v6ops Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

C. Ma
C. Xie
China Telecom

2 March 2026

Requirements for Provider Edge in IPv6-only Underlay Networks
draft-ma-v6ops-pe-ipv6only-reqs-00

Abstract

This document defines functional, protocol, and operational requirements for Provider Edge (PE) devices operating in a multi-domain network environment where the underlay is exclusively based on IPv6. These requirements ensure consistent service delivery, interoperability, and efficient operations across autonomous domains while supporting IPv4-as-a-Service (IPv4aaS).

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 3 September 2026.

Copyright Notice

Table of Contents

1. Introduction

The evolution towards IPv6-only underlay networks presents a compelling opportunity to simplify network architecture and reduce operational overhead. In large-scale service provider environments, this underlay often spans multiple administrative domains (e.g., different Autonomous Systems or organizational boundaries). A framework for such a multi-domain IPv6-only underlay, supporting services like IPv4-as-a-Service (IPv4aaS), is described in [I-D.ietf-v6ops-framework-md-ipv6only-underlay].

In this architecture, the Provider Edge (PE) device serves as the critical nodal point where customer service attachment intersects with the multi-domain IPv6-only underlay. Its role extends beyond traditional PE functions to include key responsibilities in inter-domain service routing, protocol-agnostic data encapsulation, and cross-domain service assurance.

This document specifies the requirements for PE devices operating in this specific context. These requirements ensure that PEs from different vendors and domains can interoperate seamlessly to deliver end-to-end services across a pure IPv6 underlay.

It should be noted that this document does not cover all the requirements for PE devices, it only introduces the parts directly related to IPv6-only. By specifying these requirements, this document aims to provide a common recommendation for equipment vendors and network operators, facilitating the development, deployment, and interoperability of multi-domain IPv6-only PE Routers. This will ultimately contribute to the successful establishment and operation of multi-domain IPv6-only network.

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 defined in this draft:

* Multi-domain Network: A network comprising two or more administratively separate domains (e.g., Autonomous Systems) interconnected to provide end-to-end services.

* IPv6-only Underlay: A network infrastructure that exclusively uses IPv6 for forwarding and control-plane protocols, with no native IPv4 forwarding path.

* Inter-domain PE: A PE device located at the boundary between two administrative domains, responsible for exchanging service information and forwarding traffic across domains.

* Intra-domain PE: A PE device located within a single administrative domain.

* IPv4-as-a-Service (IPv4aaS): A service model where IPv4 connectivity is provided to customers over an IPv6-only underlay network, using encapsulation or translation techniques.

* PE: Provider Edge, defined in [RFC4026]

3. Architectural Reference Model

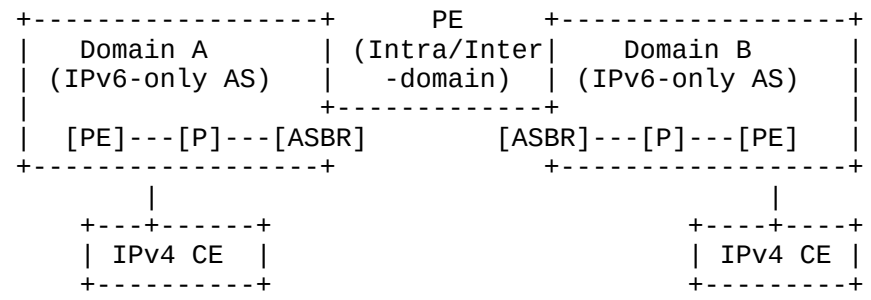

Figure 1: Multi-domain IPv6-only Underlay
with PE Service Attachment Points

In the reference model (Figure 1):

* Each domain operates an internal IPv6-only IGP (e.g., IS-ISv6 or OSPFv3) .

* Domains are interconnected via Inter-domain PEs (acting as ASBRs) using IPv6 EBGP sessions.

* Service layer routes (e.g., IPv4 VPN routes) are exchanged via MP-BGP over the IPv6 underlay.

* Customer IPv4 traffic is encapsulated within IPv6 (e.g., with an SRv6 header) for traversal across the multi-domain underlay.

4. Core Requirements for PEs

4.1. Inter-domain Routing Distribution Requirements

REQ-1: A PE device, when acting as an Inter-domain PE, MUST be able to exchange service layer routes (e.g., VPN-IPv4/VPN-IPv6 NLRIs [RFC4364] with PEs in other domains using MP-BGP sessions where the next-hop is an IPv6 address. This MUST be supported over IPv6-only underlay transport.

REQ-2: A PE device MUST support the necessary MP-BGP extensions and procedures for advertising IPv4-to-IPv6 (and vice versa) service mappings across domain boundaries. This capability is FUNDAMENTAL for enabling services like IPv4-as-a-Service (IPv4aaS) in an IPv6-only underlay. It allows a downstream PE to correctly associate a received IPv4 service prefix with its corresponding IPv6 underlay path identifier or tunnel endpoint. The specific encoding (e.g., a dedicated BGP attribute, extended community, or a new AFI/SAFI) and procedures for distributing these mappings MUST follow [I-D.ietf-idr-mpbgp-extension-4map6]. PE devices MUST be able to install forwarding state based on these received mappings.

REQ-3: A PE device MUST correctly process and enforce policies based on BGP Extended Community attributes, specifically Route Target, to identify and isolate service routes and their associated mappings belonging to different customers or domains.

REQ-4: A PE device MAY support additional mechanisms for inter-domain traffic steering and policy enforcement beyond basic connectivity. If supported, such mechanisms (e.g., BGP-based mechanisms for segment routing or other tunnel technologies) SHOULD operate transparently across the IPv6-only underlay. Crucially, the absence of any specific traffic steering mechanism MUST NOT break basic IPv4aaS service connectivity established through REQ-1 and REQ-2.

4.2. Data Plane Forwarding & Encapsulation Requirements

REQ-5: A PE device MUST support at least one IETF-standardized encapsulation method for carrying non-IPv6 customer traffic (e.g., IPv4, Ethernet) across the IPv6-only underlay. Possible mechanisms include, but are not limited to, IPv6-in-IPv6 tunneling [RFC2473], Generic Routing Encapsulation (GRE) over IPv6, MPLS-in-UDP/IPv6 [RFC7510], or SRv6 [RFC8986]. The choice of encapsulation MUST be derivable from the control-plane information (e.g., the service mapping advertised per REQ-2).

REQ-6: As an ingress PE, the device MUST be capable of encapsulating received customer packets (e.g., IPv4) with an appropriate IPv6 header, based on the forwarding state installed via the control plane (driven by REQ-2). As an egress PE, the device MUST be capable of decapsulating the outer IPv6 header to recover the original customer packet and forward it to the correct service instance.

REQ-7: Due to the increased packet size from encapsulation, a PE device MUST implement Path MTU Discovery for IPv6 [RFC8201]. It MUST either handle fragmentation appropriately or signal MTU issues back to the source (e.g., using ICMPv6 Packet Too Big messages) in a way that functions correctly across domain boundaries, considering the encapsulated path.

4.3. Operations & Management Requirements

REQ-8: A PE device MUST support standardized OAM mechanisms that operate end-to-end across multiple domains over the IPv6-only underlay. This includes, but is not limited to, IPv6-based Bidirectional Forwarding Detection (BFD)[RFC5880]and ICMPv6-based traceroute. The OAM packets MUST follow the same encapsulation path as data traffic.As a forwarding node in IPv6-only networks, PE router shall accept centralized policy scheduling from the network controller and implementing automated configuration through the NETCONF protocol/YANG model. Meanwhile, the PE supports syslog and SNMP Trap alarm mechanisms, enabling rapid fault location and security incident tracing through standardized logs.

REQ-9: A PE device MUST be able to map the Quality of Service (QoS) markings from the customer packet (e.g., IPv4 DSCP) to the encapsulating IPv6 header’s Traffic Class field (or equivalent), and ensure this policy is consistently applied and honored as the packet traverses different domains.

REQ-10: A PE device MUST provide robust isolation between traffic from different customers or services across the shared IPv6 underlay. This MUST be enforceable at inter-domain boundaries, typically implemented through the use of distinct VPN routing/forwarding instances or logically separated encapsulation identifiers derived from the control plane.

5. Relationship with Existing Frameworks

This document provides the concrete device-level requirements that implement the architecture described in the framework document [I-D.ietf-v6ops-framework-md-ipv6only-underlay]. It relies upon and extends the principles established in existing VPN specifications (e.g., [RFC4364], [RFC7432]) and places critical new requirements on MP-BGP for service mapping (REQ-2) to enable operation in a multi-domain, IPv6-only context.

6. Security Considerations

Security in a multi-domain IPv6-only environment introduces specific considerations for PE devices:

o BGP Session Security: All inter-domain BGP sessions MUST be secured using mechanisms such as IPsec [RFC4301]. Route Origin Authorization (ROA) validation via RPKI [RFC6810] SHOULD be implemented where feasible.

o Mapping Advertisement Integrity: The service mapping advertisements (REQ-2) are critical for correct forwarding. Mechanisms MUST be in place to ensure their integrity and authenticity, potentially leveraging existing BGP security measures.

o OAM Protection: OAM protocols MUST be rate-limited and protected from misuse that could lead to denial-of-service attacks across domains.

o Encapsulation Overhead and Inspection: The addition of encapsulation headers should be considered in network capacity planning and security monitoring. Techniques for deep packet inspection may need to be adapted to operate on encapsulated traffic within the underlay.

7. IANA Considerations

This document has no IANA actions.

8. Acknowledgements

The authors would like to thank the contributors and reviewers of this document for their valuable input and feedback.

9. References

9.1. Normative 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>.

[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>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <https://www.rfc-editor.org/info/rfc2473>.

[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>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/info/rfc4301>.

[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>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, DOI 10.17487/RFC6810, January 2013, <https://www.rfc-editor.org/info/rfc6810>.

[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>.

[RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", RFC 7510, DOI 10.17487/RFC7510, April 2015, <https://www.rfc-editor.org/info/rfc7510>.

[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>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

[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>.

9.2. Iormative References

[I-D.ietf-v6ops-framework-md-ipv6only-underlay] Xie, C., Ma, C., Li, X., Mishra, G. S., and T. Graf, "Framework for Multi-domain IPv6-only Underlay Network and IPv4-as-a-Service", Work in Progress, Internet-Draft, draft-ietf-v6ops-framework-md-ipv6only-underlay-19, 5 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-framework-md-ipv6only-underlay-19>.

Authors' Addresses

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

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