

GridFTP Protocol Improvements

Status of this Document

This document is a Global Grid Forum Draft summarizing the analysis of GridFTP v1.0 protocol and experience of using the protocol in building distributed grid applications. This document is an EXPERIMENTAL document in that it provides suggestions for improvements in a protocol, based on results of experience with that protocol.

Copyright Notice

Copyright (C) Global Grid Forum (2004). All Rights Reserved.

1 Abstract

GridFTP protocol has become popular data movement tool used to build distributed grid-oriented applications. GridFTP v1.0 protocol [gftp] extends FTP protocol defined by RFC959 [rfc959] and other IETF documents by adding certain features designed to improve performance of data movement over wide area network, to allow the application to take advantage of "long fat" communication channels, to help build distributed data handling applications.

Several groups have developed independent implementations of the GridFTP protocol for different types of applications. This document summarizes the experience gained by these groups and their proposals for possible protocol improvements. The goal of these improvements is to develop more robust, reliable and scalable protocol for bulk file-oriented data transfer over wide and local area networks.

Contents

GridFTP Protocol Improvements 1

1 Abstract 1

2 GridFTP Points of Improvements 3

 2.1 Unidirectional Data Transfer in Extended Block Mode 3

 2.2 Ordering of PASV/SPAS and STOR/RETR Commands..... 3

 2.3 Possible Disconnection of Idle Control Channel Socket by Some Firewalls . 3

 2.4 Control of Contents and Frequency of Feedback From the Server 4

 2.5 Unreliable End-of-File Communication in Stream Mode..... 4

 2.6 Data Integrity Verification..... 4

 2.7 Structured Directory Listings..... 4

 2.8 IPV6 Support 4

 2.9 Packed Transfers 4

 2.10 Flexible Striping 5

3 References 5

4 Full Copyright Statement 5

5 GGF Intellectual Property Statement 5

2 GridFTP Points of Improvements

The following are the issues identified by GGF GridFTP Working Group as possible points of GridFTP protocol improvements.

2.1 Unidirectional Data Transfer in Extended Block Mode

GridFTP v1.0 requires that in extended block mode data is sent in the same direction as data channel connection is established. In other words, GridFTP server can send data only in active mode and receive data only in passive mode. This makes it impossible to use extended block mode in presence of firewall, etc.

There are several possible solutions of this problem:

- Pre-negotiation of number of data channels to be used for the transfer
- Modification of the protocol possibly introducing new transfer mode along with existing Stream, Block and Extended Block modes.

2.2 Ordering of PASV/SPAS and STOR/RETR Commands

As defined in RFC959, in passive mode server must reply to PASV command with the address of the data socket *before* it receives STOR or RETR command, and thus before it receives the name of the file to be transferred. In case of distributed server, this is not always possible. This makes it very difficult to implement passive mode in distributed FTP server in scalable way.

There are many proposed solutions for this problem:

1. Delayed passive option for PASV command which would allow to defer answer to the PASV command until STOR/RETR is received and include data socket address into (unused) answer to STOR/RETR
2. Introduction of PRET command as discussed in GridFTP v1.0 document which would carry attributes of the file which is about to be transferred and issuing PASV after PRET
3. Introduction of new pair of commands, GET and PUT which would essentially combine functionalities of (1) and (2) into single command/reply, preserve semantics of STOR/RETR and eliminate the need for PASV

2.3 Possible Disconnection of Idle Control Channel Socket by Some Firewalls

This problem is inherited from RFC959 FTP protocol. Some firewall software drops idle TCP connections. In some applications, such as disk cache in front of tape storage, data existing in name space is not always immediately accessible. In these cases, after issuing the STOR/RETR command, the client must wait for relatively long time before data transfer can even start. This makes the control channel socket connection idle for long time, and the firewall can drop it. The same may apply to data channel as well.

Proposed solution for control channel seems to be easier than for data channel. Performance data proposed in GridFTP v1.0 draft can be periodically sent over the control channel to keep it alive. As for data channel, some sort of keep-alive noise could be sent in the direction opposite to data transfer.

2.4 Control of Contents and Frequency of Feedback From the Server

Currently GridFTP server sends performance markers at fixed 5 second interval and restart markers come at a predefined block size. There should be some way to allow the interval between restart and performance markers to be set. Also it may be useful to allow this to be extensible so that other transfer event data could be returned as well, for example if the end host was a mass storage system and it were staging a file, it might send back ETA or % done markers.

Possible solutions would be to use FEAT/OPTS mechanism or introduction of new command (tentatively TREV for or TRansfer EVent).

2.5 Unreliable End-of-File Communication in Stream Mode

As specified by RFC959, during data upload, the server is supposed to treat end of data socket as end of file. This makes it impossible for the server to distinguish between normal end of file and abnormal client shutdown in the middle of data transfer.

A possible solution for this problem is to introduce EOF command, which would be sent by the client to confirm that the entire file was sent successfully over the data socket.

2.6 Data Integrity Verification

In order to protect data from transmission errors, some data integrity verification mechanism should be introduced on the level of FTP protocol, in addition to TCP packet checksums. Some sort of CRC or another form of digital signature should be calculated either over each block of data in block-oriented transfer modes such as Block and Extended Block or over whole file in Stream mode.

2.7 Structured Directory Listings

RFC959 does not specify format of directory listing sent in response to LIST command. Usually it is well suited for human reading, but not for computer processing. The IETF FTPEXT working group [ftpext] is developing extensions to FTP such as new MLST and MLSD commands, which would have easy to parse output. This proposal should be adopted by GridFTP protocol.

2.8 IPV6 Support

IPV6 uses significantly different addressing schema than IPV4. RFC 2428 [rfc2428] introduces EPRT and EPSV commands suitable for IPV6. Endorsement of the extensions proposed by RFC2428 should be considered.

2.9 Packed Transfers

In cases when it is necessary to transfer large number of small files, time spent on transfer initiation over the control channel and establishing data channel may create significant inefficiency in overall transfer performance. Partially this problem may be solved by reusing of data channels without opening and closing them for each file.

Another solution to consider is to pack multiple files into single file, transfer this file and then unpack it on the receiving end. ESTO and ERET commands already provisioned by GridFTP protocol seem to be useful in organizing such packed transfers.

2.10 Flexible Striping

There should be a way to dynamically and flexibly control striping algorithms.

3 References

- gftp W.Allcock, et al, GridFTP: Protocol Extensions to FTP for the Grid, published January 2004 as GFD-R.020. See <http://www.ggf.org/documents/>
- rfc959 J.Postel, J.Reinolds, File Transfer Protocol (FTP) <http://www.ietf.org/rfc/rfc0959>
- rfc2428 M.Allman, S.Ostermann, C.Metz, FTP Extensions for IPv6 and NATs <http://www.ietf.org/rfc/rfc2428>
- ftpext The FTPEXT working group of the IETF, see <http://www.ietf.org/html.charters/OLD/ftpext-charter.html>. As of early 2003 this group completed its chartered work and submitted an Internet Draft (Robert Elz, Paul Hethmon, Extensions to FTP) to the IETF RFC editor.

4 Full Copyright Statement

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

5 GGF Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available;

neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director.