

RAL FTS Feedback

Dr Matthew P. Hodges
(M.P.Hodges@rl.ac.uk)

October 18, 2006

- 1 Deployment at RAL
 - FTS Box, Channel and VO Agents
- 2 Operational Issues
 - Information System
 - Database Cleanup and Mixed Case Issues
 - Usage Patterns Changing
 - Srmcopy Channels
- 3 Monitoring
 - Nagios and Log Files
 - Ganglia
 - Command Line
- 4 Wishlist
 - Installation
 - Operation
 - Monitoring

Thanks to `fts-support@cern.ch`, and in particular Paolo Badino and Gavin McCance.

FTS Box(en)

Everything currently on one box (hitting limit; more below).

Channel Agents

About 70 channel agents in total. Full SC4 model:

- Channels to RAL. {T1, UKIT2, STAR} → RAL.
- Channels to UKI T2s. {RAL, STAR} → UKIT2.

RAL supports about 20 UKI T2s (also Estonia for CMS).

VO Agents

RAL supports about 20 VOs. Only LHC + dteam/ops active.

About 90 agents in total.

Information System

- Not publishing: added `config_rgma_client` to `FTS_FUNCTIONS` in `/opt/glite/yaim/scripts/node-info.def` and reran `yaim`.
- Timeouts: increased timeout and reduced frequency of polling in `/opt/bdii/bdii.conf` and increased freshness, `cache_ttl` and timeout parameters in `/opt/lcg/etc/lcg-info-generic.conf`. (Otherwise had constant load from info provider scripts.)
- Passing SFT/SAM tests (rgma-gin problem: restart when too many open files).

Database Cleanup

Implementing the periodic cleanup helped reduce load on FTS and Oracle backend. Currently have about 6000 jobs in `t_job` table and 120,000 jobs in `t_file` table.

Mixed Case Issues

Needed DB-level intervention to uppcase src/dest channel parameters. Fixed now (everything uppcase internally)?

Usage Patterns Changing

More activity from ATLAS (DQ2), and CMS (PhEDEx) for CSA06. General high load on server seen at times.

- Tomcat: high load killed tomcat server. Increased `maxThreads` from 150 to 300; correct remedial action (maybe the restart was sufficient)? High rate of starts/fails due to broken SE around this time.
- Redeployment: split web service from channel/VO agents to start with. Also split agents?

Srmcopy Channels

Lots of requests for performance reasons (must have dCache at destination). Problems with:

- Stale transfer files and URLCOPY_GUC_MAXTRANSFERS exceeded. Fixed.
- Infinite looping (DB deadlock?). Fixed, but not yet released?

Currently we are operating only urlcopy channels.

Nagios

Basic check that channel/VO agents are running (using `/etc/init.d/transfer-agents status`). Doesn't address problem with channels that have a running process that is not working correctly (e.g., as seen with srmcopy channels). Other checks (tomcat, rgma).

Log Files


We copy the log files to `/var/lib/tomcat5/webapps/glite-data-transfer-fts/` so that users trying to debug transfer problems can access them via web service (with certificate).

Ganglia


We have plots of overall activity (start/active/done/fail/transfer rates), and per-site, per-channel, and per-VO breakdowns of activity. Helps identify some problems. No rates for srmcopy channels (not relevant at present), and will only work on single FTS deployment (at least the VO agent and channel agent logs must be available; also transfer logs for file sizes). Rewrite to query DB directly?

```
http://ganglia.gridpp.rl.ac.uk/cgi-bin/ganglia-fts/  
fts-page.pl?r=day
```


All Active transfers:




All Active transfers to/from RAL-LCG2:



All Active CMS transfers to/from RAL-LCG2:



All transfer rates for CMS transfers to/from RAL-LCG2:



Command Line

- Most users and transfer tests rely on checking status with `glite-transfer-status`. Smaller load generated from queries after deploying DB cleanup mechanism.
- Hard to get a quick overview with `glite-transfer-list` Active and `glite-transfer-status` combination.
(Monitoring at SARA, <http://winnetou.matrix.sara.nl/monitoring/datatransfer/>, looks promising.)
- List transfer IDs by DN and/or submission date?


Installation

- Yaim could handle the `glite-transfer-channel-add` step and add default VO shares?


Operation

- Site managers: map DNs to sites, and allow user to manage any channel involving site? (Currently use service admins.)
- Blacklist SEs: can't isolate problem by making channels Inactive if problem SE matches STAR; if VO-specific problem, then set VO share to zero, but what if more than one SE per site (e.g., for ATLAS at RAL, dcache SE is full but dcache-tape SE is not)?
- Allow switching between urlcopy and srmcopy channels by channel manager. Either urlcopy or srmcopy channel agent is active? Handover could be tricky.
- Controlling maximum number of Active transfers on the service to limit load? (VO fair shares of overall service?)

High load due to failing SE:



Ganglia plots:



Monitoring

- Relying on home-grown tools for overall status.
- FTS Dashboard: SARA monitoring good start; channel management; graphical overview.