



Enabling Grids for E-scienceE

File Transfer Service

Paolo Tedesco (paolo.tedesco@cern.ch)

SRM2.2 Deployment Workshop

14 November 2007

www.eu-egee.org




Information Society
and Media



- **File Transfer Service (FTS) is a data movement service**
- **Balance site resources usage**
- **Prevent network overload**
- **Prevent storage overload**
- **Jobs prioritization**
- **Service monitoring and statistics**

- SRM 2.2 support
- Improved monitoring capabilities
 - Reliability of the 'overall transfer service'
- Certificate delegation
 - Security issue: no more MyProxy passphrase in the job
- Better database model
 - Improved performance and scalability
- Better administration tools
 - Make it easier to run the service
- Placeholders for the future functionality
 - Minimise the impact of future upgrades

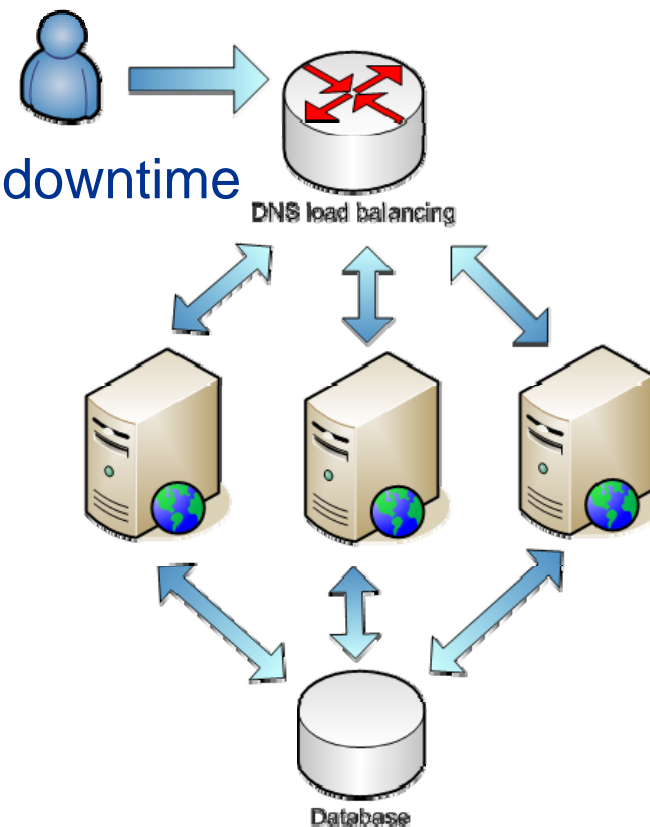
- **Single direction management queue for transfer jobs**
 - Point to point (between two sites)
 - Catch all
 - More complex groupings on the way
- **Not tied to a physical network path**
- **A channel defines:**
 - Transfer protocol (gridftp, srmcopy)
 - Transfer parameters
 - VO shares
 - Transfer priorities



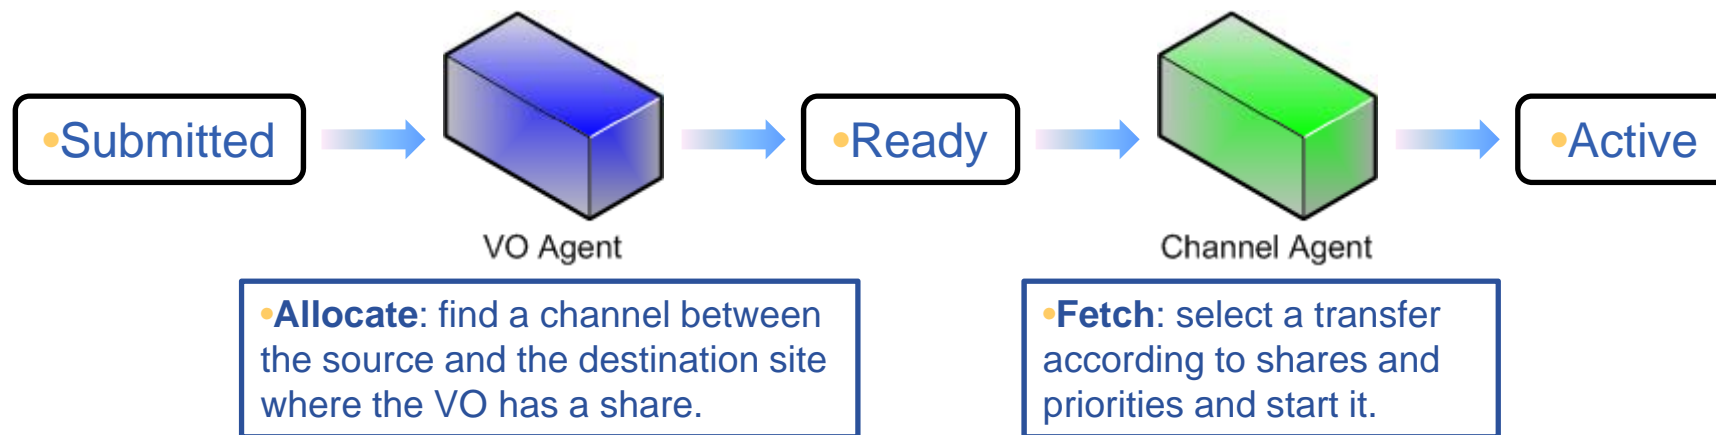
- Decoupled components

- 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
- **Job submission / tracking API**
- **Service / channel management API**



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




- One agent for each VO and 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 the FTM contains:**
 - A Gridview publisher for site-to-site data rates and volumes
 - To verify agreed service levels
 - Basic summaries about observed errors per channel / VO
 - The more detailed information views are still being prototyped
- **Pluggable architecture**
 - Incremental approach
 - Easy to extend
 - Add external contributions

- Used at tier0-export for sending GGUS tickets daily to problematic sites
- Once FTM is deployed, all sites can use this tool to get an 'overview' of what the FTS has measured about them
- Summary data will be published in WLCG monitoring framework



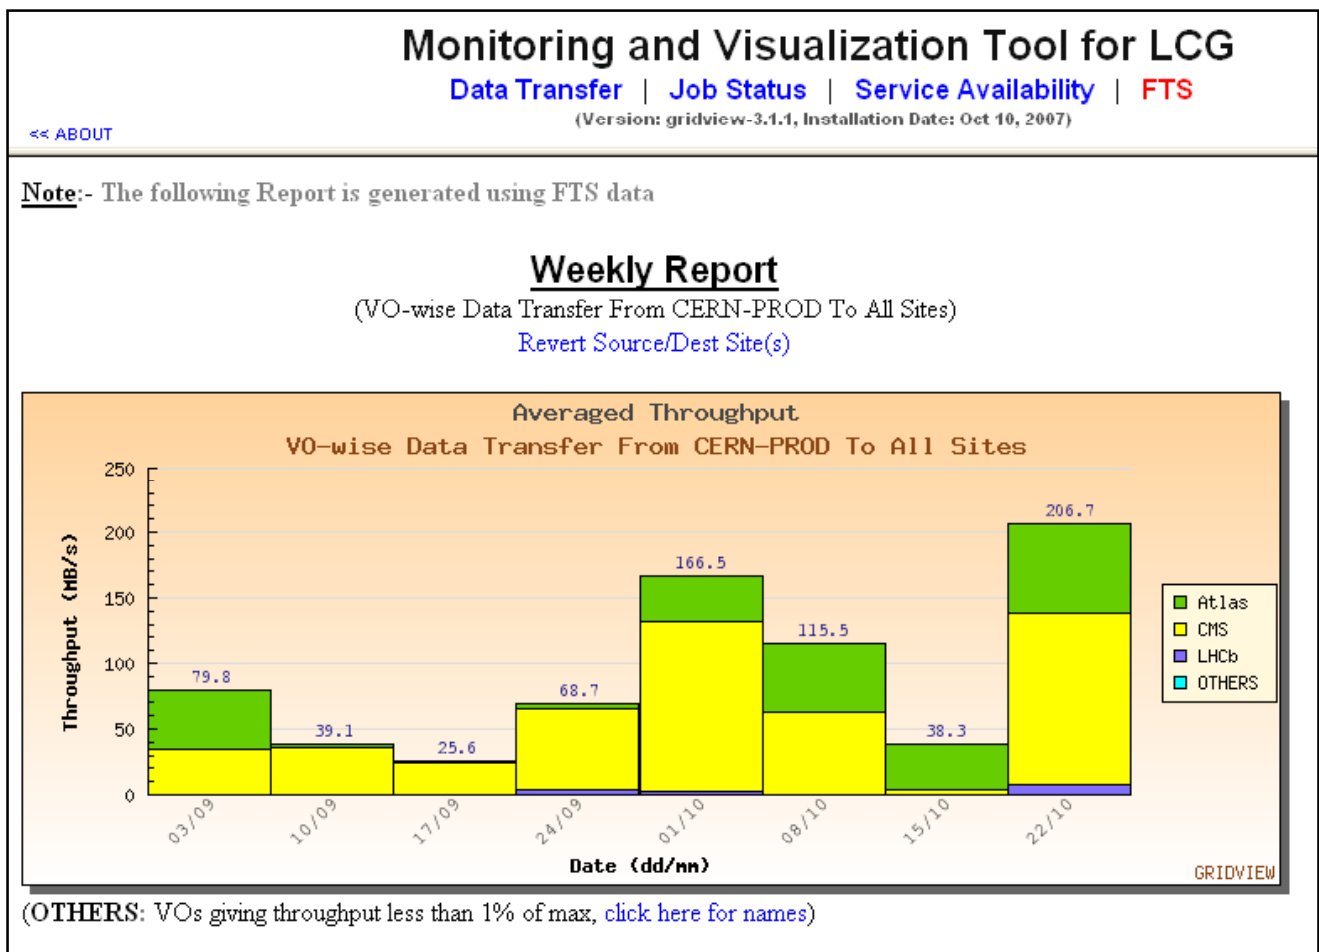
FTS Report

Statistics concerning all the transfers performed during last week:
Between 2007-09-24 00:00:00 +02:00 and 2007-10-01 00:00:00 +02:00

Channel Name	VO Name	Total	% Failures	# Succ.	# Fail.	1st Failure Reason	% 1st Failure Reason	2nd Failure Reason	% 2nd Failure Reason	Avg. Size (KB)	Avg. Duration (sec)	Avg. Tx Rate (MB/sec)	Avg. Tx Rate (Gbit)	Tx Bytes (Gbit)
CERN-BNL	[AII]	12	100	0	12	SOURCE_PREP_SRM_ERROR	100		0	0	0	0	0	0
CERN-CERN	[AII]	375	84	61	314	SOURCE_PREP_SRM_ERROR	1	99	DEST_PREP_SRM_ERROR	1	0.06	72.61	1.2	3.36
CERN-INFN	[AII]	62844	72	17482	45362	DEST_PREP_SRM_ERROR	52	SOURCE_PREP_SRM_ERROR	44	0.49	94.36	3.22	8581.73	8926.87
CERN-SARA	[AII]	31745	71	8343	22402	DEST_PREP_SRM_ERROR	59	SOURCE_PREP_SRM_ERROR	38	0.05	42.8	1.1	465.01	466.64
atlas	[AII]	15309	53	7184	8116	SOURCE_PREP_SRM_ERROR	78	DEST_PREP_SRM_ERROR	22	0	39.07	0.02	3.57	3.58
hep	[AII]	95445	87	2159	93286	DEST_PREP_SRM_ERROR	81	SOURCE_PREP_SRM_ERROR	15	0.21	55.23	4.72	461.44	463.96
CERN-B2P3	[AII]	27668	48	14522	13144	SOURCE_PREP_SRM_ERROR	87	DEST_PREP_SRM_ERROR	11	0.41	184.75	3.81	598.84	599.33
CERN-FIC	[AII]	21802	47	16775	15027	SOURCE_PREP_SRM_ERROR	87	DEST_PREP_SRM_ERROR	9	0.3	100.21	1.27	5107.95	5193.69
CERN-INDO	[AII]	8710	38	5571	3139	SOURCE_PREP_SRM_ERROR	85	Transfer	14	0.05	169.38	0.52	253.77	254.52
CERN-RAL	[AII]	23429	33	15678	7751	DEST_PREP_SRM_ERROR	51	SOURCE_PREP_SRM_ERROR	43	0.21	78.68	0.65	3272.35	3279.81
CERN-GRIDKA	[AII]	19025	29	13582	5443	SOURCE_PREP_SRM_ERROR	58	DEST_PREP_SRM_ERROR	31	0.01	69.65	0.25	178.19	178.71
CERN-TRIUMF	[AII]	31070	26	23096	7974	SOURCE_PREP_SRM_ERROR	96	Source SRM	3	0.01	48.04	0.13	145.92	146.27
CERN-ABCC	[AII]	2012	20	1602	411	DEST_PREP_SRM_ERROR	75	Transfer	12	2.46	1887.81	1.44	2947.83	4055.84
CERN-FNAL	[AII]	11429	7	10684	765	Transfer	76	REMOTE_ERROR	19	1.14	232.8	4.75	12104.6	12104.6

Click on the Channel Name to show the VO details

- Transfer volume / rates will be available for entire transfer network



- Prototyping on FTS pilot service
 - e.g. TOP errors per-site

TOP errors for CERN-INFN for last day sorted by total amount			
place	error_id	number	error sample
1	23	34	421 Timeout (900 seconds): closing control connection
2	35	8	425 Can't open data connection, timed out() failed.
3	67	4	CastorStagerInterface.c:2507 Device or resource busy (errno=[num], serrno=[num])
TOP errors for CERN-GRIDKA for last day sorted by total amount			
place	error_id	number	error sample
1	124	42	426 Transfer aborted, closing connection :failed in recvEBlock
2	41	16	failed to contact on remote SRM [srm]. Givin up after 3 tries
3	81	16	failed to prepare Destination file in 180 seconds
TOP errors for CERN-IN2P3PPS for last day sorted by total amount			
place	error_id	number	error sample

- Soon to have ability to extract example transfer logs for your top problems

- **Transfers are run using the clients' X509 credentials**
 - delegated by the client to the service (impersonation)
- **Full audit on all operations**
- **VOMS credentials (attribute certificates) used (and renewed as necessary) in FTS 2.0**
- **Roles**
 - VO manager
 - List all the transfers of the VO
 - Change transfer priorities
 - Channel manager
 - Change the state of the channel (e.g. for maintenance)
 - Manage VO shares on the channel
 - List all the transfers on the channel

- **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

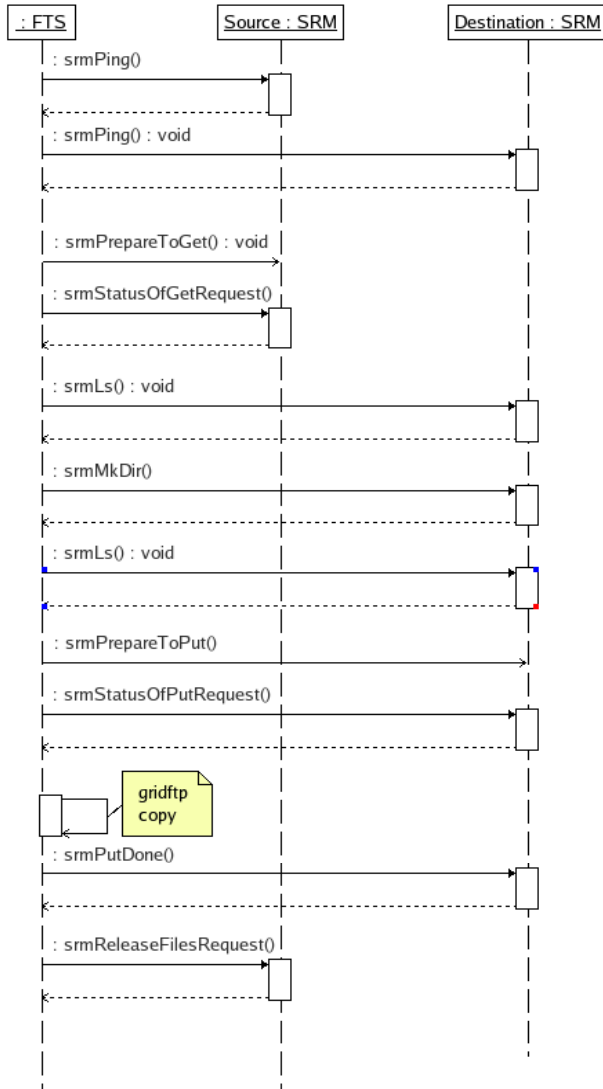
You are VO manager for 1 VOs.
You are VO manager for VO <dteam> because your cert contains the following
principal: /DC=ch/DC=cern/.../CN=Paolo Tedesco

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 or destination site**
 - Administrative convenience
 - Not a requirement

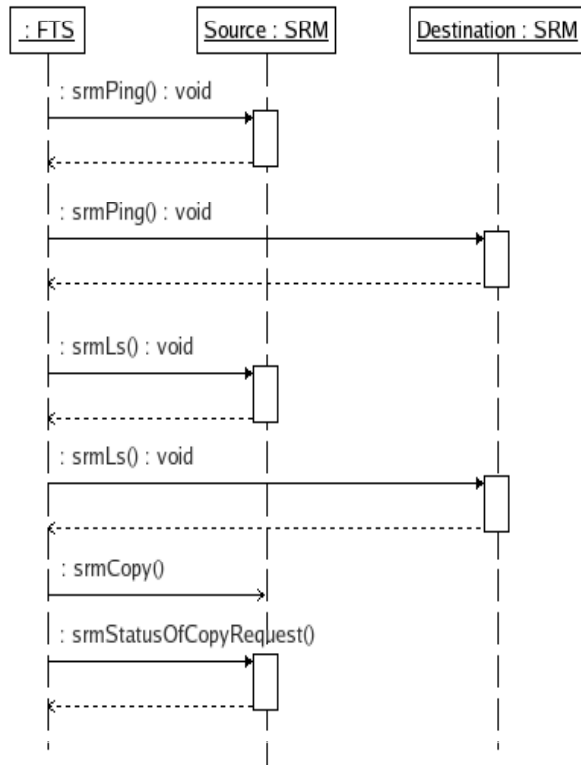


Urlcopy transfer



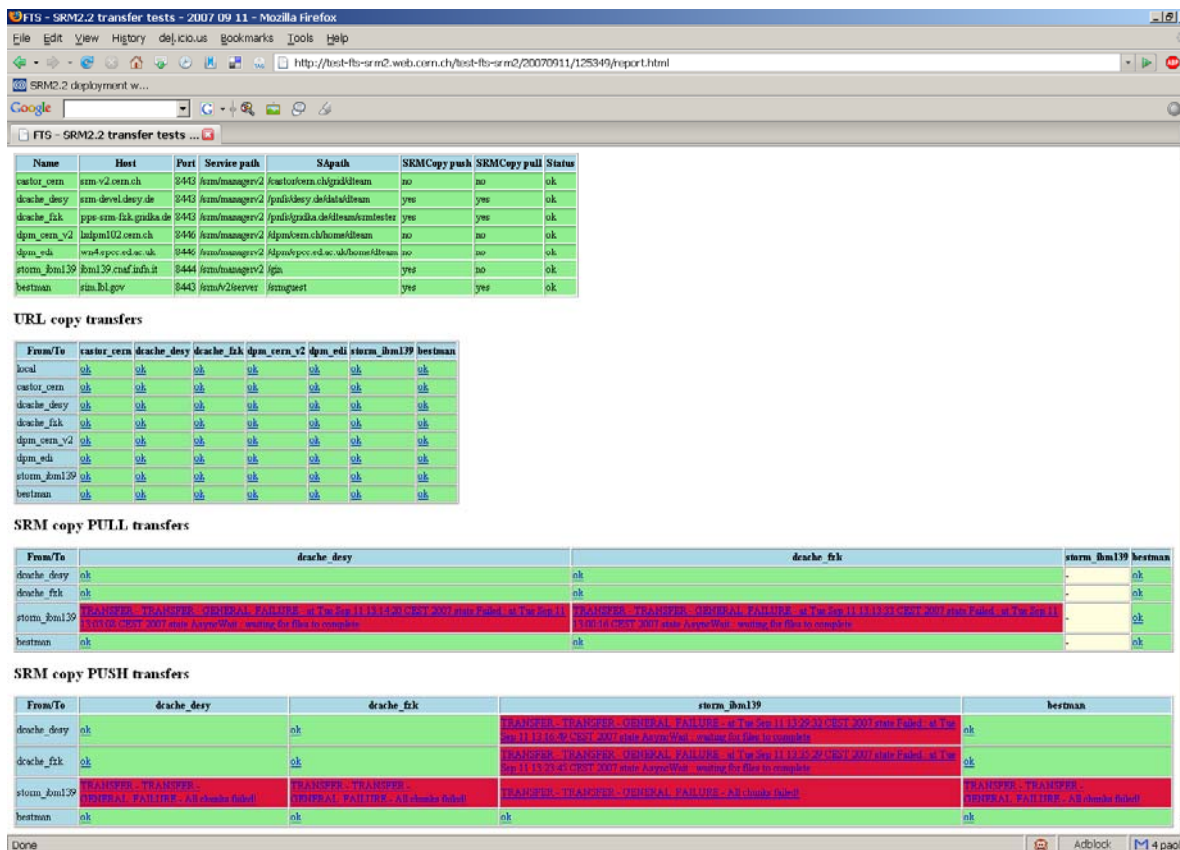
- **Ping source and destination**
- **PrepareToGet**
- **Ls + MkDir: destination directory**
 - ensure existence of target directory
- **Ls: destination file**
 - used for cleanup
- **PrepareToPut**
- **Copy using gridftp**
- **PutDone**
- **ReleaseFilesRequest**

Srmcopy transfer



- **Ping source and destination**
- **Ls + Mkdir: destination directory**
 - ensure existence of target directory
- **Ls: source file**
 - make sure source file exists
- **Ls: destination file**
 - used for cleanup
- **Copy using SrmCopy**

- <http://test-fts-srm2.web.cern.ch/test-fts-srm2/>
- Set of transfers using a list of endpoint definitions
 - urlcopy
 - srmcopy push
 - srmcopy pull
- Transfer logs
- soap messages



FTS - SRM2.2 transfer tests - 2007/09/11 - Mozilla Firefox

SRM2.2 deployment w...

Name	Host	Port	Service path	S4path	SRMCopy push	SRMCopy pull	Status
castor_orm	orm.v2.cern.ch	8443	/srm/manager/2	/castor/orm.ch/gladstream	no	no	ok
dcache_desy	stm-devel.desy.de	8443	/srm/manager/2	/ndf-kassy.de/data/stream	yes	yes	ok
dcache_fbk	ppp-srm.fbk.gsi.it	8443	/srm/manager/2	/ndf-kassy.de/data/stream	yes	yes	ok
dpqm_orm_v2	isipm102.cern.ch	8446	/srm/manager/2	isipm102.cern.ch/home/stream	no	no	ok
dpqm_edi	wa4-spc.ed.ac.uk	8446	/srm/manager/2	isipm102.cern.ch/home/stream	no	no	ok
storn_ilm139	ilm139.cisf.infn.it	8444	/srm/manager/2	/srm	yes	no	ok
bestman	jan.kil.gov	8443	/srm/2/server	/srmgsst	yes	yes	ok

URL copy transfers

From/To	castor_orm	dcache_desy	dcache_fbk	dpqm_orm_v2	dpqm_edi	storn_ilm139	bestman
local	ok	ok	ok	ok	ok	ok	ok
castor_orm	ok	ok	ok	ok	ok	ok	ok
dcache_desy	ok	ok	ok	ok	ok	ok	ok
dcache_fbk	ok	ok	ok	ok	ok	ok	ok
dpqm_orm_v2	ok	ok	ok	ok	ok	ok	ok
dpqm_edi	ok	ok	ok	ok	ok	ok	ok
storn_ilm139	ok	ok	ok	ok	ok	ok	ok
bestman	ok	ok	ok	ok	ok	ok	ok

SRM copy PULL transfers

From/To	dcache_desy	dcache_fbk	storn_ilm139	bestman
dcache_desy	ok	ok	ok	ok
dcache_fbk	ok	ok	ok	ok
storn_ilm139	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	ok	ok
bestman	ok	ok	ok	ok

SRM copy PUSH transfers

From/To	dcache_desy	dcache_fbk	storn_ilm139	bestman
dcache_desy	ok	ok	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	ok
dcache_fbk	ok	ok	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	ok
storn_ilm139	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete	TRANSFER - TRANSFER - GENERAL FAILURE - At Tue Sep 11 11:53:59 CEST 2007 state Failed - At Tue Sep 11 11:53:59 CEST 2007 state AsyncWait - waiting for file to complete
bestman	ok	ok	ok	ok

- **Things to do to improve the overall service delivered to the users**
- **Manage the channels that involve your site**
 - even if the FTS server itself is on another site
- **Understanding service failures**
 - Failures of the FTS service
 - Failures of the underlying SRM services

- **For tier-1 and tier-2 sites**
 - A lot of the transfers will be coming from an FTS at another site
 - Know who the (other) FTS service administrators are
 - Know which channels have been setup which involve your site
 - Make sure that that your local admin(s) have been set as channel managers for these channels
- **Follow intervention procedures**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsProcedures20>
 - Especially for ‘delicate’ operations (changing channel type...)
- **Pause the channels if needed**
 - Interventions (scheduled or not) on the SRM
 - Inform your T1 that you are pausing/restarting the channel
 - `glite-transfer-channel-set -S Inactive CHANNELNAME`
 - `glite-transfer-channel-set -S Active CHANNELNAME`

- **Use file as plugin**
 - Mostly ‘static’ information
 - /opt/glite/etc/services.xml
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerServicesXml20>
- **“No site for host...”**
 - Check that endpoint is published correctly
 - Check the “volist” parameters
 - Update services.xml file
- **“No channel found, channel closed for your VO...”**
 - Check source and destination site
 - Verify that a channel is defined
 - Glite-transfer-channel-list command
 - Verify that your vo has a share defined on the channel
 - Verify that the share is not null
- **“No SRM method factory found...”**
 - Check the endpoint version in the services.xml file

```

<service name="httpg://lxdpm101.cern.ch:8446/srm/managerv2"
  <parameters>
    <endpoint>httpg://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...)

- **Hardware failures**
 - Recovery procedures in case of hardware failures
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsProcedures20>
 - Use load balancing
- **Software failures**
 - Well tested, but still bugs / missing features
 - Please report any issue you see
 - Ask your ROC
 - Check / submit tickets to GGUS
 - Contact fts-support@cern.ch
 - Workarounds usually published on fts-users@cern.ch
 - Patches come from the normal gLite release mechanism

- **Service degradations for which the cause is harder to determine**
 - “Transfer failure” (gridFTP isn’t always obvious about the cause).
 - <https://twiki.cern.ch/twiki/bin/view/LCG/TransferOperationsGridFTPErrors>
 - Networking problems / firewall issues
 - Problems on SRM-copy channels (FTS doesn’t have much logging about what went wrong)
- **These require ‘expert’ involvement and investigation**
 - Experts at tier-1: on local FTS, Castor, dCache, networking
 - Experts on tier-2: SRM, networking
- **Goal: move as much of this as possible to the ‘automatic’ system**
 - Packaging ‘canned problems’ takes time and experience with the problem

- **Agents and transfer log files**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerAdmin20>
 - Agents logs (channel and VO, according to the phase)
 - /var/log/glite
 - *glite-transfer-channel-agent-urlcopy-<channel-name>.log*
 - *glite-transfer-channel-agent-srmcopy-<channel-name>.log*
 - *glite-transfer-vo-agent-<vo-name>.log*
 - Search for the job id
 - Transfer logs
 - On the related channel agent machine
 - /var/tmp/glite-url-copy-edguser
 - *<channel-name>completed*
 - *<channel-name>failed*
 - *<channel-name>lost*

- **Focus continues upon service monitoring and making it easier to run the transfer service**
- **Continue testing SRM 2.2 support**
- **SL4 and VDT 1.6 support**
- **Closer integration with experiment software frameworks**
- **Incrementally improve service monitoring and admin**
 - Better service admin tools being released
 - Monitoring work now prototyped and running on pilot
- **Site grouping in channel definition (“clouds”)**
- **SRM/gridFTP split**
- **Notification of job state changes**