draft-masotta-tftpexts-windowsize-opt-00.txt
Intended status: Proposed Standard
INTERNET-DRAFT
Expires: Aug 07, 2012

Patrick Masotta/Vercot

March 2012

# TFTP Windowsize Option

## 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 lock-step schema is not the most efficient for use on a LAN.

This document describes a TFTP option which allows the client and server to negotiate a windowsize of consecutive blocks to send as an alternative for replacing the lock-step single block schema. The TFTP Option Extension mechanism is described in [2].



## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Distribution of this draft is unlimited. It is intended to become the new BCP 42 obsoleting RFC 6195. Comments should be sent to the DNS Extensions Working Group mailing list <dnsext@ietf.org>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Windowsize Option Specification

The TFTP Read Request or Write Request packet is modified to include the blocksize option as follows. Note that all fields except "opc" are NULL-terminated.

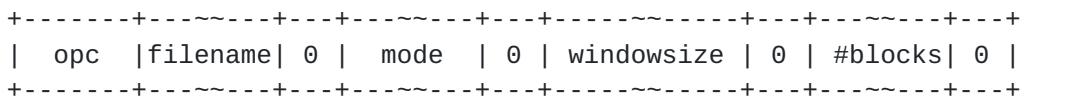

opc
The opcode field contains either a 1, for Read Requests, or 2, for Write Requests, as defined in [1].

filename
The name of the file to be read or written, as defined in [1].

mode
The mode of the file transfer: "netascii", "octet", or "mail", as defined in [1].

windowsize
The Windowsize option, "windowsize" (case in-sensitive).

#blocks
The number of blocks in a window, specified in ASCII. Valid values range between "1" and "65464" blocks, inclusive. The windowsize refers to the number of consecutives blocks transmited before stop and wait for the reception of the ack of the last block transmited.

For example:

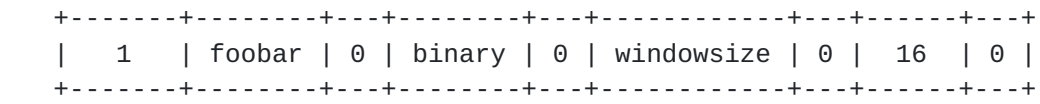

is a Read Request, for the file named "foobar", in binary transfer mode, with a window size of 16 blocks (as blocksize is not defined the 512 Bytes per block default applies).

If the server is willing to accept the windowsize 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 window with a number of blocks length less than the negotiated windowsize is the final window. If the windowsize is greater than the amount of data to be transfered, the first window is the final window. If the amount of data to be transfered is an integral multiple of the windowsize, 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 windowsizes and a fixed blocksize of 1456 bytes. The tests were run on a lightly loaded Gigabit Ethernet, between two Toshiba Tecra Core 2 Duo 2.2 Ghz, in "octet" mode, on 180MB files.

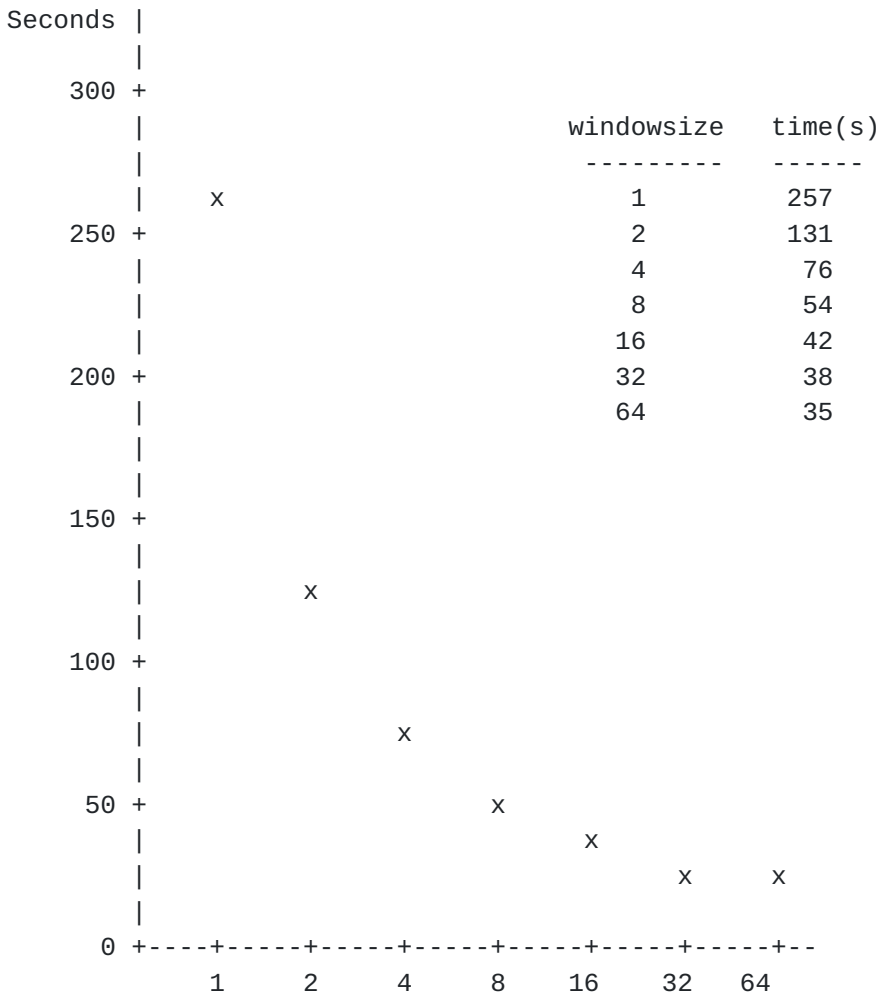


| windowsize | time(s) |
|---|---|
| 1 | 257 |
| 2 | 131 |
| 4 | 76 |
| 8 | 54 |
| 16 | 42 |
| 32 | 38 |
| 64 | 35 |

The comparisons between transfer times (without a gateway) between the standard lock-step schema and the negotiated windowsizes are:

| | |
|---|---|
| 1 | -0% |
| 2 | -49% |
| 4 | -70% |
| 8 | -79% |
| 16 | -84% |
| 32 | -85% |
| 64 | -86% |

As was anticipated, the transfer time decreases with the use of a windowed schema. The reason for the reduction in time is the reduction in the number of the required synchronous acknowledgement exchanged.

## Error Handling

In case of an error detection the whole windowsize window is retransmited.

## Security Considerations

The basic TFTP protocol has no security mechanism. This is why it has no rename, delete, or file overwrite capabilities. This document does not add any security to TFTP; however, the specified extensions do not add any additional security risks.

## References

[1] Sollins, K., "The TFTP Protocol (Revision 2)", Request for Comments 1350 (STD 33), October 1992.

[2] Malkin, G., Harkin, A., "TFTP Option Extension", RFC 1782 March 1995.

## Authors' Addresses

Patrick Masotta
300 W 11th Avenue, #9A
Denver, CO 80204

EMail: masotta@vercot.com

## Copyright and IPR Provisions