Accessing CE's with tokens

A.Tsaregorodtsev,
Aix Marseille Univ, CNRS/IN2P3, CPPM
DIRAC-Rucio Workshop, KEK, Tsukuba
19 October 2023

DIRAC
THE INTERWARE

Outline

- Pilot model reminder
- Pilot credentials
  - IAM
  - Check-in
- Status per CE
  - HTCondorCE
  - ARC
  - Cloud
- Multi-VO case
- Conclusions

Pilot model

- Pilot Factories (Site Directors) submit pilot jobs to Computing Elements
  - Using appropriate protocol and credentials
- Computing Elements authenticate and authorize pilot jobs to run :
  - Map jobs onto local Unix account applying local policies
  - Choose appropriate queues
  - Execute pilot jobs
  - Account for the consumed resources
- Pilot job communicates with the central DIRAC services to steer user's jobs execution
  - Getting job descriptions and sandboxes
  - Running user applications
  - Reporting job status
  - Uploading job results

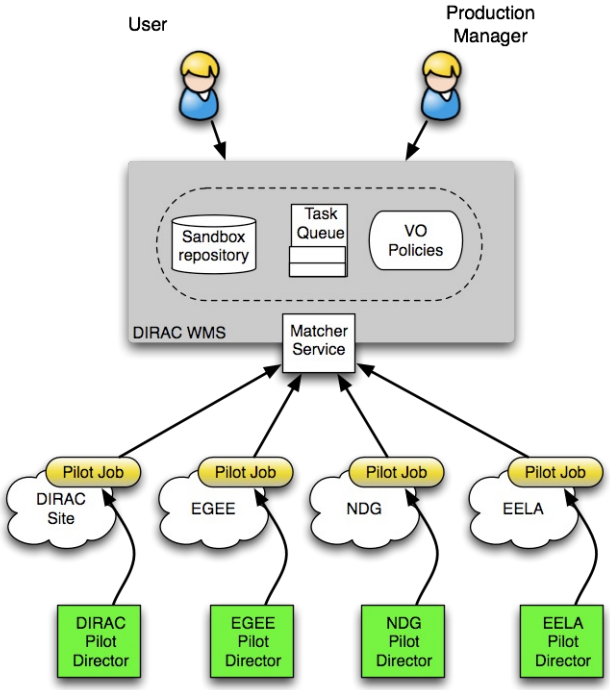

Pilot credentials

- So far, X.509 certificate proxies are used to authenticate pilots to sites as well as for the pilot communication with the central services
  - The certificate DN carries the user identity information
  - VOMS extension in the proxy carries the VO information
  - This mode will be still available in the 8.X DIRAC releases
- The use of robot or service X.509 certificates for submitting pilots is quite common
  - The robot identity represents the user community at the sites
- With the introduction of OIDC/OAuth2 tokens, the pilot submission procedure in DIRAC remains very similar

Pilot credentials

- The DIRAC server is registered as a client of an authorization service (IdP) with certain properties
  - The client must be enabled to get compute.* scopes
    - compute.create, compute.read, compute.modify, compute.cancel
- The token used to submit pilots is obtained with the ***client_credentials*** autorization flow
  - Simple: single request to the authorization service with the client ID and the secret
- The client credentials represent the whole user community (or several communities) pretty much like a robot certificate

Example pilot token


```
{
  "exp": 1696955839,
  "iat": 1696952239,
  "jti": "5f68a37e-f583-4740-a0c5-1b94c73bdb3c",
  "iss": "https://aai-dev.egi.eu/auth/realms/egi",
  "aud": "https://wlcg.cern.ch/jwt/v1/any",
  "sub": "17e1338a-61bd-43ba-9020-cab6624ff045@egi.eu",
  "typ": "Bearer",
  "azp": "196ac932-a367-4554-bfcf-c57d5c0a5e17",
  "scope": "compute.modify compute.create compute.cancel compute.read",
  "clientHost": "2001:660:5009:34:134:158:34:67",
  "nbf": 0,
  "preferred_username": "service-account-196ac932-a367-4554-bfcf-c57d5c0a5e17",
  "clientAddress": "2001:660:5009:34:134:158:34:67",
  "client_id": "196ac932-a367-4554-bfcf-c57d5c0a5e17"
}
```

WLCG IAM tokens

- Analysing pilot tokens at the CEs – IAM case
  - Tokens must have compute.* scopes in order to pass initial filtering
  - Local user/group mapping is done based on the ***sub*** and ***iss*** token fields
    - **sub** defines the local account
    - **iss** defines the VO
    - Enough for single-VO authorization service
  - Multi-VO WLCG IAM IdP services are not available yet – no need yet to provide additional VO/group information in the token

Check-In tokens

- The Check-In tokens can now be dressed with compute.* scopes to pass initial filtering at the sites
- Check-In is a Multi-VO IdP
  - The same **iss** value is used for several communities
    - Not enough for the local mapping – extra information is needed
  - Check-In enables adding group information to the pilot token
    - Example scope:
      eduperson_entitlement:urn:mace:egi.eu:group:biomed:role=pilot#aai.egi.eu
    - Needs definition of VO-specific policies in the Check-In - the client becomes a « member » of a VO
  - This Check-In functionality is not in production yet

Check-In tokens

- The Check-In pilot token must be unambigously mapped onto local account/group
- A special plug-in is developed that can be called by both HTCondorCE or ARC service while mapping:
  - https://github.com/EGI-Federation/check-in-validator-plugin
- The plug-in uses more complete token information:
  - **iss, sub, scope, audience, eduperson_entitlement**
- As a result the plug-in produces the necessary mapping to local account/group
- The token can be limited for the use with a single site by specifying the **audience** scope.
  - Not really used so far
  - Audience value format is to be defined (URL or ID or …)

Communication with DIRAC WMS

- Tokens are only used for pilot submission
  - The pilot tokens are not available in the job scope
- The pilots are sent to CEs together with pilot X.509 certificate in its bundle.
  - This certificate is used to communicate with the DIRAC central services
  - The pilot certificate becomes an internal DIRAC WMS implementation detail which can be replaced by another solution in the future

HTCondorCE case status

- Tokens are enabled starting with HTCondorCE 9.0
  - All the HTCondorCE sites can use tokens now
- Using SCITOKENS method
- Using **iss** and **sub** for user/group mapping
- Sites must be configured to accept DIRAC client token **sub**
  - Needs special arrangement between the sites and the DIRAC service admins
  - A general method to provide VO client token information to sites is discussed, e.g. as part of the VO-card
- All the LHCb HTCondorCE sites are configured to use tokens for pilot submission (*Alexandre*)
  - Easy case – single-VO IdP
- Several EGI HTCondorCE sites are accepting biomed VO jobs – as if a single-VO IdP case
- Tests of using Check-In plugin to use group information is successful on several sites
  - As soon as all the HTCondorCE sites install and configure the Check-In plugins, DIRAC will be ready to use them for all the VO's managed by Check-In

ARC case status

- Pilot submission with tokens is enabled when using the ARC REST interface
  - The token is added as the **Bearer** header in the http request
  - The user/group mapping is done based on the **iss** and **sub** token fields
- The pilot submission is demonstrated to work with ARC 6
  - Testing with ARC 7 is in progress
- The Check-In plugin should be suitable also for ARC CE's
  - This is still to be tested

Cloud sites status

- Multiple Cloud sites enable user access with tokens to their web dashboards
  - EGI FedCloud sites
  - Openstack
- DIRAC WMS manager authenticates to the cloud site as a member of a specific project/tenant corresponding to a particular VO
- The manager then defines Application Credentials linked to the project that will be used to create VMs with CloudComputingElements (*see Daniela's talk*)
  - No tokens are used while VM creation
  - Application Credentials are only available for Openstack clouds

Conclusions

- Pilot submission with tokens is fully operational for HTCondorCE's for the simple case of a single-VO IdP
  - Both IAM and Check-In case
- Pilot submission with tokens is demonstrated for ARC6 CE's. Tests with ARC7 are in progress
- The multi-VO Check-In dev service was demonstrated to support pilot submission to HTCondorCE sites with the special plug-in installed
- Work is in progress to demonstrate multi-VO DIRAC pilot submission to ARC7 CE's