



## *Updates and Plans*

*FTS@EGI Workshop 2021*

Mihai Patrascoiu  
on behalf of the FTS team

# Open Source software for reliable and large-scale data transfers within WLCG



## Features:

- x509 and OIDC tokens
- Staging + Archiving
- Multihop transfers
- Transfer Optimizer
- Cloud support
- Python bindings + CLI clients



**Intuitive**



**Robust**



**Flexible**



**Adaptive**



## *Developed at CERN*

- Mihai Patrascoiu (Project Leader)
- Steven Murray (Service Manager)
- João Lopes (C++ / Python developer)
- João Afonso (C++ / Python developer)

## *Projects under FTS umbrella*

- FTS (Server + QoS daemon)
- FTS-REST (Submission server)
- FTS-REST-clients (Python & CLI)
- FTS-Monitoring (Django Web UI)
- webFTS (Submission Web page)

## *...but not limited to CERN*

Thanks go to all our contributors  
and users!

## Data Management Clients

- Gfal2 (Grid file access library)
- Gfal2-python (Python bindings)
- Gfal2-util (Python CLI)
- Davix (Grid HTTP client)

**8 WLCG instances**

**16 non-WLCG instances**

**~36 Virtual Organizations**

**5 Token Identity Providers**

Transferred in total in 2019: 1.29 EBs and 1.10B files

Transferred in total in 2020: 1.01 EBs and 976M files

Transferred in total in 2021: 0.78 EBs and 880M files (so far)

| ~65% by CERN FTS instances



# FTS activity throughout 2021

- Major service stability improvements
- Development highlights
- New CI and Build platforms for DMC projects

New release stack: FTS v3.11.0 + Gfal2 v2.20.0 + Davix 0.8.0

Available in FTS/DMC production repositories

Ongoing deployment to production instances

# Service stability

- New database deployment: main + read-only replica
- Separation of service accounts
- Moved Web Monitoring component to read-only replica
  - Different caching patterns as to the server and qos daemons
  - Encapsulation when things go wrong
- Improved long-running staging query, which affected the overall DB performance
- Move to MySQL 8 (on-going, CERN-only so far)
- Improved internal monitoring + alarming

# Development highlights

- Swift protocol support in Gfal2 & Davix (FENIX collaboration)
  - thanks attributed to Shiting Long
- CS3API protocol support in Gfal2 & Davix (GSoC project)
  - thanks attributed to Rahul Chauhan
- Davix with libcurl backend
  - enable at runtime using `DAVIX_USE_LIBCURL=1` variable
- FTS tape destination file report (needed by WLCG experiments)
- Gfal2 SE-issued token retrieval (highly needed by WLCG)

# DMC CI & Build platforms

- Moved all DMC packages from Jenkins to gitlab-CI
- Platforms supported:
  - CC7, CentOS8
  - Fedora33, Fedora34
  - Fedora-Rawhide
- Fully Python3 compliant
- Dropped all support for platforms older than CC7

*DMC = Data Management Clients*  
*DMC Documentation ([link](#))*



# Plans for 2022

## Service stability

- Uniform deployment setup (MySQL8, read-only replica)
- Submission protection (throttling and limits)
- Profiling for core FTS operations

## WLCG activity

- IPV6 TPC monitoring & control (HTTP & XRoot)
- XRootD traffic tagging
- Better OIDC token management (fine-grained scopes, easier AT refresh)

## Core development

- Release Python3 FTS-REST-Flask
  - deprecate old clients (Python2 & C++)
- Finalize Swift support in FTS
- FTS on MySQL8 (not just CERN-only)
- Exploring per-site (multiple SEs) limits
- Benchmark FTS with different queueing backends
- Revise FTS solution for SMEs

# Conclusion

- Welcome new team members → optimistic outview on 2022 effort
- Run3 around the corner → FTS focus on service stability
  - Conscious decision on protecting the service more
  - Better DB setup, better health monitoring & user-oriented alarms
- We warmly welcome R&D projects and external contributions
  - Swift and CS3API as great examples of 2021
- FTS is a core component in the WLCG strategy for Run3 and Run4
  - It's time to think about Run4 scalability
- We acknowledge using FTS at an SME requires a large amount of pre-setup

# Thank you!

 [cern.ch/fts](https://cern.ch/fts)

 [fts-devel@cern.ch](mailto:fts-devel@cern.ch)

 [cern.ch/fts3-docs](https://cern.ch/fts3-docs)

 <https://gitlab.cern.ch/fts/fts3>

 [cern.ch/dmc-docs](https://cern.ch/dmc-docs)

 <https://gitlab.cern.ch/dmc/gfal2>

# Storage-Element issued tokens [backup]

Storage-Element issued token = exchange of a x509 proxy certificate for a bearer token issued by the SE, for a given path and a given set of capabilities

- SE-issued token retrieval now part of core Gfal2 functionality
- Enables full support for FTS token-based SRM+HTTPs transfers
- Essential step needed for decommissioning GridFTP from tapes
- Deployed on FTS3-Pilot (soon on all production instances)
- Feature also exported via Python bindings and CLI