

Introduction to FTS

Paolo Tedesco (paolo.tedesco@cern.ch)

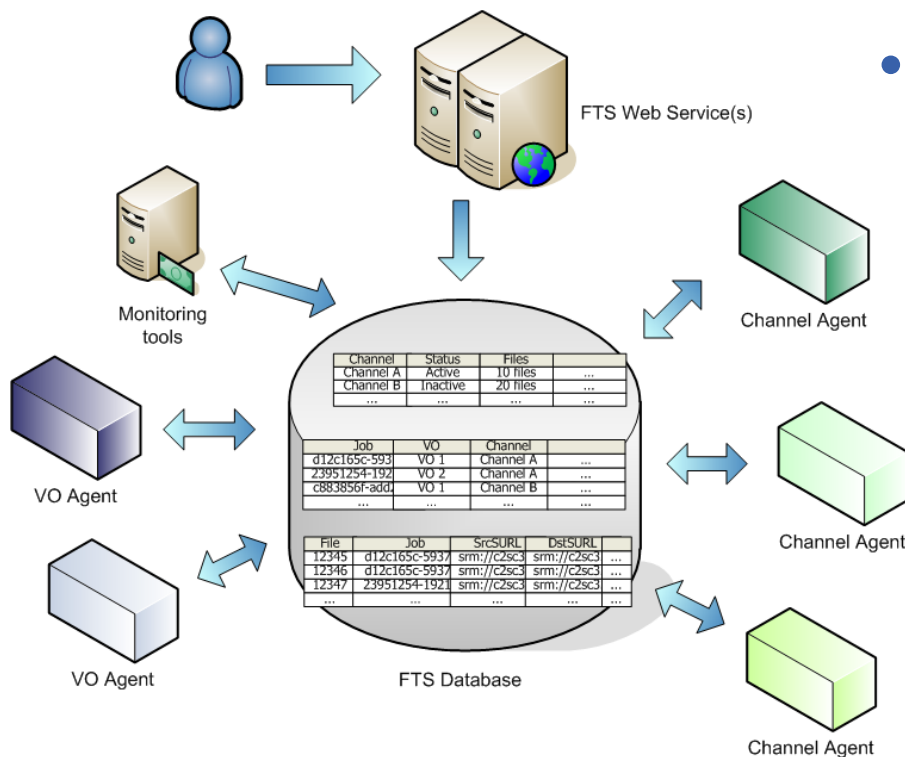
White Area Lecture

3 October 2008

- File Transfer Service (FTS) is a data movement service
 - Used by experiment frameworks
 - High volume data streams need to be managed
- Balance site resources usage
- Prevent network overload
- Prevent storage overload
- Jobs prioritization
- Service monitoring and statistics

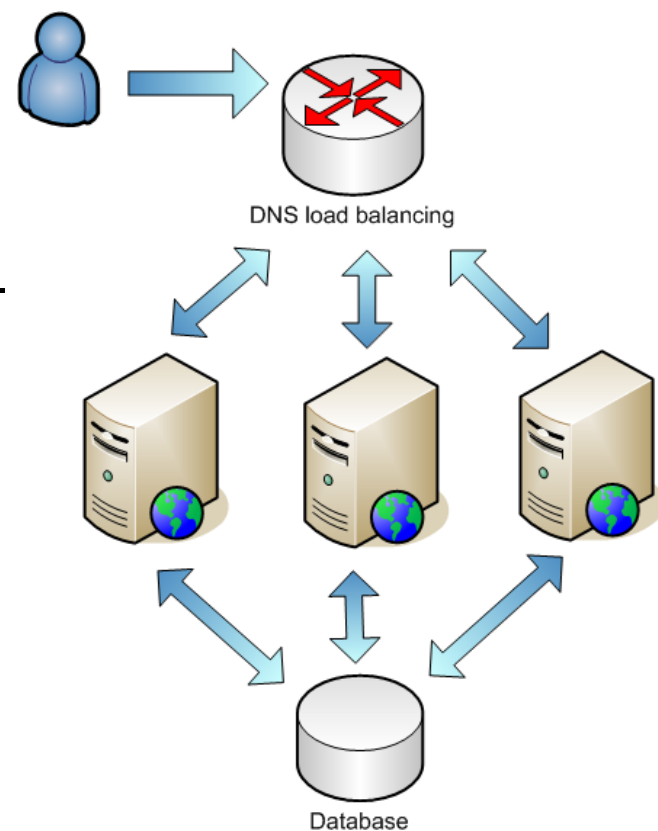
- Single direction management queue for transfer jobs
- Not tied to a physical network path
- An endpoint (source or destination) can be defined as:
 - single site
 - CERN - RAL
 - site group (formerly known as “cloud”)
 - CERN - [EU], where [EU] = INFN, IN2P3, PIC ...
 - catch all
 - CERN - *

- A channel defines:
 - Transfer protocol
 - gridftp, srmcopy
 - Transfer limits
 - maximum number of concurrent transfers
 - Transfer parameters
 - TCP streams, buffer size
 - timeouts
 - VO shares
 - “quota” of a VO as percentage
 - share calculation mechanism (absolute, normalized, normalized on active)
 - Transfer priorities
 - among the same VO



- Decoupled components
 - Web service
 - VO agents
 - Channel agents
 - Monitoring service

- Stateless
- Load-balanced
 - scalability
 - software upgrade with zero user-visible downtime
 - graceful failover if one node dies
- APIs
 - Job submission / tracking
 - Service / channel management



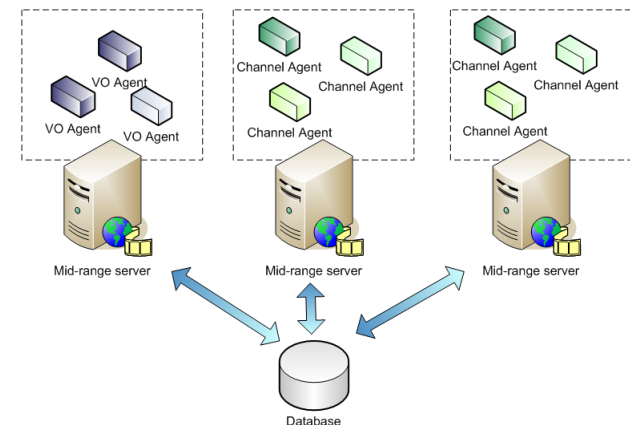
- Daemons that run a set of tasks
- Each task operates on a particular state



Allocate: find a channel between the source and the destination site where the VO has a share.

Fetch: select a transfer according to shares and priorities and start it.

- One agent per VO / channel
- Independent from each other
- Split across multiple nodes



- File Transfer Monitor (FTM) node
 - Schema additions for FTS database
 - Periodically queries FTS database
 - Creates weekly/daily/hourly summaries
 - Publishes summaries into gridview
- Currently publishing data rates for site to site transfers
- Pluggable architecture
 - Incremental approach
 - Easy to extend
 - Add external contributions

- Transfers are run using the clients' X509 credentials
 - delegated by the client to the service (impersonation)
- Full audit on all administrative operations
 - i.e. modifying channel parameters, VO shares, adding and removing VO or channel managers
 - all operations recorded in the logs, not (yet) in the DB
- VOMS credentials (attribute certificates) used (and renewed as necessary) in FTS ≥ 2.0

- Roles
 - VO production manager
 - Channel administrator
 - Service manager
- Machine certificate of web service is by default a service manager
- Get a list of your roles with `glite-transfer-getroles`

```
~ $ glite-transfer-getroles
```

```
Your current clientDN is: /DC=ch/DC=cern/.../CN=Paolo Tedesco
```

```
You are authorised to submit to this service because your cert contains the following principal: /DC=ch/DC=cern/.../CN=Paolo Tedesco
```

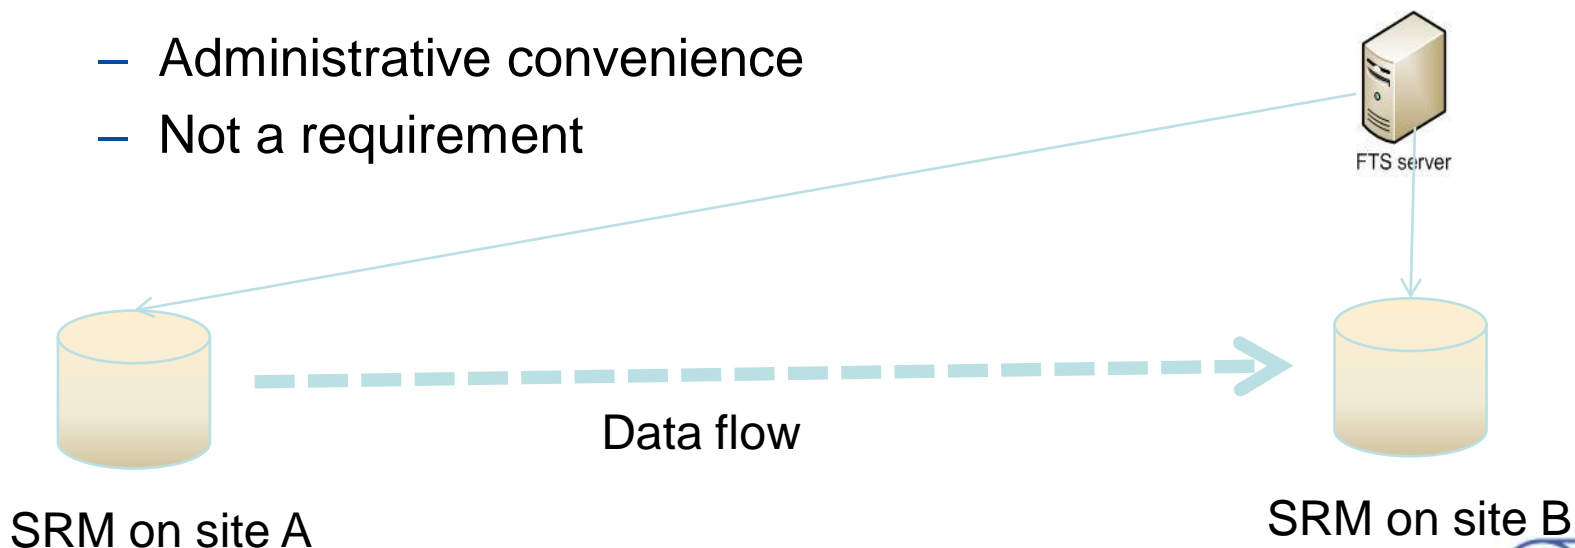
```
You are VO manager for 1 VOs.
```

```
You are VO manager for VO <dteam> because your cert contains the following principal: /DC=ch/DC=cern/.../CN=Paolo Tedesco
```

```
You are channel manager for 1 channels.
```

```
You are channel manager for channel <CERN-CERN> because your cert contains the following principal: /DC=ch/DC=cern/.../CN=Paolo Tedesco
```

- FTS uses the SRM interface to set up/run/finalize the transfer
- Actual copy
 - 3rd party gridFTP
 - SrmCopy
- FTS server is typically located on either the source of destination site
 - Administrative convenience
 - Not a requirement



SRM on site A

SRM on site B

- File as plugin
 - /opt/glite/etc/services.xml
- Cache BDII information
 - mostly 'static'
 - avoid heavy queries to information system
 - add resilience to glitches
- Updated by the sd2cache package
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsServicesXmlGliteSd2Cache>

```

<service name="http://lxdpm101.cern.ch:8446/srm/managerv2">
  <parameters>
    <endpoint>http://lxdpm101.cern.ch:8446/srm/managerv2</endpoint>
    <type>SRM</type>
    <version>2.2.0</version>
    <site>CERN-PROD</site>
    <wsdl>unset</wsdl>
    <volist>
      <vo>atlas</vo>
      <vo>cms</vo>
      <vo>dteam</vo>
    </volist>
    <param name="atlas:SEMMountPoint">/dpm/cern.ch/home/atlas</param>
    <param name="cms:SEMMountPoint">/dpm/cern.ch/home/cms</param>
    <param name="dteam:SEMMountPoint">/dpm/cern.ch/home/dteam</param>
  </parameters>
</service>

```

Check if the version is published correctly (No SRM Method factory...)

This is the site used for channel resolution (No channel found...)

Check that the entry for the service is defined and contains your vo in the volist (No site for host...)

- Agents Configuration files

```
(root@fts109) /opt/glite/etc/glite-data-transfer-agents.d $ ls  
  
glite-transfer-channel-agent-srmcopy-CERN-FNAL.log-properties  
glite-transfer-channel-agent-srmcopy-CERN-FNAL.properties.xml  
glite-transfer-channel-agent-urlcopy-CERN-NDGF.log-properties  
glite-transfer-channel-agent-urlcopy-CERN-NDGF.properties.xml  
glite-transfer-vo-agent-DTEAM.log-properties  
glite-transfer-vo-agent-DTEAM.properties.xml
```

- under /opt/glite/etc/glite-data-transfer-agents.d
- agent type: channel / vo
- transfer type: urlcopy / srmcopy (channel agents)
- channel or VO name
- properties and log-properties

- xml files loaded by agents at startup
- database connection parameters
- channel and VO settings
- typically configured via YAIME
- check installation and configuration guides
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerInstall21>
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsYaimValues21>
- many parameters moved to database

- Web service interfaces:

```
https://prod-fts-ws.cern.ch:8443/glite-data-transfer-fts/services/
```

- ChannelManagement

- glite-transfer-channel-add, glite-transfer-channel-setvolimit...

- FileTransfer

- glite-transfer-submit, glite-transfer-list...

- gridsite delegation

- glite-delegation-init, glite-delegation-destroy...

- Service location (-s)
- Service details (-v)
- Verbose flag (--verbose)

```
~ $ glite-transfer-submit -s ../FileTransfer -v --verbose ...  
# Using endpoint https://prod-fts-ws.cern.ch:8443/glite-data-  
transfer-fts/services/FileTransfer  
# Service version: 3.3.2  
# Interface version: 3.3.0  
# Schema version: 3.1.0  
# Service features: glite-data-transfer-fts-3.3.2-  
# Client version: 3.4.2  
# Client interface version: 3.3.0  
Server supports delegation. Delegation will be used by default.  
Remaining time for local proxy is 11 hours and 59 minutes.  
Not bothering to do delegation, since the server already has a  
delegated credential for this user lasting longer than 4 hours.  
Remaining time on server for this credential is 11 hours and 57  
minutes.  
7f3bbf5f-896f-11dd-824c-e41662be838f
```

service details

verbose output

job id

create a proxy

```
~ $ voms-proxy-init ...
```

submit job

```
~ $ glite-transfer-submit
```

```
srm://srm-dteam.cern.ch:8443/srm/managerV2?SFN=/foo
srm://srm-dteam.cern.ch:8443/srm/managerV2?SFN=/bar
7f3bbf5f-896f-11dd-824c-e41662be838f
```

```
~ $ glite-transfer-status -l 7f3bbf5f-896f-11dd-824c-
e41662be838f
```

```
Failed
```

job state

poll status

```
Source:      srm://...?SFN=/foo
```

```
Destination: srm://...?SFN=/bar
```

```
State:       Failed
```

```
Retries:     1
```

```
Reason:      SOURCE error during PREPARATION phase:
```

```
[GENERAL_FAILURE] blah blah...
```

```
Duration:    0
```

file details (-l option)

```
~ $ glite-transfer-submit ...
```

```
212d1116-8a19-11dd-ad25-efb5df4948c3
```

restrict on own DN

```
~ $ glite-transfer-list -u "$MY_DN"
```

```
212d1116-8a19-11dd-ad25-efb5df4948c3
```

Ready

restrict on VO

```
~ $ glite-transfer-list -o dteam
```

```
009efc11-94db-11dc-af46-d1a42b7e2b52
```

Ready

```
00dea939-9737-11dc-af46-d1a42b7e2b52
```

Ready

```
0161bdbf-94bb-11dc-bfac-c534861988d9
```

Ready

```
0383706c-9412-11dc-bfac-c534861988d9
```

Ready

```
03ad21ee-9498-11dc-b96e-86619d08b795
```

Ready

```
042d77f9-973d-11dc-bfac-c534861988d9
```

Ready

```
0573c634-94bf-11dc-bfac-c534861988d9
```

Ready

```
0a789eea-8de0-11dc-bfac-c534861988d9
```

Active

```
0b1ddf3d-8de0-11dc-b96e-86619d08b795
```

Active

```
0f83bea9-9850-11dc-af46-d1a42b7e2b52
```

Ready

```
~ $ glite-transfer-channel-list -s ../ChannelManagement
```

```
CERN-CERN
```

```
CERN-ASGC
```

```
...
```

```
~ $ glite-transfer-channel-list -x CERN-CERN
```

```
Channel: CERN-CERN
```

```
Between: CERN-PROD and CERN-PROD
```

```
State: Active
```

```
Contact: service-challenge-cern@cern.ch
```

```
Bandwidth: 0
```

```
Nominal throughput: 10
```

```
Number of files: 20, streams: 1
```

```
TCP buffer size: default
```

```
Message: Restarted for Castor interve
```

```
Last modification by:
```

```
/DC=ch/DC=cern/OU=computers/CN=fts114.cern.ch
```

```
Last modification time: 2008-01-23 15:10:12
```

```
Number of VO shares: 6
```

```
VO 'alice' share is: 100 and is limited to 5 transfers
```

```
VO 'atlas' share is: 100 and is limited to 20 transfers
```

```
...
```

detailed listing

source and dest sites

max num of transfers

VO shares and limits

create channel

```
~ $ glite-transfer-channel-add -f 20 -T 1 -S Active  
CERN-CERN CERN-PROD CERN-PROD
```

set shares and limits

```
~ $ glite-transfer-channel-setvoshare CERN-CERN alice 100  
~ $ glite-transfer-channel-setvolimit CERN-CERN alice 5
```

list and add managers

```
~ $ glite-transfer-channel-listmanagers CERN-CERN  
  
/DC=ch/DC=cern/OU=Organic Units/.../CN=Paolo Badino  
/DC=ch/DC=cern/OU=Organic Units/.../CN=Gavin Mccance  
  
~ $ glite-transfer-channel-addmanager CERN-CERN  
/DC=ch/DC=cern/OU=Organic Units/.../CN=Paolo Tedesco
```

```
~ $ glite-transfer-listvomangers -s ../FileTransfer  
dteam
```

FileTransfer interface

```
/DC=ch/DC=cern/OU=Organic Units/.../CN=Paolo Badino  
/DC=ch/DC=cern/OU=Organic Units/.../CN=Gavin Mccance
```

```
~ $ glite-transfer-addvomanager dteam  
/DC=ch/DC=cern/OU=Organic Units/.../CN=Paolo Tedesco
```

requires VO admin role

```
~ $ glite-transfer-setpriority 7f3bbf5f-896f-11dd-824c-  
e41662be838f 5
```

fts-support@cern.ch

Configuration issues, software problems...

ftscern.support@cern.ch

Operational issues related to CERN FTS services.

fts-users@cern.ch

Open list. Announcements are also sent to this list.