EPJ Web of Conferences **245**, 03012 (2020) https://doi.org/10.1051/epjconf/202024503012
CHEP 2019

CERN's Identity and Access Management

A journey to Open Source

Asier Aguado Corman[1], *Daniel* Fernández Rodríguez[1], *Maria V.* Georgiou[1], *Julien* Rische[1], *Ioan Cristian* Schuszter[1], *Hannah* Short[1,*], and *Paolo* Tedesco[1,**]

[1]European Organization for Nuclear Research (CERN)

Abstract. Until recently, CERN had been considered eligible for academic pricing of Microsoft products. Now, along with many other research institutes, CERN has been disqualified from this educational programme and faces a 20 fold increase in license costs. CERN's current Authentication and Authorization Infrastructure, dating from 2008, comprises multiple Microsoft services from the web Single-Sign-On to the Accounts Database. Replacing these core components is an opportunity to rebuild the CERN infrastructure using the latest technologies and concepts and to respond to evolving requirements of the community. It is also the appropriate moment to consider the alignment of CERN's and the Worldwide LHC Computing Grid's approaches to identity management, to create a more consistent environment for operators, developers and users. 2019 saw the launch of an Alpha version of CERN's next generation Authentication and Authorization Infrastructure, focusing on free and open source products and responding to the limitations experienced by the current system. We describe the new solution and focus on key changes.

1 Introduction

CERN's Authentication and Authorization Infrastructure (AAI) enables secure authentication of approximately 60,000 users to roughly 15,000 online services, ranging from scientific platforms to financial applications. CERN's users are authorized via their membership in over 60,000 groups which provide fine-grained access control, ensuring the protection of sensitive data and functionality. Group memberships change frequently, new services are continually rolled out and users come and go, resulting in a large scale, complex landscape for CERN's identity management.

In 2019, CERN IT rolled out an Alpha version of a new AAI. It represents a dramatic shift from the previous Kerberos [1] and Microsoft based system and brings significant advantages in usability for end-users and software maintainers, the protection of personal data and the convergence of services towards token based authorization.

Although the catalyst for urgent change was the announcement of a license fee price increase of Microsoft products in 2018, a strong case for overhauling CERN's AAI had developed independently over previous years. As of 2019, CERN and the Worldwide LHC Computing Grid (WLCG [2]) operate separate AAIs that are technically decoupled, resulting

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

in researchers managing multiple credentials. For the last 20 years, CERN has operated a Kerberos based system, whereas the WLCG opted for X.509 [3] Certificates. As alternative technologies have matured, both CERN and WLCG have become interested in moving to token based authorization[1]. An opportunity has arisen to harmonise user access for WLCG and CERN based on a common protocol that has become industry standard, OAuth2.0 [4]. In parallel, an increased focus on Data Privacy, particularly driven by the European Union's General Data Protection Regulation (GDPR) [5], requires a new authorization model that safeguards users' private information.

Instead of identifying a like-for-like replacement for the previous software stack, the following principles were adopted during the design of CERN's new AAI: identify suitable alternatives based on use cases not products, prioritise free and open source software, stick to standards, contribute back to the community, and stay in line with users' and service's requirements. A modular architecture of predominantly Open Source components emerged, supporting common standards. Those working on the project contribute back to the open source repositories used, and engage in groups to share their knowledge, such as the Federated Identity Management for Research (FIM4R) Community [6].

A staged approach was taken to development, deployment and migration, see Figure 1. The Single-Sign-On (SSO), i.e. unified authentication to CERN's online services, was addressed first, to be followed by the Directory Services, i.e. account management. After migrating services to the new SSO, a process anticipated to take several years, changing the Directory Services should be largely transparent to users and many software maintainers.

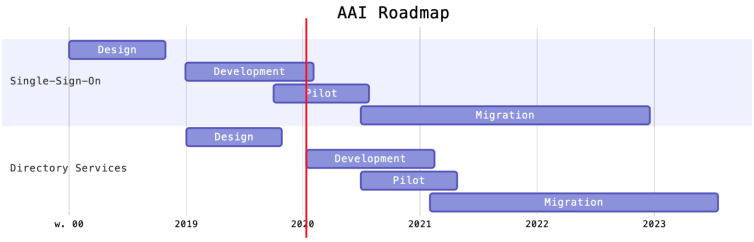

Figure 1. Expected timeline. Upper section describes Single-Sign-On, lower describes Directory Services (i.e. Accounts).

2 Previous Architecture

The architecture of the previous authentication and authorization infrastructure (see Figure 2) at CERN is centered around Active Directory [7], which acts as a central repository for user accounts, used for authentication, and groups, used for authorization. The creation of user accounts is automated through Microsoft Forefront Identity Manager (FIM). FIM is a service that periodically polls the CERN personnel database (Foundation) and applies a set of rules to enforce the CERN lifecycle policies for computing accounts. When a new member is entered in the CERN personnel database, FIM creates an account for the newcomer, and when a person's contract ends FIM blocks and deletes the account after a grace period. After its initial introduction, FIM has also been configured to manage the lifecycle of other computing resources so that automatic actions are taken when the owner's contract at CERN ends. A

[1] Authorization conveyed in signed tokens, in this case JSON Web Tokens (RFC 7519)

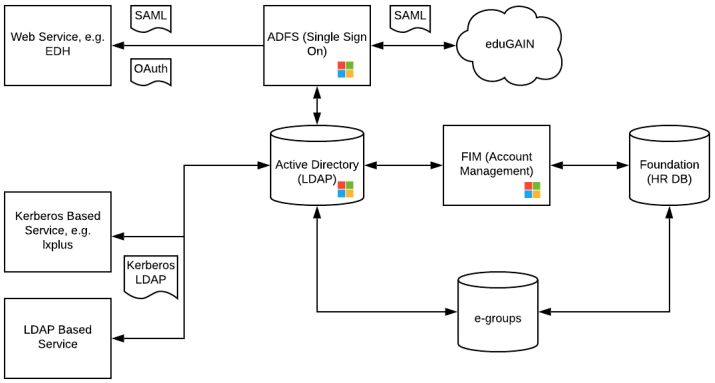

Figure 2. Previous, largely Microsoft-based architecture

custom application, called E-Groups, is the repository for groups information. Active Directory is accessed by Kerberos and LDAP based services to authenticate users and perform authorization decisions by querying groups information. SSO functionality is provided by Microsoft Active Directory Federation Services (ADFS [8]). Web applications use ADFS for users authentication and authorization.

3 Future Architecture

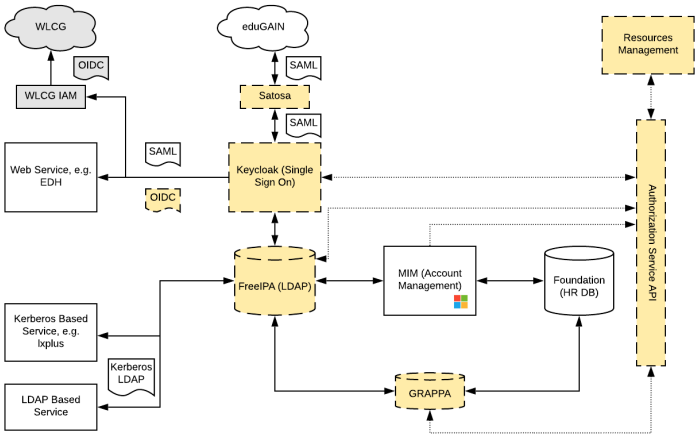

Figure 3. Future architecture, including many open source components. Administration portals and minor services are excluded for simplicity. Yellow boxes with dashed edges in Figure 3 represent components that have been added, replaced or changed significantly, and are explained in more detail below. Grey boxes represent the integration of WLCG services behind the new CERN SSO.

The future Identity and Access Management system replaces the high cost Microsoft components with suitable alternatives. FIM, the Forefront Identity Manager, is maintained

in the new system as it provides core functionality, has been found to be reliable and does not represent a large expense. All the components are managed via a REST API, the Authorization Service API, which takes care of application registration, group management, user management and additional functionality. Replacement components were identified based on several criteria, including their ability to achieve the scalability and performance required based on our knowledge of operating the previous system. We focused our search on open source tools with an active and frequent release cycle, backed by a responsive support channel where possible. We prioritised software licensed under a free and open source model to minimise any eventual vendor lock-in . To maximise compatibility with downstream systems, we selected software that supports common standards such as OpenID Connect, SAML2 and Kerberos, and that runs on GNU/Linux.

3.1 Components

3.1.1 Authorization Service API

The Authorization Service API is the core component for authorizing users in CERN's Identity and Access Management architecture. It provides a REST interface for applications and users to query and manage Applications, Identities, Groups, Roles and Resources. Although Keycloak (see Section 3.1.6) and other products also offer authorization features (e.g. Roles in Keycloak), we developed our own solution to achieve a loosely coupled architecture, where Authentication and Authorization providers could be independently replaced or upgraded in the future. The software has been developed in C#, using the .NET Core 2.2 framework. It has been recently upgraded to .NET Core 3.1. The Authorization Service contains a Workflows Engine which enables the creation of Rules linked to model objects. These Rules can be defined to trigger actions during the Create, Update and Delete operations of model objects, such as Identities or Groups. The actions are performed by running jobs in a scheduler. We decided to use Hangfire [9] as a job scheduler, as it is a native background processing library for .NET Core. The Authorization Service also keeps audit logs and a history of changes from all the operations performed in API calls using database triggers included into the Authorization Service database schema. The Authorization Service provides an Identity based authorization model, where Identities, and not the underlying Accounts, are members of authorization groups. Application administrators must define a Role for every privilege level in their application, which are linked with Identities through Groups. Another concept that has been introduced to the Authorization model is Levels of Assurance (LoA). We assign an LoA to every account provider: CERN has the highest, while a social provider with public registration has the lowest. An application Role will have an LoA assigned to it, e.g. "administrator" role with "CERN accounts", or "user" role with "social accounts". The associations between Account, Identity , Group, Role and Application are represented in Figure 4.

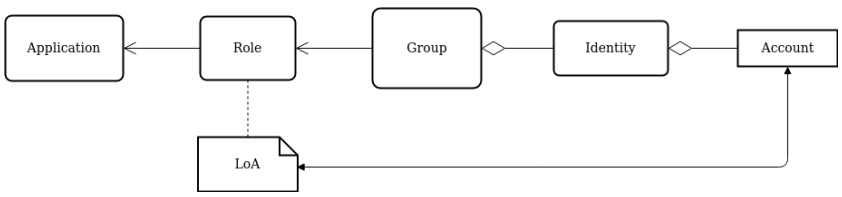

Figure 4. Authorization model associations

3.1.2 Resource management

Resources Management, a component of the Authorization Service, provides external services with updates regarding the lifecycle state of a resource. We define a resource as any kind of metadata that can be owned by an Identity, for example: websites, groups, different accounts. Services may define the fields a resource should have using a JSON schema [10] definition. Periodically, resource creations, updates or deletions are propagated to the external systems integrated with a specific type of resource. The actions performed by these systems are determined by the resource Service owner; Services that want to benefit from the resource lifecycle need only implement a simple CRUD API that is called when changes happen. Resources go through the following common lifecycle:

- An identity is subscribed to a GRAPPA group (see Section 3.1.5) and is allowed to create resources of a specific type, either personal or official. The identity creates a resource.
- When the identity leaves the role that grants access to the creation of a specific resources (typically when they leave the organisation), personal resources are deleted and official resources are typically assigned to the identity's supervisor.

3.1.3 MIM and FIM

Microsoft Identity Manager (MIM) [11] is the new version of Forefront Identity Manager (FIM) used in the previous system. MIM and FIM are synchronization engines, that can read data from a system, such as the CERN personnel database, apply a set of configured rules and export the processed data to another system, such as the Authorization Service. Dedicated connector processes, called Management Agents, are responsible for the communication with different data sources (Oracle databases, LDAP stores, REST APIs etc). For scalability and simplicity of configuration, multiple instances are used. One MIM instance is dedicated to synchronise CERN personnel data to the Authorization Service. Other instances are dedicated to synchronising data from the Authorization Service to external systems such as FreeIPA (see Section 3.1.4). Other external systems will include custom APIs to manage computing resources (discussed in Section 3.1.2).

3.1.4 FreeIPA

FreeIPA is an open-source solution for identity and authentication management (IdM) that provides centralized authentication, authorization, and account information by storing data about objects (e.g., users, groups, hosts, services) necessary to manage the security aspects of a network of computers [12]. It is built on top of open-source components and standard protocols; MIT Kerberos [1], LDAP 389 Directory Server [13] and Dogtag [14]. Currently, CERN is using solely Active Directory and our goal is to use FreeIPA as the primary identity management solution. This requires the migration of all the fore-mentioned objects to the new IdM, which include more than 140,000 hosts, approximately 60,000 users, 60,000 groups and 1000 services managed by Kerberos. The migration needs to be thoroughly tested and executed progressively as it is essential that it causes as few issues as possible to the running services and the working users. Due to the complexity of the migration, certain key features will be kept from the previous system, namely:

- Services and clients domain names unmodified (e.g. `lxplus.cern.ch`)
- User accounts and user groups will be replicated in FreeIPA with an updated LDAP schema. A subset of them will be added to AD only after fulfilling specific criteria.

- Full Kerberos authentication support for MIT Kerberos clients (but with mandatory configuration update)
- Host initial credentials installation (e.g., Kerberos keytab, private key) automatic for containers/VMs and manual for others (e.g, `cern-get-keytab`)

A great challenge posed are the Windows clients where Microsoft's operating system is deeply integrated with their identity manager, Active Directory. As a consequence, Windows clients cannot integrate properly with an IdM that follows different architecture principles. This is particularly the case when interacting with services that are part of a different Kerberos realm on the same DNS domain (that's why for this migration we chose to avoid renaming the domain names of all services). Moreover, domain users' login on Windows and roaming profiles are not supported by FreeIPA, as they require a dedicated LDAP service. Hence, we decided that we would continue to support Active Directory for a limited set of objects (e.g., users, hosts) that have a critical need.

3.1.5 E-Groups Replacement, GRAPPA

E-Groups are the legacy system used across CERN to authorize users' access to applications, based on their membership in various groups. E-groups are available inside SSO tokens or by querying Active Directory and are heavily used at CERN. E-Groups authorizations work in a recursive manner, meaning that one e-group can be a member (or parent) of several other e-groups, which gives it much flexibility for a multitude of tasks. Two types of groups can be defined:

- *static groups* - managed by a person and, optionally, by another group of administrators. Members are added by the privileged users, either through the Web User Interface or via a SOAP API.
- *dynamic groups* - populated dynamically based on criteria defined by their creators. Certain criteria which relate to personal data are restricted and require explicit approval from higher management structures, such as age, gender or nationality, to name a few.

GRoups for APPlications Authorization (GRAPPA) replaces and improves E-Groups. It provides several solutions that work around the existing issues in the authorization infrastructure when coupled with the new authorization model in Figure 4. One challenge we faced was enabling the gradual migration of users from E-Groups to GRAPPA. Most of the current static and dynamic E-Groups had to be migrated and kept in sync in GRAPPA. For dynamic groups, an ElasticSearch [15] instance provides us with data from the HR database, allowing us to query for people and match them against the data present in the database. Periodic jobs ensure the synchronization of members by adding and removing them from the groups, according to the criteria defined in the group definition. The static groups and their members are synchronized from Active Directory to the new GRAPPA groups, with functionality built-in allowing users to specify that a group and its members should be replicated to Active Directory (or FreeIPA in future) once E-Groups is phased out and GRAPPA becomes the effective owner of all authorization groups. We also set functionality in place to allow users to synchronize groups created in GRAPPA to the existing E-Groups, in case they need to be used in legacy authorization applications.

3.1.6 Keycloak

Keycloak is the open source Identity and Access Management solution selected to replace Microsoft's ADFS for SSO. Keycloak supports the two main protocols required by CERN

services; SAML2, used by many older applications, and OpenID Connect, increasingly used by modern applications. Keycloak was deployed on CentOS 7.7 on virtual machines running on OpenStack [16], with configuration managed by Puppet [17]. We are running Keycloak in clustered mode with a set of HAProxy [18] load balancers in front for high availability. All Keycloak nodes are connected to a MySQL8 database configured to provide a resilient, distributed, and highly available service. We use a Keycloak Metrics Service Provider Interface (SPI) [19] to add a metrics endpoint to Keycloak, which is scraped by Prometheus [20] and made available on Kibana [21].

3.1.7 Satosa

Satosa [22] was chosen to bridge between Keycloak and eduGAIN. Satosa is a Service Provider proxy written in python and is maintained by the wider Research and Education Identity Federation community. It natively supports the community's policies and profiles, as well as multilateral federation (federation with SAML metadata that contains multiple entities) which is currently not supported by Keycloak. Participation in wider community efforts, such as REFEDS [23], the GEANT Project [24] and eduGAIN meetings, allowed the early identification of this Open Source tool.

3.1.8 WLCG

The WLCG Authorization Working Group was established in 2017 and contains many of the key players from IGTF [25], AARC [26], e-Infrastructures and experiments, including individuals involved in SciTokens [27], INDIGO IAM [28], EGI Checkin [29]. The group has converged on a token schema to be used within the WLCG, and identified a software stack to provide the functionality. It is envisaged that this system will be fully integrated with the new CERN Authentication system and manage the grid-specific functionality itself. This separation of functionality is both to avoid special cases in CERN's authentication infrastructure and also to allow the software to be used by WLCG Virtual Organisations (VOs) who may wish to integrate an alternative Authentication source due to having little affiliation with CERN.

4 Key Changes

The following represent the key changes in CERN's AAI from a user's perspective:

- The concept of roles has been introduced to discontinue the practice of sending all authorization groups to all applications. Application owners define roles for their application (e.g. Admin), assign groups to the role (e.g. e-group my-service-admins) and configure Multifactor and LoA requirements.
- The suite of user facing services has been redesigned from scratch, with an emphasis on modern workflows, and will be continually improved over the pilot period. In particular, the Login page has been significantly simplified. Input was sought from the CERN Community, with many suggestions being incorporated into the final design.
- The adoption of Industry standard technologies, OAuth2.0 [4] and OpenID Connect (OIDC) [30] using JSON Web Tokens (JWT) [31], greatly simplifies integration with downstream services.
- The Identity Model offers the ability to link accounts, improving user experience when switching institutes or transitioning to retirement.

5 Conclusion

2019 saw an important milestone in CERN's Identity and Access Management services, the release of an Alpha version of a new Authentication and Authorization Infrastructure based on Open Source software and standards. It is a first step in a long journey that will result in an environment that is more more intuitive to both CERN researchers and software maintainers. Initial analyses indicate that the new Single-Sign-On system compares favourably with the previous Microsoft based system, and feedback on integration from service owners has been largely positive. Feedback on other components will be sought as they enter pilot phase. The project is set to continue over the coming years, in parallel with efforts in WLCG, as both move towards a token based infrastructure.

References

[1] B.C. Neuman, T. Ts'o, IEEE Communications magazine **32**, 33 (1994)
[2] J. Shiers, Computer physics communications **177**, 219 (2007)
[3] *X.509 Public-key and attribute certificate framework*, https://www.itu.int/rec/T\-REC-X.509-201910-I/en, accessed: 2020-03-02
[4] *OAuth2*, https://oauth.net/, accessed: 2020-03-02
[5] *General Data Protection Regulation*, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX\%3A32016R0679, accessed: 2020-03-02
[6] *Federated Identity Management for Research*, https://fim4r.org, accessed: 2020-03-02
[7] *Active Directory Domain Services*, https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/active-directory-domain-services, accessed: 2020-18-02
[8] *Active Directory Federation Services*, https://docs.microsoft.com/en-us/windows-server/identity/active-directory-federation-services, accessed: 2020-18-02
[9] *Hangfire*, https://www.hangfire.io, accessed: 2020-18-02
[10] *JSON Schema*, https://json-schema.org/, accessed: 2020-10-02
[11] *Microsoft Identity Manager*, https://docs.microsoft.com/en-us/microsoft-identity-manager/microsoft-identity-manager-2016, accessed: 2020-18-02
[12] *FreeIPA*, https://www.freeipa.org/page/About, accessed: 2020-03-02
[13] *389 Directory Server*, https://directory.fedoraproject.org/, accessed: 2020-03-02
[14] *Dogtag*, https://www.dogtagpki.org/wiki/PKI_Main_Page, accessed: 2020-03-02
[15] *Elasticsearch*, https://www.elastic.co/elastic-stack, accessed: 2020-10-02
[16] *Openstack*, https://www.openstack.org, accessed: 2020-21-02
[17] *Puppet*, https://puppet.com, accessed: 2020-21-02
[18] *HA Proxy*, http://www.haproxy.org, accessed: 2020-21-02
[19] *Keycloak Metrics SPI*, https://www.keycloak.org/extensions.html, accessed: 2020-21-02
[20] *Prometheus*, https://prometheus.io, accessed: 2020-21-02
[21] *Elastic Search Kibana*, https://www.elastic.co/kibana, accessed: 2020-21-02
[22] *Satosa*, https://github.com/IdentityPython/SATOSA, accessed: 2020-03-02
[23] *The Research and Education Federations Group (REFEDS)*, https://refeds.org, accessed: 2020-03-02
[24] *GÉANT Project (GN43)*, https://www.geant.org/Projects/GEANT_Project_GN4\-3/Pages/Home.aspx, accessed: 2020-03-02
[25] *Interoperable Global Trust Federation*, https://www.igtf.net, accessed: 2020-03-02

[26] *AARC (Authentication and Authorisation for Research and Collaboration)*, https://aarc-project.eu, accessed: 2020-21-02
[27] *SciTokens*, https://scitokens.org, accessed: 2020-21-02
[28] *INDIGO Identity and Access Management*, https://github.com/indigo-iam, accessed: 2020-03-02
[29] *EGI Check-in*, https://www.egi.eu/services/check-in/, accessed: 2020-21-02
[30] *OpenID Connect*, https://openid.net/connect/, accessed: 2020-03-02
[31] *JSON Web Tokens*, https://jwt.io/introduction/, accessed: 2020-03-02