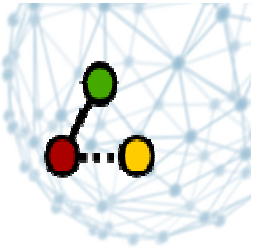

FTS3
Developers status update and road-map

Michail Salichos
IT/SDC

9 Oct 2013

FTS3 overview

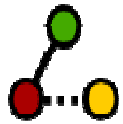

- Moving into production
  - CERN and RAL instances ready
  - Dedicated MySql database required
- Pilot services running for > 1 year at many sites
  - used for verification and certification
- Reached the level of stability and efficiency expected by the experiments?
  - Will run a “service challenge” to compare performance with FTS2

FTS3 status

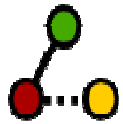

- In EPEL 6 (fts-*) + our continuous integration repo (stable)
- Heavily used by ATLAS for prod transfers
- WLCG FTS3 task-force still actively involved
  - Schedule of demos will be reduced, new functionality will be presented when available
- Installed at CERN, RAL, PIC, KIT, ASGC, BNL, IN2P3 and PNL
  - both prod transfers & testing
  - avg weekly transfer volume from RAL only: ~1.5PB

FTS3 status (2)

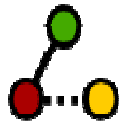

- **Latest stable version: 3.1.26**
- **Protocols support**
  - SRM, GridFTP, HTTP, xroot
- **DB back-ends**
  - Oracle, MySql (SQLite on-demand)
- **Deployment**
  - Horizontal scalability
- **Clients**
  - FTS2 clients compatibility
  - FTS3 CLI clients with new features (file/job metadata, etc)
  - Standard clients, REST-style interface for transfer submission and status retrieval
- **Resource management**
  - transfer auto-tuning / adaptive optimization
  - VO shares (+activity shares), endpoint centric configuration using JSON
  - smart transfer retry mechanism based on error classification
  - multiple replicas support
  - session/connection reuse (GridFTP, SSL, SRM KeepAlive)
  - job priorities
  - blacklisting users (DN) and SEs
  - Bulk staging files from archive
- **Logging**
  - debug mode transfer logging – GridFTP control channel info
- **The list goes on … https://svnweb.cern.ch/trac/fts3**

FTS3 usage

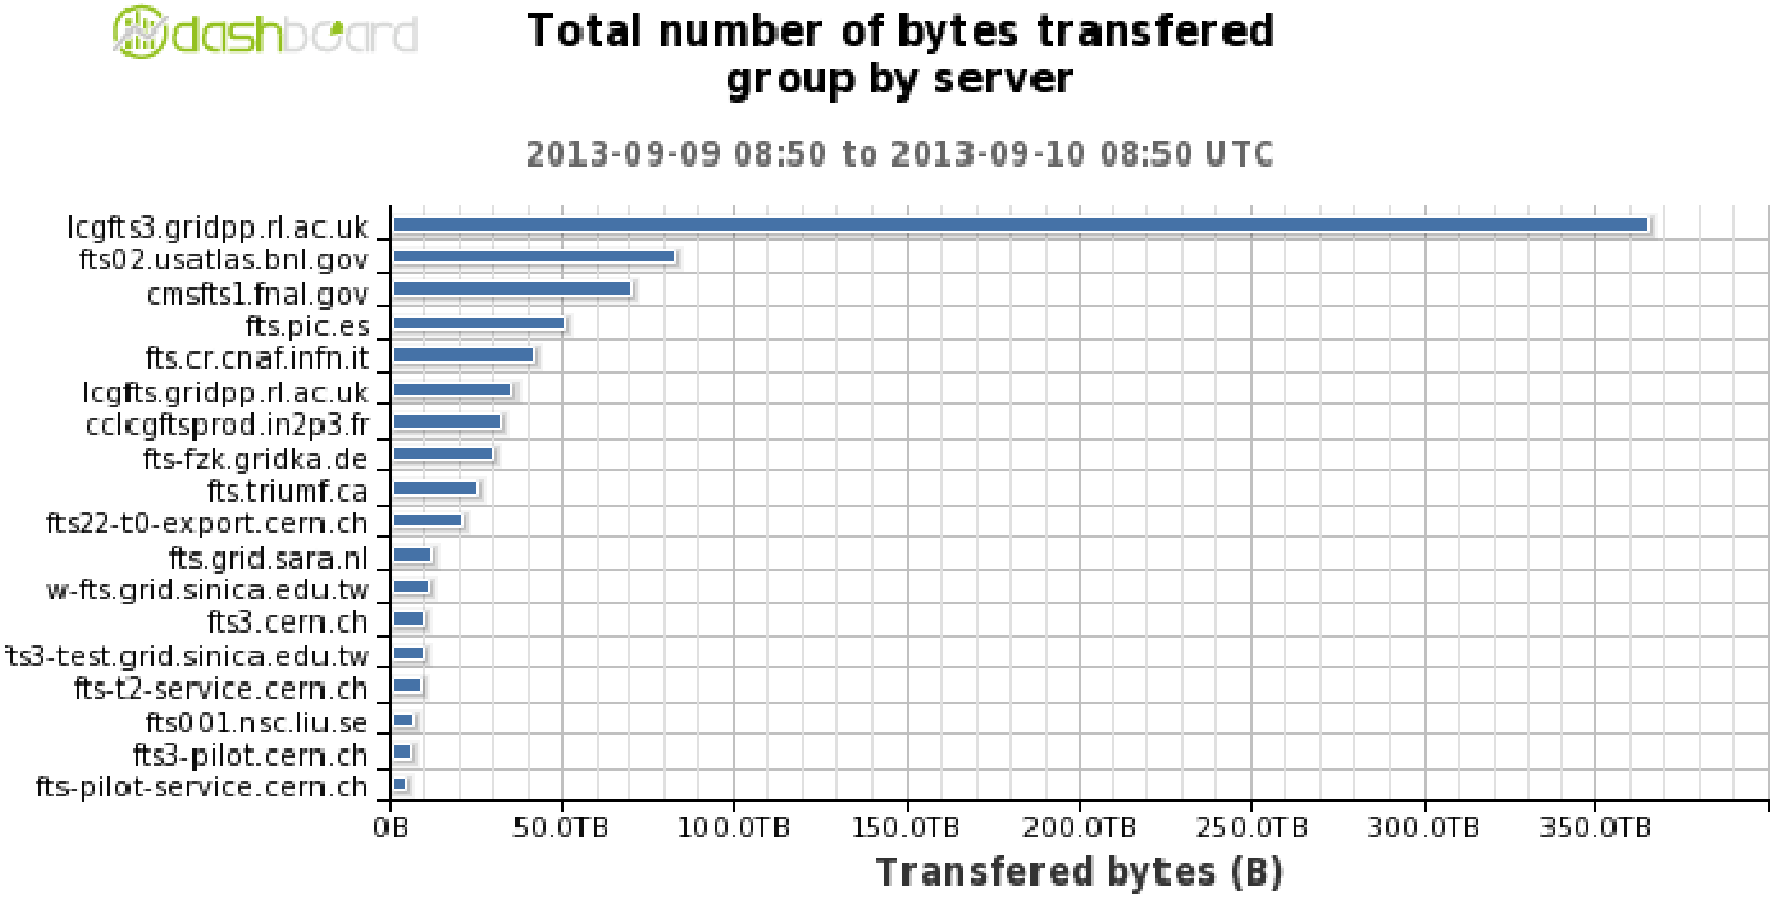
dashboard
Total number of bytes transfered
group by server
2013-09-09 08:50 to 2013-09-10 08:50 UTC
lcgfts3.gridpp.rl.ac.uk
fts02.usatlas.bnl.gov
cmsfts1.fnal.gov
fts.pic.es
fts.cr.cnaf.infn.it
lcgfts.gridpp.rl.ac.uk
cclcgftsprod.in2p3.fr
fts-fzk.gridka.de
fts.triumf.ca
fts22-t0-export.cern.ch
fts.grid.sara.nl
w-fts.grid.sinica.edu.tw
fts3.cern.ch
fts3-test.grid.sinica.edu.tw
fts-t2-service.cern.ch
fts001.nsc.liu.se
fts3-pilot.cern.ch
fts-pilot-service.cern.ch
0B
50.0TB
100.0TB
150.0TB
200.0TB
250.0TB
300.0TB
350.0TB
Transfered bytes (B)

FTS3 usage (2)

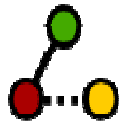

FTS2 vs FTS3 (CERN prod & RAL FTS3) last 7 days	FTS2	FTS3
Number of installations	13	2
Number of VMs (web service and VO/channel agents)	~38	8
Transfer volume	4.6PB	1.8PB
Number of files	~8.7M	~1.3M

FTS3 usage (3)

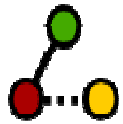

FTS2 vs FTS3 (CERN prod & RAL FTS3) last 7 days	FTS2	FTS3
VOs and sites	ATLAS/CMS/LHCb prod + debug transfers T1/T2/T3	**ATLAS:** Russian protoTier1 34 T2s out of total ~ 80 T2 (Total sites ~130) **CMS:** Currently used for Debug transfers. Total CMS sites active in Debug ~85 FROM all sites TO CERN FROM all sites TO selected T1s (4 out of 8 T1s: CCIN2P3, RAL, ASGC, JINR) FROM T1 RAL TO selected T2s (~10 out of ~50 T2s) FROM all T2s TO selected T2s in UK/EE/FI (~5 out of 50 T2s) **LHCb:** all transfers in/out at CERN, RAL, CNAF and 3 T2D sites total of sites with storage are CERN + 6 T1s + 3 T2Ds

FTS3 monitoring

- Dashboard transfers UI
  - http://dashb-wlcg-transfers.cern.ch/ui/
- FTS3 standalone web-app
  - http://www-ftsmon.gridpp.rl.ac.uk/fts3/ftsmon/#/
  - http://fts3.cern.ch/fts3/ftsmon/#/
  - https://fts3-pilot-mon.cern.ch/fts3/ftsmon/#/
- Nagios probes (CERN pilot)
  - http://fts3-pilot-mon.cern.ch/nagios/

FTS3 global monitoring

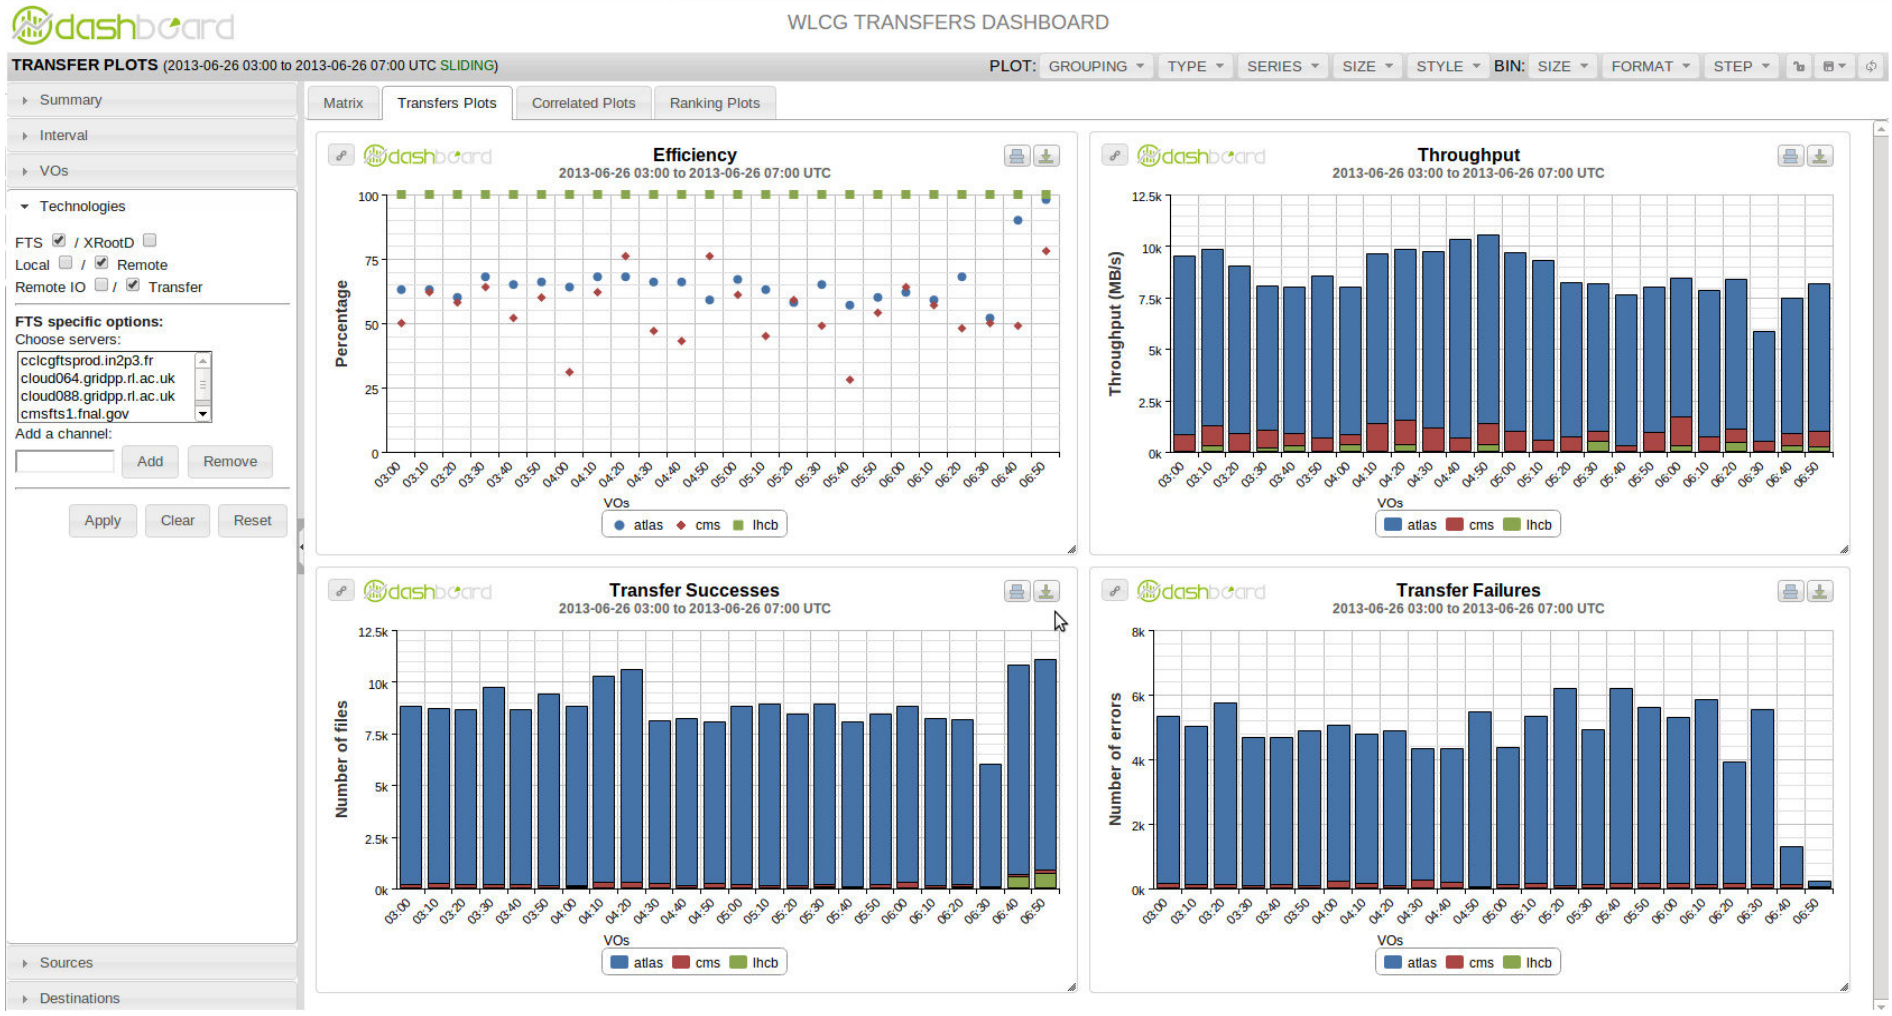

FTS3 global monitoring (2)

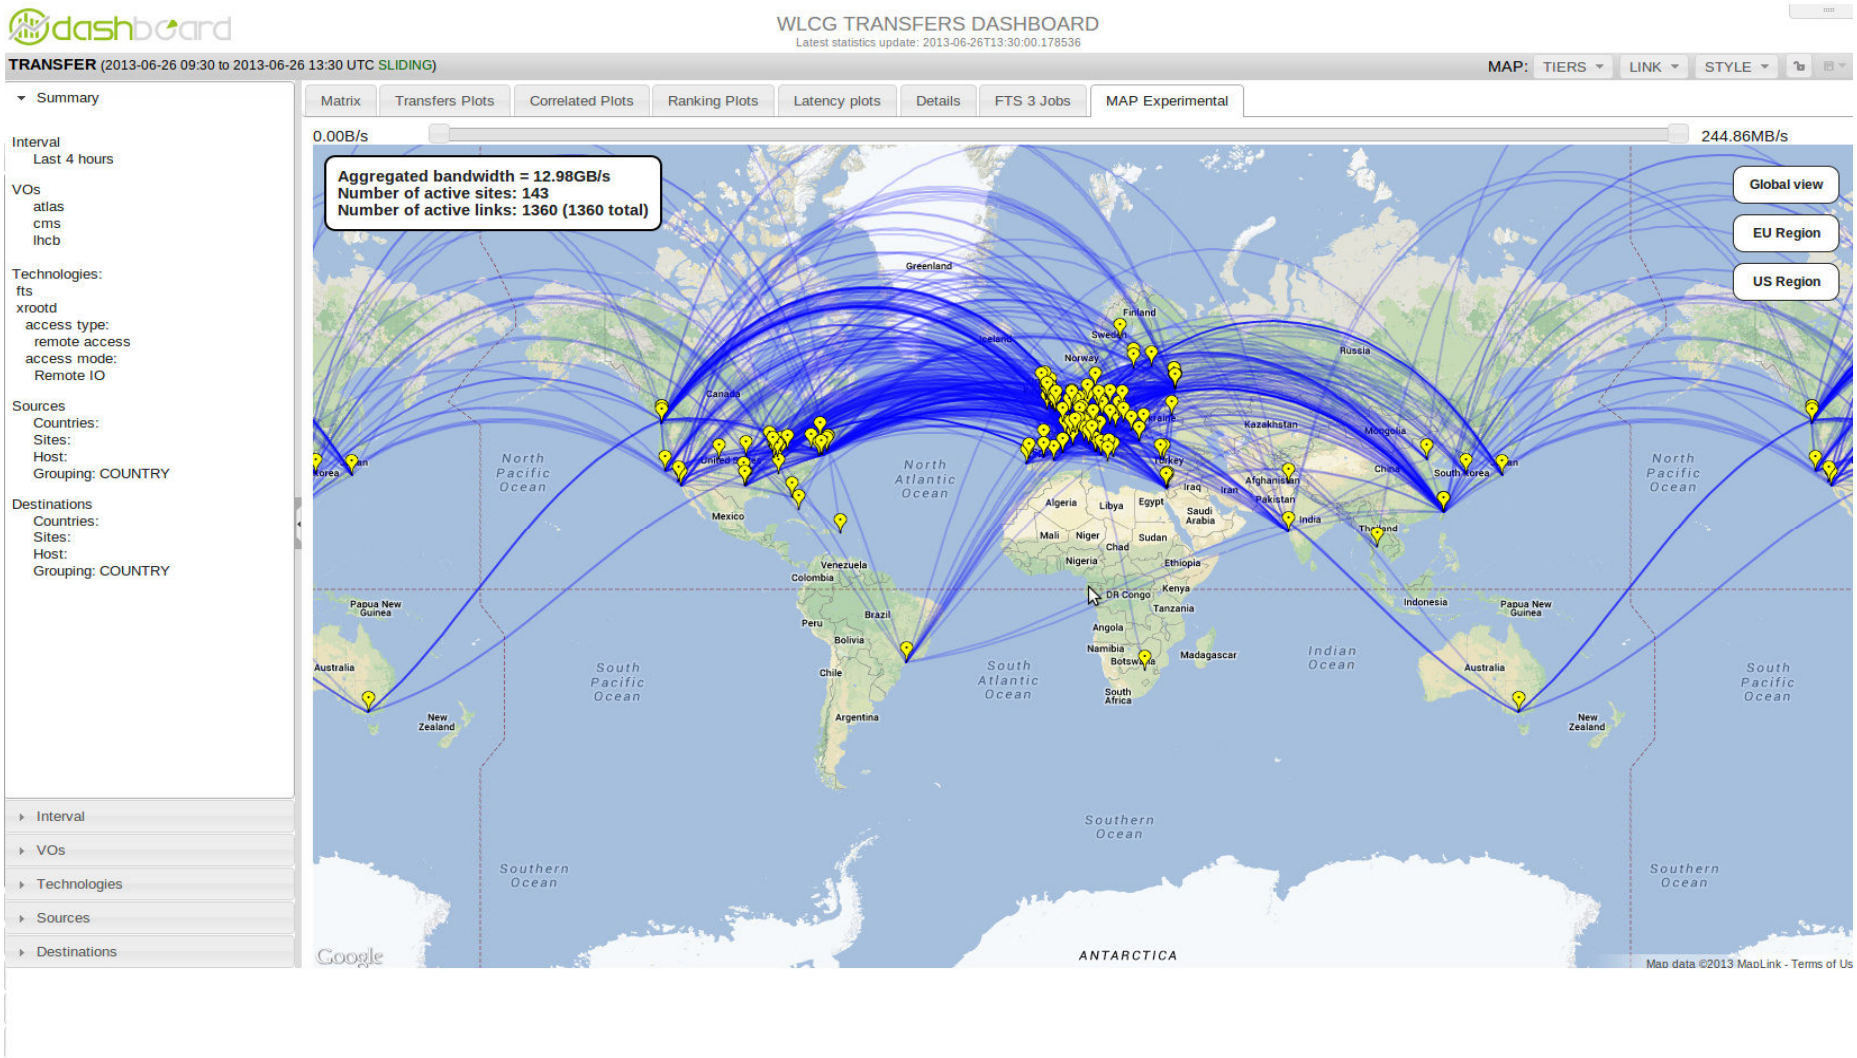

FTS3 standalone monitoring

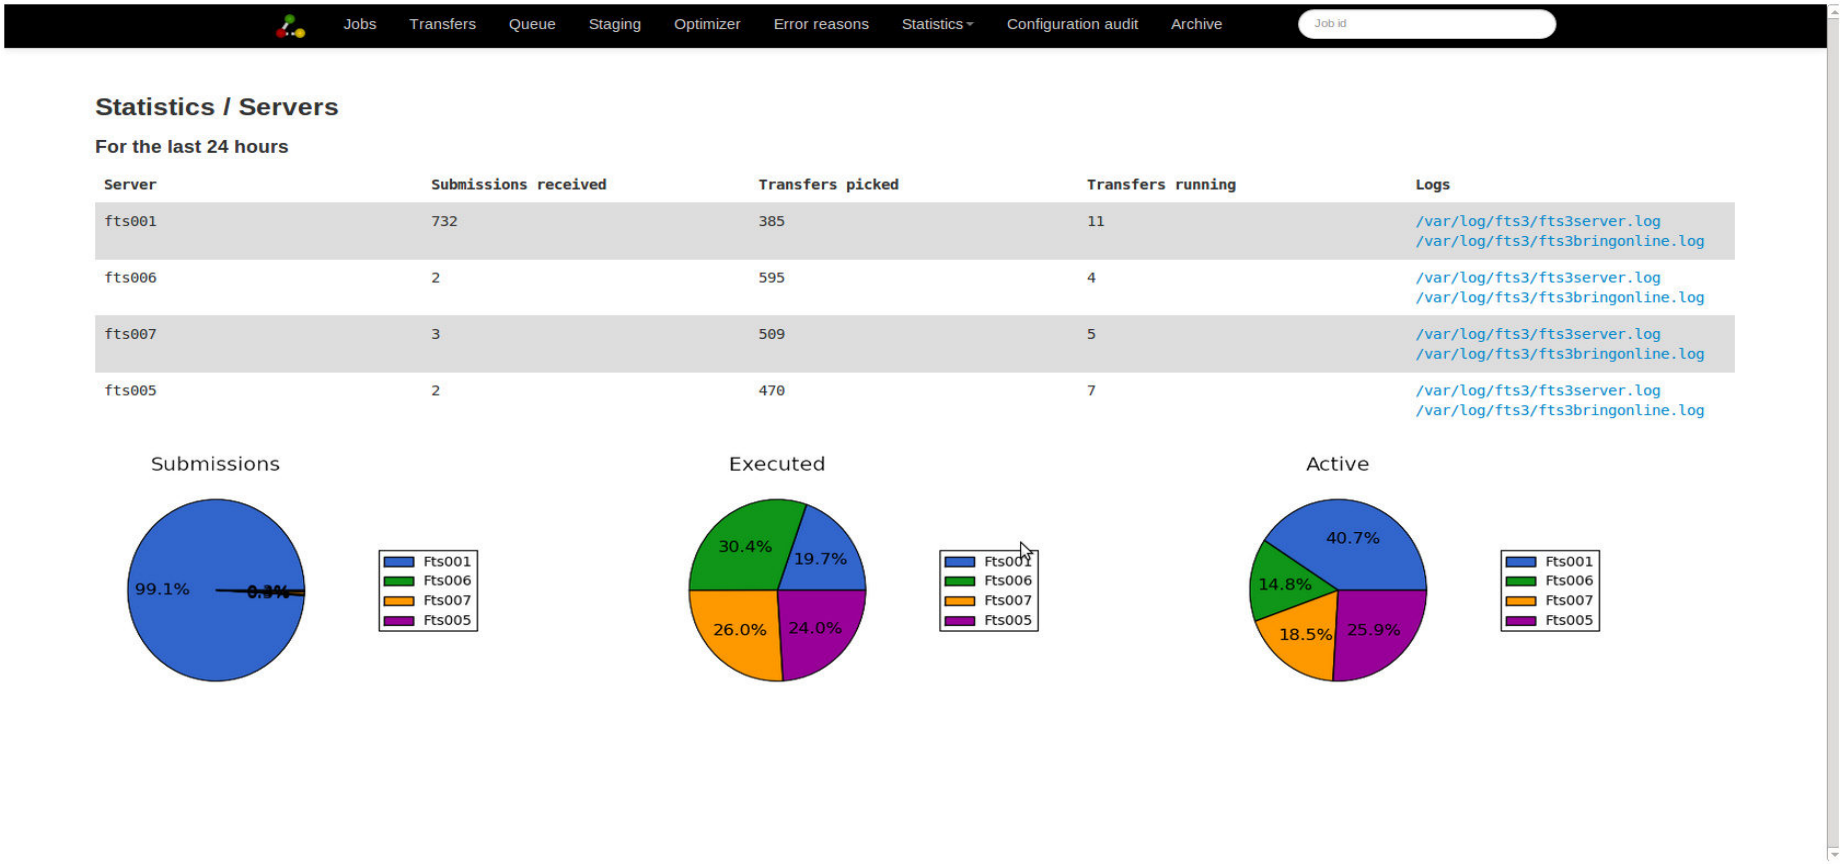

FTS3 evaluation & testing

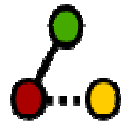

- LHC VOs (ATLAS, CMS, LHCb)
- EGI/EUDAT against globus GridFTP, dCache GridFTP and GridFTP interface for iRODS (Griffin)
- Many other VOs already tested it successfully: snoplus.snolab.ca, ams02.cern.ch, vo.paus.pic.es, magic, T2K, NA62, etc
- To be tested:
  - Exceed 1M file transfers per day using 1 instance only
  - http & xroot third-party transfers
  - FTS3 clients and new features (session reuse, job and file metadata info, multiple replicas, etc)

FTS3 experiences

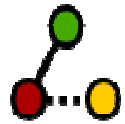

- Over the last year:
  - MySql performance tuning needed (InnoDB pool size, increase VM RAM, use query cache, etc)
  - Not optimized queries, many re-written
  - Redundant and duplicate indexes found
  - ATLAS spotted many errors – all have been fixed

FTS3 roadmap

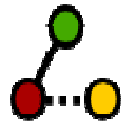

- Road-map entirely determined by experiment requirements and prioritization

- What's next:
  - Global scheduling and shared VO configuration across distributed FTS3 servers
  - Multi-hop transfers
  - Bulk file deletions
  - ATLAS VO shares per activity (primary, production, secondary, tier0, tier1, etc)
  - Integration and testing of perfSonar information (throughput & ping tests) for transfer optimization
  - deeper integration with archival storage and include high performance file management capabilities (deletes, renames...)

- Road-map -> https://svnweb.cern.ch/trac/fts3/roadmap

FTS3 summary

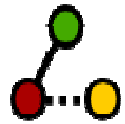

- Dev team is confident that a single FTS3 instance can cope with existing FTS2 load of all T0/T1 installations
  - On top of MySql
- Dev team will ensure that both deployment models will be supported
  - Distributed (similarly to FTS2)
  - Single instance

Questions ?