



Tech Info Library

Open Transport 1.1 and Fragmenting TCP/IP Packets (3/96)

Revised: 3/28/96
Security: Everyone

Open Transport 1.1 and Fragmenting TCP/IP Packets (3/96)

=====
Article Created: 25 March 1996
Article Reviewed/Updated: 28 March 1996

TOPIC -----

I have a network administrator commenting that Open Transport 1.1, which comes with the System 7.5 Update 2.0, will not allow the fragmenting of TCP/IP packets.

MacTCP 2.0.6 does not set the Don't Fragment (DF) bit and allows packet fragmentation, however, Open Transport 1.1 does set the DF bit. This may be an issue because some TCP/IP devices reject this. Can you confirm or deny this as an issue without traces?

DISCUSSION -----

Open Transport supports a new IP feature called MTU Path Discovery. The MTU Path Discovery algorithm does set the DF bit on all packets as long as the minimum datagram size (576) is not the current MTU. The DF bit is set on all packets so that in the event that the route changes, which could cause the MTU to change, the sender will receive an ICMP "packet to large" datagram from a router and adjust the MTU. The only time the DF bit is not set is when the minimum (576) size datagram is being used.

We are not aware of any problems using MTU Path Discovery or setting the DF bit with any IP implementations.

Here is a summary of how this all fits together from rfc 1191:

1. Introduction

When one IP host has a large amount of data to send to another host, the data is transmitted as a series of IP datagrams. It is usually preferable that these datagrams be of the largest size that does not require fragmentation anywhere along the path from the source to the destination. (For the case against fragmentation, see [5].) This datagram size is referred to as the Path MTU

(PMTU), and it is equal to the minimum of the MTUs of each hop in the path. A shortcoming of the current Internet protocol suite is the lack of a standard mechanism for a host to discover the PMTU of an arbitrary path.

NOTE: The Path MTU is what in [1] is called the "Effective MTU for sending" (EMTU_S). A PMTU is associated with a path, which is a particular combination of IP source and destination address and perhaps a Type-of-service (TOS).

The current practice [1] is to use the lesser of 576 and the first-hop MTU as the PMTU for any destination that is not connected to the same network or subnet as the source. In many cases, this results in the use of smaller datagrams than necessary, because many paths have a PMTU greater than 576. A host sending datagrams much smaller than the Path MTU allows is wasting Internet resources and probably getting sub-optimal throughput. Furthermore, current practice does not prevent fragmentation in all cases, since there are some paths whose PMTU is less than 576.

2. Protocol overview

In this memo, we describe a technique for using the Don't Fragment (DF) bit in the IP header to dynamically discover the PMTU of a path. The basic idea is that a source host initially assumes that the PMTU of a path is the (known) MTU of its first hop, and sends all datagrams on that path with the DF bit set. If any of the datagrams are too large to be forwarded without fragmentation by some router along the path, that router will discard them and return ICMP Destination Unreachable messages with a code meaning "fragmentation needed and DF set" [7]. Upon receipt of such a message (henceforth called a "Datagram Too Big" message), the source host reduces its assumed PMTU for the path.

The PMTU discovery process ends when the host's estimate of the PMTU is low enough that its datagrams can be delivered without fragmentation. Or, the host may elect to end the discovery process by ceasing to set the DF bit in the datagram headers; it may do so, for example, because it is willing to have datagrams fragmented in some circumstances. Normally, the host continues to set DF in all datagrams, so that if the route changes and the new PMTU is lower, it will be discovered.

Unfortunately, the Datagram Too Big message, as currently specified, does not report the MTU of the hop for which the rejected datagram was too big, so the source host cannot tell exactly how much to reduce its assumed PMTU. To remedy this, we propose that a currently unused header field in the Datagram Too Big message be used to report the MTU of the constricting hop. This is the only change specified for routers in support of PMTU Discovery.

This article was published in the Information Alley on 27 March 1996.

Article Change History:

28 Mar 1996 - Added Information Alley notation.

Copyright 1996, Apple Computer, Inc.

