

File Transfer Service – FTS

EOS 2024 Workshop

João Lopes on behalf of the FTS team

Thursday 14th March 2024

Outline




- **What is FTS?**
- **Components of FTS**
- **How does FTS work?**
- **How to use FTS?**
- **Who uses FTS?**
- **2023 Operations overview**

What is FTS?

FTS is an open-source software for large scale queuing and reliable execution of file transfers

Capabilities:

- Orchestration of Third-Party Copies (TPCs)
- Streams transfers through itself if TPC is not supported
- Tape storage operations via the WLCG HTTP Tape REST API, SRM and XRootD
- Support for Cloud based storage
- Certificate and token authentication



What is FTS?



Intuitive

- Simple REST API for submitting transfers.
- Convenient command-line tools and Python bindings
- Minimal configuration requirements



Robust

- Automated verification of destination file sizes and checksums
- Automated retries of failed file transfers



Flexible

- Multi-protocol support: HTTP, SRM, XRootD, GridFTP, S3, ...
- Web-based real-time monitoring
- Web-based runtime configuration



Adaptative

- Runtime optimization to maximize throughput without overloading storages
- Activity based bandwidth reservation

What is FTS?



FTS team:

- Mihai Patrascoiu (Project Leader)
- Steven Murray (Service Manager)
- João Lopes (C++ / Python developer)
- Shubhangi Misra (Token authorization/Web developer)

Projects under FTS umbrella:

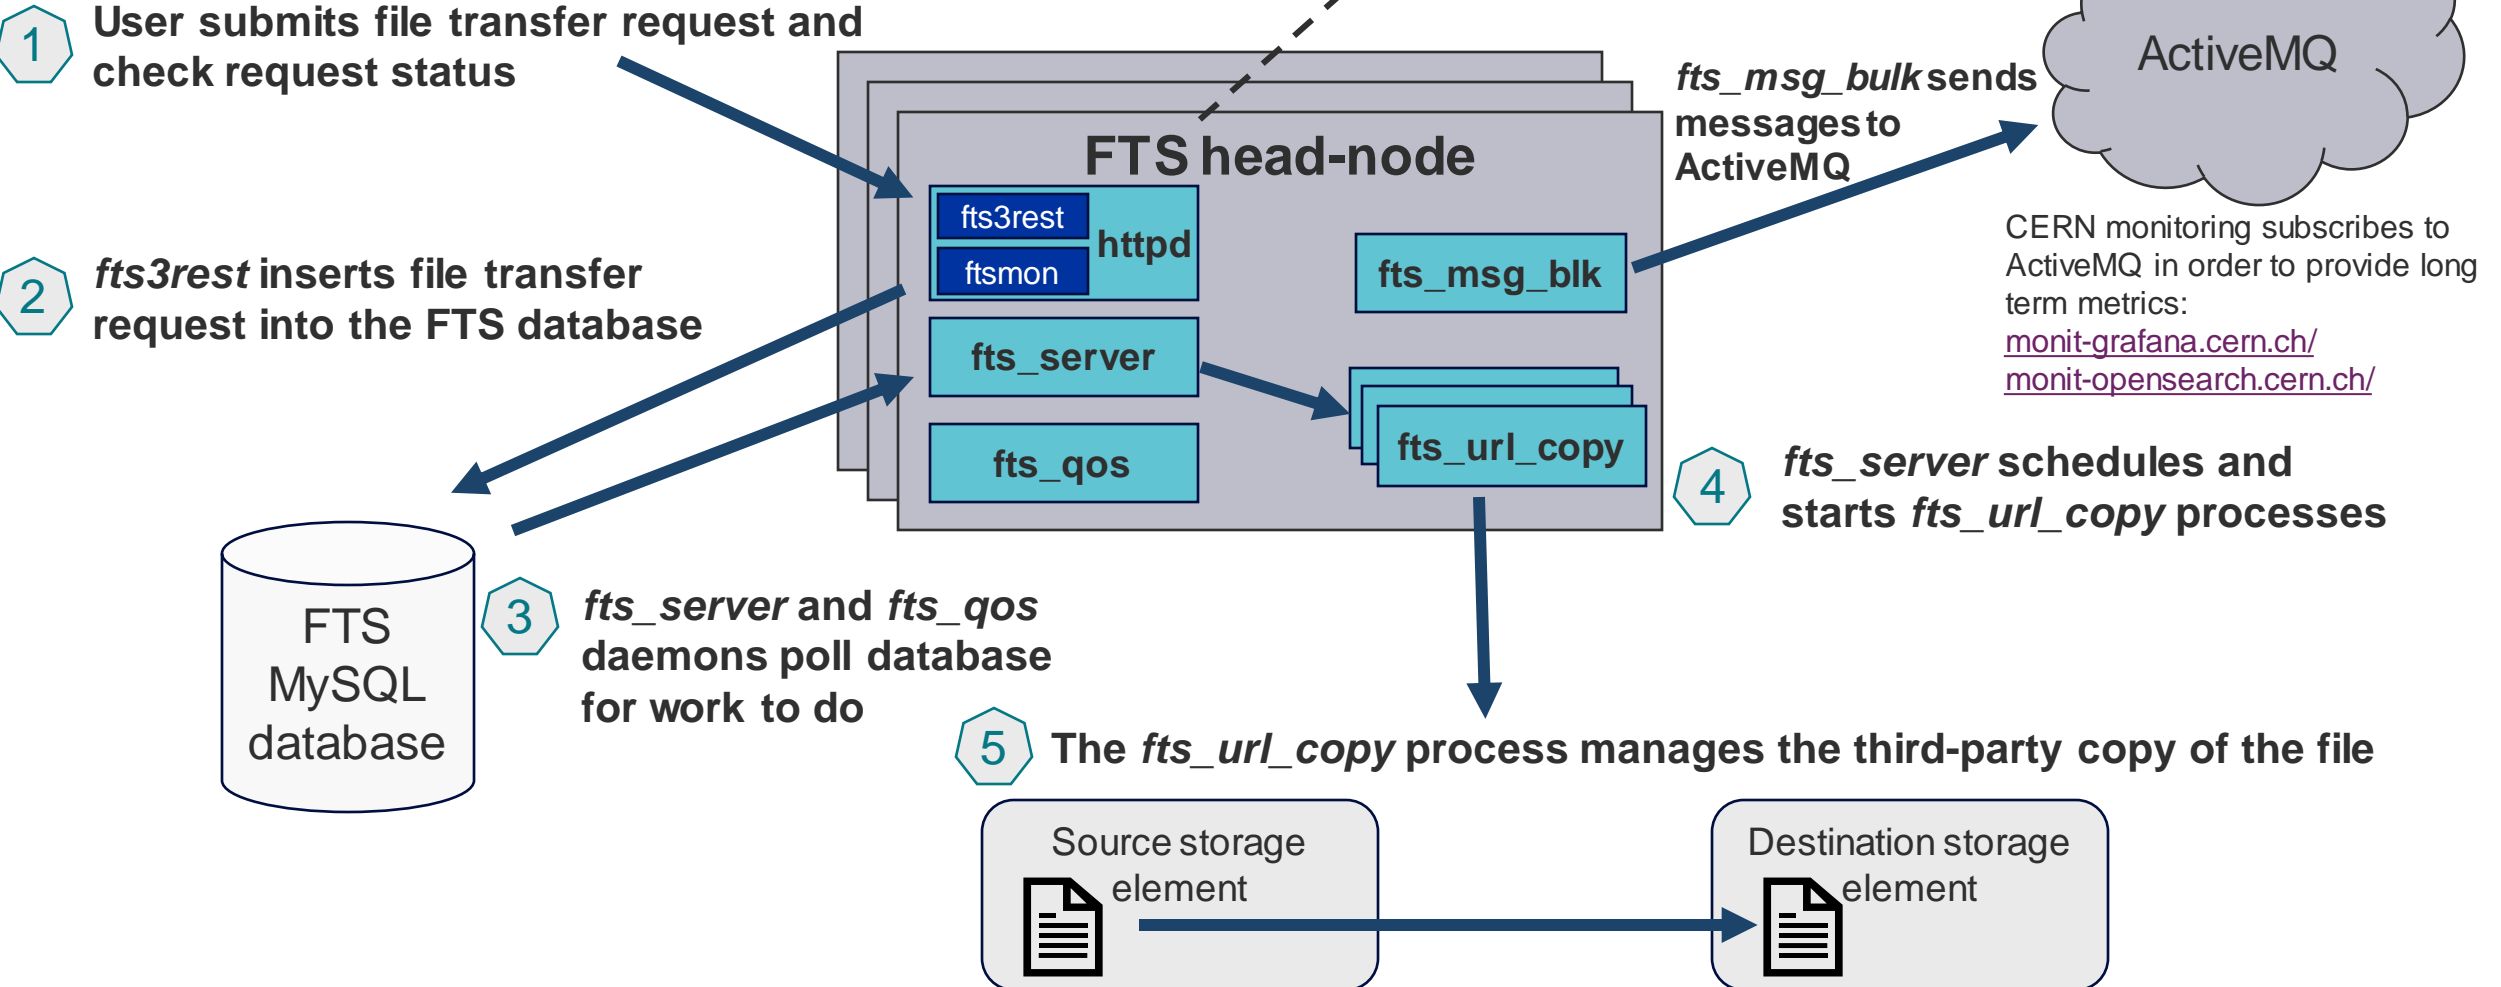
- [FTS](#) (Transfer and Tape daemons)
- [FTS-REST](#) (Transfer submission server)
- [FTS-REST-Client](#) (Python & CLI tools)
- [FTS-Monitoring](#)

Data management clients (DMC):

- [Gfal2](#) (Grid file access library)
- [Gfal2-python](#) (Python bindings)
- [Gfal2-util](#) (CLI tools)
- [Davix](#) (Grid HTTP client)
- [SRM-IFCE](#) (needed by Gfal2 SRM)
- [CGSI-GSOAP](#) (needed by Gfal2 SRM)

How does FTS work?

An FTS instance is made of multiple head-nodes



How does FTS work?

FTS uses Gfal2

- **All FTS ↔ storage interaction is done indirectly via the Gfal2 library**
- **Gfal2 (Grid File Access Library) provides a common top-level file API**
... but supports multiple protocols behind-the-scenes
- **Supported protocols include:**
 - HTTP/Webdav
 - Cloud storage (S3, Swift, GCloud)
 - Xrootd
 - SRM
 - GridFTP
 - Local file
- **GFAL2 can be used standalone via CLI – useful for debugging:**

```
$ gfal-copy https://eospublic.cern.ch:443/<file>  
           https://eosatlas.cern.ch:443/<file>
```

How to use FTS?




Submitting transfers to FTS:

- FTS provides a REST API for transfer submissions and querying its status
- Dedicated CLI clients (`$ fts-rest-transfer-submit`)
- Python 3 bindings (`$ python3 -c 'import fts3; transfer = fts3.new_job(..)'`)
- Direct JSON submission via ``/jobs`` endpoint

```
$ fts-rest-transfer-submit -s https://fts3-pilot.cern.ch:8446/  
    https://eospublic.cern.ch:443/<path> https://eosatlas.cern.ch:443/<path>  
  
$ fts-rest-transfer-status -s https://fts3-pilot.cern.ch:8446/  
    d4e3dc36-f7c2-46f7-8f40-70981d9d539c  
  
$ curl -X POST --cert <cert> --data=submission.json https://fts3-pilot.cern.ch:8446/jobs
```


How to use FTS?

Monitoring transfers progress – FTS Web Monitoring:



Overview

Find all transfers between two storages

Find specific job id

Showing 1 to 50 out of 1600 from the last 1 hour

First Previous 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 ... Next Last

Source	Destination	V0	Submitted	Active	Staging	S.Active	Archiving	Finished	Failed	Cancel	Rate (last 1h)	Thr.
+ davs://webdav.data.net2.mghpcc.	srm://ccsrm.in2p3.fr	atlas	9867	53	-	-	-	21	7	-	75.00 %	104.41 MiB/s
+ davs://bohr3226.tier2.hep.manch	davs://gdav1.physik.uni-mainz.de	atlas	7042	6	-	-	-	23	82	-	21.90 %	0.16 MiB/s
+ davs://ccdavatlas.in2p3.fr	davs://cceos.ihep.ac.cn	atlas	4495	2	-	-	-	-	30	102	0.00 %	-
+ davs://tbit00.nipne.ro	davs://rdr.echo.stfc.ac.uk	atlas	3085	11	-	-	-	81	144	-	36.00 %	129.13 MiB/s
+ davs://webdav.grid.surfsara.nl	davs://se.hpc.utfsm.cl	atlas	3048	8	-	-	-	-	117	41	0.00 %	-
+ davs://dcgftp.usatlas.bnl.gov	davs://tbit00.nipne.ro	atlas	2821	3	-	-	-	8	7	-	53.33 %	7.45 MiB/s
+ davs://eos.grif.fr	davs://atlaswebdav-kit.gridka.de	atlas	2699	2	-	-	-	-	-	-	0.00 %	-
+ davs://tbit00.nipne.ro	davs://atlaswebdav-kit.gridka.de	atlas	2323	26	-	-	-	261	52	-	83.39 %	286.97 MiB/s
+ davs://webdav.data.net2.mghpcc.	srm://storm-fe.cr.cnaf.infn.it	atlas	2017	47	-	-	-	174	16	-	91.58 %	267.66 MiB/s

How to use FTS?

Monitoring transfers progress – FTS Web Monitoring:

Generated at 12:47:01 AM (ftsatlas-011.cern.ch) Overview Jobs Optimizer Error reasons Statistics Configuration Job id v3.12.1

atlas davs://eosatlas.cern.ch davs://rdr.echo.stfc.ac.uk 1 hour Apply Reset More filters

Showing 1 to 50 out of 462

First Previous 1 2 3 4 5 6 7 8 9 10 Next Last

Job id	Submit time	Job state	VO	Source SE	Destination SE	Files	Priority
Fdf99dca-df2f-11ee-96ec-fa163e39e2d9	2024-03-10T22:46:11Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	3	N
8ffc2038-df30-11ee-8d4e-fa163e2cda95	2024-03-10T22:46:41Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	66	3 N
1835c308-df30-11ee-b89f-fa163e017123	2024-03-10T22:46:55Z	ACTIVE	atlas		davs://rdr.echo.stfc.ac.uk	75	3 N
21ede7c2-df30-11ee-ba79-fa163ea7ee69	2024-03-10T22:47:12Z	ACTIVE	atlas		davs://rdr.echo.stfc.ac.uk	50	3 N
22088b0e-df30-11ee-8f30-fa163ec3b00c	2024-03-10T22:47:13Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	48	3 N
33b8e416-df30-11ee-b0a7-fa163e52bca9	2024-03-10T22:47:41Z	ACTIVE	atlas		davs://rdr.echo.stfc.ac.uk	48	3 N
3b33096e-df30-11ee-8411-fa163e2cda95	2024-03-10T22:47:54Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	43	3 N
45742d64-df30-11ee-b0a7-fa163e52bca9	2024-03-10T22:48:11Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	53	3 N
470629a2-df30-11ee-9295-fa163e5ce271	2024-03-10T22:48:14Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	57	3 N
57957598-df30-11ee-9219-fa163e39e2d9	2024-03-10T22:48:42Z	FINISHED	atlas		davs://rdr.echo.stfc.ac.uk	55	3 N
5a88dde-df30-11ee-b0a7-fa163e52bca9	2024-03-10T22:48:51Z	FINISHED	atlas	davs://eosatlas.cern.ch	davs://rdr.echo.stfc.ac.uk	1	3 N

Generated at 12:45:46 AM Overview Jobs Optimizer Error reasons Statistics Configuration Job id v3.12.1

Transfer 'db4307da-df2f-11ee-8c85-fa163e017123' FINISHED

VO: atlas

Delegation ID: bad9d909de813e31
 Submitted time: 2024-03-10T22:45:13Z
 Job finished: 2024-03-10T23:08:18Z
 Priority: 3
 Bring online: -1
 Archive timeout: -1

Received by fts-atlas-003.cern.ch
 Job expires: 2024-03-17T22:45:13Z
 Overwrite flag: Y
 Job type: N
 Cancel flag:
 Pin lifetime: -1
 Target QoS:

Metadata:

```
{"issuer": "rucio", "multi_sources": false, "auth_method": "certificate"}
```

Files transferred	Bytes transferred	Submission time	Start time	Running time	Avg. file throughput	Current job throughput
47 out of 47	7.03 GiB	2024-03-10T22:45:13Z	2024-03-10T22:45:14Z (+1s)	10 s	13.91 MiB/s	-

Showing 1 to 47 out of 47

SUBMITTED DELETE READY STAGING ARCHIVING ACTIVE STARTED CANCELED FAILED 47 FINISHED NOT_USED

First Previous 1 Next Last

File ID	File State	File Size	Throughput	Remaining	Start Time	Finish Time	Staging State	Staging Start	Staging End	Archiving Start	Archiving End
6299063636	FINISHED	49.90 MiB	18.19 MiB/s	-	2024-03-10T22:45:19Z	2024-03-10T22:45:24Z	-	-	-	-	-

davs://eosatlas.cern.ch:443/eos/atlas/atlasdatadisk/rucio/mc23_13p6TeV/93/6d/EVNT.37212728._005868.pool.root.1?copy_mode=pull

davs://rdr.echo.stfc.ac.uk:1094/atlas/datadisk/rucio/mc23_13p6TeV/93/6d/EVNT.37212728._005868.pool.root.1

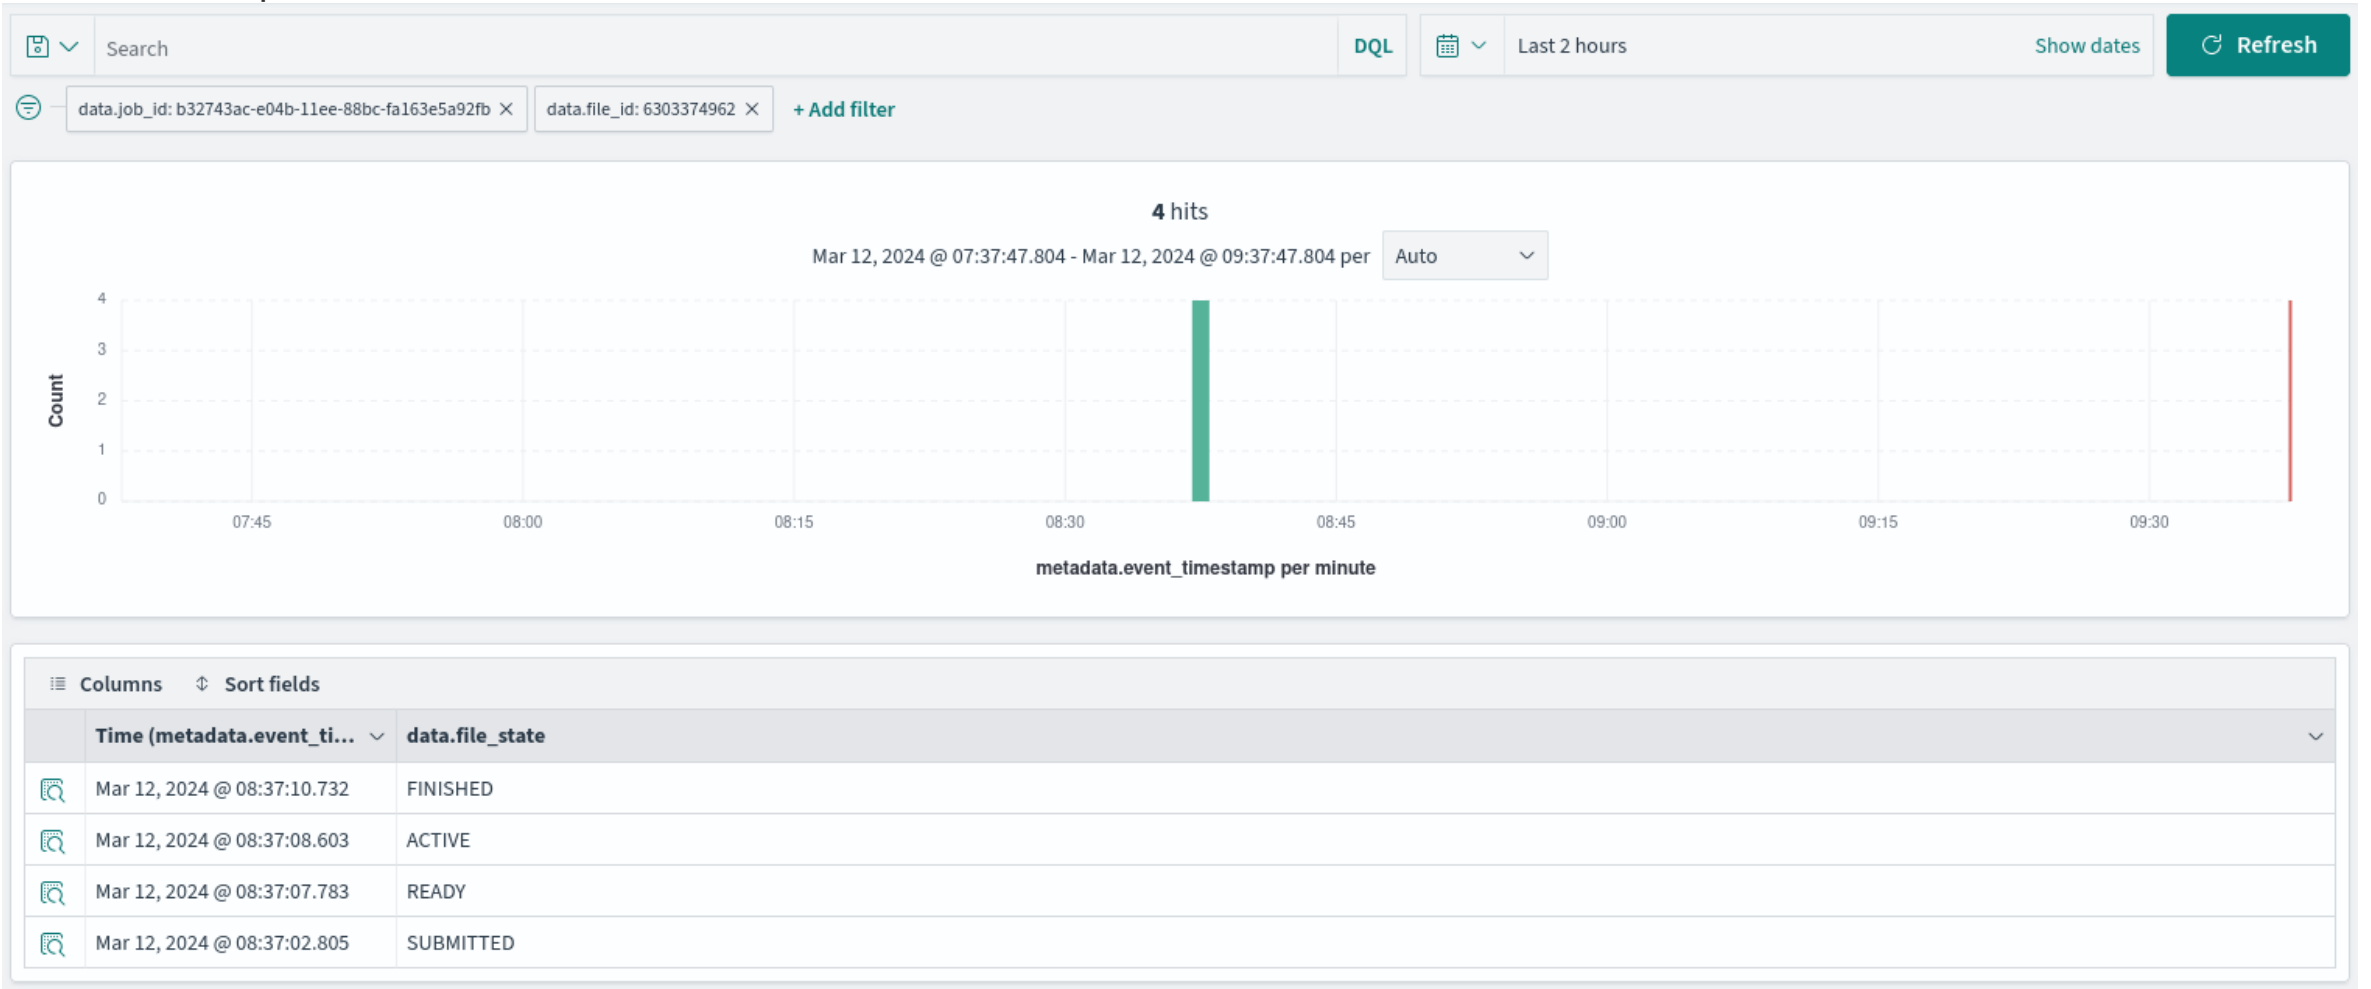
File transfer logs



How to use FTS?

Monitoring transfers progress – Monitoring messages & ActiveMQ

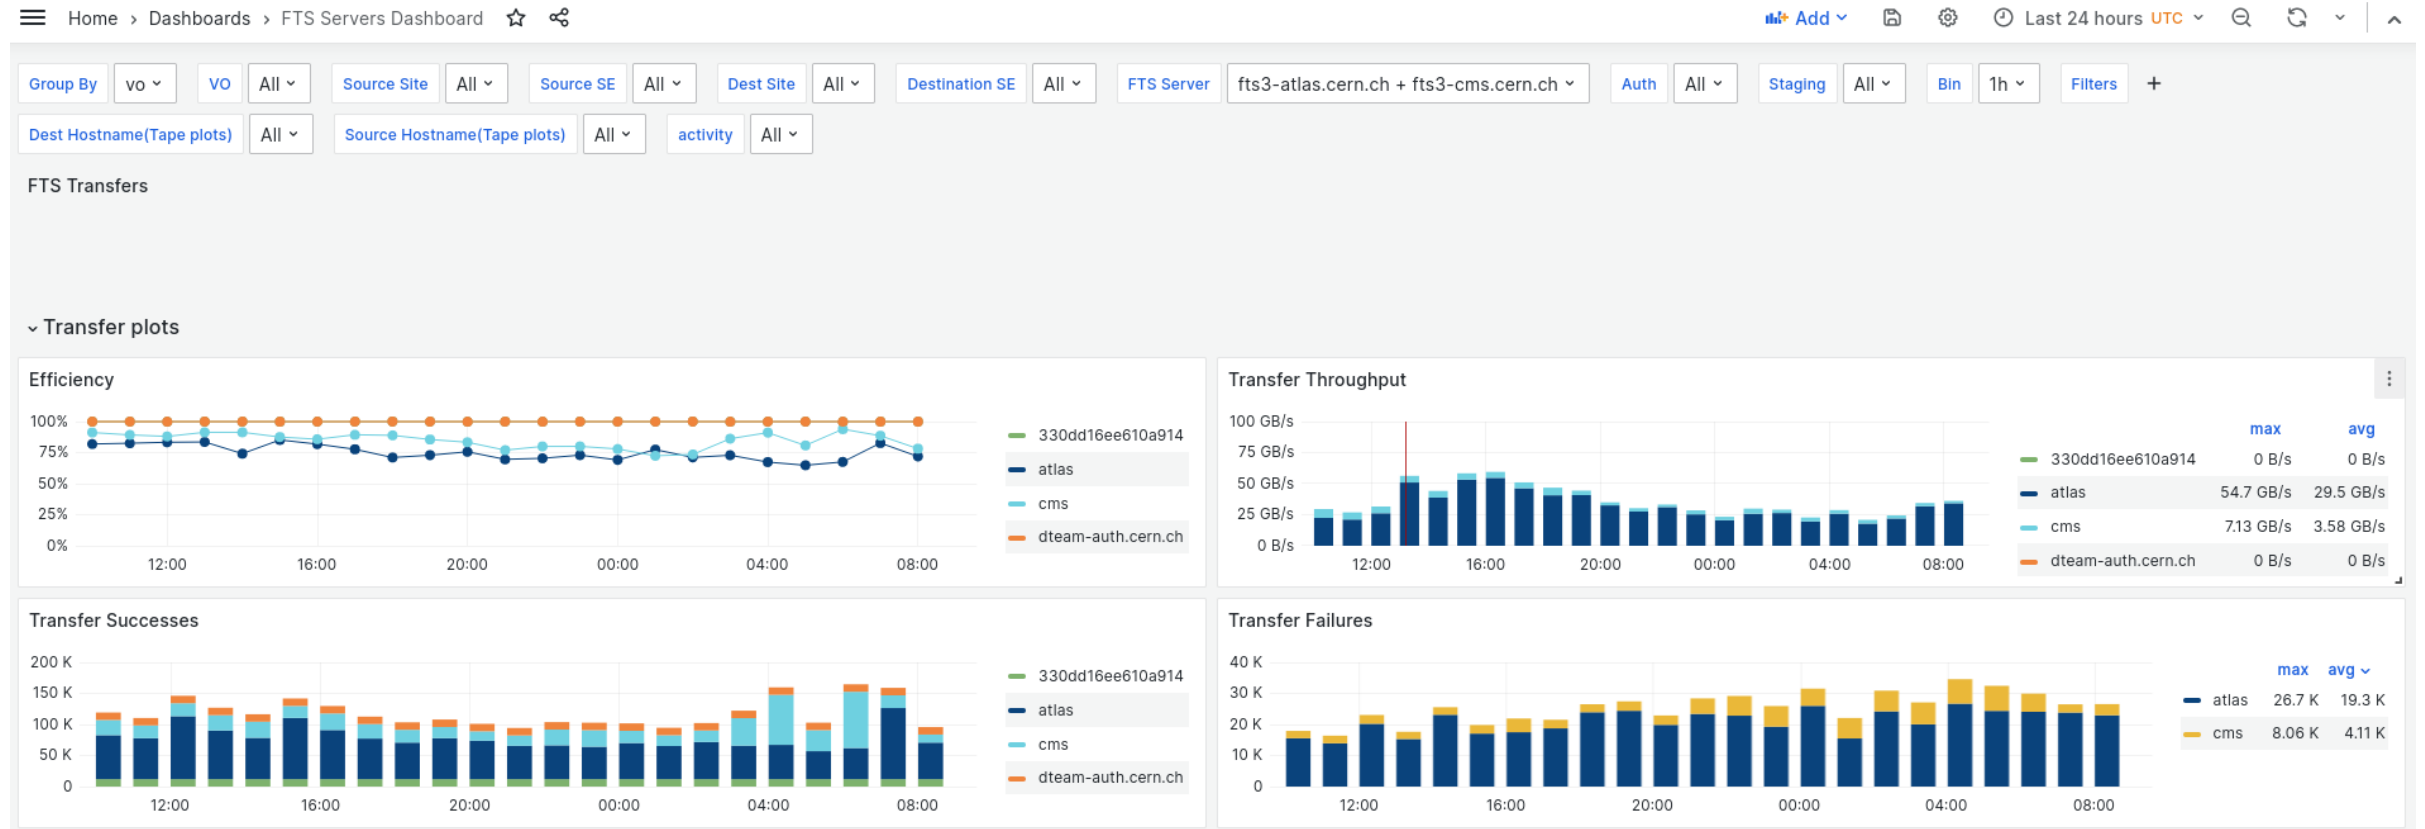
- From monit-opensearch.cern.ch/



How to use FTS?

Monitoring transfers progress – Grafana monitoring @ CERN

- From monit-grafana.cern.ch



Who uses FTS?

1. Scientific experiments and communities



2. Scientific collaborations




3. Scientific frameworks



Data moved by CERN FTS (2023)

- 823 Million transfers
- ~1.4 Exabytes of data

Volume Transferred / Number of Transfers



Where to find us?

- 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 / dmc-devel@cern.ch
- **Announcements:** fts3-steering@cern.ch
- **FTS Service Managers:** fts-ops-forum@cern.ch
- **User support:** fts-support@cern.ch / Service Now (CERN) / GGUS (experiments, WLCG)



cern.ch/fts3-docs/



cern.ch/dmc-docs/



cern.ch/fts

XRootD and FTS Workshop @ STFC UK Sep 9 – 13, 2024

XRootD and FTS Workshop @ STFC UK

Sep 9 – 13, 2024
The Cosener's House
Europe/London timezone

Overview

Call for Abstracts

The XRootD and FTS workshop brings together the XRootD and FTS developers and people from Academia, Research, and Industry to discuss current and future data access activities as related to the XRootD framework and to the FTS project.

Presentations focus on achievements, shortcomings, requirements and future plans for both XRootD and FTS projects.

This year we decided to bundle together these two partially overlapping communities. during the workshop there will be additional sessions focusing on review of backend filesystems, on experiments' plans for the LHC Run3 data analysis and remote data access, and on support for efficient remote data access for non-HEP VOs.

<https://indico.cern.ch/event/1386888/>

