Open Science Grid

MORGRIDGE INSTITUTE FOR RESEARCH

Update on the Token Transition

Brian Bockelman
OSG Technology Area Coordinator
Associate Scientist,
Morgridge Institute for Research

Token Transition

- Recall we are in the middle of three (related) transitions.
  - Moving the bulk data protocol from GridFTP to HTTP-TPC.
  - Retiring the (former) Globus Toolkit from our software stack.
  - Moving from GSI authorization to tokens.
- The above is the approximate order too!
  - ATLAS is at 1/3 traffic over HTTP-TPC and 1/2 sites preferring HTTP-TPC.
  - The Toolkit retirement should happen next year.
  - Given the number of interlocking pieces, GSI-to-tokens will finish ~2024.

Tokens – A Primer

- 'GSI' is based on using X.509 ("certificate") credentials for authentication.
  - It has several extensions to allow delegation, impersonation, group assertions.
  - It is based on the concept of *identity mapping*. Sites map the global identity (the CA-issued X.509) credential and map it to a local identity.
    - The credential is authorized to do whatever the local identity can do.
- We use 'capability-based bearer tokens':
  - Capabilities: Describes a specific authorized operation.
  - Bearer tokens: Any holder of the token is authorized.
  - Example: Signed statement saying "The bearer of this token is authorized to write ATLAS files."
  - We are using JSON Web Tokens (JWTs).
    - We support the 'SciTokens Profile' and the 'WLCG Common Token Profile'; both are agreed-upon ways to interpret token contents.

Tokens – A Primer

Open Science Grid

Encoded PASTE A TOKEN HERE


```
eyJraWQiOiJyc2ExIiwiYWxnIjoiUlMyNTYifQ.
eyJ3bGNnLnZlciI6IjEuMCIsInN1YiI6IjI3MjM
0ODQzLWZlZGYtNDJjOC1iYjgxLWExNjk1YmJkN2
MyOCIsImF1ZCI6Imh0dHBzOlwvXC93bGNnLmNlc
m4uY2hcL2p3dFwvdjFcL2FueSIsIm5iZiI6MTYx
NDc0NDA4Miwic2NvcGUiOiJvcGVuaWQgb2ZmbGl
uZV9hY2Nlc3Mgc3RvcmFnZS5yZWFkOlwvIHN0b3
JhZ2UubW9kaWZ5OlwvIHdsY2ciLCJpc3MiOiJod
HRwczpcL1wvd2xjZy5jbG91ZC5jbmFmLmluZm4u
aXRcLyIsImV4cCI6MTYxNDc0NzY4MiwiaWF0Ijo
xNjE0NzQ0MDgyLCJqdGkiOiIyNjhjYmQ1OC0wZj
EwLTRmYjctOWU5Mi1kZmVlMjVlZmJhMTIiLCJjb
GllbnRfaWQiOiJiMGQ4N2Q0Yi0wMjFkLTRmN2Yt
OTc0Yy1iY2E2YThlM2JlNDgifQ.XR5NrLGvMWh_
rRmTuV60OyoJTL0mY7EoF1y2MgumAObWOh5KkCw
ojrsAYEtJjaCEFr63oEK3DGB7ykDUpprh3nJjXf
```


Decoded EDIT THE PAYLOAD AND SECRET

HEADER: ALGORITHM & TOKEN TYPE


```
{
  "kid": "rsa1",
  "alg": "RS256"
}
```


PAYLOAD: DATA


```
{
  "wlcg.ver": "1.0",
  "sub": "27234843-fedf-42c8-bb81-a1695bbd7c28",
  "aud": "https://wlcg.cern.ch/jwt/v1/any",
  "nbf": 1614744082,
  "scope": "openid offline_access storage.read:/
storage.modify:/ wlcg",
  "iss": "https://wlcg.cloud.cnaf.infn.it/",
  "exp": 1614747682,
  "iat": 1614744082,
  "jti": "268cbd58-0f10-4fb7-9e92-dfee25efba12",
  "client_id": "b0d87d4b-021d-4f7f-974c-bca6a8e3be48"
}
```

IMPORTANT

Toolkit Retirement
!=
GSI Retirement

Examples:

XRootD has its own GSI implementation – we will use it for years.

HTCondor's SSL implementation can support client X.509 authentication (but *not* GSI proxies).

I expect some VOs will internally use GSI for much longer!

OSG 3.6 Compatibility

- Yesterday I had a talk that went through the interoperability scenarios. Summary:
  - OSG 3.6 has no supported GridFTP implementation. Data transfers need to move to GridFTP before OSG 3.5 EOL.
    - OSG 3.6 *does* have GSI/VOMS support for HTTP-TPC transfers via XRootD.
  - OSG 3.5 supports both token- and GSI-based pilot submissions to the HTCondor-CE.
    - OSG 3.6 *only* supports token-based submission.
    - More work to do on the pilot side currently!

Token Transition

- A quick reminder of the timeline:
  - 2017: the SciTokens project was funded by NSF to support use of capability tokens for the authentication & authorization infrastructure.
  - 2017: Globus announced the end-of-life for the Globus Toolkit.
  - (late 2017): OSG helped fork the Globus Toolkit to the Grid Community Toolkit.
  - 2019: OSG announced a timeline for transitioning to tokens for AAI.
  - 2021 (6 days ago): OSG 3.6 released without any Globus Toolkit dependencies.
  - February 2022 (planned): OSG 3.5 support ends – no Toolkit in supported projects.

OSG / OSG-LHC Timeline

- A few expected milestones coming up this year:
  - April 2021: OSG migrates hosted CE submission to SciTokens.
  - May 2021: (WLCG) LHC transfers migrate to HTTP-TPC.
  - June 2021: WLCG token issuers are available.
  - Oct 2021: (WLCG) LHC experiments give sites of receiving SciTokens-based pilots.
- We have been coordinating closely with FNAL SCD and WLCG.
  - Next few slides are part of the WLCG timeline draft (purposely left in various Google Doc markup!).

WLCG Timeline

Milestone ID	Date	Description	Dependencies	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.	None	WLCG Ops, Storage providers
M.2	May 2021	WLCG hosts “CE and pilot factory hackathon”	None	Pilot framework providers
M.3	July 2021	Production IAM Instances Available	None	WLCG Ops~~AAI~~, IAM, CERN IT
M.4	Oct 2021	Pilot job submissions may be performed with tokens	M.3	Experiments, pilot framework providers, OSG/EGI, sites

	Nov 2021	Coordinate documentation update with VOMS Admins and VO secretariats		Experiments, IAM, CERN IT
M.5	Dec 2021	VOMS-Admin shutoff. IAM is sole authz provider (including for VOMS server)	M.3	WLCG Ops, CERN IT~~AAI~~
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
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

Key Risks

- I view technology risk as relatively low: many of the key technologies are in place and working reasonably well.
  - One item to note is the token issuer. Current technologies are (a) IAM from INFN, (b) CILogon (subscription service), or (c) directly issue the token (simple cases only).
- There is effort risk:
  - We rely on external software providers to fix bugs and implement features. We contribute significantly to many projects but cannot do it all.
  - Collaborations have *many* competing priorities. Given the number we support, there's a risk at least one doesn't make this a priority in the next year.
- There is planning risk: We have found groups often don't start building timelines unless we push them on it.
  - Without knowing timelines and going through the exercise, there's higher risk of an unknown dependency.

Impact of Risks

- I worry about two impacts if the risks are realized -
  - Schedule: PATh / IRIS-HEP incurs additional months of OSG 3.5 support because stakeholders are not ready in time.
  - Coordination / Cohesiveness: There's a risk the community will fragment, require multiple solutions, and "re-implement GSI with tokens".
- Mitigation strategies:
  - We've written an NSF proposal to get coordination / technical effort to help with the above.
  - Coordinated closely with FNAL SCD and WLCG.
  - In the past we've worked to spread technical knowledge:
    - FNAL hosted Brian, Jim, and Todd in 2019 for a "SciTokens day".
    - WLCG held an in-person and a virtual hackathon in 2020.
    - Is this a useful model? Who else would be a good target?

What I need from the OSG Council

- Here are the ways the OSG Council can help:
  - Help identify missing use cases or potential problems. E.g., what communities are either unaware of this transition or are far behind?
    - We are happy to be consultants to build timelines. The problem is much more tractable when you start to break it down!
  - Help provide constructive input to existing timelines. "WLCG transition is impossible" is less actionable than "we are late on WLCG milestone M.5; it needs to push back 5 months".
    - We are happy to provide training and 'hackathon' support for developers.
  - Help provide effort to move the community forward. Particularly, FNAL has been very active in providing development effort on tokens.
    - We have 'shovel-ready' projects!

Questions?

support@opensciencegrid.org

This material is based upon work supported by the National Science Foundation under Grant No. 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.