



Enabling Grids for E-scienceE

Transfer Service Status and Directions

Gavin McCance

*SC4 WLCG Service Workshop
10 February 2006*

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




- **Development focus**
- **Operational requirements**
- **Status of SRM support**
- **Status of new features**
- **Roll-out plan**
- **Summary**

- **FTS is a fabric service**
 - i.e. it's a management tool for the sites to allow them to easily control, prioritise and service the transfer requests of their VO's
- **Functionality-wise it's not very exciting**
 - It does **reliable** point-to-point movement of files
 - Plus some interaction with the experiment framework if you want to write it (for FTS 1.5)
 - The management and monitoring part is the “exciting part”
- **We deploy it to manage the major point-to-point transfers that feature in the experiment computing models**

- **Channel is a management concept**
 - It doesn't have to correspond to the exact network topology
 - Although the two are often the same (e.g. T0 to T1 transfers)
- **They are trivial to set-up**
 - Make sure the SRM endpoints are known
 - Run the command to define the channel in your DB
 - Run the command on one of your machines to create the agent that will service the channel
 - Start the agent
- **SRM endpoints are currently held in a static file**
 - Moving this to use standard BDII SRM information now
- **Management is the issue**
 - Having one channel to every tier-2 would be hard
 - Pending improvement to group sites together for better channel management:
 - “switch off all my tier-2's just now”
 - “increase the rate of traffic from the UK just now”

What's been achieved so far



- **Have shown stable running for a short while**
 - But with workarounds and significant support overhead
- **The goal of stable running for long periods is still a challenge!**

- **The focus must be upon the stability of the overall transfer service**
 - Operational issues
 - Monitoring and alarms
 - Procedures for recovery
- **Work towards limited functional additions**
 - SRM v2 support
 - Deployment of some things that are already included in FTS 1.5



Enabling Grids for E-scienceE

Operational Requirements

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



- **Service team and FTS developers has identified 44 issues that need addressing**
- **Of which 21 are critical to stable operations of the transfer service**

- **Many are “updating existing things”**
 - e.g. troubleshooting guide, procedures
- **Some are “hardening” of existing stuff**
 - e.g. properly integrating and fixing ad-hoc monitoring tools
- **Some are fair pieces of work**
 - e.g. support SRM v2, better monitoring tools

- Listed here: <https://edms.cern.ch/document/700073/1>
- **Gathered from a variety of sources**
- **LCG service tick-list:**
 - OSG/EGEE operations workshop:
<http://agenda.cern.ch/fullAgenda.php?ida=a054670>
- **WLCG “dashboard”**
 - <https://uimon.cern.ch/twiki/bin/view/LCG/WlcgScDash>
- **Requirements from GGUS**
- **CC integration requirements**
 - CC hooks still needed (coming from IT-FIO)
- **Debugging requirements**
 - “I submitted job X and it didn’t work? Why?”
- **Operational experience running FTS service for last few months**
 - We know what’s still painful

- **1. FTS agent**
 - Various service / configuration issues to fix
 - Use BDII information system to find SRMs
 - DB performance issues
 - MyProxy security issue
 - High availability plan from WLCG
- **2. FTS url-copy**
 - Support for SRM v2
 - Fix SRM copy for v1 (and write it for v2)
- **3. FTS portal**
 - Debugging / tracing information
 - Channel / management system overview
 - High availability plan from WLCG

- **4. FTS service monitoring**
 - Upgrade daily plots
 - Current status “live” plots
 - Service / site-component functional tests
- **5. CC integration**
 - Proper CC sensors and alarms to monitor the basic parts of the service
 - Improve 1st level procedure for operators
- **6. Procedures and documentation**
 - Improve and update the various guides
 - Interoperability guides
 - Understand impact of service outages
 - Improve the troubleshooting guide
 - Update the “check-my-fts” script
 - Improve operating procedures



Enabling Grids for E-scienceE

SRM support status

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



- **We believe we understand SRM put / get / 3rd party copy reasonably well**
 - This has taken a significant time and effort
 - Issue: differences between SRM implementations, particularly on errors
 - Issue: non-atomic SRM operations
 - put fails leaving a file entry, advisoryDelete fails, retry fails with “file exists” (workaround exists for FTS 1.5)
- **We do not understand SRM copy**
 - dCache and Castor do not agree on the request state-machine
 - Very easy to overload the system
 - No experience of running SRM copy hard: bugs will persist until we do
 - Some sites have volunteered!
 - Issues running the service, since we don't really get back enough information from SRM about problems
 - Next slide...
 - **This will take time and effort**

- This is tricky when there are problems:

- e.g. SRM copy fails to get file from Castor:

```
2006-01-24 16:18:27,721 [ERROR] - FINAL:TRANSFER: Failed on SRM copy: Failed To Perform srm copy. Error in srm__copy: SOAP-ENV:Client - CGSI-gSOAP: Error reading token data: Success
```

- **Retry fails saying file already exists:**

```
2006-01-24 16:20:29,811 [ERROR] - FINAL:TRANSFER: Failed on SRM copy: Failed SRM copy context id=-2121790037 Error is
RequestFileStatus#-2121790036 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790035 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790034 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790033 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790032 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790031 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790030 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790029 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790028 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
RequestFileStatus#-2121790027 failed with error:[ GetStorageInfoFailed : file exists, cannot write ]
```

- **We have no idea of the sequence of operations that the SRM is doing and which ones of them failed**

- We only see the final error message
- A failure can leave the files' states bad in the SRMs, so retry does not help

- **This is a problem for resolving service issues**

- **We lose much of our monitoring**
 - e.g. we can't monitor the current status of the gridFTP transfer since the SRM is doing it

- **Less information: failure classes**
 - For 3rd party copy, we distinguish 16 different error classes (without parsing error messages) – since we know when in the process the failure occurred.
 - This is helpful for debugging service issues.
 - For SRM copy, we get a “Failed” back
 - Error messages could be parsed
 - *But are generally not nice (often java stack traces)*
 - *Are not common between SRM vendors*
 - *Often change between versions*

- **Why do transfers work better with SRM copy?**
- **Since FTS has limited information to debug service problems when using SRM copy, the sites will see an increased support load (*even if the problem is not at their end of the transfer*)**
- **Given all sites achieved target rates using SRM 3rd party copy (and 8 out of 11 sites achieved their nominal rates), is it worth the effort?**

- **We believe we understand SRM put / get / 3rd party copy reasonably well for version 1**
 - This has taken a significant time and effort
 - Issues still exist in differences in SRM version, particularly on errors
 - Issues still exist with non-atomic SRM operations – put fails leaving a file entry, advisoryDelete fails, retry fails with “file exists” (workaround exists for FTS 1.5)
- **SRM version 2 is hopefully better specified**
 - Nevertheless, it would be prudent to expect problems, particularly interoperability problems
 - Especially in the case where there are service problems
- **This will take time to code: just starting now**
 - ...and a good time to test to understand the new behaviour of the SRM v2 implementations
- **It will not be for the throughput phase.**
- **We will try for the service phase, but integration timelines are tight.**

- **Would really rather not**
 - Neither SRM v1 or v2 has a getVersion method
 - Can look up version in information system?
- **Tooling problem with gSOAP**
 - Talking to two SRM versions from same executable is hard because of error handling in gsoap
 - Can probably work our way round the problems
- **Not clear how SRM copy works when there are two different version of SRM at either end?**
 - I would expect it not to work in an interoperable way the first time round (it doesn't know the version)
 - What is the plan for managing this?



Enabling Grids for E-scienceE

New features

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



- **Manpower, staff turnover:**
 - Team of people who support and develop the code and procedures reduced from 4 people to 2.
 - This is being addressed: 1 new developer just starting
 - Lost our lead developer on monitoring work
 - This will affect the monitoring work
 - Lost our lead developer on SRM / gridFTP layer
 - This will affect the SRM v2 work
 - Our new developer will focus in this area
- **Our main effort is in service support and stability rather than coding new features**
 - We will focus on operational improvements and deploying the features already in FTS 1.5

1. **Planning and routing**
2. **Monitoring**
3. **File integrity checking**
4. **Retry logic and Hold states**
5. **FTA Cataloguing**

-
6. **Staging support**
 - Should do the planning early


- **Items 1 to 5 seen by experiments as required for SC4**

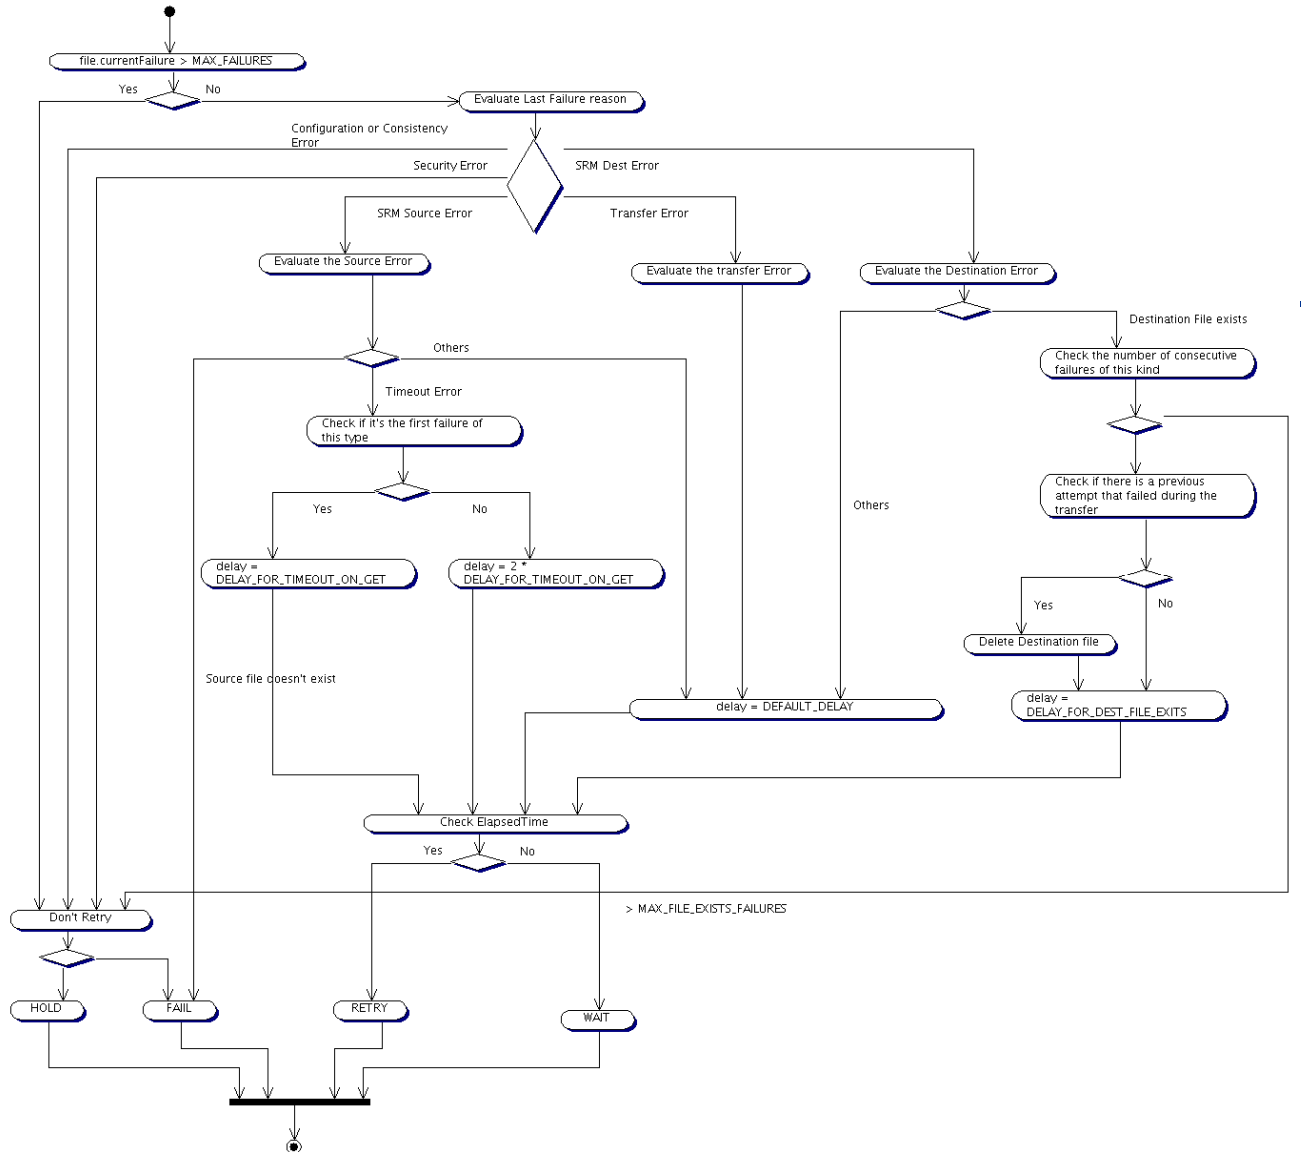
1. Planning and routing
2. Monitoring
3. **File integrity checking**
4. Retry logic and Hold states
5. FTA Cataloguing

-
6. Staging support
 - Should do the planning early

- Items 1 to 5 seen by experiments as required for SC4

- **This has been investigated**
- **The FTS server is remote from both the source and the destination**
 - It cannot checksum the files on disk
- **If and when the SRM vendors support checksumming, we can easily call the method and compare the values**
 - This is unlikely to happen before SC4
- **Poor man's checksumming 😊**
 - We already compare the file size

1. Planning and routing
 2. Monitoring
 3. File integrity checking  (requires SRM changes)
 4. **Retry logic and Hold states**
 5. **FTA Cataloguing**
-
6. Staging support
 - Should do the planning early
- Items 1 to 5 seen by experiments as required for SC4






- **Released in gLite 1.5**
 - Pre-transfer hooks, post transfer hooks (i.e. experiment code) running inside the VO agent part of FTS
 - Original plan was for “Cataloguing” operations
 - In fact, can be used for other pre/post operations
 - Also specific retry policies if the default isn’t good enough

- **Deployment issue is not solved**
 - How do we get the VO code to run on the VO agent node?
 - we don’t know how yet
 - See proposal on next slide...

- **ATLAS evaluating these hooks now**
 - Hooks to contact software framework before and after transfer
 - Good news: This should allow ATLAS to remove much of the code currently running on the VO box

- **Idea here is to absorb some of the functionality currently in expt. software frameworks into the FTS**
 - Bad news: it looks kinda like the VO box from security and deployment view
- **Where do we run the VO agent?**
 - A plain default VO agent (running default non-experiment-specific plug-ins) can run on a service node (maybe the same node as the rest of the FTS service)
 - This is how we run VO agents just now
 - If it has experiment plug-ins installed, the requirements are fairly similar to the VO box
 - Probably run the VO agent on the “VO box” for a while?
 - Once the experiment plug-ins have stabilised, take the stable software back onto a site-managed service node.

1. Planning and routing
 2. **Monitoring**
 3. File integrity checking  (requires SRM changes)
 4. Retry logic and Hold states  (still to work out deployment)
 5. FTA Cataloguing (plug-ins) 
-
6. Staging support
 - Should do the planning early
- Items 1 to 5 seen by experiments as required for SC4

- **Covered by the large list of operational requirements**
- **We have some monitoring now**
 - Still not good enough – we still use a lot of effort running the service
- **Monitoring is one of the top priorities identified by the service team as well**
 - Improvements underway

1. **Planning and routing**

2. Monitoring (improvements underway)

3. File integrity checking (requires SRM changes)

4. Retry logic and Hold states
(still to work out deployment)

5. FTA Cataloguing (plug-ins)

6. **Staging support**

- Should do the planning early

- Items 1 to 5 seen by experiments as required for SC4

- Routing (multi-hop) transfers
- “Central point” to submit everything
- We have worked in December/January to develop a prototype to determine the issues and the effort involved
- **Unfortunately it’s too costly to proceed with it**
 - It will undermine our effort to maintain a stable service
- **The intermediate “cache” storage would need to be managed at the T1**
 - Deployment impact is not understood
- **The “central point” would create a significant operational problem**
 - Hard to make this scale well

- The problem has been reviewed by various people since
- We believe we can cover the primary production use-cases in the computing models using FTS point-to-point (i.e. no multi-hop)
 - T0 to T1
 - T1 to T1
 - T2 to associated T1, T1 to associated T2
 - With reasonable changes (allowing “site groups” in channel definitions)
 - T2 to non-associated T1 (point-to-point)
 - T2 to T2 (point-to-point)
 - Point-to-point traffic will be routed over whatever link is available
 - For production routes, this will be the OPN

- **Providing a central service for all data transfer submissions**
 - Scaling problems
 - We don't have the effort for it
- **But..**
 - we can provide a fat client or other tool that makes the decision for you of where to submit based the job content
 - all the information to make the decision is in BDII
- **Maybe this is less of a problem given we can do all transfers in a single FTS hop**

1. **Planning and routing**  (possible solution – to discuss)

2. **Monitoring**  (improvements underway)

3. **File integrity checking**  (requires SRM changes)

4. **Retry logic and Hold states**  (still to work out deployment)

5. **FTA Cataloguing (plug-ins)** 

6. **Staging support**

- Should do the planning early

- **Items 1 to 5 seen by experiments as required for SC4**



Enabling Grids for E-scienceE

Roll-out plan

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



- **We've been running 1.4 for some time now at CERN (several months)**
 - Very few (if any) T1 sites have been running it
 - Because we didn't tell you – we wanted to iron out operational problem (and bugs) with it first, otherwise the support load would have killed us
 - We found lots of problems - but now it's stable
- **I think this is a fairly good model...**
 - Although “several months” is too long...
- **gLite 1.5 FTS is already out with new features**
 - It will have operational problems
 - We'd like to find the bad ones before we give 1.5 to you

- **Sites update to 1.4 as soon as possible**
 - Upgrade and install guides on the Wiki
 - 1.4 clients are being distributed in LCG 2.7.0
- **CERN begins update to gLite 1.5 now**
 - Discover the pit-falls
 - Make the workarounds and the fixes
 - Update the operating procedures
- **Deploy 1.5 before SC4 at sites**
 - Perhaps one site would like to be an early customer
 - Warning! Support load for the site...
- **Clients will be backwards / forwards compatible and are distributed with LCG 2.7.0 (and then gLite 3.0)**



Enabling Grids for E-scienceE

Summary

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



- **Manpower issues**
 - Some turnover (and reduction) in staff
 - This is slowing things down
- **SRM v1 3rd party copy mostly understood**
 - Can obtain good rates / fair stability with this
- **SRM copy not really understood**
 - Harder to debug problems
 - Consequent higher support load on T1 sites
 - Is it needed?
- **Still to develop SRM v2 support and test properly at high transfer rates**
 - This will take time and effort: just starting now
 - Work is common with SRM v2 support for LCG utils / GFAL
 - Goal is for the SC4 service phase
- **Mixed SRM version support is not understood**
 - What is the plan for this?
 - Do we need this?

- **Development goals**
 - Stability, stability, stability...
 - Need to work on operational issues or overall service will not be stable
- **Functional improvements**
 - SRM v2 support
 - Plug-in support from gLite 1.5 – still need to understand / agree deployment model for this
 - “Site grouping” to allow easier management of “point-to-point” channels
 - Point-to-point production channels will go over the OPN
 - Less well used channels (e.g. random T2 to random T2) will likely go over general network
- **Roll out plan**
 - Suggest sites upgrade to 1.4 – it’s understood now
 - We’ll work on understanding the new features of 1.5








Enabling Grids for E-scienceE

Backup slides

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



1. **Planning and routing**  (possible solution – to discuss)
 2. **Monitoring**  (improvements underway)
 3. **File integrity checking**  (requires SRM changes)
 4. **Retry logic and Hold states**  (still to work out deployment)
 5. **FTA Cataloguing (plug-ins)** 
-
6. **Staging support**
 - Should do the planning early
- **Items 1 to 5 seen by experiments as required for SC4**

- **Why do transfers work better with SRM copy?**
- **Since FTS has limited information to debug service problems when using SRM copy, the sites will see an increased support load (*even if the problem is not at their end of the transfer*)**
- **Given all sites achieved target rates using SRM 3rd party copy (and 8 out of 11 sites achieved their nominal rates), is it worth the effort?**