

DOMA-BDT

2025 Status Report

Mihai Patrascoiu (CERN - IT Storage) &
Petr Vokac (Czech Technical University in Prague)

DOMA-BDT: *The Road So Far*

- Takes birth January 2022 ([link](#))
- Continuation of the DOMA-TPC Working Group
- Steered by B. Bockelman, M. Arsuaga-Rios, P. Vokac

- **Topics of interest**
 - Archive Management (SRM → HTTPs)
 - Data Access w/ Tokens (Tokens Smoke Tests)
 - Network Utilisation (Packet Marking)
 - WLCG DataChallenge'24

- **Successful contributions over the years, achieving proposed objectives**

DOMA-BDT: *The Road So Far*

- **Discussion moved towards recurrent topics**
 - Discussion place for token-based transfers status
 - Discussion place for Tape REST API status
 - Discussion place for Packet Marking status
 - **New:** Improving WebDav transfer error messages
 - AOB
- **Activity slowed down after DC'24 (last meeting on 17th April 2024)**

DOMA-BDT: Reawakening

- **New coordination by Mihai Patrascoiu and Petr Vokac**
(warm salute to the first team of steerers!)
- **First meeting held on 26th February 2025**

| **Monthly meeting**, every 3rd Wednesday of the month (@16:00)

⇒ **Clear need for continued monitoring and discussion on data access and transfers, as the technologies evolve**

- **Communication: [wlcg-doma-bdt @ cern.ch](mailto:wlcg-doma-bdt@cern.ch)**
- **Have a topic to present? Get in touch with us!**

DOMA-BDT: Mandate

The DOMA Bulk Data Transfers Working Group proposes to serve as the discussion forum for data access and transfer technologies.

As part of its duties, it will guide the discussion for the design of new and/or evolving solutions, as well as follow their implementation and deployment across the WLCG.

DOMA-BDT: Role in the Ecosystem

- **DOMA-BDT continues to serve as a technical discussion forum**
- **Targets topics of interest in data management, spanning the full technology stack:**
 - > **Storage, Compute, AuthNZ, FTS, Experiment frameworks**
- **Targets activities of interest in the WLCG community:**
 - > **DataChallenge'27, mini-DataChallenges**
- **Technical progress follow-up & conformity compliance (i.e.: HTTP-TPC evolution, Tape + Tokens “smoke tests”)**
- **Reports to DOMA General & WLCG OTF**

DOMA-BDT “Renewed” Topics of Interest

Immediate

- Publish DOMA-BDT Mandate document
- Redo DOMA-BDT TWiki / Documentation ([current link](#))
- Revive “WLCG JWT Compliance test” routine
- Continue “Packet Marking” follow-up
- Assist in mini-DataChallenges
- Assist in DataChallenge’27 preparedness
- Evolution of HTTP-TPC
 - Improve PerformanceMarkers, provide expected checksum, investigate # of streams, improve error reporting
- Token adoption and compliance for Tape access

DOMA-BDT “*Exotic*” Topics of Interest

- **Pre-signed URL storage access**
- **Impact of HTTP/3 (TCP → UDP)**
- **Alternative checksum support**
- **Redesign performance markers**
- **Transfers between storages with no HTTP-TPC**

DOMA-BDT “Exotic” Topics: pre-signed URLs

- [Pre-signed URLs idea](#) presented already before DC24
 - Conclusion: investigate “only” after DC24
- **Current status:**
 - LHCb interested
 - ALICE tokens provide similar behavior
 - ATLAS investigates per-file token solutions
 - High token rates require performant and reliable service
 - Still to define right token capabilities, lifetimes and rates
 - Batch data transfers w/ OIDC tokens not main interest for EGI AARC
- **Goal: Understand advantages / disadvantages and role of pre-signed URLs in our workflows**

DOMA-BDT “*Exotic*” Topics: HTTP/3

- **Appearance of HTTP/2 posed problems already**
- **Much bigger change with HTTP/3 (QUIC protocol)**
(e.g.: TCP → UDP)

- **Understand benefits and drawback of new protocols**
- **Understand impact on our whole infrastructure**
(network, storage access)

- **Goal: Answer basic questions such as “Adopt or Avoid”**

DOMA-BDT “Exotic” Topics: checksums

- **Adler32 not the most popular algorithm for clouds**
(usually `crc32c`, `crc32`, `md5`, `sha1`, `sha256`)
- **Current “solution”**: no validation for storages w/o `adler32` support
- **Software (often) supports different / multiple checksums**
- **Goal: Understand usage of alternative checksums (where necessary)**
Understand performance impact of multiple checksums

DOMA-BDT “Exotic” Topics: PerfMarkers

- HTTP-TPC `PerfMarker` structure “adopted” from GridFTP
- Problems:
 - Several fields no longer used or just confusing (drop them)
 - Not enough details about transfer source & destination
- Opportunity:
 - More details for network link monitoring
 - Socket `TCP_INFO` details (e.g. retransmission)
 - Continuous link testing by production transfers
 - Test of right link by design (storage ↔ storage)
 - Detect faulty disk nodes
- Goal: Combined with packet marking effort, gain good insight on network behavior

DOMA-BDT “*Exotic*” Topics: HTTP-TPC service

- **HTTP-TPC is proprietary HTTP COPY extension**
(supported only by our storage implementations)
- **Tricky to use industry standard without additional layers**
 - Uploads (HTTP-TPC push) / Downloads (HTTP-TPC pull)
 - No way to TPC between storages with no COPY support
(multihop not ideal)
- **Investigation: HTTP-TPC gateways (implemented as separate service)**
 - Running close to the storage / suitable places on network
 - Common implementation of `PerfMarker` features
 - Small fraction of data transfers (relatively light service)

Discussion

(if time allows, first DOMA-BDT
HTTP-TPC extension proposal)

DOMA-BDT

Proposal: Data Integrity in WLCG (TPC)

by Hugo Gonzalez Labrador

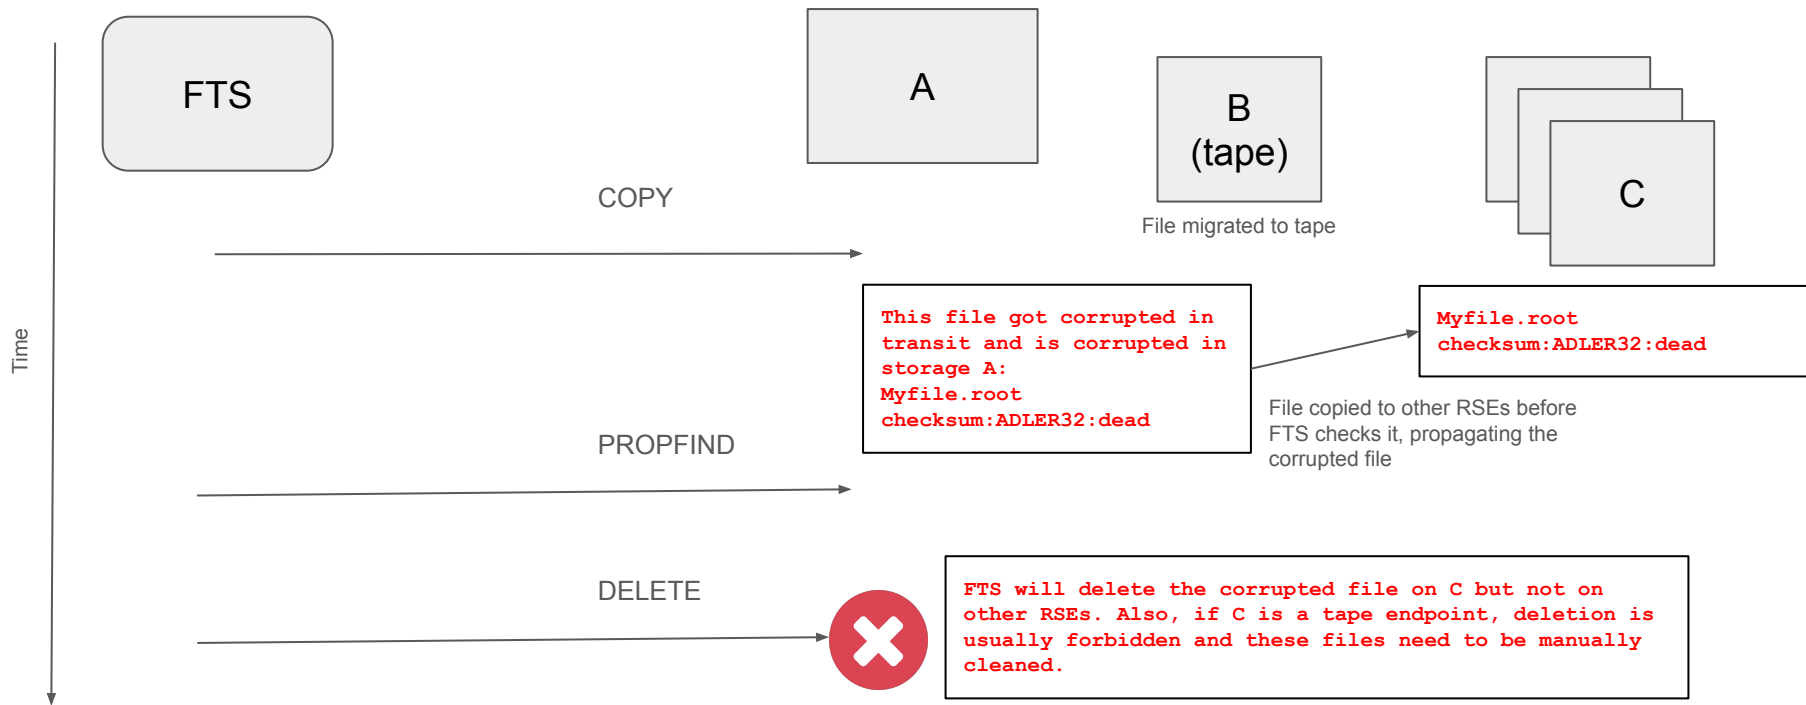
CERN - IT Storage and Data Management
Rucio Core team

Presenter: Mihai Patrascioiu
(CERN - IT Storage and Data Management)

Problem context

- In HTTP-TPC, files are copied directly between source → dest
 - TPC pull: active party downloads the file
 - TPC push: active party uploads the file
- FTS supervises the HTTP-TPC process
- Performs file verification at the end (file size & checksum match)
- Bad files are cleaned up from the destination
 - (Tokens: Cannot know which transfer will fail
 - require delete permissions beforehand !)
- Corrupted files may still end up on tape storages (rarely)
- Degree of impact varies by experiment (large for CMS Tape)

Failure Case Example



HTTP-TPC Data Integrity Proposal

- Evolve the HTTP-TPC to include checksum verification as part of the copy operation
- Move storages to RFC#9530 (governs checksum requests)
- New header "Repr-Digest" for the HTTP COPY command
- When storages receive this header, they MUST perform checksum verification
- Should move away from current checksum RFC#3230 (obsolete)

Proposal presented a DOMA-BDT, 16th April 2025 ([link](#))

Full details: <https://github.com/labkode/wlcg-tpc>

RFC#9530

- **Supersedes RFC#3230**

- **Headers**

`Content-Digest`

`Want-Content-Digest`

`Repr-Digest`

`Want-Repr-Digest`

- **Published**

February 2024 (Cloudflare and Team Digitale [Italian Government])

- **RFC is short and easy to read**

- **Digest value format is well-defined: [base64 structured field](#)**

```
Repr-Digest: adler=:<DIGEST>=:
```

RFC#9530: Pragmatic view

- **Changes to the codebase should be small (checksum support already exists):**
 - **Support new headers**
 - **Support new encoding & decoding (structured base64 field)**
- **HTTP responses left to implementers**
 - **re-use existing logic**
- **Note: "adler32" becomes just "adler"**
(for simplicity, suggest to accept both interchangeably)

HTTP-TPC & RFC#9530: Pragmatic view

- Storages would need to:
 - Implement RFC#9530
 - Accept "Repr-Digest" header in the HTTP COPY command
 - Perform checksum verification when writing the file

HTTP GET with "Want-Repr-Digest" header

HTTP PUT with "Repr-Digest" header

(Tokens: Storage cleans file
on failed PUT / GET !!)

- FTS to send the "Repr-Digest" header in the HTTP-TPC request
- Can storage developers evaluate magnitude of change?
(difficulty, estimated timeline)
- How to track implementation in DOMA-BDT?
(e.g.: can open tickets in preferred issue-tracking system)

Thank you!