



Enabling Grids for E-scienceE

Tutorial - Debugging of FTS transfer problems

Paolo Badino (paolo.badino@cern.ch)

Gavin McCance (gavin.mccance@cern.ch)

WLCG Asian Tier-2 Workshop - Mumbai, 2 December 2006

www.eu-egee.org
www.glite.org



- **FTS Overview**
- **Tier-2 Perspective**
 - What transfers will I see and from where?
 - What control do I have over my channels?
How do I set up the client software?
- **Debugging problems**
 - What errors/tickets will I see and how do I handle them?
 - What else do I need to check for my site?
- **Reporting & Monitoring**



- **Transfer service tickets are generally of three categories:**
- **Information service tickets**
 - Publishing bad SRM information or not enough information
- **Explicit FTS problems or user problems**
 - Various possible configuration problems, failures, submitting to the wrong endpoint, etc
- **SRM problems**
 - (“FTS says my job failed with ...”)

- Any problem through GGUS will be assigned first to your ROC
 - And they will then involve the tier-2 admins if necessary
- Tier-2's should expect to process **information service problems** at their site and **SRM problems** at their site
- **FTS server problems and general user problems** will generally be taken by the tier-1 running the FTS server (or their ROC).

- **Make sure your SRMs are published in information system**
 - FTS only runs transfers to production SRMs published in BDII and GOCDB
 - In particular:
 - Make sure the SRM is registered in GOC-DB
 - make sure there is a Glue 1.2 **GlueService** entry, as well as the normal GlueSE entry
 - make sure the the **GlueServiceForeignKey GlueSiteUniqueID** is set to the correct GOC-DB name for your site (ideally upper-case)
 - make sure the VOs that you support are published in both the **GlueSARoots** and in the **GlueService** entry

- **Please check for your site now...**
 - SFT on the way...

- **A variety of service problems may be reported on the FTS server itself, most of which are not the responsibility of the tier-2**
 - **BUT: Your tier-1 site should tell you if there is a problem on their FTS which will affect the rate to your site**
- **Client problems also happen from time to time**
- **If you receive mails / complaints indicating the following problems(...), please ask the user to submit a ticket to GGUS**
 - These should be handled by your ROC
 - Feel free the short-cut if it's clear what the problem is and you can fix it 😊

- **Cannot submit to web-service**
- **USER USING WRONG ENDPOINT:**
Failed to determine the interface version of the service: getInterfaceVersion: SOAP fault: "http://xml.apache.org/axis/":Server.NoService - The AXIS engine could not find a target service to invoke! targetService is FileTransfer
- **USER USING WRONG HOST:**
Failed to determine the interface version of the service: getInterfaceVersion: SOAP fault: SOAP-ENV:Client - CGI-gSOAP: Could not open connection ! (TCP get host by name failed in tcp_connect())
- **BAD GRID CREDENTIALS / CA CERTS / CRL LISTS ON CLIENT OR SERVER:**
Failed to determine the interface version of the service: getInterfaceVersion: SOAP fault: SOAP-ENV:Client - CGI-gSOAP: GSS Major Status: Authentication Failed
GSS Minor Status Error Chain:
(null)
- **USER NOT IN FTS SUBMIT MAPFILE (generated from VOMS):**
submit: You are not authorised to submit jobs to this service
- **If your users report this to you, and the solution isn't obvious, please ask them to submit to GGUS, and the request should be handled by your ROC**

- **The transfer fails with error:**

"No site found for host srm.random.site"

"No Channel found or VO not authorized for transferring between TIER2SITE and TIER1SITE"

- **These are likely information service problems**

- SRM not published
- VO not published for a given SRM
- Site name incorrect in GlueService

- **..or FTS configuration problems:**

- VO not given a share on the channel
- Channel not defined for this transfer

- **Same again – ask users to submit to GGUS if the problem isn't obvious**

- Feel free to shortcut with your tier-1 admin if you know what's wrong

- **SRM problems: These will be reported by FTS if the retries all fail**
- **If the failure was on write, you will see an error with “SRM put”:**
 - Failed on SRM put: SRM getRequestStatus timed out on put" and "Failed on SRM put: Failed SRM put on httpg://srm.cern.ch:8443/srm/managerv1 ; id=770376972 call. Error is CastorStagerInterface.c:2457 Device or resource busy (errno=0, serrno=0)
- **If the failure was on read, you will see an error with “SRM get”:**
 - FINAL:SRM_SOURCE: Failed on SRM get: Failed SRM get on httpg://srm.cern.ch:8443/srm/managerv1 ; id=759760078 call. Error is specified file(s) does not exist%
- **Sometimes the failure is on the gridFTP transfer:**
 - FINAL:TRANSFER: Transfer failed. ERROR the server sent an error response: 425 425 Cannot open port: java.lang.Exception: Pool manager error: No write pool available for <sc_tape:SC@osm>
- **The message should contain the error string that was returned by the SRM or gridFTP server to the FTS - and hopefully this is sufficient to start investigating the problem**

- Otherwise...
- Detailed logging. If it's not obvious, the tier-1 FTS admin can give you the FTS transfer log to correlate with your SRM logs
- The successful call pattern to your SRM is:
 - SRM get on source with source SURL
 - SRM getStatus to get the source TURL
 - SRM put on destination with dest SURL
 - SRM getStatus to get the dest TURL
 - globus_gass_copy_url_to_url (VDT gridFTP client) using TURLs
 - SRM setStatus("Done") on source SURL
 - SRM setStatus("Done") on dest SURL

29164

GSIFTP

```

srm://srm.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
srm://dcache.gridpp.rl.ac.uk/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:33,027 [INFO ] - STATUS:START
2006-06-13 13:09:33,028 [WARN ] - Getting source from SRM: ; srm://srm.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
2006-06-13 13:09:33,028 [DEBUG ] - contacting srm for get at ; http://srm.cern.ch:8443/srm/managerv1
2006-06-13 13:09:33,042 [DEBUG ] - Entered SrmUtil::constructor
2006-06-13 13:09:33,043 [INFO ] - STATUS:BEGIN:SRM_GET
2006-06-13 13:09:33,043 [DEBUG ] - Now calling requestTurlFromSurl ; srm://srm.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
2006-06-13 13:09:33,043 [DEBUG ] - Entered SrmUtil::requestTurlFromSurl
2006-06-13 13:09:33,043 [DEBUG ] - Entered SrmUtil::get
2006-06-13 13:09:33,043 [DEBUG ] - Entered SrmUtil::checkService - srm ping
2006-06-13 13:09:33,326 [DEBUG ] - Exiting SrmUtil::checkService - srm__ping to 0.0.0.0
2006-06-13 13:09:34,083 [DEBUG ] - Exiting SrmUtil::get - srm__get to 0.0.0.0
2006-06-13 13:09:34,083 [DEBUG ] - Request Status Ready
2006-06-13 13:09:34,083 [INFO ] - SRM get request ID: 679085254
2006-06-13 13:09:34,083 [DEBUG ] - Exiting SrmUtil::prepareRequest
2006-06-13 13:09:34,083 [DEBUG ] - got file status state Ready
2006-06-13 13:09:34,083 [DEBUG ] - Exiting SrmUtil::requestTurlFromSurl - finishing
2006-06-13 13:09:34,083 [DEBUG ] - Entered SrmUtil::finishRequest
2006-06-13 13:09:34,083 [DEBUG ] - TURL is for SRM get gsiftp://lxfsrc6208.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
2006-06-13 13:09:34,083 [INFO ] - STATUS:END:SRM_GET
2006-06-13 13:09:34,083 [DEBUG ] - TURL = ; gsiftp://lxfsrc6208.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
2006-06-13 13:09:34,083 [DEBUG ] - Entered SrmUtil::setFileStatus
2006-06-13 13:09:34,083 [DEBUG ] - setting file status to Running
2006-06-13 13:09:34,355 [DEBUG ] - Exiting SrmUtil::srm__setFileStatus to 0.0.0.0
2006-06-13 13:09:34,355 [DEBUG ] - SIZE = 1048576000
2006-06-13 13:09:34,355 [WARN ] - Prepare put to SRM: ; srm://dcache.gridpp.rl.ac.uk/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:34,355 [DEBUG ] - contacting srm for put at ; http://dcache.gridpp.rl.ac.uk:8443/srm/managerv1
2006-06-13 13:09:34,355 [DEBUG ] - Entered SrmUtil::constructor
2006-06-13 13:09:34,356 [INFO ] - STATUS:BEGIN:SRM_PUT
2006-06-13 13:09:34,356 [DEBUG ] - Now calling requestTurlFromSurl ; srm://dcache.gridpp.rl.ac.uk/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:34,356 [DEBUG ] - Entered SrmUtil::requestTurlFromSurl
2006-06-13 13:09:34,356 [DEBUG ] - Entered SrmUtil::put
2006-06-13 13:09:34,356 [DEBUG ] - Entered SrmUtil::checkService - srm ping
2006-06-13 13:09:34,939 [DEBUG ] - Exiting SrmUtil::checkService - srm__ping to 0.0.0.0
2006-06-13 13:09:35,799 [DEBUG ] - Exiting SrmUtil::put - srm__put to 0.0.0.0
2006-06-13 13:09:35,799 [DEBUG ] - Request Status Pending
2006-06-13 13:09:35,799 [INFO ] - SRM put request ID: -2147439870
2006-06-13 13:09:35,799 [DEBUG ] - Exiting SrmUtil::prepareRequest
2006-06-13 13:09:35,799 [DEBUG ] - got file status state Pending
2006-06-13 13:09:35,799 [DEBUG ] - Exiting SrmUtil::requestTurlFromSurl - not finished yet
2006-06-13 13:09:35,799 [DEBUG ] - Entered SrmUtil::updateTurlFromSurlRequest
2006-06-13 13:09:35,799 [DEBUG ] - Entered SrmUtil::getRequestStatus
2006-06-13 13:09:36,533 [DEBUG ] - Exiting SrmUtil::srm__getRequestStatus to 0.0.0.0
2006-06-13 13:09:36,533 [DEBUG ] - Request Status Active
2006-06-13 13:09:36,533 [DEBUG ] - got file status state Ready
2006-06-13 13:09:36,533 [DEBUG ] - Exiting SrmUtil::updateTurlFromSurlRequest - finishing
2006-06-13 13:09:36,533 [DEBUG ] - Entered SrmUtil::finishRequest
2006-06-13 13:09:36,533 [INFO ] - SRM put on http://dcache.gridpp.rl.ac.uk:8443/srm/managerv1 ; id=-2147439870 call returned an error:
2006-06-13 13:09:36,533 [INFO ] - This error is ignored since we did receive a TURL:gsiftp://gftp0444.gridpp.rl.ac.uk:2811/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:36,533 [DEBUG ] - TURL is for SRM put gsiftp://gftp0444.gridpp.rl.ac.uk:2811/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:36,533 [INFO ] - STATUS:END:SRM_PUT
2006-06-13 13:09:36,533 [DEBUG ] - TURL = ; gsiftp://gftp0444.gridpp.rl.ac.uk:2811/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:36,533 [DEBUG ] - Entered SrmUtil::setFileStatus
2006-06-13 13:09:36,533 [DEBUG ] - setting file status to Running
2006-06-13 13:09:37,246 [DEBUG ] - Exiting SrmUtil::srm__setFileStatus to 0.0.0.0
2006-06-13 13:09:37,246 [INFO ] - STATUS:BEGIN:TRANSFER
2006-06-13 13:09:37,246 [WARN ] - Starting gsiftp transfer
TURL source = gsiftp://lxfsrc6208.cern.ch/castor/cern.ch/grid/dteam/sc4/055/file-086.dat
TURL dest = gsiftp://gftp0444.gridpp.rl.ac.uk:2811/pnfs/gridpp.rl.ac.uk/data/dteam/sc3/2006-06-13/file-5485-cbeddd91-1227-4a86-849b-a198fb9cedec.dat
2006-06-13 13:09:37,246 [DEBUG ] - Calling gsiftp getfilesize
FILE SIZE = 1048576000
2006-06-13 13:09:38,436 [DEBUG ] - Calling gsiftp transfer
2006-06-13 13:10:12,313 [INFO ] - Transferred 253165568 out of 1048576000 bytes. Inst: 8241.07 KB/s. Avg: 8241.07 KB/s
2006-06-13 13:10:42,372 [INFO ] - Transferred 662110208 out of 1048576000 bytes. Inst: 13312.00 KB/s. Avg: 10776.53 KB/s
2006-06-13 13:11:12,430 [INFO ] - Transferred 1019609088 out of 1048576000 bytes. Inst: 11598.67 KB/s. Avg: 11051.19 KB/s
2006-06-13 13:11:14,601 [INFO ] - GSIFTP: source: set up FTP mode. DCAU disabled. Streams = 5, TCP buffersize = 0
GSIFTP: dest: set up FTP mode. DCAU disabled. Streams = 5, TCP buffersize = 0
2006-06-13 13:11:17,440 [DEBUG ] - Gsiftp transfer done.
2006-06-13 13:11:18,956 [DEBUG ] - filesize at destination = 1048576000
2006-06-13 13:11:18,956 [INFO ] - STATUS:END:TRANSFER
2006-06-13 13:11:18,956 [WARN ] - Done with gsiftp transfer

```

Times are UTC

SRM request IDs are provided

All SRM calls and returned statuses are logged

gridFTP transfer markers are reported

WLCG Asian Tier-2 Workshop -
Debugging FTS problems

- The unsuccessful call pattern is (for example):
 - SRM get on source with source SURL
 - SRM getStatus to get the source TURL
 - SRM put on source with dest SURL
 - SRM getStatus to get the dest TURL
 - globus_gass_copy_url_to_url FAILS
 - SRM setStatus("Done") on source SURL
 - SRM setStatus("Done") on dest SURL
 - SRM advisoryDelete on dest SURL
 - It always attempts to cleanup and delete the destination SURL upon failure (although when things are going wrong it doesn't always succeed).

- **FTS is a front-end service**
 - It uses other services (SRM, GridFTP, MyProxy ...)
 - Users will always contact the FTS administrators (even when the problem is somewhere else)
 - Can be used as a tool for monitoring the overall transfer service
 - since it measures the status of all the transfers
 - The service operation needs good error reporting
 - The service team needs to be able to get a consolidated view of the overall service status
 - i.e. the overall state of all the components involved

- **Reporting:**
 - Produce a consolidated view of the service status
 - Measure the obtained service level
 - Understand the actual problems
- **Monitoring:**
 - Provide immediate feedback on current state of the transfer service
 - Alarms
 - Act proactively to avoid problems
- **We need both !**

- **We have a lot of information stored in the DB that we can extract, e.g.:**
 - Summaries of
 - channel activity
 - failure rate
 - type of failures
 - which site failed?
 - average file size, average duration, bytes written, ...
 - The new DB schema version (FTS 2.0) stores a lot more information to help understand the system

- **We need a way to present it**
 - (work in progress...)


FTS Report - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://pcitgm02.8081/channeltransfer/day/

CERN DataMgmt Links Development Interesting Stuff

FTS Report



Disclaimer
This page contains a report generated from information stored in the FTS Database and is intended for reporting purposes only. Since the format will probably change in the future, it's therefore recommended not to use parsing robots on it.

Statistics concerning all the transfers performed yesterday
Between 2006-10-12 08:00:00 +02:00 and 2006-10-13 08:00:00 +02:00

CERN-*

Channel Name	VO Name	Total	% Failures	# Succ.	# Fail.	1st Failure Reason	% 1st Failure Reason	2nd Failure Reason	% 2nd Failure Reason	Avg. Size (GB)	Avg. Duration (sec)	Avg. Tx Rate (MB/sec)	Eff. Tx Bytes (GB)	Tx Bytes (GB)
CERN-PIC	[All]	12262	73.97	3192	9070	Dest SRM	56.22	Other	37.53	0.53	263.03	1.62	1700.41	1700.41
	atlas	8932	99.92	7	8925	Dest SRM	57.13	Other	38.08	0	220	0	0	0
	cms	208	0	208	0					2.7	767.55	3.64	561.26	561.26
	dteam	974	0.51	969	5	Other		80 Source SRM	20	0.95	356.31	2.88	923.83	923.83
	lhcb	2145	6.53	2005	140	Source SRM	99.29	Other	0.71	0.11	165.85	0.81	215.32	215.32
	ops	3	0	3	0					0	202.67	0	0	0
CERN-RAL	[All]	8699	59.26	3544	5155	Other	84.91	Source SRM	14.88	0.85	478.22	2.59	3026.81	3027.57
	alice	1155	82.6	201	954	Other	99.58	Dest SRM	0.31	1.86	1805.05	1.11	372.95	372.95
	atlas	4512	88.52	518	3994	Other	84.85	Source SRM	15.15	1.79	1428.94	1.49	926.26	926.57
	cms	227	3.08	220	7	Dest SRM	85.71	Source SRM	14.29	2.53	348.65	10.08	555.61	555.61
	dteam	1077	3.99	1034	43	Other	86.05	Source SRM	9.3	0.95	276.64	4.01	980.47	980.91
	lhcb	1725	9.1	1568	157	Source SRM	99.36	Other	0.64	0.12	146.03	1.16	191.52	191.52
	ops	3	0	3	0					0	27	0.01	0	0
CERN-SARA	[All]	8792	42.55	5051	3741	Dest SRM	83.77	Source SRM	12.22	1.34	108.02	15.4	6777.95	6784.92
	alice	3134	15.12	2660	474	Source SRM	57.17	Dest SRM	41.14	1.66	109.53	18.43	4426.44	4430.29
	atlas	2018	53.32	942	1076	Dest SRM	72.4	Source SRM	16.54	1.15	144.44	9.42	1085.07	1087.6
	dteam	3488	61.32	1349	2139	Dest SRM	98.74	Other	0.98	0.93	81.91	14.66	1260.74	1261.32
	lhcb	148	35.14	96	52	Dest SRM	92.31	Other	3.85	0.06	76.1	0.93	5.7	5.7
	ops	4	0	4	0					0	97.25	0.02	0	0
CERN-INFN	[All]	11492	42.31	6630	4862	Dest SRM	43.85	Other	37.7	1.13	395.77	3.21	7514.29	7614.84
CERN-CERN	[All]	1536	39.71	926	610	Source SRM	58.36	Dest SRM	15.9	0.07	287.71	0.38	67.89	69.08
CERN-ASCC	[All]	6851	23.54	5238	1613	Source SRM	50.84	Other	28.89	1.14	1098.6	1.08	5955.81	6080.58
CERN-GRIDKA	[All]	12755	21.38	10028	2727	Source SRM	64.36	Other	32.53	0.87	371.97	3.19	8762.02	8767.53
CERN-TRIUMF	[All]	2244	20.63	1781	463	Other	61.77	Source SRM	31.1	1.04	395.15	3.63	1847.25	1917.13
CERN-BNL	[All]	13975	19.42	11261	2714	Source SRM	69.97	Other	24.24	0.44	190.38	3.41	4951.59	4960.34
CERN-IN2P3	[All]	11697	13.76	10087	1610	Source SRM	48.57	Other	47.45	1.22	296.21	5.33	12329.63	12329.63
CERN-FNAL	[All]	917	4.58	875	42	Transfer	97.62	Other	2.38	0	379.88	0	0	0

Click on the Channel Name to show the VO details

- **FTS should be a tool which makes it easier to run the transfer service**
- **We need to integrate with existing monitoring and operations tools, providing direct feedback to the people running the service (notifications, CIC on duty alarms, GGUS tickets)**
- **We cannot meet the WLCG service availability targets without better monitoring tools**
 - The current distributed service takes too much ‘expert’ supervision to discover and debug **inter-site** problems
- **Provide current status**
 - e.g. Per-site “top-5 errors on your site” web-page automatically generated by the FTS
- **A lot of work has already been done by the T1 sites for this**
 - e.g. plumbing onto local monitoring systems
 - http://winnetou.matrix.sara.nl/monitoring/datatransfer/fts_transfers.tar.gz.
- Extract the “common stuff” into the base FTS release

- **You should expect to be assigned tickets for:**
 - Information publishing problems
 - SRM problems for investigation
- **Bump other problems to your ROC via GGUS**
 - Encourage users to make use of GGUS to report problems!
- **Check you're publishing your SRMs into BDII!**

- **Monitoring Tools**



- **All release information and guides:**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsRelease15>
- **All the tutorial material is at:**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsTutorial>
- **FTS procedures (upgrading, moving, cleaning)**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsProcedures15>
- **FTS FAQ**
 - <https://twiki.cern.ch/twiki/bin/view/EGEE/DMFtsSupport>
- **Workplan**
 - <https://twiki.cern.ch/twiki/bin/view/EGEE/DMFtsWorkPlan>
- **Support list**
 - support@ggus.org for user support
 - fts-support@cern.ch for administrator support (closed support)
 - fts-users@cern.ch for community support