



Enabling Grids for E-scienceE

FTS Administrators Tutorial for Tier-2s

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**



- **gLite File Transfer Service is a data movement fabric service**
 - It is a multi-VO service, used to balance usage of site resources according to VO and site policies

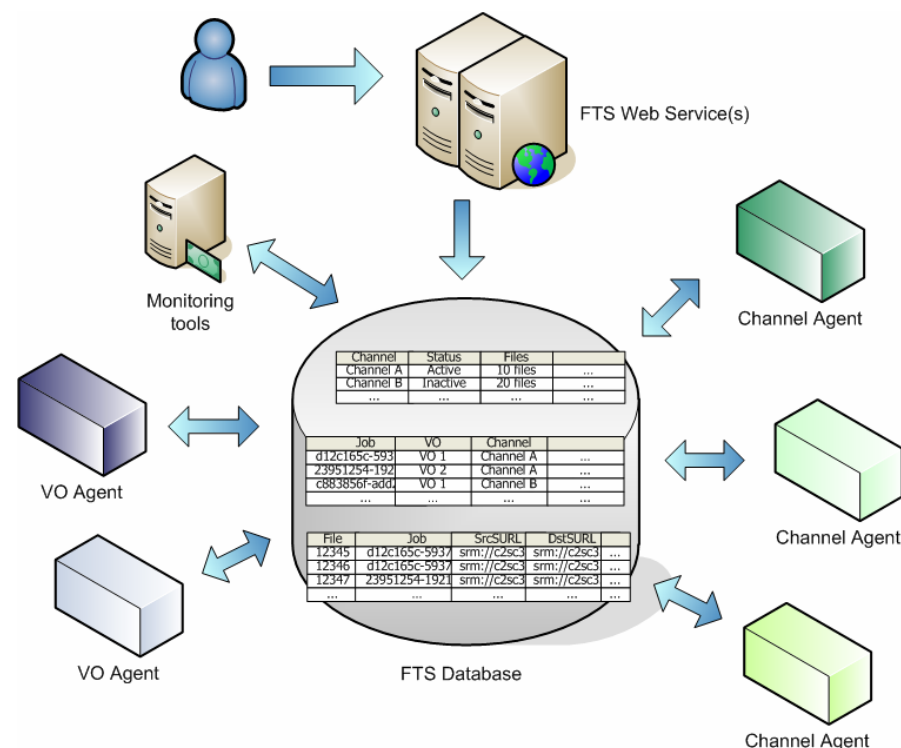
- **Why is it needed ?**
 - For the user, the service it provides is the reliable point to point movement of Storage URLs (SURLs)
 - For the site manager, it provides a reliable and manageable way of serving file movement requests from their VOs
 - For the VO manager, it provides ability to control requests coming from his users
 - Re-ordering, prioritization,...
 - **The focus is on the “service”**
 - **It should make it easy to do these things well**

- **FTS Web Services**
 - User: **FileTransfer**
 - Administration: **ChannelManagement**

- **File Transfer Queue**
 - Oracle DB

- **File Transfer Agents**
 - VO Agents
 - Channels Agents

- **Monitoring Tools**



- **For management ease, the service supports splitting jobs into multiple “channels”**
 - Submitted jobs are assigned to a suitable channel for serving
- **A channel may be:**
 - A point to point network link
 - Dedicated channels
 - (e.g. we manage all the T0 to T1 links in LCG on a separate channel)
 - Various “catch-all” channels
 - Non-dedicated channels
 - (e.g. everything else coming to me, or everything to one of my tier-2 sites)
 - More flexible channel definitions are on the way (but not there yet)
- **Channels are uni-directional**

- **“Channel”**: It’s not a great name
 - Isn’t tied to a physical network path
 - It’s just a management concept
 - “Queue” might be better 😊

- **All file transfer jobs on the same channel are served as part of the same queue**
 - Inter-VO priorities for the queue
 - Intra-VO priorities within a VO

- **Each channel has its own set of transfer parameters**
 - Number of concurrent transfers, number streams, TCP buffer size, etc

- **The WLCG model assigns each FTS sever responsibilities of what transfers it is supposed to manage**
 - Channels allow you to split up the management of the service as you see fit

- **What use-cases does FTS support and how do we deploy the servers to do this?**
- **Primary use-cases from experiment computing models**
 - tier-0 export to tier-1
 - tier-1 import into tier-0
 - tier-1 to tier-1 data transfer
 - tier-2 upload of data to associated tier-1: MC upload
 - tier-1 push of data to associated tier-2: AOD
- **Secondary use-cases**
 - non-associated tier-2 to tier-1: backup MC upload
 - non-associated tier-1 to tier-2: AOD?
 - tier-2 to tier-2: ?

- **Deploy only at tier-0 CERN and tier-1 sites**
 - Ease of operations – put the service where the support is
 - Simplifying the client job of “who do I submit to?”
 - Does lead to some “odd” channel definitions and introduces some restrictions on who can control what
- **Which FTS servers are responsible for which transfers?**
 - There are a few basic rules
- **Described in:**
 - <https://twiki.cern.ch/twiki/pub/LCG/FtsServerInstall15/SC4FTSsetupplan.doc>
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerDeployExampleTier1>

- 1. Tier-0:** If it involves CERN, then CERN's FTS does the transfer
 - This covers tier-0 to tier-1 and tier-1 to tier-0.
- 2. Tier-1 sites:** if you are the destination, your FTS is responsible for running the transfer
 - (if you're the source, the other end is responsible)
- 3. Tier-2 sites:** if you are the destination, your associated tier-1's FTS is responsible for running the transfer
 - Regardless of who is the source
- 4. We prioritise control of writing over control of reading**

- **Primary use-cases from experiment computing models**
 - tier-0 export to tier-1: CERN
 - tier-1 import into tier-0: CERN
 - tier-1 to tier-1 data transfer: the destination tier-1
 - tier-2 upload of data to associated tier-1: the associated tier-1
 - tier-1 push of data to associated tier-2: the associated tier-1
- **Secondary use-cases**
 - non-associated tier-2 to tier-1: the destination tier-1
 - non-associated tier-1 to tier-2: the tier-1 associated to the destination tier-2
 - tier-2 to tier-2: the tier-1 associated to the destination tier-2

- We prioritise control of writing over control of reading
- All **inbound transfers** are controlled by your associated tier-1 site
 - You can control this
- The majority of your **outbound transfer** will be upload to your associated tier-1 site
 - You can control this – if your T1 set up an explicit channel for you 😊
- You should expect some **outbound traffic** transferring to other tier-1 sites or to other tier-2 sites
 - These will be controlled by the other tier-1 sites
 - We offer no easy way to shut this off because of the way the channels are defined

- All inbound transfers are controlled by your associated tier-1 site
 - You can control this
- For each of its associated tier-2's, a tier-1 sets up a channel **STAR-TIER2**
 - This channel will match *any* source to your site
- Additionally, the associated tier-1 may wish to manage separately its transfers to the tier-2
 - This would be an explicit **TIER1-TIER2** channel
 - If this channel is not defined, traffic from your associated tier-1 will be matched on the general STAR-TIER2 channel

- **Upload to your tier-1. The tier-1 will define, either:**
 - **STAR-TIER1**. To match any traffic to the tier-1. In this case, you have limited control, since switching you off as a source switches off everyone else as well
 - Or... an explicit **TIER2-TIER1** channel. This allows you to control the reading from your site to your tier-1
- **Transfer to other tier-1s or tier-2s:**
 - This will be managed by the non-associated tier-1 server.
 - Although “possible”, it is not practical to expect that tier-1 to manage a distinct channel for upload from all potential tier-2 sites
 - So in practice, you have no control of these transfers

- **Admin CLI:**
 - use the ChannelManagement port-type
 - **glite-transfer-channel-***
- **User CLI:**
 - Use the FileTransfer port-type
 - **glite-transfer-***
- **Common Options:**
 - s FTS endpoint
 - v Verbose
 - h Help

- Site admins should use BDI to find the relevant FTS

```

root@fts103:~
[root@fts103 root]# glite-transfer-channel-list -v
# Using endpoint https://prod-fts-ws.cern.ch:8443/glite-data-transfer-fts/services/ChannelManagement
# Service version: 3.2.0
# Interface version: 3.1.0
# Schema version: 2.2.0
# Service features: glite-data-transfer-fts-3.2.0-15
# Client version: 3.3.0
# Client interface version: 3.1.0
ASCC-CERN
BARI-CERN
BNL-CERN
CERN-ASCC
CERN-BNL
CERN-CERN
CERN-DESY
CERN-FNAL
CERN-GRIDKA
CERN-IN2P3
CERN-INFN
CERN-NDGF
CERN-PIC
CERN-RAL
CERN-SARA
CERN-STAR
CERN-TRIUMF
DESY-CERN
FNAL-CERN
GRIDKA-CERN
IN2P3-CERN
INFN-CERN
NDGF-CERN
PIC-CERN
RAL-CERN
SARA-CERN
STAR-CERN
TRIUMF-CERN
[root@fts103 root]#
  
```

S
e
e
e
v
T
u
S

70

it's

Setting up the client for users

- T
- ir
- G

```

root@fts103 root]#
root@fts103 root]#
root@fts103 root]#
root@fts103 root]#
root@fts103 root]#
root@fts103 root]#
root@fts103 root]#
root@fts103 root]# glite-transfer-channel-list -s CERN-PROD
ASCC-CERN
BARI-CERN
BNL-CERN
CERN-ASCC
CERN-BNL
CERN-CERN
CERN-DESY
CERN-FNAL
CERN-GRIDKA
CERN-IN2P3
CERN-INFN
CERN-NDGF
CERN-PIC
CERN-RAL
CERN-SARA
CERN-STAR
CERN-TRIUMF
DESY-CERN

root@fts103 root]# glite-transfer-discovery CERN-PROD TAIWAN-LCG2
https://prod-fts-ws.cern.ch:8443/glite-data-transfer-fts/services/FileTransfer
root@fts103 root]# glite-transfer-submit -s https://prod-fts-ws.cern.ch:8443/glite-data-transfer-fts/services/FileTransfer ...
root@fts103 root]#
    
```


- **The initial channel setup at the tier-1**
- **Scheduled intervention on your SRM**
- **Unscheduled intervention on your SRM**
- **Changing the channel properties**

```

root@fts103:~
root@fts103:~
[root@fts103 root]# glite-transfer-channel-listmanagers CERN-ASCC
/C=CH/O=CERN/OU=GRID/CN=Andrea Sciaba 8968
/C=CH/O=CERN/OU=GRID/CN=Antonio Delgado Peris 4680
/C=CH/O=CERN/OU=GRID/CN=Gavin Mccance 1838
/C=CH/O=CERN/OU=GRID/CN=Harry Renshall 3515
/C=CH/O=CERN/OU=GRID/CN=James Casey 8576
/C=CH/O=CERN/OU=GRID/CN=Jean-Philippe Baud 7183
/C=CH/O=CERN/OU=GRID/CN=Maarten Litmaath 1689
/C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
/C=CH/O=CERN/OU=GRID/CN=Patricia Mendez Lorenzo 3183
/C=CH/O=CERN/OU=GRID/CN=Patricia Mendez Lorenzo-ALICE
/C=CH/O=CERN/OU=GRID/CN=Roberto Santinelli 7205
/C=CH/O=CERN/OU=GRID/CN=Simone Campana 7461
/C=CH/O=CERN/OU=GRID/CN=Sophie Lemaitre 2268
/C=IT/O=INFN/OU=Personal Certificate/L=Pisa/CN=Flavia Donno/Email=flavia.donno.infn.it
/C=TW/O=AS/OU=CC/CN=Chih-Chiang Chang/Email=ccchang@beta.wsl.sinica.edu.tw
/C=TW/O=AS/OU=CC/CN=Horng-Liang Shih/Email=hlshih@gate.sinica.edu.tw
/C=TW/O=AS/OU=CC/CN=Hung-Che Jen/Email=whifang@gate.sinica.edu.tw
/C=TW/O=AS/OU=PHYS/CN=Di Qing/Email=Di.Qing@cern.ch
/O=Grid/O=CERN/OU=cern.ch/CN=David Hugh Smith
[root@fts103 root]# █

```

e.g. `glite-transfer-channel-addmanager STAR-TIER2`
`"/C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032"`

- If the channel commands don't work:

```

pcitgm02.cern.ch - PuTTY
-bash-2.05b$ ./glite-transfer-getroles -s https://prod-fts-ws.cern.ch:8443/glite-data-transfer-fts/services/FileTransfer
You current clientDN is: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032

You are not authorised as Service Admin on this service.

You are authorised to submit to this service because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032

You are VO manager for 5 VOs.
You are VO manager for VO <alice> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are VO manager for VO <atlas> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are VO manager for VO <cms> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are VO manager for VO <dteam> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are VO manager for VO <lhcb> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032

You are channel manager for 26 channels.
You are channel manager for channel <TRIUMF-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <STAR-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-FNAL> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-GRIDKA> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-IN2P3> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <ASCC-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <BNL-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-ASCC> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-BNL> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-DESY> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-INFN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-NDGF> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-PIC> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-RAL> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-SARA> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <CERN-TRIUMF> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <DESY-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <FNAL-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <GRIDKA-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <IN2P3-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <INFN-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <NDGF-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <PIC-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <RAL-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
You are channel manager for channel <SARA-CERN> because your cert contains the following principal: /C=CH/O=CERN/OU=GRID/CN=Paolo Badino 3032
-bash-2.05b$

```

```

root@fts103:~
[root@fts103 root]# glite-transfer-channel-list CERN-ASCC
Channel: CERN-ASCC
Between: CERN-PROD and TAIWAN-LCG2
State: Active
Contact: lcg-sc@lists.grid.sinica.edu.tw
Bandwidth: 10
Nominal throughput: 0
[root@fts103 root]# glite-transfer-channel-setvoshare CERN-ASCC dteam 10
[root@fts103 root]# glite-transfer-channel-list CERN-ASCC
Channel: CERN-ASCC
Between: CERN-PROD and TAIWAN-LCG2
State: Active
Contact: lcg-sc@lists.grid.sinica.edu.tw
Bandwidth: 10
Nominal throughput: 0
Number of files: 65, streams: 20
Number of VO shares: 4
VO 'dteam' share is: 10
VO 'atlas' share is: 100
VO 'cms' share is: 100
VO 'ops' share is: 100
[root@fts103 root]#

```

- The share is decided as:
- Find all the VOs with pending jobs. Total their shares. e.g.
 - LHCb and dteam have jobs
 - Atlas has no job
 - Total = 100 + 25 = 125
 - dteam gets $25/125 = 1/5$
 - LHCb get $100/125 = 4/5$
- The share is calculated point-in-time w.r.t. current jobs. It does not use past data.

LHCb	100
Atlas	100
dteam	25

- **I'm fixing / upgrading / etc my SRM (scheduled intervention). How to stop all transfers:**
 - Know in advance which of your tier-1 channels affects you (and which you are able to control)
 - 10 to 15 minutes before the intervention:
 - Inform your tier-1 that you are pausing the channel(s)
`glite-transfer-channel-set -S Inactive <CHANNEL_NAME>`
- **After the intervention:**
 - Inform your tier-1 that you are restarting the channel(s)
`glite-transfer-channel-set -S Active <CHANNEL_NAME>`

- **My SRM has just gone down (or otherwise become unavailable). Unscheduled intervention:**
 - The same procedure as a scheduled intervention. If the SRM is down, set the channels that involve you directly to **Inactive**.
 - Do this **as soon as you can** after the problem with the SRM is noticed to minimize the Failed jobs
 - you may find that your tier-1 site has noticed the problem first and switched you off already
 - Inform the tier-1 site that you are doing it.
 - Set the channels back **Active** once the problem is fixed, informing the tier-1.

- **A dedicated pool node is unavailable (or you need to schedule an intervention). How to stop all transfers for just one VO:**

- The same procedures described before, but you can turn off a VO
- Instead of modify the channel status, you can set the VO share to 0.

`glite-transfer-channel-list <CHANNEL_NAME>`

(in order to get the current value of the share)

`glite-transfer-channel-setvoshare <CHANNEL_NAME> <VO_NAME> 0`

- Inform the tier-1 site that you are doing it.
- Set the share back to the previous value once the problem is fixed, informing the tier-1.

`glite-transfer-channel-setvoshare <CHANNEL_NAME> <VO_NAME>
<OLD_SHARE_VALUE>`

- **Changing the transfer rate on the channel:**
 - If you need to lower or increase the transfer rate to your site
 - Again, know in advance which of your tier-1 channels affects you (and which you are able to control). This is the same as those you can switch Active/Inactive
 - Check the current number of concurrent files:


```
glite-transfer-channel-list <CHANNEL_NAME>
```
 - Set the number of concurrent files, as you need:


```
glite-transfer-channel-set -f 20 <CHANNEL_NAME>
```
 - You can also change the number of streams for URLCOPY transfers


```
glite-transfer-channel-set -T 5 <CHANNEL_NAME>
```
- **Usually you need to experiment a bit to find the optimum setting**

- **FTS is a Reliable Data Movement service and a Management tool**
 - Channel concept
- **You can control your inbound transfers and (most) outbound transfers**
- **The majority of your transfers should be controlled by your Tier-1's FTS server**
- **The client tools should be setup using BDII**
- **Reviewed common service operations**



- **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

- **Agreed at recent FTS workshop (Sept. 2006)**
- **Two main checkpoints:**
 - Delegation + Schema updates
 - Code is mostly ready now, but not backwards compatible
 - Need to work on backwards compatibility on the client level
 - *For internal test end November / start December*
 - SRM v2.2
 - Should be ready for test end November
 - *Will be tested with new SRM implementations on validation cluster from end November*
 - But expect several incremental releases as we track and understand the SRM implementations

- **Development focus is service stability and making the service easier to run**
 1. Improve the service reporting and monitoring
 2. SRM/gridFTP communication split (improves the stability)
 3. More flexible channel definitions
 - To make it easier to meet the needs of CMS and Alice
 4. Site blacklisting
 - To avoid clogging up shared channels with bad sites