



FTS3 deployment model and FTS2 decommissioning

Nicolò Magini, Alessandro Di Girolamo,
Michail Salichos
IT/SDC

February 11th, 2014




FTS3 scale testing

- Progressively increasing traffic managed by RAL FTS3. Currently:
 - ~80% of all ATLAS DDM transfers
 - ~20% of CMS PhEDEx transfers
 - Will also add ASO (user data) test transfers
 - LHCb – 0% (currently on CERN FTS3)
 - Will also be switched temporarily if needed

FTS3 scale testing

- Revealed issues with leaking file descriptors at high submission rate, fixed.
- Continuing to increase load...
 - All VO transfers (production and tests)
 - Order of ~ 3M transfers/day

FTS transfer overview (by FTS instance)

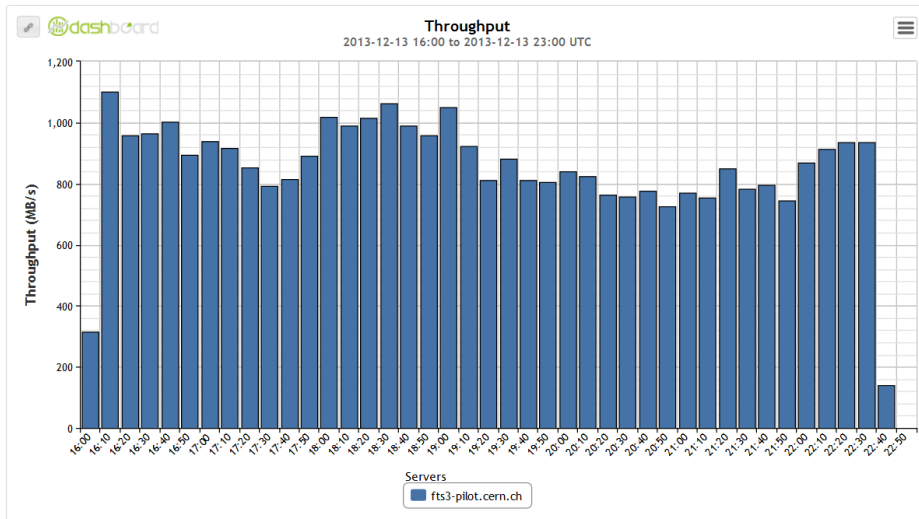


FTS3 transfer performance

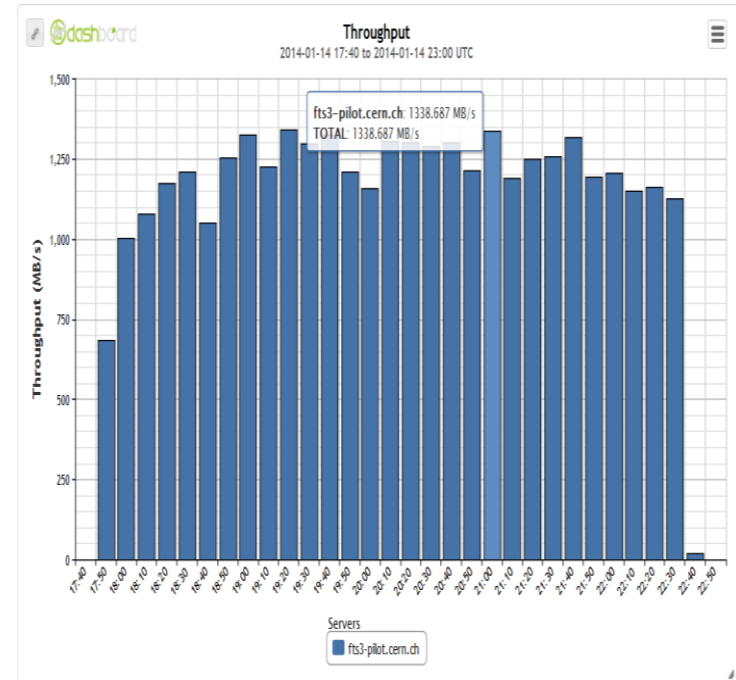
- Testing procedure in place:
 - verify FTS3 performances under controlled environment
 - with different configurations
 - after updates
 - The procedure:
 - Submit jobs to fts3-pilot.cern.ch
 - Monitor metrics in Dashboard
 - Delete destination files, repeat...

FTS3 transfer performance: one example

20*1000 files, 1GB/file



Fixed config (60 files, 20 streams)



Auto-config optimizer

✓ **FTS3 Optimizer can increase throughput**

FTS3: understanding optimizer

- Ongoing: adding detailed optimizer monitoring

FTS3 configuration management for WLCG

- Default configuration is auto-configuration, left to optimizer
- Configuration is managed centrally by team of FTS and experiment contacts
 - Settings for specific endpoints, in particular **max active files** protection, can be applied if needed, e.g. site request
 - Procedures under discussion (for now email to FTS3 steering)
 - Central team takes care of adjusting conf to VO policy

FTS3 deployment proposal

- Limited number of FTS3 servers needed to handle full scale (potentially even 1!)
- Multiple servers useful for service resilience
 - Minimize impact in case one instance unavailable
- Proposal:
 - Deploy very few servers (3?) in production:
 - Volunteers? CERN and RAL are obvious candidates...
 - The 3 LHC VOs could distribute their load among the servers. Just an example:
 - ATLAS & CMS “round-robin” between all servers
 - LHCb on CERN FTS3 100%, with RAL as backup

Questions?

backup



FTS3 configuration management for WLCG

- !!!!! Minimize as much as possible the configuration
 - If optimizer not optimal, we need to understand it and fix it
- Settings for specific endpoints, in particular **max active files** protection, can be applied on request
 - Central team takes care of adjusting conf to VO policy
 - e.g. if policy is round-robin among N servers, apply $maxActive=M/N$ on each FTS3 server to get M maximum transfers to endpoint
 - Shared endpoints: apply VO share, and then set the limit
- No need to automatically sync between servers, each server can fetch central config

Data Transfers (last7days) one ATLAS T2

Browser address bar: &grouping.dst=(site,cloud)&grouping.src=(site,cloud)&p.grouping=activity&p.type=sum&tab=transfer_plots

ATLAS dashboard

ATLAS DDM DASHBOARD 2.1.1

TRANSFER (2014-02-04 00:00 to 2014-02-11 00:00 UTC SLIDING)

Summary

Interval: Last 7 days

Activities:

- Data Brokering
- Data Consolidation
- Debug
- Functional Test
- Group Subscriptions
- Production
- Staging
- T0 Export
- Upload/Download (Job)
- Upload/Download (User)
- User Subscriptions

Sources:

- Tiers:
- Clouds:
- Countries:
- Federations:
- Sites: ROMA1
- Tokens:
- Grouping: SITE CLOUD

Destinations:

- Tiers:
- Clouds:
- Countries:
- Federations:
- Sites:
- Tokens:
- Grouping: SITE CLOUD

Interval

Activities

Sources

Destinations

