




Science & Technology Facilities Council  
Rutherford Appleton Laboratory



**GridPP**  
UK Computing for Particle Physics

# **DOMA**

## **Third-party Copy**

Tim Adye

Rutherford Appleton Laboratory

GridPP 42 Collaboration Meeting  
The Cosener's House, RAL  
24<sup>th</sup> April 2019



**ATLAS**

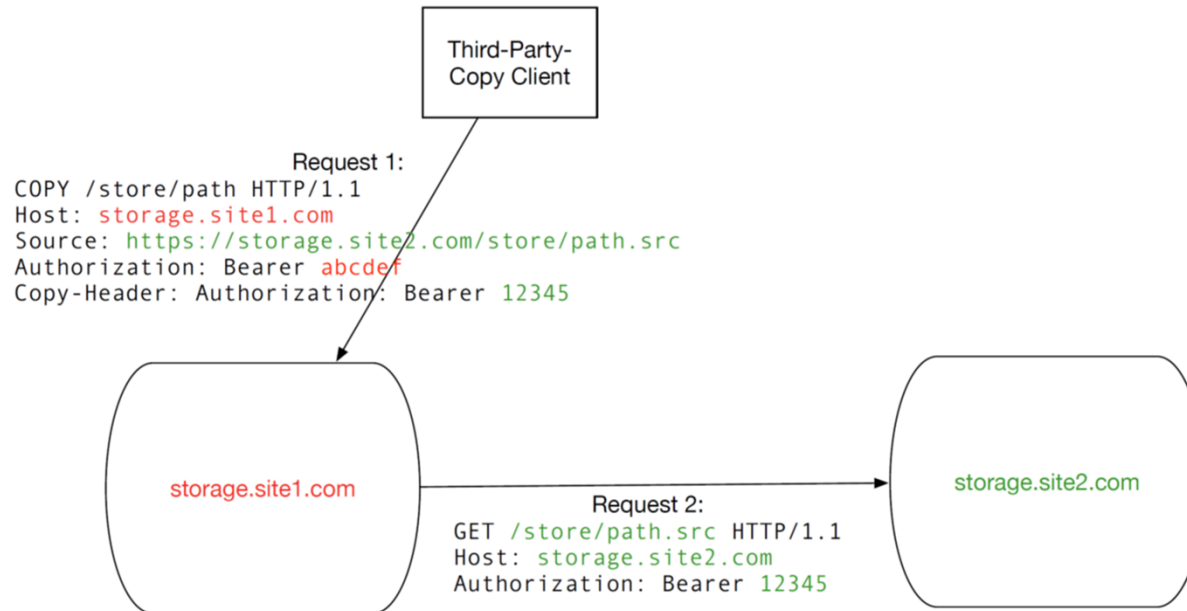
- WLCG DOMA project (Data Organization, Management, and Access):
  - TPC: Third-Party Copy
  - QoS: Quality of Service
  - ACCESS: data access, content delivery, and caching
- Concentrate on TPC in this talk
  - Two talks on XCache later in this session

Three phases:

1. **Prototype / implementation**: demonstrate viability of protocols
  - Ensure all storage implementations have a valid alternate in production
  - **Mostly complete**, but still some issues with some storage systems ← see next slide
2. **Early deployment**: ensure rollout of alternates at all sites with >3 PB storage
  - Preliminary deadline 30 June 2019
3. **Widespread deployment**: rollout to remaining WLCG sites

# Third-Party Copy

- Third-party copy allows
  - moving data from **site 1** to **site 2** controlled by a client at **site 3**
  - Eg. FTS at RAL requests transfer from **SLAC** to **Glasgow**



- Currently implemented in **GridFTP**, but we need other options as we move away from the Globus Toolkit

# TPC implementation

- Developing TPC for two protocols:
  - XRootD — can now delegate X509 credentials for the transfer A→B
  - HTTP / WebDAV — uses token-based authentication
- and different storage technologies:
  - XRootD — recommend v4.9.0 for TPC. Implements both XRootD and HTTP protocols
  - dCache — HTTP TPC fully supported. XRootD with delegation in v5.0
  - DPM — requires DOME. Works with XRootD v4.9.0
  - EOS — HTTP supported. XRootD doesn't yet support delegation
  - StoRM — HTTP works in latest release
  - Echo — XRootD fix implemented and some success with TPC ← see later
- Documentation building up for each case on [TWiki](#)

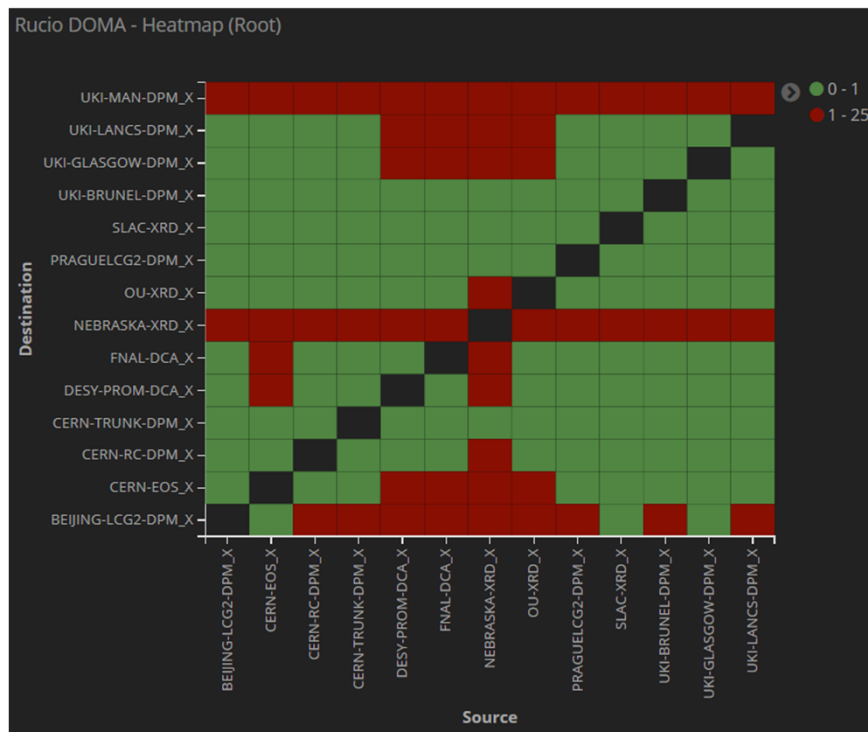
# TPC production / testbed / development sites

[DOMA TPC TWiki](#) gives details of current endpoints

- [RAL \(Echo & Castor & S3\)](#)
- Prague
- [Glasgow](#)
- [Brunel](#)
- [Manchester](#)
- [Lancaster](#)
- CERN DPM trunk
- CERN DPM release candidate
- DESY developer testbed
- DESY dedicated testbed
- AGLT2
- BNL
- [Imperial College](#)
- PIC
- CERN EOS pre-production
- INFN-T1
- [QMUL](#)
- NERSC
- SLAC
- Oklahoma University
- Nebraska
- Bonn
- FNAL
- Beijing
- DynaCloud CERN
- DynaCloud CERN (Grid instance)
- IN2P3
- Brussels
- Florida
- SURFSara (production & test)
- NDGF (production & preprod)
- LRZ-LMU
- University of Victoria
- TRIUMF

# Functional tests

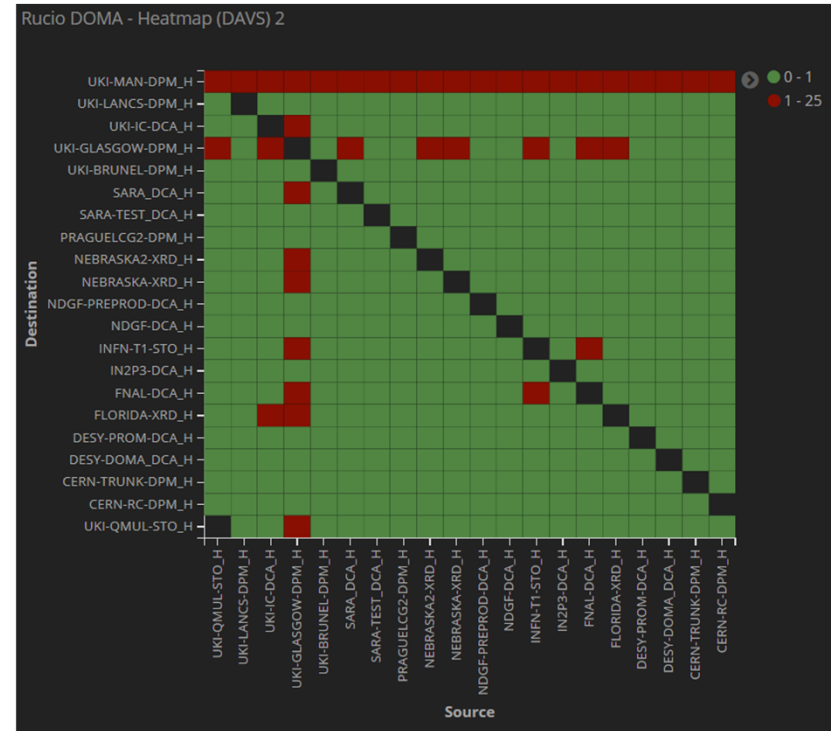

- Use Rucio to manage test transfers between many test sites
  - XRootD and HTTP
  - Monitoring displayed in Kibana
- Connectivity status shown in Kibana matrix plots



**XRootD transfer status**

(on 16/04/2019)

# Functional tests – HTTP



## HTTP transfer status


(on 16/04/2019)

- Also running nightly “smoke tests” to quickly identify problems
  - Performs a series of atomic tests for each entry point
  - email reports



# Scale tests

- Recently started scale tests of HTTP transfers between some sites
  - regularly moving > 500 TB/week across test sites



# RAL Echo XRootD

- XRootD access to Echo used in production from worker nodes at RAL
  - Uses an XRootD plugin, **libXrdCeph**
- Now testing TPC access to allow use also in transfers between sites
  - Two test gateways at RAL:
    1. DOMA TPC test gateway using production Echo instance (XRootD 4.9.1 rc1)
    2. ALICE testbed using Tier-1 Ceph development cluster (XRootD 4.8.5)
- TPC works better with latest XRootD versions, 4.9.0 or 4.9.1

# Echo XRootD development

- Worked round XRootD double-slash (//) feature
  - XRootD expects to be backed by a filesystem, so thinks nothing of sanitising path by adding a leading / or removing multiple /s.
    - This does not always work for an object store like Echo, where path is an exact string lookup
  - Already using XRootD options to disable much of this behaviour for production services
    - TPC complicates this. Can't expect other sites (2<sup>nd</sup> or 3<sup>rd</sup> parties) to have these same options.
  - Main issue is that XRootD expects (or adds) a leading / in the path (ie. // between host and path).
    - Thought to fix this at a lower level by removing leading /s in libXrdCeph.
    - Found an equivalent fix by specifying a transform in the CMS TFC name mapping
- ALICE requires XRootD TPC access already for production use
  - Currently working on using ALICE authentication tokens

# Conclusion

- DOMA TPC group has something working for most storage systems
  - Working on getting systems updated and documenting procedures
  - Still working on Echo, but making good progress
- Can now start to roll out in production for more systems