



Token Transition Workshop Summary

Brian Lin
OSG Software Area Coordinator
University of Wisconsin–Madison





OSG Token Transition Workshop

- Oct 14 - 15, 2021
 - Virtual
 - 166 registrants
 - 138 unique zoom attendees
- <https://indico.fnal.gov/event/50597/>
- Workshop Goals
 - Provide a space for OSG staff and communities to share their progress
 - **Discuss the OSG timeline and revisit if necessary**
 - Educate the OSG community on tokens and how to use them
 - Help sites accept token-based pilots



OSG Software Plans

- OSG CEs and GridFTP hosts are acutely affected by the February 2022 milestone
- After February 2022, XRootD will continue to work with GSI/VOMS proxies
- OSG CEs should be updated to HTCondor-CE from OSG 3.5 “upcoming”
- Sites using GridFTP should install XRootD Standalone from OSG 3.6 when released



What's next?

- Oct 2021 (WLCG milestone): pilot submissions may be performed with tokens
- Dec 2021 (WLCG milestone): WLCG VOMS admin servers retired in favor of IAM
- **Feb 2022: OSG 3.5 end-of-life.**
 - GCT packages are no longer distributed in OSG Yum repositories
 - Remaining OSG packages have no GCT dependencies
 - XRootD supports X.509 and VOMS proxies without GCT dependencies
- Mar 2022 (WLCG milestone): all storage endpoints on the WLCG provide support for tokens
- Mar 2023 (WLCG milestone): WLCG experiment data stageout and reads performed with tokens
- Mar 2024 (WLCG milestone): X.509 client authentication becomes optional

Oct 14, 2021

OSG Token Transition Workshop

3

Slide 3 from Brian Lin's talk: [“OSG Software Bearer Token Transition Plans”](#)

Open Science Pool

- The Open Science Pool (OSPool) is an HTCondor resource pool based on GlideinWMS
- 52/137 HTCondor-CEs reporting to the OSG Central Collector accept OSPool token-based pilots
- 95% of execution points are reporting back to the pool using HTCondor IDTOKENS
- Storage tokens are automatically generated per user job and have been in production for ~4 years

Use Case 1: Submitting pilots to HTCondor-CE

We purposely generate a **SciToken per CE** instead of one per VO:

- The CE SciToken is only accepted at a single host. All others reject it. Minimizes the “blast radius” of a stolen CE token.
- Each CE token also receives a unique subject and a unique token identifier.
- Scopes limit the token to being used for job submission.
- 1 hour expiration: token only needs to create the session with the CE. Continuously renewed; does not need to travel with the pilot.

```
{
  "sub": "vofrontend-SLATE_US_MNSU_DISCOVERY",
  "scope": "compute.read compute.modify compute.create
compute.cancel",
  "wlog_ver": "1.0",
  "aud": "osg-ce.msu.edu:9619",
  "jti": "3f776538-e4c6-4781-86f2-b7c734684ae6",
  "iss": "https://scitokens.org/osg-connect",
  "exp": 1634188686,
  "iat": 1634177886,
  "nbf": 1634177886
}
```



Community Report: ATLAS

- Harvester successfully submitting token-based pilots to BNL & MWT2
- Aiming to update all central Harvesters to support tokens by end of October
- Move to storage & user tokens in 2022-2025
- Token → CE user mappings
 - Production vs analysis pilots are mapped to separate users based on
 - Need a mapping for SAM/ETF tests

Bot Clients for Submission


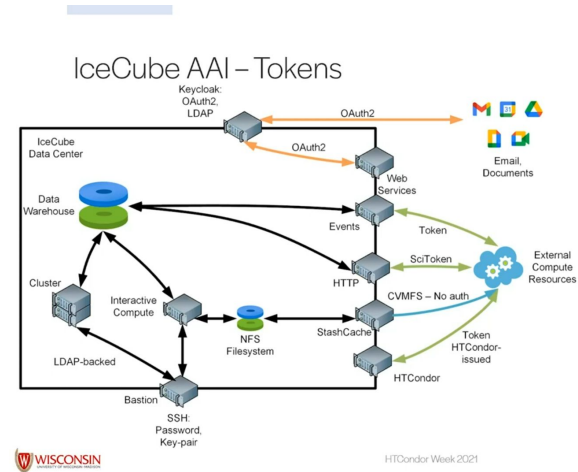
- The clients are registered in ATLAS IAM
 - Use **client credentials flow** authorization
 - Can directly get access token. Agreed to be optimal for harvester use case
 - No need of refresh tokens. No need to run odic-agent
 - 2 separate clients, for mapping `production_vs_analysis` to different accounts on the sites
 - `production:harvester-compute-production` (7dee38a3-6ab8-4fe2-9e4c-58039c21d017)
 - `production:harvester-compute-analysis` (750e9609-485a-4ed4-bf16-d5cc46c71024)
 - E.g. In CondoCE `etc/condor-ce/mapfiles.d/10-scitokens.conf`:
`SCITOKENS /https://atlas-auth.web.cern.ch/,7dee38a3-6ab8-4fe2-9e4c-58039c21d017/ atlasprd`
`SCITOKENS /https://atlas-auth.web.cern.ch/,750e9609-485a-4ed4-bf16-d5cc46c71024/ atlasplt`
 - Analogous to voms proxy roles. Less change of model for now, easier for the sites
 - Ideally mapping to only one account is OK, as ATLAS jobs, both user and production, run in singularity containers
 - Will need another client as "lcgadmin" role for SAM/ETF tests
- Assigned with WLCG compute scopes:
 - `"compute.read compute.cancel compute.modify compute.create"`
 - Token privileges are restricted to communication with CEs

15

Slide 15 from FaHui Lin's talk: "[Token Transition Report: ATLAS](#)"

Community Report: IceCube

- Moving away from everything in LDAP → Keycloak
- Will allow for offloading of business services (email, calendar)
- Storage access using SciTokens
- HTCondor pool will use IDTOKENS
- To come: federated identity and SSH with OAuth2




- Many changes being AAI
- Offloading services (Email, documents)
- HTTP- and Events-based transfer (SciToken, JWT-based)
- StashCache only for certain applications

Slide 11 from Benedikt Reidel's talk: "[Token Transition Report: IceCube](#)"




Community Report: LIGO

- Pilot submission through OSG-managed GlideinWMS
- Storage access tokens will be managed through a combination of HTCondor, Kerberos, and Vault



HTCondor Local Issuer



- SciTokens generated by HTCondor credmon
- `iss` in the token set to <https://scitokens.org/ligo>
- Static website based on <https://github.com/scitokens/ligo>
- Private key configured into HTCondor
- Public keys manually added to <https://scitokens.org/ligo/oauth2/certs>
- OSG XRootD configured to trust <https://scitokens.org/ligo>
- OSG XRootD configured to map scopes to file system paths

Slide 20 from Ron Tapia's talk: "[Token Transition Report: LIGO](#)"



Community Report: CMS

- Starting the campaign to update site CEs to accept token-based pilots
- Making good progress on moving from GridFTP → WebDAV-TPC
- In the late stages of testing IDTOKENS for the global pool



Conclusions

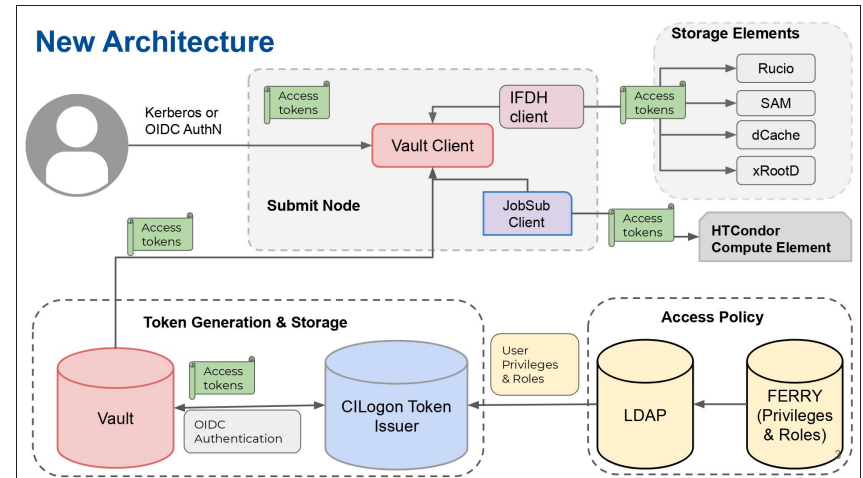
- (At least) **Three main areas** undergoing important changes in auth/authz due to GSI/GridFTP retirement
- **WebDAV endpoints** for HTC transfers are being deployed
 - Campaign is in a **good shape**
 - Solution for **T1 TAPE transfers** with srm+http is **being tested** and verified
- **Ready to start campaign for migration to SCITOKEN for HTCondorCEs**
 - **Hopefully not an invasive campaign**: update software to the latest version and add mapping for CMS SCITOKENS
 - Maybe a bit more difficult for EU sites (usage of Argus for mapping and software not taken from OSG repo)
- Migration of **Global pool to IDTOKENS** also on our radar
 - CRAB server (TaskWorker) still uses GSI and the Argus server for mapping purposes.

13

Slide 13 from Marco Mascheroni's talk: ["Token Transition Report: CMS"](#)

Community Report: FNAL

- Federated identity project will start with OAuth/JWT for FNAL users then extend access to users at partner institutions
- User tokens will be based on Vault, the CILogon token issuer, and FERRY + LDAP
- FNAL will maintain separate token issuers for each of their international VOs but are working through sub-VO user mappings



Slide 2 from Mine Altunay's talk: "[Token Transition Report: FNAL](#)"



Community Report: BNL

- Working on upgrading and testing token support in HTCondor-CEs
- DUNE's dCache will be upgraded to support FNAL/WLCG tokens
- Different experiments have varied token requirements and timelines
- The experiments using tokens will mostly all be using different token issuers

Token Status - HEP

ATLAS

- Will need to use WLCG Tokens (IAM)
- Compute: Successfully Tested HTCondor CE w/ Harvester
 - Currently upgrading all CE's to handle tokens
- Storage: no plans yet will work with rest of ATLAS

Belle II

- Follow lead from Host lab KEK

DUNE

- Will need to use FNAL Tokens
- Compute:
 - Currently upgrading all CE to handle tokens. Need to test with DUNE jobs
- Storage:
 - upgrading our dCache instance to latest Golden release (7.2)
 - will configure for FNAL Tokens and WLCG Tokens (for testing)
 - Collaborate with FNAL on Rucio/FTS transfers between FNAL - BNL



Rucio + FTS

- Rucio has had token support since 2019
 - Next iterations on token support include fine-grained scopes for security and data embargo purposes
- FTS supports OIDC-provided tokens and SE-issued tokens obtained by FTS
 - Can accept access tokens and exchange them for refresh tokens from the IdP. Currently the same access tokens are used between source and destination.
 - Can take user X.509 proxies and exchange them for SE-issued tokens



Remarks

- Properly scoped and audience limited tokens
 - Plan is to return these as part of the `list_replicas` query to the Rucio server
 - Will probably also need alternative REST endpoint to request SE-tokens for users who do not use Rucio clients to download data
- Rucio as the central token-issuing component
 - Tokens can be very specifically scoped/audience restricted
 - Gives us unique opportunity to limit data access (very interesting for non-HEP sciences)
 - Data embargos based on projects/scopès/datasets/metadata
 - Risk needs to be properly assessed
 - Not only Rucio, but entire token-driven software stack
 - Security audits

2021-10-15

8

Slide 8 from Martin-Stefan Barisits' talk: "[Rucio Token workflows](#)"



Token Education

- [“The Future of Bearer Tokens”](#) by Jim Basney
 - Principles of least-privilege auth; capability-based authorization
 - JWTs, OAuth2, OpenID Connect
 - New project to assist with the transition to tokens: <https://sciauth.org/>
- [“Token AuthZ and Validation for Site Services”](#) and a Bearer Tokens Hands-On by Derek Weitzel
 - Token technical details: token structure, attributes, and how to decode them
 - Jupyter notebook to familiarize site administrators with tokens



HTCondor-CE Hackathon

- OSG staffed a “hackathon” room to help site administrators troubleshoot and verify token-based pilot submissions to their CEs
- **OSG staff verified token-based pilot submission to 8 HTCondor-CEs**
- Common issues:
 - HTCondor-CE uses default CA certificate locations for its SSL handshake used to authenticate incoming tokens:
<https://htcondor.com/htcondor-ce/v5/configuration/authentication/#configuring-certificates>
 - A WLCG token compatibility bug in the scitokens-cpp package. Fixed in 0.6.2 available in EPEL; HTCondor packaging to be updated to require the appropriate version



Summary

- **OSG communities are making good progress to meeting the February 2022 Grid Community Toolkit end-of-life in the OSG**
- **Next Steps**
 - Communities should continue coordinating with US-based sites to accept token-based pilots by February 2022
 - Communities need to plan out their token → site user mappings
 - OSG Security + EGI CSIRT tabletop exercises
 - HTCondor-CE improvements
 - Payload audit logging
 - Provide more flexibility in token → user mapping (scope-based?)

Questions?

This material is based upon work supported by the National Science Foundation under Grant Nos. 1836650 and 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.