

HTTP Third-Party Copy: Getting rid of GridFTP

Diego Davila, Brian Bockelman

Benedikt Maier, Fernando Garzon, Felipe Gomez-Cortes, Lisa Paspalaki

Stephan Lammel, Eric Vaandering, Nick Smith

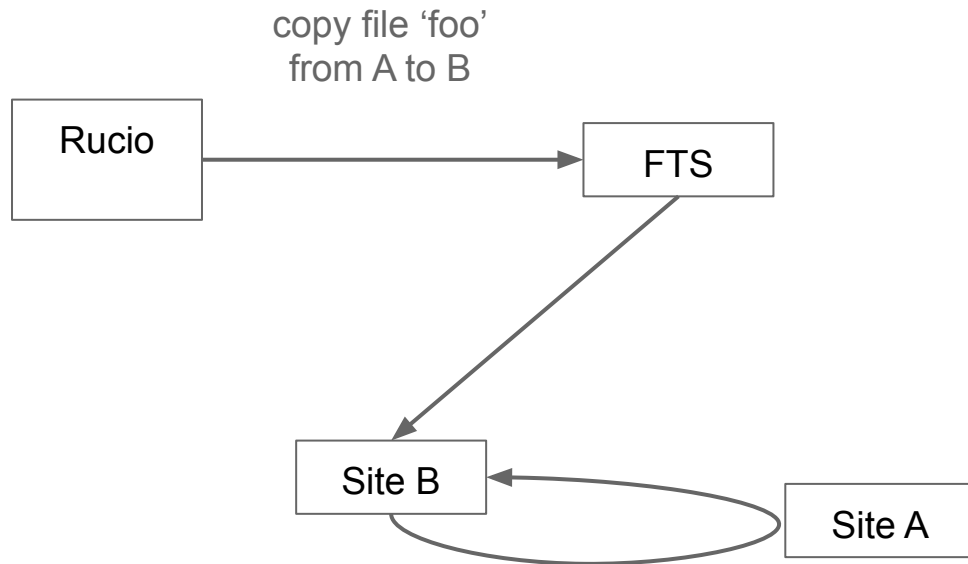
March, 2021



Let's start by breaking down the title into 3 different concepts:

1. Third Party Copy (TPC)
2. GridFTP
3. HTTP-TPC

Third Party Copy (TPC)



FTS*
orchestrates the
copy **without**
streaming the
file

(*)FTS - File Transfer Service

GridFTP

In simple terms:

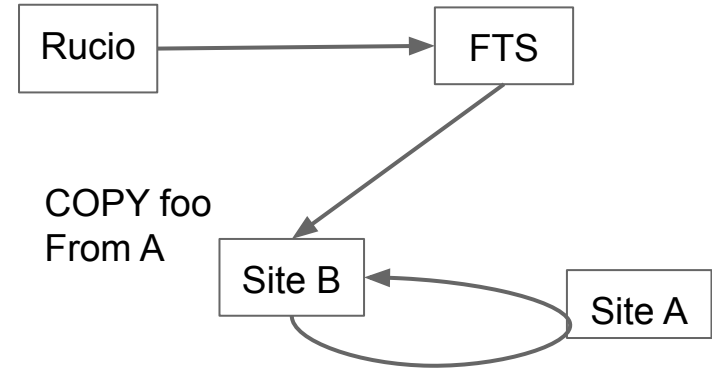
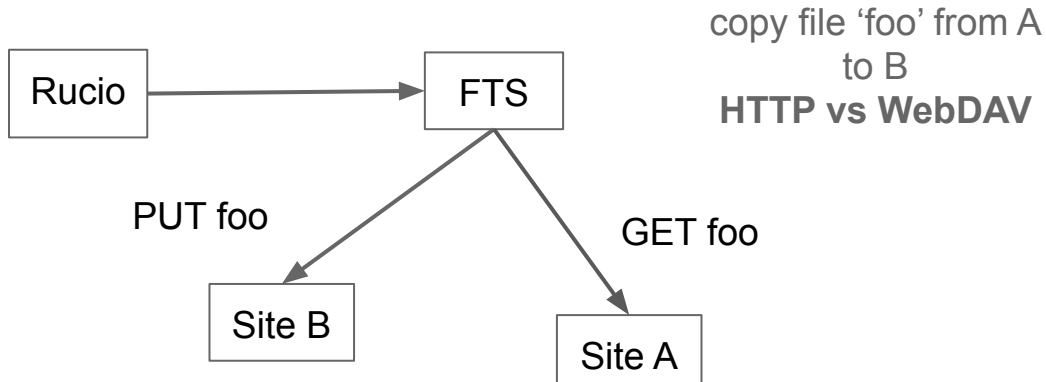
- An implementation of gsiFTP protocol which is an extension of FTP(*) to be used in the grid, i.e. FTP + GSI(**) authentication
- The most popular implementation used by the LHC experiments CMS and ATLAS to do TPCs i.e. move files around sites.

The issue: it heavily relies on the Globus Toolkit which announced its [end-of-support](#) on January, 2018

(*)FTP - File Transfer Protocol
(**)GSI - Grid Security Infrastructure

HTTP-TPC


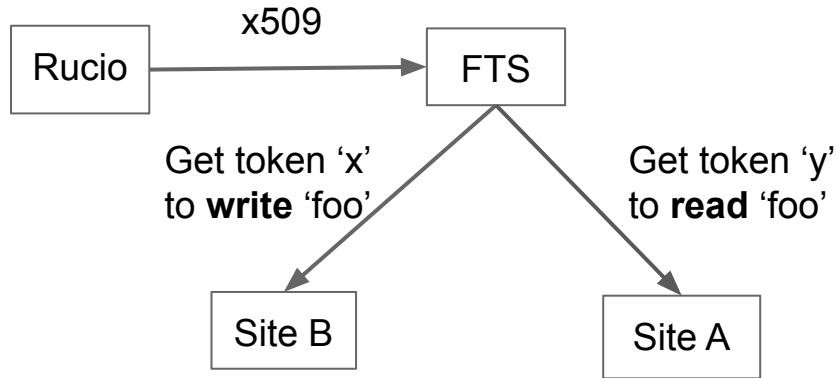
- The idea of doing TPC using the HTTP protocol
- HTTP by itself does not support TPC as we know it (without streaming)
- **WebDAV(*)**: an extension of the HTTP protocol that adds COPY to the verbs supported by HTTP.



(*)WebDAV: Web Distributed Authoring and Versioning

WebDAV

- **Bonus:** WebDAV supports tokens!
- A step further on getting rid of GSI
- Currently a mix of x509 and tokens are used for TPC



The adoption of WebDAV from the CMS perspective

This was done in 2 phases:

1. Testing

- Picked sites a guinea pigs
- Focused on debugging issues
- Integration with PhEDEx is managed by sites
- Rules in old TFC(*) (xml)

2. Getting the bulk of the sites

- Publicly announced to sites
- More CMS teams on board
- Focused on a systematic approach
- Integration with Rucio is centrally managed (can be automated)
- Rules in new TFC (json)

(*)TFC: Trivial File Catalog

The process from the CMS site's perspective

1. Minimal testing

- a. Open ticket to the site asking for a WebDAV endpoint
- b. Manual testing and debugging of the endpoint
- c. Make sure the new TFC (storage.json) is up to date

2. LoadTests and validation of the TFC

- a. The new TFC is used to create a `_Test` instance of the site
- b. LoadTests are enabled on that `_Test` instance
- c. Monitoring of the LoadTests

3. Production mode

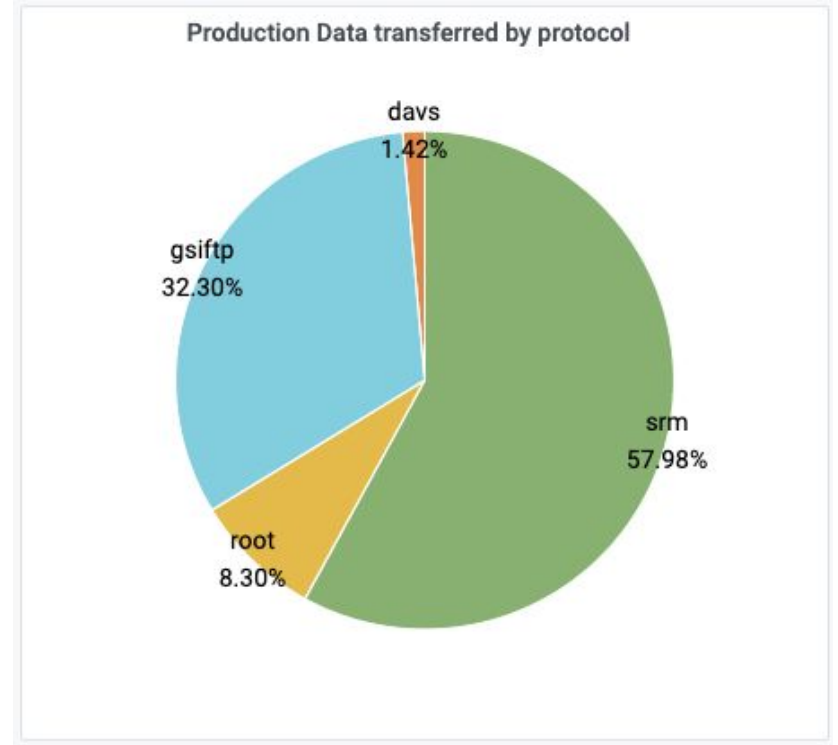
- a. Once the LoadTests seem fine the production instance is configured.

Current Status: CMS

- 35 sites reported to have a WebDAV endpoint
- 30 have passed manual tests
- 12 have production reads enabled in Rucio
- 3 have production reads and writes enabled in Rucio

Status per site can be seen here:

<https://twiki.cern.ch/twiki/bin/view/CMS/WebDAVStatus4Sites>



Production transfers by protocol in the last 30 days

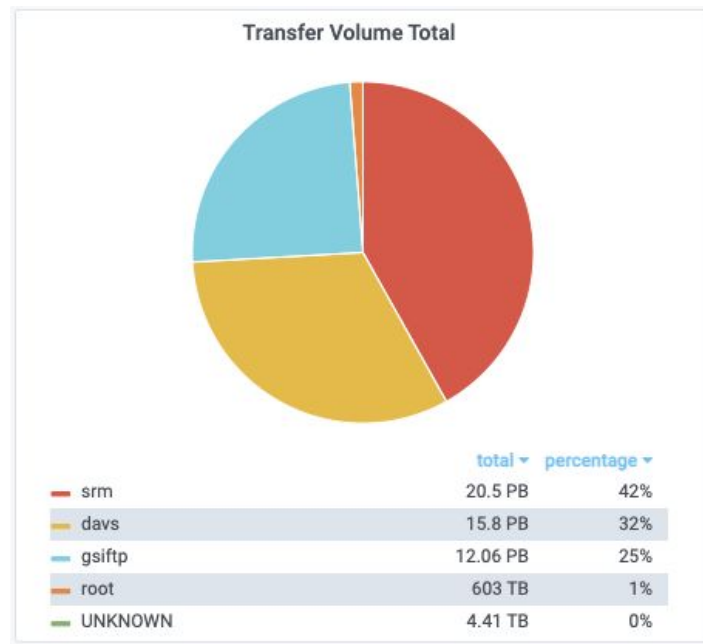
Current Status: ATLAS

Sites using WebDAV in production:

- T1 - 11/12
- T2 - 29/60
- T3 - 8/24

Status per site can be seen here:

<https://atlasdistributedcomputing-live.web.cern.ch/tpc/>



Production transfers by protocol in the last 30 days

Time for Questions