

Network Working Group
Request for Comments: 1783
Updates: 1350
Category: Standards Track

G. Malkin
Xylogics, Inc.
A. Harkin
Hewlett Packard Co.
March 1995

TFTP Blocksize Option

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.

Abstract

The Trivial File Transfer Protocol [1] is a simple, lock-step, file transfer protocol which allows a client to get or put a file onto a remote host. One of its primary uses is the booting of diskless nodes on a Local Area Network. TFTP is used because it is very simple to implement in a small node's limited ROM space. However, the choice of a 512-byte blocksize is not the most efficient for use on a LAN whose MTU may 1500 bytes or greater.

This document describes a TFTP option which allows the client and server to negotiate a blocksize more applicable to the network medium. The TFTP Option Extension mechanism is described in [2].

Blocksize Option Specification

The TFTP Read Request or Write Request packet is modified to include the blocksize option as follows:

```
+-----+---~~~+---+---+---~~~+---+---+---~~~+---+---+---~~~+---+
|  opc  |filename| 0 |  mode  | 0 | blksize| 0 | #octets| 0 |
+-----+---~~~+---+---+---~~~+---+---+---~~~+---+---+---~~~+---+
```

opc

The opcode field contains either a 1, for Read Requests, or 2, for Write Requests, as defined in [1].

filename

The name of the file to be read or written, as defined in [1]. This is a NULL-terminated field.

mode

The mode of the file transfer: "netascii", "octet", or "mail", as defined in [1]. This is a NULL-terminated field.

blksize

The Blocksize option, "blksize" (case insensitive). This is a NULL-terminated field.

#octets

The number of octets in a block, specified in ASCII. Valid values range between "8" and "65464" octets, inclusive. This is a NULL-terminated field.

For example:

```
+-----+-----+---+-----+---+-----+---+-----+---+
|  1   | foobar | 0 | binary | 0 | blksize| 0 |  1432 | 0 |
+-----+-----+---+-----+---+-----+---+-----+---+
```


is a Read Request, for the file named "foobar", in binary transfer mode, with a block size of 1432 bytes (Ethernet MTU, less the UDP and IP header lengths).

If the server is willing to accept the blocksize option, it sends an Option Acknowledgment (OACK) to the client. The specified value must be less than or equal to the value specified by the client. The client must then either use the size specified in the OACK, or send an ERROR packet, with error code 8, to terminate the transfer.

The rules for determining the final packet are unchanged from [1]. The reception of a data packet with a data length less than the negotiated blocksize is the final packet. If the blocksize is greater than the size of the packet, the first packet is the final packet. If amount of data to be transferred is an integral multiple of the blocksize, an extra data packet containing no data is sent to end the transfer.

Proof of Concept

Performance tests were run on the prototype implementation using a variety of block sizes. The tests were run on a lightly loaded Ethernet, between two HP-UX 9000, in "octet" mode, on 2.25MB files. The average (5x) transfer times for paths with (g-time) and without (n-time) a intermediate gateway are graphed as follows:



The comparisons between transfer times (without a gateway) between the standard 512-byte blocksize and the negotiated blocksize are:

1024	2x	-32%
1432	2.8x	-42%
2048	4x	-54%
4096	8x	-71%
8192	16x	-80%

As was anticipated, the transfer time decreases with an increase in blocksize. The reason for the reduction in time is the reduction in the number of packets sent. For example, by increasing the blocksize from 512 bytes to 1024 bytes, not only are the number of data packets halved, but the number of acknowledgement packets is also halved (along with the number of times the data transmitter must wait for an ACK). A secondary effect is the efficiency gained by reducing the per-packet framing and processing overhead.

Of course, if the blocksize exceeds the path MTU, IP fragmentation and reassembly will begin to add more overhead. This will be more noticeable the greater the number of gateways in the path.

Security Considerations

Security issues are not discussed in this memo.

References

- [1] Sollins, K., "The TFTP Protocol (Revision 2)", STD 33, RFC 1350, MIT, July 1992.
- [2] Malkin, G., and A. Harkin, "TFTP Option Extension", RFC 1782, Xylogics, Inc., Hewlett Packard Co., March 1995.

Authors' Addresses

Gary Scott Malkin
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803

Phone: (617) 272-8140
EMail: gmalkin@xylogics.com

Art Harkin
Internet Services Project
Information Networks Division
19420 Homestead Road MS 43LN
Cupertino, CA 95014

Phone: (408) 447-3755
EMail: ash@cup.hp.com

