




## *5th Rucio Community Workshop*



**FTS**  
File Transfer Service

## *FTS 2023: Plans and Direction*

Mihai Patrascoiu  
on behalf of the FTS team

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

## Features:

- TPC Orchestration
- Tape Operations (over multiple protocols)
- Certificate and token auth
- Multihop transfers
- Transfer Optimizer
- Cloud support
- Python bindings + CLI clients



...and many others

# FTS Team



- Mihai Patrascoiu (Project Leader) [CERN]
- Steven Murray (Service Manager) [CERN]
- João Lopes (C++ / Python developer) [CERN]
- Ed Dambik (C++ developer) [ATLAS / Indiana University Bloomington, USA]
- Eraldo Silva Junior (Python developer) [ATLAS / LHCb / CERN / CBPF, Brazil]
- *Tom Hepworth* (C++ / Python / Grafana guru) [2022 CERN Summer Student]

...and thanks to many other past and present contributors



# FTS Ecosystem



## *Projects under FTS umbrella*

- |                  |                       |                  |
|------------------|-----------------------|------------------|
| - FTS            | (Server + QoS daemon) | [C++]            |
| - FTS-REST       | (Submission server)   | [Python, Flask]  |
| - FTS-clients    | (Python & CLI)        | [Python]         |
| - FTS-Monitoring | (Django Web UI)       | [Python, Django] |
| - webFTS         | (Submission Web page) | [PHP]            |

## *Data Management Clients*

- |                |                            |                |
|----------------|----------------------------|----------------|
| - Gfal2        | (Grid file access library) | [C++]          |
| - Gfal2-python | (Python bindings)          | [C++, BoostPy] |
| - Gfal2-util   | (Python CLI)               | [Python]       |
| - Davix        | (Grid HTTP client)         | [C++]          |
| - SRM-IFCE     | (SRM interface for Gfal2)  | [C, gsoap]     |
| - CGSI-GSOAP   | (gsi interface for Gfal2)  | [C, gsoap]     |

FTS and DMC clients  
published to PyPi, EPEL,  
Debian\*

\*Special thanks to Mattias Ellert





# submissions by issuer (last 30d)



*All submissions vs Rucio submissions (past 30d)*

# FTS – v3.12.0 series

- Python2 → Python3 migration
- FTS-REST → FTS-REST-Flask rewrite
- Jump to MySQL8
- Tape REST API (**v3.12.2 RC**, September 2022)
- WebMonit “linkinfo” feature (Eraldo contribution)
- HTTP-TPC IPv4/6 flag reporting (Ed contribution)
- Improved support for S3v4 transfers
- Eviction configurable per-StorageEndpoint
- Multihop “--overwrite-hop” feature
- Deprecation notice in C++ clients
- GDPR / OC11 compliance  
(VO name generation for certs w/o VOMS data)
- ActiveMQ format documentation up-to-date
- All DMC & FTS clients published to EPEL

... and many others



# Web Monitoring “linkinfo” feature

- Allows transfer operators to quickly identify why a link is not progressing (e.g.: number of transfers for a <src> → <dst> link is “subpar”)
- A pain-point in a transfer operator’s day-to-day life, many times requiring assistance from the FTS team
- The answer lies in the storage and link limits, which have now been made fully transparent and easily accessible via the Web Monitoring (e.g.: fts3-atlas.cern.ch:8449)
- Successfully initiated Eraldo in the world of FTS development !

The screenshot displays the 'linkinfo' web monitoring interface. It features a table with three columns: Source, Destination, and Link. The Source column contains details for 'Active transfers: 114', 'Outbound limit: 100', and 'Config type: Specific'. The Destination column shows 'Active transfers: 105', 'Inbound limit: 300', and 'Config type: Specific'. The Link column lists 'Active transfers: 96', 'Min limit: 10', 'Max limit: 100', and 'Config type: Generic'. Below the table is an 'Optimizer' section with the following text: 'Active transfers: 96', 'Decision: 150', and 'Description: Good Link Efficiency, Current Average Throughput Is Larger Than The Preceding Average. Hit Upper Range Limit'. At the bottom of the interface are two buttons: 'Storage Config' and 'Configure Link'. A 'Close' button is located in the bottom right corner of the window.

Source	Destination	Link
Active transfers: 114 Outbound limit: 100 Config type: Specific	Active transfers: 105 Inbound limit: 300 Config type: Specific	Active transfers: 96 Min limit: 10 Max limit: 100 Config type: Generic

**Optimizer**

Active transfers: 96  
Decision: 150  
Description: Good Link Efficiency, Current Average Throughput Is Larger Than The Preceding Average. Hit Upper Range Limit


Storage Config | Configure Link

Close

# IPv6 reporting for HTTP-TPC transfers

- HTTP-TPC IP version extracted via Davix, propagated to FTS
- Change implemented by Ed across the many layers: FTS <> Gfal2 <> Davix

*Change deployed in August 2022 (at CERN)*





# Immediate pipeline: evolve IPv6 reporting

## Objective:

A way to tell between “not ipv6” and “no info available”

## Approach:

- Introduce a new co-existing field `<ipver>`
- Existing dashboards / scripts won't be impacted by the addition of a new field
- Concerned dashboards / scripts **will have to adopt** the new field
- The `ipv6` is no longer guaranteed long-term longevity  
(e.g.: may disappear in 1y+)

```
ipv6  = true | false (boolean type)
ipver = ipv4 | ipv6 | ipv4+6 | unknown
      (string type)
```

# Immediate pipeline: metadata for tape

## Objective:

Allow clients to pass tape metadata for staging and archiving (per file)



```
--staging-metadata / {"staging_metadata"} = <metadata>  
--archive-metadata / {"archive_metadata"} = <metadata>*
```

*Staging metadata* – Passed as specified in the Tape REST API

*\*Archive metadata* – Submission ground is ready, but no action yet

*(data will sit in FTS DB, workflow to be defined)*

# FTS – Packages and Platforms

## Server

- `fts-server` / `fts-rest-server` / `fts-monitoring`
- Only available on CC7 ! (move to next platform once Linux strategy decided)
- Packages only available via the FTS repositories

## Clients

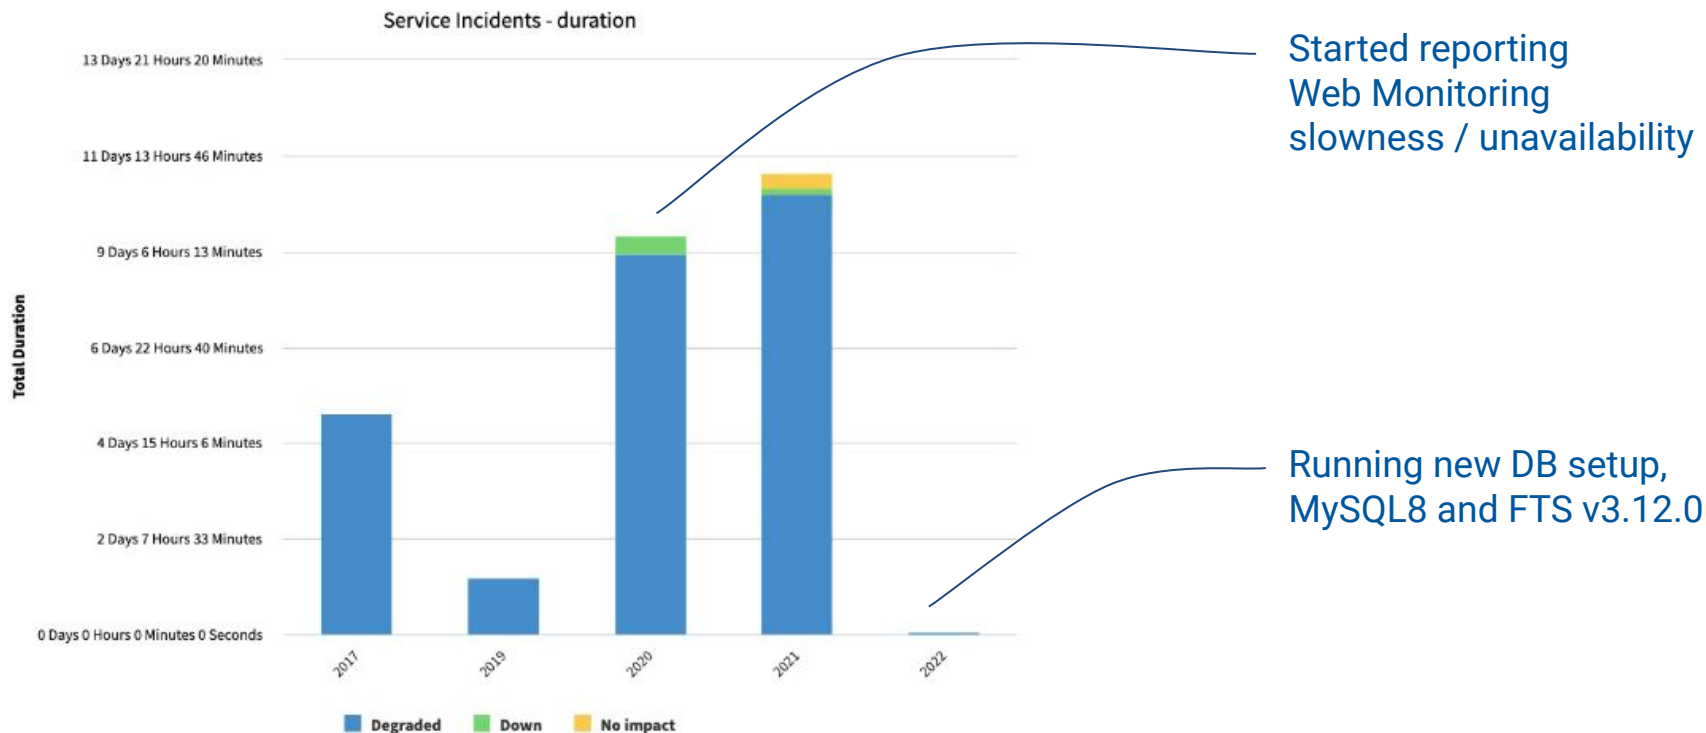
- `fts-rest-client` / ~~`fts-client`~~ (deprecated C++ client)
- `gfal2` / `gfal2-python` / `gfal2-util` / `davix`
- Packages available via EPEL (covers CC7, CS8, CS9, Fedora)

# FTS – Releases: change of approach

- Starting with FTS v3.12.0 + MySQL8: schema upgrades are no longer disrupting to the service operations
- **The FTS team will treat schema upgrades as non-disruptive to the service !**
  - No more need to bundle many schema upgrades into one big change
  - Allows us to release more often, including new features in patch releases
- We encourage all FTS sites to upgrade to v3.12.0 and MySQL8 (**or equivalent\***)

*\* MySQL remains the CERN FTS deployment choice, but FTS will not break other implementations (e.g.: MariaDB) on purpose*

# FTS – Service @ CERN



# What's next?

# FTS - Token support for exascale transfers

*Native support for a scalable infrastructure based on tokens for the critical workflows of all the LHC experiments*

## New person joining FTS team in Q1 2023

- Will liaison with WLCG AuthZ WG
- Active FTS participation in the next AuthZ workflow

- Assess current FTS software stack w.r.t. tokens
- Discuss and review main workflows with all stakeholders
- Design and develop fully integrated token support for all supported protocols
- Design and develop automated token refresh mechanism
- Design and develop fully integrated token support for cold storage
- Test individual components and service scalability
- Develop and deploy a functional testing framework
- Deploy a prototype service with full token support
- Validate main experiments' workflows
- Highlight and address shortcomings and scalability issues

# FTS - ALTO Collaboration

Collaboration between Dr. Richard Yang's group (Yale University) and FTS team


## Goal

- Incorporate networking usage information from ALTO
- Improve FTS Scheduling and Optimizer decision
- Full details explained in Richard Yang's talk

| ALTO/TCN: Rucio/FTS Control with Deeper Network Visibility ([link](#))



# FTS - The error reporting storm



Too many times, HTTP errors are not clear enough

→ FTS is in the middle of it, assisting sites and users in debugging

Need a better way of HTTP error reporting across the grid !

→ Will collaborate with DOMA-BDT on achieving this

# FTS - Through the user's lens

## Overview

Showing 1 to 50 out of 1631 from the last 1 hour



	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...	Next	Last
Source	Destination	VO	Submitted	Active	Staging	S.Active	Archiving	Finished	Failed	Cancel	Rate (Last 1h)	Thr.						
+ davs://fal-pygrid30.lancs.ac.uk	davs://xgate.hec.lancs.ac.uk	atlas	86054	32	-	-	-	431	-	-	100.00 %	194.59 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://atlaswebdav-kit.gridka.de	atlas	14443	10	-	-	-	12	33	-	26.67 %	4.07 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://tech-gftp.hep.technion.ac.il	atlas	10939	12	-	-	-	20	29	-	40.82 %	5.17 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://se-goegrid.gwdg.de	atlas	9087	10	-	-	-	15	12	-	55.56 %	5.20 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://goliath100.farm.particle.in2p3.fr	atlas	8320	9	-	-	-	11	27	-	28.95 %	3.74 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://eosatlas.cern.ch	atlas	7008	10	-	-	-	12	33	-	26.67 %	5.41 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://bohr3226.tier2.hep.manche.ac.uk	atlas	6612	10	-	-	-	7	29	-	19.44 %	2.17 MiB/s						
+ davs://gdav1.physik.uni-mainz.de	davs://clrlcgse01.in2p3.fr	atlas	5723	11	-	-	-	16	12	-	57.14 %	4.66 MiB/s						

Our Documentation and Web Monitoring is too developer oriented

→ This causes confusion for our end users

Will need volunteers to “walk us” through the pages and highlight weak points !

# FTS - Next development focus?



Service running

FTS Scheduler rewrite

Scalability concerns

Better service health  
monitoring

Reduce legacy code

Community requests

FTS global config

Transfer limits per Site

Aggressive Optimizer

Improve Cloud support

Improve error reporting

- Need a way to prioritise between service and stakeholder requests

*\*Lists not exhaustive*

# FTS - Revive the steering meetings

Last FTS Steering Meeting: *too long ago*

Ultimately, need a way to understand and prioritise what's most important to our communities...

→ ...so why not simply ask?

## Steps

- FTS to reach out to main stakeholder communities asking for requirements
- Gather and structure stakeholder requirements
- Organize the next FTS Steering meeting
- Negotiate with stakeholders on items

# Conclusions

- Last years show a maturing FTS deployment at CERN
  - confident in FTS abilities during Run3
- From v3.12.0: releases and schema changes are no longer painful
  - allows us to release more frequently
  - better meet client demands
- HTTP error reporting must be addressed
  - we must allow our users to understand what goes on without needing expert intervention
- FTS continuously evolves to meet community demands
  - revive steering discussions and establish priority items

# Thank you!



- Issue tracking: JIRA [FTS](#) / [DMC](#)
- Code: Gitlab (CERN) → mirrored on Github
  - 📄 <https://gitlab.cern.ch/fts/fts3>
  - 📄 <https://gitlab.cern.ch/dmc/gfal2>
  - 📄 <https://gitlab.cern.ch/dmc/davix>
- **Non-formal communication:** ~IT-FTS (Mattermost, CERN)
- **E-mail:** [fts-devel@cern.ch](mailto:fts-devel@cern.ch) / [dmc-devel@cern.ch](mailto:dmc-devel@cern.ch)
- **Announcements:** [fts3-steering@cern.ch](mailto:fts3-steering@cern.ch)
- **FTS Service Managers:** [fts-ops-forum@cern.ch](mailto:fts-ops-forum@cern.ch)
- **User support:** [fts-support@cern.ch](mailto:fts-support@cern.ch) / Service Now (CERN)  
/ GGUS (experiments, WLCG)
- Documentation:
  - 📄 [cern.ch/fts3-docs](http://cern.ch/fts3-docs)
  - 📄 [cern.ch/dmc-docs](http://cern.ch/dmc-docs)
  - 🌐 [cern.ch/fts](http://cern.ch/fts)

