6man
Internet-Draft
Updates: 2460, 2780 (if approved)
Intended status: Standards Track
Expires: August 26, 2013

B. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies Co., Ltd

February 22, 2013

Transmission of IPv6 Extension Headers
draft-carpenter-6man-ext-transmit-02

Abstract

Various IPv6 extension headers have been defined since the IPv6 standard was first published. This document updates RFC 2460 to clarify how intermediate nodes should deal with such extension headers and with any that are defined in future. It also specifies how extension headers should be registered by IANA, with a corresponding minor update to RFC 2780.

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 August 26, 2013.

Table of Contents

1. Introduction and Problem Statement

In IPv6, an extension header is any header that follows the initial 40 bytes of the packet and precedes the upper layer header (which might be a transport header, an ICMPv6 header, or a notional "No Next Header").

An initial set of IPv6 extension headers was defined by [RFC2460], which also described how they should be handled by intermediate nodes, with the exception of the hop-by-hop options header:

"...extension headers are not examined or processed by any node along a packet's delivery path, until the packet reaches the node (or each of the set of nodes, in the case of multicast) identified in the Destination Address field of the IPv6 header."

This provision allowed for the addition of new extension headers, since it means that forwarding nodes should be completely transparent to them. Thus, new extension headers could be introduced progressively, used only by hosts that have been updated to create and interpret them. The extension header mechanism is an important part of the IPv6 architecture, and several new extension headers have been defined since RFC 2460.

Unfortunately, experience has showed that the network is not transparent to these headers. The main reason for this is that by design, some firewalls attempt to inspect the transport header or payload. This means that they need to traverse the chain of extension headers, if present, until they find the transport header (or an encrypted payload). Unfortunately, because not all IPv6 extension headers follow a uniform TLV format, this process is clumsy and requires knowledge of each extension header's format.

The process is potentially slow as well as clumsy, possibly precluding its use in nodes attempting to process packets at line speed. The present document does not intend to solve this problem, which is caused by the fundamental architecture of IPv6 extension headers. This document focuses on clarifying how the header chain should be traversed in the current IPv6 architecture.

If they encounter an unrecognised extension header type, some firewalls treat the packet as suspect and drop it. Unfortunately, it is an established fact that several widely used firewalls do not recognise some or all of the extension headers defined since RFC 2460. It has also been observed that certain firewalls do not even handle all the extension headers in RFC 2460, including the fragment header [I-D.taylor-v6ops-fragdrop], causing fundamental problems of connectivity. This applies in particular to firewalls that attempt

to inspect packets statelessly at very high speed, since they cannot take the time to reassemble fragmented packets, especially when under a denial of service attack.

Other types of middlebox, such as load balancers or packet classifiers, might also fail in the presence of extension headers that they do not recognise.

A contributory factor to this problem is that, because extension headers are numbered out of the existing IP Protocol Number space, there is no collected list of them. For this reason, it is hard for an implementor to quickly identify the full set of defined extension headers. An implementor who consults only RFC 2460 will miss all extension headers defined subsequently.

This combination of circumstances creates a "Catch-22" situation [Heller] for the deployment of any newly designed extension header. It cannot be widely deployed, because existing firewalls will render large parts of the Internet opaque to it. However, most firewalls will not be updated to allow the new header to pass until it has been proved safe and useful on the open Internet, which is impossible until the firewalls have been updated.

The uniform TLV format now defined for extension headers [RFC6564] will improve the situation, but only for future extensions. Some tricky and potentially malicious cases will be avoided by forbidding very long chains of extension headers that need to be fragmented [I-D.ietf-6man-oversized-header-chain]. This will alleviate concerns that stateless firewalls cannot handle a complete header chain as required by the present document.

However, these changes are insufficient to correct the underlying problem. The present document clarifies that the above requirement from RFC 2460 applies to all types of node that forward IPv6 packets and to all extension headers defined now and in the future. It also requests IANA to create a subsidiary registry that clearly identifies extension header types, and updates RFC 2780 accordingly. Fundamental changes to the IPv6 extension header architecture are out of scope for this document.

Also, Hop-by-Hop options are not handled by many high speed routers, or are processed only on a slow path. This document also updates the requirements for processing the Hop-by-Hop options header to make them more realistic.

1.1. 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 [RFC2119].

2. Requirement to Transmit Extension Headers

2.1. All Extension Headers

Any node along an IPv6 packet's path, which forwards it for any reason, SHOULD do so regardless of any extension headers that are present, as described in RFC 2460. Exceptionally, if this node is designed to examine extension headers for any reason, such as firewalling, it MUST recognise and deal appropriately with all defined IPv6 extension header types. The list of currently defined extension header types is maintained by IANA (see Section 4) and implementors are advised to check this list regularly for updates.

RFC 2460 requires destination hosts to discard packets containing unrecognised extension headers. However, intermediate forwarding nodes MUST NOT do this by default, since that might cause them to inadvertently discard traffic using a recently defined extension header, not yet recognised by the intermediate node.

As mentioned above, firewalls that discard packets containing extension headers are known to cause connectivity failures and deployment problems. Therefore, it is important that firewalls can parse all defined IPv6 extension headers and are able to behave according to the above requirements. If a firewall discards a packet containing a defined IPv6 extension header, it MUST be the result of a configurable firewall policy, and not just the result of a failure to recognise such a header. This means that the discard policy for each defined type of extension header MUST be individually configurable. The default configuration SHOULD allow all defined extension headers. Firewalls MUST be configurable to allow packets containing unrecognised extension headers, but such packets MUST be dropped by default.

The IPv6 Routing Header Types 0 and 1 have been deprecated and SHOULD NOT be used. However, as specified in [RFC5095], this does not mean that the IPv6 Routing Header can be unconditionally dropped by forwarding nodes. Packets containing undeprecated Routing Headers SHOULD be forwarded by default. At the time of writing, these include Type 2 [RFC6275], Type 3 [RFC6554], and Types 253 and 254 [RFC4727]. Others may be defined in future.

2.2. Hop-by-Hop Options

The IPv6 Hop-by-Hop Options header SHOULD be processed by intermediate nodes as described in [RFC2460]. However, it is to be expected that high performance routers will either ignore it, or assign packets containing it to a slow processing path. Designers planning to use a Hop-by-Hop option need to be aware of this likely behaviour.

As a reminder, in RFC 2460, it is stated that the Hop-by-Hop Options header, if present, must be first.

3. Security Considerations

Firewall devices MUST conform to the requirements in the previous section in order to respect the IPv6 extension header architecture. In particular, packets containing specific extension headers are only to be discarded as a result of a configurable policy.

When new extension headers are defined in the future, those implementing and configuring firewalls will need to take account of them. It is to be expected that this process will be slow. Until it is complete, the new extension will fail in some parts of the Internet. This aspect needs to be considered when deciding to standardise a new extension.

4. IANA Considerations

IANA is requested to clearly mark in the Assigned Internet Protocol Numbers registry those values which are also IPv6 Extension Header types, for example by adding an extra column to indicate this. This will also apply to any IPv6 Extension Header types defined in the future.

Additionally, IANA is requested to replace the existing empty IPv6 Next Header Types registry by an IPv6 Extension Header Types registry. It will contain only those protocol numbers which are also marked as IPv6 Extension Header types in the Assigned Internet Protocol Numbers registry. The initial list will be as follows:

- 0, Hop-by-Hop Options, [RFC2460]
- 43, Routing, [RFC2460], [RFC5095]
- 44, Fragment, [RFC2460]
- 50, Encapsulating Security Payload, [RFC4303]
- 51, Authentication, [RFC4302]

- o 60, Destination Options, [RFC2460]
- o 135, MIPv6, [RFC6275]
- o 139, HIP, [RFC5201]
- o 140, shim6, [RFC5533]

The references to the IPv6 Next Header field in [RFC2780] are to be interpreted as also applying to the IPv6 Extension Header field.

5. Acknowledgements

This document was triggered by mailing list discussions including John Leslie, Stefan Marksteiner and others. Valuable comments and contributions were made by Dominique Barthel, Lorenzo Colitti, Fernando Gont, Bob Hinden, Ray Hunter, Suresh Krishnan, Marc Lampo, Michael Richardson, Dave Thaler, Joe Touch, and others.

Brian Carpenter was a visitor at the Computer Laboratory, Cambridge University during part of this work.

This document was produced using the xml2rfc tool [RFC2629].

6. Change log [RFC Editor: Please remove]

draft-carpenter-6man-ext-transmission-02: clarifications following WG comments, recalibrated firewall requirements, 2013-02-22.

draft-carpenter-6man-ext-transmission-01: feedback at IETF85: clarify scope and impact on firewalls, discuss line-speed processing and lack of uniform TLV format, added references, restructured IANA considerations, 2012-11-13.

draft-carpenter-6man-ext-transmission-00: original version, 2012-08-14.

7. References

7.1. 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.

[RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For

Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.

[RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", RFC 4727, November 2006.

[RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, December 2007.

[RFC5201] Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol", RFC 5201, April 2008.

[RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.

[RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

[RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", RFC 6564, April 2012.

7.2. Informative References

[Heller] Heller, J., "Catch-22", Simon and Schuster , 1961.

[I-D.ietf-6man-oversized-header-chain] Gont, F. and V. Manral, "Security and Interoperability Implications of Oversized IPv6 Header Chains", draft-ietf-6man-oversized-header-chain-02 (work in progress), November 2012.

[I-D.taylor-v6ops-fragdrop] Jaeggli, J., Colitti, L., Kumari, W., Vyncke, E., Kaeo, M., and T. Taylor, "Why Operators Filter Fragments and What It Implies", draft-taylor-v6ops-fragdrop-00 (work in progress), October 2012.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, March 2012.

Authors' Addresses

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

Email: jiangsheng@huawei.com