

v6ops Working Group
Internet-Draft
Intended status: Informational
Expires: 12 April 2026

C. Xie
C. Ma
China Telecom
X. Li
CERNET Center/Tsinghua University
G. Mishra
Verizon Inc
T. Graf
Swisscom
9 October 2025

Framework of Multi-domain IPv6-only Underlay Network and IPv4-as-
a-Service
draft-ietf-v6ops-framework-md-ipv6only-underlay-14

Abstract

For the IPv6 transition, IPv6-only is considered as the final stage, where only IPv6 protocol is used for transport while maintaining global reachability for both IPv6 and IPv4 services. This document illustrates a framework of multi-domain IPv6-only underlay network from an operator's perspective. In particular, it proposes stateless address mapping as the base for enabling IPv4 service data transmission in an multi-domain IPv6-only environment (i.e., IPv4-as-a-Service). It describes the methodology of stateless IPv4/IPv6 mapping, illustrates the behaviors of network devices, analyzes the options of IPv6 mapping prefix allocation, examines the utilization of SRv6, and discusses the security considerations.

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 12 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	3
1.1. Requirements Language	4
2. Terminology	4
3. Scenarios	6
4. IPv6-only Deployment in Multi-domain Network	7
5. IPv4/IPv6 Address Mapping for IPv4aaS	9
5.1. IPv4/IPv6 Address Mapping	9
5.2. End-to-End IPv4 Service Delivery	10
6. Framework Introduction	11
6.1. Overview	11
6.2. Mapping Rule Processing at the Control Layer	12
6.3. Packet Processing at the Forwarding Layer	13
7. IPv6 Mapping Prefix Allocation	15
8. Applicability of SRv6 for Multi-domain IPv6-only Network	16
9. Security Considerations	16
9.1. Authenticity and Integrity of Packets	16
9.2. BGP-4 and Multiprotocol Extensions for BGP-4	17
10. IANA Considerations	17
11. Acknowledgements	17
12. Contributors	17
13. References	18
13.1. Normative References	18
13.2. Informative References	18
Authors' Addresses	20

1. Introduction

IPv6 capabilities have been widely deployed over the past decade, with IPv6 traffic growing at a faster rate than IPv4. As of 2022, most IPv6 deployments rely on dual-stack [RFC4213]. However, dual-stack has long-term drawbacks, including duplicated network resources and states, as well as increased operational complexity from maintaining both protocol stacks. For instance, when broadband users experience access service failure, operators must determine whether the failure stems from IPv4 or IPv6, effectively doubling troubleshooting efforts. Once IPv6 adoption becomes dominant, transitioning to IPv6-only could reduce resource overhead and simplify operations.

In 2016, the IAB stated that it “expects the IETF to no longer mandate IPv4 compatibility in new or updated protocols, with future IETF work focusing on IPv6 optimization” [IAB-statement]. To ensure service continuity after IPv4 address exhaustion, operators must deploy IPv6 while maintaining access to the global IPv4 Internet—commonly referred to as IPv4-as-a-Service (IPv4aaS)—a logical approach for IPv6-only networks.

The network infrastructure of large operators typically consists of at least an access section and a backbone section. The access section serves customer by delivering access links, assigning addresses, and enabling two-way data transmission. The backbone section, also known as the backbone network, is usually a multi-domain network comprising interconnected autonomous systems (i.e., ASes), each with a full-mesh or partial-mesh topology. The backbone network is sometimes referred to as the underlay network. Accordingly, IPv6-only deployment involves two key sections: IPv6-only in the access section and IPv6-only in the backbone section.

For the IPv6-only deployment in the access section, to date various transition technologies like 4G4XLAT [RFC6877], MAP-T [RFC7599], MAP-E [RFC6333] and DS-Lite [RFC6333] have been developed and deployed [RFC9313]. These solutions allocate only IPv6 addresses to customer terminals or networks, addressing IPv4 address exhaustion on the user side while enabling access to both IPv4 and IPv6 Internet services. [I-D.ietf-v6ops-6mops] illustrates a deployment scenario called “an IPv6-Mostly network”, when IPv6-only and IPv4-enabled endpoints coexist on the same network (network segment, VLAN, SSID etc). It allows IPv6-capable devices to remain IPv6-only while the network is seamlessly supplying IPv4 access to those that require it.

However, the current IPv6-only ecosystem remains incomplete, particularly in backbone networks. For large-scale network operators, a comprehensive multi-domain IPv6-only framework is needed to integrate multiple related technologies to ensure seamless IPv4 and IPv6 data transmission. A key objective is enabling efficient cross-domain IPv4 service delivery across IPv6-only network, utilizing tunneling or translation to forward data between edge devices. With such a framework, IPv4-IPv6 packet conversion relies on stateless address mapping at the edges, meaning no user-specific state or translation tables are required for packet processing. In addition, for any IPv4 traffic flow traversing a multi-domain IPv6-only network, there should be no IPv4-to-IPv6 conversion point along the data path.

This document presents a framework for building a large-scale IPv6-only network from network operators' perspective. As it is described at a high level, this framework is not meant to replace existing IPv6-only technologies but rather to leverage and remain compatible with them. When transmitting IPv4 service data, this framework enables end-to-end tunneling across multiple network providers, and therefore is different from existing SRv6/MPLS VPN solutions which are for a single NP. In addition, networks implemented under this framework can integrate with other IPv6-only access methods.

Unless otherwise stated, the term "IPv6-only network" in this document refers specifically to "IPv6-only underlay network". This document does not propose new IPv6 transition mechanisms or IPv4aaS solutions.

1.1. 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.

2. Terminology

The following terms are used in this document:

- * Multi-domain IPv6-only underlay network: IPv6-only underlay network which consists of multiple ASes operated by single or multiple operators.
- * AS: Autonomous System.

- * UE: User Equipment, e.g., mobile phone.
- * CLAT: Customer-side translator (Section 1 of [RFC6877]).
- * CE: Customer Equipment.
- * DC: Data Center.
- * IXP: Internet Exchange Point.
- * MD: Mapping rule Database.
- * NP: Network Provider.
- * NSP: Network-Specific Prefix.
- * P: Provider Router.
- * PE: Provider Edge (Section 5.2 of [RFC4026]).
- * IPv4-embedded IPv6 address: IPv6 address used to represent IPv4 nodes in an IPv6 network. This address includes a 32-bit embedded IPv4 address and are also referred to as IPv6-mapped addresses ([RFC6052]).
- * IPv4-embedded IPv6 packet: An IPv6 packet created through encapsulation or translation of an IPv4 packet, where the source and destination IPv4 addresses are statelessly mapped to corresponding IPv6 addresses.
- * PLAT: Provider-side translator (Section 1 of [RFC6877]).
- * ADPT: Adapter in PE, a function entity that implements the two-way IPv4 and IPv6 packet conversion for IPv4 service delivery over IPv6-only network.
- * Conversion point: A function that converts between IPv4 and IPv6 realms, such as the translation (XLAT) function defined in [RFC6144]
- * RAN: Radio Access Network.
- * WKP: Well-Known Prefix.

3. Scenarios

For the purpose of the framework described in this document, an IPv6-only network is defined as a cluster of inter-connected ASes where the underlay network is IPv6, with no IPv4 present. An IPv6-only network may interconnect with external networks, including IPv4-only networks.

As a general network infrastructure, an IPv6-only network should support the following scenarios,

Scenario 1: IPv6 user accessing an IPv4 server, i.e., an IPv6-only user accesses to an IPv4-based service hosted in a data center or other places.

Scenario 2: IPv4 user accessing an IPv4 server, i.e., an IPv4-only user accesses to an IPv4-based service hosted in a data center or other places.

Scenario 3: IPv6 user accessing an IPv6 server, i.e., an IPv6-only user accesses to an IPv6-based service hosted in a data center or other places.

Scenario 4: IPv4 user accessing an IPv6 server, i.e., an IPv4-only user accesses to an IPv6-based service hosted in a data center or other places.

Scenario 5: DC-to-DC, i.e., an IPv6-only network provides communications between servers hosted in different data centers, regardless of whether they are IPv4, IPv6 or IPv4/IPv6 dual-stack.

Scenario 6: Transit for neighbor networks, i.e., an IPv6-only network interconnects multiple isolated IPv4-only networks, enabling IPv4 packet transmission over the IPv6 infrastructure.

Scenario 7: Mobile transport network, mobile operators can utilize IPv6-only to connect the RAN and 5G core network, delivering the specific connectivity services required for 5G application operations.

It should be noted the scenarios listed above represent only a subset of those supported by IPv6-only networks, which are at least as capable as current dual-stack networks.

4. IPv6-only Deployment in Multi-domain Network

This framework is designed to assist large operators in deploying IPv6-only in a multi-domain network. Large-scale operators typically manage network infrastructure composed of multiple interconnected autonomous systems (ASes), that's why it is called "Multi-domain Underlay Network" in this document. These ASes often support different functions, such as metro area networks (MANs), backbone networks, 4G/5G mobile core networks, DCs, and may be administered by separate departments or operators with different routing and security policies. In a multi-domain network, edge nodes are commonly referred to as Provider Edge (PE) routers. The ingress PE is the router where a packet enters the network, while the egress PE is the router where it exits. Internal nodes are typically called Provider (P) routers.

To facilitate the illustrating the framework from the perspective of operators, Figure 1 shows a multi-domain network, namely N1, which consists of three interconnected ASes, i.e., AS1, AS2, and AS3. Among them, AS1 and AS2 are operated by operator NP-1, AS3 is operated by operator NP-2. Routers located outside the backbone but directly connected to it are referred to as Customer Edge (CE) routers. AS1 of NP-1 provides connectivity services to mobile, home broadband, and enterprise customers, represented by CE1, CE2, and CE3, respectively. [RFC8585] defines IPv4 service continuity requirements for IPv6 CE routers, extending the basic IPv6 CE specifications to support IPv4 delivery in IPv6-only access networks. Additionally, service instances in DCs must support cross-site communication, whether on-premises or in external data centers. A multi-domain network must facilitate these data center connections. Network N1 should support at least two connectivity modes for data centers: the first is the mode between data center and individual users, for instance, the user of CE1 accesses the service hosted in DC1, the second is the mode between data centers, for instance, communications between service instances hosted in DC1 and DC2 respectively.

Regarding external interconnection, NP-3 is a neighboring operator of NP-2. AS4 of NP-3 is an IPv4-only network and does not support IPv6. AS4 and AS3 are interconnected via BGP.

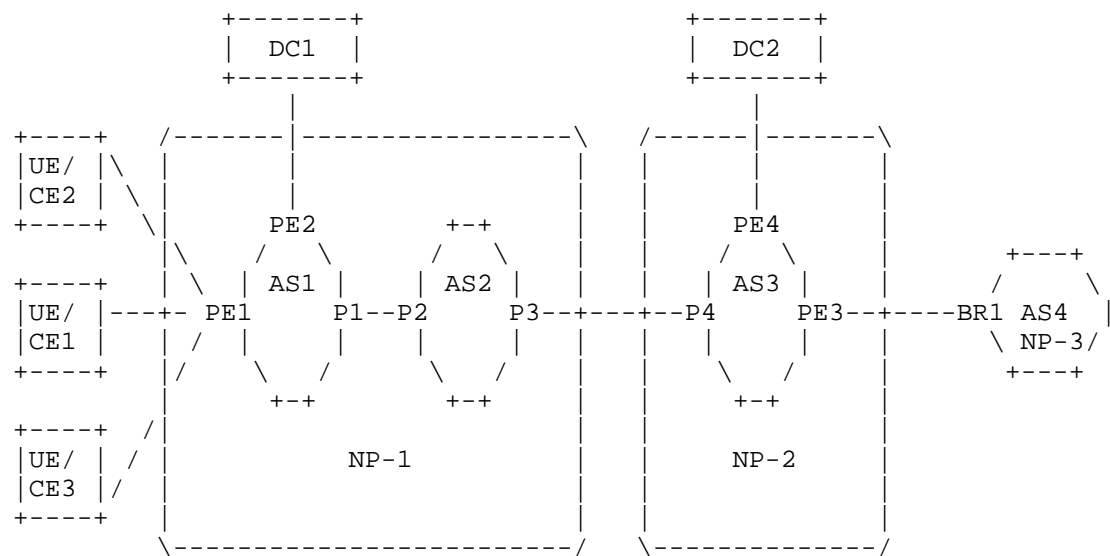


Figure 1. An Example of Multi-domain Underlay Network

For network N1, transitioning from dual-stack to IPv6-only involves gradually disabling some or all IPv4 protocol instances, making IPv6 the primary network-layer protocol. Specifically, core P routers, such as P1, P2, P3 and P4, run only IPv6, while PE routers, such as PE1, PE2, PE3 and PE4, support IPv4 on interfaces facing IPv4 client networks and IPv6 on interfaces facing the core, requiring them to handle both address families. Network N1 transports packets that originate and terminate outside the network. These packets enter the IPv6 network at a PE router, traverse the network, and exit through another PE router to continue their path.

In addition, the IPv6-only deployment in network N1 must ensure the continued operation of remaining IPv4 services without degrading user experience. As some Internet services may remain IPv4-based even in an IPv6-dominated environment. Therefore, an IPv6-only network should support access to IPv4-only services, as well as native IPv6 services. [RFC6992] describes a routing scenario where IPv4 packets are transported over an IPv6 network, based on [RFC7915] and [RFC6052], along with a separate OSPFv3 routing table for IPv4-embedded IPv6 routes in the IPv6 network. Since it is based on the OSPF protocol, it supports IPv4aaS within a single AS.

For a multi-domain network operator, deploying IPv6-only without a unified framework may lead to independent adoption of IPv6 transition approaches across different ASes. This can result in multiple IPv6-only islands interconnected by IPv4 links between domains. With

independent deployment in different domains, the network may contain multiple IPv4-IPv6 packet conversion points with varying functionalities. In such cases, IPv6 packets converted from IPv4 packets may need to revert to IPv4 at the egress of one AS, then back to IPv6 in the next domain. The number of conversion gateways increases with the number of ASes. Excessive IPv4-IPv6 conversion gateways introduce network complexity and higher capital expenditures (CAPEX). Thus, a unified framework is required to define network edge behavior for IPv4 service delivery and eliminate unnecessary IPv4/IPv6 conversion gateways within the multi-domain network.

5. IPv4/IPv6 Address Mapping for IPv4aaS

5.1. IPv4/IPv6 Address Mapping

To support IPv4aaS in a multi-domain IPv6-only network, the framework proposes each PE device to be allocated and identified by at least one IPv6 mapping prefix, denoted by Pref6(PE). Each device will also have one or more associated IPv4 address blocks which are extracted from local IPv4 routing table or address pool. The mapping relationship between an IPv4 address block and its corresponding IPv6 prefix is referred to as a mapping rule, which will have at least the following data structure.

IPv4 address block: Pref6(PE)

It should be noted that the mapping rule contains not only the data structure above, but also other necessary information to support IPv4 service delivery over IPv6-only network, the detailed structure definition of the mapping rule is out of the scope of this document.

The mapping rule of destination address will determine the direction of IPv4 service data transmission in a multi-domain IPv6-only network. When the mapping rule corresponding to the destination address of a given IPv4 packet is available, the ingress PE can generate corresponding IPv6 source and destination addresses from its IPv4 source and destination address as below,

- * The IPv6 source address is derived by appending the IPv4 source address to its local IPv6 mapping prefix, i.e., Pref6(ingress PE).
- * The IPv6 destination address is derived by appending the IPv4 destination address to the Pref6(egress PE) in the mapping rule.

Since mapping rule adopts prefix-level mapping, there is no need to maintain user-related status or translation tables for packet transformation at the PE devices.

[RFC6052] illustrates the algorithmic translation of an IPv4 address to a corresponding IPv6 address, and vice versa, using only statically configured information. With this approach, IPv4-embedded IPv6 addresses are composed by concatenating the prefix, the 32 bits of the IPv4 address, and the suffix (if needed) to obtain a 128-bit address. The prefixes can only have one of the following lengths: 32, 40, 48, 56, 64, or 96.

For IPv4 service delivery across a multi-domain IPv6-only network, it is proposed that IPv4 address is located at the last 32 bits of the IPv6 address, most significant bits first. The bits between IPv6 mapping prefix and IPv4 address can be set to zero and are reserved for future extensions. Examples of such representations are presented in Table 1.

IPv6 mapping prefix	IPv4 address	IPv4-embedded IPv6 address
2001:db8::/32	192.0.2.33	2001:db8::192.0.2.33
2001:db8:100::/40	192.0.2.33	2001:db8:100::192.0.2.33
2001:db8:122::/48	192.0.2.33	2001:db8:122::192.0.2.33

Table 1. Text Representation of IPv4-Embedded IPv6 Address

Prior to IPv4/IPv6 packets conversion, the mapping rule for the destination address needs to be obtained remotely in advance. To meet this requirement, specific mechanism of mapping rule exchange needs to be designed, the exchange can be within or across domains. Using the mapping rule exchange mechanism, an egress PE can inform other PEs that an IPv4 packet with a destination address within the IPv4 address block of a mapping rule should be forwarded to the egress PE identified by the corresponding IPv6 mapping prefix.

5.2. End-to-End IPv4 Service Delivery

To enable IPv4 service data forwarding in a multi-domain IPv6-only network, IPv4 packets must be converted to IPv6 packets-either at the UE/CE or at the edge PE of the network. Consider the network case of section 4, when one ingress PE, e.g., PE1, receives an IPv4 packet (destined for a remote IPv4 network) from a client-facing interface, it queries its mapping rule database(i.e., MD) in PE1 to find the rule that best matches the packet's destination IPv4 address. The IPv6 mapping prefix in the rule identifies the corresponding egress PE. In this case, the ingress and egress PEs belong to different autonomous systems (ASes), the ingress PE1 resides in NP-1, while the egress PE3 is in NP-2. The ingress PE converts the IPv4 destination address into an IPv6 address using PE3's IPv6 mapping prefix and forwards the IPv6 packet to PE3. Upon receiving the IPv6 packet, PE3

extracts the original IPv4 source and destination addresses from the IPv4-embedded IPv6 addresses and reconstructs the IPv4 packet. The packet is then forwarded based on the IPv4 routing table maintained at the egress PE. In this case, there are only two IPv4-IPv6 conversion actions, which occur in PE1 and PE3 respectively, the IPv6 data path for this process is between PE1 and PE3, as shown in Figure 2.

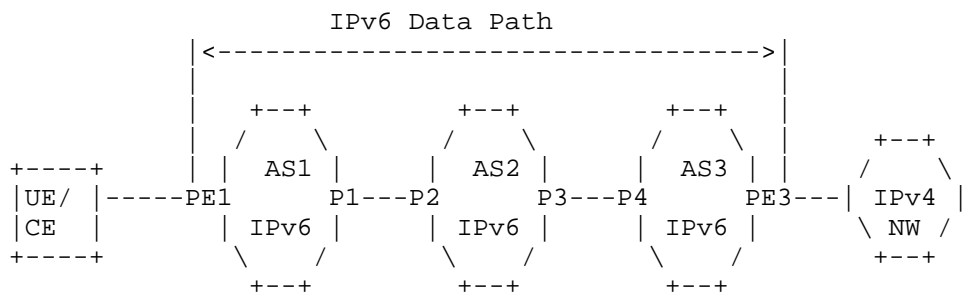


Figure 2. End-to-end IPv4 Service Delivery from Ingress to Egress

It should be noted that P3 and P4 are Ps from the perspective of the framework defined by this document, but they are PEs from the NPs' perspective for they are at the edge of the NPs' networks.

It can be seen that such a multi-NP tunneling requires only two-end NPs to support this solution for it to work, it has no specific requirements for NPs in the middle of the path, as long as it supports IPv6.

6. Framework Introduction

6.1. Overview

This section describes the framework of multi-domain IPv6-only underlay network from NP' perspective. As shown in Figure 3, the whole framework comprises PE devices at the edge, core P devices and customer-side IPv4 routers. In this framework, the PE devices are responsible for implementing stateless IPv4/IPv6 packet conversion, and they are required to support the following functions,

1. RP(Rule Processing):

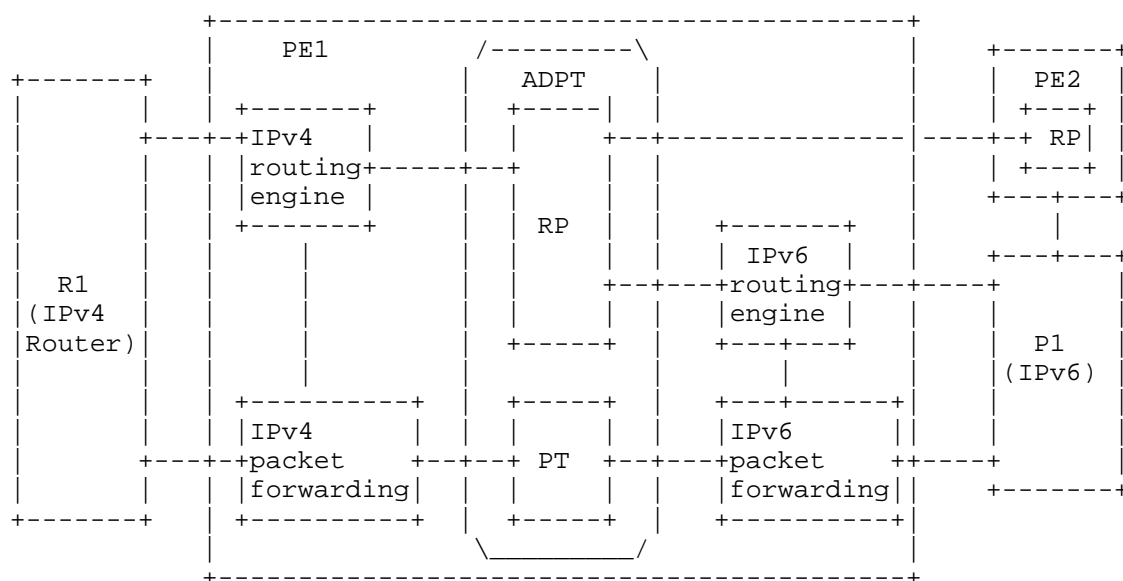
- * Manage the mapping relationships between IPv4 address blocks and their corresponding IPv6 prefixes.
- * Exchange of mapping rules and associated routing information between PEs at the routing layer.

2. PT(Packet Transformation):

- * Generate IPv4-embedded IPv6 packets using either translation or encapsulation for IPv4 data delivery.

The combination of the two functions above forms an adaptation (ADPT) within PE for transmitting IPv4 service data in IPv6-only network.

There is no specific requirements to the P devices.



RP: Rule Processing
PT: Packet Transformation

Figure 3. Multi-domain IPv6-only Underlay Network Framework

6.2. Mapping Rule Processing at the Control Layer

In this framework, IPv4/IPv6 address mapping rules are processed at the control layer, which includes two processes,

1. Mapping Rule Generation

In Figure 3, IPv4 routing engine of PE1 receives an IPv4 BGP route advertisement sent by IPv4 router R1 and generates an IPv4 route entry. The RP function of PE1 extracts IPv4 address blocks from its IPv4 BGP routing instance and generates device-specific mapping rules by combining them with its IPv6 mapping prefix.

Each PE stores all mapping rule entries in its local MD database, whether locally generated or received from other PEs. The RP function also supports management operations such as insertion, modification, and deletion of mapping rules.

If certain IPv4 address blocks are not explicitly announced by any egress PEs to the ingress PE, the ingress PE will lack corresponding mapping rules. To address this, the framework introduces a default egress PE, which advertises a default IPv6 mapping rule containing a default mapping prefix to all other PEs. The format of the default IPv4 address mapping rule is as follows:

0.0.0.0/0: Pref6(PE)

2. Mapping Rule Transport

For IPv4 service delivery, IPv4/IPv6 address mapping rules need to be exchanged between PEs at the routing layer. This process must occur before IPv4 data transmission begins, otherwise, IPv4 traffic packets will be dropped due to the lack of a corresponding IPv6 mapping prefix for the IPv4 destination address.

When a mapping rule is generated locally, PE1 converts it into a data structure compatible with the IPv6 routing system and forwards it to the IPv6 routing engine. In the opposite direction, when PE1 receives a routing update from the IPv6 routing engine, it extracts the embedded mapping rule and store it locally for further processing. It can be seen that, based on the address mapping rule received, the PE who has sent the mapping rule (i.e., Egress PE) is visible to the ingress PE. To enable mapping rule transmission at the routing layer, extensions to MP-BGP or other control protocols would be required, a typical approach is [I-D.ietf-idr-mpbgp-extension-4map6].

6.3. Packet Processing at the Forwarding Layer

In this framework, Forwarding Layer provides data forwarding capability to IPv6 packets, including native IPv6 packets and IPv4-embedded IPv6 packets. IPv4-embedded IPv6 packets can be generated using either translation or encapsulation for IPv4 data delivery.

1. Translation

Translation refers to the packet conversion from one protocol format to the other. When the PT function receives an IPv4 packet from the IPv4 packet forwarding module, it queries the local MD database which store all the address mapping rule, if a mapping rule exists for the IPv4 destination address, it will generate corresponding IPv6 source and destination addresses from the IPv4 addresses, following the procedure described in Section 5.1. This process should comply with [RFC7915].

Upon receiving the IPv6 packet, the egress PE firstly checks whether its destination IPv6 prefix matches its own IPv6 mapping prefix. If not, it discards or forwards it as a regular IPv6 packet. Otherwise, it the egress PE extracts the original IPv4 source and destination addresses from the IPv4-embedded IPv6 addresses and reconstructs the IPv4 packet, this process should comply with [RFC7915]. The IPv4 packet is then forwarded based on the IPv4 routing table maintained at the egress PE.

2. Encapsulation

The address mapping process for encapsulation follows the same procedure as translation: When the PT function receive an IPv4 packet from the IPv4 forwarding module, it queries the local MD database which store all the address mapping rule, if a mapping rule exists for the IPv4 destination address, it will generate corresponding IPv6 source and destination addresses from the IPv4 addresses, following the procedure described in Section 5.1. This process should comply with [RFC7915].

Upon receiving the IPv6 packet, the egress PE firstly checks whether its destination IPv6 prefix matches its own IPv6 mapping prefix. If not, it discards or forwards it as a regular IPv6 packet. Otherwise, the egress PE decapsulation it by removing outer IPv6 header and restores the IPv4 packet. The IPv4 packet is then forwarded based on the IPv4 routing table maintained at the egress PE.

For IPv4-embedded IPv6 packets, the Pref6 portion of the destination address identifies the network egress, regardless of whether translation or encapsulation was used. Therefore, packet forwarding can be performed by P devices based solely on the Pref6 part of the destination address.

Although this document illustrates the framework of multi-domain IPv6-only network operated by a single operator, this multi-domain model can naturally be extended to IPv6-only network which is operated by multiple operators.

7. IPv6 Mapping Prefix Allocation

With this framework, a specific IPv6 address range, i.e., IPv6 mapping prefix, is used to represent an IPv4 address block by stateless mapping as illustrated in section 5.1, there are two options to allocate IPv6 mapping prefixes:

1) Well-Known Prefix (WKP)

A specific WKP can be allocated from the global IPv6 address prefix, e.g., 64:ff9b::/96, or an IPv6 address prefix specifically assigned for this purpose.

Pros:

This can offer two key advantages for IPv6 mapping prefix allocation. First, operators are not required to dedicate IPv6 address prefixes from their own resources for mapping IPv4 addresses. Second, they can easily control the range of IPv6 mapping routes, such as applying routing restrictions at network boundaries to prevent leakage into external networks.

Cons:

When the PE device converts an IPv4 address to an IPv6 address using a Well-Known Prefix, the IPv4 portion of the IPv4-embedded IPv6 address is used for routing the packet. Consequently, multiple specific routes with prefix lengths greater than 96 may be inserted into the FIB of P routers in an IPv6-only network. However, most networks do not support such fine-grained routes with prefix lengths exceeding 96.

2) Network Specific Prefix (NSP)

NSP refers to a dedicated IPv6 prefix (or prefixes) assigned from operator's available IPv6 address pool to each PE for IPv4 addresses mapping. The assigned IPv6 mapping prefix differs per PE.

Pros:

In a multi-domain network, the length of the IPv6 mapping prefix can be adjusted to meet IPv6 routing requirements. The IPv6 mapping prefix is part of a larger and routable IPv6 address block assigned to the PE, so this approach is unlikely introduce new routing entries or impact the global IPv6 routing system. For IPv4-embedded IPv6 packet, the P devices can forward them as regular IPv6 packets without requiring a specific FIB entry for the mapping prefix.

Cons:

If the operator has not implemented specific address prefix planning and policy configuration, interworking between operators may result in the same IPv4 address block receiving multiple NSP prefixes from different operators. This can generate multiple IPv6 mapping routes, potentially increasing the size of IPv6 routing tables (including FIB and RIB).

As specified in Section 5.1, each PE must be assigned at least one IPv6 mapping prefix. This prefix serves as the fundamental information for forwarding IPv4-embedded IPv6 packets to the correct egress PE. Operators should carefully determine the IPv6 mapping prefix length during implementation. The length of all the IPv6 mapping prefixes is recommended to be the same, to avoid unnecessary processing cost and complexity induced by the prefix length diversity.

8. Applicability of SRv6 for Multi-domain IPv6-only Network

SRv6 [RFC8986] enables network operators to specify a packet processing program by encoding a sequence of Segment IDs (SIDs) in the IPv6 packet header. It can also use specific SIDs (e.g., DT4 or DX4) to transport IPv4 packets across an IPv6-only network from one PE device to another. However, SRv6 is a generally single-NP solution and not suitable for a multi-NP scenario, for the current specification assumes SRv6 deployment within a trusted domain, this may limit the joint deployment of IPv6-only by multiple operators. In addition, if Traffic Engineering (TE) is needed in the case of Single-NP deployment, SRv6 can be used. With SRv6, IPv6-only network can optimize network performance by controlling traffic paths to avoid congestion and balance load.

9. Security Considerations

Besides regular security checks on configured mapping rules, the following two aspects need to be considered as well.

9.1. Authenticity and Integrity of Packets

In this framework, for each egress PE, they assume that all ingress PEs are legal and authorized to convert the received IPv4 packets into IPv6 packets and send them into IPv6-only network. If IPv6 packets cannot guarantee its authenticity or integrity, then there may be a spoofing attack. Some faked ingress PEs can send IPv6 data converted from IPv4 to attack the egress PE. After the egress PE recovers the received IPv6 packets into IPv4 packets, they are routed based on the destination IPv4 address and enter the Internet. Since

the PE in this framework is stateless, the effect of the attack is limited.

9.2. BGP-4 and Multiprotocol Extensions for BGP-4

The framework allows BGP to propagate mapping rule information over an IPv6-only underlay network, BGP is vulnerable to traffic diversion attacks. The ability to advertise a mapping rule adds a new means by which an attacker could cause traffic to be diverted from its normal path. Such an attack differs from pre-existing vulnerabilities in that traffic could be forwarded to a distant target across an intervening network infrastructure (e.g., an IPv6 core), allowing an attack to potentially succeed more easily since less infrastructure would have to be subverted. The security issues already exist in BGP-4 and MP-BGP for IPv6, the same security mechanisms are applicable.

10. IANA Considerations

There are no other special IANA considerations.

11. Acknowledgements

The authors would like to thank Brian E. Carpenter, Mohamed Boucadair, Bob Harold, Fred Baker, Xipeng Xiao, Giuseppe Fioccola, Vasilenko Eduard, Zhenbin Li, Jen Linkova, Ron Bonica, Shuping Peng, Jingrong Xie, Eduard Metz, Wu Qin, Dhruv Dhody, Nick Buraglio, Linda Dunbar, Weiqiang Cheng, Aijun Wang, Daryll Swer, Tim Wicinski, David 'equinox' Lamparter, Tianran Zhou and Huaimo Chen for their review and comments.

12. Contributors

Guoliang Han
Indirection Network Inc.
China
Email: guoliang.han@indirectionnet.com

Ruoyu Zhao
Beijing TC Group
China
Email: api_zhao@126.com

Linjian Song
Alibaba Cloud
China

Email: linjian.slj@alibaba-inc.com

13. References

13.1. Normative References

- [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>>.
- [RFC3587] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format", RFC 3587, DOI 10.17487/RFC3587, August 2003, <<https://www.rfc-editor.org/info/rfc3587>>.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, DOI 10.17487/RFC4026, March 2005, <<https://www.rfc-editor.org/info/rfc4026>>.
- [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>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/info/rfc6877>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/info/rfc7915>>.
- [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>>.

13.2. Informative References

[I-D.ietf-idr-mpbgp-extension-4map6]

Xie, C., Dong, G., Li, X., Han, G., and Z. Guo, "MP-BGP Extension and the Procedures for IPv4/IPv6 Mapping Advertisement", Work in Progress, Internet-Draft, draft-ietf-idr-mpbgp-extension-4map6-04, 14 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-mpbgp-extension-4map6-04>>.

[I-D.ietf-spring-srv6-security]

Buraglio, N., Mizrahi, T., tongtian124, Contreras, L. M., and F. Gont, "Segment Routing IPv6 Security Considerations", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-security-07, 18 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-security-07>>.

[I-D.ietf-v6ops-6mops]

Buraglio, N., Caletka, O., and J. Linkova, "IPv6-Mostly Networks: Deployment and Operations Considerations", Work in Progress, Internet-Draft, draft-ietf-v6ops-6mops-02, 26 August 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-6mops-02>>.

[IAB-statement]

"IAB statement", <<https://www.iab.org/2016/11/07/iab-statement-on-ipv6/>>.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <<https://www.rfc-editor.org/info/rfc4213>>.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, DOI 10.17487/RFC6144, April 2011, <<https://www.rfc-editor.org/info/rfc6144>>.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<https://www.rfc-editor.org/info/rfc6333>>.

[RFC6992] Cheng, D., Boucadair, M., and A. Retana, "Routing for IPv4-Embedded IPv6 Packets", RFC 6992, DOI 10.17487/RFC6992, July 2013, <<https://www.rfc-editor.org/info/rfc6992>>.

- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<https://www.rfc-editor.org/info/rfc7599>>.
- [RFC8585] Palet Martinez, J., Liu, H. M.-H., and M. Kawashima, "Requirements for IPv6 Customer Edge Routers to Support IPv4-as-a-Service", RFC 8585, DOI 10.17487/RFC8585, May 2019, <<https://www.rfc-editor.org/info/rfc8585>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [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>>.
- [RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. Liu, "Operational Implications of IPv6 Packets with Extension Headers", RFC 9098, DOI 10.17487/RFC9098, September 2021, <<https://www.rfc-editor.org/info/rfc9098>>.
- [RFC9313] Lencse, G., Palet Martinez, J., Howard, L., Patterson, R., and I. Farrer, "Pros and Cons of IPv6 Transition Technologies for IPv4-as-a-Service (IPv4aaS)", RFC 9313, DOI 10.17487/RFC9313, October 2022, <<https://www.rfc-editor.org/info/rfc9313>>.

Authors' Addresses

Chongfeng Xie
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China
Email: xiechf@chinatelecom.cn

Chenhao Ma
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China
Email: machh@chinatelecom.cn

Xing Li
CERNET Center/Tsinghua University
Shuangqing Road No.30, Haidian District
Beijing
100084
China
Email: xing@cernet.edu.cn

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

Thomas Graf
Swisscom
Binzring 17
CH- CH-8045 Zurich
Switzerland
Email: thomas.graf@swisscom.com