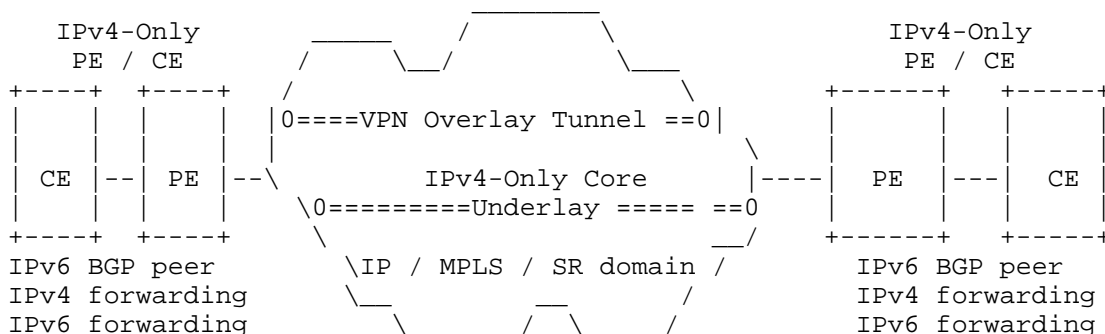


Figure 5: E2E VPN Solutionv4

7. IPv6-Only Edge Peering Design ALL SAFI

7.1. IPv6-Only Edge Peering Packet Walk ALL SAFI

The IPv6-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 IPv6-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 scenarios 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 IPv6-Only P (Provider) core network or IPv4-Only P (Provider) core network. With this IPv6-Only Edge peering design, the Software Mesh Framework is not extended beyond the Provider Edge (PE) and continues to terminate on the PE router.

7.2. IPv6-Only PE Design ALL SAFI over IPv4-Only Core Packet Walk

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

In the scenario where IPv4 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 IPv4 packets that ingress the CE from within the CE AS are over an IPv6-Only interface and are forwarded to an IPv4 NLRI destination prefix learned from the Pure Transport Single IPv6 BGP Peer. In the IPv6-Only Edge peering architecture the PE is IPv6-Only as all PE-CE interfaces are IPv6-Only. However, on the CE, the PE-CE interface is the only interface that is IPv6-Only and all other interfaces may or may not be IPv6-Only. Following the data plane packet flow, IPv4 packets are forwarded from the ingress CE to the IPv6-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 software 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 software 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 6: IPv6-Only PE Design ALL SAFI 6to4 Software - IPv6 Edge over an IPv4-Only Core

7.3. IPv6-Only PE Design ALL SAFI over IPv6-Only Core Packet Walk

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

In the scenario where IPv4 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 IPv4 packets that ingress the CE from within the CE AS are over an IPv6-Only interface and are forwarded to an IPv4 NLRI destination prefix learned from the Pure Transport Single IPv6 BGP Peer. In the IPv6-Only Edge peering architecture the PE is IPv6-Only as all PE-CE interfaces are IPv6-Only. However, on the CE, the PE-CE interface is the only interface that is IPv6-Only and all other interfaces may or may not be IPv6-Only. Following the data plane packet flow, IPv4 packets are forwarded from the ingress CE to the IPv6-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 software IPv6-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 4to6 software IPv6-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 7: IPv6-Only PE Design ALL SAFI 4to6 Software - IPv4 Edge over an IPv6-Only Core

8. IPv4-Only Edge Peering Design

8.1. IPv4-Only Edge Peering Packet Walk ALL SAFI

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 scenario 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.

8.2. IPv4-Only PE Design ALL SAFI over IPv4-Only Core Packet Walk

6to4 software 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 software 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 8: 6to4 Software - IPv6 Edge over an IPv4-Only Core Packet Walk

8.3. IPv4-Only PE Design ALL SAFI over IPv6-Only Core Packet Walk

4to6 software 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 software 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 software 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 9: 4to6 Software - IPv4 Edge over an IPv6-Only Core Packet Walk

9. IPv4 Next Hop Encoding

- * 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 for Unicast SAFI 1, Multicast SAFI 2 and BGP-LU SAFI 4, and 24/48 byte 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 SAFI 2 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 SAFI 2 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 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.

- * Some of the major vendors across platforms even support a variety of different encodings as well in some cases on the same platform the control plane BGP encoding and hardware programming is even different and does not match up.
- * For interoperability if a vendor does not support the new next hop encoding, it would continue to use the IPv4 mapped IPv6 address format until the P2P send / receive neighbors MP-BGP MP_REACH BGP capability exchange is for the new IPv4 Next hop encoding codepoint.

10. IPv6-Only PE Design / IPv6-Only PE Design Test Results

Listed below are the following IPv6-Only PE Design, design scenario's that have been tested with test results related to the two of the Most Common SAFIs used today listed below:

- * <AFI/SAFI> ,
- * IPv4 Unicast <1/1>, IPv6 Unicast <2/1> ,
- * VPN-IPv4 <1/128>, VPN-IPv6 <2/128> ,

IPv6-Only PE Design Proof of concept interoperability testing of the 2 most common SAFI with 12 of the most common design use cases between the 4 vendors Cisco, Juniper, Nokia and Huawei.

Cisco, Juniper, 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, XRd (XR Docker)

Cisco: All testing completed tests 1-12 on XRd 7.10.11 Global table and VRF and IPv4 MPLS core and IPv6 SRv6 core

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

All Testing 1-4 completed for Unicast SAFI 1/1 IPv4-Unicast, 2/1 IPv6-Unicast

All Testing 5-12 completed for IP VPN SAFI 1/128 IPv4-VPN, 2/128 IPv6-VPN

Tested v4 edge over v6 core in a virtual setup using vMX platform and 20.4R2 and LDPv6 as underlay, but there were some data plane forwarding issues. Tested same setup on latest release 21.4 and it worked. Investigating what the minimum version is for this setup to work.

Tested on above Juniper platforms. Completed IPv6-Only PE design functionality test with PE-CE IPv6 peer carrying IPv4 and IPv6 prefixes control plane validation and data plane forwarding plane validation and verified end to end reachability CE to CE forwarding plane with Default Per-CE label allocation mode. Tested with IPv4-Only Core and IPv6-Only Core and proved that the IPv6-Only PE design solution works. Both IPv4 and IPv6 packets were forwarded identical functionality of Dual Stack without having IPv4 address configured.

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

All Testing 1-4 completed for Unicast SAFI 1/1 IPv4-Unicast, 2/1 IPv6-Unicast

All Testing 5-12 completed for IP VPN SAFI 1/128 IPv4-VPN, 2/128 IPv6-VPN

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

Not Completed

Juniper Mx960 24.2R1-S1.10 and Nokia IXR-X3 25.7.R1 interoperability testing completed over IPv4 MPLS core and IPv6 SRv6 Core, tests 1-12 completed successfully. No issues.

IPv4 Only PE Design, IPv6-Only PE Design Test Cases 1-12

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

Test Cases 1-4

* IPv4 Core:

- AFI/SAFI IPv4-Unicast SAFI 1/1 IPv6-Unicast SAFI 2/1
- Test-1 Global table (6PE)
- Test-2 IP VPN

* IPv6 Core:

- AFI/SAFI IPv4-Unicast 1/1 SAFI IPv6-Unicast SAFI 2/1

- Test-3 Global table (4PE)
- Test-4 IP VPN

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

Test Cases 5-12

* IPv4 Core:

- AFI/SAFI IPv4-VPN SAFI 1/128 IPv6-VPN SAFI 2/128
- Test-5 Global table 6PE Option-B
- Test-6 Global table 6PE Option-C
- Test-7 IP VPN Inter AS Option-B
- Test-8 IP VPN Inter AS Option-C

* IPv6 Core:

- AFI/SAFI IPv4-Unicast SAFI IPv6-Unicast SAFI
- 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

11. IPv6-Only PE Design Testing Details

Listed below are the following IPv6-Only PE Design, design scenario's that have been tested with test results related to the 3 of the Most Common SAFIs used today listed below:

<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>

11.1. Test-1 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE)




Figure 10: Test-1 E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core (6PE)

11.2. Test-2 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only Core

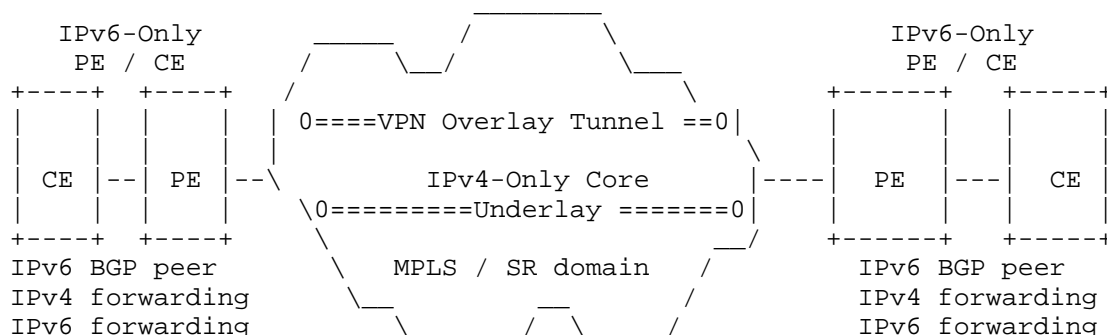


Figure 11: Test-2 E2E IPv6-Only PE-CE, VPN over IPv4-Only Core

11.3. Test-3 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core (4PE)




Figure 12: Test-3 E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core (4PE)

11.4. Test-4 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core

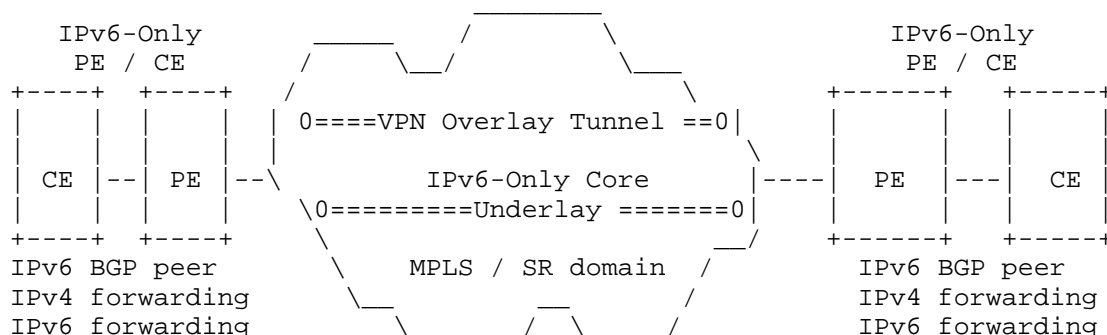


Figure 13: Test-4 E2E IPv6-Only PE-CE, VPN over IPv6-Only Core

11.5. Test-5 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-B

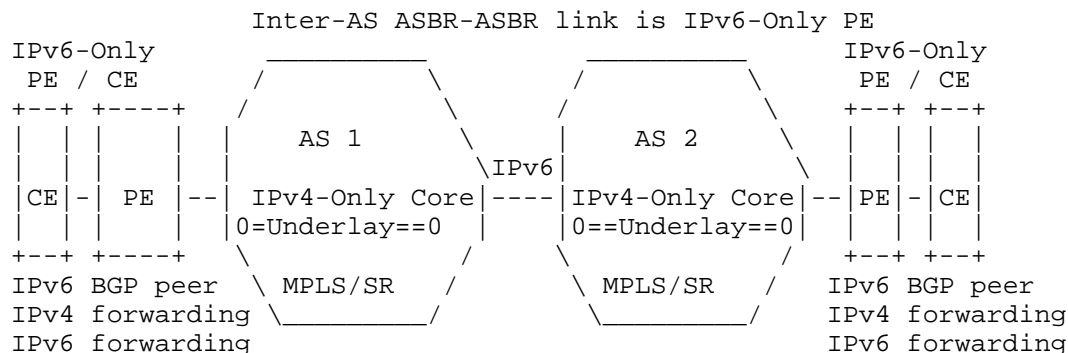


Figure 14: Test-5 E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-B

11.6. Test-6 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-C

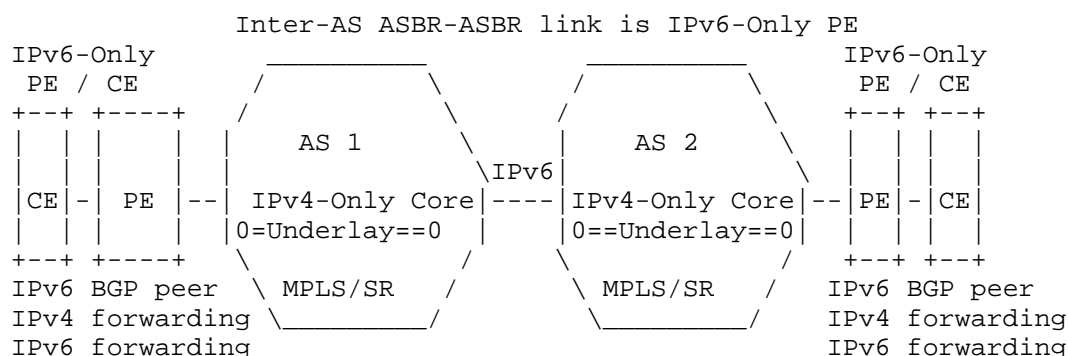


Figure 15: Test-6 E2E IPv6-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-C

11.7. Test-7 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only - Inter-AS Option-B




Figure 16: Test-7 E2E IPv6-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-B

11.8. Test-8 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C




Figure 17: Test-8 E2E IPv6-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C

11.9. Test-9 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B

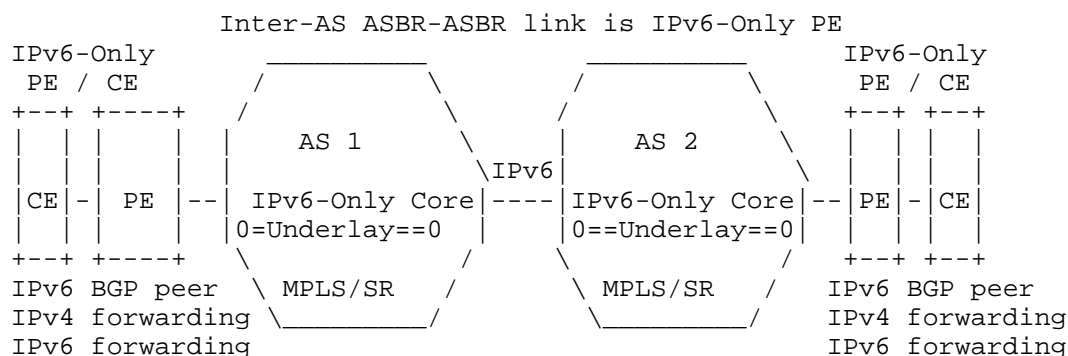


Figure 18: Test-9 E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter- AS Option-B

11.10. Test-10 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C




Figure 19: Test-10 E2E IPv6-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C

11.11. Test-11 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B




Figure 20: Test-11 E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B

11.12. Test-12 IPv6-Only PE Design All SAFI E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C




Figure 21: Test-12 E2E IPv6-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C

11.13. IPv6-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 22: Ping and Trace Test Case IPv6-Only PE Design

12. IPv4-Only PE Design Testing Details

Listed below are the following IPv4-Only PE Design, design scenario's that have been tested with test results related to the 3 of the Most Common SAFIs used today listed below:

<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>

12.1. Test-1 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE)




Figure 23: Test-1 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE)

12.2. Test-2 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only Core

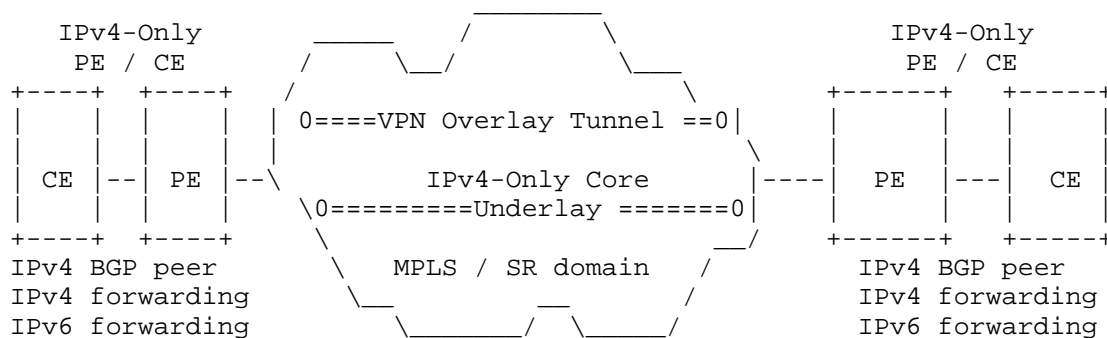


Figure 24: Test-2 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core

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

12.3. Test-3 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE)




Figure 25: Test-3 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core (4PE)

12.4. Test-4 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core

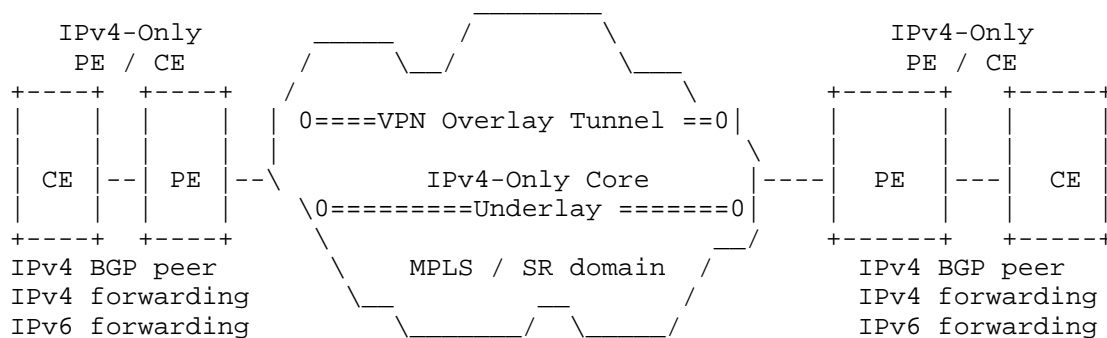


Figure 26: Test-4 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core

12.5. Test-5 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-B

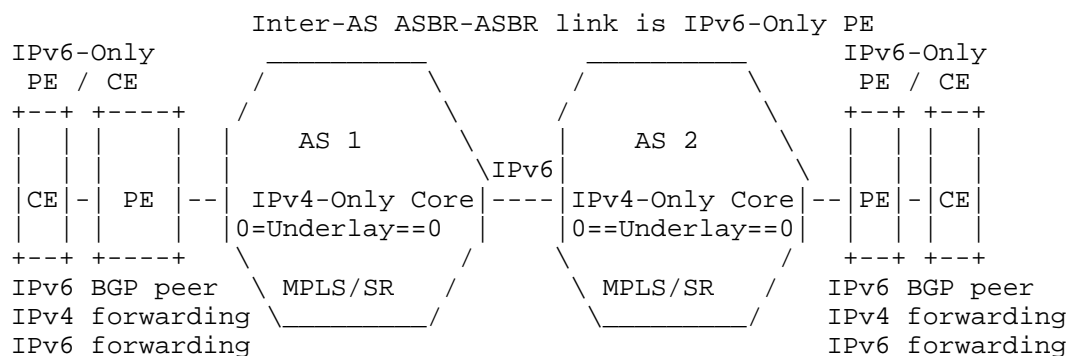


Figure 27: Test-5 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-B

12.6. Test-6 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core(6PE) - Inter-AS Option-C

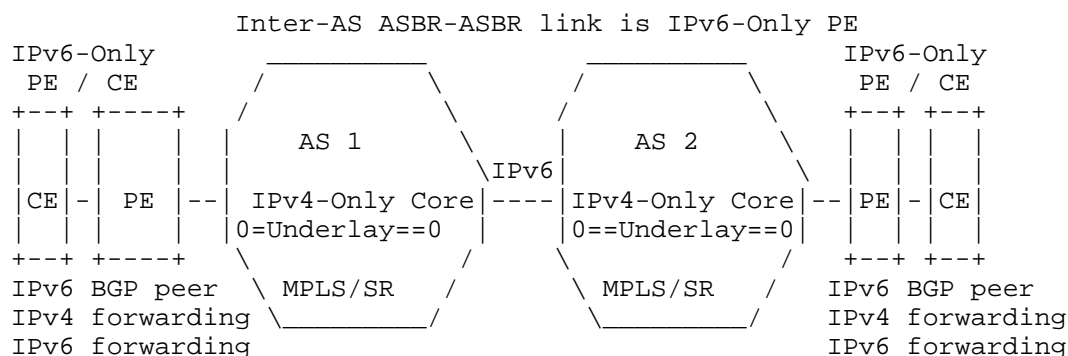


Figure 28: Test-6 E2E IPv4-Only PE-CE, Global Table over IPv4-Only Core (6PE) - Inter-AS Option-C

12.7. Test-7 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only - Inter-AS Option-B




Figure 29: Test-7 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-B

12.8. Test-8 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C

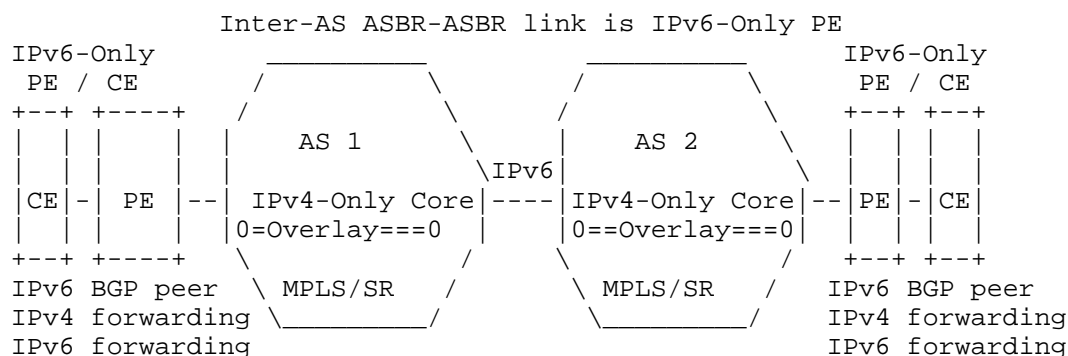


Figure 30: Test-8 E2E IPv4-Only PE-CE, VPN over IPv4-Only Core - Inter-AS Option-C

12.9. Test-9 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B

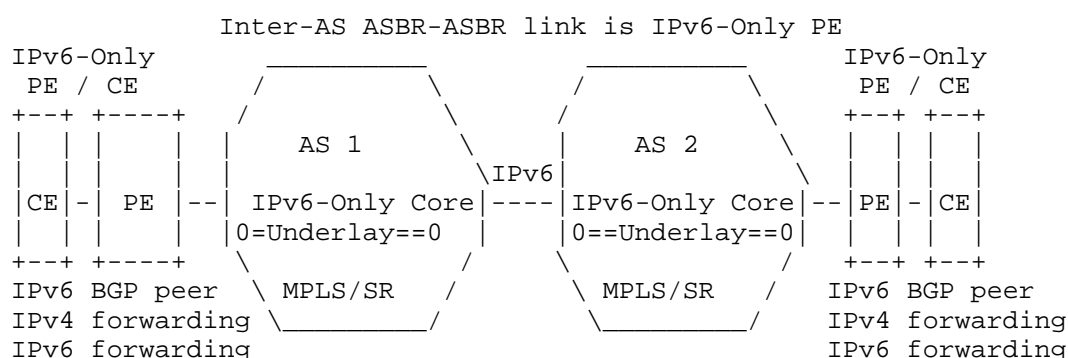


Figure 31: Test-9 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-B

12.10. Test-10 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C

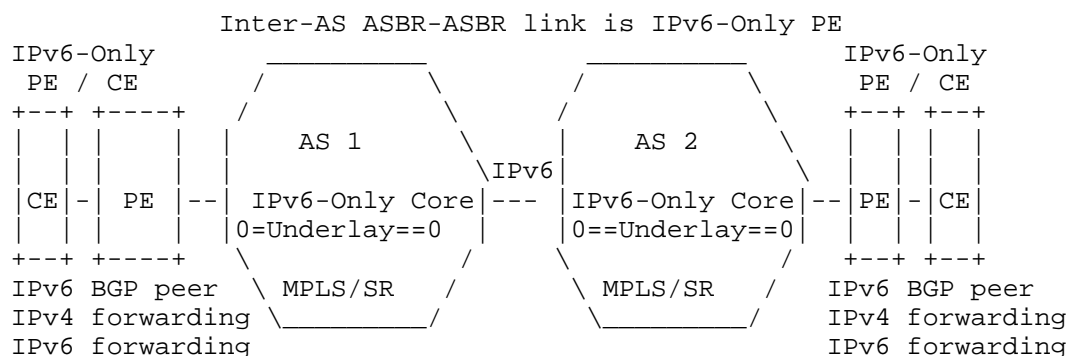


Figure 32: Test-10 E2E IPv4-Only PE-CE, Global Table over IPv6-Only Core - Inter-AS Option-C

12.11. Test-11 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B




Figure 33: Test-11 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-B

12.12. Test-12 IPv4-Only PE Design All SAFI E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C

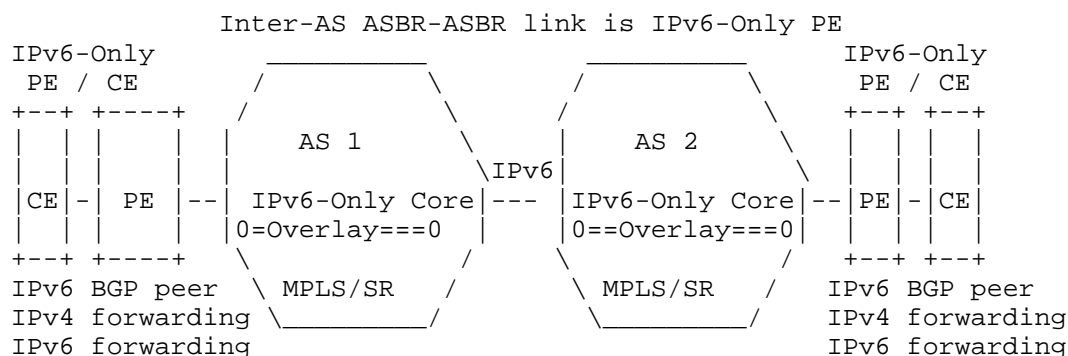


Figure 34: Test-12 E2E IPv4-Only PE-CE, VPN over IPv6-Only Core - Inter-AS Option-C

12.13. 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 35: Ping and Trace Test Case

13. SRv6 Applicability to IPv6 Only PE Design

SRv6 [RFC8986] applicability to IPv6-Only PE design

SRv6 [RFC8986] Full 128 bit SID and SRv6 Compression
[I-D.ietf-spring-srv6-srh-compression] C-SID Next C-SID (uSID)
endpoint flavor and Replace SID C-SID (G-SID) endpoint flavor is
fully supported for all 12 IPv6-Only PE Design use cases.

SRv6 Next SID utilizes a Next C-SID uSID carrier for the micro sid for up to 5 hops of steering without requiring an SRH. SRv6 Next SID uN Node sid endpoint function for uSID F3216 format shift 16 bits and forward to the next node allows the vanilla IPv6 data plane to be used to seamlessly stitch. across inter-as inter domain hops.

The following SRv6 Compression C-SID Next C-SID (uSID) endpoint flavor 4 use cases will be fully tested by Cisco, Juniper, Nokia.

1. Global Table 4PE Test Scenario

Global Table IPv6 Only PE Design where Inter-AS PE-PE core boundary AFI/SAFI 1/1 IPv4 Unicast 2/1 IPv6 Unicast (4PE) and are carried over a single IPv6 peering. In this use case the Inter-AS PE-PE nodes are SRv6 Next SID endpoint behavior capable and perform the steering function Un endpoint behavior shift 16 bits and forward function. This SRv6 compression use case is applicable to IPv6 Only PE Design Test-9 - IPv6 Core 4PE Global table BGP overlay stitching in SRv6 terms is Next Hop Unchanged Global BGP overlay stitching. Use case where Inter-AS PE-PE Edge boundary AFI/SAFI 1/1 IPv4 Unicast 2/1 IPv6 Unicast are carried over a single IPv6 peer the peering.

* Global Table (4PE) Test Scenario

- SRv6 Test Scenarios (Global)
 - 1. SRv6 enabled on inter-as PE-PE nodes (SRv6 aware inter domain SRv6 fabric stitching and active traffic steering)
 - 2. SRv6 disabled on inter-as PE-PE nodes (SRv6 unaware inter domain IPv6 Data plane stitching without traffic steering)

2. VPN Overlay Test Scenario

VPN Overlay IPv6 Only PE Design where Inter-AS PE-PE core boundary AFI/SAFI 1/128 IPv4 VPN 2/128 IPv6 VPN and are carried over a single IPv6 peering. In this use case the Inter-AS PE-PE nodes are SRv6 Next SID endpoint behavior capable and perform the steering function Un endpoint behavior shift 16 bits and forward function. This SRv6 compression use case is applicable to IPv6 Only PE Design Test-11 - IPv6 Core L3 VPN Inter-AS Option B which in SRv6 terms is Next Hop Unchanged VPN overlay stitching. Use case where Inter-AS PE-PE Edge boundary AFI/SAFI 1/128 IPv4 VPN 2/128 IPv6 VPN are carried over a single IPv6 peering.

* IP VPN Scenario

- SRv6 Test Scenarios (IP VPN)

- 1. SRv6 enabled on inter-as PE-PE nodes (SRv6 aware inter domain SRv6 fabric stitching and active traffic steering)
- 2. SRv6 disabled on inter-as PE-PE nodes (SRv6 unaware inter domain IPv6 Data plane stitching without traffic steering)

14. Typed and Untyped SAFI

Typed SAFI are Not Applicable to this specification

Typed SAFI are SAFI where the SAFI uses the Route Types in which case there is a NLRI "Route Type" field followed by NLRI Data field containing the IPv4 and IPv6 NLRI. With Non Typed SAFI the IPv4 and IPv6 NLRI is encoded directly within the NLRI making it directly data plane dependent with CP-DP control plane RIB programming followed by data plane FIB programming. This Typed SAFI case exists where the control plane procedures exist in which case the data plane is instantiated making the IPv4 and IPv6 NLRI for the Typed SAFI data plane agnostic. IPv4 Only PE design has a single IPv4 Peer that carries the IPv4 and IPv6 NLRI and the interface only has an IPv4 address and requires a vendor specific knob for IPv6 forwarding. IPv6 Only PE design has a single IPv6 Peer that carries the IPv4 and IPv6 NLRI and the interface only has an IPv6 address and requires a vendor specific knob for IPv4 forwarding. Since typed routes are data plane agnostic and so are not directly correlated to the IPv4 and IPv6 protocol forwarding makes the Typed SAFI listed below not applicable to this specification.

* Typed SAFI

- The following AFI/SAFI are not applicable (N/A) to this specification
- 1. MCAST-VPN [RFC6514]
- 2. MCAST-VPLS [RFC7117]
- 3. EVPN [RFC7432]
- 4. BGP-LS [RFC7752]

15. IPv4-Only PE Design and IPv6-Only PE Design ALL AFI/SFI Operational Considerations

With a single IPv4 Peer or 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 allocation mode where both IPv4 and IPv6 prefixes have the same label in no table lookup pop-n-forward mode should be taken into consideration.

16. IANA Considerations

New IANA capability codepoint is requested for next hop encoding for IPv4 next hop applicable to all SAFI for IPv4-Only PE design Inter-AS scenarios PE-PE, PE-CE as well as existing Intra-AS PE-RR peering Scenarios.

17. 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].

18. Acknowledgments

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

19. Contributors

The following people contributed substantive text to this document:

Mohana Sundari
EMail: mohanas@juniper.net

20. References

20.1. 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-09, 24 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-bgp-multicast-09>>.
- [I-D.ietf-bess-ipv6-only-pe-design]
Mishra, G. S., Mishra, M. P., Tantsura, J., Madhavi, S., Yang, Q., Simpson, A., and S. Chen, "IPv6-Only PE Design for IPv4-NLRI with IPv6-NH", Work in Progress, Internet-Draft, draft-ietf-bess-ipv6-only-pe-design-04, 20 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-ipv6-only-pe-design-04>>.

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

Rao, D. and S. Agrawal, "BGP Color-Aware Routing (CAR)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-car-16, 20 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-16>>.

[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-nvo3-22, 6 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-nvo3-22>>.

[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-19, 28 March 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-rpd-19>>.

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

Dunbar, L., Hares, S., Majumdar, K., Raszuk, R., and V. Kasiviswanathan, "BGP UPDATE for SD-WAN Edge Discovery", Work in Progress, Internet-Draft, draft-ietf-idr-sdwan-edge-discovery-25, 22 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sdwan-edge-discovery-25>>.

[I-D.ietf-idr-segment-routing-te-policy]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-26, 23 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-26>>.

[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-51, 23 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsvr-bgp-spf-51>>.

[I-D.ietf-spring-srv6-srh-compression]

Cheng, W., Filsfils, C., Li, Z., Decraene, B., and F. Clad, "Compressed SRv6 Segment List Encoding (CSID)", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-23, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-srh-compression-23>>.

[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-05, 13 February 2025, <<https://datatracker.ietf.org/doc/html/draft-mpmz-bess-mup-safi-05>>.

[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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [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>>.

20.2. 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-17, 6 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-dynamic-cap-17>>.

- [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 and IPv6-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.

* Unicast

* 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 and IPv6-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 and IPv6-Only PE Design listing of Inter-AS PE-PE ALL applicable SAFI. All Typed SAFI are listed as Not Applicable (N/A) Here we show the catagorization grouping by columns of the SAFI into 5 use case categories.

* Unicast

* Multicast

* L1-L2 VPN

* Tunnel

* BGP Policy

SAFI	Description	Unicast	Multicast	L1-L2	Tunnel	BGP	Reference
Value				VPN		Policy	
1	NLRI Unicast	Yes	No	No	No	No	[RFC4760]
2	NLRI Multicast	No	Yes	No	No	No	[RFC4760]
4	NLRI MPLS Laels	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	N/A	N/A	N/A	N/A	N/A	[RFC7117]
9	BGP-SFC	No	No	No	No	No	[RFC9015]
64	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 N/A N/A N/A N/A N/A [RFC7432] +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ -----+-----+ 71 BGP-LS N/A N/A N/A N/A N/A [RFC7752] +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ -----+-----+ 72 BGP-LS-VPN N/A N/A N/A N/A N/A [RFC7752] +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ -----+-----+ 73 SR-TE Policy No No No No Yes [I-D.ietf-idr-segment-routi ng-te-policy] SAFI +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
---	--	--	--	--	--	--	--	--	--

74 discovery]	SD-WAN Capabilities 	No	No	No	No	Yes	[[I-D.ietf-idr-sdwan-edge-di
75	Routing Policy SAFI 	No	No	No	No	No	[[I-D.ietf-idr-rpd]
77]	Tunneled Traffic Flowspec 	No	No	No	Yes	No	[[I-D.ietf-idr-flowspec-nvo3
78 t]	MCAST-TREE SAFI 	No	Yes	No	No	No	[[I-D.ietf-bess-bgp-multicas
80	BGP-LS-SPF 	No	No	No	No	No	[[I-D.ietf-lsvr-bgp-spf]
83	BGP CAR 	N/A	N/A	N/A	N/A	N/A	[[I-D.ietf-idr-bgp-car]
84	BGP CAR VPN 	N/A	N/A	N/A	N/A	N/A	[[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 Multicast VPN 	N/A	N/A	N/A	N/A	N/A	[[RFC6513]
132	Route Target Constrains 	Yes	No	No	No	No	[[RFC4684]
133	Dissemination 	No	No	No	No	Yes	[[RFC8955]

IPv4-Only PE Design and IPv6-Only PE design supports 20 / 32 IANA SAFI's of which the 7 not supported do not use AFI=1 IPv4 or AFI=2 IPv6 and 5 are N/A.

Authors' Addresses

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

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

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

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

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

Shuanglong Chen
Huawei Technologies
Email: chenshuanglong@huawei.com