

WLCG Authorisation

from X.509 to Tokens

Brian Bockelman⁶, Andrea Ceccanti³, Ian Collier⁴, Linda Cornwall⁴, Thomas Dack⁴, Jaroslav Guenther¹, Mario Lassnig¹, Maarten Litmaath¹, Paul Millar², Mischa Sallé⁵, Hannah Short^{1,}, Jeny Teheran⁷, and Romain Wartel¹*

¹European Organization for Nuclear Research (CERN)

²Deutsches Elektronen-Synchrotron (DESY)

³INFN

⁴UKRI-STFC

⁵Nikhef

⁶Morgridge Institute for Research

⁷Fermi National Accelerator Laboratory

Abstract. The WLCG Authorisation Working Group formed in July 2017 with the objective to understand and meet the needs of a future-looking Authentication and Authorisation Infrastructure (AAI) for WLCG experiments. Much has changed since the early 2000s when X.509 certificates presented the most suitable choice for authorisation within the grid; progress in token based authorisation and identity federation has provided an interesting alternative with notable advantages in usability and compatibility with external (commercial) partners. The need for interoperability in this new model is paramount as infrastructures and research communities become increasingly interdependent. Over the past two years, the Working Group has made significant steps towards identifying a system to meet the technical needs highlighted by the community during staged requirements gathering activities. Enhancement work has been possible thanks to externally funded projects, allowing existing AAI solutions to be adapted to our needs. A cornerstone of the infrastructure is the reliance on a common token schema in line with evolving standards and best practices, allowing for maximum compatibility and easy cooperation with peer infrastructures and services. We present the work of the group and an analysis of the anticipated changes in authorisation model by moving from X.509 to token based authorisation. A concrete example of token integration in Rucio is presented.

1 Introduction

This paper describes ongoing work by the WLCG Authorisation Working Group [2] to transition the Worldwide LHC Computing Grid's authorisation model from X.509 Certificates [10] and certificate proxies [11] to Tokens. The Tokens are defined as JSON Web Tokens [14], to be provisioned over OpenID Connect [13] and OAuth2 [12] workflows.

*e-mail: hannah.short@cern.ch

1.1 Motivation

When X.509 was chosen in the early 2000s, there was no mature alternative. The Globus Toolkit [20], chosen by WLCG, supported X.509 and provided a functional solution for distributed authentication and authorisation, when coupled with policies controlled by the Interoperable Global Trust Federation (IGTF) [21]. Despite providing a solution, certificate management proved a steep learning curve for researchers and the authorisation model adopted (although sufficient at the time) is outdated by today's standards regarding the protection of privacy of user data.

In the meantime, alternative authentication and authorisation technologies have emerged and X.509 is no longer considered the most suitable option. It is increasingly common that many off-the-shelf software choices and commercially hosted services offer inbuilt options for federated authentication, using standards such as SAML or OpenID Connect (OIDC) [13]. Integrating grid software with such third party systems has therefore become progressively more difficult, since a lot of X.509 developments are deeply integrated in the grid software stack. Users are now well accustomed to web-based authorisation workflows that use OAuth2 [12] or OIDC to delegate access rights, meaning that adoption of such technologies would present a more user-friendly experience. These factors came together to build a strong motivation for redesigning the WLCG authorisation model based on JSON Web Tokens over OIDC and OAuth2.

1.2 WLCG Authorisation Working Group

The WLCG Authorisation Working Group was formed in the summer of 2017, at a time when multiple activities were independently beginning to seriously consider token based authorisation. SciTokens [1] had produced a working system in the US focusing on capability-based authorization and an increasing number of efforts were ongoing in Europe to enable token based workflows, in particular in the context of the INDIGO DataCloud project [3] and EGI [4]. Experts from multiple domains came together in this group to chart a path towards token based authorisation for WLCG. Work to enhance software was supported by several European Commission Projects: EOSC-Hub [5], EOSC Pilot [6] and AARC2 [7]. The group's objective was to understand and meet the requirements of a future-looking authorisation service for WLCG experiments. In particular, effort focused on addressing usability, i.e. removing the need for users to manage certificates themselves and enable more intuitive authentication methods. A second objective was to simplify integration with third party services and software by adopting industry standards. A significant effort has been made to be as inclusive as possible, by including resource and service providers during the construction of the schema and the testing phase, as well as making all outputs available to the wider community such as through the FIM4R [8] workshop at Fermilab in September 2019. Documentation and a list of meetings are maintained on the WLCG Authorisation Working Group Twiki [2].

2 Milestones

Several key milestones have been achieved since the WLCG Authorisation Working Group was formed. In particular, a competitive study of two pilot solutions was made, resulting in an identified technical solution (INDIGO IAM [19]) in March 2019. A schema was negotiated between multiple stakeholders over the course of a year and was published in September 2019.

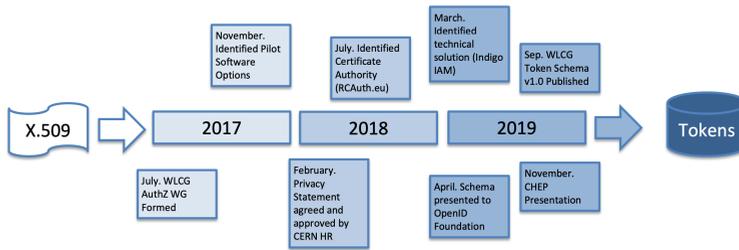


Figure 1. Key milestones achieved between 2017 and 2019 towards a Token based authorisation scheme. Note, it is not implied that the final goal has been reached; much additional effort is required to fully transition from X.509 to Tokens.

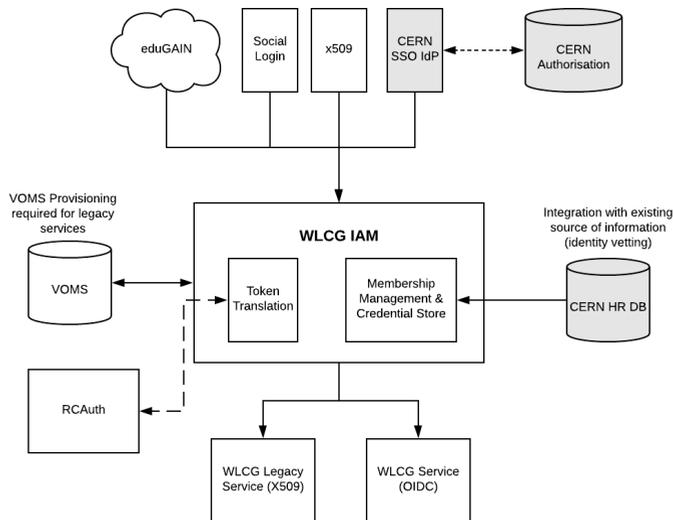


Figure 2. Design of the WLCG Token Authentication and Authorisation Infrastructure. Grey boxes represent configurable components, in this case tuned for the CERN environment.

3 Technical Design

The INDIGO IAM [19] software was chosen as the core of WLCG’s future, token based authentication and authorisation infrastructure. In Figure 2, we see the WLCG IAM (a WLCG Instance of INDIGO IAM) in the centre, proxying multiple authentication channels, centralising authorisation, providing token translation to and from X.509, and serving as a uniform token issuer for downstream WLCG Services. The Token Translation is performed using the RCAuth.eu [22] online Certificate Authority and users can request a certificate from the WLCG IAM User Interface.

This architecture is in line with guidelines from the AARC Project [7], including the AARC Blueprint Architecture.

3.1 Usage Flow

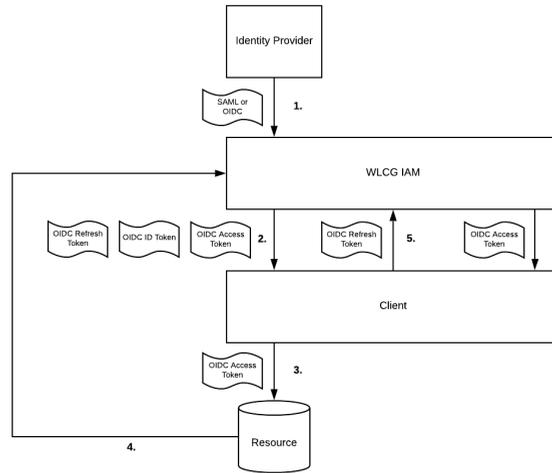


Figure 3. Anticipated Token Flow.

It is envisaged that there will be a small number of registered clients for WLCG and a much larger number of unregistered Resources. The user's authentication from their Identity Provider is technically decoupled from the tokens issued by the WLCG IAM. A user may authenticate with SAML [18], OAuth2 [12], OIDC [13], X.509 [10] or potentially any future protocol that is accepted by WLCG IAM. A more detailed description of IAM [29] and the work done to address WLCG requirements is given in a dedicated contribution to this journal [30]. Figure 3 describes the anticipated Token Workflow, the numbered steps are explained below:

1. A SAML [18] or OIDC [13] credential is sent from an Identity Provider to WLCG IAM[19]
2. WLCG IAM sends three OIDC tokens to a registered Client, the ID Token, Access Token and Refresh Token
3. The Access Token is used to authorise access to downstream Resources (such as storage)
4. The Resource validates the token against known trust roots. The trust roots may be cached in advance to avoid a high number of round trips to the token issuer.
5. When the Access Token expires, the Client may use the Refresh Token to request a newly valid Access Token from WLCG IAM
6. WLCG IAM returns an Access Token to the Client, which may repeat steps 3 to 6

4 Schema

The WLCG Common JWT Profiles document was published on September 25th, 2019 [9]. This document describes how WLCG users may use the available geographically distributed

resources without X.509 credentials. In this model, clients are issued with bearer tokens; these tokens are subsequently used to interact with resources. The tokens may contain authorization groups and/or capabilities, according to the preference of the Virtual Organisation (VO), applications and relying parties.

Wherever possible, the document builds on existing standards when describing profiles to support current and anticipated WLCG usage. In particular, three major technologies are identified as providing the basis for this system: OAuth2 [12] (RFC 6749 and RFC 6750), OpenID Connect [13] (OIDC) and JSON Web Tokens [14] (RFC 7519). Additionally, trust roots are established via, among others, the OpenID Discovery [15] mechanism or OAuth2 Authorization Server Metadata (RFC 8414 [16]). This document provides a profile for OAuth2 Access Tokens and OIDC ID Tokens.

| Common Claims | OIDC ID Token Claims | OAuth2 Access Token Claims |
|---|----------------------|---|
| sub | auth_time | scope (<i>inspired by OAuth token exchange draft</i>) |
| exp | Standard Claims [13] | |
| iss | | |
| acr | | |
| aud | | |
| iat | | |
| nbf | | |
| jti | | |
| eduperson_assurance (<i>REFEDS [17]</i>) | | |
| wlwg.ver (<i>WLCG [9]</i>) | | |
| wlwg.groups (<i>WLCG [9]</i>) | | |

Table 1. WLCG Token Claims. Claim origin is specified in brackets. Where unspecified, the origin is RFC7519 [14] or OpenID Connect core [13].

4.1 Authorisation

Two forms of Authorisation exist in parallel in the WLCG Common JWT Profile Schema; Group based Authorisation and Capability based Authorisation. Groups are semantically equivalent to existing VOMS groups, i.e. a group contains a list of users. A capability is the ability to perform an action, optionally at a specific path, e.g. "*the bearer of this token is authorised to write to /data*". Extended examples are included in the Schema [9].

5 WLCG Token Issuer

In late 2019, a WLCG Token Issuer was set up at the Istituto Nazionale di Fisica Nucleare (INFN) to provide a stable platform against which software enhancement could be tested. The Token Issuer was made available at <https://wlwg.cloud.cnaf.infn.it/>, primarily to facilitate tests planned by the WLCG DOMA Working Group. Clients are able to register to receive authentication tokens from the issuer over OIDC, and a web based platform is available to visualise tokens and their content. A VOMS attribute authority linked to this IAM instance has also been deployed, to allow users to get VOMS attribute certificates for the wlwg VO.

The objective is to demonstrate IAM interoperability with the current X.509 WLCG AAI and enable a gradual transition to tokens.

6 WLCG Tokens working example: Rucio

Rucio [27] is one of the main Data Management tools used by WLCG experiments and is a key service to be enhanced to accept WLCG Tokens. During 2019, the Rucio authentication and authorization mechanism was extended to support JSON Web tokens using the Open ID Connect protocol (following the OAuth 2.0 specifications). The implementation is based on the OIDC certified [23] *pyoidc* [24] library and follows the WLCG specifications [9] (subject to certain changes, feedback from which will be fed back into the next version of the WLCG Token Profile specification). To perform operations with Rucio *1.22.0.dev3* a user can now log in via the authorization code flow, which has been implemented in both the Rucio WebUI and Rucio command line interface (CLI). Several CLI login workflows are supported - including username and password, or session validation using a browser - to fulfil a variety of researcher use cases. The Rucio REST API has been configured as a OAuth2 protected resource and accepts tokens from trusted OIDC providers (including the WLCG Token Issuer discussed in Section 5).

6.1 Rucio and Downstream Storage

Tests have been successfully performed to demonstrate Token based authorization from Rucio to the File Transfer Service (FTS) [28]. When a Rucio user requests a file transfer, Rucio can use a JWT issued for a user or a token issued to the Rucio Admin user (synonymous with pilot jobs). In order to allow fine grained access control downstream (Rucio → FTS3), token exchange and token refresh grant flows were also implemented. The first functional tests of a third party copy were performed (Rucio → FTS3 → dCache) and a version *1.22.0.dev3* of Rucio supporting all these new features is currently deployed on the DOMA instance [25]. Additional tests with other downstream storage elements are planned.

7 Conclusion

The WLCG Authorisation Working Group is well on the way to enabling Token based Authorisation for WLCG Workflows. This work has been possible thanks to the participation of multiple individuals and projects, who together have been able to identify a technical solution and define a token schema. Requirements from Virtual Organizations and other research infrastructures outside of WLCG have been included as a continuous cooperative process to guarantee interoperability. Preliminary token based workflows have been tested with many more set to come in 2020. Further efforts are required to complete the enhancement of middleware, and to ensure that the existing WLCG policies are adhered to in the years to come.

Acknowledgements

This work was achieved thanks to support from the following European Commission funded projects; EOSC-Hub [5], EOSC Pilot [6] and AARC2 [7]. EOSC-hub receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 777536. EOSCPilot.eu has received funding from the European Commission's Horizon 2020 research and innovation programme under the Grant Agreement no 739563. AARC2 is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement number 730941. This material is based upon work supported by the National Science Foundation under Grant No. 1836650.

References

- [1] Withers, Alex, Bockelman, Brian, Weitzel, Derek, Brown, Duncan, Gaynor, Jeff, Basney, Jim, Tannenbaum, Todd, and Miller, Zach. SciTokens: Capability-Based Secure Access to Remote Scientific Data. Proceedings of PEARC '18. July 18. <https://doi.org/10.1145/3219104.3219135>
- [2] WLCG Authorisation Working Group Twiki, <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG>
- [3] INDIGO DataCloud, <https://www.indigo-datacloud.eu>
- [4] EGI, <https://www.egi.eu>
- [5] EOSC Hub, <https://www.eosc-hub.eu>
- [6] EOSC Pilot, <https://eoscpilot.eu>
- [7] Authentication and Authorisation for Research and Collaboration (AARC), <https://aarc-project.eu>
- [8] Federated Identity Management for Research and Collaboration, <https://fim4r.org>
- [9] Altunay, Mine, et al. (2019, September 25). WLCG Common JWT Profiles (Version 1.0). Zenodo. <http://doi.org/10.5281/zenodo.3460258>
- [10] Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, RFC 5280, <https://tools.ietf.org/html/rfc5280>
- [11] Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile, RFC 3820, <https://tools.ietf.org/html/rfc3820>
- [12] OAuth2.0, RFC 6749, <https://tools.ietf.org/html/rfc6749>
- [13] OpenID Connect, https://openid.net/specs/openid-connect-core-1_0.html
- [14] JSON Web Tokens, RFC 7519, <https://tools.ietf.org/html/rfc7519>
- [15] OpenID Connect Discovery, https://openid.net/specs/openid-connect-discovery-1_0.html
- [16] OAuth 2.0 Authorization Server Metadata, RFC 8414, <https://tools.ietf.org/html/rfc8414>
- [17] Research and Education Federations Group, REFEDS, <https://refeds.org>
- [18] Security Assertion Markup Language, SAML, RFC 7522, <https://tools.ietf.org/html/rfc7522>
- [19] INDIGO Identity and Access Management (IAM), <https://indigo-iam.github.io/docs>
- [20] Globus Toolkit, <http://toolkit.globus.org>
- [21] Interoperable Global Trust Federation (IGTF), <https://www.igtf.net>
- [22] RCauth, <https://rcauth.eu>
- [23] Certified OpenID Connect Implementations <https://openid.net/developers/certified/>
- [24] Python OIDC library <https://github.com/OpenIDC/pyoidc>
- [25] Rucio DOMA instance <https://rucio-doma.cern.ch>
- [26] System for Cross-domain Identity Management <http://www.simplecloud.info/>
- [27] Rucio - Scientific Data Management <https://rucio.cern.ch>
- [28] File Transfer Service <http://information-technology.web.cern.ch/services/file-transfer>
- [29] Andrea Ceccanti, Enrico Vianello, Marco Caberletti, Francesco Giacomini. Beyond X.509: token-based authentication and authorization for HEP. EPJ Web Conf., 214 (2019) 09002 <https://doi.org/10.1051/epjconf/201921409002>
- [30] Andrea Ceccanti, Enrico Vianello, Marco Francesco Giacomini. Beyond X.509: token-based authentication and authorization in practice. Submitted for publication in this journal