

Table of Contents

GridFTP errors knowledge base.....	1
data_connect_failed() failed: authentication failed.....	1
Description.....	1
Scope and Impact.....	1
Solution.....	1
451 Local resource failure: malloc: Cannot allocate memory.....	1
425 Can't open data connection.	2
425 Can't open data connection. timed out() failed.....	3
425 Can't open data connection. data_connect_failed() failed: (message).....	3
426 Data connection. data_write() failed: Handle not in the proper state.....	3
451 Timeout (900 seconds): closing control connection.....	4

GridFTP errors knowledge base

This page describes the GridFTP errors seen on production tier-0 export transfers and details the causes of these errors.

data_connect_failed() failed: authentication failed

Description

GridFTP tries to read the user's `~/globus` directory, checking for certificates. If the user's home directory does not exist, there is no problem (ie. `dteam001`, `cms002` on the diskservers).

If the user's home directory *does* exist, the `~/globus` directory needs to be "listable". Otherwise, `gridftp` will complain like this:

```
error: the server sent an error response: 425 425 Can't open data connection. data_connect_failed
GSS Major Status: General failure
GSS Minor Status Error Chain:
```

```
accept_sec_context.c:115: gss_accept_sec_context: Error with gss context
globus_i_gsi_gss_utils.c:327: globus_i_gsi_gss_create_and_fill_context: Error with GSI credential
globus_gsi_system_config.c:3493: globus_gsi_sysconfig_get_cert_dir_unix: Could not find a valid t
globus_gsi_system_config.c:313: globus_i_gsi_sysconfig_create_cert_dir_string: Could not find a v
```

Scope and Impact

This problem affects users of non-LHC VO's (compass!), who are likely to be mapped to AFS accounts.

Solution

Check ACLs with `fs la ~/globus`. If necessary, run `fs sa -dir ~/globus -acl system:anyuser l`

451 Local resource failure: malloc: Cannot allocate memory

The error message `451 Local resource failure: malloc: Cannot allocate memory` in `/var/log/gridftp.log` has a non-obvious background. It *could* be caused by problems on the other end of the transfer...

>

> *Does someone know how to interpret the gridftp error message*

> *'451 Local resource failure: malloc: Cannot allocate memory'?*

David Smith explains:

```
It really is a malloc() failure - but the reason (at least for all 6 malloc errors in today's log
on that disk server) is a logic error during the premature termination of the loop the gridftp se
enters while it is sending out the data:
```

```
There is a timeout of 20 minutes inactivity on the data channel once it is connected: if that is
the transfer should abort. However the way g_send_data() (from globus_hack.c) works at the moment
after the timer expires the rest of whatever file is being sent is read into memory (without wait
send it anywhere). The 3 different files involved with the 6 instances of the malloc failure I lo
were all with ~3.6Gb files. This may not actually quite exhaust the memory on the machine, but pr
```

TransferOperationsGridFTPErrors < LCG < TWiki

caused the heap for the single ftp process to reach its maximum size, just because of the limited addressable memory.

The 6 problem transfers were all to various disk servers at in2p3. Checking one of the FTS logs (include the performance markers) there was apparently no data transferred at all - so perhaps the side was waiting for some condition before being able to accept any bytes. (Free space, free move

And indeed: note the 20 minute gap in the log entries below!

```
May 4 07:34:28 lxfsrc6606 gridftpd[17904]: connection from fts109.cern.ch [128.142.160
May 4 07:34:29 lxfsrc6606 gridftpd[17904]: <--- 220 lxfsrc6606.cern.ch CASTOR GridFTP
1069715860-42) ready.
May 4 07:34:29 lxfsrc6606 gridftpd[17904]: <--- 334 Using authentication type GSSAPI;
May 4 07:34:29 lxfsrc6606 gridftpd[17904]: <--- ADAT (13749 bytes)
May 4 07:34:30 lxfsrc6606 gridftpd[17904]: <--- ADAT (109 bytes)
May 4 07:34:30 lxfsrc6606 gridftpd[17904]: <--- ADAT (413 bytes)
May 4 07:34:30 lxfsrc6606 gridftpd[17904]: client identity is: "/C=CH/O=CERN/OU=GRID/C
May 4 07:34:31 lxfsrc6606 gridftpd[17904]: <--- 235 GSSAPI Authentication succeeded
May 4 07:34:31 lxfsrc6606 gridftpd[17904]: GSSAPI user "/C=CH/O=CERN/OU=GRID/CN=Miguel
May 4 07:34:31 lxfsrc6606 gridftpd[17904]: <--- 331 GSSAPI user "/C=CH/O=CERN/OU=GRID/
May 4 07:34:32 lxfsrc6606 gridftpd[17904]: USER :globus-mapping:
May 4 07:34:32 lxfsrc6606 gridftpd[17904]: PASS password
May 4 07:34:32 lxfsrc6606 gridftpd[17904]: <--- 230-No directory! Logging in with home
May 4 07:34:33 lxfsrc6606 gridftpd[17904]: stager-mapping: STAGE_HOST: <c2atlast0rh>;
RFIO_USE_CASTOR_V2: <YES>
May 4 07:34:33 lxfsrc6606 gridftpd[17904]: <--- 230 User atlas002 logged in.
May 4 07:34:33 lxfsrc6606 gridftpd[17904]: FTP LOGIN FROM fts109.cern.ch [128.142.160.
May 4 07:34:33 lxfsrc6606 gridftpd[17904]: failed to stat umask-mapfile "/etc/castor/g
May 4 07:34:33 lxfsrc6606 gridftpd[17904]: FEAT
May 4 07:34:34 lxfsrc6606 gridftpd[17904]: <--- 211-Extensions supported:
May 4 07:34:34 lxfsrc6606 gridftpd[17904]: <--- REST STREAM
May 4 07:34:35 lxfsrc6606 gridftpd[17904]: <--- ESTO
May 4 07:34:35 lxfsrc6606 gridftpd[17904]: <--- ERET
May 4 07:34:36 lxfsrc6606 gridftpd[17904]: <--- MDTM
May 4 07:34:37 lxfsrc6606 gridftpd[17904]: <--- SIZE
May 4 07:34:38 lxfsrc6606 gridftpd[17904]: <--- PARALLEL
May 4 07:34:38 lxfsrc6606 gridftpd[17904]: <--- DCAU
May 4 07:34:39 lxfsrc6606 gridftpd[17904]: <--- 211 END
May 4 07:34:40 lxfsrc6606 gridftpd[17904]: TYPE Image
May 4 07:34:41 lxfsrc6606 gridftpd[17904]: <--- 200 Type set to I.
May 4 07:34:41 lxfsrc6606 gridftpd[17904]: MODE 0
May 4 07:34:42 lxfsrc6606 gridftpd[17904]: <--- 200 MODE E ok.
May 4 07:34:42 lxfsrc6606 gridftpd[17904]: SIZE //castor/cern.ch/grid/atlas/t0/perm/T0
May 4 07:34:42 lxfsrc6606 gridftpd[17904]: <--- 213 3840000000
May 4 07:34:43 lxfsrc6606 gridftpd[17904]: <--- 200 Opts successful.
May 4 07:34:43 lxfsrc6606 gridftpd[17904]: DCAU
May 4 07:34:44 lxfsrc6606 gridftpd[17904]: <--- 200 DCAU N
May 4 07:34:44 lxfsrc6606 gridftpd[17904]: PORT
May 4 07:34:45 lxfsrc6606 gridftpd[17904]: ok: PORT 193.48.99.86,30018 from fts109.cer
May 4 07:34:46 lxfsrc6606 gridftpd[17904]: <--- 200 PORT command successful.
May 4 07:34:46 lxfsrc6606 gridftpd[17904]: RETR //castor/cern.ch/grid/atlas/t0/perm/T0
May 4 07:35:02 lxfsrc6606 gridftpd[17904]: using rfio: file not local
May 4 07:35:02 lxfsrc6606 gridftpd[17904]: <--- 150 Opening BINARY mode data connectio
May 4 07:55:35 lxfsrc6606 gridftpd[17904]: <--- 451 Local resource failure: malloc: Ca
May 4 07:55:35 lxfsrc6606 gridftpd[17904]: Local resource failure: malloc
May 4 07:55:37 lxfsrc6606 gridftpd[17904]: <--- 221 You could at least say goodbye.
May 4 07:55:43 lxfsrc6606 gridftpd[17904]: FTP session closed
```

425 Can't open data connection. .

Error on side sending data: Means the attempt to establish the data connection(s) with the peer timed out, by default after 2 minutes. Normally the sending side initiates a TCP connection towards the receiving side. A TCP connect may fail with an error after about 2 minutes, so it is possible that a failure rather than a timeout may sometimes be reported.

451 Local resource failure: malloc: Cannot allocatememory

425 Can't open data connection. timed out() failed.

Error on side receiving data: Means the attempt to establish the data connection(s) with the peer timed out, by default after 2 minutes. Normally the receiving side listens on a variable numbered port, awaiting data connection(s) from the sending side. In a third party copy the receiving side is contacted before the sending side, so this timeout may indicate the client is slow to contact the source site or that the source site is slow to initiate the connection to the destination, or that there is a connectivity or firewall problem.

425 Can't open data connection. data_connect_failed() failed: (message) .

Error on side sending OR receiving data: Means the attempt to establish the data connection(s) with the peer had an error. See the example above for the particular case of the message "authentication failed".

426 Data connection. data_write() failed: Handle not in the proper state

Is rather generic: means that the side sending out data encountered an error while sending the data to the ftp subsystem. However often this simply indicates that the TCP data connection(s) closed - probably the peer closed them although network problems have also been known to cause the connection to reset. The message '150 Opening BINARY mode data connection' is usually logged just before the error and indicates the point at which the data connection(s) were established. Normally the side sending the data will initiate the connection to the receiving side, this arrangement is required when using extended mode (e.g. when more than one parallel stream). In the example below the connection was indeed made from CERN out to the peer, as can be seen from the earlier 'PORT' command.

So in the example below the peer (131.169.80.47) accepted data connections on port 20002 but probably the connections were broken after 15 seconds. The short duration of the connection could indicate the peer deliberately closed the connection rather than a generic connectivity issue.

David

Example:

```

May 16 04:45:08 lxfsra3001 gridftpd[5523]: connection from fts108.cern.ch [128.142.160.
May 16 04:45:08 lxfsra3001 gridftpd[5523]: <--- 220 lxfsra3001.cern.ch CASTOR GridFTP S
May 16 04:45:08 lxfsra3001 gridftpd[5523]: <--- 334 Using authentication type GSSAPI; A
May 16 04:45:08 lxfsra3001 gridftpd[5523]: <--- ADAT (13749 bytes)
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- ADAT (109 bytes)
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- ADAT (413 bytes)
May 16 04:45:09 lxfsra3001 gridftpd[5523]: client identity is: "/C=CH/O=CERN/OU=GRID/CN
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 235 GSSAPI Authentication succeeded
May 16 04:45:09 lxfsra3001 gridftpd[5523]: GSSAPI user "/C=CH/O=CERN/OU=GRID/CN=Maarten
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 331 GSSAPI user "/C=CH/O=CERN/OU=GRID/C
May 16 04:45:09 lxfsra3001 gridftpd[5523]: USER :globus-mapping:
May 16 04:45:09 lxfsra3001 gridftpd[5523]: PASS password
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 230-No directory! Logging in with home=
May 16 04:45:09 lxfsra3001 gridftpd[5523]: stager-mapping: STAGE_HOST: <castorpublic>;
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 230 User dteam001 logged in.
May 16 04:45:09 lxfsra3001 gridftpd[5523]: FTP LOGIN FROM fts108.cern.ch [128.142.160.8
May 16 04:45:09 lxfsra3001 gridftpd[5523]: failed to stat umask-mapfile "/etc/castor/gr
May 16 04:45:09 lxfsra3001 gridftpd[5523]: FEAT
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 211-Extensions supported:
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- REST STREAM
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- ESTO
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- ERET
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- MDTM

```

TransferOperationsGridFTPErrors < LCG < TWiki

```
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- SIZE
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- PARALLEL
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- DCAU
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 211 END
May 16 04:45:09 lxfsra3001 gridftpd[5523]: TYPE Image
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 200 Type set to I.
May 16 04:45:09 lxfsra3001 gridftpd[5523]: MODE 0
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 200 MODE E ok.
May 16 04:45:09 lxfsra3001 gridftpd[5523]: SIZE //castor/cern.ch/grid/dteam/sc4/044/fil
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 213 1048576000
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 200 Opts successful.
May 16 04:45:09 lxfsra3001 gridftpd[5523]: DCAU
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 200 DCAU N
May 16 04:45:09 lxfsra3001 gridftpd[5523]: PORT
May 16 04:45:09 lxfsra3001 gridftpd[5523]: ok: PORT 131.169.80.47,20002 from fts108.cer
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 200 PORT command successful.
May 16 04:45:09 lxfsra3001 gridftpd[5523]: RETR //castor/cern.ch/grid/dteam/sc4/044/fil
May 16 04:45:09 lxfsra3001 gridftpd[5523]: using rfio: file not local
May 16 04:45:09 lxfsra3001 gridftpd[5523]: <--- 150 Opening BINARY mode data connection
May 16 04:45:24 lxfsra3001 gridftpd[5523]: <--- 426 Data connection. data_write() faile
May 16 04:45:24 lxfsra3001 gridftpd[5523]: Data connection. data_write() failed: Handle
May 16 04:45:24 lxfsra3001 gridftpd[5523]: <--- 221 You could at least say goodbye.
May 16 04:45:24 lxfsra3001 gridftpd[5523]: lost connection to fts108.cern.ch [128.142.1
May 16 04:45:24 lxfsra3001 gridftpd[5523]: FTP session closed
```

451 Timeout (900 seconds): closing control connection.

Hi,

Since the last 'export meeting' I'm not sure if anyone had been looking at the:

```
ERROR the server sent an error response: 421 421 Timeout (900 seconds): closing control connection
```

errors (on the castor side), but since that meeting FTS had a GGUS ticket opened asking about it decided to check into so we could at least answer the ticket. I looked at one particular instance rate at which they are seen is variable, the day I was looking at there were ~100.

I'll check one or two more instances of the error to see if there is any other failure paths - but happening:

[Example transfer was a CMS transfer, from gsiftp://lxfsra2008.cern.ch//castor/cern.ch/cms/store/LoadTest07_CERN_74 to in2p3 (dCache), started at about 08:36:04 on 27 Apr, error reported at about later]. It is clear the error was coming from the castor gridftp side, as the dcache gridftp does at least not at 900 seconds.

(1) There is a bug in the gridftp code in retrieve() in ftpd.c:

```
/* no command */
if (cmd == NULL)
{
    fin = fopen(name, "r"), closefunc = fclose;
    st.st_size = 0;
}
else
[...]
```

```
if (fin == NULL)
{
    if (errno != 0)
        perror_reply(550, name);
[...]
```

```
return;
```

The 'fopen' above is defined as a preprocessor macro for 'rfio_fopen64'. Under error conditions it returns an error number in serrno, rather than errno. So the later perror_reply() may not be called, and returning a reply to the ftp client. The ftp server reenters its command loop, waiting for commands

426 Data connection. data_write() failed: Handle not in the proper state

TransferOperationsGridFTPErrors < LCG < TWiki

while the client is waiting for a result from its 'RETR/ERET' command. The server eventually times out the command which is why we get the 900 second timeout error.

(2) The reason for the `rfio_fopen64` failure (for the job which I checked) is a timeout between the `rfio_fopen64` call and the `rmmaster` for the submission of the transfer job:

(From the `dlf` logging: time, hostname, subreq id, parameters)

```
27-04-2007 08:36:04.523124: c2cmssrv03.cern.ch: 463199d4-0000-1000-98c8-9ac26901000:
SYSTEM=stager_db_process
STRING=Forcing filesystem(s): "lxfra2008:/srv/castor/03/"
File=stager_db_service.c
Line=1301
errno=0
serrno=0
```

```
27-04-2007 08:38:04.533415: c2cmssrv03.cern.ch: 463199d4-0000-1000-98c8-9ac26901000:
ERROR=rmllog
STRING=usersend2rmd: rcv error : Timed out File=stager_db_service.c
Line=2008
errno=115
serrno=1004
```

The time difference is the expected 120 seconds from:

```
RM TIMEOUT 120
```

(3) The `rmmaster` logged this:

```
[...]
Apr 27 08:36:04 rmmaster[3139,73289]: Crmd_processuser_enterjob: ... rfs : lxfra2008:
Apr 27 08:36:04 rmmaster[3139,73289]: Crmd_processuser_validate_usergroup: Entering
Apr 27 08:36:04 rmmaster[3139,73289]: Crmd_processuser_validate_usergroup: Leaving
Apr 27 08:36:04 rmmaster[3139,73289]: Crmd_processuser_enterjob: Job No 463199d4-0000-1000-98c8-9ac26901000
being submitted as PSUSP to LSF DATE=20070427083604.533435 HOST=c2cmssrv06.cern.ch LVL=Usage FACILITY=
PID=3139 TID=73289 MSG="Job is being submitted as PSUSP to LSF" REQID=46272f11-0000-1000-b928-9500
NSHOSTNAME=N/A NSFILEID=0 FUNC="Crmd_processuser_enterjob" JOB_ID="463199d4-0000-1000-98c8-9ac26901000"
Apr 27 08:36:04 rmmaster[3139,73289]: Crmd_process_user_lsf_request: Entering
Apr 27 08:36:04 rmmaster[31705]: Crmd_process_user_lsf_doit: Entering
Apr 27 08:36:04 rmmaster[31705]: Crmd_process_user_lsf_doit: Asking for 1 hosts:
Apr 27 08:36:04 rmmaster[31705]: Crmd_process_user_lsf_doit: Host No 0: lxfra2008
Apr 27 08:36:04 rmmaster[31705]: Crmd_process_user_lsf_doit: Processing job enter
[... next log lines from process 31705 ...]
Apr 27 08:39:12 rmmaster[31705]: Crmd_process_user_lsf_report_ok: Entering
Apr 27 08:39:12 rmmaster[31705]: Crmd_process_user_lsf_report_ok: Leaving
[...]
```

I can't see anything in the code that should take time apart from the `lsb_submit()` - so I think it takes just over 3 mins, to submit the job. At that point I haven't gone further - e.g. to see why LSF messages are unusual. Do you think it would be useful to try to do that, or perhaps it is well known that the reason is

Yours,
David

Last edit: DavidSmith on 2007-10-22 - 15:28

Number of topics: 1

Maintainer: GavinMcCance, JanVanEldik

This topic: LCG > TransferOperationsGridFTPErrors

Topic revision: r5 - 2007-10-22 - DavidSmith



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback