v6ops Working Group
Internet-Draft
Intended status: Informational
Expires: 30 November 2026

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

29 May 2026

Framework for Multi-domain IPv6-only Underlay Network and IPv4-as-a-Service
draft-ietf-v6ops-framework-md-ipv6only-underlay-23

Abstract

For the IPv6 transition, IPv6-only is considered the final stage where only IPv6 protocol is used for transport while maintaining global reachability for both IPv6 and IPv4 services. This document introduces a framework for a multi-domain IPv6-only underlay network from the perspective of network operators. In particular, it proposes stateless address mapping as the basis for enabling IPv4 service data transmission in a 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, and discusses the security considerations. This framework is not intended to replace existing IPv6-only technologies, but rather to leverage or remain compatible with them.

Status of This Memo

Table of Contents

1. Introduction 3
1.1. Requirements Language 4
2. Terminology 4
3. IPv6-only Deployment in Multi-domain Network 5
4. IPv4/IPv6 Address Mapping for IPv4-as-a-Service 8
4.1. IPv4/IPv6 Address Mapping 8
4.2. End-to-End IPv4 Service Delivery 9
5. Framework Introduction 10
5.1. Overview 10
5.2. Address Mapping Rule Processing 11
5.3. Packet Conversion and Transmission 12
6. IPv6 Mapping Prefix Allocation 13
7. Operational Considerations 14
8. Security Considerations 14
8.1. Authenticity and Integrity of Packets 14
8.2. Stateless IP/ICMP Translators 15
8.3. Issues Related to MP-BGP 15
9. IANA Considerations 15
10. Acknowledgements 16
11. Contributors 16
12. References 16
12.1. Normative References 16
12.2. Informative References 17
Authors' Addresses 18

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, Network Providers (NPs) have to determine whether the failure stems from IPv4 or IPv6, effectively doubling the troubleshooting effort. Once IPv6 adoption becomes dominant, transitioning to IPv6-only can 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, network operators (NPs) require that the network maintains access to the global IPv4 Internet when deploying IPv6. This practice is commonly referred to as IPv4-as-a-Service and is a logical approach for IPv6-only networks.

The network infrastructure of large NPs typically consists of at least an access section and a backbone section. The access section serves customers by delivering access links, assigning addresses, and enabling two-way data transmission. The backbone section, also known as the backbone network, is typically a multi-domain network comprising interconnected autonomous systems (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 such as 464XLAT [RFC6877], MAP-T [RFC7599], MAP-E [RFC7597], 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] describes a deployment scenario referred to as "an IPv6-Mostly network", where 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 NPs, 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 to enable efficient IPv4 service delivery across a multi-domain IPv6-only network, utilizing tunneling or translation to forward data between edge PE 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, there will be no IPv4-to-IPv6 conversion gateway along the data path.

Unless otherwise stated, the term “IPv6-only network” in this document refers specifically to “IPv6-only underlay network”. This document presents a framework for building a large-scale IPv6-only network from the perspective of network operators. 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, nor does it propose any new IPv6 transition mechanisms or IPv4-as-a-Service solutions. When transmitting IPv4 service data, this framework enables end-to-end tunneling or translation across multiple network providers, and therefore is different from existing SRv6/MPLS VPN solutions which are for a single NP. Unlike IP-in-IP tunnel mechanism, where tunnels often operate alongside routing protocols that constitute a control plane (e.g., routing adjacencies over the tunnel), this framework includes a specific mechanism for distributing IPv4-to-IPv6 mapping rules across domains.

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:

* AS: Autonomous System.

* CE: Customer Edge.

* DC: Data Center.

* IPv4-embedded IPv6 address: An IPv6 address embedded with a 32-bit IPv4 address used to represent an IPv4 node in an IPv6 network. ([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.

* MAN: Metro Area Network.

* Multi-domain IPv6-only underlay network: IPv6-only underlay network which consists of multiple ASes operated by single or multiple network providers.

* MR-DB: Mapping Rule DataBase.

* NP: Network Provider.

* NSP: Network-Specific Prefix.

* P: Provider Router.

* PE: Provider Edge, a device at the edge of the IPv6-only underlay network, providing the functionality required for IPv4-as-a-Service.

* UE: User Equipment, e.g., mobile phone.

3. IPv6-only Deployment in Multi-domain Network

This framework is designed to assist large NPs in deploying IPv6-only networks in a multi-domain environment. Large-scale NPs usually manage network infrastructure comprising multiple interconnected Autonomous Systems (ASes). This is referred to as a "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, Data Centers (DCs), and may be administered by separate departments or NPs with different routing and security policies. In a multi-domain network environment, 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.

As some Internet services may remain IPv4-based even in an IPv6-dominated environment, an IPv6-only network needs to support access to IPv4-only services, as well as 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 IPv4-as-a-Service within a single AS.

To facilitate the illustration of the framework from the perspective of NPs, Figure 1 shows a multi-domain network, namely NP-1, which consists of three interconnected ASes, i.e., AS1, AS2, and AS3. Among them, AS1 and AS2 are operated by NP-1, AS3 is operated by 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 often require cross-site communication, whether on-premises or in external data centers. A multi-domain network needs to facilitate these data center connections. Network NP-1 needs to support at least two connectivity modes for data centers: The first is between a data center and individual users, for instance, a user of CE1 accesses a service instance hosted in DC1; The second is between data centers, for instance, communications occur between service instances hosted in DC1 and DC2, respectively.

Regarding external interconnection, NP-3 is a neighboring network of NP-2. AS4 of NP-3 is an IPv4-only network and does not support IPv6. The interconnection protocol between AS3 and AS4 is BGP that supports IPv4 route advertisement.

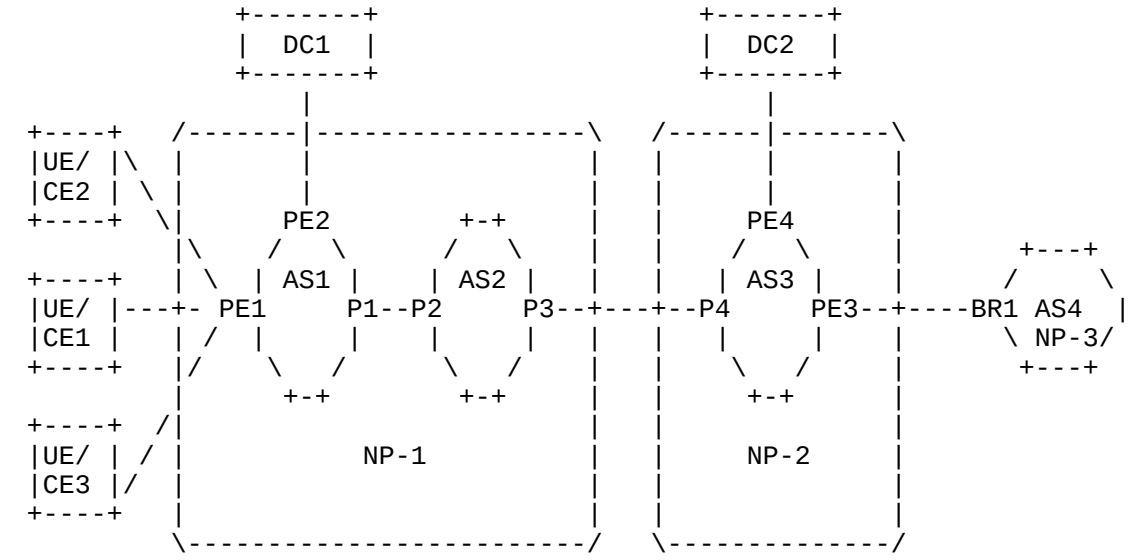

Figure 1: An Example of Multi-domain Underlay Network

For network NP-1, 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. Furthermore, the network may operate multiple IPv4-IPv6 packet conversion gateways with varying functionalities. For a given AS, incoming IPv4 packets are converted to IPv6 at an ingress, then reverted to IPv4 at an egress. When the IPv4 packets reach the next AS, another round of IPv4 -> IPv6 -> IPv4 packet conversion is carried out. Excessive IPv4-IPv6 conversion gateways introduce network complexity and increases 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.

For IPv6-only deployment guided by a unified framework, IPv4 protocol instances are gradually disabled and IPv6 will be the primary network-layer protocol. Specifically, core P routers, such as P1, P2, P3, and P4, operate only IPv6 protocol, while PE routers, such as PE1, PE2, PE3, and PE4, support IPv4 protocol on interfaces facing IPv4 client networks and IPv6 on interfaces facing the core, requiring them to handle both address families. Network NP-1 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.

4. IPv4/IPv6 Address Mapping for IPv4-as-a-Service

4.1. IPv4/IPv6 Address Mapping

To support IPv4-as-a-Service in a multi-domain IPv6-only network, the framework proposes that each PE device be allocated and identified by at least one IPv6 mapping prefix, denoted by Pref6(PE). Each PE 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 called an address mapping rule, which can be represented at a minimum by the following data structure.

IPv4 address block: Pref6(PE)

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

The address mapping rule for the destination address will determine the direction of IPv4 service data transmission in a multi-domain IPv6-only network. When the address 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 its address mapping rule.

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

This framework proposes to algorithmically translate an IPv4 address to a corresponding IPv6 address, and vice versa, using only statically configured information. The IPv6 address translated from an IPv4 address is called an IPv4-embedded IPv6 address, which has been defined in [RFC6052]. As shown in Section 2.2 of [RFC6052], IPv4-embedded IPv6 addresses are composed of a variable-length prefix, the embedded IPv4 address, and a variable-length suffix. Table 1 shows examples of such representations, it is the same as Table 1 of [RFC6052], except that the first column of the latter is

"Network-Specific Prefix". Note that [RFC7915] also allows the use of unicast addresses without u-bit (as long as they're not derived from an IEEE MAC-layer address).

IPv6 mapping prefix	IPv4 address	IPv4-embedded IPv6 address
2001:db8::/32	192.0.2.33	2001:db8:c000:221::
2001:db8:100::/40	192.0.2.33	2001:db8:1c0:2:21::
2001:db8:122::/48	192.0.2.33	2001:db8:122:c000:2:2100::
2001:db8:122:300::/56	192.0.2.33	2001:db8:122:3c0:0:221::
2001:db8:122:344::/64	192.0.2.33	2001:db8:122:344:c0:2:2100::
2001:db8:122:344::/96	192.0.2.33	2001:db8:122:344::192.0.2.33

Table 1: Representation Examples of IPv4-Embedded IPv6 Address

Prior to IPv4/IPv6 packet conversion, an ingress PE needs to obtain the address mapping rule for the destination address within or across domains. To meet this requirement, a specific mechanism of address mapping rule exchange needs to be designed, so an egress PE can inform other PEs that an IPv4 packet with a destination address being within a specific IPv4 address block can be forwarded to itself directly.

4.2. End-to-End IPv4 Service Delivery

To enable IPv4 service data forwarding in a multi-domain IPv6-only network, IPv4 packets need to be converted to IPv6 packets - either at the UEs/CEs or at the PEs located at the edge of the network. Consider the network case of Section 3, when an ingress PE, e.g., PE1, receives an IPv4 packet (destined for a remote IPv4 network, e.g., NP-3) from a client-facing interface, it queries its mapping rule database (i.e., MR-DB) 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 reside in different autonomous systems (ASes): the ingress PE (PE1) is in AS1 of NP-1, while the egress PE (PE3) is in AS3 of 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 to NP-3 based on the IPv4 routing table maintained at PE3. In this case, the IPv6 data path is between PE1 and PE3, there are only two IPv4-IPv6 conversion actions, which occur in PE1 and PE3

respectively, as shown in Figure 2.

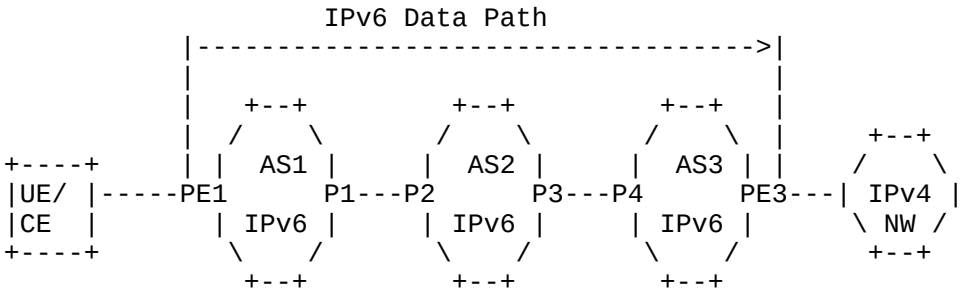

Figure 2: End-to-end IPv4 Service Delivery from Ingress to Egress

It should be noted that P3 and P4 are P routers from the perspective of the framework defined by this document, although they are the edges of providers' networks.

[I-D.ietf-idr-mpbgp-extension-4map6] can be implemented in PE1 and PE2 to support the MR-DB operations in the framework. In addition, for the mapping rules to propagate from PE3 to PE1 across the network, the intermediate BGP speakers (e.g., route reflectors / ASBRs) that propagate the relevant NLRI will support [RFC8950]. [I-D.ietf-idr-mpbgp-extension-4map6] provides IPv4-to-Pref6 mappings processing at each PE device. [RFC8950] specifies the extensions necessary to allow the advertising of IPv4 NLRI or VPN-IPv4 NLRI with a next-hop address that belongs to the IPv6 protocol. It allows gradual deployment of the functionality of advertising IPv4 reachability via an IPv6 next hop without any flag day or any risk of traffic black-holing.

5. Framework Introduction

5.1. Overview

This section outlines the multi-domain IPv6-only underlay network framework from the perspective of network operators. As shown in Figure 1, the framework consists of edge PE devices, core P devices, and customer-side IPv4 routers. The PE devices are responsible for performing stateless IPv4/IPv6 packet conversion and will support the following functions:

1. Address Mapping Rule Processing

* Generate and manage the address mapping rules

* Exchange the address mapping rules across an IPv6-only network

2. Packet Conversion

* Generate IPv4-embedded IPv6 packets using either translation or encapsulation

* Recover IPv4 packets from IPv4-embedded IPv6 packets via translation or decapsulation.

5.2. Address Mapping Rule Processing

Within PE devices, IPv4/IPv6 address mapping rules are processed at the control layer, which includes two processes,

1. Address Mapping Rule Generation and Management

For IPv4 service delivery, IPv4/IPv6 address mapping rules need to be generated. In the network shown in Figure 1, when PE3 receives an IPv4 BGP route advertisement from an IPv4 router, e.g., BR1, it extracts IPv4 address blocks and generates address mapping rules by combining them with its own IPv6 mapping prefix. All the address mapping rules, whether locally generated or received from other PEs, are stored in its local MR-DB. PE devices also support rule management operations, such as insertion, modification, and deletion of address mapping rules.

If the address mapping rule of a certain IPv4 address block has not been received by the ingress PE, the IPv4 service data destined to that IPv4 address block will not be forwarded to the correct egress PE. To mitigate this issue, the framework introduces a default egress PE, which advertises a default address mapping rule to all other PEs. The format of the default address mapping rule is as follows:

0.0.0.0/0: Pref6(PE)

With the availability of a default egress PE, the ingress PE can deliver the IPv4 packets to the default egress PE when it does not obtain the address mapping rule for that IPv4 address block.

2. Address Mapping Rule Exchange

The address mapping rules generated at one PE device need to be sent to other PE devices. This process can be implemented through the routing layer. When an address mapping rule is generated locally, the PE device will convert it into a data structure and forward it to the IPv6 routing engine for transmission. In the opposite direction, upon receiving a routing announcement with an address mapping rule from a neighboring IPv6 router, the PE device extracts and stores it in its MR-DB.

To enable address mapping rule transmission at the routing layer, extensions to MP-BGP [RFC4760] or other protocols would be required. A typical approach is illustrated in [I-D.ietf-idr-mpbgp-extension-4map6]. It should be noted that address mapping rule exchange can be implemented by non-routing mechanisms as well, however, this is outside the scope of this document.

Based on the received mapping rule, the ingress PE can identify the appropriate egress PE (i.e., the Pref6(PE) to use as the destination prefix).

5.3. Packet Conversion and Transmission

In this framework, the forwarding layer of PE devices provides data forwarding capability to 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 ingress PE receives an IPv4 packet from its neighboring IPv4 network, it queries the local MR-DB which stores all the address mapping rules. If an address mapping rule for the IPv4 destination address is found, it will generate corresponding IPv6 source and destination addresses from the IPv4 addresses, following the procedure described in Section 4.1. This process complies with [RFC7915].

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

2. Encapsulation

The address mapping process for encapsulation follows the same procedure as translation: When the ingress PE receives an IPv4 packet from its neighboring IPv4 network, it queries the local MR-DB which stores all the address mapping rules, if an address mapping rule for the IPv4 destination address is found, the ingress PE will generate corresponding IPv6 source and destination addresses from the IPv4 addresses, following the procedure described in Section 4.1.

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

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

In summary, both translation and encapsulation rely on the same control-plane mechanisms (the MR-DB and address mapping rules) and impose identical requirements on the IPv6-only core network (forwarding based on the Pref6 part of the destination address).

Although this document illustrates the framework for a multi-domain IPv6-only network operated by multiple NPs, this framework is also applicable to an IPv6-only network operated by a single NP.

6. IPv6 Mapping Prefix Allocation

As shown in Section 4.1, IPv6 mapping prefixes are used by PEs to generate address mapping rules for any IPv4 address blocks. These prefixes are to be allocated from the Network Specific Prefix (NSP). An NSP refers to a dedicated IPv6 prefix (or a set of prefixes) assigned from NP’s available IPv6 address pool for IPv4 address mapping. For an IPv6-only network, one or more distinct IPv6 mapping prefixes are assigned to each PE device. For any PE device, its IPv6 mapping prefix is part of a larger and routable IPv6 address prefix previously assigned, if Pref6(PE) is allocated from a prefix that is already advertised with a next hop that reaches that PE (or its site), additional FIB entries in the IPv6 core may be avoided. For any IPv4-embedded IPv6 packets, a P device can forward them as

regular IPv6 packets without requiring a specific FIB entry for each IPv6 mapping prefix.

For NPs, having all IPv6 mapping prefixes share the same length during deployment will likely avoid the unnecessary processing cost and complexity caused by prefix length diversity.

7. Operational Considerations

One of the major purposes of this framework is to address how to support remaining IPv4 service delivery in a multi-domain IPv6-only network. When IPv4 packets are encapsulated in IPv6, the resulting packet is 40 bytes larger than the original; when translated, the IPv6 header is 20 bytes larger than the IPv4 header. In a multi-domain network traversing multiple ASes, MTU differences between domains are likely to cause silent packet drops or performance degradation - a fundamental operational concern for any IPv4-over-IPv6 framework. Therefore, network operators must handle MTU and fragmentation issues, for example, by configuring appropriate MSS clamping, ensuring consistent MTU across domains, or following the recommendations in Section 1.4 of [RFC7915] for general MTU/ fragmentation handling, and Sections 4.2 and 5.2 for specific ICMP translation handling.

8. Security Considerations

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

8.1. Authenticity and Integrity of Packets

In this framework, as the receiver of IPv4-embedded IPv6 packets, each egress PE assumes that all ingress PEs are legal and authorized to send IPv4-embedded IPv6 packets to it. After the egress PE receives IPv4-embedded IPv6 packets, it will convert them into IPv4 packets and forward them into the IPv4 Internet. If IPv6 packets cannot guarantee their authenticity or integrity, then there may be a spoofing attack. A malicious ingress PE could send IPv6 packets converted from IPv4 packets to attack an egress PE. Since the PEs in this framework are stateless, even when receiving large-volume traffic flows, they will not increase mapping session counts within the device like a stateful NAT device would, thus avoiding significant consequences. Even if no per-flow state exists, it should be acknowledged that in some extreme cases PEs can possibly be overwhelmed by bandwidth exhaustion, packet-rate/CPU exhaustion, MR-DB lookup pressure, or queue/buffer exhaustion. To mitigate these issues, measures such as rate limiting, ACLs between PEs, or provision validation can be applied for actual deployment.

8.2. Stateless IP/ICMP Translators

In this framework, the Stateless IP/ICMP Translation Algorithm is used to translates between IPv4 and IPv6 packet headers. As stated in Section 8 of [RFC7915], the use of stateless IP/ICMP translators does not introduce any new security issues beyond the security issues that are already present in the IPv4 and IPv6 protocols and in the routing protocols that are used to make the packets reach the translator. Other security considerations can be found in [RFC6052].

8.3. Issues Related to MP-BGP

The framework allows MP-BGP protocol as one approach to propagate mapping information over an IPv6-only network. However, MP-BGP has inherent vulnerabilities, such as route hijacking, where malicious alterations to routing announcements can redirect service traffic from its intended path or even steer it toward unauthorized destinations.

When the capability to advertise address mapping rules via BGP is introduced, attackers may alter the IPv6 mapping prefix within these rules, leading to improper delivery of IPv4 service traffic over an IPv6-only network. Such an attack differs from pre-existing vulnerabilities in that traffic could be forwarded to a remote target across an intervening network infrastructure (e.g., an IPv6 core), allowing an attack to potentially succeed more easily since less infrastructure needs to be compromised. To mitigate this risk, [I-D.ietf-sidrops-moa-profile] proposes an approach by leveraging RPKI [RFC6480] architecture to verify the authenticity of the address mapping rule associated with an IPv4 address block.

This framework proposes a default rule 0.0.0.0/0: Pref6(PE), which sends unknown IPv4 traffic (i.e. IPv4 traffic without definite IPv6 address mapping rule) to a "default egress PE". This approach has obvious implications, such as traffic attraction (posing a DoS concentration risk) and becoming a "catch-all" hijack target if rule distribution is compromised. To mitigate these issues, the default rule (0.0.0.0/0:Pref6(PE)) is OPTIONAL and is only used within a single administrative domain unless explicit bilateral policy exists. Operators can consider applying monitoring and rate-limiting/ACLs on the default egress PE, and avoid exporting the default rule across inter-domain boundaries.

9. IANA Considerations

This document has no IANA action.

10. 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.

11. 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

12. References

12.1. 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>.

[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>.

[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>.

12.2. 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-4map6-05, 3 December 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-mpbgp-extension-4map6-05>.

[I-D.ietf-sidrops-moa-profile]
Xie, C., Dong, G., Li, X., Huston, G., and D. Ma, "A Profile for Mapping Origin Authorizations (MOAs)", Work in Progress, Internet-Draft, draft-ietf-sidrops-moa-profile-03, 11 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-moa-profile-03>.

[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-07, 2 March 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-6mops-07>.

[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>.

[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>.

[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>.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <https://www.rfc-editor.org/info/rfc6480>.

[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>.

[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>.

[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>.

[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- 8045 Zurich
Switzerland
Email: thomas.graf@swisscom.com