

Implementation Guidelines for parsing IPv6 Extension Headers
draft-kampanakis-6man-ipv6-eh-parsing-01

Abstract

IPv6 is widely used on the internet today and is expected to be deployed more as more devices (i.e. home automation) get interconnected. The IPv6 header format allows for the use of Extension Headers (EH). EHs could be chained together with very few existing guidelines by the IPv6 protocol on how devices should parse them, which open room for security concerns and inconsistencies. This document presents guidelines for parsing IPv6 EHs with a goal of providing a common and consistent parsing methodology for IPv6 implementers among the industry.

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 February 6, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Parsing EH chains	3
4.	Parsing malformed EHs	5
5.	Other guidelines	6
6.	Security Considerations	7
7.	Updates	7
8.	Acknowledgements	8
9.	Normative References	8
	Author's Address	8

1. Introduction

As defined in [[RFC2460](#)], the IPv6 protocol was designed to have a constant header length and allows for the use of IPv6 Extension Headers (EH) which could be used to carry routing or other information for intermediary or end-devices. For example, Destination Option (DestOpt) EH are used by Mobile IPv6 end-hosts and Hop-by-Hop (HbH) EHs are used by intermediate routing devices. Multiple EHs can also be stacked together in an IPv6 header.

Because of the various possible combinations of EHs in an IPv6 packet, it is not clear to implementers how headers should be evaluated and parsed by intermediary and end-devices. [[RFC2460](#)] describes some IPv6 EH recommendations of the order and allowed occurrences of headers, but it does not provide other guidance on how EHs should be parsed. Experience has shown that, based on the receiving device vendor and operating system (OS), there can be inconsistencies in the receiver's behaviour when EHs are chained together in an IPv6 packet. This document presents how EHs in an IPv6 packet should be parsed in order to provide a consistent standard behaviour for IPv6 enabled intermediary (i.e. routers) or end-devices.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. Parsing EH chains

[[RFC2460](#)] defined various IPv6 EHs and other documents defined new EHs for various applications. Even though it includes some guidelines, [[RFC2460](#)] does not provide strong recommendations on the exact restrictions and order of EH headers present in an IPv6 Header. Thus, it is can be challenging for vendors of intermediate or end-devices to consistently parse and distinguish between which EH chains in an IPv6 Header are legitimate or not.

Since [[RFC2460](#)] does not limit the maximum EH chain size, [[RFC7112](#)] presents the implications of oversized IPv6 Header chains and how these should be fragmented and parsed by intermediate devices and hosts. Also, it defines the maximum allowed size of a EH chain.

Guideline 1: intermediate devices that process IPv6 header chains in order to enforce policies SHOULD be able to parse the IPv6 Header at

least up to the maximum header processing length supported by the device hardware acceleration in order to enforce policies or packet legitimacy. When intermediate devices are configured to match certain EHs, the EH SHOULD be matched regardless of its position in the EH chain unless the header is malformed and dropped according to guidelines in this document. If multiple EHs of the same type exist in the EH chain, all SHOULD be parsed by the intermediate device until matched in order to enforce the matching policy. When an EH chain exceeds the hardware acceleration processing limit the packets MAY be software processed in order to enforce policies. To protect against Denial of Service Denial of Service (DoS), intermediate devices MAY provide rate-limiting mechanisms that limit resources consumed by software processed packets. Intermediate devices that are not enforcing policies based by matching IPv6 EHs, might not parse EH chains other than HbH when present ([\[RFC7045\]](#)).

Guideline 2: From [Section 5 of \[RFC7112\]](#), when parsing an EH chain,

- o a host MUST drop non-initial fragments that include parts of an IPv6 EH chain and initial fragments whose last EH Next Header field is not an Upper Layer (UL) protocol (i.e. TCP) or the No Next Header. The host SHOULD send an ICMPv6 error message when dropping the fragment.
- o an intermediate device that processes IPv6 header chains in order to enforce policies MAY drop (if not configured to allow) non-initial fragments that include parts of an IPv6 EH chain or an initial fragment whose last EH Next Header field is not an Upper Layer (UL) protocol (i.e. TCP) or the No Next Header.

[Section 4 of \[RFC2460\]](#) defines that the HbH EH can only occur once in an IPv6 Header chain and has to be the first header.

Guideline 3: Thus, when parsing an EH chain,

- o an end-host MUST discard packets that have an HbH EH that is not first in the chain.
- o an intermediate device that processes IPv6 header chains in order to enforce policies MUST discard packets that have an HbH EH that is not first in the chain ([Section 2.2 of \[RFC7045\]](#)). If the device is not parsing the EH chain, it might not discard the packets.

Guideline 4: According to [Section 4.1 of \[RFC2460\]](#), when an end-host or an intermediate system that does not inspect IPv6 header chains in order to enforce policies processes an IPv6 packet with multiple DestOpt EHs,

- o if there is no RH EH in the packet, then the final destination host alone SHOULD process all the DestOpt EHs in the packet.
- o if there is a RH in the packet
 - * if the node is the final destination, then it SHOULD process all the DestOpt EHs in the packet.
 - * if the node is one of the destinations in the RH, it SHOULD only process the DestOpt EHs in the chain that are before the RH EH.

In general, intermediate devices that process IPv6 header chains in order to enforce policies should follow [Section 2 of \[RFC7045\]](#) in order to transmit IPv6 packets with EHs.

Note about EH order: [\[RFC2460\]](#) mentions

Each extension header should occur at most once, except for the Destination Options header which should occur at most twice (once before a Routing header and once before the upper-layer header).

But, later it says

IPv6 nodes must accept and attempt to process extension headers in any order and occurring any number of times in the same packet, except for the Hop-by-Hop Options header which is restricted to appear immediately after an IPv6 header only. Nonetheless, it is strongly advised that sources of IPv6 packets adhere to the above recommended order until and unless subsequent specifications revise that recommendation.

Thus, there is no other specific EH order end-hosts or intermediate devices can enforce following current specifications.

4. Parsing malformed EHs

IPv6 EHs that are malformed MUST be efficiently dropped. Malformed EHs could contain incorrect Hdr Ext Len or Opt Data Len fields.

Guideline 5: When parsing an EH chain, a host or an intermediate device that processes IPv6 header chains in order to enforce policies MUST discard packets that contain EHs with inaccurate Hdr Ext Len. To ensure that, while parsing the chain, each EH Hdr Ext Len SHOULD be used to determine the next EH in the chain. If the data of the last EH or the UL Header plus its data does not align with the original IPv6 Header Payload Length the packet MUST be dropped unless another

policy step already discarded it.

Especially for DestOpt and HbH EHs, they can carry a variable number of type-length-value (TLV) encoded Options. Each option contains an Opt Data Len field for the data length in the option. Malformed packets with Options that contain inaccurate length MUST be dropped by hosts and intermediate devices that are parsing these Options unless another policy step discarded the packet.

Guideline 6: When parsing EH Options, the Opt Data Len field points to the end of the Option parsed. If the Opt Data Len of the last Option in the EH does not align with the Hdr Ext Len of the EH that contains the Option, the packet MUST be discarded. The guidelines for actions based on the Option Type value described in [Section 4.2 of \[RFC2460\]](#) MUST still be followed. A host or intermediate device that is not parsing the Option ([\[RFC7112\]](#)) SHOULD NOT discard the packet.

The following algorithm shows how an end-host or intermediary device that enforces policies (ACLs, stateful inspection and firewalling) by matching on IPv6 EHs could be implementing Guideline 6 and 7:

```
while (IPv6 EH chain not parsed) {  
    [TODO: write pseudocode]  
}
```

5. Other guidelines

Guideline 7: Unknown EHs Since IPv6 Extension Headers and upper-layer protocols share the same namespace (the "Assigned Internet Protocol Numbers" registry/namespace), it is impossible to tell whether an unrecognized Next Header value refers to an IPv6 Extension Header to to an unrecognized upper-layer protocol, [\[RFC6564\]](#) notwithstanding. Thus, when parsing an IPv6 header chain:

- o As defined in [Section 4 of \[RFC2460\]](#), end-hosts should discard packets with unknown EH types and send an ICMP Parameter Problem message to the source of the packet, with an ICMP Code value of 1 ("unrecognized Next Header type encountered") and the ICMP Pointer field containing the offset of the unrecognized value within the original packet.
- o Intermediate devices that process IPv6 header chains in order to enforce policies MUST NOT attempt to continue parsing an IPv6 header chain after encountering an unrecognized Next Header value. Per [Section 2.1 of \[RFC7045\]](#) such systems MUST be configurable to allow such packets, but the default configuration MAY drop such

packets.

Guideline 8: End-hosts or intermediate systems that do not process IPv6 header chains in order to enforce policies MUST process an RH in an IPv6 packet only when the node's address is the same as the packet's IPv6 Destination Address ([Section 4.4 of \[RFC2460\]](#)).

Guideline 9: Unknown RH

- o End-hosts that receive packets with a RH with an unrecognised Routing Type value MUST ignore the RH if Segments Left is zero. If Segments Left is non-zero, the host MUST drop the packet and send an ICMP Parameter Problem ([Section 4.4 of \[RFC2460\]](#)).
- o Intermediate devices that process IPv6 header chains in order to enforce policies SHOULD forward (unless configured to drop) standardised and undeprecated RH ([Section 2.1 of \[RFC7045\]](#)).

Guideline 10: Intermediate or end-devices performing re-assembly MUST silently discard overlapping IPv6 fragments ([Section 4 of \[RFC5722\]](#)). More info on errors and the corresponding messages to be generated by a host or intermediate device performing re-assembly can be found in [Section 4.5 of \[RFC2460\]](#).

Guideline 11: As defined in [\[RFC6946\]](#), "atomic" IPv6 fragments SHOULD be "re-assembled" from the contents of that sole fragment. End-hosts and intermediary devices performing re-assembly SHOULD conform to [\[RFC6946\]](#).

6. Security Considerations

No new security exposures or issues are raised by this document. This document describes how IPv6 EHs should be parsed by intermediate and end-devices in a consistent manner. Guidelines presented in this document also leverage recommendations described in [\[RFC2460\]](#), [\[RFC6946\]](#), [\[RFC5722\]](#), [\[RFC7112\]](#) and [\[RFC7045\]](#).

Implementers following a common document for parsing and matching IPv6 EHs can ensure that no network policies are bypassed due to inconsistent processing of IPv6 EHs.

7. Updates

version -00: Initial submission.

version -01 updates:

- (1) Addressed comment about MAY NOT being ambiguous language.
- (2) Addressed Mike's comment about wording in various sections and guilines 4, 7 and 8.

8. Acknowledgements

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", [RFC 5722](#), December 2009.
- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", [RFC 6564](#), April 2012.
- [RFC6946] Gont, F., "Processing of IPv6 "Atomic" Fragments", [RFC 6946](#), May 2013.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", [RFC 7045](#), December 2013.
- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", [RFC 7112](#), January 2014.

Author's Address

Panos Kampanakis
Cisco Systems
170 West Tasman Dr.
San Jose, CA 95134
US

Email: pkampana@cisco.com

