

Storage & Tokens

[Data Challenge 2024 Workshop](#)

2023-11-10

[Petr Vokáč](#), Stephan Lammel

Transfers with WLCG JWT tokens during DC24

- [Proposal for DC24](#)
- Optional goal – tokens can be turned on/off at any point
 - Storage sites can participate in DC24 regardless of token support
 - X.509 VOMS proxies
 - It up to data distribution frameworks to use tokens when possible
 - deletions
 - HTTP-TPC
 - only when both parties support WLCG JWT tokens
 - stress ATLAS IAM with production load
- DC24 can use tokens only if there are supported by full chain
 - **storage -> sites** (many sites with different implementation)
 - FTS – good progress [presented yesterday](#) (few instances)
 - data distribution frameworks status (one instance per experiment – simpler to upgrade)
 - Rucio(?), Dirac(?)

Storage implementations & WLCT JWT tokens

- [WLCG JWT token profile](#) is "HEP proprietary" format
 - only friendly storage implementations understand content of our tokens
 - WLCG tokens, SciTokens, EGI AARC tokens, ..., Alice tokens
 - nightmare for site administrators
 - GUT profile – [mailing list](#) ([xkcd](#))
 - Storage
 - dCache – minimal version 8.2.23 (June 2023), important updates in 8.2.36 / 9.2.4
 - avoid obsolete scitoken plugin
 - StoRM – v1.3.1? (from SE ops point of view important updates in latest version)
 - XRootD – 5.5.4 minimal usable version (Mar 2023), latest recommended
 - EOS – recent 5.2 version still with issues (scitoken library?)
 - DPM/Dynafed – no token support (still in use by few sites)
- Basic / [generic documentation](#) for WLCG SE with tokens configuration (dCache info will be updated today)

ATLAS tokens for SE deployment status

- Still just one site for "each" storage implementation
 - PRAGUELCG2 (dCache)
 - CERN-PROD (EOS)
 - INFN-T1 (StoRM)
- [Namespace reorganization](#) proposal killed
 - clients have to deal with different storage scope PATH prefix
 - e.g. ATLASDATADISK RSE – storage.read:/atlasdatadisk for PRAGUELCG2 but storage.read:/BNLTOD1 for BNL
 - Rucio have no problem to deal with different prefixes but it's a bit messy for capability based model
 - common audience: fqdn (e.g. se1.farm.particle.cz, doesn't match WLCG JWT profile root://se1.farm.particle.cz:1094)
- SAM/ETF storage probes with tokens
 - not yet ready for storage without well-organized namespace
 - recent presentation mentioned that [ATLAS storage probes needs to be rewritten](#) from scratch
 - temporarily provide similar tests used to test tokens with [HTCondor-CE](#) & [ARC-CE](#)
 - allow sites to validate their SE token configuration
- Manual tests by experts available for ATLAS IAM users in the atlas/policy/storage-test group and client with storage.*:/ scopes
 - more complex without well-organized namespace, simple "storage.*:/atlasdatadisk" scope doesn't work everywhere
 - these users will be able to create 10+ different top level directories like BNLTOD1 at each ATLAS SE with token support

site configuration
campaign via GGUS
for sites involved in
DC24

CMS Site Storage Token Status

- tested config for CMS with few US sites early on
- developed detailed SAM/ETF probe
 - to not only check token access works
 - but provide detailed information that helps site admins understand/debug/correct any detected issue
 - test every 15 minute and also on-demand on ETF web site
 - for both WebDAV and XRootD storage endpoints
- about 90 active sites in the CMS computing grid
- many DPM sites
 - DPM does not support WLCG/IAM/SciTokens
 - all but 6 completed migration to other storage technologies
- four technologies used by CMS sites:
 - native xrootd, STORM, dCache, and EOS

CMS Site Storage Token Status

- **native xrootd was token-ready first:**
 - 12 of 14 Tier-0,1,2 sites are fully IAM-issued token ready
- **SToRM**
 - all 3 Tier-0,1,2 sites are fully ready
- **dCache**
 - 3 of 32 Tier-0,1,2 sites fully ready
 - deployment started only last Friday (improved no-access config for token without storage scopes)
- **EOS**
 - 0 of 7 Tier-0,1,2 sites fully ready (v5.2 now just available)
- **also working with Tier-3 sites but not counting them in token readiness**

CMS Site Storage Token Status

- **most sites have been very responsive**
- **so far commissioning has been quite smooth**
- **delay mainly from storage technology readiness for WLCG/IAM/SciTokens**
- **we expect to have storage at all CMS Tier-0,1,2 sites token ready by the end of this year (assuming no EOS v5.2 surprise)**
 - **issues found even with 5.2**
 - **may not be ready with EOS by end of year**

SE token support summary

- ATLAS
 - 3 site with token support
 - GGUS campaign for sites participating in DC24
- BelleII
 - 3 sites with token support (33%)
- CMS
 - 18 out of 49 sites (37%) already support tokens
 - deploy / configure tokens on each SE
- DUNE
 - would like to use tokens, tokens supported at FNAL
- LHCb
 - would like to use tokens, 0 sites

BACKUP

ATLAS SE namespace & capabilities

- ~~[Namespace reorganization](#)~~ proposal killed
- more complicated implementations, additional storage specific details must be provided, e.g.
 - `davs://dcgftp.usatlas.bnl.gov:443/pnfs/usatlas.bnl.gov/BNLT0D1/rucio/user/jhinds/9f/9c/user.jhinds.34215941._000022.CxAOD.root`
- unable to determine **issuer prefix** (must be in CRIC/Rucio protocol configuration)
- can't use same storage scope `storage.read:/atlasdatadisk` to access ATLASDATADISK on each storage
- at least we agreed on **common audience** for ATLAS storage