WLCG Token based Authentication & Authorisation
Supporting Information for VO Questionnaires.
Authored by the WLCG Authorisation Working Group

Background

Current WLCG Authentication & Authorisation Infrastructure

WLCG users today are authenticated via x509 certificates (User Certificates) issued by Certificate Authorities accredited by the Interoperable Global Trust Federation (IGTF). They are authorised according to their VO groups and roles, which are maintained in VOMS. The authorisation attributes are added as extensions to certificate proxies generated by the user using their User Certificate. The resulting token is called a VOMS proxy.
The typical user flow is as follows

The user registers at the CERN User Office and with their experiment secretariat, has their identity verified and proves their affiliation to an experiment. The user is entered into the CERN HR database.
The user gets a certificate from any eligible Certificate Authority
The user registers with VOMS Admin for their VO and presents their certificate to authenticate. Their email address is matched with their record in the CERN HR database to confirm that their identity has been verified. The VO manager approves the registration. From that time on, the HR ID is used as the unique ID in the VOMS DB, allowing the e-mail address to be changed independently at either end.
VOMS attributes, such as groups and roles, are maintained by VO Managers.
The user can submit jobs to the grid or access data on the grid by saving their User Certificate locally, running voms-proxy-init to generate the VOMS proxy, and running the relevant grid command; their proxy will get delegated to grid services as needed.

Motivation for Change

There is growing interest in using more modern and flexible authentication and authorisation models, in both the research and commercial sectors. The following points highlight the principal reasons why tokens are being considered as the successor of the x509 model.

Software and systems are increasingly enabling protocols such as OAuth2.0 [https://oauth.net/2/ (https://oauth.net/2/)] and OIDC [https://openid.net/connect/ (https://openid.net/connect/)]. Storage systems used by the WLCG community are already being enhanced to receive these credentials thanks to efforts such as SciTokens [https://scitokens.org (https://scitokens.org)]. Adopting updated standards presents a benefit to Infrastructure developers (since robust, user-friendly libraries exist in multiple languages) and facilitates smooth integration and interoperability between tools and organisations.
User certificate management is increasingly unintuitive as early career researchers are becoming used to standard web based authorisation flows. A better user experience can be facilitated by an Authentication and Authorisation Infrastructure (AAI) in line with the AARC Project's [https://aarc-project.eu (https://aarc-project.eu)] Blueprint Architecture [https://aarc-project.eu/architecture/ (https://aarc-project.eu/architecture/)].
The need to expose user VO membership and certificate information to Grid services in a manageable way via VOMS and VOMS Admin may lead to violations of data protection requirements such as those of the European General Data Protection Regulation (GDPR).

WLCG Authorisation Working Group

The WLCG Authorisation Working Group was set up to identify a solution for token based authentication and authorisation that improves the user experience and alleviates certain data protection concerns. The group is sending this document and related questionnaire to WLCG VOs to collect additional information. The short term objective is to combine the feedback with the WLCG AAI prototypes and present recommendations to the WLCG Management

Board in early 2019.
The group collected a list of requirements between 2017 and 2018: https://twiki.cern.ch/twiki/pub/LCG/WLCGAuthorizationWG/WLCG_Authorisation_Requirements.pdf (https://twiki.cern.ch/twiki/pub/LCG/WLCGAuthorizationWG/WLCG_Authorisation_Requirements.pdf). This document forms the basis for design decisions taken.
What is Token based Authentication & Authorisation?
Generally speaking, what we mean is an Infrastructure where users are authenticated and authorised via JSON web tokens. “JSON Web Tokens are an open, industry standard RFC 7519 method for representing claims securely between two parties.” [https://jwt.io/introduction/ (https://jwt.io/introduction/)]
Tokens are encoded strings of data, issued by an issuer with which the client (receiver) has a trusted relationship. They can be easily passed around.

OAuth2.0 is an authorisation protocol that often uses JWT to convey the data. Open ID Connect (OIDC) is a layer on top of OAuth2.0 that authenticates (rather than authorises) users. In short, if an application just wants to authenticate a user it would often use OIDC. If an application is looking to authorise the user on a deeper level (i.e. authorise access to a second application on behalf of the user) it would use a variety of OAuth2.0 flows.
Proposed Architecture
Overview
An Authentication and Authorisation Infrastructure (AAI) architecture was designed following the principles of the AARC Blueprint Architecture [https://aarc-project.eu/architecture/ (https://aarc-project.eu/architecture/)], which includes a component that centralises group membership (VO user management, groups and roles), token translation (for example from a token to a VOMS proxy certificate for legacy systems) and token provisioning (passing the tokens to the users in a user friendly process for command line interaction). This central component is commonly termed a service-provider/identity-

provider proxy (IdP-SP-Proxy); from the perspective of an internal service it is the sole identity provider and from the point of view of an external identity provider it is a single service.
The medium term aim is for WLCG services to enable token based authentication and authorisation (OAuth2.0 and OIDC) when undergoing enhancements, and for the new AAI to facilitate this transition by providing the necessary token translation. Each VO would manage an instance of the new WLCG AAI, which would replace VOMS-Admin for user membership. There is an extended offer that these instances can be hosted by CERN IT, in the CDA group; hosting sensitive data of this type must be done at CERN according to new HR data protection practices.
In the long term it is hoped that VOMS can be retired.

Interaction with CERN
Currently CERN and WLCG operate two entirely separate authorisation models. CERN is built on Kerberos and SAML for authentication, with e-groups for authorisation. WLCG is built on x509 with VOMS for authorisation. There is no link between the two. The long term vision is for both CERN and WLCG to move to token based authentication & authorisation in an interoperable manner. Specific grid functionality will remain under the control of WLCG and is is envisaged that the new user management services will be integrated with CERN SSO to provide a unified user experience. Identity vetting will continue to rely on established procedures at the CERN Users' Office for the foreseeable future.
Potential Workflow
It is envisaged that the future workflow would be more-or-less as follows

The user registers at the CERN User Office and with their experiment secretariat, has their identity verified and proves their affiliation to an experiment. The user is entered into the CERN HR database.
The user registers with the WLCG AAI for their VO by authenticating through their Identity Provider, initially their account is disabled. Their

account is enabled once their email address is matched with a record in the CERN HR database to confirm that their identity has been verified.
VOMS attributes, such as groups and roles, are maintained by VO Managers.
The user completes a user friendly process to associate their local environment with the WLCG AAI (e.g. uploads SSH key).
The user runs a command (or includes it in a script) that provisions a token (and proxy for legacy systems) locally and can submit jobs as usual.

The notable difference is that there is no need for the user to get a certificate.
Specific Information
The following information provides more detailed discussion around key points of token based authentication and authorisation.
Groups, Roles & Capabilities
The VOMS authorisation model revolves around groups and roles. The group and role information, as well as personal information, is included in the VOMS proxy and interpreted by the WLCG service receiving it, meaning the authorisation decision point is made at the service. An alternative authorisation model sends capabilities in credentials, rather than roles and groups. A capability is a statement of what a token bearer is allowed to do, e.g. “The bearer of this token is allowed to access file X”. Capability based tokens typically contain minimal information about the identity behind the token.
The following tables describe simplified token content:

Group/Role Token Attribute
Value

Name
Jane

Home Organisation
Edinburgh

Groups
[testvo-user, testvo-testgroup]

Roles
[researcher]

Role
user

Capability Token Attribute
Value

VO
testvo

Opaque identifier
asd6768yuasd

Capabilities
[allowed to read file X, allowed to write to directory Y]

In the proposed WLCG AAI there is the possibility to configure both types of authorisation model. A compromise somewhere in the middle is also possible, e.g. including information that identifies the user plus capabilities rather than (or in addition to) roles and groups. The WLCG WG is interested to know which model best fits your workflows.
Token renewal
In OAuth2.0 Token based authorisation, two tokens are typically issued to the relying party for authorisation throughout long-lived workflows. An access token (that grants access to a resource) and a refresh token (that can be used to generate more access tokens). Best practices recommend issuing short lived access tokens (e.g. 20 minutes) and a longer refresh token - when access beyond the short lifetime is required the refresh token can be sent to the Token Issuer,

validated and a new access token generated. In this way access tokens can be used without validation by the target resource since the lifetime is sufficiently low to pose minimal risk to security. The process of refreshing a token requires a round trip to the token issuer. The WLCG Authorisation WG would like to understand at which points an access token would be required, this is in order to estimate workload. It is also necessary to understand the maximum lifetime of a job. The following is a list of anticipated points when the access token is required.

Access Token Required

Job is submitted

Data is fetched from storage

Output is returned to user

...

Current status

Two pilot solutions are currently under development, they will be analysed in the December pre-GDB to assess their compliance with the identified WLCG Requirements [https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG#Reference_Documents (https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG#Reference_Documents)]
A WLCG Token Schema is under finalisation. This is being done in collaboration with relevant community members already deploying token based solutions, in the hope that we can converge on a WLCG profile for interoperability.

Resources

Twiki https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG

(https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG)

Beyond X.509: Token-based Authentication and Authorization for HEP (CHEP Plenary)

https://indico.cern.ch/event/587955/contributions/3012583/

(https://indico.cern.ch/event/587955/contributions/3012583/)