04-06-2018

Deliverable DJRA1.2: Scalable, integrated authorisation models for SPs

Deliverable DJRA1.2

Contractual Date:	31-03-2018
Actual Date:	04-06-2018
Grant Agreement No.:	730941
Work Package:	JRA1
Task Item:	JRA1.2
Lead Partner:	KIT
Document Code:	DJRA1.2
Authors:	Marcus Hardt (KIT), David Hübner (DAASI), Jens Jensen (STFC), Christos Kanellopoulos (GÉANT), Nicolas Liampotis (GRNET), Mikael Linden (CSC), Shiraz Memon Jüelich), Stefan Paetow (JISC), Mischa Sallé (Nikhef), Diego Scardaci (EGI), Uros Stevanovic (KIT), Davide Vaghetti (GARR), Niels van Dijk (SURFnet)

Abstract

Over the last years, federated access has established itself as an enabler to access and share resources in a user-friendly and, at the same time, privacy-preserving way, as it allows users to use their existing credentials verified by their home organisations. However, authorisation still poses challenges, particularly in the context of international research collaboration, as it requires the exchange of community-managed attributes across research infrastructures and e-infrastructure. This document describes common authorisation models that can be employed by Service Providers (SPs) in order to control access to resources in such an environment. These common models are based on a thorough analysis of use cases collected from the research communities participating in the pilot activities of AARC. The analysis includes describing the different authorisation functions, including management, evaluation and enforcement of policies and their mapping to elements of the AARC Blueprint Architecture. The types of attributes that are commonly used for evaluating authorisation policies are also elaborated on.

Table of Contents

Table of Figures

Executive Summary

The key reason for using any means of Authentication and Authorisation Infrastructures (AAI) is - simply speaking - to provide access to only the right people. Through national identity federations and eduGAIN, the identity providers (IdPs) of the users' home organisations provide a well-established authentication service and, in some cases, they are also used as authoritative sources of information for determining access to resources. One such example is the "common-lib-terms" entitlement attribute value, which is used by home organisations to signal eligibility for accessing publisher resources. However, in the context of international research collaboration, authorisation is typically based on information managed by the collaboration, for instance group/project membership and role information. Therefore, we cannot expect home organisations to manage collaboration-specific authorisation attributes. Thus, the challenge is to allow a certain group of people to manage access rights to resources that belong to different groups of people for users that are under the control of yet another, very diverse, group of people.

Investigating the authorisation aspects in such a multi-domain federated environment is one of the key objectives of AARC2. To this end, we engaged with the research communities participating in the pilot activity of the project in order to describe their authorisation approaches. Following a series of unstructured interviews and meetings we observed a diverse set of authorisation requirements. The observed range covers fields from Astrophysics, which largely benefits from data being publicly available, over particle physics, where large teams work in a competitive environment, to several LifeScience fields, that have to ensure safety and privacy of their information and often include committees that decide - on a per dataset level -who is given access.

For the analysis of the different authorisation use cases, each community has provided a description of their authorisation model. We use the terminology from well-established authorisation concepts, such as PEP, PDP, PAP, and PIP. Even though these concepts are also commonly used for XACML it should be emphasised that we neither mandate the use of XACML nor do we oppose it. We solely use its nomenclature, because it is concise, expressive, and generally well understood. To avoid misunderstanding we reference this model "PPP-Model", because of the many P*P acronyms involved.

We conclude with the five different models for authorisation that we found. These are explained and associated with the respective communities that make use of it. The Annexes contain additional information on authorisation patterns, technologies, as well as additional technical information for some of the use-cases.

.

1 Introduction

The goal of this document is to derive common patterns regarding the use of authorisation in distributed federated infrastructures and to identify the logical position of different authorisation components in the Blueprint Architecture.

For this we collected input from several different communities that we describe in section 3. To do this, we need to use a common language. We rely on concepts for authorisation from RFC 2753 [RFC2753] and RFC 2904 [RFC2904]. Please note that these concepts are also used in the context of XACML. However, we do not mandate or limit ourselves to XACML in any way because they are well understood in general. This model allows for distinguishing between administration, definition, retrieval, evaluation and enforcement of authorisation policies. Since the model defines several elements that start with 'P' (e.g. PIP, PAP, PDP, ...) we refer to this authorisation model at the "PPP-model".

For reference we also provide an Annex with several short sections on authorisation models (Appendix A) and one on technologies (Appendix B). There, short sections contain links to more extensive information.

Appendix C collects the details of the use-case descriptions including the architectures of the authorisation. These are the basis for the resulting observations and patterns that we describe in section 4.

2 Terms and definitions

This section describes the concepts that we used to describe authorisation architectures. We understand this standard to be completely technology agnostic and use it only for descriptive purposes. In case we encounter a feature of a use-case that we cannot describe with this standard, we will explicitly emphasise this in the text.

RFC 2753 [RFC2753] and RFC 2904 [RFC2904] describe a general Authorisation Framework model to define Policy based Admission Control. They form the architectural base for the e**X**tensible **A**ccess **C**ontrol **M**arkup **L**anguage (XACML), an OASIS standard. XACML adds the Policy Administration Point, PAP to the earlier RFCs. The resulting architecture is outlined in Figure 2.1.

This model assumes that a Policy Enforcement Point (PEP) is responsible for protecting access to one or more resources. When a resource access is attempted, the PEP sends a description of the attempted access to a Policy Decision Point (PDP) in the form of an authorisation decision request. The PDP evaluates this request against its available policies and attributes and produces an authorisation decision that is returned to the PEP. The PEP is responsible for enforcing the decision.

In producing its description of the access request, the PEP may obtain attributes from online Attribute Authorities (AA) or from Attribute Repositories into which AAs have stored attributes. Attribute authorities act as policy information points

(PIP) in the sense of RFC2753/RFC2904. The PDP may augment the PEP's description of the access request with additional attributes obtained from PIPs.

The PDP may obtain policies from online Policy Administration Points (PAP) or from Policy Repositories into which PAPs have stored policies.

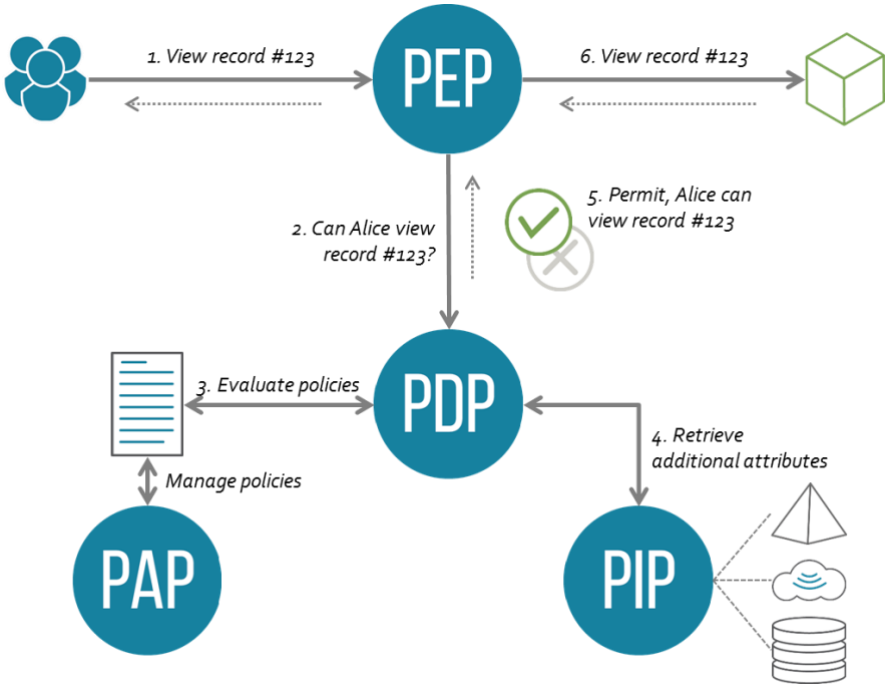

Figure 2.1: The PPP-Model (source: Wikipedia, CC-BY-3.0, image credit David Brossard)

3 Real life architectures for authorisation

This chapter contains use-case descriptions contributed by the pilot communities of the AARC project. Each use-case section uses the same structure to aid comparison, and is limited to two pages to allow easy navigation. The sections start by giving a short general description of the authorisation challenge addressed by the given architecture.

3.1 CLARIN

3.1.1 General description

The CLARIN infrastructure [CLARIN] is a distributed infrastructure, consisting of many independent centres. Federated identity management and single sign-on have been standardised on SAML, including eduGAIN support, in the CLARIN service provider federation (SPF) [CLARIN-SPF]. The main authorisation approach is ABAC (please refer to A.3), based on:

1. The standardised set of attributes defined for the CLARIN infrastructure. These attributes are obtained directly from the user's home organisation identity provider.

2. And possibly extended with information on signed licenses, categorised into CLARIN PUB(lic), CLARIN ACA(demic) or CLARIN RES(tricted). This information is typically managed by the CLARIN centres, which for the common case equals the actual service providers.

3. And possibly extended with other information specific to the CLARIN centre hosting the resource or service. This information is typically managed by the CLARIN centres, close to the actual service providers.

3.1.2 Architecture and flow

Every centre, aka SP, implements its own authorisation approach, including administration of their authorisation policies. For accessing resources or services hosted at some of the centres or services, it is usually sufficient for users to just be authenticated, while in other cases users are required to include signed licenses information and/or have group / entitlement information, managed locally at the specific centres.

The Language Archive (TLA) [TLA] **[AARC-G002]** Expressing group membership and role information (AARC-G002); https://aarc-project.eu/guidelines/aarc-g002/

[AARC-G006] Best Practices for managing authorisation (AARC-G006); https://aarc-project.eu/guidelines/aarc-g006/

[AARC-G036] Roles, responsibilities and security considerations for VOs (AARC-G036); https://aarc-project.eu/guidelines/aarc-g036/

[AARC-MJRA1.1] Existing AAI and available technologies for federated access (MJRA1.1); https://aarc-project.eu/wp-content/uploads/2016/01/MJRA1.1-Existing-AAI-and-available-technologies.pdf

[ALF04] Alfieri R. et al. (2004) VOMS, an Authorization System for Virtual Organizations. In: Fernández Rivera F., Bubak M., Gómez Tato A., Doallo R. (eds) Grid Computing. Lecture Notes in Computer Science, vol 2970. Springer, Berlin, Heidelberg

[ARGUS] Argus Authorization Service; http://argus-documentation.readthedocs.io/en/stable/

[CLARIN] Common Language Resources and Technology Infrastructure (CLARIN); https://www.clarin.eu/content/clarin-in-a-nutshell

[CLARIN-SPF] CLARIN Service Provider Federation; https://www.clarin.eu/content/service-provider-federation

[CORNWALL2004] Cornwall, Linda A., et al. "Authentication and authorisation mechanisms for multi-domain grid environments." Journal of Grid Computing 2.4 (2004), pp. 301-311.

[DYKE2016] Dyke, S., Kirby, E.,Shabani, M., Thorogood, A.,Kato, K., Knoppers, B. Registered access: a 'Triple-A' approach. *European Journal of Human Genetics* volume 24, pages 1676–1680 (2016); https://www.nature.com/ejhg/journal/v24/n12/full/ejhg2016115a.html

[EGI-REG] EGI AAI entitlement registry; https://wiki.egi.eu/wiki/URN_Registry:aai.egi.eu

[EPOS] European Plate Observing System (EPOS); https://www.epos-ip.org/

[NIST800-162] Guide to Attribute Based Access Control (ABAC) Definition and Considerations, (NIST Special Publication 800-162), 2014; https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-162.pdf

[PERUN] Perun, Identity and Access Management System; https://perun.cesnet.cz/web/

[SANDHU1996] Sandhu, R., Coyne, E.J., Feinstein, H.L. and Youman, C.E. (August 1996). "Role-Based Access Control Models" (PDF). IEEE Computer. IEEE Press. 29 (2): 38–47. doi:10.1109/2.485845

[TLA] is an example of a repository with a sophisticated authorisation approach. The Flat repository (based on Fedora and Islandora [TLA-FLAT]) is used to manage group membership for its users, based on the user-id, keeping track of signed licences and managing the repository authorisation policies. Users authenticated into the CLARIN SPF are tracked in Flat based on their associated attributes. Each user belongs to one or more groups and might have signed zero or more license agreements. This information is aggregated and combined with an XACML authorisation policy to make an authorisation decision. Flat provides PAP and PIP interfaces and enforces the XACML policies, thus also acting as the PDP and PEP.

3.1.3 Architecture described in the reference model

In the general CLARIN architecture, authorisation is implemented at each individual centre. Therefore, each centre provides a PEP, PDP and PAP implementation specific to their setup and repository and optionally additional attribute authorities (again specific to that centre) can be used as policy information points (PIPs).

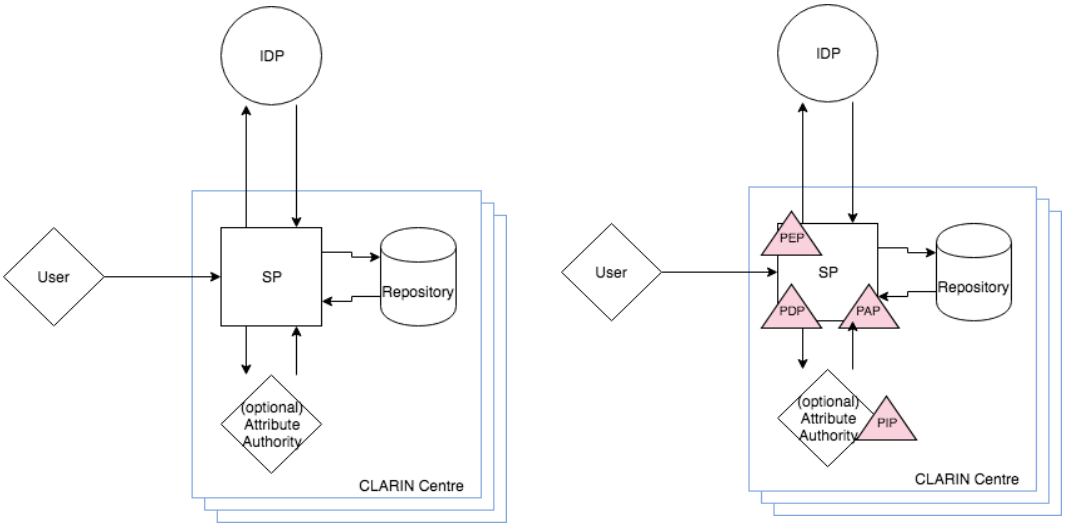

Figure 3.1: CLARIN authorisation model

3.1.4 Mapping to the Blueprint Architecture

Applying this architecture to the BPA results in the PDP, PEP and PAP all located in the green authorisation box, overlapping with the red end services box, since each of these components is implemented at each centre. The PIP is put in the blue user attributes box, but functionally a separate AA component can be running at the individual end services.

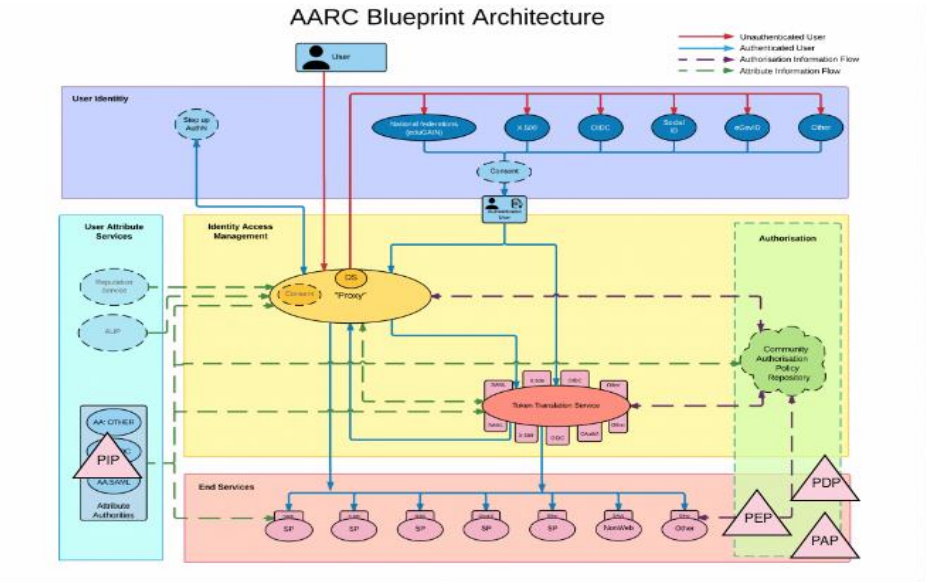

Figure 3.2: Mapping the CLARIN authorisation model to the Blueprint Architecture

3.2 DARIAH

3.2.1 General description

DARIAH provides an AAI and services for researchers in the digital humanities. The DARIAH subproject *TextGrid* runs an authorisation infrastructure to allow researchers access to various resources (e.g. TextGrid Repository). Anyone who can authenticate in eduGAIN is allowed to create *projects* in the TextGrid Repository. The project creator can assign roles to authenticated users and control permissions on their project in this way. The general authorisation approach therefore follows the RBAC model (see Appendix A.2). The TextGrid authorisation model has recently been combined in DARIAH to a more general approach that combines the RBAC model with OAuth2 technology (see Appendix B2). The PDP consist of an OAuth2 Authorisation Server and the OpenRBAC engine.

3.2.2 Architecture and flow

A project in the TextGrid Repository consists of a set of predefined roles. The project creator can assign these roles to other authenticated users. The PDP database stores, for each resource, a mapping of roles to permissions on operations (which is a predefined set of the verbs *manage, create, read, write, delete*). The PDP API offers all functions the NIST RBAC standard describes, with the most frequently used function *checkAccess* that returns a permit/deny decision based on a (resource, user, operation) triple. As illustrated in Figure 3.3, the flow consists of the following steps:

1. The TG-Lab Client applications request a token from the TG-Auth authorisation server.

2. Authentication of the user.

3. The TG-Auth authorisation server issues a so-called SID token (comparable to an access token in OAuth2).

4. TG-Lab requests the intended action at TG-CRUD (file and metadata management service). This requests contains the (resource, user, operation) triple with the user's identity being encoded in the SID token.

5. The checkAccess function at the PDP is called with this triple.

6. The PDP looks up policy information (roles and permissions) in an attached LDAP store, using the user information from the SID token.

7. The PDP informs TG-CRUD about the policy decision (permit or deny)

8. TG-CRUD enforces the policy decision and only performs the operation if it was permitted.

3.2.3 Architecture described in the reference model

- The TG-CRUD service fulfills the role of a *PEP* and enforces the policy decision made by the TG-Auth PDP. Only if this decision is positive, the requested operation is performed.

- The TG-Auth PDP acts as a *PDP*. Based on the requested triple (operation, resource, user), a policy decision is made using the RBAC approach.

- The user store (LDAP server) acts as both a *PIP* and *PAP*. Initially during authentication user attributes are provided by this LDAP server. In addition to that, this data store also contains the roles for all users and the mapping from roles to permissions. This information is then received by the TG-Auth PDP.

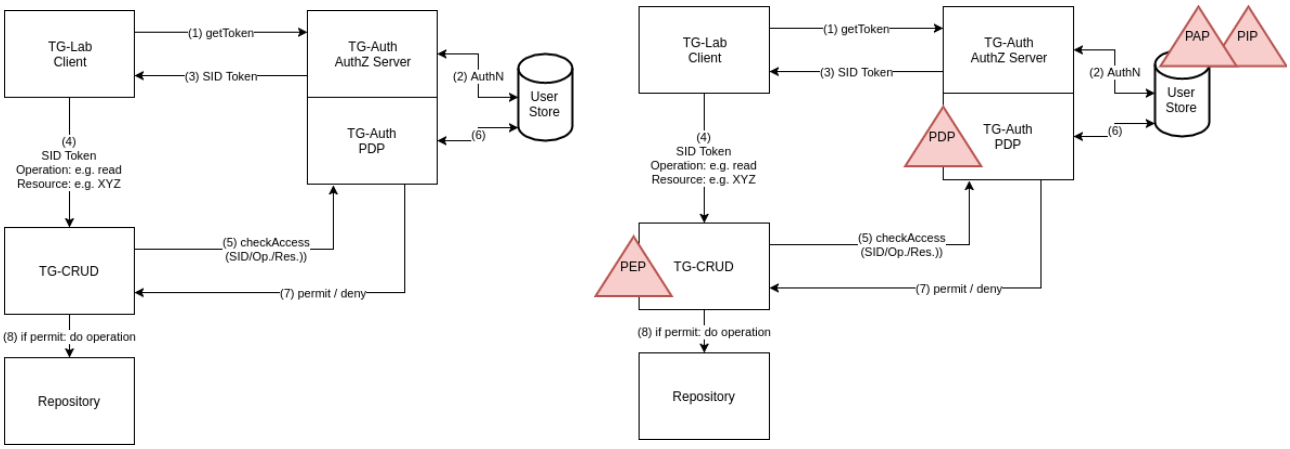

Figure 3.3: DARIAH AAI authorisation model

3.2.4 Mapping to the Blueprint Architecture

TG-CRUD (the *PEP*) is located in the "End Services" layer and hence relies on both authentication and authorisation information provided by additional components. The *PDP* and the attached part of the LDAP directory, which contains policy information (the *PAP*), can be considered as central components that only deal with authorisation and are therefore located in the green-yellow overlap region of the BPA. The user attribute part of the LDAP directory, which serves as a *PIP* for authorisation purposes is clearly located in the "User Attribute" layer. Since the BPA distinguishes between logical and not necessarily physical components, the *PAP* and *PIP* parts of the LDAP are located in different regions.

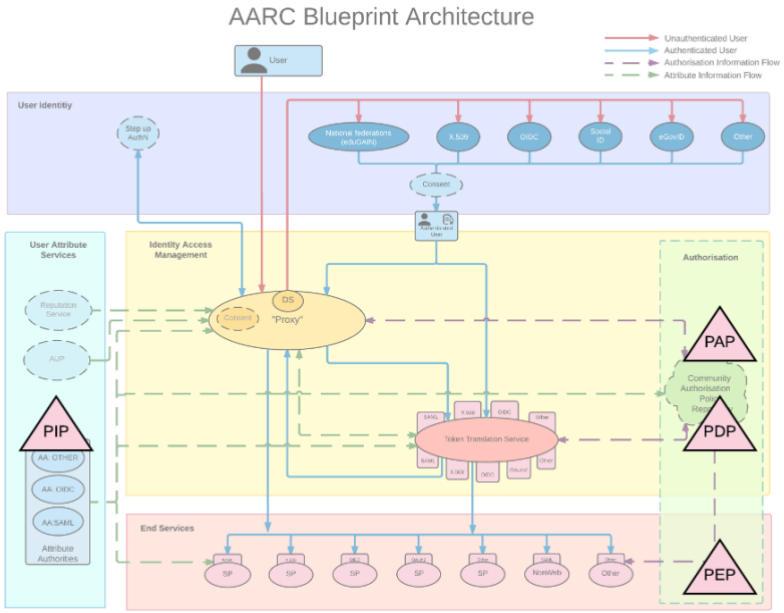

Figure 3.4: Mapping the DARIAH AAI authorisation model to the Blueprint Architecture

3.3 EGI Check-in-enabled services

3.3.1 General description

EGI is an e-infrastructure that provides resources and services to diverse research disciplines. The EGI Check-in service is an Identity and Access Management solution that makes it easy to secure access to services and resources. Through Check-in, users are able to authenticate with the credentials provided by the IdP of their Home Organisation (e.g. via eduGAIN), as well as using social identity providers or other selected external identity providers. Check-in provides an intuitive interface for communities to manage their users and their respective groups, roles and access rights. For communities operating their own group management system, Check-in has a comprehensive list of connectors that allows to integrate their systems as externally managed Attribute Authorities, as illustrated in Figure 3.5.

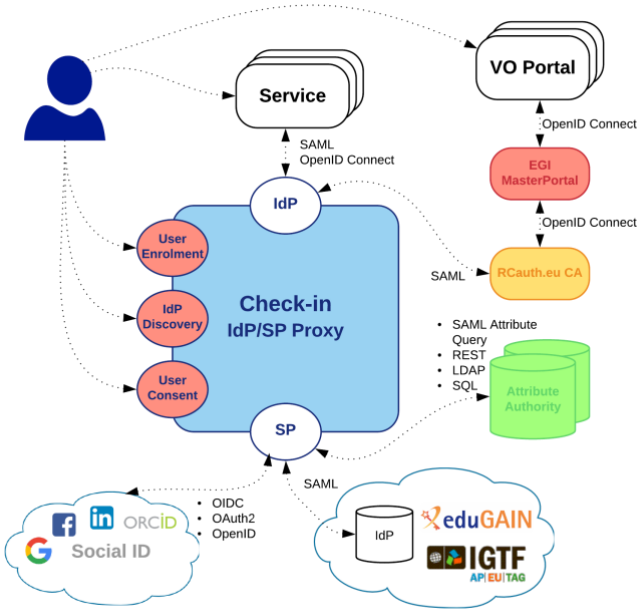

Figure 3.5: EGI Check-in AAI service architecture

3.3.2 Architecture and flow

EGI Check-in supports authorisation decisions based on the combination of different types of information, including:

- identity attributes asserted by the IdP of the user's home organisation;
- VO/group membership and role information aggregated from one or more community-managed attribute authorities;
- assurance information associated with the authenticated identity.

Based on the information above, Check-in makes available two types of attributes that can be used by SPs to control access to resources, namely, entitlements and assurance. Entitlements can either refer explicitly to a set of rights/capabilities of the user to access specific services/resources, or implicitly by conveying the user's VO/group membership and role information (group- and/or role-based access control). Attributes carrying assurance information can be used by SPs to decide how much to trust the assertions made by Check-in and its attribute sources.

3.3.3 Architecture described in the reference model

As already stated, Check-in aggregates authorisation-related information from different sources, such as the authenticating IdP and the community-managed attribute authorities. Following the attribute aggregation, we can distinguish between two flows, as illustrated in Figure 3.6:

A. Check-in evaluates the incoming authorisation request against the policies it has been configured with and returns an entitlement that represents the right of the authenticated user to access that particular resource. For example, the "urn:mace:egi.eu:aai.egi.eu:rcauth" value is used to indicate that the holder of this entitlement is eligible for accessing the RCauth.eu Online CA service. The EGI AAI URN registry [EGI REG] lists all supported entitlement values. Check-in acts as a PDP in this flow.

B. Check-in passes the aggregated information onto the service that the user is trying to access and the service is then responsible for making the appropriate authorisation decision. Check-in is merely passing on information and therefore acts as a centralised PIP.

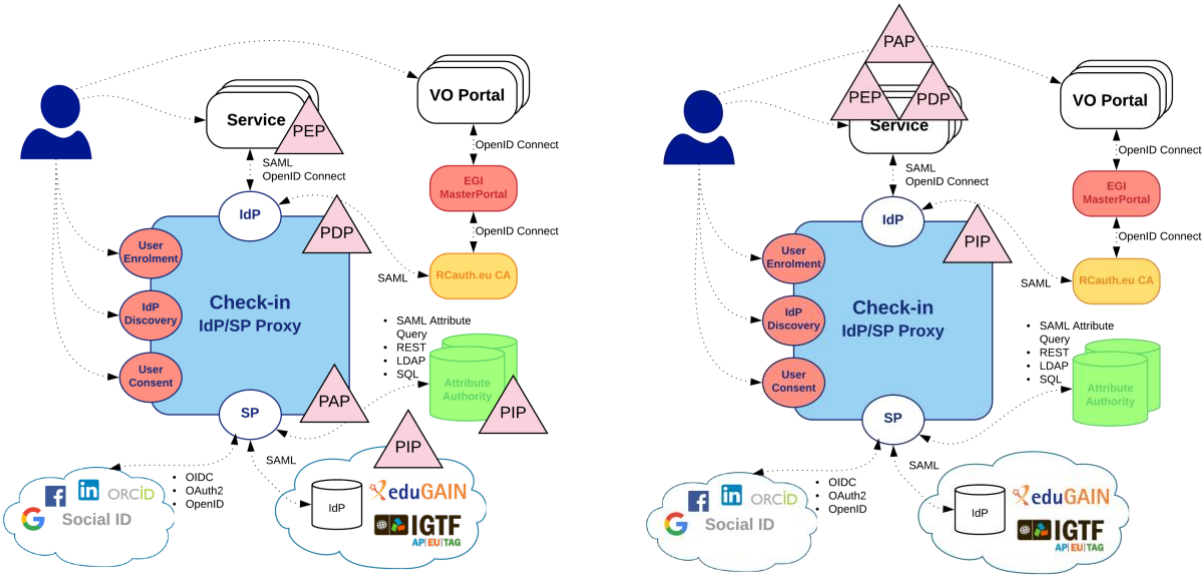

Figure 3.6: EGI Check-in authorisation models. Left: Centralised policy management and decision making. Right: Centralised policy information point

3.3.4 Mapping to the Blueprint Architecture

Figure 3.7 illustrates how the two authorisation models described in Section 3.3.3 are mapped to the elements of the AARC Blueprint Architecture.

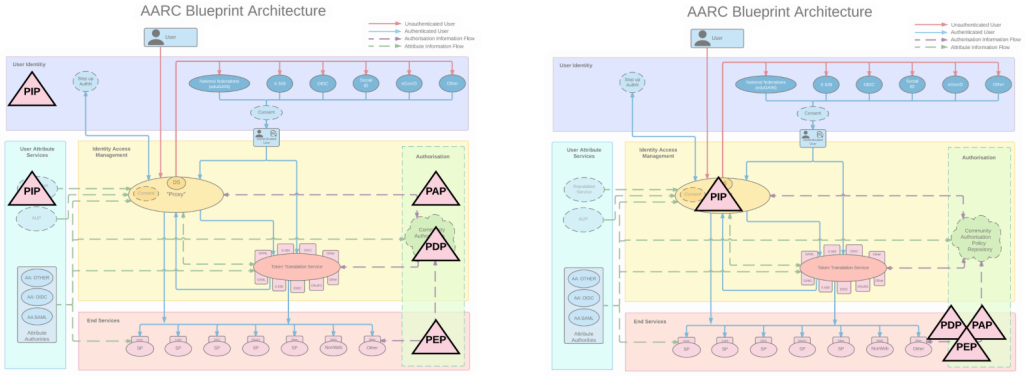

Figure 3.7: Mapping the EGI Check-in AAI authorisation models to the Blueprint Architecture. Left: Centralised policy management and decision making. Right: Centralised policy information point

3.4 ELIXIR

3.4.1 General description

ELIXIR is the European research infrastructure for biological data. ELIXIR is distributed and consists of 21 national nodes who each can provide a handful of services, or more. Some services are simple collaboration tools (like wiki) but the high-end services can be data archives, compute clouds and workflow systems with sophisticated AAI needs.

Much of the biological data (such as plant and marine) is publicly available but the datasets donated to research by human patients are typically sensitive and require careful authentication and management of the researchers' access rights. Although the sensitive datasets may be replicated to several data centres globally, researchers' access rights to them are typically defined centrally by the dataset owner, e.g. the organisations who belonged to the project that gathered the samples.

3.4.2 Architecture and flow

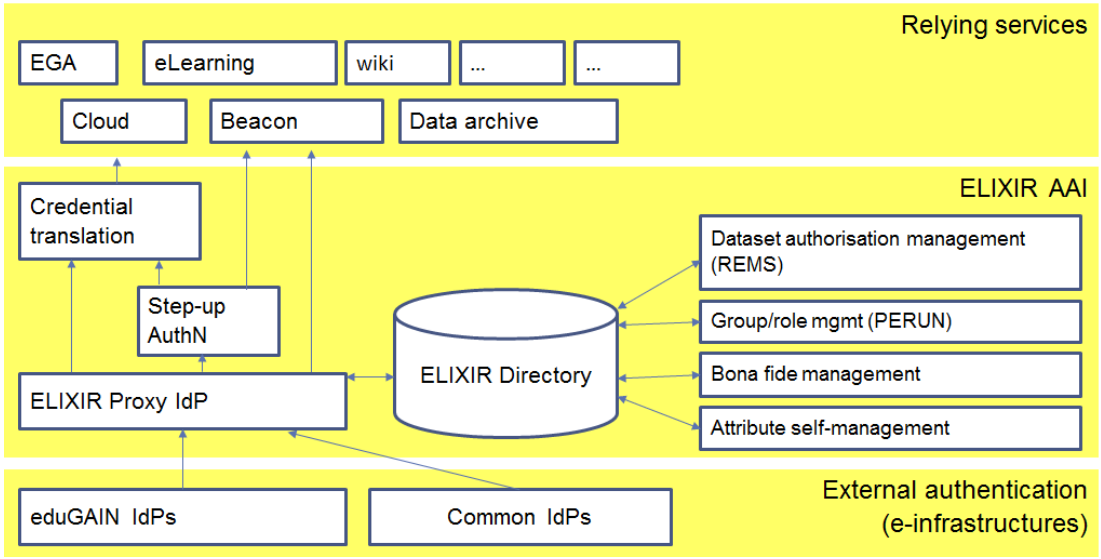

Figure 3.8: ELIXIR AAI architecture

Figure 3.8 illustrates the general ELIXIR AAI architecture. The services relying on ELIXIR AAI are in the upper part of the figure and the authentication providers (like eduGAIN, ORCID and Google) in the lower part of the figure. In the middle are the ELIXIR AAI service components; the services for user authentication on the left and the services to "decorate" the authenticated users with extra attributes (describing their access rights to relying services) on the right.

A three-tier access model pertaining to access to sensitive human data is emerging in Life sciences:

1. **Public access**; the datasets are publicly available. No authentication and separate authorisation needed.

2. **Registered access**; the datasets are available to people who demonstrate they are bona fide researchers (cf. [AARC-MJRA1.1] section 2.1.2-4.) i.e. researchers in good standing. This step needs to be done only once and gives the researcher access to all datasets (and other services) that belong to the registered access tier. The exact mechanism to register the bona fide researcher status is subject to discussion; the approach made by ELIXIR relies on:

    a. The user's Home Organisation claiming the person is a researcher,

    b. A person qualifying through (a) above vouching for the user being a bona fide researcher, or

    c. The user demonstrating they have publications in recognised scientific journals

    The three alternative approaches are complemented by attestations the person needs to make to claim the bona fide status, for instance "I refrain from trying to re-identify individuals from the datasets". (For more information see [DYKE2016])

3. **Controlled access**; access to datasets is based on the researcher presenting a data access application to the dataset owner. This is the classical approach which has the downside that it is slow and labour-intensive for the dataset owner. Recently electronic tools have been developed to automate the process; for a short overview, see the section on REMS below (link).

3.4.3 Architecture described in the reference model

Figure 3.9 illustrates a specific access control enforcement scenario has been selected for further analysis. A researcher has received access rights to a controlled access dataset (using the REMS Dataset authorisation management tool in the above figure, "PAP") and the access rights are stored in the Central EGA (European Genome-phenome Archive) service. The user wants to use their access rights in a private cloud (client e.g. CSC) which already possess a copy of the dataset.

The user launches their web browser and is authenticated by the ELIXIR AAI which in turn fetches the user's permissions from the Central EGA server ("PDP") and assembles them into an access token which describes the user's permissions. The access token is then presented to the server that controls access to the datasets ("PEP").

Figure 3.9: ELIXIR AAI authorisation use case

3.4.4 Mapping to the Blueprint Architecture

As illustrated in Figure 3.10, the components described in the previous section can be projected to the blueprint architecture as follows:

- Central EGA service is the authoritative source of user's permissions to datasets (PAP and PDP)
- The user interface (PAP) for the dataset owner to configure the users' dataset permissions is the Dataset Authorisation Management service (REMS)
- The user attributes (such as their name and ORCID identifier) available (PIP) in ELIXIR AAI may help the dataset owner in making the decision to grant access rights to datasets. Furthermore, fresh information on the

researcher's Home Organisation is important because the dataset owner typically couples researchers' permissions to their continuing affiliation with their Home Organisation.

- The enforcement of the access rights (PEP) is done in the relying service

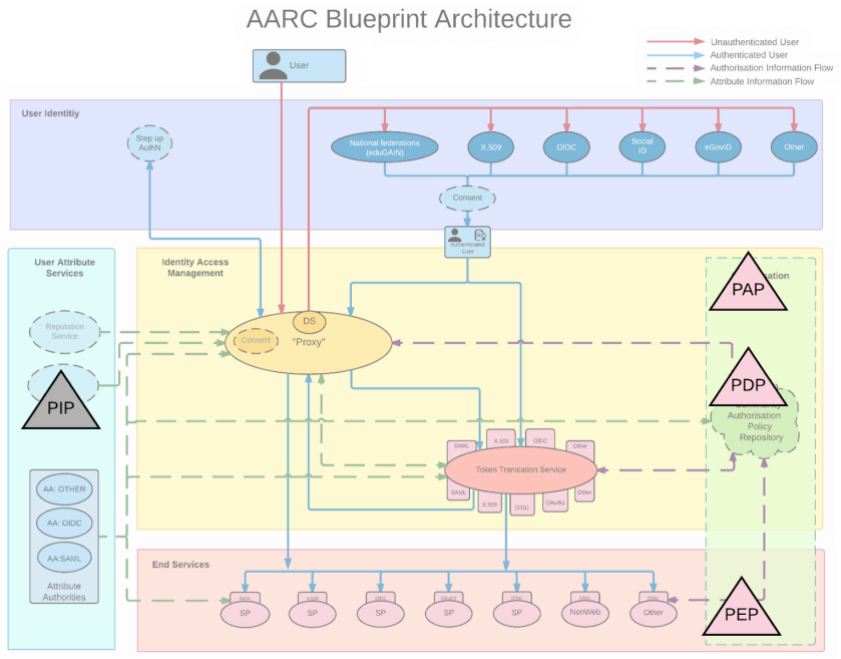

Figure 3.10: Mapping the ELIXIR AAI authorisation model to the Blueprint Architecture

3.5 EPOS

3.5.1 General Description

The main aim of EPOS (European Plate Observing System) [EPOS] is to coordinate, collect, archive high-quality Earth Science data across Europe. By definition, EPOS is a distributed Research Infrastructure where Data, Data Products, Software and Services (DDSS) are provided by different communities in the domain of the solid Earth sciences. In this framework, EPOS envisage the construction of a central hub called "Integrated Core Services" (ICS-C) which aggregates all DDSS from the various disciplines. From the technical viewpoint, DDSS are provided by a distributed network of endpoints (Thematic Core Services, TCSs) which needs to use heterogeneous authorisation mechanisms. EPOS enables cooperation of about 2000 users coming from academia, industry and society.

3.5.2 Architecture and flow

The EPOS AAI allows authentication through external IdPs (e.g. eduGAIN), while it also provides an internal EPOS IdP for the "homeless" users.

In a typical scenario, the user will:

- Log in at the ICS-C level,
- search the data it requires,
- download the data from a TCS.

In more advanced scenarios, the user will:

- transfer the data to a TCS science gateway,
- download additional data files (if necessary, from other TCSs),
- run an analysis via the science gateway,
- reuse the results and pass them in to other software packages for additional analysis,
- access some external visualisation engine.

3.5.3 Architecture described in the reference model

The EPOS AAI has to take into account the requirements of both: EPOS Central Hub services (ICS-C) and Thematic Core Services (TCS). At the Central Hub level, ICS-C employs the general policies which apply for all TCSs. Additionally, the TCSs can extend the level of policies adding their own, specific requirement. For example, some TCSs are gathering data coming from private companies. There will be classes of users whose role will depend on data policies:

- open access – all the datasets and software are available for everyone. Data search does not require authentication, however data access does;
- embargoed data – access requires authorisation for a given period of time, then data reaches status of open data;

Embargoed data are subject to strict access rules for a given period of time. After the embargo for the data is lifted, the data will be available for the remaining users without any restrictions.

The general set of attributes about a user will be provided by EPOS attribute DB at the ICS-C level. Additional set of attributes can be employed at the TCS level.

3.5.4 Mapping to the Blueprint Architecture

The EPOS AAI is an implementation of the AARC Blueprint Architecture that comprises an SP-IdP-Proxy component acting as a central hub between Identity Providers (both the external IdPs in eduGAIN and the internal EPOS IdP) and the ICS-C/TCS services. As illustrated in Figure 3.11, authorisation policies are managed at the service level (see PAP element).

The authorisation policies are generally based on attributes centrally provided by the EPOS attribute DB through the SP-IdP-Proxy (central PIP). However, in some cases, services can retrieve additional attribute sets in order to evaluate access requests (see additional PIP(s) connected with services). Authorisation policies are evaluated at the service level (PDP) and the service is responsible for enforcing the authorisation decision (PEP).

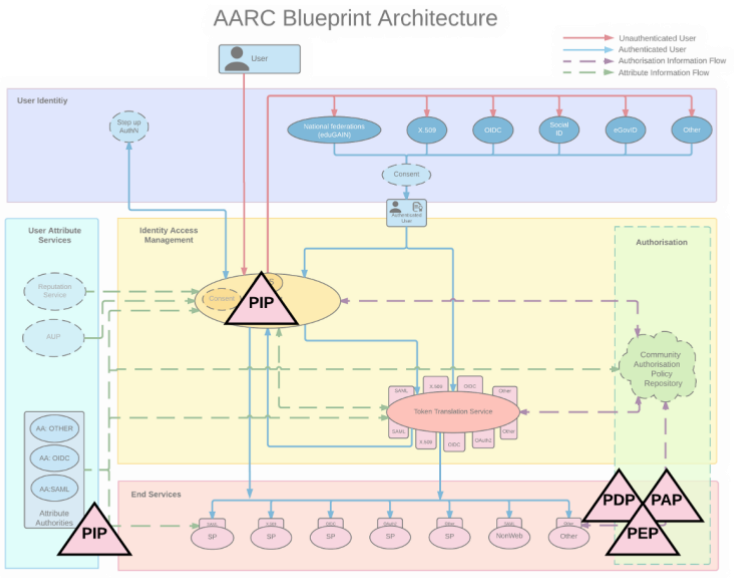

Figure 3.11: Mapping the EPOS AAI authorisation model to the Blueprint Architecture

3.6 EUDAT B2ACCESS

3.6.1 General description

The EUDAT "collaborative data infrastructure" provides access to a wide range of data services, such as B2SAFE (storage), B2SHARE (sharing), B2DROP (dropbox), B2FIND (metadata and discovery), etc. In order to implement federated identity management, the EUDAT project surveyed in 2012 the available technologies and their maturity. The technology eventually matured into a distinct service called B2ACCESS.

3.6.2 Architecture and flow

As illustrated in Figure 3.12, the typical flow is as follows:

- Users access a service, e.g. B2SHARE
- If they are not authenticated, they are given the option to log in through B2ACCESS

- B2ACCESS, in turn, presents a discovery service which directs to the selected IdP.
- Once users have authenticated with the IdP, they are registered with B2ACCESS (if they haven't used it before) or they get redirected back to their service if they are already registered.
- Once authenticated, B2ACCESS can add attributes about the user, e.g. groups or roles.
- At this stage, credential conversion is possible: since the EUDAT services were built on a variety of software products, the credential can be converted to OIDC, SAML, or an X.509 certificate.

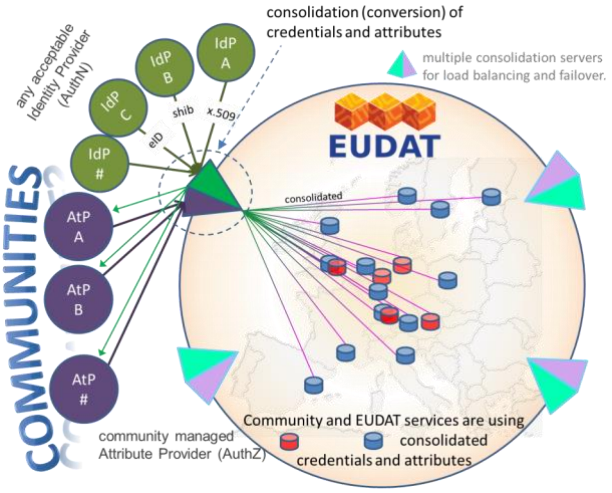

Figure 3.12: EUDAT AAI high-level design (J.Reetz, et al., 2012)

3.6.3 Architecture described in the reference model

As illustrated in Figure 3.13, B2ACCESS is the central AAI component in EUDAT; in the BPA it is essentially both the SP-IdP-Proxy and the credential conversion component, as well as an attribute authority.

- It acts as a PIP because it maintains attributes on behalf of the users (usually used for authorisation, e.g. memberOf, or unique id.) It makes these attributes available to services in the federation.
- It acts as a PAP insofar as it gives authorised users (by default federation administrators, but delegation is possible) the means of assigning these attributes.
- It acts as a PDP in the limited sense that it decides whether a credential conversion is permitted. For example, not all users are authorised to obtain RCauth certificates. *It does not* act as a PDP for services in general.

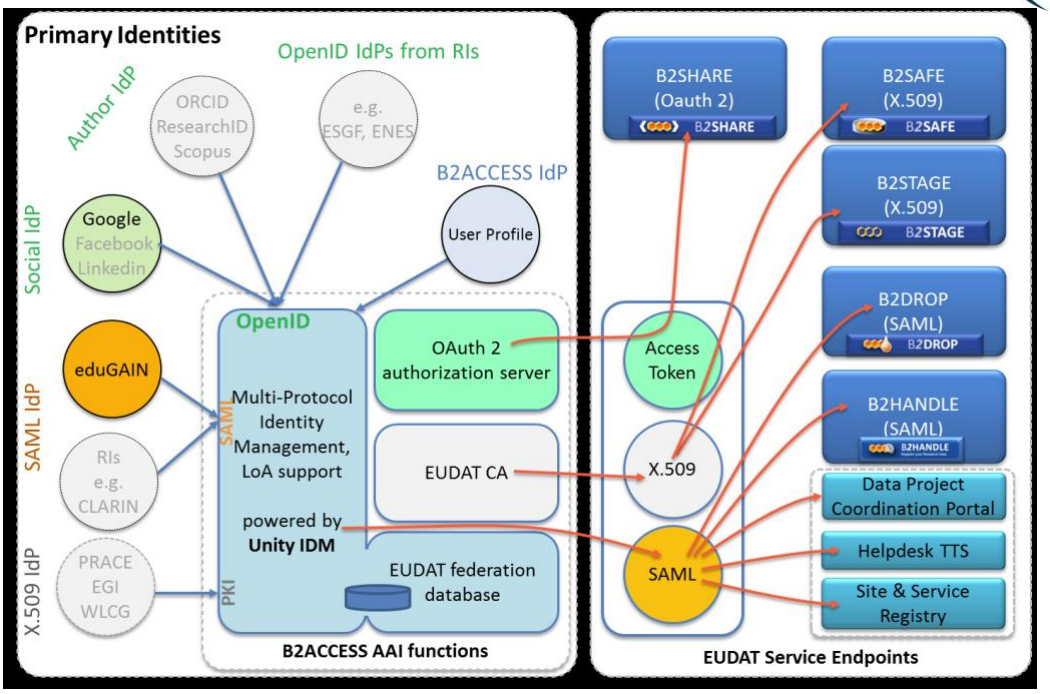

Figure 3.13: EUDAT AAI architecture (W Elbers, J Jensen, S Memon, et al.)

It should be noted that towards the end of the EUDAT2020 project, a full XACML infrastructure (see Figure 3.14) was deployed and tested. However, it was not fully integrated with the services, nor was it used in production.

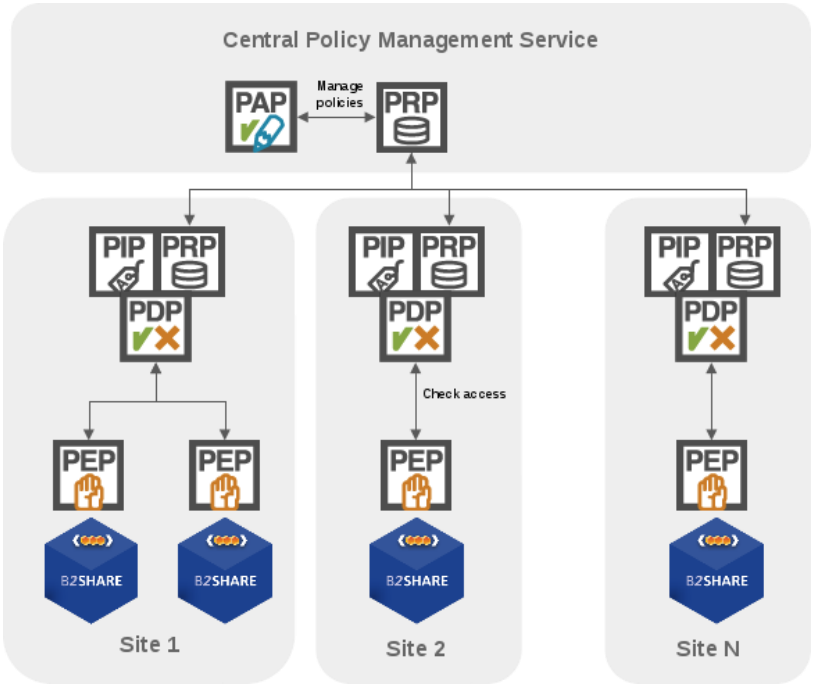

Figure 3.14: EUDAT XACML-based authorisation model

The XACML-based (proof-of-concept) architecture is based on a two-level hierarchy; at the top level, rules are defined as policy sets for each type of EUDAT service. The service-specific policies are managed by their service administrators at this level. Service administrators create or update policies through the central (read-write) PAP. The central policy repository (PRP) pushes the policies or policy sets to the site PRP - resides at a lower part of the hierarchy. Consequently, the site's PRP receives and executes the updates in a consistent manner. For each EUDAT site, a full XACML stack with a PEP for each service (or a group of closely co-located services) is deployed, and a single PDP for the site together with a local, read-only PRP.

In case of a user trying to access an EUDAT service, the service-specific PEP sends the user authorisation request to the PDP (which has access to the policies from the site PRP only) to evaluate the access decision requests, e.g. for a B2SHARE PEP, it will request only B2SHARE policy sets. However, the attributes required by the site PDP are usually sent by the B2ACCESS (shown in Figure 3.13) via the PEP, along with every authorisation request. The attributes contain roles, organisation or entitlements information.

3.6.4 Mapping to the Blueprint Architecture

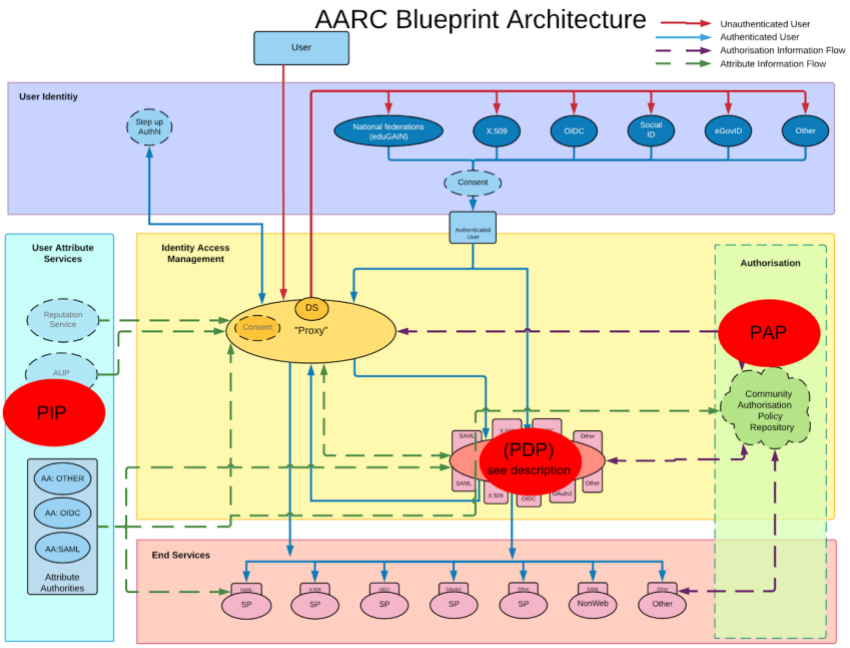

Figure 3.15: Mapping the EUDAT AAI authorisation model to the Blueprint Architecture

3.7 GEANT eduTEAMS

3.7.1 General description

GÉANT is a fundamental element of Europe's e-infrastructure, delivering the pan-European GÉANT network for scientific excellence, research, education and innovation. Through its integrated catalogue of connectivity, collaboration and identity services, GÉANT provides users with highly reliable, unconstrained access to computing, analysis, storage, applications and other resources, to ensure that Europe remains at the forefront of research.

Built on top of eduGAIN, eduTEAMS aims to provide an AAI solution for enabling communities to access and share resources using federated identities. It enables to integrate users from a wide range of environment, connecting them to specific services such as instruments, and also to other generic services such as storage and compute provided by any infrastructure provider or commercial entity.

3.7.2 Architecture and flow

The platform provides a full implementation of the AARC Blueprint Architecture comprised by an IdP/SP Proxy component (eduTEAMS Proxy) that acts as the gateway to Identity Providers in eduGAIN, a Membership Management Service (eduTEAMS MMS) for managing attributes and groups and onboarding members, an IdP Discovery Service (eduTEAMS DS), a Metadata Service (eduTEAMS MDS) and a guest identity service (eduTEAMS Identity Hub) that allows collaborations to engage with users outside of the eduGAIN community.

provides an overview of the functional components in the eduTEAMS architecture

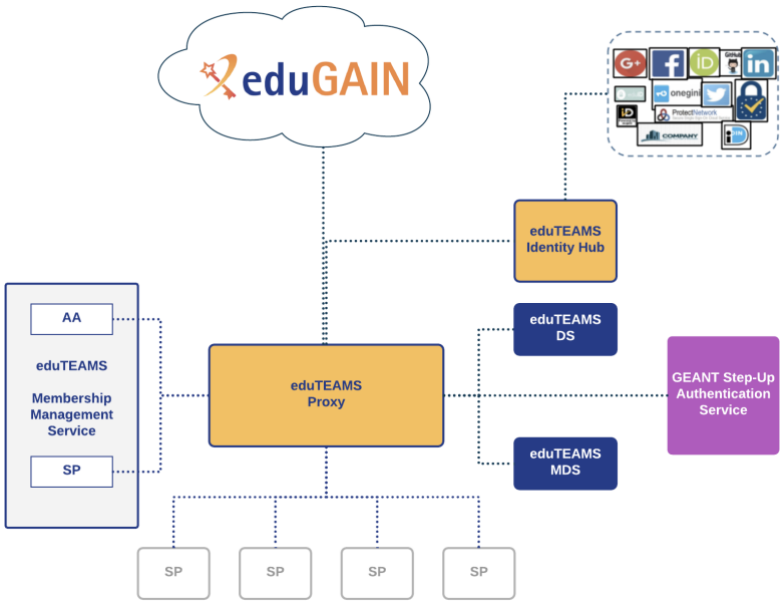

Figure 3.16: Overview of the functional components of the eduTEAMS architecture

3.7.3 Architecture described in the reference model

The eduTEAMS Membership Management Service acts a PIP. It provides information about authenticated users based on which policy decision can be made and enforced by other components of the platform.

In a similar manner, the eduTEAMS MDS is also acting as PIP, providing metadata information about Identity Providers and Service Providers to the rest of the components of the platform. In its basic configuration, it aggregates information from other PIPs, but in more advanced scenarios it can be configured to exclude, change or enrich the aggregated information with more information derived from local information stores.

The eduTEAMS Proxy can act as a PIP, a PAP, a PDP and/or a PEP depending on its specific configuration. For example, the eduTEAMS Proxy can provide information about the authenticated session of the user or inject new attributes to the user's identity, so that downstream SPs can use this information for making policy decisions, thus acting as PIP. In addition, the operators of the eduTEAMS Proxy can define policy configurations that affect how users access services, thus it acts also as a PAP. An example of the eduTEAMS Proxy acting as a PDP, is when it uses information from the eduTEAMS MDS based on which it can signal to the downstream SPs, whether the IdP used in a given authenticated session meets or not specific requirements. An example of the eduTEAMS Proxy acting as a PEP, is when it is configured to prevent access to specific services or to interrupt the flow of authenticated users and initiate alternate flows (e.g. registration for new users). In this case, the eduTEAMS Proxy is enacting on specific policy configuration, thus acting as a PEP.

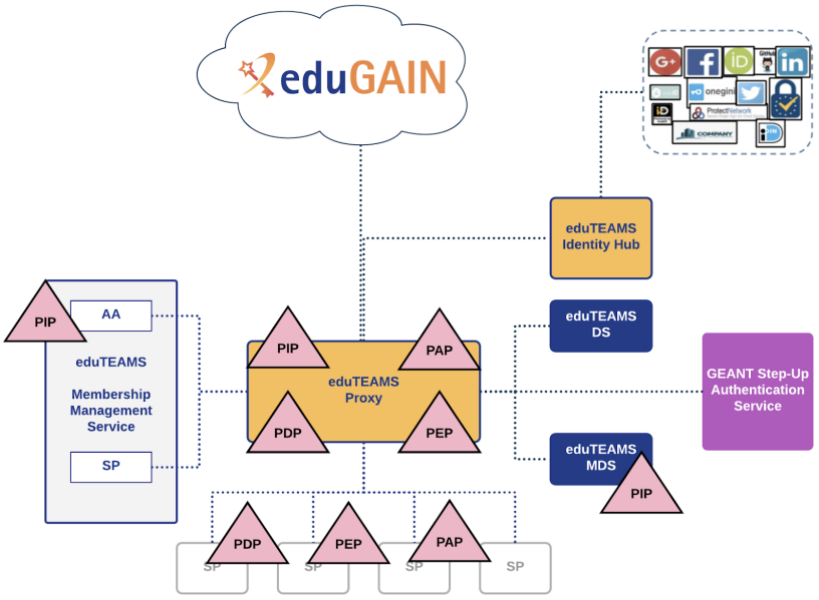

Figure 3.17: eduTEAMS components and their roles as defined in the authorisation model

3.7.4 Mapping to the Blueprint Architecture

From the perspective of the BPA, eduTEAMS Membership Management Service is an attribute authority and SP providing access to the user registration portal. The eduTEAMS Proxy acts as an IdP/SP proxy and as a token translation service between SAML to OIDC. Along with the eduTEAMS DS and the eduTEAMS MDS, they comprise the Identity Access Management Layer. The eduTEAMS Identity Hub is a (guest) Identity Provider. The authorisation layer in the AARC BPA is not directly mapped to an existing service component on the eduTEAMS platform, but we can say that is part integral part of the eduTEAMS Proxy and the SPs connected to the platform.

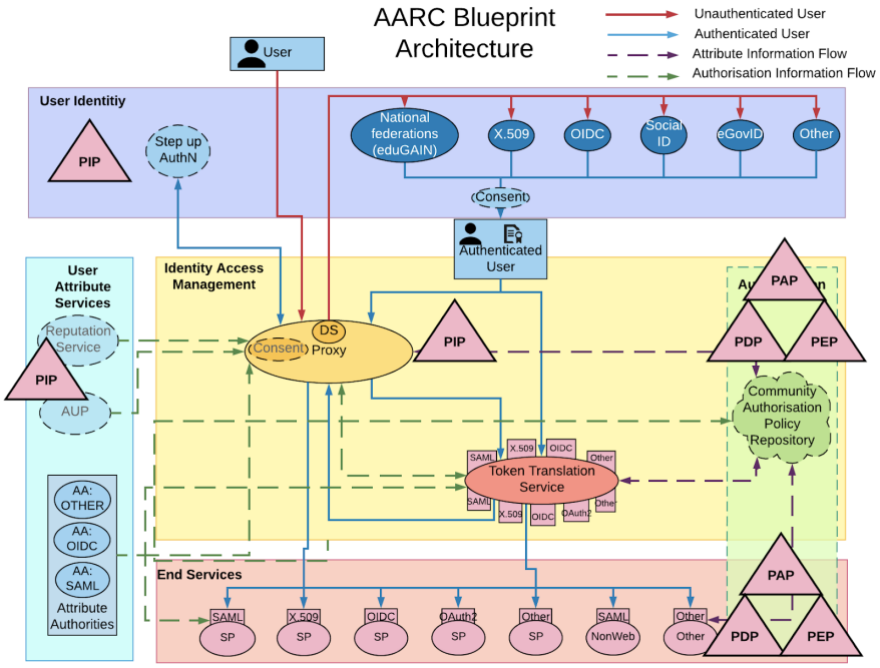

Figure 3.18: Mapping the eduTEAMS authorisation model to the Blueprint Architecture

3.8 LIGO Scientific Collaboration

3.8.1 General description

The LIGO Scientific Collaboration along with the Virgo collaboration analyses data generated by the LIGO, Virgo and Geo600 detectors to search for gravitational waves. It enables collaboration between approximately 1000 researchers in 20 countries. Data, documents, and other resources need to be shared amongst many different combinations of groups of researchers. We use the ABAC authorisation model along with SAML and X509 certificate technologies. The production environment utilises a single internal IdP, but the pilot will make use of a SAML proxy and federated identities.

3.8.2 Architecture and flow

In a typical web-based workflow, users will:

1. Request access to a resource within a web application.

2. The application will redirect users to the internal identity provider (or to the institution IdP via the SAML proxy)

3. Once authenticated, the user will be redirected back to the application and the user's information will be added to the application environment, in particular the isMemberOf attribute will hold the (typically dozens) of groups that the user belongs to.

4. These group permissions include home institution, collaboration membership, subgroup membership, and subgroup management roles. These group memberships are managed in the user management application, my.ligo.org, the Grouper application, and distributed via LDAP.

5. The web-application controls access to resources via configuration settings that determine rules from these roles, and also from user management of the resource access controls.

For X509 and SSH access:

A. The application server will use the LDAP and attribute rules to create an access control list.

B. Users will then authenticate to the ssh server using their X509 certificate or their ligo.org password and will then be compared against the ACL to decide whether they can login or access the resource.

3.8.3 Architecture described in the reference model

The final decision and enforcement of access rights is completed at the service providers indicated. They take information about the user from my.ligo.org, that we indicate with the PIP label. The policy rules can be either in the service provider configuration, or as part of the user-assigned rules. In both cases, the decision is made at the service provider; therefore the PEP, PDP and PAP all exist on the service. However, in some cases there is also some pre-processing of the groups and attributes by Grouper so we add a second PAP label here.

In the case of X509 + GSISSH access, the user's grid subject DNs are stored on the LDAP server (PIP) and compiled into a grid-mapfile by the LGMM application using rules decided by the operator. When users log in then the gsissh server checks against the grid-mapfile (PEP and PDP).

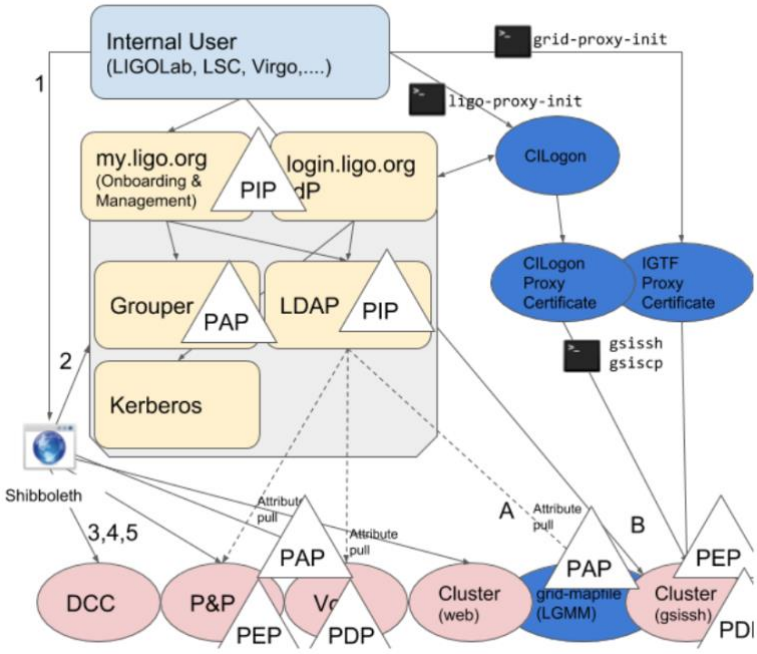

Figure 3.19: LIGO Scientific Collaboration authorisation model

3.8.4 Mapping to the Blueprint Architecture

The final decision and enforcement of access rights is completed at the service providers. They take information about the user (PIP) that exist in the user attribute area and the groups that they have been assigned to and are free to combine it in new ways. In the case of access control lists scenario, the ACLs may be constructed by one element and then used at the service level by a different element. This mapping also supports the inclusion of a PAP role within the authorisation section.

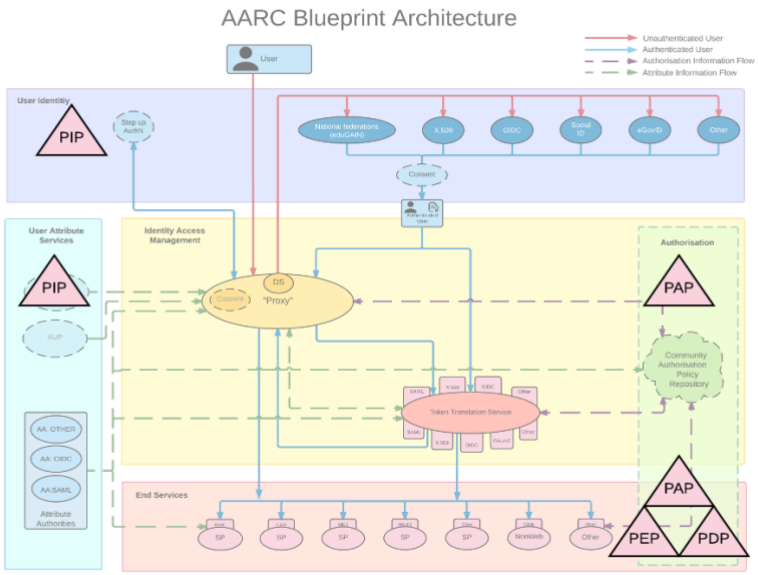

Figure 3.20: Mapping the LIGO Scientific Collaboration authorisation model to the Blueprint Architecture

3.9 WLCG Compute and storage facilities (VOMS / Token based)

3.9.1 General description

WLCG is the computing platform for the CERN based Large Hadron Collider. This particle accelerator operates four major experiments that operate in parallel to test a range of scientific hypotheses in particle physics. The authorisation systems have to provide access to data and compute resources for roughly 10,000 scientists. Fine grained read and write authorisation is required, respecting the need for confidentiality between certain experiments. Virtual organisation membership reflects experiment membership and is decided per community. Resources are provided by NGIs (in Europe) and by OSG (in the US). To reflect membership, a role-based access control scheme (RBAC, see A.3) is in place. It is enabled via the VOMS (see B.1) and the XACML based ARGUS (see B.8) technologies.

3.9.2 Architecture and flow

The access control mechanisms currently employed follow the model established for the European grid computing, which reflects a contract between a VO and resource centres [CORNWALL2004]. The model recognises three main levels of authorisation policies that are taken into account:

- Infrastructure level (WLCG in collaboration with participating infrastructures): covers security and policy aspects, such as central banning of users, acceptable assurance profiles for IdPs and/or CAs.

- VO level: provides information on membership, group/role in the VO(s) of the user,

- Resource centre level: gives additional control to resource centre to authorise access to its resources. E.g. a list of supported VOs, and filtering required by local laws.

Ultimately, authorisation decisions are made by the resource centre that provides the requested service; the decisions can override any suggestions from higher levels. For instance, a resource centre can decide to ignore the central banning information, or to reject access for selected users even if they are valid VO members.

The model expects that the user is assigned an identifier (the X.509 subject name of an IGTF certificate), while the attributes are assigned to the users by their VO(s).

Information on the level of the infrastructure is maintained either using dedicated tools, like resource management (e.g. VOMS [ALF04], Perun [PERUN]) or an authorisation service (e.g. Argus [ARGUS]).

VOMS acts as an attribute authority asserting VO-related attributes for the user. This implements a "Push model" that conveys RFC3281 attribute certificates [RFC3281] from the user to the service as extensions to user "proxy" certificates (RFC3820 [RFC3280], not to be confused with the architecture SP-IdP-Proxy).

3.9.3 Architecture described in the reference model

WLCG provides attributes about the user via VOMS. VOMS corresponds to a PIP. In addition to VOMS, the Argus infrastructure is a hierarchical model that allows administration of policies (PAP) at several levels. All decisions about authorisation are made at the worker nodes (PEP) and encoded such into the system that they can be enforced (PEP) by the underlying operating system, as illustrated in Figure 3.21.

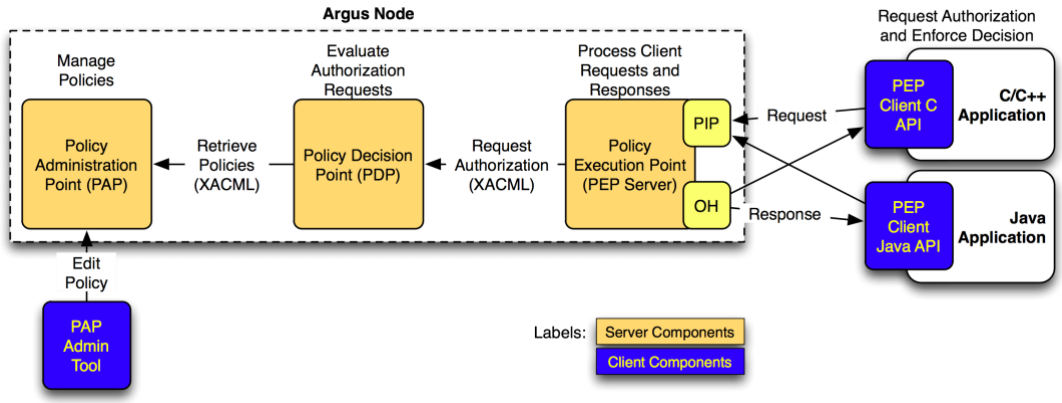

Figure 3.21: Argus authorisation service components

3.9.4 Mapping to the Blueprint Architecture

User attributes in our VOMS/PIP are located in the attribute layer. The process for obtaining attributes is external to the authorisation model and therefore greyed out in the diagram below. The authorisation layer holds the hierarchy of policy

administration points (PAP), that are used at the service level (the overlap of services and authorisation layers) for decision making and enforcement (PDP, PEP).

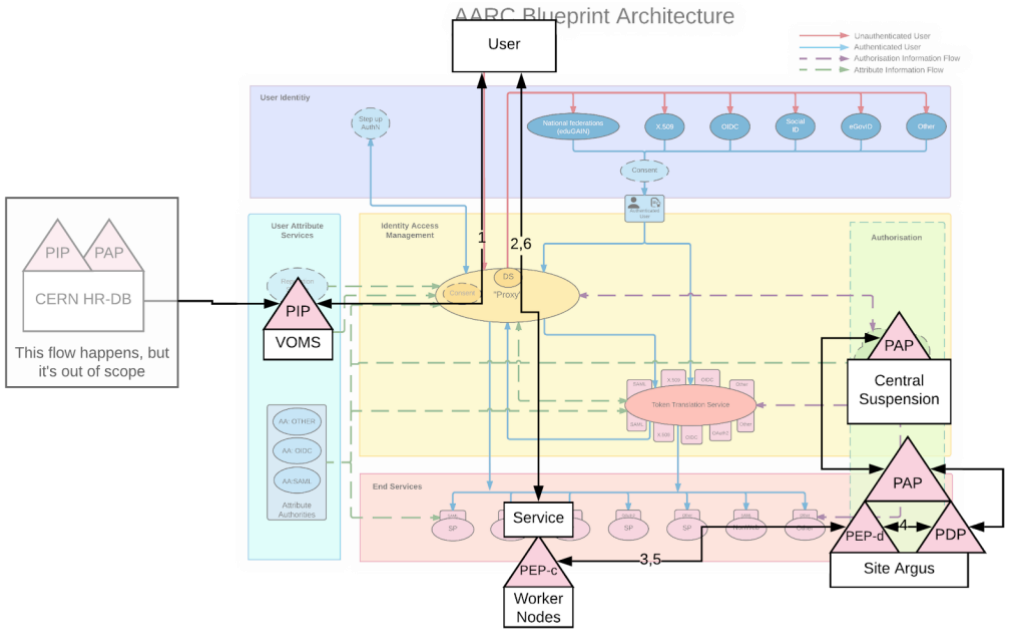

Figure 3.22: Mapping the WLCG authorisation model for compute and storage facilities to the Blueprint Architecture

4 Observed Models for Authorisation

This chapter presents the different authorisation models we have identified based on the authorisation use cases presented in Chapter 3. We observed five different models that can be summarised as follows:

1. Resource-local policy management and decision making
2. Centralised policy information point
3. Centralised policy management and decision making
4. Hierarchical policy management and decision making
5. Distributed policy enforcement

The remainder of this chapter provides a more detailed description of the observed models and discusses common types of attributes that can be used for authorisation.

4.1 Analysis of common authorisation models

This section provides an analysis of the five different authorisation models we observed.

4.1.1 Resource-local policy management and decision making

The user requests access to a resource protected by the SP. After successful authentication of the user at their home IdP, the SP retrieves additional information about the user (e.g. group membership and roles) or the context (e.g. assurance information) from different attribute sources (PIPs), such as group management systems and IdP metadata repositories. This information is used by the SP to determine if the user is authorised to access the service. Therefore, in this model, the SP serves three functions, i.e. PAP, PDP, and PEP.

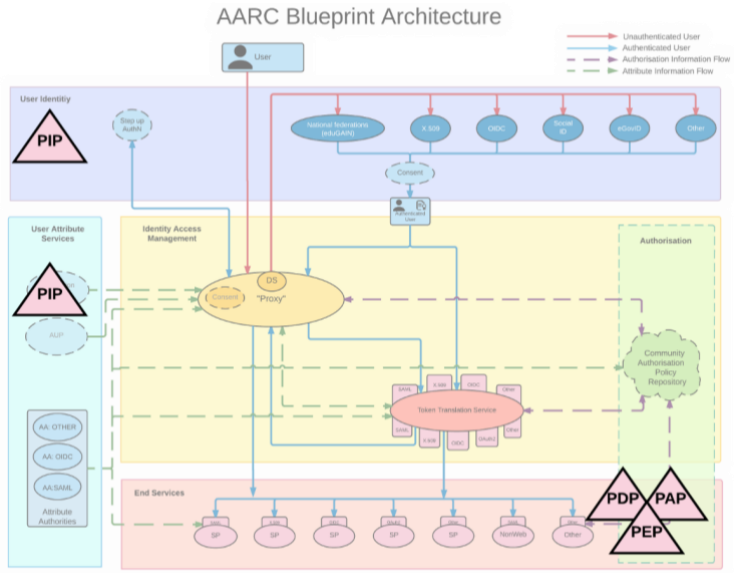

Figure 4.1: Resource-local policy management and decision making

As illustrated in Figure 4.1, this model does not require an SP-IdP-Proxy-based architecture.

Example use cases: CLARIN, EPOS, LIGO Scientific Collaboration

4.1.2 Centralised policy information point

This model assumes an SP-IdP-Proxy-based architecture whereby the proxy component acts as an SP towards the authentications providers and as an IdP towards the end services. The user requests access to a resource protected by the SP. After successful authentication of the user at their home IdP, the SP-IdP-Proxy retrieves additional information about the user (e.g. group membership and roles) or the context (e.g. assurance information) from different attribute sources, such as group management systems and IdP metadata repositories. The proxy supplements the attributes from the home IdP with information from the additional attribute sources and pushes the aggregated attribute set to the SP. Therefore, while the SP-IdP-Proxy is collecting information from multiple sources, it acts as a single PIP towards the SP. The SP uses the information from the proxy to determine if the user is authorised to access the resources. Thus, the SP serves three functions, i.e. PAP, PDP, and PEP, but it does not need to implement complex technical solutions for supporting multiple PIPs as in the case of resource-local policy management and decision making.

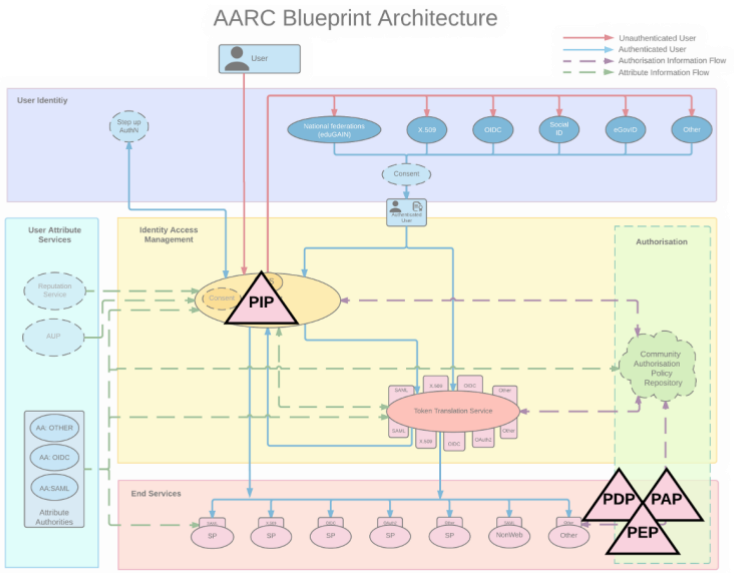

Figure 4.2: Centralised policy information point

Example use cases: EGI Check-in, EPOS, EUDAT B2ACCESS, GÉANT eduTEAMS, DARIAH AAI, LifeWatch, EISCAT_3D, CTA

4.1.3 Centralised policy management and decision making

This model assumes an SP-IdP-Proxy-based architecture whereby the proxy component acts as an SP towards the authentication providers and as an IdP towards the end services. The user requests access to a resource protected by the SP. After successful authentication of the user at their home IdP, the SP-IdP-Proxy retrieves additional information about the user (e.g. group membership and roles) or the context (e.g. assurance information) from different attribute sources (PIPs), such as group management systems and IdP metadata repositories. The proxy uses this information to evaluate the incoming request against policies it has been configured with and returns a decision to the SP. Therefore, the proxy acts as both a central PAP and PDP while the connected SPs serve as PEPs. Please note that in the figure, PAP and PDP are inside the light-yellow proxy layer, and as such are proxy functionality, even though they are not located inside the dark-yellow proxy ellipsis.

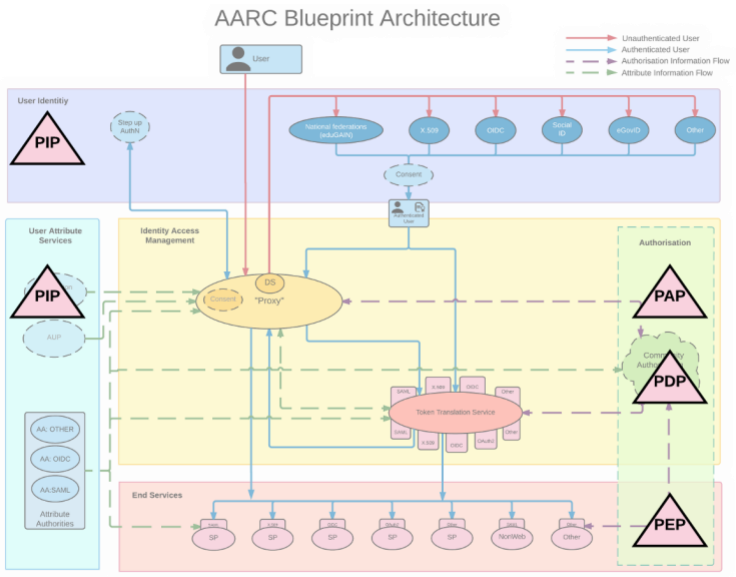

Figure 4.3: Centralised policy management and decision making

Example use cases: DARIAH AAI, EGI Check-in, ELIXIR AAI, GÉANT eduTEAMS, EUDAT B2ACCESS

4.1.4 Hierarchical policy management and decision making

In this authorisation model, the PAPs are deployed in a hierarchical way, as illustrated in Figure 4.4. An example hierarchical structure could have a central PAP at the root level connected to national PAPs, which are in turn connected to site-specific PAPs. This structure would allow each site-specific PAP instance to import policies from the national instance, which in turn can import policies from the central PAP. The PAP at the root level is often used for centrally suspending users.

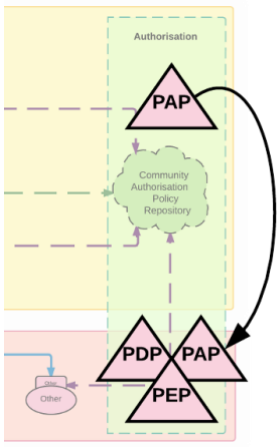

Figure 4.4: Hierarchical policy management and decision making

Example use cases: EPOS, EUDAT B2ACCESS, WLCG

4.1.5 Distributed policy enforcement

In this authorisation model, the policy enforcement functionality is distributed following a client-server architecture. There is a PEP server (PEP-s) component which is responsible for handling authorisation requests from different lightweight PEP clients (PEP-c). These PEP client libraries are used to authorise requests from the application side, and to enforce decisions locally. The PEP server ensures the integrity and consistency of the authorisation requests received from the PEP clients. The distributed policy enforcement approach can be applied as an extension to other authorisation models, including the centralised and the hierarchical policy management decision making described above.

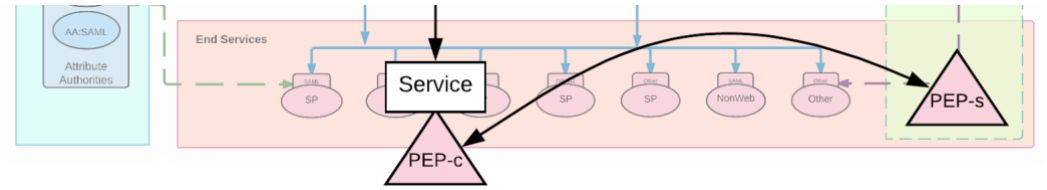

Figure 4.5: Distributed policy enforcement

Example use cases: WLCG

4.2 Common authorisation attributes

This section presents commonly used types of attributes that can be used for authorisation based on the analysed community-specific authorisation use cases.

4.2.1 Affiliation and group/project information

Cases described here are based on the human genome data (European Phenome-Genome Archive) and BBMRI, but stand as examples for additional communities.

The relying services (such as EGA) often give access to users not personally but as a user affiliated with their Home Organisation. This involves a user in his role as a representative of a given home organisation or project member (i.e. as in RBAC). This poses requirements on the SP-IdP-proxy. Especially when a user departs from their home-organisation their access rights must be closed swiftly. An approach to implement this is observing changes in user's eduPersonAffiliation attribute whose freshness (RAF -> ATP) becomes therefore important for the infrastructure.

Several communities have adopted an approach where users can link several external (Home Organisation) identities to their community identity. In such circumstances, the community needs to pay attention to which one the user's permissions are actually coupled. If the user's affiliation to that Home Organisation ceases their related permissions must be closed in the service, although the user may still be a legitimate holder of an active community ID.

In some circumstances the segregation of permissions associated to two different parallel Home Organisations must be also enforced in a dynamic way. For instance, the user has permissions to dataset X as a researcher of university A and permission to dataset Y as a researcher of university B. When they now log in to a compute environment, they must not see both dataset X and Y at the same time. Instead they need to first indicate "for this session I log in as a representative of university A" and then dataset X is available for them but dataset Y is not, and vice versa.

A very similar case is access granted based on a project membership. When logged in to the compute environment via one project role a user may access only data belonging to that project, but not files of the other projects they are otherwise permitted to access. (c.f. Dynamic segregation of duties).

4.2.2 Access context information

Assurance frameworks, such as implemented by the REFEDS Assurance Framework (RAF) or the NIST SP-800-63-3, decompose the assurance about a person into independent components. The RAF, for example, defines components such as ID-Proofing (how well is a user known), authentication strength (how strongly was a user authentication, password, security token, combinations) and more. These fine-grained assurance profiles allow infrastructures and communities to control access to resources based on the assurance profile of the user.

In some cases, infrastructures or communities find it useful to use a priori knowledge about ID-providers to make an authorisation decision. This is for example used by Services that need to authorise parts of their services differently, if users come from social network IdPs or from an academic home organisation.

Examples:

- Services that require use of a 2nd factor to be accessed (i.e. increase of the authentication component).
- IGTF certificates that are only issued if a certain level of identity vetting has been met.
- RCAuth.eu puts requirements on the user's assurance profile, (identifier-uniqueness) and in addition on the security policies supported by the identity provider that authenticated a user.

It should be noted that an authorisation decision is not typically based solely on the information from an assurance-related PIP. In most use cases, assurance information is combined with other types of attributes about the user, such as group or affiliation information.

5 Conclusions

This document provided an analysis of authorisation use-cases collected from the research communities participating in the pilot activities of AARC. To this end, we engaged with representatives from the communities through a series of unstructured interviews and meetings in order to describe their authorisation approaches using well-established concepts, such as PDP, PAP, PEP and PIP, the so-called "PPP model". Through these concepts we were able to model different authorisation functions, including management, evaluation and enforcement of policies, in a technology-agnostic way. Mapping the components of the analysed community-specific authorisation approaches to the AARC Blueprint Architecture allowed for identifying five common authorisation models: i) Resource-local policy management and decision making; ii) Centralised policy information point; iii) Centralised policy management and decision making; iv) Hierarchical policy management and decision making; and v) Distributed policy enforcement.

The resource-local policy management and decision making approach is the basic model (which does not require an AARC BPA-compliant AAI), whereby each service is responsible for evaluating access requests and making decisions by aggregating information from one or more attribute sources (PIPs). The centralised PIP model is an evolved version of the basic model, in which the SP-IdP-Proxy acts as a single PIP so that services don't need to implement complex technical solutions for supporting multiple sources of authorisation attributes. In the centralised policy management and decision making model, the decision as to whether a user can access a specific service can be taken centrally and then communicated to the service by adding a service-specific entitlement/capability to the user's attributes. The hierarchical policy management approach is an extension of the previous models allowing for managing policies at different levels, thus enabling policies at lower levels to override or extend the policies defined in upper levels. Lastly, the distributed enforcement policy enforcement is suitable for use-cases requiring the coordinated execution of compute/storage tasks across multiple worker nodes following a client-server architecture.

Appendix A Authorisation patterns

This appendix briefly discusses different authorisation patterns. For further information the reader is referred to [WINDLEY2005].

A.1 ACL

An access control list (ACL) [RFC4949], with respect to a computer system, is a list of subject IDs and associated permissions attached to an object. An ACL specifies which users or system processes are granted access to objects, as well as what operations are allowed on given objects.

Each entry in a typical ACL specifies a subject and an operation. For instance, if a file object has an ACL that contains (Alice: read,write; Bob: read), this would give Alice permission to read and write the file, and Bob to only read it.

For data interchange, and for "high level comparisons", ACL data can be translated to XACML

A.2 RBAC

The Role-based access control (RBAC) [SANDHU1996] is an approach to restricting system access to authorised users, and it is sometimes referred to as role-based security. The model is based on the roles that can be assigned to the users of the system and on the privileges associated with such roles.

Within an organisation, roles are created for various job functions, and the permissions to perform certain operations are assigned to specific roles. Members or staff (or other system users) are assigned particular roles, and through those role assignments acquire the system permissions to perform particular functions. Since users are not assigned permissions directly, but only acquire them through their role (or roles), management of individual user rights becomes a matter of simply assigning appropriate roles to the user's account; this simplifies common operations, such as adding a user, or changing a user's department. This is more scalable than access control lists; see [AARC-G036] for further discussion.

More sophisticated RBAC models would take time into account, e.g. to allow a user who goes on holiday to temporarily delegate a role to a colleague.

A.3 ABAC

Attribute-based access control (ABAC) [NIST800-162] defines an access control paradigm whereby access rights are granted to users through the use of policies. The policies are statements that bring together attributes to express what can happen and what is not allowed: in ABAC there can be granting or denying policies, local or global policies, and they can also be written in a way that overrides other policies.

One standard that implements attribute- and policy-based access control is XACML.

The fact that the access is based on (potentially) arbitrary attributes of the subject and object, as well as the environmental conditions, allows for greater flexibility and broader set of fine-grained access control policies as compared to the predefined assignment of roles or groups to the user in RBAC.

Unlike role-based access control (RBAC), which employs predefined roles that carry a specific set of privileges associated with them and to which subjects are assigned, the key difference with ABAC is the concept of policies that express a complex boolean rule set that can evaluate many different attributes.

A.4 Capability based

Capability-based security is a concept in the design of secure computing systems, one of the existing security models. It was defined by Dennis and Van Horn in 1966 as “a token, ticket, or key that gives the possessor permission to access an entity or object in a computer system”

In our context programs and operating systems are understood as federated services. A user program on a capability-based operating system must use a capability to access an object. In a system with capabilities, the mere fact that a user program possesses that capability entitles it to use the referenced object in accordance with the rights that are specified by that capability. In theory, a system with capabilities removes the need for any access control lists or similar mechanisms by giving all entities all and only the capabilities they will actually need.

The user does not access the data structure or object directly, but instead via a handle. In practice, it is used much like a file descriptor in a traditional operating system (a traditional handle), but to access every object on the system. Programs possessing capabilities can perform functions on them, such as passing them on to other programs (delegation), converting them to a less-privileged version, or deleting them.

Appendix B Technological overview

This section briefly presents authorisation technologies relevant for this document. For a more detailed description please refer to [AARC-MJRA1.1].

B.1 RFC 3820 proxies with VOMS

X.509 is an ITU-T standard for a public key infrastructure (PKI), also known as PKIX (PKI X509). For delegating user authentication, a special type of certificate has been introduced, a so-called proxy certificate, standardised in RFC3820, where the certificate is signed by the (private key belonging to the) user's end-entity certificate (EEC) or another proxy certificate, instead of the CA. In practice proxy certificates are often combined with attribute certificates (signed by the private key belonging to the previous certificate level in the chain), in particular as used by VOMS.

B.2 OAuth2

OAuth (OAuth2.0, RFC6749) is an open standard for access delegation, commonly used as a way for Internet users to grant websites or applications access to their information on other websites but without giving them the passwords.

The OAuth 2.0 authorisation framework enables a third-party application to obtain limited access to an HTTP service, either on behalf of a resource owner by orchestrating an approval interaction between the resource owner and the HTTP service, or by allowing the third-party application to obtain access on its own behalf. This mechanism is used by companies such as Amazon, Google, Facebook, Microsoft and Twitter to permit the users to share information about their accounts with third party applications or websites.

B.2.1 JWT

JSON Web Token (JWT) is a JSON-based open standard (RFC 7519) for creating access tokens that assert some number of claims. For example, a server could generate a token that has the claim "logged in as admin" and provide that to a client. The client could then use that token to prove that it is logged in as admin. The tokens are signed by the server's key, so the client and server are both able to verify that the token is legitimate. The tokens are designed to be compact, URL-safe and usable especially in web browser single sign-on (SSO) context. JWT claims can be typically used to pass identity of authenticated users between an identity provider and a service provider, or any other type of claims as required by

business processes the tokens can also be authenticated and encrypted JWT relies on other JSON-based standards: JWS (JSON Web Signature) RFC 7515 and JWE (JSON Web Encryption) RFC 7516.

B.2.2 SciTokens

The SciTokens project (https://scitokens.org/) is applying the OAuth 2.0 and JSON Web Token standards to authorisation in distributed scientific computing infrastructures, where remote compute jobs need security credentials to access data and other resources. In common practice today, those credentials are identity tokens, carrying the identity of the individual researcher or the virtual organisation (VO), enabling the job to act on behalf of that researcher or VO when accessing remote resources like file servers. Using identity tokens in this way creates significant risk of abuse, since they are used by jobs that are running on remote, less trusted systems and if stolen, these tokens provide wide access to the attacker. SciTokens follows the OAuth 2.0 model of delegating capabilities (access tokens) with specific scopes that indicate the specific resources needed by each job. Using signed JSON Web Tokens makes the capabilities self-describing, so they can be validated locally by widely distributed resource providers, in contrast to other approaches that require a central OAuth token server for validation.

B.2.3 Macaroons

Macaroons are bearer tokens similar to cookies but with contextual caveats.

1. Macaroons implement capability based authorisation using a hashing method

2. Macaroons are based on (symmetric) secrets which means that the service that issues the macaroon is the service which will accept the macaroon. Otherwise the two services would have a very strong trust and key-management issues.

3. Caveats are used to narrow capabilities down and may be added by third parties. They are easy to add, but hard (cryptographically hard) to remove.

A consequence of 2. is that it is computationally cheap to create a macaroon. It should be easily possible to generate macaroons in the kHz range with a single server. So, creating a macaroon per request is feasible. This is in contrast to JWT, which probably need caching.

One consequence of 3. is that any user can create a more limited macaroon without contacting a central service, so it scales well. This is a bit like with proxy certificates and quite different from OAuth2/JWT.

Another consequence of 3. is that each caveat describes what the user **can't do**. This may be somewhat counter-intuitive, but comes from the fact that anyone can add a caveat (create a new macaroon with extra caveat), so the caveats don't authorise activity, but the opposite: restrict activity.

B.2.4 SAML attributes / OpenID Connect claims

Both SAML attribute assertions and OpenID Connect claims can be used to transport information that may be used for authorisation.

OpenID Connect Clients use scope values, as defined in Section 3.3 of OAuth 2.0 [RFC6749], to specify what access privileges are being requested for Access Tokens. The scopes associated with Access Tokens determine what resources will be available when they are used to access OAuth 2.0 protected endpoints.

B.3 XACML (technological view, implementation of RFC 2753)

XACML is an implementation of the architecture specified in RFC 2753. XACML itself defines the content of some of the messages necessary to implement this model, but deliberately confines its scope to the language elements used directly by the PDP and does not define protocols or transport mechanisms. Full implementation of the usage model depends on use of other standards to specify assertions, protocols, and transport mechanisms. XACML also does not specify how to implement a Policy Enforcement Point, Policy Administration Point, Attribute Authority, Context Handler, or Repository, but XACML artifacts can serve as a standard format for exchanging information between these entities when combined with other standards.

One standard suitable for providing the assertion and protocol mechanisms needed by XACML is the OASIS Security Assertion Markup Language (SAML), Version 2.0 [SAML]. SAML defines schemas intended for use in requesting and responding with various types of security assertions. For that purpose an extension to the SAML standard was defined, the XACML SAML Profile, that can be used to convey security information within a system that uses XACML. This profile specifies extensions to the SAML profile that allow to transport security information (i.e. XACML artifacts) between components:

- PEP to PDP to request a XACML authorisation decision
- PDP to PED to convey such a decision
- PDP to PAP to query for XACML policy information
- PAP to PDP to respond to such a policy query

In addition to that the usual mechanisms defined in the SAML standard can be used to query for and convey attributes from e.g. an external attribute authority to the PDP, that then might come to a decision based on this information.

See e.g. the sections on XACML and Argus in [AARC-MJRA1.1].

B.4 Posix-related points

Unix-like and otherwise POSIX-compliant systems, including Linux-based systems and all macOS versions, have a simple system for managing individual file permissions. This includes methods to assign permissions or access rights to specific users and groups of users. These systems control the ability of the users to view, change, navigate, and execute the contents of the file system.

Posix file access in itself does not convey authorisation information. In itself it may be used to represent and enforce authorisation decisions. The fact that many services that are provided today rely on Posix systems underneath encourages to translate external authorisation information into Posix filesystem permissions.

Existing systems that implemented this are mostly the HEP experiments around CERN and other X.509 based infrastructures. They use locally deployed mapping files to map users to local (often pooled) Posix accounts.

B.5 REMS tool

REMS (Resource Entitlement Management System) is an open source tool developed at CSC - Finnish IT Centre for Science. REMS is designed to manage access rights to datasets and other resources. An applicant (authenticated by a SAML Identity Provider) logs in to REMS, identifies the dataset(s) she wants to apply for access rights to herself and her research groups, commits to the dataset's terms of use and submits the application. REMS follows the workflow configured for the dataset and circulates the application to one or more persons for review and approval. As a result, REMS outputs a stream of (user-ID, dataset-ID) tuples for the PDP and the downstream PEP(s). REMS is currently used in ELIXIR for controlled access datasets, BBMRI for biobank samples and CLARIN for language material belonging to the RES (restricted) category. For a short overview on REMS, see the REMS leaflet (pdf).

Appendix C Technical parts of use cases

This appendix contains technical details about the analysed use-cases that were omitted from the main part of the document for the sake of brevity.

C.1 EGI Check-in

C.1.1 Resource-specific entitlements

A resource-specific entitlement represents the right of a user to access a particular resource. For example, the `urn:mace:egi.eu:aai.egi.eu:rcauth` value is currently being used to indicate that the holder of this entitlement is eligible for accessing the RCauth.eu Online CA service. The EGI AAI URN registry lists all supported entitlement values.
Note that the resource-specific entitlements are meant to be used to grant access to specific EGI central services rather than distributed services, such as HTC or cloud resources, for which authorisation is typically based on group membership.

C.1.2 VO/Group-related entitlements

The eduPersonEntitlement values follow [AARC-G002] for expressing VO/group membership and role information and are thus formatted as follows:
`urn:mace:egi.eu:group:<VO>[[:<GROUP>][:<SUBGROUP>*]][:role=<ROLE>]#<GROUP-AUTHORITY>`
where:

- <VO> is the name of the Virtual Organisation
- <GROUP> is the name of a group in the identified <VO>; specifying a group is optional
- zero or more <SUBGROUP> components represent the hierarchy of subgroups in the <GROUP>; specifying sub-groups is optional
- the optional <ROLE> component is scoped to the rightmost (sub)group; if no group information is specified, the role applies to the VO
- <GROUP-AUTHORITY> is a non-empty string that indicates the authoritative source for the entitlement value. For example, it can be the FQDN of the group management system that is responsible for the identified group membership information

C.1.3 Levels of Assurance

Based on the authentication method selected by the user, EGI Check-in assigns a Level of Assurance (LoA), which is conveyed to the SP through either the eduPersonAssurance attribute and the Authentication Context Class (AuthnContextClassRef) of the SAML authentication response, or using the acr claim in the case of OIDC services. While the EGI AAI currently distinguishes between three LoA levels, namely Low, Substantial and High, it is planned to support the REFEDS Assurance Framework (RAF), which allows for both a composite assurance level/profile and for assurance component values to be expressed. In the RAF, it is the component values that play the principle role in expressing assurance information, and the composite profiles (e.g. "Cappuccino" and "Espresso") are the result of a specific combination of assurance components.

C.2 Elixir

C.2.1 Bona Fide researcher in other research communities

The concept of the Bona Fide research is also used in communities other than Life Sciences.

The Human Brain Project (HBP) offers different accounts for researchers with different rights:

- HBP Identity Accounts are available by invitation. Identity accounts allows access to many of the tools produced by the HBP, including the HBP Collaboratory.
- Full HBP Membership requires that to be a member of a lab in one of the HBP Partner Institutions. HBP Membership grants privileged access to the Platforms of the project. People be invited in one of three ways:
    - Invitation by a current HBP Identity account holder.
    - Contact a HBP SubProject manager to receive an invitation.
    - Request an invitation by sending a short email describing your interest in the HBP Platforms to platform@humanbrainproject.eu.
- The HBP Collaboratory and the HBP Platforms are subject to some restrictions on their use. In most cases these restrictions are due to limited computing or storage capacity powering the Platform service offerings. There may be reasons to expand resource allocations for particular services if a strong scientific case can be made for the increased allocation. Access to resources is given on a per project basis or to partnering projects.

C.3 GEANT eduTEAMS

C.3.1 Group membership and roles

In eduTEAMS the functionality of managing user registration, groups and roles is provided by the eduTEAMS MMS, which comes in three flavours. Communities can choose the flavour that best matches their needs (e.g. some communities prefer a simplified environment that allows them to perform operation in an easy and quick manner, while other communities have more complex structures and require more fine grained control and this needs to be supported by the MMS). Regardless of the flavours of the MMS used, expression of group membership and role information is following the AARC guidelines [AARC-G002].

The URN format used is the following:

```
urn:geant:eduteams.org:group:<GROUP>[:<SUBGROUP>*][:role=<ROLE>]#mms.eduteams.org
```

In standalone deployments for a community, the namespace is changed to reflect a namespace managed by the community and the attribute authority is adjusted accordingly.

C.3.2 Handling of authorisation

eduTEAMS recognises that not all SPs can operate in the same way. Some SP operators require full control, while other SP operators would prefer to rely on a eduTEAMS so provide them with a simple entitlement or even handle authorisation for them. In this regard, eduTEAMS support three strategies, which can be implemented on per SP basis:

- By default eduTEAMS releases all attributes retrieved by the eduTEAMS MMS to the SPs. Attribute release policies can be configured per SP. It is up to the SPs to evaluate the SAML assertion or the OIDC claims retrieved from eduTEAMS and decide whether a user is allowed to access and use their resources.
- eduTEAMS can be configured to release service specific entitlements, based on which the service providers can allow or prevent the user from accessing their resources
- eduTEAMS can be configured to interrupt the user authentication flow and prevent the user from accessing specific services if some requirements are not met. For example, if eduTEAMS can have a policy configuration for given SP that requires multi-factor authentication before the user is allowed to access that SP.

C.3.3 Levels of Assurance

eduTEAMS supports requesting and expressing required levels of assurance. There are three ways to supports SPs requiring specific level of assurance:

- SPs can express the need for a specific level of assurance in their metadata
- SPs can request a specific level of assurance during the authentication request
- eduTEAMS can have a policy configuration for an SP to require a specific level of assurance.

eduTEAMS supports the expression of levels of assurance based on the information received by the Identity Provider used in a given authentication session. The assurance level is singled to the SAML SPs using the Authentication Context class and the eduPersonAssurance attribute, while to OIDC services using the acr claim.

With the new GEANT Step-up authentication service which is currently under piloting, eduTEAMS will be able use that service in combination with the information provided by the Identity Providers.

C.3.4 Metadata Service

eduTEAMS is using internally a metadata service, which is a full MDX implementation. eduTEAMS MDS aggregates IdP and SP metadata from various sources, such as eduGAIN, federation feeds, but also directly from SPs and IdPs that might be connected directly to eduTEAMS. eduTEAMS MDS can process the aggregate feeds and apply policy configuration on them, allowing the eduTEAMS operators to blacklist, tag or introduce new IdP/SP entities. The eduTEAMS MDS is used primarily by the eduTEAMS Proxy, but also by other components of eduTEAMS, in order to get information about the entities they are interacting with. So, for example, in each authentication flow, the eduTEAMS Proxy queries the eduTEAMS MDS to get information about the Identity Provider during the authentication. The most common scenarios at the moment are three:

- Introducing new SPs/IdPs that are not available from eduGAIN

- Blacklisting IdPs or SPs
- Check whether an IdP used supports certain policies (e.g. R&S and SIRTFI)

References

[AARC-G002]	Expressing group membership and role information (AARC-G002); https://aarc-project.eu/guidelines/aarc-g002/
[AARC-G006]	Best Practices for managing authorisation (AARC-G006); https://aarc-project.eu/guidelines/aarc-g006/
[AARC-G036]	Roles, responsibilities and security considerations for VOs (AARC-G036); https://aarc-project.eu/guidelines/aarc-g036/
[AARC-MJRA1.1]	Existing AAI and available technologies for federated access (MJRA1.1); https://aarc-project.eu/wp-content/uploads/2016/01/MJRA1.1-Existing-AAI-and-available-technologies.pdf
[ALF04]	Alfieri R. et al. (2004) VOMS, an Authorization System for Virtual Organizations. In: Fernández Rivera F., Bubak M., Gómez Tato A., Doallo R. (eds) Grid Computing. Lecture Notes in Computer Science, vol 2970. Springer, Berlin, Heidelberg
[ARGUS]	Argus Authorization Service; http://argus-documentation.readthedocs.io/en/stable/
[CLARIN]	Common Language Resources and Technology Infrastructure (CLARIN); https://www.clarin.eu/content/clarin-in-a-nutshell
[CLARIN-SPF]	CLARIN Service Provider Federation; https://www.clarin.eu/content/service-provider-federation
[CORNWALL2004]	Cornwall, Linda A., et al. "Authentication and authorisation mechanisms for multi-domain grid environments." Journal of Grid Computing 2.4 (2004), pp. 301-311.
[DYKE2016]	Dyke, S., Kirby, E.,Shabani, M., Thorogood, A.,Kato, K., Knoppers, B. Registered access: a 'Triple-A' approach. *European Journal of Human Genetics* volume 24, pages 1676–1680 (2016); https://www.nature.com/ejhg/journal/v24/n12/full/ejhg2016115a.html
[EGI-REG]	EGI AAI entitlement registry; https://wiki.egi.eu/wiki/URN_Registry:aai.egi.eu
[EPOS]	European Plate Observing System (EPOS); https://www.epos-ip.org/
[NIST800-162]	Guide to Attribute Based Access Control (ABAC) Definition and Considerations, (NIST Special Publication 800-162), 2014; https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-162.pdf
[PERUN]	Perun, Identity and Access Management System; https://perun.cesnet.cz/web/
[SANDHU1996]	Sandhu, R., Coyne, E.J., Feinstein, H.L. and Youman, C.E. (August 1996). "Role-Based Access Control Models" (PDF). IEEE Computer. IEEE Press. 29 (2): 38–47. doi:10.1109/2.485845
[TLA]	The Language Archive (TLA); http://tla.mpi.nl/
[TLA-FLAT]	Fedora Language Archiving Technology (FLAT); https://github.com/TLA-FLAT/FLAT
[RFC2753]	A Framework for Policy-based Admission Control, RFC 2753; https://tools.ietf.org/html/rfc2753
[RFC2904]	AAA Authorization Framework; https://tools.ietf.org/html/rfc2904
[RFC3280]	Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, RFC 3280; https://www.ietf.org/rfc/rfc3280.txt
[RFC3281]	An Internet Attribute Certificate Profile for Authorization, RFC 3281; https://www.ietf.org/rfc/rfc3281.txt
[RFC4949]	Internet Security Glossary, Version 2, RFC 4949; https://tools.ietf.org/html/rfc4949
[WINDLEY2005]	Windley P. "Authorisation patterns". Digital Identity. Unmasking Identity Management Architecture (IMA). O'Reilly Media Inc, 2005.

Glossary

AA	Attribute Authority
AAI	Authentication and Authorisation Infrastructure
AARC	Authentication and Authorisation for Research and Collaboration
API	Application Programming Interface
AS	Authorisation Server
AUP	Acceptable Use Policy
AuthN	Authentication
AuthZ	Authorisation
BPA	Blueprint Architecture
CA	Certification Authority
DN	Distinguished Name
ECP	Enhanced Client Protocol
eduGAIN	International interfederation service interconnecting research and education
EGI	European Grid Infrastructure
EI	e-Infrastructures
eP	eduPerson
ePA	eduPersonAffiliation
ePE	eduPersonEntitlement
ePO	eduPersonOrcid
ePPN	eduPersonPrincipalName
ePSA	eduPersonScopedAffiliation
ePUID	eduPersonUniqueId
FIM	Federated Identity Management
FQDN	Fully Qualified Domain Name
FTP	File Transfer Protocol
GDPR	General Data Protection Regulation
GSI	Grid Security Infrastructure
GSS-API	Generic Security Service Application Program Interface
GUI	Graphical User Interface
HO	Home Organisation
HPC	High-Performance Computing
HTTP	Hypertext Transfer Protocol
Iaas	Infrastructure as a Service
IAM	Identity Access Management
IdP	Identity Provider
IGTF	Interoperable Global Trust Federation
IR	Incident Response
JRA1	Joint Research Activity 1, Architectures for an integrated and interoperable AAI

JSON	JavaScript Object Notation
PKI	Public Key Infrastructure
LCMAPS	Local Credential Mapping Service
LDAP	Lightweight Directory Access Protocol
LIGO	Laser Interferometer Gravitational-Wave Observatory
LoA	Level of Assurance
MACE	Middleware Architecture Committee for Education
NREN	National Research and Education Network
OAuth2	The industry-standard protocol for authorisation
OIDC	OpenID Connect
OIDCre	OpenID Connect for Research and Education Working Group
OP	OpenID Connect Provider
OS	Operational Security
PAM	Pluggable Authentication Modules
POSIX	Portable Operating System Interface
PR	Participant Responsibilities
RBAC	Role-Based Access Control
RC	Research Collaboration
RCauth	The white-label Research and Collaboration Authentication CA Service for Europe
RI	Research Infrastructures
R&E	Research and Education
R&S	Research and Scholarship
RP	Relying Party
TTS	Token Translation Service
UC	User Community
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
URN	Uniform Resource Name
VLE	Virtual Learning Environment
VM	Virtual Machine
VO	Virtual Organisation
VOMS	Virtual Organisation Membership Services
XACML	Extensible Access Control Markup Language