Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2013

F. Templin, Ed.
Boeing Research & Technology

July 03, 2012

IPv6 Network Fragmentation for Tunnels
draft-generic-6man-tunfrag-00.txt

Abstract

IPv6 intentionally deprecates fragmentation by routers in the network. An exception to this requirement occurs when there is an IPv6/IPv4 protocol translator in the path, where in-the-network fragmentation may be unavoidable when the IPv4 network path includes a restricting link. Unfortunately, recent investigations have shown that IPv6 Path MTU Discovery (PMTUD) interacts poorly with tunnels to the point that tunneled packets can be best accommodated only when the tunnel ingress is permitted to perform fragmentation. This document therefore updates the IPv6 protocol specification to enable in-the-network fragmentation by routers that configure tunnel endpoints.

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 January 4, 2013.

Table of Contents

1. Introduction

IPv6 intentionally deprecates fragmentation by routers in the network. An exception to this requirement occurs when there is an IPv6/IPv4 protocol translator in the path, where in-the-network fragmentation may be unavoidable when the IPv4 network path includes a restricting link. Unfortunately, recent investigations have shown that IPv6 Path MTU Discovery (PMTUD) interacts poorly with tunnels [I-D.generic-v6ops-tunmtu] to the point that tunneled packets can be best accommodated only when the tunnel ingress is permitted to perform fragmentation. This document therefore asserts an additional case in which in-the-network IPv6 fragmentation is permitted.

2. Problem Statement

The current "Internet cell size" is effectively 1500 bytes (i.e., the minimum MTU configured by the vast majority of links in the Internet) and should therefore also be the minimum MTU assigned to tunnels. However, due to issues with PMTUD this size can only be accommodated when the tunnel ingress is permitted to perform fragmentation. The tunnel ingress can perform fragmentation on the outer packet following encapsulation and can instead (or in addition) perform "tunnel fragmentation" via an encapsulation mid-layer inserted between the inner and outer header. In both cases reassembly would be performed by the tunnel egress.

Unfortunately, the tunnel ingress may not always know the size of the reassembly buffer configured by the egress, and the burden for reassembly on the egress may be excessive - especially if the egress must service many destinations. The third alternative therefore is to permit the tunnel ingress to perform fragmentation on the inner packet before encapsulation in which case reassembly would be performed by the final destination. This document therefore updates the IPv6 protocol specification [RFC2460] to enable in-the-network fragmentation by tunnel routers.

3. IPv6 Updates

Section 4.5 of [RFC2460] includes the clause:

"(Note: unlike IPv4, fragmentation in IPv6 is performed only by source nodes, not by routers along a packet's delivery path -- see section 5.)"

This text is already ignored by IPv6/IPv4 translators which must perform fragmentation when the IPv4 path beyond the translator

includes a restricting link. This document proposes that this text also be relaxed in the case of tunnels that must perform fragmentation.

Section 5 of [RFC2460] states:

"In response to an IPv6 packet that is sent to an IPv4 destination (i.e., a packet that undergoes translation from IPv6 to IPv4), the originating IPv6 node may receive an ICMP Packet Too Big message reporting a Next-Hop MTU less than 1280. In that case, the IPv6 node is not required to reduce the size of subsequent packets to less than 1280, but must include a Fragment header in those packets so that the IPv6-to-IPv4 translating router can obtain a suitable Identification value to use in resulting IPv4 fragments. Note that this means the payload may have to be reduced to 1232 octets (1280 minus 40 for the IPv6 header and 8 for the Fragment header), and smaller still if additional extension headers are used."

This document proposes to add the following text immediately after the above:

"In response to an IPv6 packet that is sent to an IPv6 destination that is located beyond a tunnel, the originating IPv6 node may receive an ICMP Packet Too Big message reporting a Next-Hop MTU less than 1280. In that case, the IPv6 node is not required to reduce the size of subsequent packets, but must include a Fragment header in any packets that are larger than this MTU value but no larger than 1500 bytes. Note that the node is permitted to continue to send packets larger than 1500 bytes without including a Fragment header, but should implement [RFC4821] to ensure that the packets are reaching the final destination."

An example tunnel protocol that invokes this new clause appears in: [I-D.templin-intarea-seal].

4. IANA Considerations

There are no IANA considerations for this document.

5. Security Considerations

The security considerations for [RFC2460] apply also to this document.

6. Acknowledgments

This method was inspired through discussion on the IETF v6ops and NANOG mailing lists in the May/June 2012 timeframe.

7. References

7.1. Normative References

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

7.2. Informative References

[I-D.generic-v6ops-tunmtu]
Templin, F., "Operational Issues with Tunnel Maximum Transmission Unit (MTU)", draft-generic-v6ops-tunmtu-08 (work in progress), June 2012.

[I-D.templin-intarea-seal]
Templin, F., "The Subnetwork Encapsulation and Adaptation Layer (SEAL)", draft-templin-intarea-seal-42 (work in progress), December 2011.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
USA

Email: fltemplin@acm.org