

# Table of Contents

<b>WLCG Transfers Dashboard.....</b>	<b>1</b>
Links.....	1
Milestones.....	1
Delivered.....	1
In progress.....	5
Future.....	5
Under evaluation.....	6
VO requests.....	6
ATLAS requests.....	6
Minutes from meeting 13/02/2012 by Jaroslava Schovancova.....	6
Email 30/05/2012 from Alessandro Di Girolamo.....	7
CMS requests.....	8
Minutes from meeting 21/11/2011 by Andrea Sciaba.....	8
Minutes from meeting 20/03/2012 by David Tuckett.....	8

# WLCG Transfers Dashboard

This page documents the development of the WLCG Transfers Dashboard, which is a component of WLCGTransferMonitoring.

## Links

- Production User Interface: <http://dashb-wlcg-transfers.cern.ch/ui/>
- Integration User Interface: <http://dashb-wlcg-transfers-dev.cern.ch/ui/>
- Administrator Interface (<http://dashb-wlcg-transfers.cern.ch/ai/>):
  - ◆ Consistency Checks: <http://dashb-wlcg-transfers-dev.cern.ch/ai/consistency.html>
  - ◆ FTS Collector Monitoring: <http://dashb-wlcg-transfers-dev.cern.ch/ai/ftsmonitoring.html>
  - ◆ XRootD Collector Monitoring:  
<http://dashb-wlcg-transfers-dev.cern.ch/ai/xrootdmonitoring.html>
  - ◆ Stress Test: <http://dashb-wlcg-transfers-dev.cern.ch/ai/stresstest.html>
- Support Guide
- Savannah

## Milestones

### Delivered

These milestones have already been delivered.

1. Infrastructure I: DELIVERED 18/07/2011
  - ◆ Organise virtual machine for server and integration account for database.
  - ◆ Deliverable: Server and DB access for all team members.
  - ◆ Notes:
    - ◇ DB account: lcg\_dashboard\_tfr@int6r
    - ◇ Server: dashboard46
2. Message chain I: DELIVERED 15/08/2011
  - ◆ Send messages from mock producer and consume into message tables in database.
  - ◆ Deliverable: Code in SVN including demo that can be launched from command line.
3. Message chain II: DELIVERED 15/08/2011
  - ◆ Consume multiple messages and perform bulk inserts into database, only acknowledging receipt if insert successful.
  - ◆ Deliverable: Code in SVN including demo that can be launched from command line.
  - ◆ Notes:
    - ◇ STOMP 1.0 does not allow negative acknowledgement. Closing the connection without acknowledging is the workaround. However, we propose to accept all messages and publish those that cannot be inserted into the database to an error queue on the message broker.
4. Message chain III: DELIVERED 15/08/2011
  - ◆ Configurable consumer to handle changes in message format and/or new message types.
  - ◆ Deliverable: Code in SVN including demo that can be launched from command line.
  - ◆ Notes:
    - ◇ Configuration is currently restricted to mapping a queue to a table not message fields to columns. We can make this more configurable as and when necessary.
5. Message chain IV: DELIVERED 07/09/2011

- ◆ Publish messages which fail to insert to error queue (see notes for Message chain II).
- ◆ Deliverable:
  - ◇ Code in SVN including demo that can be launched from command line.
  - ◇ Brief features and configuration documentation in SVN covering work completed so far.
- ◆ Notes:
  - ◇ A standard component provided by IT-GT "consume2db" provides this functionality. We should either use it or justify why it is not used.
- 6. Integration with FTS test instance: DELIVERED 07/09/2011
  - ◆ Consume messages from FTS test instance. Update mock producer for any changes in format.
  - ◆ Deliverable: Code in SVN including demo that can be launched from command line.
- 7. Stress test I: DELIVERED 07/09/2011
  - ◆ Test maximum consuming rate using mock producer.
  - ◆ Deliverable:
    - ◇ Code in SVN including demo that can be launched from command line.
    - ◇ Plots / tables of results in SVN.
- 8. Logging: DELIVERED 07/09/2011
  - ◆ Log message insert statistics and failures to file.
  - ◆ Deliverable:
    - ◇ Code in SVN including demo that can be launched from command line.
    - ◇ Update features and configuration documentation in SVN.
- 9. Multiple broker support: DELIVERED 07/09/2011
  - ◆ Resolve DNS alias to multiple brokers and consume from each broker.
  - ◆ Deliverable:
    - ◇ Code in SVN including demo that can be launched from command line.
    - ◇ Update features and configuration documentation in SVN.
- 10. Infrastructure II: DELIVERED 29/09/2011
  - ◆ Organise dedicated test message broker server from PES.
  - ◆ Deliverable: Server access for all team members.
- 11. Integrate UI I: DELIVERED 07/09/2011
  - ◆ Proof of concept by copying and adapting ATLAS DDM Dashboard UI.
  - ◆ Deliverable:
    - ◇ Deploy UI to integration server.
  - ◆ Notes:
    - ◇ See <http://dashboard46/ui/> (accessible inside CERN).
- 12. Rewrite mock producer and stress test as Dashboard agents. DELIVERED 15/11/2011
  - ◆ Deliverable: Demonstrate that mock producer and stress test can be launched on dashboard46 using dashb-agent-\* commands.
- 13. Present work completed to this point at the IT-ES-DNG section meeting. DELIVERED 15/11/2011
- 14. Security: DELIVERED 16/12/2011
  - ◆ Use host certificates to authenticate with message broker.
  - ◆ Deliverable:
    - ◇ Code in SVN and deployed to dashboard46.
    - ◇ Update features and configuration documentation in SVN.
  - ◆ Notes:
    - ◇ Depends on Infrastructure II milestone.
- 15. Consume messages from FTS pilot deployment: DELIVERED 16/11/2011
  - ◆ Deliverable: Transfer statistics from FTS pilot should be visible in dashboard46 UI.
  - ◆ Notes:
    - ◇ Depends on FTS pilot deployment.
- 16. Switch to using dedicated message brokers. DELIVERED 16/12/2011
  - ◆ Coordinate with Michail to use production broker alias dashb-mb.
  - ◆ Deliverable: New transfer statistics from FTS pilot should be visible in dashboard46 UI after switch.

17. Switch to using virtual queues. DELIVERED 16/12/2011
  - ◆ The use of virtual queues should mean that messages are not lost during collector downtime.
  - ◆ Coordinate with Lionel to set up virtual queues for WLCG Transfer Dashboard.
  - ◆ Implement with Lionel the best option for managing queues (e.g. queue growth due to collector downtime).
  - ◆ Deliverable:
    - ◇ Demonstrate using msg admin interface that message are not lost during collector downtime.
18. Set up dashboard46 as showcase integration service. DELIVERED 23/01/2012
  - ◆ Dashboard46 should be set up as a stable integration service to showcase the project.
  - ◆ Alarms should be put in place to monitor the service.
  - ◆ Only genuine transfers from the FTS pilot should be recorded.
  - ◆ Transfer history should be retained as follows
    - ◇ Event details: 3 months.
    - ◇ Statistics: indefinitely.
  - ◆ Deliverable: Once the service is stable, the URL will be shared with interested parties.
19. Add country as top-level filter/grouping. DELIVERED 27/01/2012
  - ◆ Currently the filtering/grouping of sources/destinations includes: sites, hosts, tokens. A higher level should also be included giving: countries, sites, hosts, tokens.
  - ◆ The default grouping will be country so as to avoid the default dimension of the matrix and plots being too large and hence unmanageable.
  - ◆ The source for classifying sites by country: <http://wlcg-rebus.cern.ch/apps/topology/>
  - ◆ Thanks to the design of the Transfers Dashboard the addition of country information can be easily performed in the web actions using a cached copy of site to country mapping.
  - ◆ Deliverable:
    - ◇ Demonstration using UI of filtering and grouping by country, in addition to existing filters and groups.
20. Make the list of VOs used in filtering dynamic. DELIVERED 27/01/2012
  - ◆ Currently the list of VOs is static. It should be made dynamic.
  - ◆ In the first instance this can be a distinct query against the archival statistics tables.
  - ◆ It should be cached at least in the browser.
  - ◆ If a performance issue is observed on the server then it should be cached also on the server.
  - ◆ Deliverable:
    - ◇ Demonstration using UI of a dynamic list of VOs.
21. VO-specific site naming in UI. DELIVERED 30/01/2012
  - ◆ When a single VO is selected in the UI the site names (both input and output) should be translated to VO-specific naming conventions
  - ◆ VO topologies can be retrieved in XML format from external sources.
  - ◆ Thanks to the design of the Transfers Dashboard translation can be easily performed in the web actions using a cached copy of the VO-specific topologies.
  - ◆ Deliverable:
    - ◇ Demonstration using UI that CMS and ATLAS naming conventions are respected both in inputs (source/destination filters) and outputs (matrix, plots).
22. Add server (aka fts endpoint) as filter. DELIVERED 03/02/2012
  - ◆ In the same way that filtering by VO is currently supported, filtering by server (aka fts endpoint) should be supported.
  - ◆ The server field should be added to the statistics tables and statistics generation procedures.
  - ◆ The server field should be added in ui, action, dao (DAVID)
  - ◆ Deliverable:
    - ◇ Demonstration using UI of filtering by server (aka fts endpoint).
23. Curl Service. DELIVERED 06/02/2012
  - ◆ Generic Curl Service that can download from a given URL and save the file locally.
  - ◆ This will be used to source files for VO-specific site-naming and country filtering/grouping.
  - ◆ The service parameters should be: interval, url, filename, accept-type.

- ◆ Deliverable:
  - ◇ Deploy service to dashboard46 and demonstrate that it updates the vo feed xml and the topology json.
- 24. Add support twiki. DELIVERED 21/02/2012
  - ◆ A support twiki page should be created to instruct dashboard team members how to manage the Transfers Dashboard in case of incident.
  - ◆ Deliverable:
    - ◇ New twiki page linked from main Dashboard twiki:  
SupportWLCGTransferDashboard
- 25. Integrate UI II (DAVID) DELIVERED 26/03/2012
  - ◆ Factor out common parts of Transfers Dashboard and ATLAS DDM Dashboard UI.
  - ◆ Deliverable:
    - ◇ New common dashboard transfers UI module: xbrowse.
    - ◇ Extension of above for ATLAS DDM Dashboard UI deployed.
- 26. Add alarms in case transfer messages are not reported from production FTS servers. DELIVERED 16/04/2012
  - ◆ This means having a list of production FTS servers in the configuration and checking that we have messages from those in the last N minutes. If not an alarm should be sent to dashb-mb-alarms.
  - ◆ This should be implemented so as not to unduly affect the performance of the collector.
  - ◆ Deliverable:
    - ◇ Demonstration that alarms are sent for an (unknown) FTS server that is not currently active, and alarms are not sent for an FTS server that is active. This can be validated against the UI.
- 27. Add reference data tables for dynamic VO and server lists. DELIVERED 16/04/2012
  - ◆ t\_vo and t\_server tables should be added. db jobs should be created to update them. dao should be updated to use new tables for reference data.
  - ◆ Deliverable:
    - ◇ Demonstration that changes show VOs and servers dynamically in the UI.
- 28. Clean-up Oracle stored procedure code. DELIVERED 16/04/2012
  - ◆ Use consistent naming for all functions and procedures.
  - ◆ Split all statistics compute/aggregate procedures into two procedures: from/to, latest (which uses from/to)
  - ◆ Deliverable:
    - ◇ Code deployed with no regression in statistics generation.
- 29. Deploy to production servers. DELIVERED 18/04/2012
  - ◆ Deploy UI and collectors to 2 machines redundantly.
  - ◆ Update support twiki.
  - ◆ Deliverable:
    - ◇ Check UI works via both machine names and alias.
- 30. Integrate UI III (DAVID) DELIVERED 23/04/2012
  - ◆ Extend xbrowse module for Transfer Dashboard UI.
  - ◆ Deliverable:
    - ◇ Code deployed with no regression in UI.
- 31. Deploy to production database instance DELIVERED 03/05/2012
  - ◆ Request production database instance.
  - ◆ Copy schema and data from integration.
  - ◆ Switch production servers to use production database instance.
  - ◆ Deliverable:
    - ◇ Both production servers should use the production database instance with no loss of data during the switch.
- 32. Add label column to t\_vo and t\_server tables for use in UI DELIVERED 03/05/2012
  - ◆ The default label column for VO would be the same as vo.
  - ◆ The default label column for server would be the hostname part of the server URL.

- ◆ The procedure that inserts into the t\_vo and t\_server tables should set the appropriate default label.
  - ◆ VO labels should be updated manually as follows: atlas -> ATLAS, cms -> CMS, lhcb -> LHCb.
  - ◆ The dao should be updated to return the name and label.
  - ◆ The GetReferenceDataAction should be simplified to just pass on the label from the db.
  - ◆ Deliverable:
    - ◇ Code deployed default and custom labels visible in UI.
33. Port admin UI to xbrowse framework DELIVERED 03/05/2012
- ◆ Extend xbrowse module for Transfer Dashboard AI.
  - ◆ Deliverable:
    - ◇ Code deployed with no regression in AI.
34. Automatic consistency check DELIVERED 03/06/2012
- ◆ Online plots showing the absolute and percentage difference between Transfers Dashboard statistics and VO-specific (e.g. PhEDEx) monitoring statistics.
  - ◆ Transfer statistics can be retrieved from PhEDEx in XML or JSON format.
  - ◆ Deliverable:
    - ◇ Add page in admin interface with plots as described above e.g. /ai/consistency.html
    - ◇ Add list of comparable links for PhEDEx and Transfers Dashboard support twiki.
35. Use db scheduled jobs instead of dashb agents for statistics generation DELIVERED 03/06/2012
- ◆ Create db scheduled jobs to call compute/aggregate statistics procedures.
  - ◆ Create db procedures to start/stop the scheduled jobs.
  - ◆ Remove corresponding dashb agent code / configuration.
  - ◆ Deliverable:
    - ◇ Jobs deployed with no regression in statistics generation.
36. Add db scheduled job to manage partitions in t\_start and t\_complete tables DELIVERED 03/06/2012
- ◆ Create db scheduled job to drop old partitions and create future partitions.
  - ◆ Create db procedures to start/stop the scheduled job.
  - ◆ Deliverable:
    - ◇ Job deployed. Database partitions management verified via SQL developer.
37. Add static topology mapping configuration for sites not in REBUS DELIVERED 17/06/2012
- ◆ This should supplement the REBUS topology used to map sites to countries.
  - ◆ This addresses [BUG:92867](#)
  - ◆ Deliverable:
    - ◇ Code deployed. No more sites in country 'n/a' in UI.

## In progress

These milestones are currently in progress. The order given here does not necessarily indicate priority.

1. Continuous code review
  - ◆ Reviewed so-far:
    - ◇ Actions
    - ◇ CurlService.
    - ◇ Oracle stored procedures.
  - ◆ Deliverable:
    - ◇ List of agreed code fixes.

## Future

These milestones have been evaluated and will be implemented. The order given here does not necessarily indicate priority.

1. Error categorisation proposal.

- ◆ Currently errors are categorised simply by t\_transfer\_complete.tr\_error\_category. This means that errors are grouped in just ~12 categories.
  - ◆ This task is to propose a better solution for grouping errors into categories so that a sample error message from a given category is representative.
  - ◆ Possible solutions:
    - ◇ Use other fields from t\_transfer\_complete such as t\_error\_code, tr\_error\_scope, t\_failure\_phase.
    - ◇ Use heuristics such as error message length with [] contents removed.
    - ◇ Use regex pattern matching.
    - ◇ ...
  - ◆ Deliverable:
    - ◇ A document / presentation showing a number of different approaches and proposing the 'best' solution. This may include statistics such as how many categories 10000 errors are split into and examples of how representative samples are for a given approach.
2. FTS queue status monitoring. (DUBNA)
    - ◆ See WLCGTransfersDashboardQueueStatus.
  3. Latency monitoring. (Alexandre)
    - ◆ See WLCGTransfersDashboardLatency.
    - ◆ See #Email\_30\_05\_2012\_from\_Alessandro
  4. Features sourced from VO requests
    - ◆ See VO requests sections below.

## Under evaluation

These milestones are being evaluated and may be implemented in the future. The order given here does not necessarily indicate priority.

1. Stress test II: POSTPONED
  - ◆ Repeat stress test I with security enabled.
  - ◆ Deliverable:
    - ◇ Plots / tables of results in SVN.
  - ◆ Notes:
    - ◇ Depends on Stress Test I and Security milestone.

## VO requests

Commonality between VO requests should be found and milestones added above.

## ATLAS requests

### Minutes from meeting 13/02/2012 by Jaroslava Schovancova

[1] <http://dashb-wlcg-transfers.cern.ch/ui>

Stephane:

- Since the information is collected from the FTS server directly, it would be interesting to identify the channels which are saturated.
  - ◆ By saturated, I mean that FTS jobs have to wait long before being processed OR they are always waiting FTS jobs in the channel.
  - ◆ My ATLAS colleagues want to transfer more and more files/GB. But I cannot provide the list of channels which are already saturated or close to be saturated.

- naming convention of sites:
  - ◆ For the moment Dashboard/FTSmon is using GOCDB names, and I believe it is the only option for the view when multiple VOs are shown.

David (re naming convention):

- Currently if a single VO is selected the naming convention of that VO is used. However, we will review whether this should be optional because it is not entirely intuitive that filters for multi-VO view do not select the same sites in single-VO view.

Discussion (re FTS saturation monitoring):

- This is very similar to the request from CMS.
- We (ATLAS) would like to have a possibility to filter/sort by the FTS server and by the FTS channel (so it is already in the todo list).
- In addition FTSmon is planning to consume information about FTS queues. This would allow to correlate the transfer latency with the state of the queue.
  - ◆ ask regularly: are there jobs in a queue? same as Simone's "programmational knowledge of status of DDM channel" request
    - ◇ each 10min, or 1hr, let's see what is better
    - ◇ even if there is no action you ask
    - ◇ over 1 month you make statistics: fraction of time when no jobs waiting, fraction of time when jobs waiting in a queue

Simone:

- list of servers: find out a nickname (e.g. hostname of the server)
- programmational knowledge of status of DDM channel (SRC+DEST pair): how many files queued/in transfer per channel
- FTS channel limits (queue length, in transfer limit)
- exposed mapping between DDM channel (SRC+DEST pair) and FTSSchannel+FTSserverHostname (info from FTS)
- exactly same view as [1] but display only rates from the gridftp part of the transfer (transfer: SRM init at SRC, gridftp, SRM init at DEST) --> almost the network part, get rid of SRM overhead
  - ◆ list of SRM endpoints, SRC SRM overhead, DEST SRM overhead for each SRM endpoint
  - ◆ UI + programmational access, want both, start with programmational one

Ale:

- some of job output which FTS server has done? would like to download the log
- occupancy of channels: not matrix, but table

### Email 30/05/2012 from Alessandro Di Girolamo

Dear David,

ATLAS observed slow transfers from IN2P3-CC to TRIUMF-LCG2. The general speed of the transfers was not globally very slow, but few transfers were very slow (0.1MB/s). The problem was that there were 2 disk servers which behave not well. These few transfers were then blocking the full channel, thus creating huge backlog. [https://ggus.eu/ws/ticket\\_info.php?ticket=82363](https://ggus.eu/ws/ticket_info.php?ticket=82363)

With the old FTS monitor "a la IN2P3-CC" we were able to spot these transfers, checking the FTS jobs that were lasting for more than N hours. Do you think it will be possible to have this info also in the new FTS monitor?

Thanks Ale

## CMS requests

### Minutes from meeting 21/11/2011 by Andrea Sciaba'

Participants: Pepe, Julia, Daniel, Andrea Sciaba', Andrea Sartirana, Tony

J: Missing aggregation wrt FTS service and channel. J: Ranking plots to be added (trivial). J: Must understand what CMS needs for transfer operations.

J: 1) CMS can already know consume messages from the broker: well documented in Twiki. The broker is configured using virtual queues. Messages are not deleted unless CMS consumes them. Some of the consumer code may be reused. This scenario does not benefit from the Dashboard. Asking the developers to get also details on the configuration of the channel. The format of the messages is documented in the Twiki. IT-ES is owner of the broker and Julia will communicate with PES and GT as needed.

P: we might avoid implementing new features in our tool and do it in the Dashboard. When will the T1's provide an FTS with the MSG info?

J: only after FTS 2.2.8 is validated. Beginning of 2012 deployment will start if validation is successful.

2) read raw events from the Dashboard API. This should not make much sense. In the next version we will have transfer-level information kept for three months.

3) you can read statistics via API and this allows you to prototype the plots you need.

A: The overhead plots are provided in the FTS monitor.

J: The most reasonable scenario is 3), we will provide API for all the needed plots.

J: Please provide all requirements and we will put them in the Twiki. I will send the Twiki, send the code of the consumer and provide the virtual queue.

P: now we are using the tool for the LHCONE tests plus other tests, they are a good environment to work on.

### Minutes from meeting 20/03/2012 by David Tuckett

Minutes from meeting 20/03/2012 by David Tuckett

Participants: Andrea Sciaba', Daniel, David, Nicolo', Pepe, Julia,

P> Reviewing previous minutes:

- We now have filtering by FTS server. Can we also have filtering/aggregation by FTS channel. [J> Yes, this is a high priority new feature, also requested by ATLAS].

D> Dashboard UI current status:

- We have extracted a common UI framework from WLCG Transfers Dashboard and ATLAS DDM Dashboard. We should finish porting the WLCG Transfers Dashboard to this framework this week allowing independent development of new features for these Dashboards. [P> Will the UI be the same?] Yes, it is an implementation change only.
- We aim to move to production server and database in the next couple of weeks.

- We aim to put in place alarms to alert when production FTS servers are not reporting transfers. [P> FTS transfers can be in bursts with no transfers for periods; it may not be easy to monitor this.]

P> Presenting slides: 20120319\_Feedback\_to\_WLCG\_Transfer\_Monitoring\_v2.pdf

- Slide 2. "Left Panel: VO/Servers" It is tedious to unselect VOs/Servers all the time. Proposal: Add ALL/NONE options. [D> Yes, we will certainly add this usability feature]. [N> Could the default selection for FTS servers be the production servers. I could provide the list we use for CMS.] [D> Yes, we would have to keep a semi-static list of production FTS servers.]
- Slide 3. "Problems with Capital Letters on filters" See bug report [BUG:92865](#)
- Slide 4. "Hosts / Tokens?" How can we filter by tier? [N> Sites are not necessarily in the same tier for all VOs but it should be fairly consistent for tier 1s.] [D> If you choose CMS as a single VO then you can use "T1\_" etc in the Sites filter. If you choose multiple VOs then there is currently no feature to filter by tier.] [J> Perhaps we could provide T0, T1, T2+ options.] [D> So for CMS there is a workaround but we will look into adding a tier filter, at least for the single VO view, which would be useful for other VOs.] Perhaps Host and Token should be renamed to SRM endpoint and SRM token.
- Slide 5. "Country filter - CMS-like?" This is already covered by the CMS naming for single VO view.
- Slide 6. "Variety of sites assigned to n/a Country" See bug report [BUG:92867](#)
- Slide 7. "Transfers CERN -> T1s (ERRORS?)" See bug report [BUG:92868](#) [N> If you need any help doing the comparison with PhEDEx come to see me.]
- Slide 8. "Transfers CERN -> T1s (Units?)" Does WLCG standardize on powers of 10 (MB/s) or powers of 2 (MiB/s)? [N> I think the WLCG standard is powers of 10 but PhEDEx still displays powers of 2.] I think we had to provide some reports to WLCG using powers of 2. [D> The data is stored in bytes and the conversion to units is done in the UI so we could easily add an option in the UI for both.]
- Slide 9. "Transfers CERN -> T1s (Site Names)" See bug report [BUG:92869](#)
- Slide 10. "Transfers CERN -> T1s (Efficiencies)" Could the efficiency plot be replaced by a quality plot as used in PhEDEx? [N> Both plots have their merits, perhaps we could show both]. [J> We already have a client-side implementation of this plot-style for SUM.] [D> So we will see how much work it is and try to do it.]
- Slide 11. "Efficiencies smoothing effect?" See bug report [BUG:92870](#)
- Slide 12. "Auto for Date Axis" The AUTO option for STEP does not seem to work when many bins are displayed. [D> We are aware of this issue. The AUTO option just means identical consecutive labels are not repeated. It is difficult to know programmatically if the labels will overlap but we will see if it is possible to make a more intelligent AUTO option.] In any case, it is useful that the user can change plot options manually.
- Slide 13. "Is 'error samples' ok?" See bug report [BUG:92868](#)
- Slide 14. "Error information" Could we have the option to expand the error samples to see the complete list of errors of a given type? [D> We store transfer details for 90 days so we could display the error details. In the ATLAS DDM Dashboard, we do this by linking to error details in another tab. We could use the same approach.] [J> I am sure VOs and Sites will want this for debugging failures.] Is the error categorization sufficient? [D> At the moment, we simply use the error code provided by FTS but this does not seem fine-grained enough. Perhaps we could use a pragmatic approach such as grouping by error code + message length as done in DDM Dashboard.] [N> There is no systematic way to do this. One approach is to group by truncated message. Another is to group by message with bracketed content removed. In any case, it will be a hack.] [D> We will investigate and try to come up with a reasonable compromise.]
- Slide 15. "BIN SIZE vs STYLE" The outline style leaves gaps between bins; this could be confusing. Could we use the basic style as default? [D> Sure, we will set the basic style as default.]
- Slides 16 & 17. "Detail: sum plots format" The labels are sometimes cut and this is worse when more grouping options are selected. [D> This is a known issue with the plotting library. When it is fixed we will upgrade the plotting library. In the meantime, the workaround is to resize the plot.]
- Slide 18. "Plots that can be added: Ch. Occupancies" Could you provide this plot based on the periodic queue status messages? [N> The status messages are sent every 10 minutes whereas some

transfers take less than 10 minutes, so you may need more frequent messages. But 10 minutes is a good place to start.]

- Slide 19. "Plots that can be added: Throughput/stream" [N> This plot is useful to see that the channels are being fully utilized and to advise sites on how to optimize the FTS configuration. The channel and number of streams is included in the transfer messages, so you should have the data to plot this.]
- Slide 20. "Plots that can be added: SRM Overheads" Do you have all the data to produce this plot? [D> The information is in the tranfer messages but we are not currently generating statistics from it. We intend to do so and then producing this plot should be straightforward.] We would also like ranking of source/destination pairs for a channel and perhaps a stacked bar plot showing rates per stream broken down by source/destination pairs. I will send links to such plots.
- Slide 21. "Plots that can be added: Tx Durations" File transfer time and the percentage lost by SRM overheads per transfer. This is particularly interesting for optimizing SRMs. [N> This may be partially redundant with the previous plot.] [D> Once we generate the statistics for the SRM overheads then it should not be too much work to add both this and the previous plot.] [N> It would also be interesting to include transfer wait time in a similar plot but this would require that Michail add submission time to the transfer message.]

N> The fine granularity of plots in the Dashboard has already been useful for debugging an FTS issue. FYI: we found that for transfers by SRMCOPY, FTS reports average time for the files transferred in a single request.

J> Thank you for the very detailed and useful review. When we have compiled the list of feature requests we should meet again to agree priorities.

-- DavidTuckett - 17-Jan-2012

---

This topic: LCG > WLCGTransfersDashboard

Topic revision: r17 - 2012-10-30 - unknown



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback