

Tape REST API:

Status and Plans

Mihai Patrascoiu

WLCG Workshop (7-9th November 2022)
Lancaster, United Kingdom



- I. Current Status
- II. Testing Phase
- III. Evolution

I. Recap on why

- Provide a modern, uniform way of managing tape data movements across **all** tape storage implementations
- Unify tape and transfer management under the **same** protocol



- Replace SRM for tape data management
- End up with something simpler

XKCD [#927](#)

I. Recap on how

- New DOMA taskforce: wlcg-tape-rest-api-discussions@cern.ch
- Storage (EOSCTA, dCache, StoRM) + FTS/Gfal2 developers involved
- Outcome: Tape REST API Reference document [v1, May 2022] ([link](#))

→ Specification agreed by all storage providers + FTS

→ Each party to develop support for the new API

I. Key points of API

- Introduces a discovery mechanism: `.well-known/wlcg-tape-rest-api`
 - A lot of potential to be explored in future API versions
- Allows sending staging metadata during staging request
 - Staging metadata is passed per-file
- Client sees uniform interaction, despite target storage being batch-oriented or file-oriented
- Tape REST API can benefit from existing effort on token transition

I. Current implementation status

- Clients
 - Gfal2 v2.21.0: released, full support
 - FTS v3.12.2: RC release, full support
- Storage providers
 - EOSCTA: version available for testing
 - dCache: version available for testing
 - StoRM: under development (part of large refactoring)

II. Testing phase

- DOMA-BDT involved in organizing the Tape REST API testbed
 - Situation diligently tracked under [TapeRestAPI Wiki](#)
- Testing may evolve into compliance-like scenario (e.g.: daily execution)
- Gfal2-based testing available: [fts/scripts/tape-rest-api](#)
 - if storage passes, it supports the subset accessed by FTS operations

II. Testing phase (cont'd)

- Integration and at-scale tests started mid-October
 - Lead by FTS foremost to validate TapeRestAPI RC release (one-off)
 - Testing between FTS (pre-production) + CTA & dCache (test endpoints)
- Integration tests goal:
 - Validate correct behavior w.r.t. staging metadata
- At-scale tests: 1M files queued for archiving, staging and file release (unpin)
 - Goal: mimic production-like usage

II. Testing phase – preliminary results

- Functionality and at-scale testing passed
- ...but there are lessons to be learned:
 - Error messages and error propagation to be improved
 - > E.g.: *Failed to query .well-known endpoint: expected 200, received 500*
 - > E.g.: Many times, no clue on what “HTTP 500” means
 - Stalling mechanisms (XRootd + EOS) interpreted as failures in HTTP
 - HTTP is not graceful when storage is under heavy load
 - > Many timeouts / lost connections
 - > Cryptic error codes

II. Testing phase – next step

- Integrate Experiment frameworks (Rucio, DIRAC) into Tape REST API testing
- Goals:
 - Ensure Experiment framework ↔ FTS integration works as expected
 - Ensure information (e.g.: staging metadata) gets propagated to MONIT
- Timeline:
 - As soon as possible, aiming for mid-November

III. Evolution – of the specification

- **Items to Review** section in the Reference Document
 - Collected by involved parties during testing phase
 - Implementations have taken unanimous approach on these items
- Review items will be formalized into a new minor version (e.g.: v1.1, v1.2).
Expected to happen only in 2023
- Production usage will tell what other revisions are needed
- Current status: keep it simple, focus on deployment

III. Evolution – of the deployment

- DOMA-General update on Tape REST API (Brian Bockelman, 29th of June 2022)

Crystal ball predictions:

- **Second half of 2022 will be for testbed instances and software verification.**
- Leading-edge T1s will have production instances in 2023 available for tests.
- Site-by-site migration during 2023; ask all sites to finish by winter shutdown 2023/2024?
 - Is there need for more aggressive timelines?

- CERN-CTA expected to run TapeRestAPI in production by end-of-year 2022
- RAL-CTA expected to run TapeRestAPI in production by February 2023
- What follows after?

Summary

- Tape REST API: simplified, uniform data management solution for tape storages
- Same protocol as data transfers:
 - Simplified experiment workflows, benefit from existing effort on token transition
- At-scale tests show cracks in how HTTP portrays being under heavy load
 - Stalling mechanism lost; Many timeout errors / lost connection ; Cryptic error codes
 - This topic warrants a discussion in itself !
- Production: FTS released autumn 2022 / CTA@CERN end-of-year 2022
CTA@RAL February 2023 (LHCb staging)
- Need a timeline to advance Tape REST API across other WLCG sites

Thank you!

- ✉ wlcg-tape-rest-api-discussions@cern.ch
- 📄 [Tape REST API Reference Document](#)
- 🌐 [Tape REST API Wiki](#)