EPJ Web of Conferences **337**, 01288 (2025) https://doi.org/10.1051/epjconf/202533701288
CHEP 2024

Supporting medium/small-sized experiments in the transition from X.509 to JWTs

Carmelo Pellegrino[1,*], *Aksieniia* Shtimmerman[1,**], *Alessandro* Pascolini[1], *Matteo* Barbetti[1], *Carmen* Giugliano[1], *Daniele* Lattanzio[1], *Lucia* Morganti[1], and *Andrea* Rendina[1]

[1]INFN-CNAF, Viale Carlo Berti-Pichat 6/2, 40127, Bologna, Italy

Abstract. X.509 certificates and VOMS proxies are still widely used by various scientific communities for authentication and authorization (authN/Z) in Grid Storage and Computing Elements. Although this has contributed to improve the scientific collaboration worldwide, X.509 authN/Z comes with some interoperability issues with modern Cloud-based tools and services.

The Grid computing communities have decided to migrate to token-based authentication, a new web technology that has proved to be flexible and secure.

The model being recently adopted by the communities is based on industrial standards such as OAuth2 and OpenID-Connect and exploits JSON Web Tokens (JWT): a compact way to securely transmit information as JSON objects.

JWT are usually short-lived and provide fine-grained authorization, based on "scopes", to perform specific actions. These scopes are embedded into the token and are specified during the request procedure so they last only until token expiration time. Scopes can be requested based on user groups and permission thus providing the possibility of restricting a group to perform only a subset of actions.

These characteristics make up to a more secure alternative to X.509 proxies.

Being largely used in industries, JWTs are also easily integrated in services not specifically developed for the scientific community, such as calendars, Sync and Share services, collaborative software development platforms, and more.

As such, JWTs suit the many heterogeneous demands of Grid communities and some of them already started the transition in 2022.

In the Italian WLCG Tier-1, located in Bologna and managed by INFN - CNAF, several computing resources are hosted and made available to scientific collaborations in the fields of High-Energy Physics, Astroparticle Physics, Gravitational Waves, Nuclear Physics and many others.

Although LHC experiments at CERN are the main users of CNAF resources, many other communities and experiments are being supported in their computing activities.

While the main LHC experiments have already planned their own transition from X.509 to token-based authN/Z, many medium/small-sized collaborations struggle to put effort into it.

The Tier-1 User Support unit has the duty of guiding users towards efficient and modern computing techniques and workflows involving data and computing resources access.

*e-mail: carmelo.pellegrino@cnaf.infn.it
**e-mail: aksieniia.shtimmerman@cnaf.infn.it

As such, the User Support group is playing a central role in preparing documentation, tools and services to ease the transition from `X.509` to JWTs. The foreseen support strategy and the related tools will be presented. Future workflow plans in view of the complete transition will also be presented.

1 Introduction

The Worldwide LHC Computing Grid (WLCG) [1] is a collaboration involving around 170 computing centres in more than 40 countries. It provides computing resources to store, distribute and analyse the ~200 PB of data expected every year from the Large Hadron Collider (LHC) experiments at CERN. WLCG employs a *Tier-based* computing architecture [2], in which computing centres are organised in a hierarchy. The Tier-0, the main datacentre, is hosted at CERN, physically close to the detectors, and is directly connected via a high-capacity network [3] to the 10 Tier-1s. A primary copy of the raw data is sent from the Tier-0 to the Tier-1s for storage preservation and analysis purposes. Thereafter, data are distributed to the over 160 Tier-2s of the Grid, for further analysis.

In this section, the INFN Tier-1, managed by CNAF, and the scientific communities being supported are described.

In the following sections, the outcome of about 6 years of experience and effort in supporting the transition from `X.509` to JWTs is presented.

1.1 INFN Tier-1

The Italian WLCG Tier-1 site is the largest INFN datacentre. It is located in Bologna and it is hosted since 2003 at CNAF [4], the National Centre of INFN, that is the Italian National Institute for Nuclear Physics. CNAF is devoted to research and development on information and communication technologies. The Centre provides the computing resources, support and IT services needed for data storage and distribution, data processing and analysis and Monte Carlo production to over 70 scientific communities.

INFN-T1 supports either *local*[1], Grid and Cloud computing. The computing farm is composed of ~2000 computing nodes providing ~60.000 cores accessed via a batch system, and another partition running Cloud services that are accessible via either IaaS or PaaS solutions. Storage resources consist of ~70 PB of disk space served via a distributed filesystem (GPFS), and ~130 PB of tape space. The latter is mainly used as a long-term storage solution for archival data.

Although the LHC experiments are the main users of the datacentre, CNAF represents the principal computing facility for many communities covering a broad scientific landscape, not only from the Physics field but also for Biomedical Health and Life science, requiring ISO-certified isolated computing zones.

1.2 Supported scientific communities

As also reported in section 1.1, over 70 scientific communities are served at INFN Tier-1.

[1] *local users* are those who connect to CNAF servers via SSH to manage their work either accessing the computing farm via a batch system and the storage via POSIX access to data.

While WLCG communities have a well esablished and structured computing model, smaller ones - ranging from ten to two hundreds persons - often lack the required effort on computing-related tasks within the community. For these, the role of the User Support team becomes increasingly relevant as a guide towards the efficient end effective use of the computing resources.

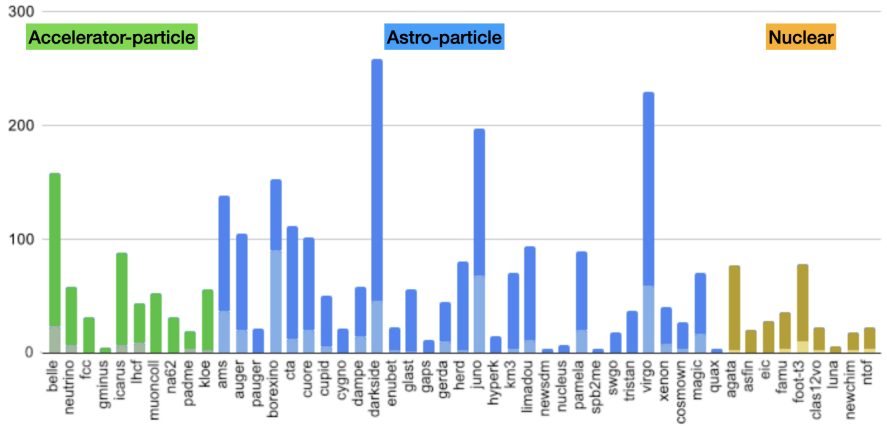

Figure 1. Number of personal CNAF accounts per scientific community. Only medium- and small-sized communities are here reported. The "pale" segments correspond to the fraction of inactive users, i.e. people who have lost access to the Centre.

To give a measure of the size of medium/small-sized scientific communities, in Figure 1 the number of users per community is reported. The pale segments correspond to the fraction of inactive users.

2 JWTs in a nutshell

In this section, we describe what a *JSON Web Token* (JWT) [5] is in analogy with `X.509` VOMS proxy certificates from an end-user perspective[2].

Typical JWTs that end-users encounter are long strings composed by three parts, each separated by a dot ("`.`"):

1. a header, represented by a `JSON` document coded in `base64` containing information about the algorithm used by the token issuer to cryptographically sign the token and the name of key used;

2. a payload, also represented by a `JSON` document coded in `base64`, containing several information regarding, for example, the identity of the user and its group membership, the scopes for which the token has been generated, the service the token is intended to be used on, and an expiration date. In Listing 1 an example of the content of a JWT payload is shown;

3. a cyptographic signature to provide a sufficient level of trust that the token is authentic, i.e. has not been tampered by an attacker.

[2]It is not the intent of this article to provide a comprehensive description of JWTs and protocols such as OpenID-Connect and OAuth2.

While `X.509` has limited support for fine-grained authorization, a JWT allows to specify very precisely which computing resource to interact with and how.

In this regards, WLCG defined a *Common JWT Profile* [6], that describes how users may access the present Grid resources without `X.509` credentials. Specifically, it defines the possible contents of the JWT payload in terms of authorization groups and capabilities. For example, in Listing 1 the WLCG JWT profile is enabled. Specifically, the token reports the user's group membership in the `wlcg.groups` claim as a list of strings and listed `compute.*` scopes allows the user to interact with a Grid Compute Element by submitting and reading batch jobs but not to remove or modify them.

Authorization is performed by the contacted service based on the information contained in the payload. This was also true for `X.509` VOMS proxies, with the difference that finer-grained authorization can be put in place, greatly increasing security.

Moreover, JWTs, in contrast to VOMS proxies, are a widely adopted industrial standard, which turns in favour of an improved interoperability between software tools and services, either proprietary and open-source, available on the market and Grid-specific middleware. Because of this, it is possible, for example, to transfer data from one Grid Storage Element to another and log into a Jupyter Hub instance using the same login to a token issuer.

Finally, JWTs are typically shorter-lived ($\leq$ 1 hour) with respect to `X.509` VOMS proxies ($\geq$ 12 hours). Also this is done in the intent to improve the security.

```
{
  "sub": "8c08e83b-5ebb-4f92-a362-c3a2a5d3ed2f",
  "iss": "https://iam-t1-computing.cloud.cnaf.infn.it/",
  "preferred_username": "budda",
  "client_id": "68a505b6-c95e-4126-8f72-019817eb6817",
  "wlcg.ver": "1.0",
  "aud": "https://wlcg.cern.ch/jwt/v1/any",
  "nbf": 1740755515,
  "scope": "openid compute.create offline_access profile
     compute.read wlcg wlcg.groups",
  "name": "Carmelo Pellegrino",
  "exp": 1740759115,
  "iat": 1740755515,
  "jti": "1abd9eb3-87f5-4b95-bf54-a0eb6d04dde3",
  "wlcg.groups": [
    "/dune",
    "/eee",
    "/foot",
    "/kloe",
    "/km3net",
    "/muone",
    "/ntof",
    "/padme",
    "/pauger",
    "/quax",
    "/swgo",
    "/user-support"
  ]
}
```

Listing 1. The payload of an actual JWT.

3 Software tools supporting the transition

In this section, the tools deployed and made available to the INFN Tier-1 users are presented.

3.1 Data transfer scripts

A recurring problem is that of transferring large amount of data files from a local storage to the INFN Tier-1 disk or tape.

While large communities heavily rely on orchestrators like RUCIO [7] or DIRAC [8] and their data catalogues to ensure correct data transfers and replicas, smaller communities often lack the effort to manage such complex tools. Also, often data transfers are performed on the spot, for example once a year during a beam-time period. In many cases, important data are first collected on some temporary media that is not provided with a Storage Element that exposes these data on the network.

For these reasons, two scripts providing a client interface towards the fleet of StoRM-WebDAV [9] servers deployed as Grid Storage Elements at INFN Tier-1 have been developed and provided to the users.

3.1.1 copy.sh

`copy.sh` [10] is a simple *Bourne Again Shell* [11] script with minimal dependencies. Indeed, apart from command line tools like coreutils, grep, sed and others available on any modern Linux distribution, it employes either `cURL` [12] or `gfal-copy` [13] to perform the actual transfer, `oidc-agent` [14], and `GNU Parallel` [15] to handle multiple parallel transfer and recovery from eventual interruptions.

The typical issues with transferring a large amount of files are that token expiration is likely reached before the end of the operations and, due to the lack of proper token renewal mechanism in `gfal-copy`, and that to be able to fully exploit the available network bandwidth multiple streams need to be orchestrated.

The script needs to be configured via shell variables and accepts a local POSIX path a command line input.

Files are recursively searched below the specified path and a *filelist* file containing the list of files to transfer is computed.

GNU parallel takes care of executing a BASH function called `copy` whose duty is to perform one single transfer.

The `copy` function first retrieves a valid JWT calling the `oidc-token` command. Then the local path of the file is translated in a valid WebDAV URI by *urlencode*-ing the file name and parent folders if necessary. Finally, it performs the file transfer by using `cURL` or `gfal-copy`. If `cURL` is used, it also recursively creates the needed parent directories.

Several versions of the script exist, each in a dedicated *git branch*, that comply with the specific use-case they served or are still serving.

It was originally prepared to transfer 1 PB of data of the MUonE experiment from CERN to CNAF, later aborted.

It has been successfully used to transfer ~70 TB of user files (~ $3 \cdot 10^6$ files) from CC-Lyon disk to CNAF disk by the Pierre Auger Collaboration and is going to be used to transfer 1 PB

of raw data and Monte Carlo productions from CC-Lyon tape to CNAF tape. In this iteration, the script also recorded the POSIX filesystem timestamp - `mtime` and `atime` - that have been later restored at the destination to preserve the history of the > 20 years long activity of the Collaboration.

3.1.2 xfer-oidc

`xfer-oidc` [16] is a script written in `POSIX` shell language. It provides all the functionalities to securely transfer a large amount of files to or from a WebDAV Storage Element, but with a really minimal set of dependencies, which makes it useful in "*extreme conditions*" such as running on one of the following environments:

- an immutable operating system
- unusual operaring system, e.g. based on BSD or AIX
- unusual CPU architecture, e.g. Power or RISC-V

where `oidc-agent` or `gfal2` are missing and for which software support is scarcer than for Linux.

It allows to register and manage an OpenID-Connect client to obtain JWTs from a token issuer such as INDIGO-IAM [17], download IGTF CA certificates, and perform basic operations on the remote WebDAV server such as transfer a file, create a folder and show the Storage Element content.

It has been used to transfer to CNAF ~ 15 TB of data of the DarkSide collaboration from a stand-alone Network Attached Storage system, located in Brazil, whose operating system is based on BSD and doesn't allow installing additional software.

3.2 HTCondor mapping plugin

One of the first issues raised with the introduction of `SCITOKENS` [18][3] is that until version 10.5 of HTCondor[4] [19] it was not possible to dynamically map a user identity, defined as the pair (*token issuer*, *token subject*). In other terms, a token issuer could identify one single Virtual Organization.

The `scitoken-mappings` [20] toolset allows the INFN Tier-1 admins to configure HTCondor to perform dynamic SCITOKEN mapping on specific pool of POSIX accounts. It employes the *SCITOKEN_PLUGINS*[21] feature of HTCondor to determine the correct pool based on the content of the `wlcg.groups` token claim.

In the first implementation, only a simple *sqlite3* database was used to persist mapping choices and configure the pool of POSIX accounts corresponding to each virtual organization.

The present implementation relies on a remote *PostGRES* database that is concurrently accessed by each of the six production INFN Tier-1 Compute Entrypoints.

All the code is written in the BASH scripting language.

3.3 Management of an instance of INDIGO-IAM

INDIGO-IAM (Identity and Access Management) [17] is the token issuer chosen by WLCG.

Typically, former VOMS Virtual Organizations have switched to using both the legacy VOMS server and a dedicated INDIGO-IAM instance that, when equipped with the VOMS-AA side-service, is capable of providing both JWTs and VOMS proxies.

[3] *SCITOKENS* are the WLCG JWTs in the HTCondor jargon

[4] HTCondor is a middleware that implements a batch system. It powers the scientific-computing farm of the WLCG INFN Tier-1.

Are an example of this, alongside many more, the LHC and the JUNO, KM3NeT and Belle-II collaborations.

In order to support medium- and small-sized scientific communities which had no legacy VOMS in the past or recently birth, a catch-all INDIGO-IAM instance, called *iam-t1-computing*[5], has been created and is managed by the INFN Tier-1 User Support team.

It gives access to storage and computing resources, also located outside CNAF, to more than one hundred users who are members of at least one of the 45 groups.

3.4 Management of an instance of MyToken

A compelling issue introduced with the adoption of JWTs for accessing the Grid Storage Elements is that, due to the short life of INDIGO-IAM tokens of typically one hour, longer-lasting Grid batch jobs may be unable to upload the outputs.

Several mechanisms to renew or obtain a new JWT have been investigated as reported in [22].

We chose to adopt myToken [23] as the solution. The choice has been driven by mostly two factors: myToken is a generic solution, not tied to a specific implentation of the batch system; it is easy to adopt for users already used to `oidc-agent`.

Beyon its ergonomicity, myToken also offers several security features to allow further access restriction to tokens, for example allowing to obtain JWTs only from a specific network segment or country.

We decided to self-host a myToken instance, which the time being, it is in private beta by two communities. Both have successfully integrated this test instance in their work flows and are using it to seemlessly access from batch jobs data stored in S3 buckets on the INFN-Cloud [24] platform based on MinIO and on the INFN Tier-1 StoRM-WebDAV Storage Elements.

References

[1] *The Worldwide LHC Computing Grid*, https://wlcg.web.cern.ch/

[2] *The LHC's worldwide computer*, https://cerncourier.com/a/the-lhcs-worldwide-computer/

[3] *LHCOPN - Large Hadron Collider Optical Private Network* , https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome

[4] *The INFN-CNAF website*, https://www.cnaf.infn.it/

[5] *The JSON Web Token RFC*, https://tools.ietf.org/rfc/rfc7519.txt, https://doi.org/10.17487/RFC7519

[6] *WLCG Common JWT Profiles*, M. Altunay, et al., 2019, https://doi.org/10.5281/zenodo.3460258

[7] *Rucio: Scientific Data Management*, Barisits, M., Beermann, T., Berghaus, F. et al., Comput Softw Big Sci 3, 11 (2019). https://doi.org/10.1007/s41781-019-0026-3

[8] *The DIRAC interware: current, upcoming and planned capabilities and technologies*, F. Stagni, et al., EPJ Web Conf. 245 03035 (2020), https://doi.org/10.1051/epjconf/202024503035

[9] *The StoRM WebDAV service* source code, https://github.com/italiangrid/storm-webdav

[10] *The copy.sh source code* https://baltig.infn.it/exp-supp/copy.sh

[11] *The Bourne Again SHell* website, https://www.gnu.org/software/bash/

[12] *The cURL website*, https://curl.se/

[5] https://iam-t1-computing.cloud.cnaf.infn.it

[13] *Gfal2 Documentation*, https://dmc-docs.web.cern.ch/dmc-docs/gfal2/gfal2.html
[14] *The oidc-agent Documentation*, https://indigo-dc.gitbook.io/oidc-agent/
[15] O. Tange (2018): GNU Parallel 2018, Mar 2018, ISBN 9781387509881, DOI https://doi.org/10.5281/zenodo.1146014
[16] *The xfer-oidc source code*, https://baltig.infn.it/cnaf-user-support/xfer-oidc
[17] *The INDIGO-IAM Documentation*, https://indigo-iam.github.io/v/current/
[18] *SciTokens: Capability-Based Secure Access to Remote Scientific Data*, A. Withers, et al., PEARC '18: Proceedings of the Practice and Experience on Advanced Research Computing: Seamless Creativity, 2018, https://doi.org/10.1145/3219104.3219135
[19] *The HTCondor* website, https://htcondor.org/
[20] *The scitoken-mapping source code*, https://baltig.infn.it/exp-supp/scitokens-mapping
[21] *The HTCondor admin manual, SCITOKENS MAPPING*, https://htcondor.readthedocs.io/en/latest/admin-manual/security.html#scitokens-mapping-plugins
[22] *Support to experiments in the transition from X.509 authN/Z to SciTokens*, International Symposium on Grids and Clouds (ISGC) 2024, https://indico4.twgrid.org/event/33/contributions/1347/
[23] *mytoken - OpenID Connect Tokens for Long-term Authorization*, G. Zachmann, PhD Thesis, Scientific Computing Center (SCC), Karlsruher Institut für Technologie (KIT), https://doi.org/10.5445/IR/1000134712
[24] *The INFN-Cloud* website, https://www.cloud.infn.it/