A Uniform Format for IPv6 Extension Headers

Abstract

   In IPv6, optional internet-layer information is encoded in separate
   headers that may be placed between the IPv6 header and the transport-
   layer header. There are a small number of such extension headers
   currently defined. This document describes the issues that can arise
   when defining new extension headers and discusses the alternate
   extension mechanisms in IPv6. It also provides a common format for
   defining any new IPv6 extension headers, if they are needed.

Copyright Notice

Table of Contents

1. Introduction

   The base IPv6 standard [RFC2460] defines extension headers as an
   expansion mechanism to carry optional internet-layer information.
   Extension headers, with the exception of the Hop-by-Hop Options
   header, are not usually processed on intermediate nodes. However,
   several existing deployed IPv6 routers and several existing deployed
   IPv6 firewalls, in contradiction to [RFC2460], are capable of parsing
   past or ignoring all currently defined IPv6 extension headers (e.g.,
   to examine transport-layer header fields) at wire speed (e.g., by
   using custom Application-specific Integrated Circuits (ASICs) for
   packet processing). Hence, one must also consider that any new IPv6
   extension header will break IPv6 deployments that use these existing
   capabilities.

   Any IPv6 header or option that has hop-by-hop behavior, and is
   intended for general use in the public IPv6 Internet, could be
   subverted to create an attack on IPv6 routers that process packets
   containing such a header or option. Reports from the field indicate
   that some IP routers deployed within the global Internet are
   configured either to ignore the presence of headers with hop-by-hop

   behavior or to drop packets containing headers with hop-by-hop
   behavior.

2. Conventions Used in This Document

   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 [RFC2119].

3. Applicability

   The base IPv6 standard [RFC2460] allows the use of both extension
   headers and destination options in order to encode optional
   destination information in an IPv6 packet. The use of destination
   options to encode this information provides more flexible handling
   characteristics and better backward compatibility than using
   extension headers. Because of this, implementations SHOULD use
   destination options as the preferred mechanism for encoding optional
   destination information, and use a new extension header only if
   destination options do not satisfy their needs. The request for
   creation of a new IPv6 extension header MUST be accompanied by a
   specific explanation of why destination options could not be used to
   convey this information.

   The base IPv6 standard [RFC2460] defines 3 extension headers (i.e.,
   Routing header, Destination Options header, Hop-by-Hop Options
   header) to be used for any new IPv6 options. The same standard only
   allows the creation of new extension headers in limited circumstances
   ([RFC2460], Section 4.6).

   As noted above, the use of any option with hop-by-hop behavior can be
   problematic in the global public Internet. New IPv6 extension
   header(s) having hop-by-hop behavior MUST NOT be created or
   specified. New options for the existing Hop-by-Hop Header SHOULD NOT
   be created or specified unless no alternative solution is feasible.
   Any proposal to create a new option for the existing Hop-by-Hop
   Header MUST include a detailed explanation of why the hop-by-hop
   behavior is absolutely essential in the document proposing the new
   option with hop-by-hop behavior.

   The use of IPv6 Destination Options to encode information provides
   more flexible handling characteristics and better backward
   compatibility than using a new extension header. Because of this,
   new optional information to be sent SHOULD be encoded in a new option
   for the existing IPv6 Destination Options header.

   Mindful of the need for compatibility with existing IPv6 deployments,
   new IPv6 extension headers MUST NOT be created or specified, unless
   no existing IPv6 extension header can be used by specifying a new
   option for that existing IPv6 extension header. Any proposal to
   create or specify a new IPv6 extension header MUST include a detailed
   technical explanation of why no existing IPv6 extension header can be
   used in the document proposing the new IPv6 extension header.

4. Proposed IPv6 Extension Header Format

   Any IPv6 extension headers defined in the future, keeping in mind the
   restrictions specified in Section 3 and also the restrictions
   specified in [RFC2460], MUST use the consistent format defined in
   Figure 1. This minimizes breakage in intermediate nodes that examine
   these extension headers.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Next Header  |  Hdr Ext Len  |                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
|                                                               |
.                                                               .
.                      Header Specific Data                     .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
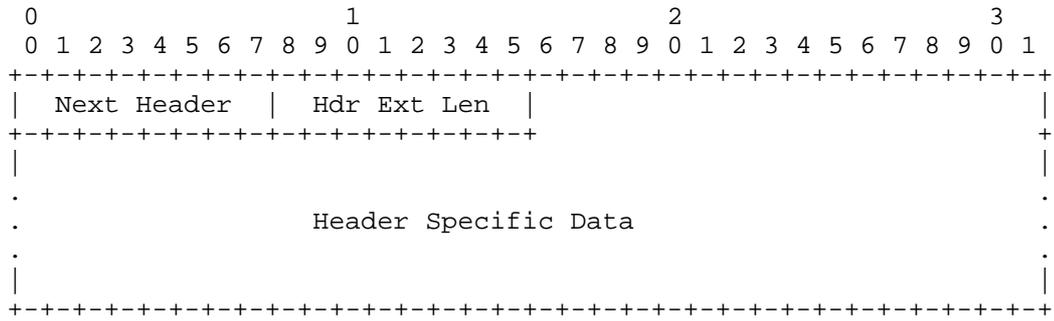
   Next Header   8-bit selector. Identifies the type of header
                           immediately following the extension header.
                           Uses the same values as the IPv4 Protocol field
                           [IANA_IP_PARAM].

   Hdr Ext Len   8-bit unsigned integer. Length of the extension
                           header in 8-octet units, not including the first
                           8 octets.

   Header Specific   Variable length. Fields specific to the
   Data   extension header.

                     Figure 1: Extension Header Layout

5. Backward Compatibility

   The scheme proposed in this document is not intended to be backward
   compatible with all the currently defined IPv6 extension headers. It
   applies only to newly defined extension headers. Specifically, the

   fragment header predates this document and does not follow the format
   proposed in this document.

6. Future Work

   This document proposes one step in easing the inspection of extension
   headers by middleboxes. There is further work required in this area.
   Some issues that are left unresolved beyond this document include:

   o There can be an arbitrary number of extension headers.

   o Extension headers must be processed in the order they appear.

   o Extension headers may alter the processing of the payload itself,
      and hence the packet may not be processed properly without
      knowledge of said header.

7. Security Considerations

   This document proposes a standard format for the IPv6 extension
   headers that minimizes breakage at intermediate nodes that inspect
   but do not understand the contents of these headers. Intermediate
   nodes, such as firewalls, that skip over unknown headers might end up
   allowing the setup of a covert channel from the outside of the
   firewall to the inside using the data field(s) of the unknown
   extension headers.

8. Acknowledgements

   The authors would like to thank Albert Manfredi, Bob Hinden, Brian
   Carpenter, Erik Nordmark, Hemant Singh, Lars Westberg, Markku Savela,
   Tatuya Jinmei, Thomas Narten, Vishwas Manral, Alfred Hoenes, Joel
   Halpern, Ran Atkinson, Steven Blake, Jari Arkko, Kathleen Moriarty,
   Stephen Farrell, Ralph Droms, Sean Turner, and Adrian Farrel for
   their reviews and suggestions that made this document better.

9. Normative References

   [IANA_IP_PARAM] IANA, "IP Parameters",
                   <http://www.iana.org/assignments/ip-parameters>.

   [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.

Authors' Addresses

   Suresh Krishnan
   Ericsson
   8400 Decarie Blvd.
   Town of Mount Royal, QC
   Canada
   Phone: +1 514 345 7900 x42871
   EMail: suresh.krishnan@ericsson.com

   James Woodyatt
   Apple Inc.
   1 Infinite Loop
   Cupertino, CA 95014
   US
   EMail: jhw@apple.com

   Erik Kline
   Google
   Mori Tower 26F
   Roppongi 6-10-1
   Minato ku
   Tokyo 106-6126
   Japan
   Phone: +81 3-6384-9635
   EMail: ek@google.com

   James Hoagland
   Symantec Corporation
   350 Ellis St.
   Mountain View, CA 94043
   US
   EMail: Jim_Hoagland@symantec.com
   URI: http://symantec.com/

   Manav Bhatia
   Alcatel-Lucent
   Bangalore
   India
   EMail: manav.bhatia@alcatel-lucent.com