



Getting FTS@CERN ready for LHC Run3

João Pedro Lopes on behalf of the FTS team

28/04/2022

Outline

- **Introduction**
- **Python3 migration**
- **New database deployment strategy**
- **Service-health monitoring**
- **2022 data challenges overview**
- **Improvements now & future**



Open source software for large scale queuing and reliable execution of file transfers.



Simple & Intuitive

Easy user interfaces for submitting transfers:
Python Clients, WebFTS and Web Monitoring



Robust

Checksums and retries are provided per transfer



Flexible

Multi-protocol support.
Multi-Authentication.
Support for tape storage



Adaptative

Priorities/Activities for transfer classification
Runtime optimization to maximize network and storage resources

FTS privacy policy notice: https://cern.service-now.com/service-portal?id=privacy_policy&se=file-transfer¬ice=fts



FTS
File Transfer Service

Developed at CERN:

- **Mihai Patrascoiu (Project Leader)**
- **Steven Murray (Service Manager)**
- **João Pedro Lopes (C++/Python Developer)**
- **João Afonso (Service Health Monitoring)**

Active contributors

- **Edward Dambik (C++/Python Developer) - Indiana University Bloomington, USA**
- **Eraldo Silva Junior (Python Developer) - CBPF, Brazil**

8 WLCG Instances

- BNL, CERN (4), FNAL, RAL, MIT

16 non-WLCG Instances

- CERN (DAQ, Public), RAL, KEK(2), Imperial (also used by CMS), PIC, MWT2, CESNET (WebFTS + RAuth), JINR, CNAF, SARA, SLAC, IHEP, Fermilab (containers), FENIX Research Infrastructure (Human Brain Project)

>30 Virtual Organizations (VO's)

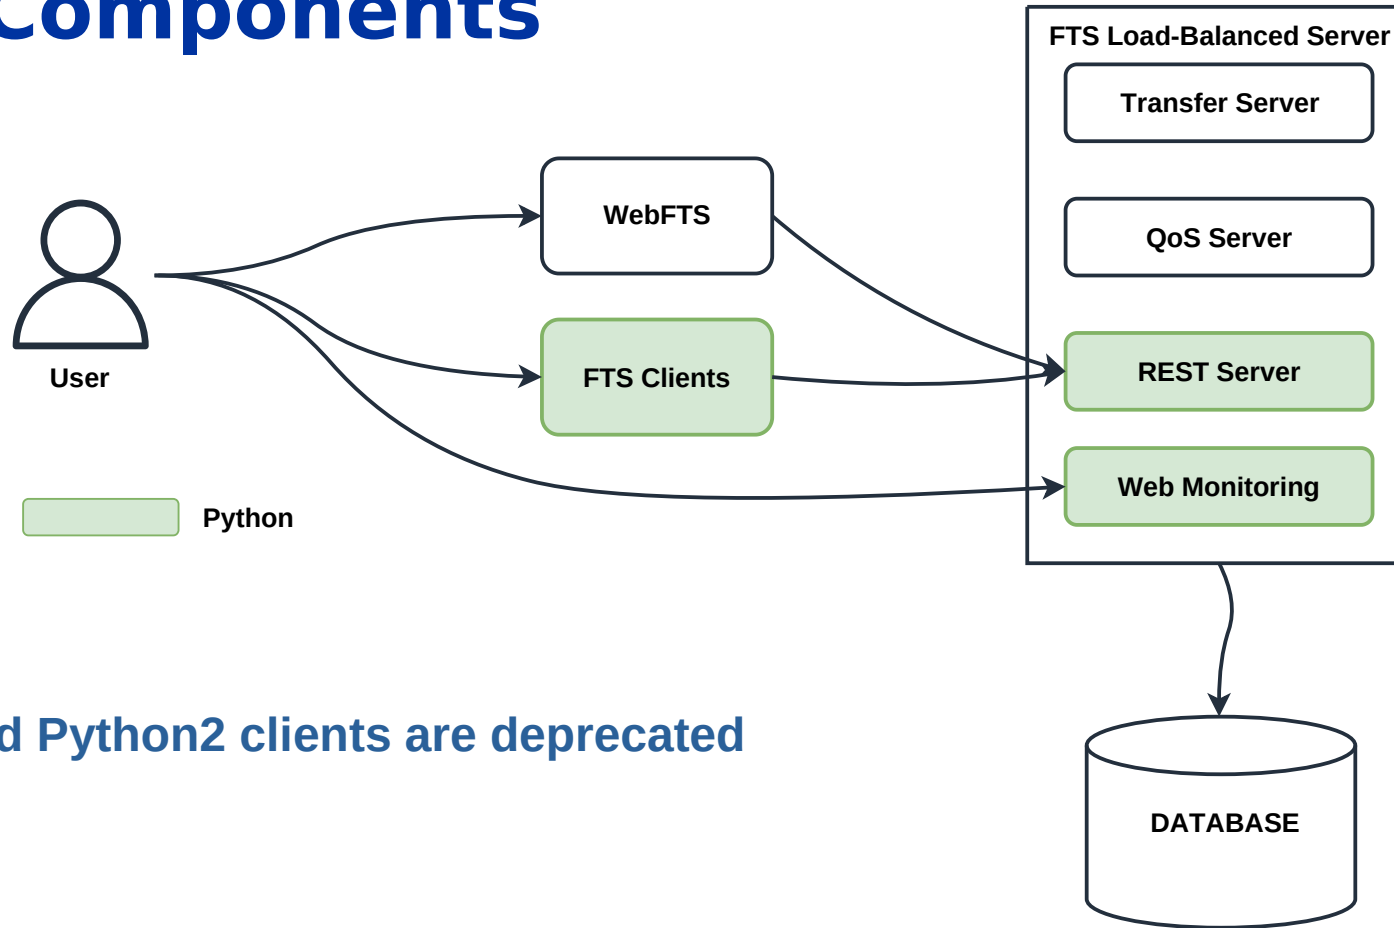
- ATLAS, CMS, LHCb, AMS, NA62, Compass, ILC, Magic, Belle, Mice, Xenon, Snoplus, Gridpp, Dune, LZ, Solidexperiment.org, SKA, Ligo, Icecube, Elixir, NP02(part of Dune), CAST, ESCAPE, Eiscat.se, Virgo, Pierre Auger Observatory, BES III, JUNO, CEPC, FENIX-RI, CTA, T2K, Project8, ICARUS, FASER, Folding@Home

CERN instances:

- Data transferred in 2021: >1.0 EBs
- Files transferred in 2021: ~1.15B files



FTS Components

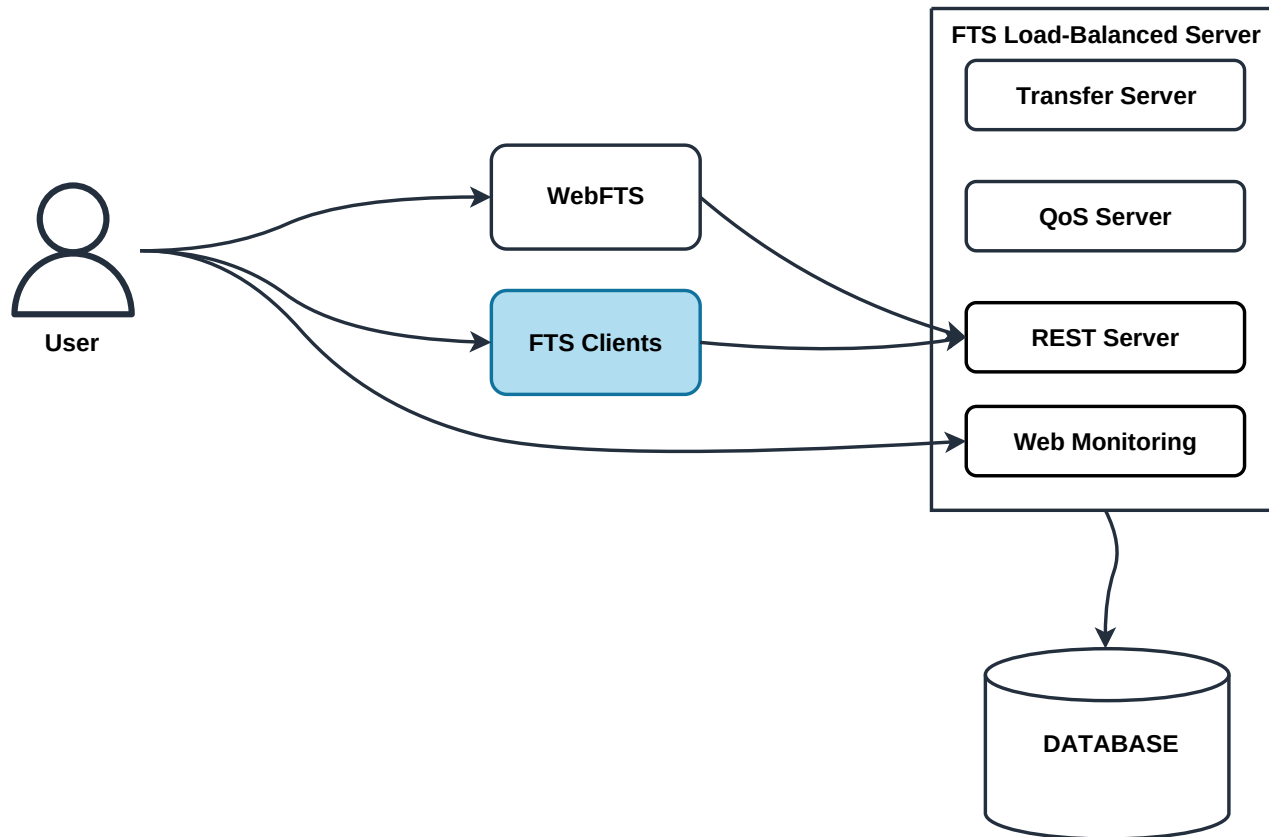


- C++ and Python2 clients are deprecated

Outline

- Introduction
- **Python3 migration**
- New database deployment strategy
- Service-health monitoring
- 2022 data challenges overview
- Improvements now & future

Python3 Migration - FTS CLIENTS



Python3 Migration - FTS CLIENTS

- Released in the FTS production repository [1]
\$ yum install fts-rest-client
- Released in PyPI [2]
\$ pip install fts3
- Already in use by Data Transfer Orchestrator (DTO) and CMS CRAB Asynchronous Stage-Out
- No functional changes for the end-users
- Python2 clients are deprecated

[1] https://fts-repo.web.cern.ch/el7/x86_64/

[2] <https://pypi.org/project/fts3/>

Python3 Migration - FTS CLIENTS

- **Examples of command line utilities:**

```
$ fts-rest-whoami -s https://fts3-pilot.cern.ch:8446
```

```
$ fts-rest-delegate -s https://fts3-pilot.cern.ch:8446
```

```
$ fts-rest-transfer-submit -s https://fts3-pilot.cern.ch:8446 <source> <destination>
```

```
$ fts-rest-transfer-status -s https://fts3-pilot.cern.ch:8446 <job-id>
```

```
$ fts-rest-transfer-cancel -s https://fts3-pilot.cern.ch:8446 <job-id>
```

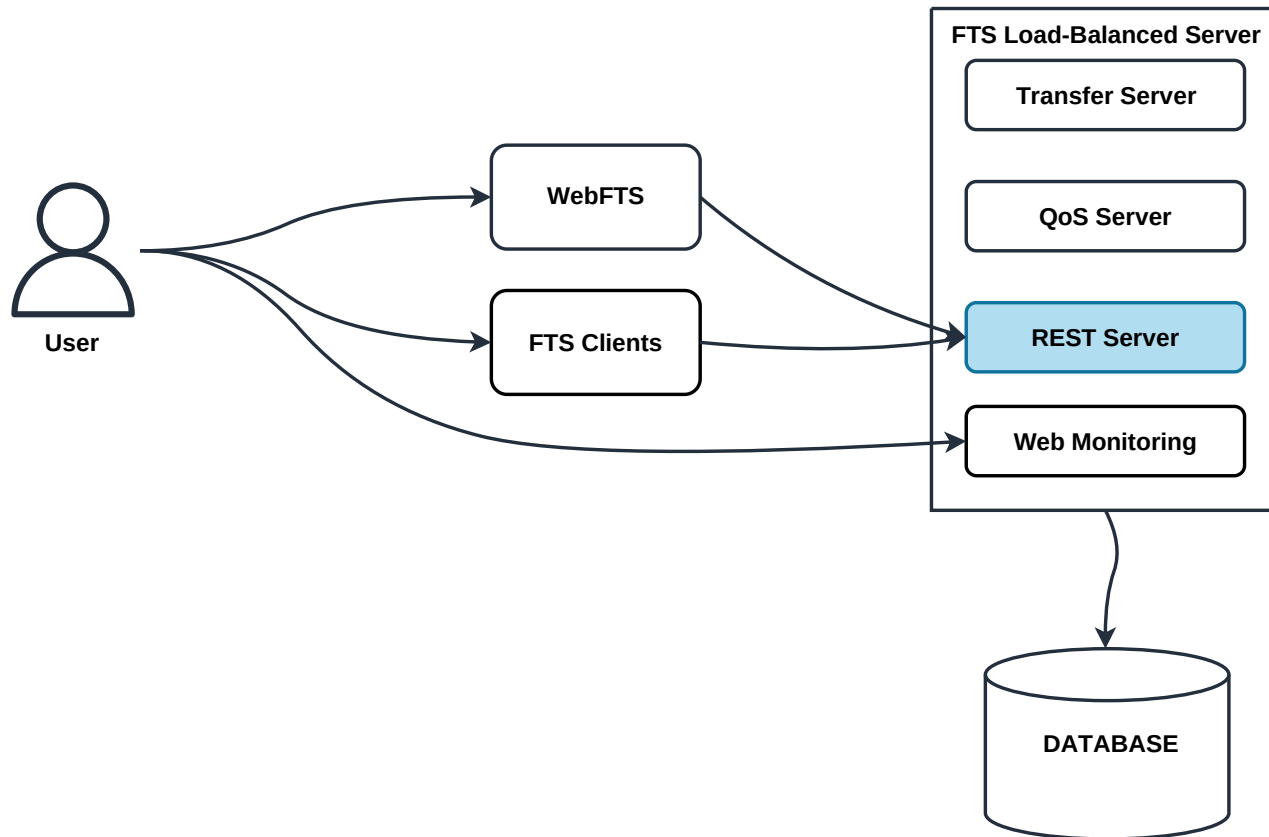
```
$ fts-rest-transfer-list -s https://fts3-pilot.cern.ch -o <vo>
```

Python3 Migration - FTS CLIENTS

- Submitting a transfer job to FTS

```
1 import fts3.rest.client.easy as fts3
2
3 context = fts3.Context(endpoint="https://fts3-pilot.cern.ch:8446")
4
5 transfer1 = fts3.new_transfer(
6     source='root://eospublic.cern.ch/eos/opstest/dteam/file.1mb',
7     destination='root://eospps.cern.ch/eos/opstest/dteam/file.1mb',
8 )
9
10 transfer2 = fts3.new_transfer(
11     source='root://eospublic.cern.ch/eos/opstest/dteam/file.100mb',
12     destination='root://eospps.cern.ch/eos/opstest/dteam/file.100b',
13 )
14
15 transfers = [transfer1, transfer2]
16
17 job = fts3.new_job(transfers)
18
19 job_id = fts3.submit(context, job)
20
```

Python3 Migration - REST Server

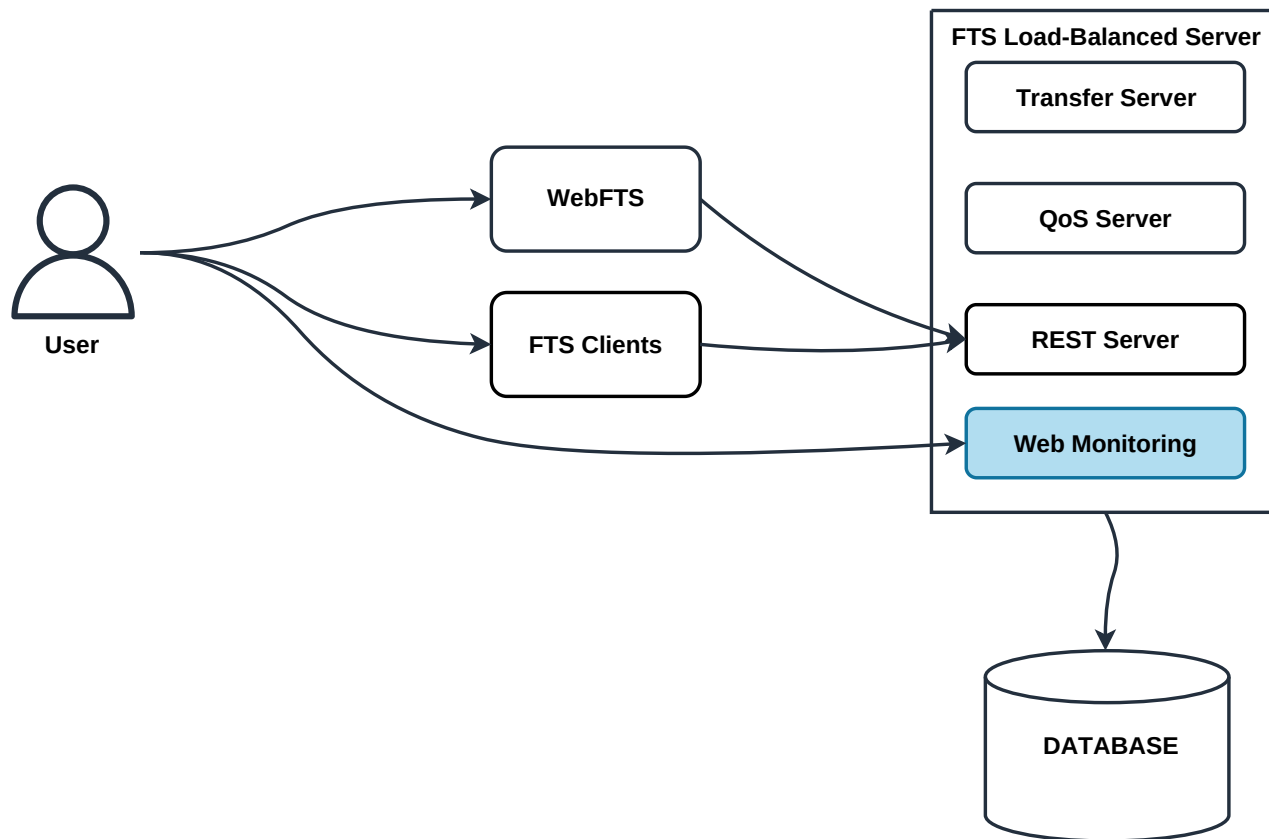


Python3 Migration - REST Server

- Official release before LHC Run 3
- Already deployed at CERN since February
- Move code base from Pylons to Flask framework
- Upgrade SQLAlchemy to v0.9.8
- Transparent for clients






Python3 Migration - Web Monitoring



Python3 Migration - Web Monitoring

- View the state of individual file transfers.
- Drill down to file transfer logs.
- View file transfer efficiency.





Generated at 3:15:07 PM (fts-cms-06.cern.ch)  Overview ▾ Jobs ▾ Optimizer ▾ Error reasons ▾ Statistics ▾ Configuration ▾

 - All - ▾ →  1 hour ▾

Overview

Showing 1 to 50 out of 1214 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://cmsdcadisk.fr	davs://se01.indiacms.cms		5818	94	-	-	-	151	5	-	96.79 %	64.70 MB/s 
+ davs://xrootd.hepgrid.org	davs://dcache-cms-webdav-wan.desy.de	cms	3544	200	-	-	-	86	3	6	96.63 %	86.73 MB/s 
+ davs://eoscms.cern.ch	davs://se01.indiacms.cms		2832	2	-	-	-	444	1	146	99.78 %	23.44 MB/s 
+ davs://eoscms.cern.ch	davs://xrootd.phy.br.cern.ch	cms	2089	13	-	-	-	-	31	-	0.00 %	27.30 MB/s 

Python3 Migration - Web Monitoring

- Official release before LHC Run 3
- Already deployed at CERN since February
- Upgrade Django to v1.11.27
- Transparent to end-users
- Will be released together with FTS REST Server



MySQL 5.7 → MySQL 8.0

Thanks to Python3 migration FTS supports MySQL8

- Python2 + SQLAlchemy prevented moving to MySQL 8.0 sooner
- New SQLAlchemy version allows upgrade to MySQL 8.0

All CERN instances have been upgraded to MySQL 8.0

- Better performance
- Online database schema updates
- Faster DB schema upgrades with no downtime required

Outline

- Introduction
- Python3 migration
- **New database deployment strategy**
- Service-health monitoring
- 2022 data challenges overview
- Improvements now & future

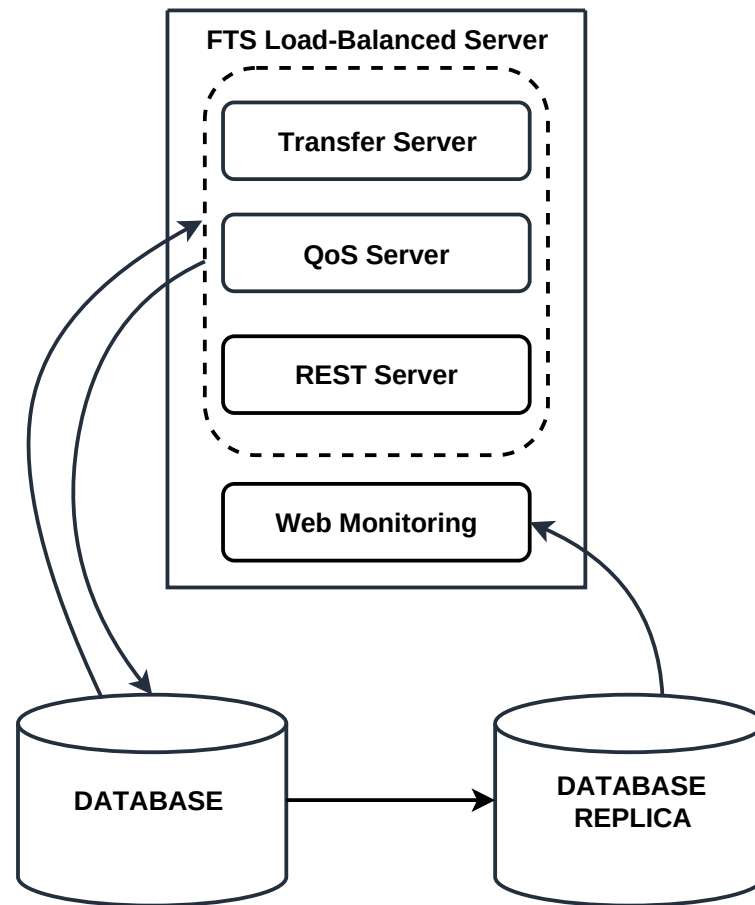
New database deployment

Service critical operations goes to the main database

Non-critical operations goes to the replica

- Replication delays have no impact on critical operations

Different database access patterns benefit from separated caching




Outline

- Introduction
- Python3 migration
- New database deployment strategy
- **Service-health monitoring**
- 2022 data challenges overview
- Improvements now & future

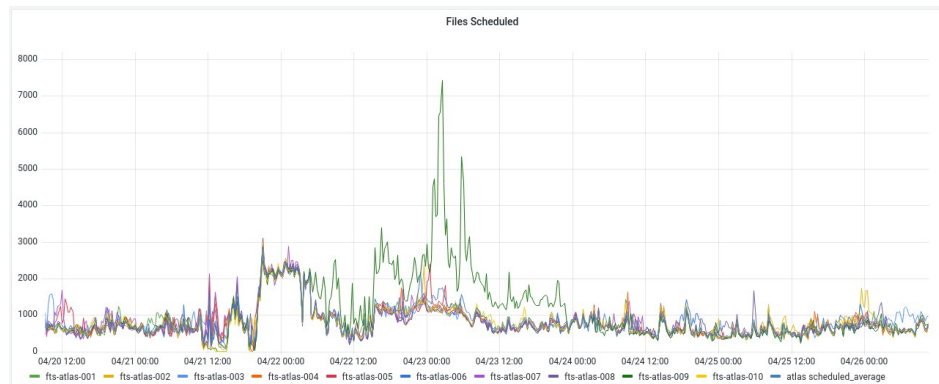

Service Health Monitoring

- FTS v3.11.2 introduced profiling logs to monitor service KPIs
- We have a Fluentbit setup which feeds data into Graphite
- Grafana dashboards are used to display metrics.


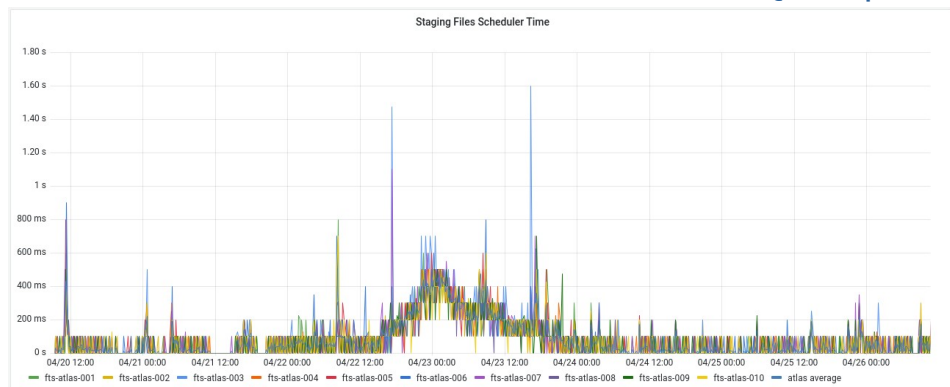


Service Health Monitoring

Transfer scheduling monitoring



QoS operations monitoring



Outline

- Introduction
- Python3 migration
- New database deployment strategy
- Service-health monitoring
- **2022 data challenges overview**
- Improvements now & future

Data Challenges Overview

FTS participated in several Tier-0 and Tier-1 commissioning tests

CERN FTS performed very well over the last year

- No major incidents reported


FTS team is confident in the service performance

- Continuously improving in collaboration with the experiments

March 2022 Data Challenges Overview

Volume: ATLAS → 42 PB; CMS → 14 PB; LHCb → 7 PB; Total → 63 PB

Transfers → 47.8 Milion




March 2022 Data Challenge

Total volume per Tier:

Tier-0 → 11.8 PB; Tier-1 → 26.6 PB; Tier-2 → 18.5 PB; Total → 63 PB

Volume Transferred / Number of Transfers



Outline

- Introduction
- Python3 migration
- New database deployment strategy
- Service-health monitoring
- 2022 data challenges overview
- **Improvements now & future**

FTS Improvements

Dramatic performance improvement due to optimised “stage-in” SQL query

Server-side configuration development to ease transition to SRM+HTTP-TPC

- New per-link configuration to choose the preferred TURL

Sanity checks improved to correct anomalies in multihop transfers

- Problem identified by the FASER experiment

Destination file report for archiving transfers

- Initially requested by CMS, allows users to decide if a file should be overwritten

Integrated with CS3 family of protocols

- e.g.: CERNBox's Reva and other compatible sync&share services

Improved transfer scheduling per activity

- Scheduling had an implicit alphabetical bias with saturated links activity

Integrated support for EGI tokens

Future Improvements

- **Support for WLCG TAPE REST API**
 - Working together with EOS and CTA teams at CERN
- **Monitoring IPv6 traffic**
- **Continuously improve the Service Health monitoring**
- **Further support for token authentication**
- **Consolidate and improve service performance**

Communication Channels

Announcements: fts3-steering@cern.ch

Support: fts-support@cern.ch

Development: fts-devel@cern.ch

For Service Managers: fts-ops-forum@cern.ch

- Please subscribe! [1]

FTS Workshop

- Fall 2022

[1] <https://e-groups.cern.ch/e-groups/Egroup.do?egroupId=10451656&tab=3>

Summary

- We have successfully migrated to Python3
- MySQL8 support will help DB service operations
- Improved our internal service monitoring
- Data challenges confirm FTS readiness
- Looking forward to continuous collaboration with experiments

Thank you!

