INTERNET-DRAFT

Patrick Masotta/Serva
Feb 2014

Intended status: Standard Track
Expires: Aug 15, 2014

TFTP Windowsize Option

draft-masotta-tftpexts-windowsize-opt-09.txt

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 early stages of nodes booting from a Local Area Network. TFTP has been always used because it is very simple to implement. 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 single block lock-step 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.

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

Introduction

While virtually unused for internet transfers today, TFTP is still massively used in network boot/installation scenarios (EFI included). The protocol inherently low transfer rate has been so far partially mitigated by the use of the blocksize negotiated extension [3]. This way the original 512 byte blocks are in practice replaced on Ethernet environments by blocks no larger than 1468 Bytes to avoid IP block fragmentation. This strategy results insufficient when transferring big files i.e. the initial ramdisk of Linux distributions or the PE images used in network installations by Microsoft WDS/MDT/SCCM. Considering TFTP looks today far from extinction this draft formally presents a negotiated natural extension that produces TFTP transfer rates comparable to the ones achieved today by modern file transfer protocols.

Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [5] and indicate requirement levels for compliant implementations.

Windowsize Option Specification

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

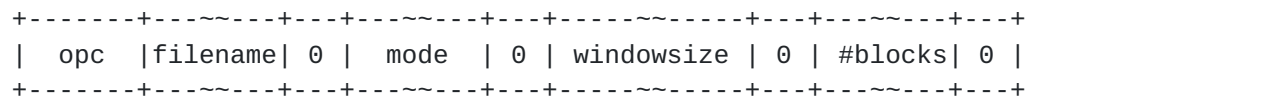

opc
: The opcode field contains either a 1, for Read Requests, or 2, for Write Requests, as defined in [1].

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 "65535" blocks, inclusive. The

windowsize refers to the number of consecutives blocks transmited before stop and wait for the reception of the acknowledgment of the last block transmited.

For example:

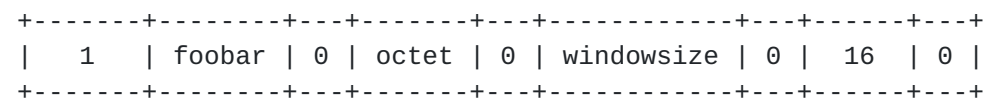

is a Read Request, for the file named "foobar", in octet transfer mode, with a window-size of 16 blocks (as option blocksize is not negotiated in this case, the 512 Bytes per block default applies).

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] and [3].
The reception of a data window with a number of blocks less than the negotiated windowsize is the final window. If the windowsize is greater than the amount of data to be transferred, the first window is the final window.

Congestion and Error Control

From a congestion control standpoint while the number of blocks in a window does not represent a threat, the rate at which TFTP UDP datagrams are sent SHOULD follow the congestion control guidelines in Section 3.1 of RFC 5405 [4].

From an error control standpoint while RFC 1350 [1] and subsequent updates do not specify a maximum number of retries for datagram retransmissions, implementations SHOULD always impose an appropriate threshold on error recovery attempts, after which a transfer SHOULD always be preventively aborted.

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, transfering a 180 MByte file.

The comparison of 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%

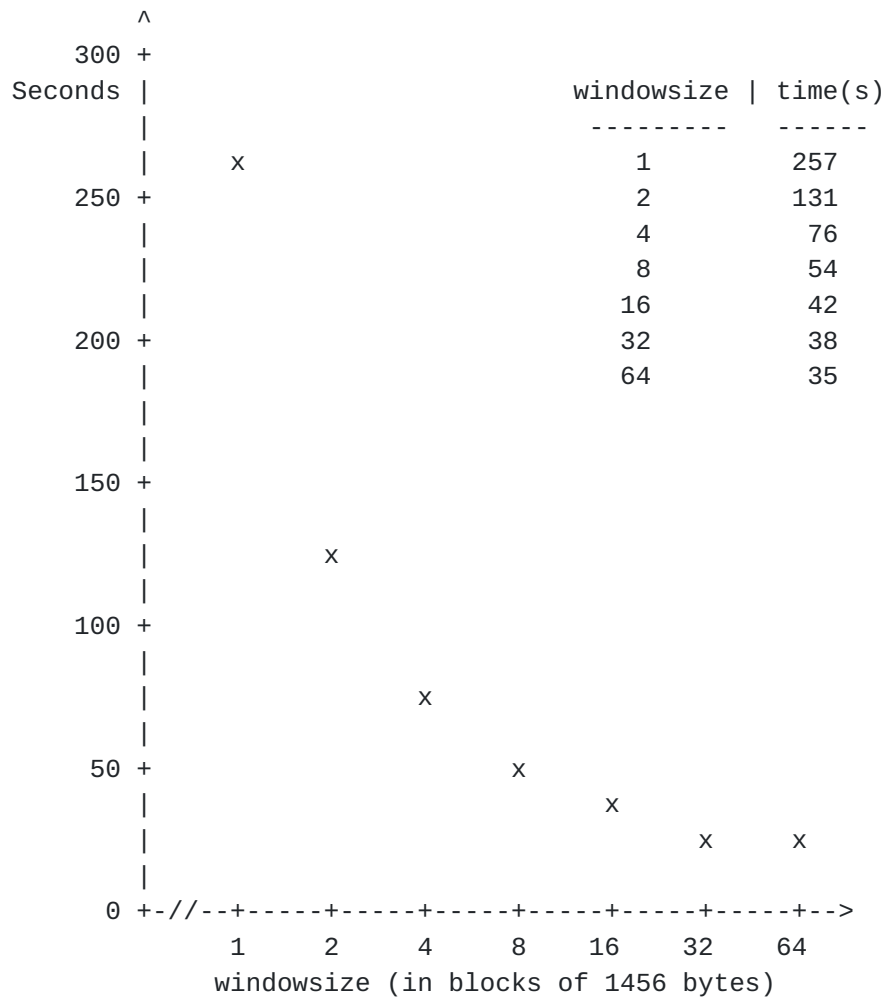

As it was expected, the transfer time decreases with the use of a windowed schema. The reason for the reduction in time is the reduction in the number of the required synchronous acknowledgements exchanged.

Comparatively the same 180 MB transfer performed over an SMB/CIFS mapped drive on the same scenario took 23 seconds.

Security Considerations

The basic TFTP protocol does not have an explicit security mechanism. However it also does not have list, rename, delete, nor overwrite capabilities either. This document does not add any security to TFTP nor the specified extension adds any additional security risk either.

IANA Considerations

This document has no actions for IANA.

Error Handling

In case of an error the last ACK received sets the beginning of the next windowsize window to send.

Section of a transfer including errors and error recovery

Normative References

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

[2] Malkin, G., Harkin, A., "TFTP Option Extension", RFC 2347 May 1998.

[3] Malkin, G., Harkin, A., "TFTP Blocksize option", RFC 2348 May 1998.

[4] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, November 2008.

[5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Patrick Masotta
300 W 11th Avenue,
Denver, CO 80204

EMail: masotta[-at-]vercot[-dot-]com

Copyright and IPR Provisions