

ALTAS

Token Transition


FaHui Lin (UTA)

ATLAS PanDA

211014 (Token Transition Workshop)

WLCG & ATLAS Plans

WLCG Token Transition Timeline

| Milestone | Date | Description | Dependn. | Teams |
|-----------|---------------|---|---------------|--|
| M.0 | Feb 2021 | Produce document with list of use cases for CMS VOMS-Admin API. | None | WLCG Ops |
| M.1 | May 2021 | WLCG baseline services include HTTP-TPC endpoints. Mind: tape services come later. | None | WLCG Ops, Storage providers |
| M.2 | June 3-4 2021 | WLCG hosts “CE and pilot factory hackathon” | None | Pilot framework providers |
| M.3 | July 2021 | Production IAM Instance(s) Available for at least 1 LHC experiment, likely CMS and possibly ATLAS | None | WLCG Ops, IAM, CERN IT |
| M.4 | Oct 2021 | Pilot job submissions <u>may</u> be performed with tokens. | M.3 | Experiments, pilot framework providers, OSG/EGI, sites, Monitoring |
| | |  | | |
| M.5 | Dec 2021 | VOMS-Admin shut off for CMS. IAM is sole authz provider for those (including for VOMS server) | M.3 | WLCG Ops, CERN IT |
| M.6 | Feb 2022 | OSG ends support for the Grid Community Toolkit | M.1, M.4 | OSG |
| M.7 | Mar 2022 | All storage services provide support for tokens | M.1 | WLCG Ops, Storage providers |
| | ? | All VO's shut off VOMS-Admin | | |
| | Sept 2022 | End of HTCondor support for GSI Auth | | |
| M.8 | Oct 2022 | Rucio transfers performed with token auth in production | M.7 | Rucio, Experiments |
| M.9 | Mar 2023 | Experiments stageout & data reads performed via tokens. | M.7 | Experiments |
| M.10 | Mar 2024 | X.509 client auth becomes optional. | M.9, M.8, M.4 | Experiments |

under discussion



ATLAS

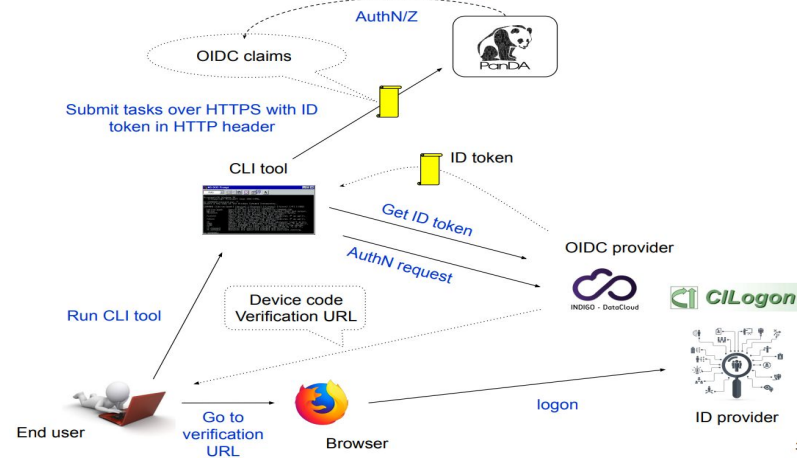
- ATLAS already has the Indigo IAM (WLCG M.3)
 - <https://atlas-auth.web.cern.ch/>
- Generally ATLAS tries to follow WLCG's timeline
 - While sometimes WLCG's timeline can be too ambitious
 - In production we don't want to use software that is not ready or widely tested at scale
 - E.g. CMS will retire current VOMS servers by end of this year (WLCG M.5)
ATLAS will maintain VOMS longer, and will join CMS months later in next year
- A few points of timeline mentioned in ATLAS S&C plenary in Sep 2021:

| ATLAS Distributed Computing | |
|---|---------|
| Transition to tokens | |
| Submission from Harvester to all HTCondor CEs with tokens | 03/2022 |
| All users move from VOMS to IAM for X509 | 12/2022 |
| All job submission and data transfers use tokens | 12/2025 |

Status of ATLAS Components

PanDA (1) - Client

- PanDA already has token authN/Z for users
 - Compliant with OIDC/OAuth2
 - Support both tokens and x509
 - Support in panda-client-1.4.45+
 - User can interact with PanDA with ID token
 - To submit, retry, kill, show tasks, ...
 - PanDA authorizes the user based on name and groups claims in ID token
- The mechanism is already working for sPHENIX and Rubin users in DOMA PanDA
 - With dedicate Indigo IAM for DOMA PanDA: <https://panda-iam-doma.cern.ch>
 - Use CILogon for ID providers
- The mechanism is already tested with ATLAS IAM
 - Migration timeline depends on the ATLAS decision. Probably during LS3 (2025)
 - Users do not want the auth change during data taking
- Ref:
 - <https://panda-wms.readthedocs.io/en/latest/architecture/iam.html> (PanDA docs)





PanDA (2) - Harvester & Pilot

- PanDA has some token authN/Z mechanism for bots
 - First attempt was to authenticate Harvesters with self-signed token. Not used in production
- Implementation for ATLAS
 - Will extent the mechanism to work with ATLAS IAM token
 - Timeline depends on ATLAS
- PanDA-Harvester
 - Internal between PanDA and central Harvesters. Easily changed if needed.
- Pilot
 - Need to access storages, aside from PanDA
 - Dependent on the mechanism to renew access token for pilots on WNs in the future
 - Still under discussion (Mechanism with Vault? Involving PanDA?)



Harvester - Job Submission

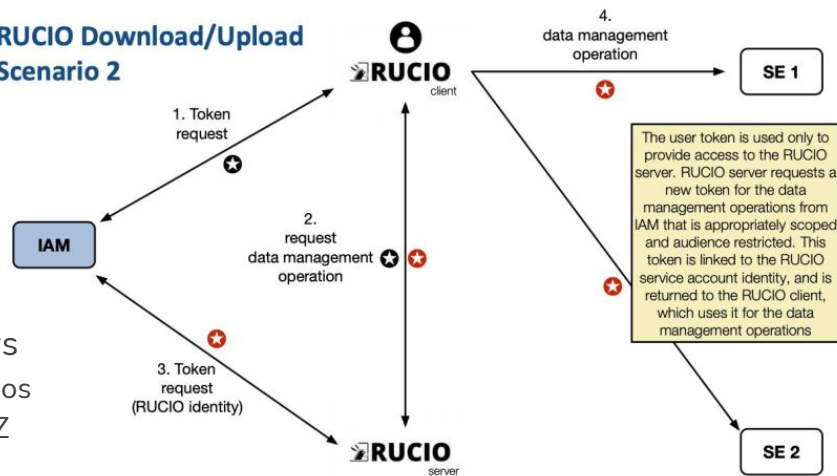
- Harvester already supports submission to HTCondorCE with token
 - Verified it works with ATLAS IAM tokens
- Details will be mentioned later in upcoming transition plan



Rucio

- Capability based authentication
 - Available since Rucio 1.22.0
- Plan: move to fine-grained token workflows
 - Now working on fine-grained token scenarios
 - Addressing open questions in WLCG AuthZ workgroup
 - Will require extensive development in Rucio
 - Security audit of our full-stack workflows will be needed
- More in Martin's talk tomorrow
 - <https://indico.fnal.gov/event/50597/contributions/225902/>

RUCIO Download/Upload Scenario 2



© Andrea Ceccanti

US Sites (1) - Services to be Migrated

| | | AGLT2 | BNL | MWT2 | NET2 | SWT2 | US HPCs |
|---------|----------------|-------|-----|------|------|------|---------|
| CE | HTCondorCE | ✓ | ✓ | ✓ | ✓ | ✓ | |
| | ARC CE | | ✓ | | | | |
| | Harvester | | | | | | ✓ |
| Storage | dCache | ✓ | ✓ | ✓ | | | |
| | XRootD | | | | ✓ | ✓ | |
| | Tape endpoints | | ✓ | | ✓ | | |
| | FTS | | ✓ | | | | |
| | Globus service | | | | | | ✓ |
| Env. | WN environment | ✓ | ✓ | ✓ | ✓ | ✓ | |



US Sites (2)

- HTCondorCE upgrade should finish by Feb 2022
 - In accordance with OSG timeline to end GCT support
 - Have performed submission tests with WLCG token issued by ATLAS IAM in MTW2 & BNL
 - Sites also need to use other tokens from OSG and FNAL, to be tested
- Components that need no token transition: CVMFS, batch systems (slurm, condor, ...)
 - Sites may upgrade batch systems along with necessary CE upgrade
- Analysis facilities (shared T3) will be deferred for the future

Upcoming Transition - Harvester & CondorCE



Short-term Goal

- Follow up WLCG's timeline M.4:
Pilot job submissions may be performed with tokens by Oct 2021
 - Submit from Harvester to HTCondor CE with the token issued by ATLAS IAM
 - But also keep x509 proxy in the job for pilot
- Upgrade all HTCondorCEs in ATLAS sites by Sep 2022
 - According to timeline of OSG & HTCondor



Requirements - Submission with Tokens

- Bot clients registered in ATLAS IAM
 - Granted with WLCG compute scopes
 - <https://github.com/WLCG-AuthZ-WG/common-jwt-profile/blob/master/profile.md#capability-based-authorization-scope>
 - See next page
- Condor:
 - Condor version 9.0.5+ or 9.1.1+:
Can submit with both token and proxy. The CE can decide which credential to authenticate
 - But not 9.3+ or later, as GSI support will be dropped
- Condor CE:
 - OSG 3.6 or HTCondor-CE 5, which supports bearer token
 - <https://htcondor.com/htcondor-ce/v5/configuration/authentication/#scitokens>
 - Set up token credential mappings for ATLAS IAM tokens
- Harvester:
 - See next pages



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-58039c21d817)
 - “**production:harvester-compute-analysis**” (750e9609-485a-4ed4-bf16-d5cc46c71024)
 - E.g. In CondorCE /etc/condor-ce/mapfiles.d/10-scitokens.conf :
`SCITOKENS /^https:\\/\\/atlas-auth.web.cern.ch\\/,7dee38a3-6ab8-4fe2-9e4c-58039c21d817/ 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



Harvester

- IAM token credential manager plugin
 - As of Harvester v0.2.9
 - Can get access tokens for all ATLAS CEs
 - Fetches list of ATLAS CEs from CRIC
 - Obtains access tokens with the bot client credential from IAM
 - Runs periodically to keep tokens up to date
 - Runs as part of Harvester service. No need to run external processes (oidc-agent, scripts/cron)
 - Default token lifetime is 1hr. Harvester updates tokens every 30 minutes (configurable)
 - Tokens are put into cephfs to share with all harvester instances
- htcondor submitter plugin support token credential submission
 - As of Harvester v0.2.8
 - It picks up the corresponding token file according to the CE chosen to submit to
 - We add a line of +ScitokensFile with placeholder of token path in submit file template:
 - Both proxy and token in submit file template:

```
x509userproxy = {x509UserProxy}
+ScitokensFile = "{tokenPath}"
```


Example of evaluated submit file of a real Harvester worker

```
executable = /cvmfs/atlas.cern.ch/repo/sw/PandaPilotWrapper/latest/runpilot2-wrapper.sh
arguments = -s MWT2_TEST -r MWT2_TEST -q MWT2_TEST -j managed -i PR --pythonversion 3 -w generic --pilot-user ATLAS --url https://pandaserver.cern.ch -d
--harvester-submit-mode PULL --allow-same-user=False --job-type=ANY --resource-type SCORE --pilotversion 2
initialdir = /cephfs/atlpan/harvester/harvester_wdirs/cern_cloud/86/53/7828653
universe = grid
log = /data/atlpan/condor_logs/grid.%(Cluster).%(Process).log
output = /data/atlpan/condor_logs/grid.%(Cluster).%(Process).out
error = /data/atlpan/condor_logs/grid.%(Cluster).%(Process).err
transfer_executable = True
x509userproxy = /cephfs/atlpan/harvester/proxy/x509up_u25606_prod
environment = "PANDA_JSID=harvester-cern_cloud HARVESTER_WORKER_ID=7828653
GTAG=https://aipanda170.cern.ch/condor_logs_2/21-09-30_09/grid.%(Cluster).%(Process).out APFMON=http://apfmon.lancs.ac.uk/api APFFID=cern_cloud
APFCID=%(Cluster).%(Process)"
+harvesterID = "cern_cloud"
+harvesterWorkerID = "7828653"

grid_resource = condor osg-gk.mwt2.org osg-gk.mwt2.org:9619
+remote_jobuniverse = 5
+remote_ShouldTransferFiles = "YES"
+remote_WhenToTransferOutput = "ON_EXIT_OR_EVICT"
+remote_TransferOutput = ""
+ioIntensity = 0
+xcount = 1
+maxMemory = 1
+remote_queue = "analy_mwt2_ucose"
+maxWallTime = 7200

delegate_job_GSI_credentials_lifetime = 0

periodic_remove = (JobStatus == 2 && (CurrentTime - EnteredCurrentStatus) > 604800)
+remote_PeriodicRemove = (JobStatus == 5 && (CurrentTime - EnteredCurrentStatus) > 3600) || (JobStatus == 1 && globusstatus != 1 && (CurrentTime -
EnteredCurrentStatus) > 86400)

+sdfPath = "/cephfs/atlpan/harvester/harvester_wdirs/cern_cloud/86/53/7828653/tmpw8rc29hc_submit.sdf"

+ScitokensFile = "/cephfs/atlpan/harvester/tokens/ce/1eae42bccfd28fd6a339c2dffe756481"

queue 1
```

Credential flow for ATLAS jobs

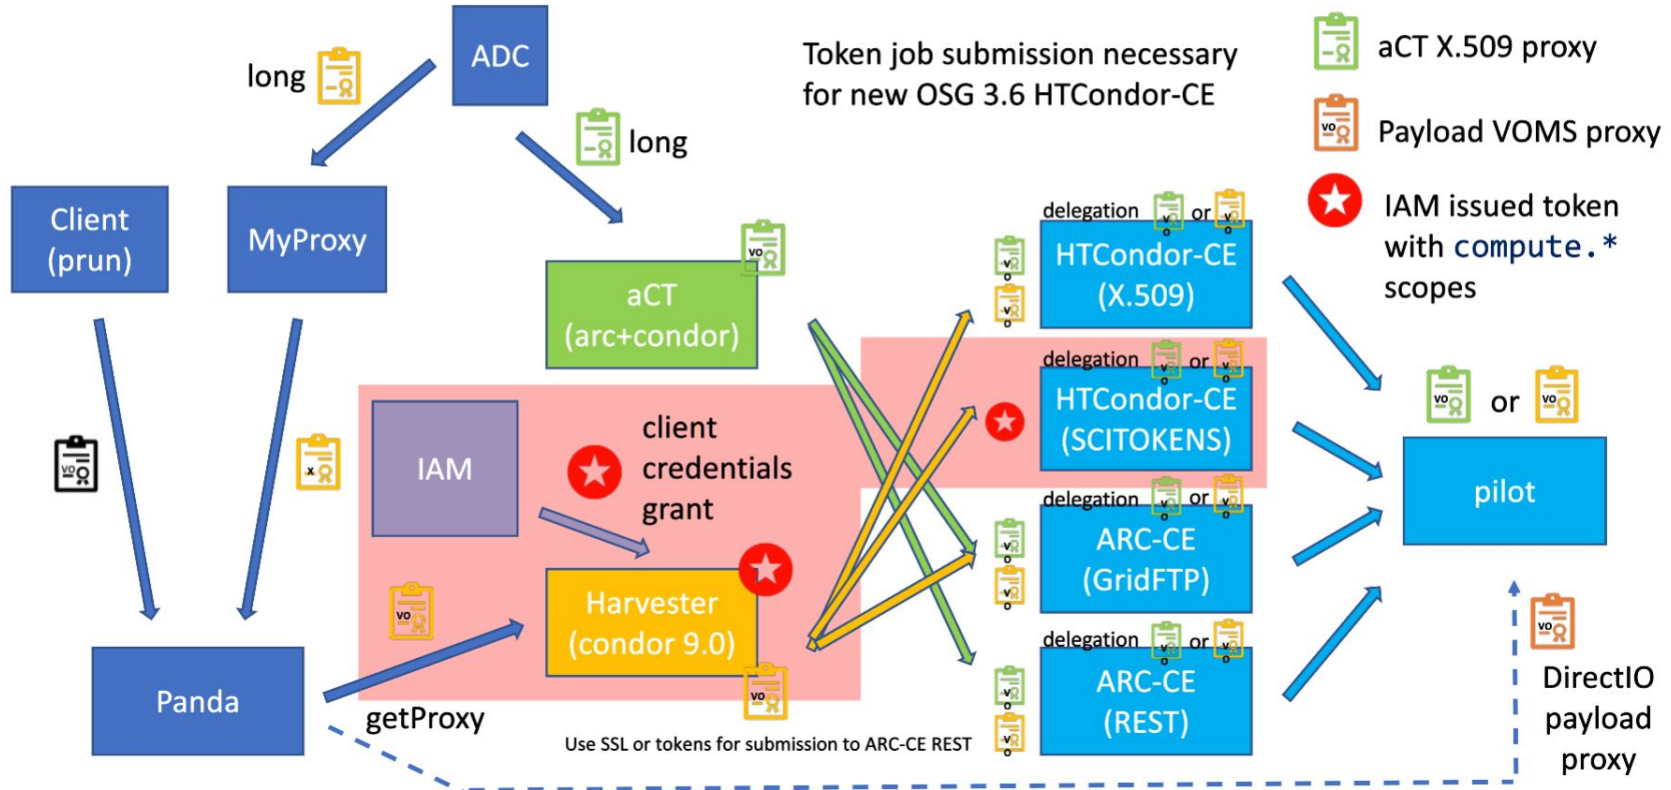


Image by courtesy of Petr Vokac



Verification

- Test with harvester v0.2.9 and its local condor v9.0.6
- CE at MWT2
 - `osg-gk.mwt2.org:9619` , set up by Lincoln
 - Test with dedicate PQ: MWT2-TEST
 - Submissions are successful with the access token. ATLAS jobs finish normally
- CE at BNL
 - `gridgk05.racf.bnl.gov:9619` , set up by Xin
 - Submissions are successful
 - Tokens are mapped correctly to production or analysis account



Upcoming Transition Plan

- Upgrade software and configurations on central Harvester and Condor instances
 - Harvester to v0.2.9; Condor to **v9.0.6** (latest stable release, still supports GSI)
 - Let Harvester always submit with both token and voms proxy
 - Harmless to ARC CE or Condor CE without token setup. Verified
- Just started this week !
 - Will take 1-2 weeks to carry out - including observation phase
 - Can finish upgrading central Harvesters and Condors by end of October
 - Will meet WLCG Token Transition Timeline M.4
- Meanwhile, ATLAS sites should take time to upgrade CEs to support token
 - No more change necessary in central (e.g. CRIC, Harvester configs, etc.)
 - US sites should move by Feb 2022, when OSG ends support for GCT
 - Other HTCCondor-CE sites are suggested to move by Jun 2022



Short-term Timeline

- For CEs at sites
 - Timeline suggested by Petr
 - US sites support tokens by Feb 2022
 - Other HTCondor-CE sites support tokens by Jun 2022
 - ARC-CE sites provide REST interface by Jun 2022
 - no tokens necessary, job submission can still use x509
 - All done before Sep 2022, EOL of condor v9.0.x (stable release using GCT)
- For central Harvester & Condors
 - Upgrade to condor v9.0.x and submit with both token and proxy, by Oct 2021 (ongoing)
 - Upgrade to condor v10.0.x (release in May 2022) , around Jun-Jul 2022
 - Next stable release. No proxy authentication. Support submission to ARC CE REST
 - Hopefully all sites will their finish CE upgrade then
 - Provide some condor instances with dev release (e.g. v9.3) if necessary



Summary



Summary

- We are making steady progress with some unknown still to be decided
- Full transition to tokens outside of OSG will still take several years
- Short term: 2021-2022
 - Job submission
 - Upgrade of CEs at sites due to OSG & HTCondor timelines
- Longer term: 2022-2025
 - VOMS to IAM
 - All storages & data transfers
 - Mechanism to renew token on WN
 - All sites (US & EU) readiness
 - Token for ATLAS users (PanDA, Rucio, storage access, ...)

Backup



Access Token issued by ATLAS IAM

- “sub”: ID of the bot client “production:harvester-compute-production”
- “aud”: put the FQDN:port of the CE
 - Identical to the endpoint on CRIC for Condor CEs
- “scope”: WLCG-defined scopes for compute resource
- “iss”: ALTAS IAM
- Lifetime: 1 hour (default)

```
{
  "sub": "7dee38a3-6ab8-4fe2-9e4c-58039c21d817",
  "aud": "osg-gk.mwt2.org:9619",
  "nbf": 1632994050,
  "scope": "compute.cancel compute.create compute.modify compute.read",
  "iss": "https://atlas-auth.web.cern.ch/",
  "exp": 1632997650,
  "iat": 1632994050,
  "jti": "d5baa0cb-518f-46b4-937d-f7ee7d954562"
}
```



Other points

- Format of aud for Condor CE
 - WLCG token spec recommends aud format like “**condor://condorce.example.com**”
 - While Condor CE's default is FQDN:port, like “condorce.example.com:9619”
Also, on CRIC this is the format of CE endpoint
 - Condor CE team will probably update default to also support WLCG format in next release (?!)
 - Harvester can do anything to follow the rules
- Condor plan to drop Grid Community Toolkit (GCT)
 - <https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki/index.cgi/wiki?p=PlanToReplaceGridCommunityToolkit>
 - As of version 9.3.0 (Oct 2021), and 10.0.0 (May 2022)
 - Will drop gridftp submitted to ARC CE, GSI authentication, proxy delegation
 - Cannot authenticate with x509 proxy
 - Stable release v9.0.x which uses GCT will end supports in Sep 2022
 - All sites needs to finish migration to authenticate token by then
- Submission to ARC CE REST interface
 - As of condor v9.1.0, it supports submission to ARC CE REST interface
 - Not related to tokens



JupyterHub on Cloud

- Jupyter/Dask cluster pools on Google cloud
 - A cloud R&D project in ATLAS distributed computing
 - For users to access resources/services not easily available on the GRID
 - <https://indico.cern.ch/event/1077170/contributions/4541620/subcontributions/351391/attachments/2324228/3958456/Cloud%20R%26D%20infrastructure.pdf> (*Amazon and Google Cloud R&D infrastructure* by F. Barreiro Megino)
- Start to integrate JupyterHub with ATLAS IAM
 - Currently trying only on a Jupyter testbed, not the main one users are using
 - Successfully to authenticate ATLAS users. Very easy to integrate and very promising
 - Exploring the possibility and understanding details about access control, etc.

