


FTS Perspective on Tokens @ WLCG OTF #1

Mihai Patrascoiu (*CERN-IT Storage*)



FTS – At the heart of all things token

- Diverse ecosystem of experiment frameworks and technology providers (IAM, storages)
- FTS common element to link stakeholder requirements with provided technical solutions



FTS Requirements for Tokens

Simple: Don't want to look inside tokens

- Only "iss", "exp" and "offline_access" claims
- FTS services multiple communities, not only WLCG

Scalable: FTS performance should not be tied to the performance of an external system

- ~30 Hz average submission rate, spikes of 150-200 Hz

Robust: Transfers continue even in the case of IAM downtime

Pragmatic: Token management can grow complex

Complete: Handles both **disk** and **tape** transfers

Token Approach: in practice

Decision

- 1 Token to submit to FTS
- **2 tokens for each DISK transfer**
(source & destination)

Advantages

- FTS system is “agnostic” to transfer token contents
- Experiments can decide
(and change) granularity

Experiment
fine-tuning

Many tokens
Fine granularity
Higher-security

Scalability risks



Token lifetime?

Fewer tokens
Coarse granularity

Policy and Security risks

Disk-to-disk Transfers: Status

- Tokens successfully deployed for DataChallenge'24, **in-use since**
- Token support deployed on all production FTS instances (available since v3.13.0)
 - Background daemon to continuously refresh tokens
 - **“Just-in-time” refreshing**: the last development for disk-to-disk token transfers
- Different approaches within the experiments (unfortunately):
 - CMS performing token tests using the standard token-refresh lifecycle
 - ATLAS performing token tests using long-lived per-file non-refreshed access tokens
 - LHCb plans to use per-file tokens
 - What token approach will DUNE have? (expected in 2025)

Tape Transfers: Status?

- **Tokens + Tape currently not defined**
- **Review WLCG Tape REST API to support tokens**
- Several FTS-specific tape + token concerns have to be addressed as well (see later slides)




WLCG Tape REST API: Status?

- Current **bulk** Tape REST API does not support tokens
- Tokens + Tape discussed for first time @ FTS & XRootd Workshop 2024
- Follow-up (first) technical discussion @ CHEP 2024 (FTS, EOS/CTA, dCache, StoRM teams present)
 - ... converging on a solution will require compromises between involved actors
 - More time needed to refine a proposal before anything can be presented
 - Warrants its own separate discussion (February 2025?)

FTS Considerations: Tokens + Tape

- Tape interaction presents 3 distinct actions:



- FTS creates the staging batches for the Tape SE (experiments submit staging per-file)
 - > **How will this look with tokens/scopes involved?**
 - > (see no other way than using a token for each file)

- Tape SEs want a queue as large as possible
 - Staging campaigns of up to 1M transfers, lasting weeks


FTS needs a valid access token for: stage, poll, abort, release, transfer

Issue: Bottleneck w/ IAM token interaction (*token-exchange*) measured by FTS in the 7-12 Hz range

> **Can we do bulk tape fast enough?**

FTS Considerations: Tokens + Tape

- Tape interaction presents 3 distinct actions:



- Need a valid access token for staging, polling, abort, release and transfer
> **Is it worth refreshing tokens whilst polling for days?**
- Tape actions are expensive. How do we protect tape from fragile token ecosystem?
(e.g.: cannot refresh token due to IAM downtime → failed tape operation)

Token Reflections

- Large token usage during DC'24, shy adoption afterwards
- A fragile system of many levers: Rucio, FTS, IAM
 - Can we ensure transfers continue even in the case of IAM downtime?
- Experience shows token workflows are complex
 - Simple approaches (token-per-file) bring scalability concerns
 - Anything else brings complicated workflows
 - Refresh workflows significantly more complex than using only ATs
 - **Is refreshing tokens the right mechanism for FTS?**
- Likely, FTS will have to support two token solutions:
 - Targeted solution for large frameworks, such as Rucio & DIRAC
 - Full token lifecycle management for smaller communities or clients

Token ~~Reflections~~ Wishes

- Token direction (still) looks unclear
 - **Need a uniting driving force at the WLCG level**
(many stakeholders: Storages, FTS, IAM, experiments)
 - **Would want to see more alignment between experiments** (token granularity)
 - **Who is responsible for enforcing file ownership? Storage or central IAM?**
- **Is WLCG JWT Token profile still respected in practice?** (many workarounds all over the Grid)
- **2025 should finalize the disk-to-disk part**
- **2025 should bring Tokens + Tape prototype**
- **2025 should bring larger token adoption for WLCG transfers**