FTS

File Transfer Service

FTS news and plans

3rd Rucio Community Workshop

Edward Karavakis
on behalf of the FTS team

FTS

File Transfer Service

- Distributes majority of LHC data across WLCG infrastructure
- 7 WLCG and 14 non-WLCG instances
- ~28 Virtual Organisations
  - ATLAS, CMS, LHCb, AMS, NA62, Compass, ILC, Magic, Belle Ⅱ, Mice, Xenon, Snoplus, GridPP, DUNE, LZ, Solidexperiment.org, SKA, Ligo, Icecube, Elixir, NP02 (part of DUNE), CAST, ESCAPE, Eiscat.se, Virgo, BES III, JUNO, Pierre Auger Observatory, CEPC
- Integrated with experiment frameworks: Rucio, PhEDEx, DIRAC
- Transferred in 2019 more than 800 million files and 0.95 Exabyte of data

FTS Core Features

Simplicity

- Easy user interaction for submitting transfers. Copy one file from one place to another
- WebFTS portal for end-users, Real Time monitoring and Web Admin

Reliability & Integrity

- Checksums and retries are provided per transfer

Flexibility & Scalability

- Multiprotocol support (HTTP, gsiftp, xrootd, SRM, S3,..)
- Different clients to access the service (REST APIs, python bindings)
- Transfers from/to different storages
- Support for bringonline
- FTS can be run "zero config"

Intelligence

- Parallel transfers scheduling and optimisation to get the most from network without burning the storages
- Priorities/Activities support for transfers classification

Plans for 2020

Tokens

- Authentication: WLCG is moving from X.509 certificates to HTTP tokens
- FTS 3.10 supports OpenID Connect
- Pre-release server and client RPM packages available
  - http://fts-repo.web.cern.ch/fts-repo/xdc/el7/x86_64/
  - Get fts-rest-cli-3.10.0 to use the client
  - Example next

Token example

- Get an access token and submit a job
  - https://github.com/indigo-dc/oidc-agent

```
$ export tok=`oidc-token wlcg`
```

```
$ fts-rest-transfer-submit \
-s https://fts3-xdc.cern.ch:8446 \
--access-token $tok \
https://prometheus.desy.de/Users/carles/source \
https://prometheus.desy.de/Users/carles/destination

Job successfully submitted.
Job id: 2fb22734-360f-11ea-9524-fa163e362acc
```

Token workflow

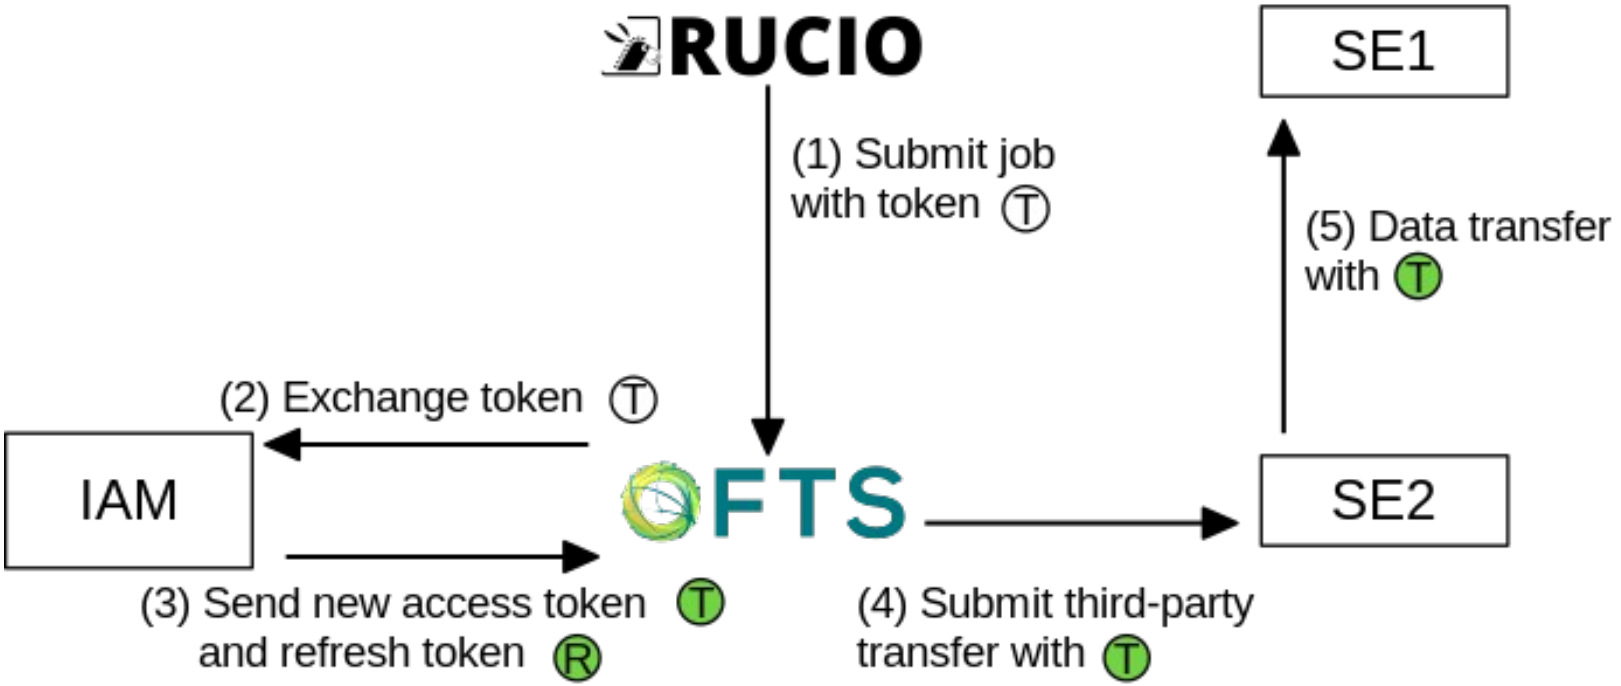

WebFTS

- A web-based file transfer & management solution for small-sized experiments
- Users can invoke data transfers on distributed infrastructures directly from within their web browser
- OIDC integration work was done at KIT
- WebFTS+OIDC will be released soon

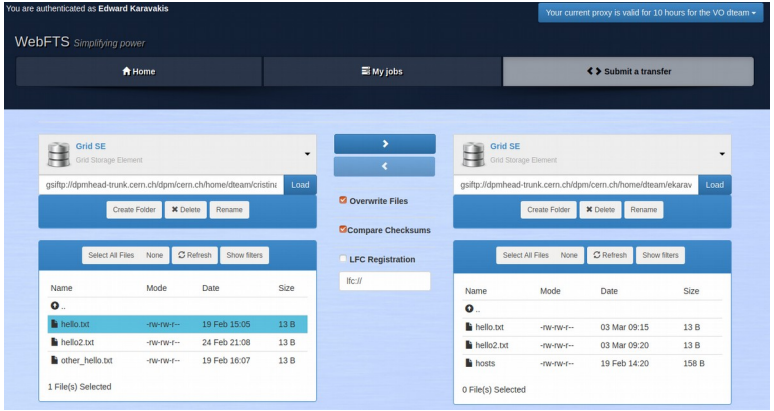

REST and monitoring plans

- FTS-REST new framework and move to Python 3
  - Python framework that FTS-REST is built upon is obsolete (pylons)
  - Already started migrating it to Flask
  - API won't change and most probably will gain performance/stability by using an up2date technology

- Web monitoring to Python 3 and latest Django framework

Scalability improvements

- First set of improvements were released already in 3.9 series
  - Adding missing indices by examining all slow queries and passing them through a profiler
  - Optimisations to get rid of expensive joins

- Work on DB table partitioning bringing significant performance gains in optimiser and scheduler
  - To be released in the coming months

Integration for tapes

- CTA is the new tape based solution at CERN that is built on top of EOS exposing an XRootD interface and supporting TPC
  - EOS+CTA integration is production ready – interface to FTS doesn't change for the experiments, everything handled transparently by FTS
    - Implemented staging via XRootD
    - Disk copy eviction on transfer completion, to better handle the reduced buffer size
    - Staging+Multihop supported (to handle data export to T1s)
  - Stress-tested during the ATLAS Data Carousel exercise

Monitoring the migration to tape

CERN Tape Archive

- When transferring to tape system, tape migration is not taken into account by FTS
  - Transfer successfully completed at the storage system disk buffer level only
  - Clients have to explicitly check on the destination storage if file is on tape

- New feature is being implemented to report a transfer as completed only when file has been migrated to tape successfully
  - Included a new “ARCHIVING” state in the state machine
  - Clients need to enable this feature when submitting the transfer
  - Implemented for both XRootD and SRM (for any other than CTA)

- Buffer-aware scheduling in the future?

Central Data Recording (CDR) using FTS

- A tool developed by NA62 delegating to FTS the data movement from the experiment facilities directly to CASTOR - transferred several Petabytes of data
  - Checks a folder for new entries
  - Calls FTS to transfer data from local GridFTP to CASTOR public
  - Checks transfer status and keeps file metadata locally
  - Checks for the ‘m’ bit and deletes data if on tape

- FTS team plans to release a version (based on the NA62 CDR code) that is generic and can be used by any small experiment for DAQ, capable of writing to CASTOR/CTA-public or to any given endpoint

Conclusion

- FTS continues to evolve with the infrastructure as WLCG's principal data movement service
- Expanding community and adoption by upcoming data intensive projects
- Various performance improvements and new features put in place in preparation for Run-3
- Full support planned for TPC and token auth
- Integration with CERN's new tape archival system CTA and migration to tape monitoring

fts-devel@cern.ch / fts-support@cern.ch

https://fts.web.cern.ch

https://fts3-docs.web.cern.ch/fts3-docs/