

IPv6 Working Group  
Internet-Draft  
Updates: [2460](#) (if approved)  
Intended status: Standards Track  
Expires: March 26, 2009

S. Krishnan  
Ericsson  
September 22, 2008

**Handling of overlapping IPv6 fragments**  
**draft-ietf-6man-overlap-fragment-00**

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 March 26, 2009.

Abstract

The fragmentation and reassembly algorithm specified in the base IPv6 specification allows fragments to overlap. This document demonstrates the security issues with allowing overlapping fragments and updates the IPv6 specification to explicitly forbid overlapping fragments.

Table of Contents

|                      |  |                   |
|----------------------|--|-------------------|
| <a href="#">1.</a>   | Introduction . . . . .                                   | <a href="#">3</a> |
| <a href="#">1.1.</a> | Conventions used in this document . . . . .              | <a href="#">3</a> |
| <a href="#">2.</a>   | Overlapping fragments . . . . .                          | <a href="#">3</a> |
| <a href="#">3.</a>   | The attack . . . . .                                     | <a href="#">3</a> |
| <a href="#">4.</a>   | Recommendation . . . . .                                 | <a href="#">5</a> |
| <a href="#">5.</a>   | Security Considerations . . . . .                        | <a href="#">5</a> |
| <a href="#">6.</a>   | IANA Considerations . . . . .                            | <a href="#">6</a> |
| <a href="#">7.</a>   | Normative References . . . . .                           | <a href="#">6</a> |
|                      | Author's Address . . . . .                               | <a href="#">6</a> |
|                      | Intellectual Property and Copyright Statements . . . . . | <a href="#">7</a> |



## 1. Introduction

Fragmentation is used in IPv6 when the IPv6 packet will not fit inside the path MTU to its destination. When fragmentation is performed an IPv6 node uses a fragment header as specified in [section 4.5](#) of the IPv6 base specification [[RFC2460](#)] to break down the datagram into smaller fragments that will fit in the path MTU. The destination node receives these fragments and reassembles them. The algorithm specified for fragmentation in [[RFC2460](#)] does not prevent the fragments from overlapping, and this can lead to some security issues with firewalls [[RFC4942](#)]. This document explores the issues that can be caused by overlapping fragments.

### 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. Overlapping fragments

Commonly used firewalls use the algorithm specified in [[RFC1858](#)] to weed out malicious packets that try to overwrite parts of the transport layer header to bypass inbound connection checks. [[RFC1858](#)] prevents an overlapping fragment attack on an upper layer protocol (in this case TCP) by recommending that packets with fragment offset 1 be dropped. While this works well for IPv4 fragments, it will not work for IPv6 fragments. This is because the fragmentable part of the IPv6 packet can contain extension headers before the TCP header, making this check less effective.

## 3. The attack

This attack describes how a malicious node can bypass a firewall using overlapping fragments. Consider a sufficiently large IPv6 packet that needs to be fragmented.

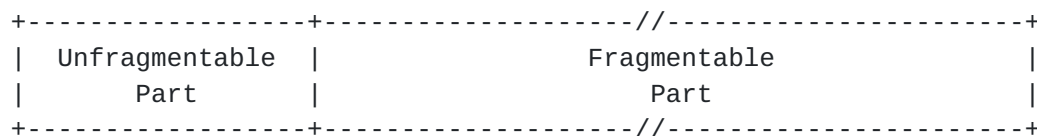


Figure 1: Large IPv6 packet

This packet is split into several fragments by the sender so that the



packet can fit inside the path MTU. Let's say the packet is split into two fragments.

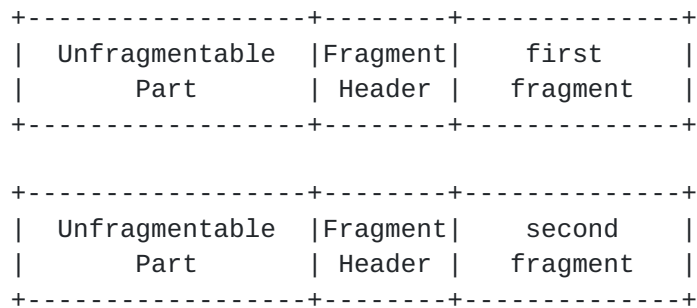


Figure 2: Fragmented IPv6 packet

Consider the first fragment. Let's say it contains a destination options header (DOH) 80 octets long and is followed by a TCP header.

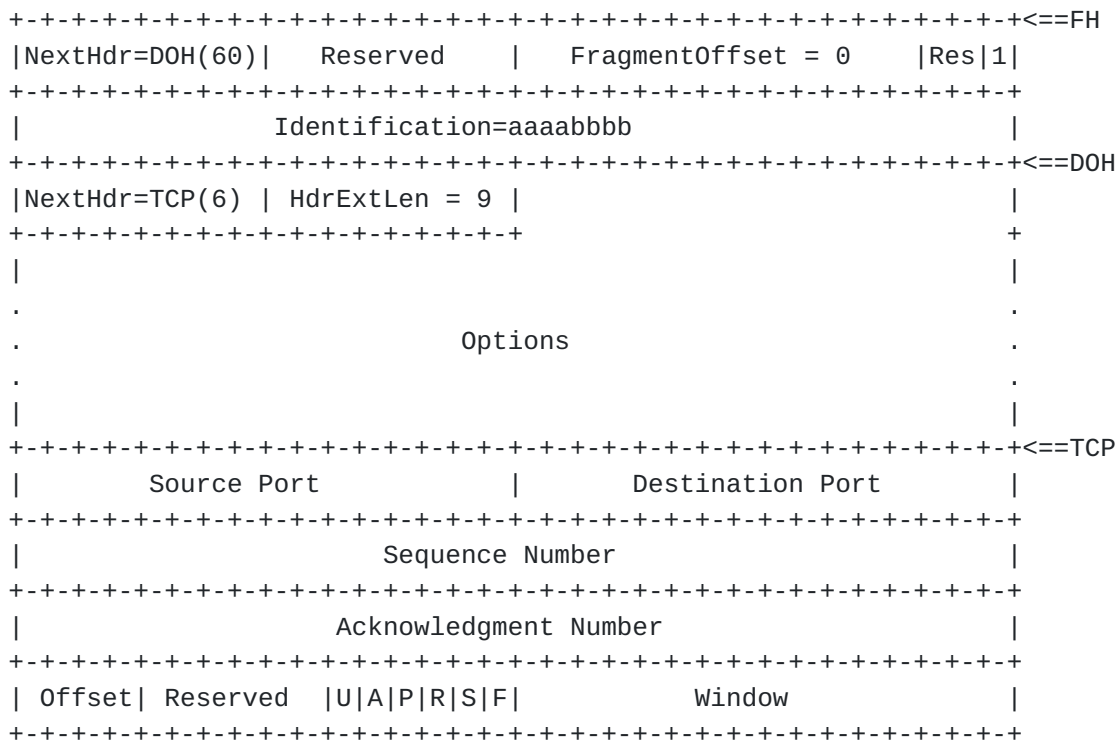


Figure 3: First Fragment

The TCP header has the following values of the flags S(YN)=0 and A(CK)=1. This makes an inspecting stateful firewall think that it is a response packet for a connection request initiated from the trusted side of the firewall. Hence it will allow the fragment to pass. It will also let the following fragments with the same Fragment Identification value in the fragment header to pass through.







## **6.    IANA Considerations**

This document does not require any action from the IANA.

## **7.    Normative References**

- [RFC1858]    Ziemba, G., Reed, D., and P. Traina, "Security Considerations for IP Fragment Filtering", [RFC 1858](#), October 1995.
- [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.
- [RFC4942]    Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/ Co-existence Security Considerations", [RFC 4942](#), September 2007.

### Author's Address

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

Email: [suresh.krishnan@ericsson.com](mailto: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](mailto:ietf-ipr@ietf.org).

