

WLCG DOMA BDT May 2023 update

Brian Bockelman, Maria Rios, Petr Vokac

TAPE REST deployment status

HTTP only
transfers

- Access to tape – last SRM dependency / almost last [Globus dependency](#)
- [TAPE REST API specification](#) – uniform tape access for CTA, dCache, StoRM
- Already supported by FTS, CTA and dCache – stress test end of 2022
 - [StoRM pre-production](#) implementation now also available on INFN-T1 testbed
- ATLAS already started to use TAPE REST for production T0/T1 tape transfers
 - CTA at CERN – [early March planning](#), enabled ([OTG0076390](#)), mid-April production
 - dCache at FZK ([GGUS:160869](#))
 - started with LOCALGROUPDISK and 8.2.17, few issues discovered & fixed
 - short path 256 -> 1024 chars changed in **8.2.22** (improved in 9.2 [dCache#7104](#))
 - [configuration](#) – [site specific JSON](#) published on well-known URL + firewall for REST endpoint
 - tuning of transfer parameters – default values updated in 8.2.20 ([WLCG BDT meeting](#))
 - FTS – different retention for transferred files compared to SRM [FTS-1913](#)
- ATLAS contacted T1 and ask them for deployment plans (preferably by end 23)
 - RAL CTA also moved to TAPE REST (this may allow to remove multi-hop via Echo in future)

Further TAPE development

- Improving data placement for TAPE data reads
 - Archival metadata tests at FZK
 - Standard way of passing archival metadata
 - CTA presentation at [ATLAS S&C #75 Week](#) in two weeks
- Tape access with tokens
 - `storage.stage` behavior clarified recently ([WLCG BDT meeting](#))
 - hoping to see news at tomorrow's [dCache Workshop](#)
 - WLCG TAPE API & Token presentations

Storage token configuration status

- CMS already [started ask sites](#) to configure tokens on production SEs
 - [CMS SAM/ETF probes](#) extended to use tokens (access wlcg-doma-tpc e-group)
 - WebDAV: [se_webdav.py](#) ([documentation](#))
 - xrootd: [se_xrootd.py](#) ([documentation](#))
 - started with EOS and XRootD sites (XRootD 5.5.4)
 - followed by dCache sites
- ATLAS
 - production dCache enabled for tokens and https&roots (PRAGUEELCG2) a month ago
 - few issues found & fixed (e.g. [issue#7129](#) with xroot protocol)
 - work in progress on (simple) ATLAS SAM/ETF probes for storage access with tokens
 - HTTP token support enabled on INFN-T1 StoRM
 - storage namespace layout for disks and tapes vs. one gateway
 - EOS ATLAS with HTTP tokens enabled
 - troubles with tokens (ZTN) enabled for roots protocols – configuration reverted
 - new [ATLAS IAM optional groups](#) used by IAM scope policies
 - group atlas/policy/storage-test - for ATLAS SAM/ETF probes and clients allowed to test SE
 - storage.* gives client access to /atlasdatadisk/SAM/ + /atlasscratchdisk/SAM/
 - temporarily(?) also to /atlas:datadisk/SAM/ and /atlas:scratchdisk/SAM/

audience:

<https://wlcg.cern.ch/jwt/v1/any>

dCache 8.2 mandatory for ATLAS & CMS

- WLCG JWT token support already in 7.2, but with different "scitoken" plugin
- uniform token support in "oidc" plugin comes with WLCG JWT support in 8.2
 - we don't want to deal with two different configurations
- avoid other [troubles with older releases](#) (e.g. related to old CaNL)
- [current status from WLCG CRIC](#):
 - 8.2.x - 48% (32 sites)
 - 7.2.x - 44% (29 sites)
 - 6.2.x - 2% (1 sites)
 - 5.2.x - 6% (6 sites)
- dCache developers recommends latest version of stable dCache for token support
 - e.g. 8.2.17 came with security fix for tokens
 - 8.2.22 include another fix for sites that rely on symlinks

Pre-signed URL discussion

- [LHCb presented](#) idea to use *industry standard / S3* way to access storage ([WLCG BDT Meeting](#))
- Still a lot of open questions
 - Our infrastructure is different than Amazon (hundreds sites / orgs)
 - It may not be possible to use (S3) standard without extension
- A lot of effort already invested in WLCG tokens
 - Not everything is yet clear but we should focus
 - resolve use-cases that are not well covered (understood) by WLCG JWT tokens (e.g. [WLCG-AuthZ-WG/#22](#))
- ATLAS & CMS will use WLCG JWT tokens for DC24 and we'll push service providers in this direction

Rucio & FTS token design

- Developers working on design document
- On June WLCG DOMA BDT agenda
 - open questions related to WLCG JWT profile

Misc

- IAM 1.8.2 release hopefully this week
 - performance limits / scalability / robustness needs to be better understood
 - CMS expects hundreds of tokens per hour
 - ATLAS roughly thousands per hour (with our current limited knowledge)
 - It is still too easy for one user to cause Halloween troubles
 - IAM built on top of OpenID Connect reference implementation [MitreID](#)
 - Database design causing performance issuer, e.g. [iam#579](#), [iam#580](#)
 - Too much state stored in DB (e.g. access tokens [iam#537](#))