

IPv6 Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 27, 2008

S. Krishnan
Ericsson
February 24, 2008

**The case against Hop-by-Hop options
draft-krishnan-ipv6-hopbyhop-02**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 27, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

The Hop-by-Hop option header is a type of IPv6 extension header that has been defined in the IPv6 protocol specification. The contents of this header need to be processed by every node along the path of an IPv6 datagram. This draft highlights the characteristics of this extension header which make it prone to Denial of Service attacks and proposes solutions to minimize such attacks.

Table of Contents

1.	Introduction	3
1.1.	Conventions used in this document	3
2.	Details of the attack	4
2.1.	Effects of the attack	4
3.	Proposed Solutions	5
3.1.	Deprecation	5
3.2.	Skipping	5
3.3.	Rate limiting	5
4.	Impact on deployed IPv6 nodes	6
5.	Security Considerations	7
6.	IANA Considerations	8
7.	Normative References	9
	Author's Address	10
	Intellectual Property and Copyright Statements	11

1. Introduction

The IPv6 base specification [[RFC2460](#)] defines the hop-by-hop extension header. This extension header carries the options which need to be processed by every node along the path of the datagram. Certain characteristics of the specification make it especially vulnerable to Denial of Service attacks. The characteristics are:

- o All the ipv6 nodes on the path need to process the options in this header
- o The option TLVs in the hop-by-hop options header need to be processed in order
- o A sub range of option types in this header will not cause any errors even if the node does not recognize them.
- o There is no restriction as to how many occurrences of an option type can be present in the hop-by-hop header.

This document details a low bandwidth Denial of Service attack on ipv6 routers/hosts using the hop-by-hop options extension header and possible ways of mitigating these attacks.

1.1. 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](#)].

2. Details of the attack

The denial of service attack can be carried out by forming an IP datagram with a large number of TLV encoded options with random option type identifiers in the hop-by-hop options header. The option type is a 8 bit field with special meaning attached to the three most significant bits. The attack is most effective when all the nodes in the path are affected, meaning we do not want any node to drop the packet and send ICMP errors regarding unrecognized options. If the two most significant bits are cleared(0), the receiving node will silently ignore the option if it does not recognize the option type. The third most significant bit is used to denote whether the option data can change en-route. If the bit is set to 1 the option data can change en route. The attack is equally effective whether or not an IPSec Authentication Header(AH) treats the option data as zero valued octets. Hence we can include this bit in generating option types. The acceptable option types would be laid out like below

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+ - - - - -
| Option Type | Opt Data Len | Option Data
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+ - - - - -
|0 0 x x x x x x|.....|.....
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+ - - - - -
```

Figure 1: Option type layout

Since the option types 0(0x00) and 1(0x01) are reserved for the Pad1 and PadN options in [[RFC2460](#)] we exclude these from the acceptable range as well. So we choose the option type identifiers for each of these options to be in the range 0x02-0x63. More option types defined by other RFCs can be excluded from the attack as and when they are allocated by the IANA. Examples are Tunnel Encapsulation limit (0x04) and Router Alert (0x05).

2.1. Effects of the attack

The attack can be used to cripple the routers by attacking the control processor rather than the forwarding plane. Since the control traffic, like the routing protocols, shares the same resources with this traffic, this kind of attack may be hard to control. On routers having separate Control and Forwarding elements only the Control traffic would be affected. For routers whose the Control and Forwarding elements are fused together this would lead to problems with forwarding packets as well.

3. Proposed Solutions

There are at least three possible solutions to handle the DoS attack mentioned in this draft. The first one is to get rid of the feature altogether and prevent the attacks. The second one is to limit the attacks to nodes that need to process hop-by-hop options. The third is to let the attacks occur, but limit the damage.

3.1. Deprecation

The first solution is to deprecate hop-by-hop options from the IPv6 specification and to stop allocation of any new ones. The existing hop-by-hop options MAY be grandfathered but new ones MUST NOT be allocated. This allows existing protocols depending on hop-by-hop options to continue working, but discourages the development of new solutions based on hop-by-hop options.

3.2. Skipping

This option allows nodes to skip over the hop-by-hop extension header without processing any of the options contained in the header. If a node receives an IPv6 datagram with a hop-by-hop header, and it does not support any hop-by-hop options at all, it can just skip over the header.

3.3. Rate limiting

A less severe (and less effective) solution is to simply rate limit packets with hop-by-hop option headers and start dropping them randomly when the CPU load becomes very high. While this solution is very simple and has no impact on deployed IPv6 nodes, it is sub-optimal. A legitimate packet with a hop-by-hop option header has the same probability of being dropped as an attack packet. Implementing the solution proposed in this draft does not preclude the use of rate limiting. In fact it gives a legitimate packet a lower probability of being dropped, since most of the obvious attack traffic would have been dropped by the receiving algorithm.

4. Impact on deployed IPv6 nodes

The proposed changes can affect all currently IPv6 nodes which need to send and receive packets with hop-by-hop options. If the deprecation option is chosen, the IPv6 stack on both sending and receiving nodes needs to be modified to not send or receive hop-by-hop options. In addition, transit nodes need to be modified as well in order to not inspect these options.

5. Security Considerations

This document highlights the possible security issues with the IPv6 hop-by-hop option header specified in [[RFC2460](#)] which can lead to denial of service attacks and suggests some changes to reduce the effect of the DoS attacks.

6. IANA Considerations

This requests IANA to stop allocation of new entries for IPv6 hop-by-hop option types.

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

Author's Address

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Email: suresh.krishnan@ericsson.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

