Network Working Group
Internet-Draft
Expires: March 1, 2009

V. Manral
IPInfusion Inc.

Operational issues with Tiny Fragments in IPv6
draft-manral-v6ops-tiny-fragments-issues-03

Status of this Memo

Copyright Notice

Abstract

IPv6 fragmentation allows fragments to be sent only by the source of a packet. The Fragment header is used by an IPv6 source to send a packet larger than would fit in the path MTU to its destination.

Firewalls generally use 5-tuples to filter out packets. However there are cases where fragmentation can be used to disguise TCP packets from IP filters used in routers and hosts. This document specifies

where tiny fragments can be issues.

Requirements Language

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].

1. Problem Statement

With many IP implementations it is possible to impose a fragment small enough to force some of a packet's Upper Layer e.g. TCP header fields into the second fragment.

This can cause all middlebox's like firewall and NAT-PT which expect the fields header information in the first fragment to not work properly.

Though the NAT Behave draft, states that NAT box should reassemble the packets, a lot of new issues can result. Keeping state could result in easy DoS attacks. Besides the jury is still out about how many NAT boxes do reassembly.

All policy based devices where packets are forwarded or sent on a tunnel based on some policy are also affected.

2. Issues with Firewalls

There are different types of firewalls and state can be created in these firewalls through different methods. Independent of the adopted method, firewalls typically look at five parameters of the traffic arriving at the firewalls:

- Source IP address
- Destination IP address
- Protocol type
- Source port number
- Destination port number

Based on these parameters, firewalls usually decide whether to allow the traffic or to drop the packets.

However in cases where the first fragment does not have the upper layer header information, the firewall is not able to get the port

information and other upper layer information, thus allowing the packets to be sent to the protected side.

This can lead to attacks to the network and the firewall not being able to block such an attack.

3. Issues with NAT-PT

NAT-PT [RFC2766] assumes that for NAPT-PT operation the ports are visible to the translator. However if the Upper Layer Header is not there in the first fragment. This causes the visibility ot the port to be lost. This can cause the translation process to fail.

When the translator gets a tiny IPv6 fragment which has to be translated to an IPv4 packet. The translator will have to reassemble the packets as the IPv4 non last fragment needs to have a datagram size of 68 octets atleast.

STD 5, RFC 791 states:

> Every internet module must be able to forward a datagram of 68 octets without further fragmentation. This is because an internet header may be up to 60 octets, and the minimum fragment is 8 octets.

4. Issues with Policy Boxes

Tiny Fragments could cause issues to Policy boxes which look further inside the packet, to make decisions.

For IPsec Security Policy Database (SPD) specifies what services are to be offered to IP datagrams and in what fashion. The draft [RFC2401bis] states:

> "Non-initial" vs "Initial" Fragments
>
> Throughout this document, the phrase "non-initial" fragments is used to mean fragments that do not contain all of the selector values that may be needed for access control. And the phrase "initial" fragment is used to mean a fragment that contains all the selector values needed for access control.

However, it should be noted that for IPv6, which fragment contains the Next Layer Protocol and ports (or ICMP message type/code or Mobility Header type) will depend on the kind and number of extension headers present.

Having tiny fragments could mean that none of the fragments would be the Initial Fragment. So any access control/ tunneling based on that may not work unless reassembly is done, or extra state like next

Header and previous header length remaining are kept across fragments.

5. Proposed solutions to the problem

a. Impose a minimum packet size for the non-last fragments. If a fragment of a lesser size is received, the packet is treated as a malformed packet and is discarded.

b. Reassemble all the fragments of the packet, translate the header fields and, glean out relevent information and then pass the original fragments ahead after modifying the relevent fields.

c. Reassemble all the fragments of the packet till we have the header fields of the upper layer , glean out relevent information and then pass the original fragments ahead after modifying the relevent fields.

d. If upper layer protocol present then the header must be there in the first fragment.

The above is just a first summary and the proposals are expected to change as the draft matures.

6. Issues with fragment size of Minimum MTU

The minimum fragment size of the non last fragment could be specified to be 1280 octets, the minimum link MTU [RFC2460].

However if the IPv6 packet has to be further tunnelled the packet may have to be fragmented. To prevent such a case a minimum packet size of the non-last fragment should be less then 1280.

7. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

8. Security Considerations

This draft outlines security issues arising if "Tiny Fragments" are sent. This draft raises no new security issues.

9. Acknowledgements

This draft borrows text heavily from draft-ietf-mip6-firewalls-03.txt and RFC1858. Thanks to Brian Carpenter, Pekka Savola, Stig Venaas,Fred Baker, Pyda Srisuresh, Senthil Sivakumar and Radhakrishnan.S for the helpful discussion.

10. References

10.1 Normative References

[RFC2460] Deering & Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC2460, December 1998

[RFC2766] Tsirtsis & Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC2766, February 2000

[RFC2401bis] Kent & Seo, "Security Architecture for the Internet Protocol", Work in Progress, September, 2005

10.2 Informative References

[RFC1858] Ziemba, Reed & Traina , "Security Considerations - IP Fragment Filtering", RFC1858, October 1995

Authors' Addresses

Vishwas Manral
IPInfusion Inc,
41, Ground Floor,
5th Cross Road,
Off 8th Main Raod,
Vasanth Nagar,
Bangalore
India

Phone: +91-80-4113-1268
Email: vishwas@ipinfusion.com

Full Copyright Statement

Intellectual Property

Acknowledgement

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