

TCP and UDP over IPv6 Jumbograms

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

1. Overview

IPv6 supports datagrams larger than 65535 bytes long, often referred to as jumbograms, through use of the Jumbo Payload hop-by-hop option [Deering95]. The UDP protocol has a 16-bit length field that keeps it from being able to make use of jumbograms, and though TCP does not have a length field, both the MSS option and the Urgent field are constrained by 16-bits. This document describes some simple changes that can be made to allow TCP and UDP to make use of IPv6 jumbograms.

2. UDP Jumbograms

To allow UDP to make use of jumbograms, either the UDP length field needs to be extended, or it needs to be ignored. Since the size of the field can't be changed, a length of zero is used to indicate that it is to be ignored, and the length in the "pseudo-header" is to be used to determine the true length of the UDP header plus data. This works because UDP length field includes the UDP header, so the minimum valid value for this field is 8.

When sending a UDP packet, if and only if the length of the UDP header plus data is greater than 65,535, set the length field in the UDP header to zero.

Note 1: The length used in the "pseudo-header" for computing the UDP checksum is always the true length of the UDP header plus data, NOT zero [RFC-1883, Section 8.1].

Note 2: An IPv6 packet that carries a UDP packet of length greater than 65,535 will necessarily carry a Jumbo Payload option in a Hop-by-Hop Options header [RFC1883, Section 4.3]). The length field in the Jumbo Payload option contains the length of the IP packet excluding the IPv6 header, that is, it contains the length of all extension headers present plus the UDP header plus the UDP data. The length field in the IPv6 header contains zero to indicate the presence of the Jumbo Payload option.

If a UDP packet is received with a length field of zero, the length of the UDP packet is computed from the length field in the Jumbo Payload option minus the length of all extension headers present between the IPv6 header and the UDP header.

3. TCP Jumbograms

Because there is no length field in the TCP header, there is nothing limiting the length of an individual TCP packet. However, the MSS value that is negotiated at the beginning of the connection limits the largest TCP packet that can be sent, and the Urgent Pointer cannot reference data beyond 65535 bytes.

3.1 TCP MSS

When determining what MSS value to send, if the MTU of the directly attached interface is greater than 65535, then set the MSS value to 65535.

When an MSS value of 65535 is received, it is to be treated as infinity. MTU discovery code, starting with the MTU of the outgoing interface, will be used to determine the actual MSS.

3.2 TCP Urgent Pointer

The Urgent Pointer problem could be fixed by adding a TCP Urgent Pointer Option. However, since it is unlikely that applications using jumbograms will also use Urgent Pointers, a less intrusive change similar to the MSS change will suffice.

When a TCP packet is to be sent with an Urgent Pointer (i.e., the URG bit set), first calculate the offset from the Sequence Number to the Urgent Pointer. If the offset is less than 65535, fill in the Urgent field and continue with the normal TCP processing. If the offset is greater than 65535, and the offset is greater than or equal to the length of the TCP data, fill in the Urgent Pointer with 65535 and continue with the normal TCP processing. Otherwise, the TCP packet must be split into two pieces. The first piece contains data up to, but not including the data pointed to by the Urgent Pointer, and the

Urgent field is set to 65535 to indicate that the Urgent Pointer is beyond the end of this packet. The second piece can then be sent with the Urgent field set normally.

Note: The first piece does not have to include all of the data up to the Urgent Pointer. It can be shorter, just as long as it ends within 65534 bytes of the Urgent Pointer, so that the offset to the Urgent Pointer in the second piece will be less than 65535 bytes.

For TCP input processing, when a TCP packet is received with the URG bit set and an Urgent field of 65535, the Urgent Pointer is calculated using an offset equal to the length of the TCP data, rather than the offset in the Urgent field.

It should also be noted that though the TCP window is only 16-bits, larger windows can be used through use of the TCP Window Scale option [Jacobson92].

4. Security Considerations

There are no known security issues involved in these changes.

5. References

[Jacobson92] Jacobson, V., R. Braden and D. Borman, "TCP Extensions for High Performance", RFC 1323, LBL, ISI and Cray Research, May 1992.

[Deering95] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, Xerox PARC and Ipsilon Networks, December 1995.

Author's Address

David A. Borman
Berkeley Software Design, Inc.
4719 Weston Hills Drive
Eagan, MN 55123
Phone: (612) 405-8194
Mailing List: ipng@sunroof.Eng.Sun.COM
Email: dab@bsd1.com

