


Current status of transfer service



FTS Workshop

16 November 2005

Gavin McCance
IT/GD, CERN

Outline



- Summary of FTS deployment in SC
- Status of current version
- Current work-plan
- Issues



Reminder: FTS?

- File Transfer Service is a **fabric** service
- It provides point to point movement of URLs
 - Aims to provide reliable file transfer between sites, and that's it!
 - Allows sites to control their resource usage
 - Does **not** do 'routing' (e.g like PheDEX)
 - Core FTS does **not** deal with GUID, LFN, Dataset, Collections
- It's a fairly simple service that provides sites with a reliable and manageable way of serving file movement requests from their VOs
- We are understanding together with the experiments the places in the software where extra functionality can be plugged in



What is the FTS?

- Hope you all know this by now ;)
- Some points to stress:
- Decouple the "FTS" software component from the "fts" service.
 - The "FTS" software component of gLite is something which does a simple well-defined thing: point-to-point movement of SURLs
 - The "fts" service (or "transfer service"?) is more
 - Service: e.g. monitoring and debugging of SRMs
 - Related functionality: how to find FTSES, etc
 - Warning: I have a tendency to talk about the two together and mix them, often in the same sentence
- We should try to focus on the overall service



Why would I want to use it?

- ...instead of using say...
- ... plain globus-url-copy or lcg-cr?
- FTS is a **fabric** service
 - The exposed functionality to the experiments is rather simple
 - Much of the work is to expose management / monitoring functionality and resource control to the sites so that they can actually operate and debug the overall service (with limited staff)
 - Includes (active) service monitoring
 - We should know before you do!
 - Is integrated into the service debugging
 - We can make tools to mine the logs or DB to debug problems at sites
 - Improvements here allow us to provide a better overall service



SC3 summary

- Throughput phase was completed using version 1.1.2 of FTS
 - A rewrite of the Radiant prototype
- 1.3 version came in for the service phase, adding
 - Multi-VO support: fair-share
 - Introducing basic VO-agent framework
 - VO production manager role
 - Better logging / monitoring
- FTS software has been ~stable operation for some months now
 - Many issues / bugs resolved
 - A lot of new ~small features added
 - Some from experiments / many on the service side
 - Try to maintain high availability of the service
 - Interventions scheduled but frequency has been a little too high!



Service deployment

- FTS deployed at CERN initially to manage the T0 to T1 transfers
 - Extended for T1 to T0 channels, still managed by CERN FTS
 - T1 site managers can control the service even though it is hosted at CERN
- Also deployed at tier-1 sites to manage T1 to T2 and T2 to T1 transfers
 - Analogous model: T1 hosts the FTS, but T2 site managers also have control
 - Reasonable experiences here - though with many versions the support load is rather high
- Used against Castor, Castor2, dCache and DPM SRMs
 - All currently using SRM v1 interface



Service deployment

- Now trying out a T1 to T1 'mesh' of transfers
 - Just a few sites
 - Interested to understand from the experiments how sparse (or not) the mesh is?
- Don't yet have a (good) deployment model for supporting T2 to T2 transfers
 - Software can 'do' it, but the challenge is to make it manageable given the resource limitations of T2 sites
 - What are the use-cases for this?



Status of 1.4.1 release

- New gLite release 1.4.1 FTS pre-deployed at CERN
 - Plan to roll out to T1 sites soon
 - Backwards-compatible clients (c.f. 1.1.2 -> 1.3!)
- Added new features
 - User can specify which MyProxy server to use
 - BDII publication of channels
 - SRM-copy support
 - MySQL support

 - Basic monitoring API
 - Cleaner configuration
 - Agent registration
 - Improvements in logging / monitoring



Current work-plan

- Items (non-exhaustive) on current work-plan
- One of the purposes of the workshop
 - Subject to discussion / change !
 - ...and in no particular order
- Staging support
 - Currently the FTS times out, so jobs fail
- SRM v2 support
 - Initially will be basic support for the early versions



Current work-plan

- **Topology discovery**
 - Where is/are the FTS endpoint(s) that can move from CERN to Glasgow?
 - Better information system integration
- **Channel halting / error recovery**
 - How many retries should we do?
 - When does it halt a channel rather than drain the queue?
- **Horde of 'service' issues**
 - Logging / monitoring / procedures / problem tracing / integration with CERN CC infrastructure
- **The work-plan is tracked here:**
 - <https://uimon.cern.ch/twiki/bin/view/EGEE/DMFtsWorkPlan>

Issues

Don't want to steal your thunder...



Issues: finding FTSeS

- Issue:
 - The FTS servers do not communicate, each manages its own (set of) channels
 - If I know I want to go from XXX to YYY, where is the server endpoint I need to talk to
- Two issues:
 - Routing: I need to go from XXX to YYY (via ZZZZ)
 - Endpoint resolution: I know I need to go from XXX to YYY, where is the endpoint that deals with the XXX-YYY channel?



Issues: staging

- Issue:
 - No explicit stage support - staging not managed as a separate state
 - SRM.get makes implicit stage but FTS times-out
 - Increasing timeout is no good since we 'block' a 'transfer slot'
 - Should have explicit "stage" step in state machine
 - Control staging separately from transfer
- It's VO specific (potentially)
 - Managed by VO agents
 - We provide default, you can override if you want to write the code e.g. if you software framework already manages its pool in a specific way



Issues: VO framework

- Description of VO plug-in framework
 - See Paolo's talk
- Retry policies and channel halting
 - How much of the retry policy can be absorbed into the transfer service
- FPS Cataloguing functionality
 - Looking to provide a more integrated overall 'service'
 - VO specific catalog functionality
- Staging
 - Potentially VO specific



Issues: various

- SRM v2 support
 - Timelines, etc
-what else ?

Summary

- ... at the end ...

