

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

E. Kline
Loon LLC
July 8, 2019

IPv6 Path MTU Option
draft-kline-6man-pmo-00

Abstract

This document describes an IPv6 Neighbor Discovery (ND) Path MTU Option (PMO) for inclusion in Router Advertisements (RAs). This allows some environments greater flexibility to support, for example, a higher MTU for on-link or intra-administrative-domain communications than for broader Internet communications.

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 <http://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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

This document describes an IPv6 Neighbor Discovery (ND) Path MTU Option (PMO) for inclusion in Router Advertisements (RAs). This allows some environments greater flexibility to support, for example, a higher MTU for on-link or intra-administrative-domain communications than for broader Internet communications.

TBD: Explain why extending [RFC4191](#) RIOs didn't look easy.

TBD: more discussion

2. Terminology

2.1. 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.

2.2. Summary of key terms

For the purposes of this document the following terms are used as described here.

2.2.1. Link MTU

The MTU of the link ([\[RFC4861\]](#)); alternatively, the MTU as it would be determined were no Path MTU Option ([Section 2.2.5](#)) present. This may be:

the value specified in an MTU Option ([Section 2.2.4](#)),

a value specified in a separate document that covers operating IP over a particular link type (e.g., [\[RFC2464\]](#)), or

a value derived by other means (e.g. administrative or a hint from a sub-IP layer mechanism).

The Link MTU MUST be the initial Path MTU used when transmitting to any link-local destination.

2.2.2. Path MTU

TBD: Path MTU

2.2.3. Path MTU Discovery

TBD: Path MTU Discovery (cite [[RFC8201](#)])

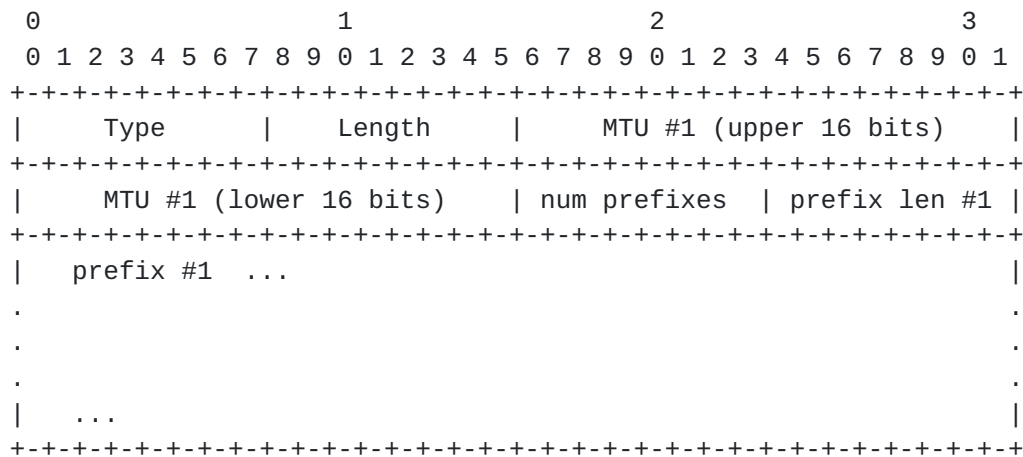
2.2.4. MTU Option

The MTU Option is defined in [[RFC4861](#)] [section 4.6.4](#). In this documented it may also be referred to as the Link MTU Option, in order to disambiguate it from this the Path MTU Option ([Section 2.2.5](#)).

2.2.5. Path MTU Option

The IPv6 ND Path MTU Option is described in this document. It provides more explicit signaling of the best initial Path MTU value for a given set of destinations when sending via the advertising router.

3. Path MTU Option Format



Fields:

Type TBD

Length The length of the option in units of 8 octets.

When not set, the receiving node MUST NOT make any assumptions of exclusive use of the specified Prefix, i.e. processing is unchanged from previous standards behavior.

Option contents The contents of the option as described below.

The Path MTU Option contents are encoded as a repeated sequence of:

4-octet MTU value

1-octet number of prefixes

sequence of prefixes to which this MTU applies

where each prefix is encoded as:

1-octet prefix length

variable length leading bits of prefix or IP address

Each sequence of octets representing a prefix uses only as many octets as required to by the prefix length (e.g. for a prefix length of 0: 0 octets are required, for prefix lengths 1 through 8: 1 octet is required, and so on).

The option is padded with zero-valued octets to the 8 octet boundary as given by the option length.

4. Option Processing Rules

Nodes compliant with this specification do not change their processing of RAs that contain no Path MTU Options. Additionally, for all Path MTU determination, an effective Path MTU learned via a Path MTU Discovery mechanism ([[RFC8201](#)]) MUST take precedence.

Any IPv6 link-local prefixes listed within a Path MTU Option MUST be ignored by the receiver and SHOULD be logged for review by an administrator.

Any MTU value lower than the IPv6 minimum MTU (i.e. 1280, [RFC8200](#) [section 5](#)), SHOULD be logged for administrator review as a

configuration error and MUST be treated by the receiver as though the IPv6 minimum MTU had been the encoded value.

The Link MTU MUST also be used for all on-link destinations, to maintain compatibility with existing behavior and expectations. For the same reason, the Link MTU SHOULD be used for destinations within any PIO prefix in the RA, even if the L bit is not set. As noted in [RFC5942], a destination may at some time be learned to be on-link, and this information may expire or be changed.

For all other destinations considered reachable via the option's advertising router, an initial Path MTU SHOULD be determined by first looking for a prefix that includes the destination in a Path MTU Option and using the corresponding MTU value. If no such prefix exists, the Link MTU SHOULD be assumed to be the default.

Note that as a matter of convenience a Path MTU Option may contain an entry for `::/0` even when the router lifetime is zero. This in no way indicates that the router will function as a default gateway. Rather, it may be used, for example, to apply a Path MTU to all prefixes listed in a set of RIOs.

5. Examples

TBD

6. Security Considerations

TBD

7. References

7.1. 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>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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>>.

- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [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>>.

7.2. Informative References

- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", [RFC 2464](#), DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [RFC5942] Singh, H., Beebe, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", [RFC 5942](#), DOI 10.17487/RFC5942, July 2010, <<https://www.rfc-editor.org/info/rfc5942>>.

Author's Address

Erik Kline
Loon LLC
1600 Amphitheatre Parkway
Mountain View, California 94043
US

Email: ek@loon.com

